

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

High Voltage Considerations for Dark Matter Searches

Permalink

<https://escholarship.org/uc/item/5t04640t>

Author

Watson, James Reed

Publication Date

2023

Peer reviewed|Thesis/dissertation

High Voltage Considerations for Dark Matter Searches

by

James Reed Watson

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Daniel McKinsey, Chair

Professor Kam-Biu Luk

Professor Eric Norman

Fall 2022

High Voltage Considerations for Dark Matter Searches

Copyright 2022
by
James Reed Watson

Abstract

High Voltage Considerations for Dark Matter Searches

by

James Reed Watson

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Daniel McKinsey, Chair

Cold dark matter is a well motivated astrophysical phenomenon which is currently unexplained by the standard model of particle physics. Presently, time projection chambers such as LUX-ZEPLIN are setting world leading limits in certain mass ranges, and possible improvements for next generation experiments are being researched. Scaling up the size of experiments is one way to improve sensitivity, but poses certain challenges from a high voltage perspective. In this dissertation I detail my contributions to the LZ dark matter direct detection experiment's understanding of its electric fields and the impacts of high voltage on the future of TPCs in general. In Chapter 3 the testing and installation of LZ's novel cathode high voltage design are detailed, and in Chapters 4,6 the electric fields simulations for the LZ TPC are documented. The sources of accidental coincidence backgrounds are examined in Chapter 7, and results of a dielectric breakdown study of liquid xenon are reported in Chapter 8. Finally the results of an extension to the LZ dark matter search to be sensitive to Planck-scale masses are reported in Chapter 5. Throughout these work the electric field and grid voltages feature prominently, demonstrating their criticality to future generations of dark matter searches.

I dedicate this thesis to my parents, Donald and Jacinda Watson, without whom I would not have the confidence I have today.

Contents

Contents	iv
List of Figures	vii
List of Tables	xiv
1 Dark Matter	1
1.1 Introduction	1
1.2 Evidence for Dark Matter	2
1.3 Dark Matter Models	9
1.4 Experiments	15
2 The LUX-ZEPLIN Experiment	20
2.1 Overview	20
2.2 Time Projection Chambers	21
2.3 Xenon	23
2.4 Backgrounds and Mitigation	26
2.5 Veto System	28
2.6 Calibrations	29
2.7 Analysis	32
2.8 Physics searches	35
3 Cathode High Voltage for LZ	38
3.1 Introduction	38
3.2 Design Considerations	39
3.3 LAr Prototype Test	43
3.4 Prototype Testing Results	67
3.5 Deployment	77
3.6 Conclusion	83
4 Electric Field Simulations for LZ	84
4.1 Introduction	84
4.2 Finite Element Simulation	84
4.3 LZElectricField simulations package	93

4.4	Comparisons to SR1 data	107
4.5	Handoff	113
5	Multiple Scatter Dark Matter Search	118
5.1	Introduction	118
5.2	Models	119
5.3	Backgrounds	124
5.4	Upper Bounds	126
5.5	Projected Limits	129
5.6	Simulation	131
5.7	Testing on LUX Run03 Data	142
5.8	Analysis on LZ SR1 Data	144
5.9	Summary	158
6	Simulations	160
6.1	Introduction	160
6.2	DMCalc	160
6.3	LZLAMA	165
6.4	BACCARAT	170
7	Accidental Coincidence Backgrounds in LZ	177
7.1	Introduction	177
7.2	Sources of Isolated Pulses	179
7.3	Analysis of SR1 Data	185
7.4	Isolated S1 Locations	201
7.5	Creation of accidentals model	203
7.6	Grid Voltage Scan	208
7.7	Machine Learning	216
7.8	Single Scatter Contamination	222
7.9	Summary	224
8	The Xenon Breakdown Apparatus	225
8.1	Introduction	225
8.2	Dielectric Breakdown Theory	225
8.3	Apparatus	232
8.4	Procedure	248
8.5	Run History	251
8.6	Analysis	254
8.7	Results	261
8.8	Summary	279
	Conclusion	282
	References	283

List of Figures

1.1	Isotopic abundances predicted by Big Bang Nucleosynthesis.	4
1.2	The Cosmic Microwave Background Radiation temperature anisotropy spectrum.	8
1.3	The matter power spectrum, inferred from a variety of observations at different distance scales.	10
1.4	Current exclusion limits for primordial black holes in mass and fraction of dark matter density.	14
1.5	Differential recoil rates for assorted WIMP models with a WIMP-nucleon cross section of 1 zb.	18
1.6	The lab frame velocity probability distribution function of the standard halo model	19
2.1	The cross section and primary operating principle of LUX-ZEPLIN’s time projection chamber.	22
2.2	$\log_{10}(S2c)$ vs. $S1c$ areas for different calibration sources, demonstrating particle identification.	28
2.3	CAD rendering of the LZ water tank and detector.	30
2.4	A Doke plot, taken from [83], which determined $W = 11.5$ eV.	31
2.5	2D histogram of reconstructed event positions, S2 radius vs. drift time, from an early ^{83m}Kr injection.	32
2.6	The nuclear recoil analysis acceptance curve for LZ’s first science run.	34
2.7	The LZ spin-independent WIMP-nucleon exclusion limit.	36
3.1	A side-view cross section of the LZ Cathode High Voltage (CHV) feedthrough.	40
3.2	Photographs of examples of traditional feedthrough designs for dual-phase TPCs.	41
3.3	CAD rendering of the LZ water tank with components relevant to the CHV feedthrough annotated.	42
3.4	The results of a finite element field simulation in Maxwell.	43
3.5	An annotated diagram of the spool can connection / warm feedthrough at the top of the water tank.	44
3.6	A photograph of Harwin pin which connects the resistors in series between the CHV grading rings.	45
3.7	A photograph of the CHV prototype test stand.	47
3.8	Schematic of the charge sensing circuit for the CHV feedthrough test.	47
3.9	The schematic for the base for the CHV photomultiplier tube.	48
3.10	CAD rendering of high voltage test structure and surrounding instrumentation.	49

3.11	The schematic of the charge sensing circuit for the CHV prototype test.	50
3.12	Results of SPICE simulations using the schematic in Fig. 3.11.	51
3.13	The noise power spectrum for Run 1 of the CHV Prototype tests.	52
3.14	A representative waveform from a CHV prototype test stand acquisition.	53
3.15	Waveform results from the CHV baseline subtraction technique.	54
3.16	The noise power spectrum (FFT) before and after the baseline subtraction technique.	55
3.17	A representative “chirp” waveform which was observed in the data.	55
3.18	The Hamamatsu PMT impulse response, obtained from CHV prototype waveforms.	56
3.19	The optimal FIR filter for the CHV PMT response.	57
3.20	An example CHV prototype test event waveform with the zero-suppressed Richardson-Lucy deconvolution algorithm applied.	59
3.21	An event waveform classified under criteria (4) and (5), showing all three channels: PMT and both charge amplifiers.	61
3.22	The CHV purity monitor charge amplifier impulse response.	62
3.23	The Pulse area spectrum for the two PMT LED calibrations.	63
3.24	Overview of the LED background subtraction procedure.	64
3.25	The R5912-02-MOD gain as a function of applied cathode voltage.	65
3.26	The PMT gain as a function of PMT temperature.	65
3.27	The measured oxygen-equivalent impurity level, obtained via the Argon triplet lifetime, over the course of CHV Run 2.	66
3.28	A photograph of an oscilloscope from a test of the purity monitor performed while bypassing the feedthrough.	67
3.29	PMT (top) and charge amplifier(bottom) data from the first day of Run 1 high voltage ramps.	70
3.30	The waveforms from terminal event in the first day of CHV Run 1’s high voltage ramps.	71
3.31	PMT time series data from the CHV Prototype - May 25 high voltage ramp.	71
3.32	PMT data from the CHV - August 14 high voltage ramp.	73
3.33	PLC and PMT data from the August 14-15 overnight zero field data.	73
3.34	PMT time series data from the August 24 (Run 2) high voltage ramp.	74
3.35	Charge amplifier data from the CHV Prototype - August 24 high voltage ramp.	74
3.36	Waveforms from the terminal event in the CHV Prototype - August 24 high voltage ramp, defining the quiet time period for Run 2.	75
3.37	Construction photos of the LZ-CHV xenon displacer in the LBL clean room.	78
3.38	Photographs from the production grading structure construction process in the LBL clean room.	79
3.39	A photograph of the completed production grading structure in the LBL clean room.	80
3.40	A photograph of the cathode high voltage delivery crate. Aluminum panels lined the box.	81

3.41	Photographs taken during the cathode high voltage feedthrough final installation at SURF.	82
4.1	The LZ 2D axisymmetric finite element mesh, produced and rendered in GMSH.	86
4.2	The zoomed-in 2D mesh for the LZ cathode, gate, and anode rings.	89
4.3	The resulting simulated electric field for the LZ forward field region from Fenics.	90
4.4	Radial electric fields and equipotentials near the cathode region.	91
4.5	Volume-weighted histograms of simulated field magnitudes in the LZ TPC.	92
4.6	The field nonuniformity (RMS magnitude) as a function of cathode and gate voltage, averaged over the analysis and engineering fiducial volumes.	94
4.7	Illustration of the principle of barycentric coordinates. Each vertex is a primary color (red,green,blue).	99
4.8	Illustration of the “dead zones”/ charge-loss regions near the TPC wall.	101
4.9	The volume-weighted charge loss probability distribution function, used for validation of simulations against data.	102
4.10	The radial component of the electric field near the TPC wall boundary, with and without a particular charge distribution.	104
4.11	The equipotentials which trace out the cathode excursions.	105
4.12	Illustration of the principle of partially transparent grids, taken from Ref. [135].	106
4.13	The funneling and time delay effect of the woven mesh grid as a function of displacement from the middle of a cell.	107
4.14	The fraction of charge loss regions near the wall ($R = 728$ mm) as a function of drift field.	108
4.15	The results of the estimated charge density on the reconstructed location of LZ PTFE walls.	111
4.16	The estimated total surface charge accumulated on the TPC walls over time.	112
4.17	The estimated surface charge density profile for the calibration data.	113
4.18	The impact of a uniform wall charge on the field lines near the walls.	114
4.19	The charge loss volume vs uniform charge density.	115
4.20	Uncorrected S2 areas as a function of drift time with simulated perfect xenon purity.	116
4.21	Electron location scatter plots at the liquid level from grid funneling simulations.	117
5.1	WIMP-nucleon collisional cross section reduction as a result of coherence loss.	122
5.2	An example e-train event waveform taken from commissioning data using the LZ event viewer.	126
5.3	An example of a split-S1 event seen in the LZ event viewer.	127
5.4	Interstellar gas heating limits on contact interactions, figure from Ref. [168].	128
5.5	The extrapolated single scatter exclusion limits for LZ (orange) and Xenon1T (blue), along with best-case scenario, background free, 100% acceptance MIMP search for LZ	130
5.6	Energy attenuation as a function of incident angle, indicating The overburden or Earth shielding effect on different reduced cross sections.	136

5.7	An illustration of the geometry involved in the transits of ultramassive dark matter particles through the earth.	137
5.8	Earth shielding effect simulations with DMCalc.	138
5.9	The velocity distribution at surf after the attenuation of earth shielding.	140
5.10	The reduced cross section scaling of MIMP simulations.	141
5.11	The analysis efficiency as a function of altitude (angle to the horizon) for the turnaround model $\sigma_{\chi p} = 10^{-30} \text{ cm}^2$	143
5.12	Simulated horizon angles of SHM particles at SURF.	143
5.13	A simulated multiple scatter colinear even waveform in LUX, made using LUX-BACCARAT.	145
5.14	Unique reconstructed quantities for multiple scatter events, analyzed for the LUX simulations.	145
5.15	A histogram of simulated MIMP velocities for a model with $\sigma_{\chi p} = 10^{-30} \text{ cm}^2$ using LZLAMA.	150
5.16	A heat map of simulated S1 and S2 areas for the multiple scatter analysis.	151
5.17	Example reconstructed tracks / individual S2 locations from SR1.	151
5.18	S1-S2 distributions using various estimates for the vertex counts, indicating the convergence towards the mean value.	154
5.19	Pulse count predictions using the linear regression model.	155
5.20	MIMP search cut acceptances as a function of the number of vertices, evaluated on the turnaround model.	156
5.21	MIMP acceptance relative to the multiplicity requirement of 2+S1s, 2+ S2s as a function of incident velocity and entry/exit points.	156
5.22	The LZ SR1 MIMP search exclusion limits for A^4 (left) and contact interactions right.	157
5.23	LZ SR1 MIMP search exclusion limits plotted alongside the results from competing analyses for A^4 SI scaling.	158
5.24	Nugget search results.	159
6.1	WIMP recoil rates generated with DMCalc	162
6.2	Velocity distributions for various dark matter halo models, produced in DMCalc.	163
6.3	Relative modulation amplitudes for WIMP masses of 20 and 50 GeV/c^2 , as a function of recoil energy, calculated with DMCalc.	164
6.4	Krypton calibration single scatter S1-S2 distribution histogram.	169
6.5	Histograms of radial uncertainty from Mercury for Krypton charge loss events as a function of S2 area in the top PMT array.	170
6.6	Independent S2 position resolution vs. S2 area scaling models.	171
6.7	The radial and azimuthal resolution as a function of radius within the TPC, as estimated through the rotated S2 position covariance matrix of ^{83m}Kr events.	172
6.8	Simulations of the gaps in the PTFE walls.	173
6.9	The S1-only spectrum which results from events near the wall, including the newly simulated PTFE pockets	174
6.10	Code evaluation outputs from code profiling programs.	176

7.1	A histogram of total SPE rate preceding S1s in single scatters passing the livetime-impacting data quality cuts	181
7.2	A 2D histogram of top-bottom asymmetry and S2 area of cathode events within the WS ROI.	182
7.3	A 2D histogram of S1 and TBA-corrected S1 areas from an SR1 α skim.	183
7.4	The field fringing background spectrum as predicted by SR1 single scatters with $S2_R > 60$ cm.	184
7.5	The isolated S1 spectrum obtained from long random data.	186
7.6	Isolated S1 rate as a function of the lower bound in <i>Top</i> : S1 area, <i>Bottom</i> : coincidence threshold, with stacked analysis selection criteria.	187
7.7	The isolated S1 rate plotted over long timescales.	188
7.8	A 2D histogram of the maximum channel area vs. S1 area distribution of isolated S1 pulses, taken in the “long random” mode.	191
7.9	The drift time independent HSC cut development. Curves are shown for $\text{MaxChannelArea} < S1^{0.4} + 1$	192
7.10	Maximum channel area ratio vs. drift time histogram for ^{83m}Kr calibration, with the High Single Channel selection criterion indicated in dashed red.	193
7.11	A histogram of the individual channel areas for low coincidence HSC and non-HSC pulses.	194
7.12	Histogram of channel contributions to the HSC pulses. The lower indices before(after) the break at 300 are the top(bottom) array.	195
7.13	Analytic TBA vs. S1 area calculations, with confidence intervals calculated with several methods, indicated for $\hat{p} = 0.5$	197
7.14	The TBA-drift time cut efficiencies as a function of S1 pulse area for 3σ confidence interval selections.	198
7.15	S1 shape decision boundaries for the FWHM RQ, as a function of S1 area.	199
7.16	The S1 shape cut acceptance, as applied to the tritiated methane injection calibration dataset, and the iS1 pulses from Pre-SR1 data.	200
7.17	Histograms of PMT hit pattern spreads vs top-bottom-asymmetry, presented for top array (left) and bottom array(right), and for HSC pulses (top) and non-HSC pulses (bottom).	202
7.18	TBA-Area bands for HSC cut failing (left) and passing(right) isolated S1s from the Pre-SR1 dataset.	202
7.19	Heatmaps of the PMT array hit centroids for isolated S1s from the Pre-SR1 long random acquisition.	203
7.20	The total chopstitch cut acceptances as a function of corrected S2 area (left) and corrected S1 area (right).	207
7.21	Chopstitch 1D histograms for corrected S1 and S2 areas, which form the inputs to the outer product which becomes the complete “accidentals pdf.”	207
7.22	The chopstitch-accidentals pdf, before smoothing.	209
7.23	The smoothed accidentals PDF, with SR1 data overlaid.	209
7.24	Single Scatter drift time spectrum in LZ, indicating the difference between physical and unphysical drift times.	210

7.25	Differential iS1 rates vs corrected S1 area for various datasets over the WIMP search ROI.	213
7.26	The decay of the isolated S1 rate following particular alterations to the grid voltages.	214
7.27	Accidentals data isolated S1 distributions during a flare-up.	214
7.28	Isolated S1 hotspots in two different grid configurations.	215
7.29	The reweighted S1 spectrum which provides the negative samples for the boosted decision trees.	216
7.30	Boosting stages for the BDT training.	219
7.31	Receiver-operating characteristic for the pulse shape cuts.	220
7.32	Event acceptance vs S1 area for the BDT and RF models.	220
7.33	WIMP mass results. A smaller simulated WIMP mass results in more large-area iS1s being removed, as expected.	221
7.34	Scan over simulated WIMP masses for the BDT selection criterion.	221
7.35	Feature importance values for the isolated S1 classifying boosted decision tree. .	223
7.36	Accidentals single scatter contamination	224
8.1	Paschen-Townsend law / predicted breakdown curves in xenon vs pressure-separation distance.	227
8.2	Piping and instrumentation (P&ID) of the XeBrA experiment.	235
8.3	The XeBrA test stand with the OCV in place.	236
8.4	The XeBrA Rogowski electrodes, in-situ.	237
8.5	A CAD rendering of the XeBrA ICV cross section.	238
8.6	A simulated electric field profile for the XeBrA electrodes.	239
8.7	A typical refrigeration cycle plot in pressure-enthalpy space.	240
8.8	simulated stressed electrode areas vs. cathode-anode distances for each of the tilts which were estimated from the photogrammetric method.	242
8.9	The schematic for the XeBrA charge sensing circuit, produced in LTSpice. . . .	244
8.10	A photograph of the XeBrA purity monitor field cage.	246
8.11	Three distinct examples of streamers observed using the high speed cameras installed in XeBrA.	247
8.12	An example time series of measurements during a subcooling cycle in Run 7. . .	251
8.13	The distribution of the time differences between breakdowns and their associated discrete voltage steps.	255
8.14	An analytical example of combining two hazard functions.	257
8.15	High voltage power supply voltage vs. time, with the identified breakdowns for a representative dataset	258
8.16	An example of a breakdown distribution from a dataset that strongly favors a two component, three-parameter Weibull PDF	259
8.17	Various categories of event acquisition waveforms in XeBrA.	262
8.18	The results of the dechirping procedure on a typical anode signal event waveform.	263
8.19	The stressed electrode area scaling relationship inferred from XeBrA.	264
8.20	XeBrA pressure scan data.	266
8.21	XeBrA estimated breakdown voltages vs. ramp rate scan results.	266

8.22	Fowler-Nordheim plots from a run of XeBrA.	269
8.23	Pulse rate evolution over the course of a single ramp.	271
8.24	A precursor event window. Dotted lines indicate locations of identified pulses.	272
8.25	Distributions of anode current preceding a breakdown for different parameter scans in Run 10, over time.	273
8.26	Comparison of anode current and SiPM rate in the 30 seconds preceding a breakdown between Runs 9 (orange, mechanically polished electrodes) and Run 10 (blue, passivated electrodes).	274
8.27	Pulse height distribution for Run 9, demonstrating frequent DAQ saturation.	275
8.28	Average current between 30 s and 120 s from the beginning of the HV ramps for each breakdown in Run 9.	275
8.29	Anode pulse Rate vs. stressed electrode area over several runs.	276
8.30	A representative PMT ramp and hold pattern.	276
8.31	A series of high-framerate photographs of a breakdown containing two apparent discharges. Four frames are shown in sequential	278
8.32	The XeBrA electrodes after hundreds of breakdowns, showing damage in the form of white pits.	279

List of Tables

3.1	Test results for the CHV feedthrough Runs 1 and 2. The “highest voltage” column is the actual recorded voltage from the HVPS, while “target voltage” indicates the voltage necessary to achieve -100 kV on the LZ cathode, when correcting for surface area and purity.	76
4.1	Parameters used in the LZ electrostatics simulation. The electrode biases are a parameter which was explored as part of this work.	88
7.1	Summary of the contributions to Isolated S1 pulses in LZ.	184
7.2	S1-only shape selection criteria; parameters for Eq. 7.10.	198
7.3	Results of the electrode grid scan as applied to the isolated S1 rates. All rates are presented after the relevant cuts, and potential changes to the ROI and extra exclusion regions are tabulated. A general trend towards lower values was seen over the course of the test.	212
8.1	Estimated cathode tilt angles for different Runs. The front and side tilt labels refer to the tilts of the electrodes in the image plane for images taken from the side and front viewports, respectively, and the total value is obtained from adding those angles in quadrature.	242
8.2	XeBrA power law fit parameters to $E_2(A) = C(A/\text{cm}^2)^n$. C is a multiplicative constant, and $-b$ is the coefficient of the SEA in the power-law scaling.	261
8.3	Purity measurements taken during the 2020–21 data-taking runs. A general trend of lower purities are observed between consecutive runs. For Runs 5–7 the purity was inferred from a cold trap method similar to the one described in Ref. [241]. The discrepancy in uncertainties is due to the averaging of the waveforms and inconsistent flash rates. In certain datasets for unknown reasons the purity monitor was able to maintain bias for longer than others, allowing for longer integration periods. This introduces a bias towards shorter lifetimes / higher impurities due to the capacitor in series with the anode requiring recharging.	267
1	Calibrated Sensor gains for each run.	309
2	XeBrA dataset results for the Weibull fits.	310

Acknowledgments

Throughout my studies I have enjoyed the support of many talented and kind individuals. They have provided me with the emotional and academic resources without which I could not have completed the challenges over the course of these six years. Those who assisted me in any capacity are innumerable, but the confidence that certain individuals had in me has remained persistent on my mind.

Firstly, I want to acknowledge my advisor, Professor Daniel McKinsey. He offered me an opportunity to begin research immediately, which allowed me to make rapid progress towards my degree. I have been supported as a graduate student researcher since early on in my education, which allowed me to participate in LZ projects despite not being on a fellowship. Professor McKinsey has given me helpful guidance throughout my education, while allowing me the freedom to experiment and choose my own directions.

My graduate career began under the guidance of Ethan Bernard. His wide range of expertise in matters of hardware can not be overstated. Coming from a background of entirely software based projects, the experimental portion of research was intimidating in the first year of graduate school. He helped me to develop into a real asset for our group, and was always available to answer my questions on any relevant topic. The late nights in the LBL clean room were tough, but Ethan's stoicism was incredibly inspiring to me.

I want to thank my excellent colleague Quentin Riffard. While many of the postdocs, scientists, and students that I have had the pleasure of working with have influenced and inspired me, Quentin stands out in particular. His straightforward honesty helped me understand aspects of scientific collaboration and career advancement. From his excellent mentorship I felt confident in directing efforts towards areas of research with long term pay-offs. He helped me to truly believe that I was making important and unique contributions, and that my value extended beyond them.

My parents, Donald and Jacinda Watson, deserve immense gratitude. They supported me throughout my undergraduate degree, which allowed me to enter into the University of California at Berkeley in the first place. During times of self doubt their pride allowed me to visualize holding my future Ph.D.

Finally, my longtime friends Tyler and Victoria Brehm have supported me from before graduate school even began. They have known me from a young age, and they entered medical school together. I was able to confide my struggles with them, knowing that they understood the convergence academic and personal aspects which contributed to them. Their receipt of their respective medical doctorates filled me with happiness as their friend, and I hope my graduation instills the same feeling with them.

Chapter 1

Dark Matter

1.1 Introduction

The standard model(SM) of particle physics has been extensively tested by a number of experiments. It is a quantum field theory consisting of three generations of quarks and leptons, two unbroken gauge symmetries $SU_C(3) \times U_Y(1)$ and a $SU_L(2)$ symmetry which is spontaneously broken by the scalar Higgs field. Using the SM, a large number of phenomena can be accurately predicted, such as the scattering amplitudes and decay widths of the myriad of existing particles. Some observations are not explained by the SM. For instance, the standard model can not account for neutrino masses due to the apparent non-existence of right handed neutrinos, necessary for the Dirac mass term. In another case, astrophysical and cosmological observations point towards the existence of copious amounts of “cold dark matter” (CDM), a gravitationally interacting substance with little to no coupling to the electromagnetic field[1]. The SM contains no viable candidates, and therefore CDM is presently one of the strongest pieces of evidence for beyond the standard model (BSM) physics.

The following sections will outline the outstanding evidence for CDM, and I will briefly explain the case for it being an extension to the standard model. Observations across a wide range of distance scales indicate that a large component of the mass-energy density of the universe is in the form of a nonrelativistic, pressureless matter field feebly coupled to the electromagnetic field. The only neutral SM particle stable on cosmological time scales is the neutrino. Neutrinos which were thermally produced in the early universe are referred to as relics, and would be nonrelativistic today. However, in order to explain the virial mass of galaxies, the maximum number density of dark matter particles possible from Fermi-Dirac statistics requires leptons with masses larger than $\approx 1 \text{ MeV}/c^2$ [2]. This is sometimes referred to as the Tremaine-Gunn bound[2]. Such a requirement excludes the upper limit on the neutrino masses set by the Planck collaboration ($\sum m_\nu < 0.12 \text{ eV}$)[3]. Dim celestial bodies such as primordial black holes (PBHs) or brown dwarfs are another logical choice, however much of the mass range has been excluded from being the *entire* dark matter density with microlensing and gravitational wave searches[4]. An exception to this exists for PBHs of approximately asteroid mass ($\sim 10^{18} \text{ kg}$).

An extension to the SM is therefore required to account for gravitational observations. Either gravitational theory must be modified under certain circumstances, or a new, practically invisible particle must exist. If CDM is massive particle, it is feebly coupled to the SM through some portal which also accounts for its initial production in the early universe. It is through this small interaction with ordinary particles that experiments aim to discover its nature. An overview of possible explanations for the required dark matter density and the methods for detection are presented below.

1.2 Evidence for Dark Matter

Galaxy Clusters

In 1933 Fritz Zwicky studied the velocities of galaxies within the Coma cluster. These galaxies exhibited an abnormal amount of dispersion in recessional velocity compared to other clusters[5]. By estimating the distance to and radius of the Coma cluster, an order of magnitude estimate of the matter density was obtained. Applying the virial theorem:

$$T = -\frac{1}{2}V, \quad (1.1)$$

where T is the average kinetic energy, and V is the average potential energy, a prediction for the velocity dispersion was found to be 80 km/s, whereas the measured velocity dispersion was 1000 km/s. The increased velocity dispersion requires a density of luminous matter over 400 times larger than that estimated from observations. This calculation was enough for Zwicky to note that the amount of “dark matter” far exceeded that of luminous matter. These estimates were later refined, but the field remained largely skeptical of the dark matter explanation. A popular explanation at the time was the existence of an intergalactic medium (IGM), which existing techniques were unable to detect. Modern studies indicate that the intergalactic medium contributes 40-50% of the baryonic matter of the universe[6].

Galactic Rotation Curves

Until the 1970s it was expected that the rotation curves of spiral galaxies asymptotically approached Kepler’s third law ($(r\omega)^2 = GM/r$) as the distance of a particular star from the galactic core (i.e. the majority of the baryonic mass), r , increased. Deviations from Keplerian velocities were noticed fairly early but not fully understood until Rubin and collaborators utilized image tubes to collect data on hundreds of rotation curves, observing the Doppler shift in the 21 cm line. The observed luminosities provided fairly accurate predictions for the velocities near the galactic nuclei, but they flattened, or even increased as a function of radius. This is indicative of the vast majority of the mass of a galaxy coming in the form of a non-luminous “halo” which extends far past the visible edge of the galaxy. Since then, more modern technology has afforded surveys out to higher z and r , confirming the initial results.

While not evidence for dark matter, an important result for understanding spiral galaxy rotation curves is the “Tully-Fischer” relationship. This is a power-law relation between the luminosity of a galaxy and the asymptotic rotation velocity. In the dark matter paradigm, this is interpreted as suggesting a proportionality (on average) between the amount of baryonic matter and the amount of dark matter in a galaxy. However, it can also be interpreted as evidence for modified gravity, since the dynamics of the rotation curves could be predicted by a more complicated theory.

Bullet Cluster

In 2006 what was referred to at the time as the “smoking gun” evidence for dark matter was observed[7]. Galaxy cluster 1E 0657-56, sometimes referred to as the “Bullet Cluster” consists of two apparent subclusters in the process of merging. The bow shock in the intergalactic medium produces copious amounts of x-rays and affords estimation of the relative velocity of the subclusters. The center of gravity, measured by the strong and weak lensing, was observed to be offset from the “center of mass,” measured by the x-rays and optical observations. This observation has given great credence to the particle dark matter theory, which explains that the bulk of the gravitating mass was in the form of matter which did not have a strong self-interaction cross section. In other words, the dark matter, making up the bulk of the mass in either subcluster, passes uneventfully through itself during the collision, while the baryonic matter experiences drag and slows down. Limits were subsequently set on the dark matter self interaction cross section[8] at

$$\frac{\sigma_{\chi}^2}{M_{\chi}} < 1.25 \frac{\text{cm}^2}{\text{g}} . \quad (1.2)$$

Big Bang Nucleosynthesis

Light elements were formed during the early universe in a short time window between when the temperature dropped below the 2.2 MeV dissociation temperature of the deuteron and the 14 minute lifetime of the neutron. In that time, nearly all of the neutrons and protons are fused into deuterons, and those deuterons then nearly completely fuse into ^4He , which yields a mass fraction $Y \approx 0.25$, consistent with observation. The primordial value of the deuterium-to-proton ratio D/H from the leftover deuterons is extremely sensitive to the baryon-to-photon ratio η . Elements including ^3He , ^6Li , and ^7Li also form, allowing for independent measurements, as shown in Figure 1.1. The densities today are parameterized today in terms of the ratio to the critical density, where $\Omega_b = \rho/\rho_c$ is the normalized density of baryons. This process of *Big Bang Nucleosynthesis*[9](BBN) estimates a value of $\Omega_b h^2 = 0.022305 \pm 0.000225$, or $\approx 4\%$ of the critical density. Here $H_0 = h \times 100 \text{ km} \cdot \text{s}^{-1} \text{Mpc}^{-1}$, absorbing the uncertainty on the Hubble constant into h . A flat universe therefore requires some missing component to constitute the remainder of the critical density $\Omega_0 h^2$, suggesting the existence of a dark matter component.




Figure 1.1: Isotopic abundances predicted by BBN, figure from [10]. The CMB bounds are consistent with observations, validating BBN. There is a noticeable discrepancy with the ^7Li abundance, known as the “Cosmological Lithium-7 Problem” [11].

Cosmic Microwave Background Radiation

Big Bang Cosmology

Under the assumptions of homogeneity and isotropy, an expansion of an N -dimensional space can be modelled as a surface of constant curvature k embedded within an $N + 1$ -dimensional space. The curvature k is related to the difference between the circumference of a circle in the curved space and a circle of identical radius in Euclidean geometry. As the universe expands, the distance between two points $R(t)$ increases with time. This expansion is factored out into a scale factor a and comoving distance r , $R(t) = a(t)r$, with $a_0 = a(t_0) = 1$ being the scale factor today. With these assumptions, one obtains the FLRW metric:

$$ds^2 = c^2 dt^2 - a(t)^2 \left[\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right], \quad (1.3)$$

with $d\Omega^2 = d\theta^2 + \sin^2 \theta d\phi^2$ being the solid angle. The curvature can be reparameterized such that the available options are $k = \{-1, 0, 1\}$. By inserting this equation into the Einstein field equations from general relativity, the Friedmann equations are obtained, which describe the evolution of the scale parameter a depending on the pressure p and density ρ of the universe.

$$\left(\frac{\dot{a}}{a}\right)^2 + \frac{kc^2}{a^2} = \frac{8\pi}{3}G\rho \quad (1.4)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi}{3c^2}G(\rho c^2 + 3p), \quad (1.5)$$

where G is the gravitational constant, ρ is the energy density, and p is the pressure. In the absence of curvature $k = 0$, we can define the Hubble expansion rate $H \equiv \dot{a}/a$, as well as the critical density ρ_c as the density which provides the Hubble rate today H_0 :

$$\rho_c = \frac{3H_0^2}{8\pi G}. \quad (1.6)$$

Under-or-over densities $\Omega = \rho/\rho_c \neq 1$ thereby correspond to changes in the curvature parameter k . Inserting a particular equation of state, or relationship between pressure and density, the Friedmann equations can be solved for $a(t)$. In the early universe radiation was initially dominant, with $p = \rho c^2/3$, followed by a period of matter domination, with $p = 0$. These lead to expansion of $a \sim t^{1/2}$ and $a \sim t^{2/3}$, respectively.

As the universe expanded, the radiation density fell as a^{-4} , whereas the energy density of the matter fell like a^{-3} . Two important events happen as a consequence of this. First, the density of radiation and matter become equal at $z_{\gamma m}$. Then, at a later time the temperature of the photon-baryon plasma drops low enough that neutral hydrogen can form, referred to as “recombination.” At this point in time the universe became quasi-transparent to light, and now the surface of last scattering can be seen in the cosmic microwave background radiation (CMBR).

While subdominant in the early history of the universe, it is worth mentioning that an additional cosmological constant term may be added to the right hand side of the Friedmann equations, turning them into the following:

$$\left(\frac{\dot{a}}{a}\right)^2 + \frac{kc^2}{a^2} = \frac{8\pi}{3}G\rho + \frac{\Lambda c^2}{3} \quad (1.7)$$

$$\frac{\ddot{a}}{a} = -\frac{4\pi}{3c^2}G(\rho c^2 + 3p) + \frac{\Lambda c^2}{3} . \quad (1.8)$$

This term provides a constant positive energy density and negative pressure. In a flat universe $k = 0$, when Ω_Λ dominates the matter and radiation terms Ω_m and Ω_r , this leads to exponential expansion of the scale factor $a(t)$.

CMB Anisotropies

The CMBR is a near-perfect Planck spectrum at 2.7 K, uniform to 10^{-4} . The small observed anisotropy is dominated by the dipole moment, corresponding to the Doppler shift of the relative motion of the earth to the CMB. After subtracting off the dipole component, and expanding into spherical harmonics, a distinct pattern of peaks and troughs is observed.

Initial small density perturbations δ grow until they cross the *sound horizon*, and oscillate thereafter. The acoustic peaks thereby represent an alternating sequence of compression and rarefaction peaks which are frozen in time at the “moment” of recombination. Smaller peaks (larger multipole l) entered their horizons and started oscillating earlier than larger peaks (smaller l).

The standard model of cosmology, Λ CDM, which includes dark energy and dark matter, explains the salient features of the acoustic peaks, shown in Fig. 1.2 While the overall model contains some degeneracies, some observables are explained relatively succinctly:

- The location of the first compression peak at $l \sim 200$. This is most sensitive to the curvature parameter and is consistent with a flat spacetime.
- The relative height of the second (rarefaction) peak. Its relative amplitude is caused by an effect known as *baryon loading*. Because the photon-baryon system started at rest, the inertia of the baryons causes the equilibrium around which the oscillations occur to shift. This causes the compression peaks to be taller than the rarefaction peaks. This effect is predicated on the presence of potential wells for the baryons to fall into. In a radiation dominated universe, the baryon-photon plasma forms these potential wells, which decay away as a^4 , which removes the loading effect. Baryons also contribute to an effect called *Silk damping*, where the finite scattering length as the plasma recombines results in the higher peaks being smoothed out relative to the lower peaks. This effect is therefore degenerate with $\Omega_m h^2$ up to two peaks, needing a third peak to determine with certainty.

- The height of the third peak relative to the second peak. This is strong evidence for dark matter[12]. As the radiation density decays away faster than the matter density, oscillations receive a strong driving force when matter-radiation equality occurs at z_{eq} . Overdensities which enter the the sound horizon during radiation domination result in acoustic peaks $\sim 5\times$ higher than those entering the horizon at matter domination. The density of matter $\Omega_m h^2$ shifts the location of radiation-matter equality, with higher densities pushing the effect to higher multipoles. By acting as a replacement potential well, dark matter maintains the baryon loading effect, as mentioned above. Radiation-matter equality occurred slightly before recombination, with $z_m \approx 3000$ and $z_{CMB} = 1100$. Thus, the fact that the third peak is boosted relative to the second is strong evidence for dark matter.
- The damping tail. The extreme multipoles are suppressed due to the Compton scattering that occurs during the time between recombination and full transparency. This provides an independent cross check of the baryon density σ_b .

The observed heights of the acoustic peaks require additional gravitational effects in excess of what the baryons themselves can provide, confirming the existence of CDM. The *Planck* survey[3] provides a high precision estimation of the CMB anisotropy. When combined with other constraints on the cosmological parameters (such as Big Bang Nucleosynthesis), it favors a flat universe with cold dark matter density $\Omega_c = 0.267$, with $\Omega_\Lambda = 0.7$ and Ω_b taking up the remaining 0.04 of the critical density.

Large Scale Structure

The relative uniformity of the CMB raises an important question, namely how the small perturbations can grow to form large scale structures such as galaxies and galaxy clusters. In other words, why nonlinear density and structures exist with $\delta_\rho > 1$, rather than modest perturbations $\delta_\rho \ll 1$ on top of a uniform matter field. Features smaller than the mean free path of a photon will be damped during recombination, and the Hubble term prevents further collapse. CDM resolves this by providing additional mass to continue collapsing into the nonlinear regime, leading to galaxies today.

The theory of gravitational collapse is described by the Jean's instability equation. Assuming that the fluid can be described by a collisionless Boltzmann equation, the density perturbations $\delta = \rho - \rho_0$ evolves according to a wave equation:

$$\frac{\partial^2 \delta}{\partial t^2} + 2H(t) \frac{\partial \delta}{\partial t} + \delta(k^2 c_s^2 - 4\pi G \rho_0) = 0 . \quad (1.9)$$

Thus, solutions with at sufficiently small k (i.e. large λ) result in solutions with $\omega^2 < 0$, leading to exponential growth or decay. This defines the Jeans length, and therefore the




Figure 1.2: The Cosmic Microwave Background Radiation temperature anisotropy spectrum. Figure from Ref. [3]. Note the relative height of the second and third peaks, which indicates baryon loading in the presence of gravitational wells formed by cold dark matter $\Omega_m h^2$.

Jean’s mass, above which the overdensities δ will grow over time.

$$\lambda_J = c_s \left(\frac{\pi}{G\rho_0} \right)^{1/2} \quad (1.10)$$

$$M_J = \frac{4\pi}{3} \left(\frac{\lambda_J}{2} \right)^3. \quad (1.11)$$

Crucially, since ρ_0 diminishes as the scale factor $a(t)$ increases, the Jeans mass increases over time during radiation domination ($c_s = c/3$ during this time). Thus, overdensities will grow until until M_J overtakes them, or equivalently when they enter the sound horizon. Larger fluctuations will grow for a longer amount of time than smaller fluctuations. Between matter-radiation equality and recombination, photon diffusion will damp out smaller fluctuations. When recombination finally occurs, the speed of sound rapidly drops, allowing for smaller fluctuations to once again grow. However, by this point larger k have been suppressed, making galaxy scale structure formation impossible. The Jean’s mass in this time frame evolves as the following:

$$M_J \approx 8 \times 10^{30} (1+z)^{-3} \Omega_b h^2 M_\odot. \quad (1.12)$$

Nonbaryonic matter solves this problem by virtue of being able to grow in the epochs between matter-radiation equality and recombination. Subhorizon fluctuations are suppressed before matter domination due to the effective “friction” term resulting from the Hubble expansion

$H = \dot{a}/a$. After matter domination the δ will grow at all scales, and after recombination the baryons will fall into the potential wells formed by the cold dark matter, leading to the growth of galaxy-scale structures. Following recombination M_J contracts further due to $c_s \propto T^{1/2}$, causing features of diminishing size to collapse. This “bottom-up” formation has observable consequences for the way matter is distributed on supergalactic scales.

Overall, the results are encapsulated in the matter power spectrum $P(k)$, the Fourier transform of the two point correlation function. The halting of growth between t_H and $t_{\gamma m}$ leads to a power law suppression for $k > 1/x_h(t_{\gamma m})$. Primordial fluctuations $P(k) \sim k$ are convolved on top of this spectrum, leading to a particular distribution which is confirmed in CMB anisotropies, galaxy clustering, and Lyman- α forests, shown in Fig. 1.3.

Overall, cold dark matter leads to a bottom-up, hierarchical distribution of structure in the universe. Interestingly, warm dark matter models, which describe particles which decouple while relativistic while being nonrelativistic in the present day, lead to predictable distortions in the power spectrum above the free-streaming scale of the dark matter particles. Particle masses above ~ 1 keV remain consistent with large scale structure formation while smoothing out extremely small features. The latter property is desirable since it resolves the “core-cusp” problem, a discrepancy between simulated and observed density profiles in galactic nuclei.

1.3 Dark Matter Models

Thermal Relics

Despite the necessary feeble coupling between dark matter and baryonic matter, some production mechanism in the early universe is assumed to be necessary. The simplest case is a $2 \rightarrow 2$ interaction between dark matter particles χ and standard model particles f , with some velocity-dependent cross section $\sigma(v)$. Combined with the additional assumption that the dark matter was initially in thermal equilibrium with the baryonic matter, one can predict the density of dark matter in the present epoch. Later, the annihilation process $2\chi \rightarrow 2SM$ freezes out when the reaction rate drops below the Hubble expansion rate. This sequence of events can be described by the Boltzmann equation:

$$\frac{dY}{dt} = T^3 \langle \sigma v \rangle [Y_{EQ}^2 - Y^2] , \quad (1.13)$$

where $Y \equiv n/T^3$ is the dimensionless “yield” of the particles in question, which is proportional to the ratio n_χ/n_γ when the universe was radiation dominated. The quantity Y_{EQ} is the equilibrium value of the yield at a particular point in time. Replacing the time coordinate with $x \equiv m/T$, and assuming that the process of freezing out occurs during radiation domination $\rho \propto T^4$, the Boltzmann equation becomes:




Figure 1.3: The matter power spectrum, inferred from a variety of observations at different distance scales. Figure taken from Ref. [13]. The slope of the tail at large wavenumber is indicative of cold dark matter. Relativistic or “warm” dark matter would result in larger free-streaming distances, resulting in increased attenuation of the matter power spectrum at smaller scales.

$$\frac{dY}{dx} = \frac{\lambda}{x^2} (Y_{EQ}^2 - Y^2) \quad (1.14)$$

$$\lambda \equiv \frac{m^3 \langle \sigma v \rangle}{H(x=1)}. \quad (1.15)$$

The new dimensionless parameter λ describes the annihilation rate relative to the universe’s expansion rate. As the universe expands and cools $a \propto 1/T$, the thermal equilibrium number density falls rapidly. For a nonrelativistic species the scaling relationship is set by the Boltzmann factor $\exp(-m/T)$. When $x \ll 1$, Y can track Y_{EQ} because the species is able to self-annihilate quickly. However, when $x > 1$, the annihilation rate drops, and eventually

$Y \gg Y_{EQ}$, where the species is referred to as being “frozen out.” The time freezeout occurs, x_f , is given approximately by equating the annihilation and Hubble expansion rates.

$$n_{EQ} \langle \sigma v \rangle = H(x_f) . \quad (1.16)$$

The solution to this equation depends on m and the velocity dependence of σ . For s -wave annihilation, σ is constant, and $x_f \approx 20$, yielding an approximate freezeout velocity of $v \approx c/4$.

The equation 1.15 lacks an analytic solution, necessitating some approximations. Assuming $Y \gg Y_{EQ}$ at the time of freezeout, and that the velocity dependence of the annihilation cross section can be written as $\lambda = \lambda_0 x^{-m}$, one obtains:

$$\frac{dY}{dx} = -\frac{\lambda_0}{x^{2+n}} Y^2 \quad (1.17)$$

$$Y_\infty = \frac{n+1}{\lambda_0} x_f^{n+1} , \quad (1.18)$$

where in the second line the differential equation is integrated from $x = x_f$ to $x = \infty$ and the asymptotic yield is assumed to be much less than the yield at freezeout, $Y_\infty \ll Y_f$. Plugging in the mass m , present temperature T_0 , and taking into consideration the change in number of relativistic degrees of freedom g between freezeout and today, one obtains

$$\Omega_\chi = \frac{\rho_\chi}{\rho_c} = \frac{4\pi^3}{45} \frac{8\pi}{3H_0^2} \frac{x_f}{\langle \sigma v \rangle} \left(\frac{T_0}{M_{pl}} \right)^3 \frac{g_\star(0)}{\sqrt{g_\star(m)}} \quad (1.19)$$

$$\Omega h^2 \approx 0.2(n+1) \frac{x_f^{n+1}}{20} \frac{10^{-26} \text{ cm}^3/s}{\langle \sigma v \rangle} . \quad (1.20)$$

One particularly popular dark matter model is the weakly interacting massive particle, or WIMP. When applying a Weak interaction annihilation for a 100 GeV scale particle, one obtains $\sigma \sim G_F^2 m_\chi^2 \approx 10^{-39} \text{ cm}^2$. This coincidence, that the approximate dark matter density is obtained naturally from a weakly interacting particle, is referred to as the “WIMP Miracle” and served to motivate the field of dark matter direct detection towards optimizing for its discovery.

There are some critical boundaries to this analysis. When using the s -channel annihilation via the Z boson to standard model particles, the matrix element is proportional to m_χ^2 . Therefore, lighter masses will underannihilate, and account for more than the required dark matter density and overclose the universe. For a classical WIMP this occurs at 2–4 GeV and is referred to as the “Lee-Weinberg bound,” below which different theories are required. These theories are varied, but revolve around methods to enhance the annihilation cross section relative to the SM-DM scattering cross section. So the Lee-Weinberg bound is not a hard limit, but rather the expected edge of the “WIMP miracle”.

Regardless of the annihilation channel, quantum mechanical scattering theory is expected to hold. Using partial wave analysis, the maximum cross section that can be achieved at a given spin is

$$\sigma_J \leq \pi \frac{2J+1}{k^2} . \quad (1.21)$$

Since the relic density is inversely proportional to $\langle \sigma v \rangle$ at the time of freezeout, and freezeout occurs at nonrelativistic velocities, this sets a maximum possible cross section consistent with the observed relic density. Since $k \approx m_\chi v$, and v is constrained by the dynamics of the Boltzmann equation to be $\approx c/4$, this equivalently sets a maximum dark matter mass under s-wave ($J = 0$) scattering at [14]

$$M_\chi < 340 \text{ TeV} . \quad (1.22)$$

This limit is known as the “unitarity bound” and is widely applicable to thermal relics. This bound may be avoided in the case of non-thermal relics (where the assumptions in the above analysis no longer apply), or composite dark matter (where the mass at freezeout differs from the mass today).

MSSM

The standard model contains ultraviolet divergences which require large degrees of fine-tuning to cancel out to the Planck scale [15]. Supersymmetry solves this issue by placing all fundamental particles into either a gauge or chiral supermultiplet. Superpartners of standard model particles have differing spin, i.e. fermions are superpartners with bosons. This transformation leads to the desired cancellation as fermion loops present in the perturbative expansion are paired with boson loops of opposite sign. The winos and binos (the superpartners of the W and B bosons), along with the higgsino can mix, as the W and B do in the standard model, to form a new state, the neutralino $\tilde{\chi}_1^0$.

The lightest supersymmetric partner (LSP), which in the minimal supersymmetric extension to the standard model (MSSM) is a dark matter candidate under certain circumstances. MSSM by convention only adds terms to the Lagrangian necessary for the UV completion. To protect against baryon number and lepton number violation, a new discrete symmetry, *R-parity*, is introduced [16]. All particles have a new quantum number:

$$P_R = (-1)^{3(B-L)+2s} . \quad (1.23)$$

Every term in the Lagrangian must preserve P_R , which guarantees the stability of the LSP, a desirable feature for a dark matter candidate. Tunable parameters include the fraction of each superpartner in the neutralino state, *e.g.* mostly bino, wino, or higgsino. Variations such as the *Constrained* MSSM [17], have been extensively tested, and the surviving parameter space in that case is limited to the single TeV scale [18].

Axions and axion-like-particles

The Strong CP problem is the apparent fine tuning of the CP-violating terms in the QCD Lagrangian to nearly zero [19]:

$$\mathcal{L}_{CP} = \bar{\theta} \frac{g^2}{32\pi^2} G_{\mu\nu} \tilde{G}^{\mu\nu} , \quad (1.24)$$

where g is the QCD coupling constant, $G_{\mu\nu}$ The neutron electron dipole moment, experimentally constrained to be $|d_n| < 1.8 \times 10^{-26} \text{e} \cdot \text{cm}$ [20], is sensitive to the value of $\bar{\theta}$. Axions are a potential solution to this problem which promotes the $\bar{\theta}$ parameter to a dynamical variable, and introduces a global $U(1)$ symmetry[21]. The axial Peccei-Quinn symmetry is spontaneously broken, which results in the Goldstone modes acquiring a mass, becoming the pseudoscalar axions ϕ_a . The axion mass and couplings are constrained to $m_a f_a \approx (10^{10} \text{eV})^2$.

The axion is a dark matter candidate by virtue of the *misalignment mechanism*[22]. There, the initial value of the θ field is displaced from the minimum of the potential, and oscillations begin when $m_a \sim 3H$. The sub-eV mass scale of axions results in enormous number densities if it constitutes the entirety of the dark matter density. These searches are therefore optimized for more wavelike dark matter models, examining resonant production. Experiments such as ADMX[23] and DM-radio[24] search for resonances in superconducting solenoids via the inverse Primakoff effect.

When the strict relationship between the m_a and f_a is relaxed, the model no longer solves the strong CP problem, but could still be a viable dark matter model. In this case the particle is an ‘‘axion-like particle’’ (ALP). ALPs can couple to electrons and lead to an ER signal in direct detection experiments[25].

Sterile Neutrinos

Neutrinos decoupled from the thermal bath in the early universe, similar to the relic photons which constitute the Cosmic Microwave Background. While the relic neutrinos constitute some matter density as they have some minimum mass splitting $\Delta m_{31}^2 \approx 2.4^{-3} \text{eV}^2$ [26], they can not explain the full density due to Tremaine-Gunn bound[2]. This bound constrains the dark matter mass to $> 1 \text{ MeV}$ from the Fermi-Dirac statistics within the galactic halo. Thermal relic neutrinos have a density today[27]:

$$\Omega_\nu h^2 = \frac{\sum_i m_{\nu i}}{94.14 \text{ eV}}, \quad (1.25)$$

where $m_{\nu i}$ is the mass of the i -th neutrino eigenstate. The necessary DM density requires an $\approx 10 \text{ eV}$ scale neutrino, in tension with the CMB result[3]. Additionally such light states result in attenuation of large scale structure formation, inconsistent with observation.

Since all active neutrinos are left handed, neutrino masses require some extension of the standard model. Sterile, right handed neutrinos provide this through the seesaw mechanism[28]. In a type-I seesaw the left- and right-handed neutrinos mix, leading to a mass matrix which, when diagonalized, forces the active species to be light and the sterile species to be heavy. The sterile neutrinos would be produced via oscillation and would never be in thermal equilibrium (so-called ‘‘freeze-in’’).

A mass of keV scale leads to *Warm Dark Matter*, which is produced relativistic but is today nonrelativistic. It suppresses structure formation on distances smaller than its free streaming distance, which for 1 keV is approximately 1 Mpc. A 3.55 keV X-ray line from nearby galaxies has also been observed[29]. However, further analysis of X-ray lines has led to constraints on sterile neutrino dark matter[30].




Figure 1.4: Current exclusion limits for primordial black holes in mass and fraction of dark matter density. Figure from [31]. A small gap of masses is still allowed to constitute the entirety of the dark matter density.

Due to the feeble couplings to the standard model, direct detection via sterile neutrino scattering off of nucleons is not competitive.

Primordial Black Holes

Black holes which form in the early universe, as opposed to stellar collapse, are known as *Primordial Black Holes* (PBHs). The mass of the PBHs is proportional to the mass contained in the Hubble horizon at the time of formation[4], which is typically in the radiation-dominated era. PBH are an attractive dark matter model as they may be small enough to lack detectable accretion disks and (fortunately) large enough to only occasionally pass close to the earth. Stellar mass PBHs $M_{PBH} > M_{\odot}$ are excluded by 21 cm line analysis [31], gravitational waves, and CMB anisotropy. Microlensing searches which exclude MACHOs[32] force much of the mass parameter space to make up only a subcomponent of dark matter, usually less than 10^{-2} . PBHs of $M < 10^{17}$ g will have evaporated by now. If PBHs do constitute the entirety of the dark matter density $\Omega_{\chi} h^2 \approx 0.12$ then their mass must lie in the range $[10^{-16}, 10^{-12}]M_{\odot}$ [4].

Dark Sector

Weak-scale interactions, mediated through the Higgs or the Z-boson, places constraints such as the above Lee-Weinberg bound, limiting dark matter to masses above a few GeV. To escape such bounds a more flexible theory is required. These generically introduce a fermion or family of fermions χ_i to populate the dark sector, along with a *portal* which would couple the DS to the SM. A vector-like portal would be the result of a new, broken U(1) symmetry, which may kinetically mix with the SM photon:

$$\mathcal{L} \supset \frac{1}{2} m_{A'}^2 A'^\mu A'_\mu + \epsilon e A' J_\mu^{EM} , \quad (1.26)$$

where $m_{A'}$ is the mass of the dark photon, A'_μ is the dark photon field, ϵ is the mixing parameter, and J_μ^{EM} is the electromagnetic current. This would result in millicharged dark matter[33]. Scalar dark sector mediators have the following additional Yukawa interactions: terms[34, 35]:

$$\mathcal{L} \supset \frac{1}{2} m_\phi^2 \phi^2 + g_\chi \phi \bar{\chi} \chi + g_n \phi \bar{n} n + g_e \phi \bar{e} e . \quad (1.27)$$

The mediator mass becomes a new parameter of the model. With heavy mediators their mass is integrated out, and the cross section becomes dependent on g^2/m_{phi}^2 . For light mediators the interaction range is longer, and therefore a dependence on recoil momentum $F(q) \propto q^{-2}$ is introduced.

Ultraheavy Dark Matter

Thermal relic dark matter requires $\langle \sigma v_{rel} \rangle \approx 3 \times 10^{-27} \text{ cm}^3/\text{s}$, which sets a minimum necessary cross section at the freezeout velocity $x \approx 25 \rightarrow v_{rel}^2/4 \approx 1/16$. Partial wave expansion saturates for s-waves at a mass of[14]

$$m_\chi \leq 340 \text{ TeV} .$$

Circumventing this limit requires some additional effects. Sommerfeld enhancement can lead to enhanced annihilation which yields the correct density today. Freeze-in models[36] can generate large masses by not being subject to thermal freeze-out. Asymmetric models[37] achieve the correct density by assigning a baryon number to dark matter, allowing the matter-antimatter asymmetry to be communicated to the dark sector. If dark matter forms bound states from constituents[38] then the problem of achieving the correct relic density is decoupled from the mass of the dark matter state observed today. More details can be found in Chapter 5.

1.4 Experiments

Production

The first prong of dark matter searches involves the production process $2SM \rightarrow 2\chi$. The procedure involves colliders and accelerators creating dark matter particles, which then es-

cape the system. The analysis consists of reconstructing the momenta of the jets in the interaction, and searching for interactions with deficits in momenta. Collider limits are highly model-dependent, but are mostly sensitive to light (<10 GeV) masses with heavy mediators[39].

Indirect Detection

Another prong of dark matter searches is the indirect route, $2\chi \rightarrow 2SM$. While dark matter froze out soon after the big bang, regions of high density may still be able to annihilate dark matter at rates high enough to be observable, but low enough to still exist today. The signal would therefore be a peak or rise in the particle spectra at m_χ . This situation may arise in the dense core of galaxies, and the detected particles can in principle be any standard model particle lighter than m_χ . Fermi-LAT [40, 41] observes an excess of gamma rays from 1 GeV out to 100 GeV coming from the galactic center. The “galactic center excess” is not claimed as dark matter due to uncertainties related to emission components above and below the galactic center (a.k.a the *Fermi* bubbles). Limits have been set using galaxy-cluster γ -ray data which thermal relic dark matter as a source out to 100 GeV[42]. However, these excesses have not been confirmed and are in tension with blank-sky measurements[43].

Direct Detection

Recoil Rates

In a dark matter direct detection experiment, one is primarily interested in the number of events per unit mass of the detector medium. The initial assumptions are that the incident particles follow some velocity distribution $f(v)$, and that the cross section can be calculated in terms of the recoil energy E_r .

$$\frac{dR}{dE_r} = \frac{n_\chi}{m_T} \int dv v f(v) \frac{d\sigma_{\chi T}}{dE_r}, \quad (1.28)$$

where n_χ is the DM number density, m_T is the mass of the recoiling nucleus, v is the DM velocity, $\sigma_{\chi T}$ is the DM-nucleus cross section, and E_r is the kinetic energy of the recoiling nucleus. From here, we make an assumption of the form of $\frac{d\sigma_{\chi T}}{dE_r}$. Working in the nonrelativistic regime, the generic cross section is given by:

$$\frac{d\sigma}{d\Omega} \approx \frac{|\mathcal{M}|^2}{64\pi E_{\text{cm}}^2} = \frac{g^4}{M_\Lambda^4} \frac{m_\chi^2 m_T^2}{(m_\chi + m_T)^2} \propto \mu_{\chi T}^2, \quad (1.29)$$

where \mathcal{M} is the matrix element, E_{cm} is the center of mass energy, and $\mu_{\chi T}$ is the reduced mass of the DM-nucleus system. Isotropic scattering implies that the recoil energies are uniformly distributed from 0 to the maximum kinematically allowed energy $rm_\chi v^2/2$, where $r = 4m_\chi m_T / (m_\chi + m_T)^2$ is the maximum fraction of DM kinetic energy which can be transferred in the backscatter case. Instead of writing limits in terms of a coupling constant g and mediator mass M_Λ , spin-independent direct detection results are quoted in terms of

a model-independent WIMP-nucleon cross section $\sigma_{\chi p}$ and scaled to the appropriate target. This allows experiments utilizing different targets to compare limits.

$$\sigma_{\chi T} = A^2 |F(q)|^2 \frac{\mu_{\chi T}^2}{\mu_{\chi p}^2} \sigma_{\chi p} \quad (1.30)$$

Where A is the atomic number of the target and $F(q)$ is the form factor of the target, encapsulating the loss of coherent enhancement. Putting this together, one obtains:

$$\frac{dR}{dE_r} = A^2 |F(q)|^2 \sigma_{\chi p} \frac{\rho_\chi}{2m_\chi \mu_p^2} \int_{v_{min}(E_r)}^{v_{esc}} dv \frac{f(v)}{v}, \quad (1.31)$$

where $v_{min}(E_r)$ is the minimum velocity of the incident particle which can result in a recoil energy E_r . The scaling features favor matching the target mass with the WIMP mass. For example, Xenon ($A \sim 131$) is most sensitive at approximately $m_\chi = 40$ GeV. The asymptotic behaviour as $m_\chi \rightarrow \infty$ can be seen, as $\mu_p \rightarrow m_p$, and the overall differential recoil rate scales as $\sigma_{\chi p}/m_\chi$. This degeneracy leads to an inability for direct detection experiments to distinguish high mass WIMP models from one another. In Xenon, masses above 100 GeV/ c^2 generate nearly identical recoil spectra. Example recoil spectra are shown in Figure 1.5.

Halo Model

The velocity distribution $f(v)$ is determined by the dynamics of the galactic halo. This is typically assumed to be a Maxwell-Boltzmann distribution with dispersion v_0 . The laboratory velocity is then the earth velocity in galactic coordinates, which is the sum of the local standard of rest, the sun's peculiar velocity, and the velocity of the earth around the sun. Together, this model is referred to as the “standard halo model” and its probability distribution is plotted in Fig. 1.6.

Crystal Experiments

Solid state detectors measure energy deposits which promote electrons from the valence band into the conduction band. Crystals may be grown with small, eV scale band gaps, and by operating at mK temperatures SuperCDMS has achieved a single e-h pair sensitivity in their 0.93 g sensor[46]. Either the ionization signal is detected, as is the case with (Super)CDMS, or the crystal is doped such that the electrons are trapped and emit a scintillation photon upon deexcitation, as in the NaI-based DAMA[47] and the calcium-tungstate based CRESST[48].

Some experiments have also started to utilize the Migdal effect in WIMP-nucleon scattering in order to search for lower masses, such as the Ge-based Edelweiss[49], which set competitive limits in the 10-100 MeV/ cm^2 range. This is a phenomenon predicted to occur for low-energy nuclear recoils, whereby the transient displacement of the recoiling nucleus within its own electron cloud results in either direct ionization or de-excitation x-rays[50].

Concurrently, Edelweiss and similar experiments have started searching for DM-electron coupling[51], which benefits greatly from the reduced energy threshold. Experiments of this

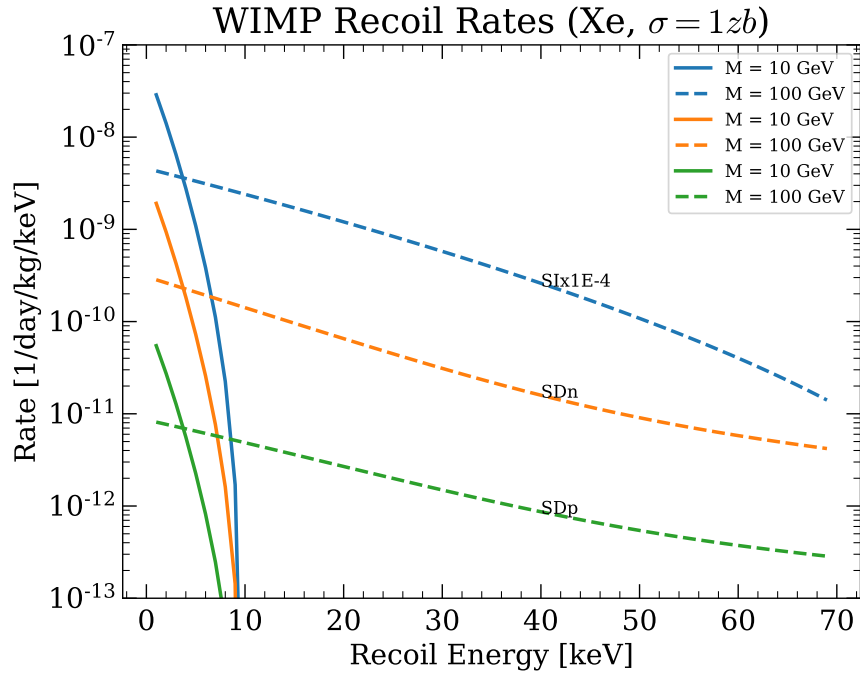


Figure 1.5: Differential recoil rates for assorted WIMP models with a WIMP-nucleon cross section of 1 zb (zeptobarn, $= 10^{-45} \text{ cm}^2$). “SI” denotes spin independent interactions, while “SDn” and “SDp” denote spin-dependent interactions with neutrons and protons, respectively. The SI limits are scaled downwards due to the fact that Xenon nuclei have at most one unpaired nucleon spin, eliminating much of the coherent enhancement present for spin-dependent limits. Each model is presented with two WIMP masses: $10 \text{ GeV}/c^2$ and $100 \text{ GeV}/c^2$. Both models show a mostly featureless falling exponential spectrum, but the larger mass has scatters out to larger recoil energies due to the better kinematic matching.

type are commonly bolometers, which detect the phonons (heat) produced by scatters. The Neganov-Luke effect produces a gain in phonons as the electrons are pulled through the crystal.

Skipper-CCDs are a leading technology in this mass range. SENSEI[52] utilizes this technology to non-destructively read the integrated charge in each pixel, which yields sub-electron noise accuracy. Their limits are world-leading below a mass of $10 \text{ MeV}/c^2$. Surface backgrounds are a problem for this technology due to their small size not benefiting from self shielding. However their low thresholds allow them to perform annual and diurnal modulation searches.




Figure 1.6: The lab frame velocity probability distribution function of the standard halo model [44, 45]. Note the Gaussian shape, peak at approximately 300 km/s, and the tail extending out to 800 km/s.

Noble Liquid Experiments

Noble liquid experiments measure charge and light signals. Time-projection chambers (TPCs) have the advantage of being monolithic and relatively straightforward to increase in target mass from kg to ton scale. They have superior kinematic matching for high mass dark matter, while having higher energy thresholds which limit sensitivity to sub- GeV dark matter. For example, LZ sets the world’s best exclusion limits for $40 \text{ GeV}/c^2$ dark matter[53]. More details can be found in Chapter 2.

Directional Detection

A smoking gun signal for dark matter would be a directional signal coming from the direction of the prevailing WIMP wind at declination 42° . In liquid noble TPCs the $4 \mu\text{m}$ thermalization length of electrons, along with the S2 position resolution, makes it challenging to determine the recoil direction. High pressure gases have the potential to balance the benefits of a dense medium (expense, energy threshold) with the ability to reconstruct long tracks and scale to extremely large sizes. The CYGNUS experiment is exploring the possibility of building a He and SF_6 based detector at atmospheric pressure[54].

Chapter 2

The LUX-ZEPLIN Experiment

2.1 Overview

LUX-ZEPLIN (LZ) is the result of a merger between the LUX[55] and ZEPLIN[56] dark matter direct detection experiments. It is primarily sensitive to WIMP dark matter, but performs other searches such as for neutrinoless double beta decay. Located at the 4850 level of the Sanford Underground Research Facility in Lead, South Dakota, USA, the overhead rock provides the equivalent of 4300 meters of water[57], which shields against cosmic ray flux.

LZ is a xenon-based dual phase time projection chamber (TPC), detecting energy deposits using prompt scintillation from electronic excited states (a.k.a the S1 light pulse) and a delayed proportional scintillation from ionized electrons (a.k.a. the S2 light pulse). The detector consists of two concentric, sealed low-radioactivity titanium containers, the inner and outer cryostat vessels (ICV and OCV, respectively), with an evacuated volume between. Inside the ICV is 10 tonnes of xenon, 7 tonnes of which is contained within the field shaping cage separating the TPC from the outer “skin” region[58]. Surrounding the OCV are acrylic tanks filled with liquid scintillator, serving as an additional neutron veto known as the Outer Detector (OD). The entirety of the OCV and OD is immersed in a large water tank which provided additional shielding against muons and radioactivity from the surrounding cavern.

Several inlets and outlets exist in the system. Xenon is continuously circulated through a purification system, and tested for purity using a cold trap-based mass spectrometer system[58]. The high voltage establishing the electric field is generated externally and is routed through the water tank and cryostat vessels (more details found in Chapter 3). Radioactive sources can be raised and lowered through *source tubes* in order to calibrate the detector response, and an evacuated conduit exists to collimate neutrons from a deuterium-deuterium fusion generator.

In this chapter I elaborate on the subsystems of LZ, provide an overview of the data collection and analysis, explain the WIMP search and its backgrounds, and enumerate the additional physics that LZ may search for.

2.2 Time Projection Chambers

LZ is a particular example of a class of detectors called Time Projection Chambers (TPCs), which consist of a set of parallel wire grids and a scintillating detection medium. An energy deposit creates a prompt scintillation signal, along with liberated electrons, which drift through the electric field some distance to the surface, where the charges are read out. The time between the two signals provides an estimate of the distance along the field line the deposit occurred, which, when combined the drifted location the electron cloud provides three-dimensional position reconstruction. The prompt scintillation signal (S1) is detected through a single-photon sensitive light detector, such as a photomultiplier tube (PMT) or silicon photomultiplier (SiPM). The ionization signal is detected through one of two methods: amplifying the charge detected on the anode wires, or the proportional scintillation (S2) as the electrons are extracted through a high electric field gap. Typically the first method relies on a secondary induction grid oriented perpendicular to the anode, in order to achieve XY position reconstruction. The second method usually occurs within a gas gap placed between the gate and anode grids. Using the S2 signal has the advantage of having electron amplification, producing dozens of detected photons for every extracted electron, and this is the method that LZ adopts. These principles are illustrated in Fig. 2.1.

In rare event searches, the rate typically scales with the number of target nuclei in the detector. This leads these TPCs to use a liquid detection medium to exploit the increased density. However, when detecting S2s this leads to an additional technical challenge of maintaining a stable liquid level in the electroluminescence region. LZ achieves this through a *weir*, which causes the liquid to flow off the sides of the TPC into the circulation system.

Of the 7 tonnes of active liquid xenon (LXe) in the LZ TPC, approximately 5.6 tonnes are encapsulated in the inner fiducial volume (see Section 2.4). The TPC detector in LZ consists of two PMT arrays located at the top and bottom, with a total of 494 3-inch Hamamatsu R11410[59] tubes, and four woven wire mesh grids. The grids are labelled as the:

1. Bottom, which protects the bottom array PMTs and establishes the Reverse Field Region (RFR) where electrons drift downwards, away from the anode. In the LZ coordinate system its height is $z = -137$ mm.
2. Cathode, biased to a large negative voltage. It establishes the electric field in the Forward Field Region (FFR), where electrons drift towards the anode and can be detected. The cathode defines the origin of the LZ coordinate system at $z = 0$ mm.
3. Gate, located at a height of $z = 1456$ mm, just below the liquid level of $z = 1461$ mm. This is biased to a small negative voltage and establishes the electroluminescence (EL) region, where S2s occur. The voltage is chosen to provide 100% transparency to drifting electrons.
4. Anode, located $z = 1469$ mm, just above the liquid level. The voltage is generally chosen to be slightly positive, and symmetric with the gate voltage about zero.




Figure 2.1: The cross section and primary operating principle of LZ’s time projection chamber. Rendering by Nicolas Angeledes. A cartoon of the resulting waveform S1 and S2 is indicated on the right hand side.

LZ lacks a grid directly below the top grid. The heights describe the location of the frames holding the grids, but external forces such as gravity and the electric fields cause the grids to deflect upwards or downwards on millimeter scales. This causes the effective LL gap, and electric field, to vary as a function of radius. The net effect is an approximately 50% larger single electron S2 size at the center of the detector than the edges.

The field between the grids is made uniform by a series of titanium field shaping rings. Encapsulating the rings is a PTFE wall, which provides structure. The inner radius of the PTFE is 728 mm. Surrounding the TPC is the hermetically sealed inner cryostat vessel (ICV) and outer cryostat vessel (OCV). Both are constructed of ultrapure titanium [60], and the ICV and OCV are separated by a vacuum for thermal insulation. A large, sealed water tank contains the OCV and associated piping, which provides additionally shielding to muons and rock gammas. The cathode high voltage feedthrough (see Chapter 3) connects through these volumes to transport power from the high voltage power supply located in the Davis cavern.

2.3 Xenon

Advantages of Xenon

Xenon was chosen for LZ for a number of positive qualities.

- **High density:** Liquid xenon has a density of 2.85 g/cm^3 at the LZ operating temperature of 174 K. This is advantageous for spin-independent (SI) couplings as more nucleons are packed into the same volume, keeping instrumentation costs down.
- **High A, Z:** The large atomic number of natural xenon ($A_{\text{ave}} = 131.293$) is similarly advantageous for SI couplings. The coherent enhancement of the scattering cross section $\sigma \propto A^2$ gives it an edge over similar noble liquids. The large Z also increases its electromagnetic stopping power, which gives it a *self-shielding* effect, where external backgrounds are stopped over short distances. This enables fiducialization, where the extremely radio-quiet center of the volume is selected for analysis.
- **Chemistry:** As a noble element, Xenon does not readily form stable molecules with other elements. This allows it to be chemically purified to extremely high purity. LZ accomplishes this using a getter system, which removes electronegative impurities. The resulting purity was observed between 5–8 ms[53], which implies an oxygen equivalent impurity of between 60 and 90 parts per trillion[61]. Note that this technique does not remove the radio-isotopes of ^{85}Kr and ^{39}Ar , which were constrained in the background model to 144 ppq and 890 ppt[62].
- **High Scintillation Yield:** Xenon has a low effective work function (13.7 eV[63]) for producing a single quanta on average. The energy threshold is primarily set by the number of coincident PMT hits, and therefore a higher scintillation yield is preferable. LXe has a scintillation yield L_y at the 122 keV ^{57}Co γ -ray line of 63 photons/ keV[64]. By comparison, LAr has a W -value for scintillation of 19.5 eV[65], and a W -value for ionization of 23.6 eV[66].
- **No problematic radioisotopes:** Xenon has several stable isotopes: ^{124}Xe , ^{126}Xe , ^{128}Xe , ^{129}Xe , ^{130}Xe , ^{130}Xe , ^{131}Xe , ^{132}Xe , ^{134}Xe , ^{136}Xe . Of these, there are no isotopes with half-lives with lifetimes long enough to not decay away during commissioning, but short enough to have high activity. The ^{124}Xe isotope undergoes $2EC$ [67] with a $t_{1/2} = 1.8 \times 10^{22} \text{ y}$, and ^{136}Xe undergoes 2β decays with a $t_{1/2}$ of $2.165 \times 10^{21} \text{ y}$ [68], and the ^{134}Xe is thought to have an (as of yet unobserved) 2β -decay with $t_{1/2} > 8.7 \times 10^{20} \text{ y}$ [69]. Between the extremely long half lives and the high Q-values of the decays, these do not pose a significant impediment to the LZ WIMP search result, with 24.3 expected electron recoils (ERs) in the ROI during the first science run (SR1)[53]. Argon, by comparison, has the cosmogenically activated ^{39}Ar with $t_{1/2} = 268 \text{ yr}$ and Q-value 565 keV. This leads to a typical event rate of 1 Bq/kg. Xenon does have short-lived cosmogenic activation isotopes ^{37}Ar and ^{127}Xe , which contribute to the early science runs[70], and are incorporated into the background model.




- **Particle ID:** LXe responds differently to electron recoils and nuclear recoils. The partitioning of the quanta into ionization and excitation channels varies between the two particle types. This allows efficient ($> 99.5\%$) rejection of ER backgrounds for the NR WIMP search.

Xenon does have some acceptable downsides. It is highly expensive, owing to its rarity, fluctuating around \$10/liter. The decay time of the singlet and triplet lifetimes are not as disparate as Argon, so pulse shape discrimination is not as exploitable. Its scintillation yield is also highly dependent on energy and electric field strength, making energy reconstruction more complicated, in general requiring the additional ionization channel to be effective at low energies.

Microphysics

Xenon responds to energy deposits with electronic excitation and ionization. Excitation results in the formation of metastable molecules known as excimers. The Xe_2^* dimer consists of two electronic states, the $^1\Sigma_u^+$ and $^3\Sigma_u^+$ states, known as the “singlet” and “triplet” states. These states de-excite to the ground state $^1\Sigma_g^+$ after some time by emitting a 175 nm VUV photon[71]. The triplet state’s transition to the ground state is forbidden, so it has a larger lifetime than the singlet. The dimer lifetimes are 3 ns and 24 ns, respectively [72].

The initial quanta of excimers N_{ex} and free electrons N_i may result in different number of photons n_γ and drifting electrons n_e . The free electrons have the ability to recombine with a Xe^+ ion and form a dimer, leading to one fewer n_e and one greater n_γ . The dimers can also undergo processes such as *Penning quenching* and *Penning ionization*, whereby the excitation energy is lost to heat or ionization. The following processes are possible:



where Xe^{**} is a doubly excited state. The recombination probability depends on many factors, including the external electric field and the initial density of ions. At high energies, the charge and light yields are well described by the *Doke-Birks* model[65]:

$$r = \frac{A(\mathcal{E})dE/dx}{1 + B(\mathcal{E})dE/dx} + C , \quad (2.6)$$

where r is the recombination fraction of ionized electrons, A , B are empirical parameters, $C = 1 - A/B$, and \mathcal{E} is the magnitude of the electric field. The linear energy transfer (LET, $\frac{dE}{dx}$), diminishes for electron recoils (ER) as a function of energy, and therefore lower energy recoils are predicted to have lower charge yields due to recombination than xenon nuclear

recoils (NRs). While this model holds at electron recoils with energy greater than 50 keV, it breaks down in the WIMP search region of interest. Experimentally, ERs below ~ 30 keV demonstrate a rise in charge yield $Q_Y = N_e/E$ as E approaches 0. This behaviour can be modelled by the Thomas-Imel Box model[73, 74], given by

$$r = 1 - \frac{\ln(1 + \xi)}{\xi} \quad (2.7)$$

$$\xi = \frac{N_i \alpha}{4a^2 v} , \quad (2.8)$$

where a is the size of the "box" the electrons are contained within, v is the velocity of the electrons, and α is a parameter fit to the data which controls the recombination rate as a function of ion and electron densities. This model is applicable when the length of the track is comparable to the thermalization distance of the electrons, $4.6 \mu\text{m}$. Both models for the partition of energy into ions and excimers are included in the NEST simulation package[64], used by LZ and other liquid noble element experiments.

WIMPs will scatter preferentially off nuclei, and most background sources will scatter off electrons. LZ uses the ratio of S1 to S2 light to distinguish the two types of recoils. In principle, low energy electron and nuclear recoils should be difficult to distinguish, as their $\langle dE/dx \rangle$ places them in the regime of the Thomas-Imel Box model. However, empirically electron recoils deposit the vast majority of their energy into ionization initially, $N_i/(N_i + N_{ex}) = 0.96$, while nuclear recoils partition their energy almost equally between ionization and excitation, $N_i = N_{ex}$. Under the recombination model above, this results in lower S2 to S1 ratios for NRs, and a corresponding separation of bands in $\text{Log}(S2)$ vs S1 space as shown in Fig. 2.2. The bulk electric field decreases the recombination fraction for a given recoil type.

In addition to the aforementioned quenching of S1 light from increased electric fields, nuclear recoils experience another form of quenching, whereby a large portion of their recoil energy goes into random atomic motion, rather than observable ionization and scintillation signals. This quenching is described by Lindhard theory[75, 73]:

$$\mathcal{L}_{\text{eff}} = \frac{kg(\epsilon)}{1 + kg(\epsilon)} , \quad (2.9)$$

where $k = 0.1394$ [64] is a proportionality constant between the electronic $\langle dE/dx \rangle$ and the velocity of the recoiling nucleus, $g(\epsilon)$ is the ratio of electronic stopping power to nuclear stopping power, and ϵ is the deposited energy. This leads to a discrepancy between E_e , the reconstructed energy assuming $\mathcal{L}_{\text{eff}} = 1$, and E_{nr} , the physical nuclear recoil energy. After calculating the Lindhard factor, predicted nuclear recoil energy spectra are commonly displayed in "electron-equivalent" energy units, keV_{ee} . The WIMP search region of interest for LZ exists between 1 keV_{ee} and 30 keV_{ee} .

2.4 Backgrounds and Mitigation

Sources

Ultimately, the LZ WIMP search depends on keeping background events low in the energy region of interest. While the PLR (Section 2.7) helps by looking at the distribution of S1 and S2 events in the TPC, the sensitivity is still determined in part by the amount of events in the NR band seen during the exposure. Sources of non-WIMP events include[62]:

- **Detector ERs:** detector components produce β and γ rays from trace radioactive elements ^{60}Co , ^{40}K , and the ^{238}U and ^{232}Th chains. These events appear along the boundary of the TPC and are removed with a fiducial cut.
- **Radon Progeny ERs:** Radon, being a noble element, dissolves into the Xe and can not be removed by gettering. While ^{220}Rn has a half life of 3.6 days, the ^{222}Rn isotope is worrisome due to the violation of secular equilibrium, with its progeny ^{210}Pb having a 22.3 yr half life[76]. The lead plates out on the surfaces, leading to β -decays later on. In both chains unstable Pb and Bi undergo β -decay, which in principle is not problematic as long as the decays can be tagged by an associated γ -ray(for Pb) or prompt daughter α (for Bi). However, ^{214}Pb and ^{212}Pb will occasionally either β -decay directly to the ground state, or emit a γ -ray which escapes without photoionization. These situations are colloquially referred to as “naked” and “semi-naked” betas, and constitute a large portion of the ER background dispersed throughout the TPC. The α decays are not a concern due to their high energies, placing them well outside the WIMP search ROI (typically $\mathcal{O}(10^4)$ phd S1s, where the upper limit for the analysis was 80 phd).
- **Intrinsic Xenon Decays:** As mentioned above, Xenon has some radioisotopes, namely ^{124}Xe , ^{127}Xe , and ^{136}Xe . The ^{127}Xe is a cosmogenic activation product and decays via electron capture with a $t_{1/2} = 36.3$ d, while the others are $2\nu 2\beta$ decays with long half lives. Another activation isotope, ^{37}Ar , is present, but unlike ^{127}Xe its rate can not be constrained by K-shell decays, and therefore in the first science run (SR1) result its rate was allowed to float with a flat prior[53].
- **Neutrinos:** Solar neutrinos contribute a uniform background of ERs, and their flux and scattering rate are constrained by other measurements. Coherent neutrino-nucleus scattering ($CE\nu NS$) from ^8B contributes a nuclear recoil background which is similar in energy to a low-mass WIMP.
- **Accidentals:**
Discussed more in Chapter 7, this background is a result of uncorrelated S1-only and S2-only signals randomly pairing up to appear in the WIMP search ROI. Multiple sources contribute to this background, and its rate is constrained by unphysical drift time (UDT) events with reconstructed Z below the cathode.

- **Neutrons:** Radiogenic neutrons can elastically scatter in the TPC, mimicking a WIMP signal. Due to the effectiveness of the outer detector, and the propensity of the neutrons to multiple scatter, these do not constitute a large background source.

The β decays to the ground state of the daughter lack associated γ rays. These are known as “naked” β decays, and they and detector ERs constitute approximately 222 of the estimated 333 background events.

Recoil Discrimination

Many of the LZ backgrounds are the result of recoiling electrons, rather than recoiling xenon nuclei. LZ passively collects ERs as it runs, and calibrates NRs using, among other sources, a collimated deuterium-deuterium (DD) fusion source. The DD neutrons elastically scatter in the LXe, forming a wide band of energies. Tritiated methane CH_3T is injected to populate the ER band down to low energies. For a given observed S1 signal, ERs have a higher S2 area on average than NRs.

The bands are defined using the quantiles of the data. The probability that an ER has a downwards fluctuation in S2/S1 ratio and ends up in the nuclear recoil band is known as the “leakage fraction” and exceeds 99.5%. While this serves as a useful metric for analyzers, the final result uses the profile likelihood ratio (PLR), discussed in Section 2.7. The ER recoils are not removed with a specific cut, but the background pdf is fit to the data, naturally informing the leakage into the NR.

Multiple Scatters

Gamma rays and fast neutrons can scatter multiple times in the detector. Dark matter will scatter at most once in the detector, owing to its extremely small scattering cross section (until masses around 10^{16} GeV/ c^2 , where there are unexcluded cross sections which unexcluded cross-sections which may lead to multiple scatters- see Chapter 5). Most backgrounds are due to relativistic particles, and result in merged S1 signals, so the determination of single vs multiple results from the number of detected S2s. As such this selection criterion is most effective for particles with vertical trajectories, as the vertical resolution is around $\sigma_Z = 2$ mm.

Fiducialization

The detector materials constitute a source of backgrounds from trace U-238 and Th-232, leading to enhanced backgrounds near the boundaries of the TPC. By removing events near the gate, cathode, and PTFE walls, the detector ER backgrounds will range out while the signal remains. This does not by itself mitigate the backgrounds from dispersed backgrounds such as Pb-214/Pb-212 betas without associated gammas, and Ar-37/Xe-127 activation peaks. The fiducial volume (FV) is a cylinder with cutouts for the cathode and gate rings, which contribute higher backgrounds. The upper and lower boundaries are 2 cm from the cathode




Figure 2.2: $\log_{10}(S2c)$ vs. $S1c$ areas for different calibration sources, demonstrating particle identification. The electron and nuclear recoil bands constructed from LZ tritium(ER, blue) and DD(NR, orange) calibration data. Figure taken from the LZ SR1 result paper[77]. Note the larger separation of the bands as the recoil energy increases.

wires, and 13 cm from the gate wires. The nominal radius is 6.1 cm from the TPC wall, but drift times below $200 \mu s$ are cut 7.7 cm from the wall, and drift times exceeding $800 \mu s$ are cut to 6.5 cm from the wall.

2.5 Veto System

Skin

The increased standoff distance of the cathode high voltage feedthrough from the cathode ring (discussed further in Chapter 3) allows for the volume of the xenon between the PTFE wall and the ICV to be instrumented with additional PMTs. These 131 additional sensors are referred to as the “skin” detector, and are divided into the “barrel” around the outer radius, and the “dome” beneath the TPC PMT array. The ICV in the barrel region surrounding the sides is coated with PTFE to increase light collection. As the ICV is grounded, there exists non-trivial electric fields within the skin region, which affects the light yield observed there. This additional detector serves as a γ -ray veto, as well as an additional high- Z shield against external backgrounds. Nuclear de-excitations from ^{127}I were observed with coincident Skin pulses 80% of the time[53].

OD

The outer detector (OD) serves a similar function to the Skin, *i.e.* it removes events with one scatter in the TPC and additional scatters elsewhere, but its operation is highly specialized for neutron tagging. It consists of acrylic tanks enveloping the OCV within the water tank, instrumented with 120 8" Hamamatsu R5912 PMTs, and filled with 17.3 tonnes of Gadolinium-loaded liquid scintillator (GdLS)[78]. A total of 23 kg of Gd was loaded into the scintillator, for a concentration by mass of 1.3 parts-per-thousand[79].

A neutron may scatter numerous times before thermalizing, after which it captures on the Gd, which has an extremely large capture cross section of $48,890 \pm 104$ barns[80]. The nucleus then emits a cascade of γ -rays with total energy of 8 MeV. It is these ERs which are used to tag the neutrons: any energy deposit above 200 keV_{ee} within 1.2 ms of the single-scatter S1 causes the OD veto to fail. AmLi calibrations were used to estimate the tagging efficiency at $\eta = 88.5\%$ [53].

The OD veto also serves as an effective muon veto. The high light yield of the linear alkyl-benzene and the high $\langle \frac{d}{E} dx \rangle$ of minimally ionizing muons leads to 20,000+ photons detected within the OD. This muon tagging also aids in the form of a *muon veto*, which removes regions of time after muons pass through the TPC where the electron/photon noise is elevated.

2.6 Calibrations

The detector response particular to LZ was characterized. In particular the S1 and S2 photon gains g_1, g_2 have to be estimated at each point in the detector, the ER and NR bands were mapped out, and the PMT single detected photo areas (phd) had to be measured in order to have a uniform gain across the PMT arrays. Note that due to the 175 nm scintillation light, the PMT photocathode will occasionally emit more than one photoelectron (phe) for every photon[82]. Due to this dpe fraction f_{dpe} , there is a conversion between the two units, with 1 phd $\approx (1 + f_{dpe})$ phe. The DPE fraction is approximately 20%.

The PMT gains were calibrated using an array of blue LEDs. These are flashed such that, on average, one in every ten flashes results in a PMT hit. The areas of the response pulses are histogrammed to identify the single and double photoelectron peak.

For g_1 and g_2 the Doke-plot method was utilized[83]. The energy of any particular event is given by

$$E = W \left(\frac{S1}{g_1} + \frac{S2}{g_2} \right), \quad (2.10)$$

where W is the effective work function, $S1$ and $S2$ are in phd, and g_1, g_2 are the photon gains per quanta. Sources of known energy are plotted in $S2/E$ vs. $S1/E$ space. Sufficient quantity of monoenergetic sources plot out the following line:

$$\frac{S2}{E} = \frac{g_2}{W} + \frac{g_2}{g_1} \cdot \frac{S1}{E}. \quad (2.11)$$




Figure 2.3: CAD rendering of the LZ water tank and detector. The TPC is shown in the middle, with the cathode high voltage feedthrough extending to the left, DD conduit extending right, and calibration source tubes running vertically. The GdLS contained within acrylic tanks is visible in green, surrounded by the OD PMTs. Figure taken from Ref. [81].

With only 2 degrees of freedom, the line cannot simultaneously determine g_2 and W . A common value in the literature for the effective work function is $W = 13.7$ eV[63], but recent measurements have indicated $W = 11.5$ eV[84, 83]. In SR1 a value of $W = 13.5$ eV was used, and the tuned g_1 and g_2 values were verified with an independent minimization procedure using NEST[85]. For LZ, many peaks of known energy are present but the ^{83m}Kr (41.5 keV), ^{129m}Xe (126 keV), and ^{131m}Xe (164 keV) peaks are prominent enough to use for this analysis. The values were determined to be $g_1 = 0.1136$ phd/photon, and $g_2 = 47.07$ phd/e[53]. An example of an application of the Doke-plot method is shown in Fig. 2.4 for another experiment[83].

As mentioned above, CH_3T is injected to calibrate the fiducial volume for electron recoils. Since tritium has an 11-year half-life for its β -decay, its injection is in theory problematic. However, it was determined in commissioning tests that methane can be efficiently removed with the LZ getter so it does not pose an issue in practice.

Monoenergetic dispersed ERs are calibrated with ^{83m}Kr injections. The Kr is formed from ^{83}Rb electron captures. The 41.5 keV total energy of the internal conversion electrons

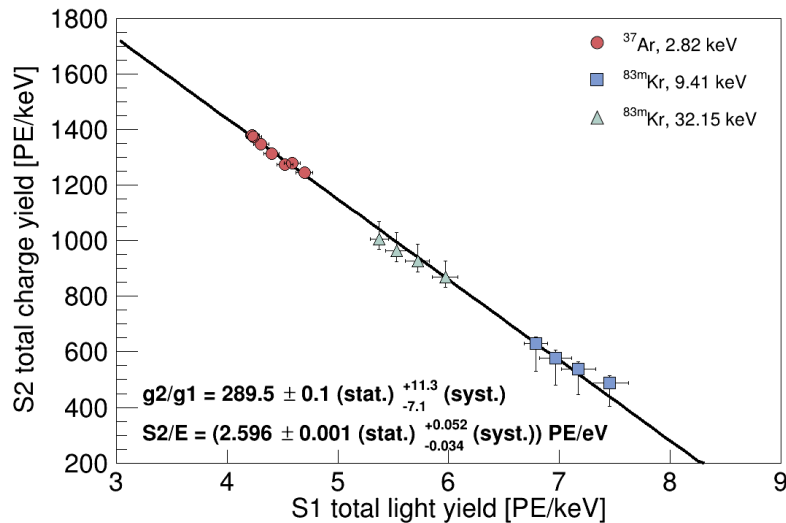


Figure 2.4: A Doke plot, taken from [83], which determined $W = 11.5$ eV.

sit outside the WIMP search ROI, but can still be used to calculate the detector yields. Larger S2s ($> 10,000$ phd) allow for excellent position resolution, which allows for mapping of the boundary of the TPC[86]. The two S1s for $^{83\text{m}}\text{Kr}$, labelled as $S1_a$ and $S1_b$, are affected by recombination in interesting ways, as the later decay occurs in the wake of the first. Increasing electric fields affect the time dependence of $S1_b/S1_a$ [87], which can allow a direct measurement of the fields at each position in the detector.

With a half life of 1.83 hr the Kr decays before adequately mixing within the LXe under certain circumstances. This allowed LZ to discover that the FV was in a “low flow” state, where little mixing occurred with the boundary layer. While advantageous from the perspective of Rn-emanation, it does raise some questions regarding electronegative impurities being effectively cleaned out by the circulation system. Flow-mapping using Rn-Po coincidences is ongoing in order to understand this region. The low flow region is seen in the Kr calibration data in Fig. 2.5

Nuclear recoil calibrations were performed with DD and AmLi. A commercial generator creates 2.45 MeV neutrons, which are guided using a hollow conduit through the LZ water tank to the OCV. The setup can be rearranged to use a neutron reflector which reduces the energy and flux further, allowing calibration down to 1.1 keV_{NR}. While this is a pure source of NRs, the scatters occur along a narrow beam at the top of the detector. To calibrate the positional dependence, an AmLi $\alpha - n$ source was prepared, which produces broad band neutrons out to 1.5 MeV. The trade offs that some amount of gammas are produced from the excited final states, and that the event rate is high, leading to copious accidental events which have to be removed for proper analysis.

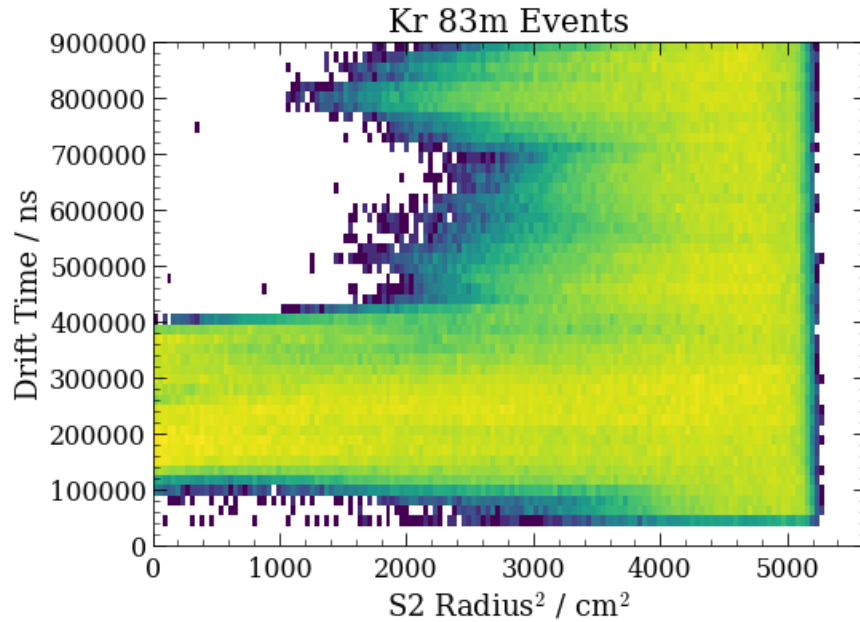


Figure 2.5: 2D histogram of reconstructed event positions, S2 radius vs. drift time, from an early ^{83m}Kr injection. The low flow region is seen in the middle, along with the detector boundaries. Single scatter events within the drift region are shown. Note that, despite the low flow in the middle of the detector, a relatively straight and clear indicator of the wall location is found.

2.7 Analysis

Data collection

Data is collected using online FPGAs with an advanced trigger which causes the digitized waveforms to be written to disk. This trigger has configurable settings for different scenarios *e.g.* calibrations or background, but during the WIMP search it is tuned to trigger an “event” for every S2. The waveforms from the assorted TPC PMTs are summed, and are filtered with a boxcar filter which efficiently detects the presence of an S2 (a separate filter exists for S1s, which under WIMP search settings does not trigger events). Triggering events are written to disk using a compression method called *Pulse-only digitization* (POD). Each PMT channel is filtered and zero-suppressed such that, as the name implies, only regions of time around likely photoelectrons signals are saved and not noise. These .evt files containing PODs and metadata are sent to the US and UK data centers for processing.

Raw event files are processed using the Gaudi-based LZ Analysis Package (LZAP). The pulsefinder runs over the summed PODs to identify discrete pulse boundaries. From these pulses reconstructed quantities (RQ)s are then calculated, *e.g.* pulse area, prompt fraction, etc. These RQs are used to classify the pulses as one of several types: S1, S2, single electron

(SE), single photoelectron (SPE), multiple photoelectron (MPE), or other. S2s and SEs have their XY positions reconstructed using the Mercury algorithm. Depending on the topology of pulses, the interaction finder then classifies the overall event into one of several scatter types: single (SS), multiple (MS), pileup (PU), or other, and calculates additional interaction-based RQs.

The RQ files output by LZAP are finally analyzed by the ALPACA framework. ALPACA is a modular framework which reads the `lzap.root` files, executes a particular analysis, and writes the results and plots to disk. It was developed so that multiple modules can be chained together into a single analysis. As opposed to LZAP, events in ALPACA are not assumed to be independent. This comes into play with the e-train veto, which removes periods of time following large S2s, amounting to a dead time fraction of 29.8%. The SR1 WS analysis proceeds by identifying single scatters within the WIMP search region of interest (ROI): $S1c \in [3, 80]$ phd, $S2c \in [600, 10^5]$. The S2 radius and drift time must lie in the fiducial volume, and OD and Skin vetoes (see Section 2.5) are applied based on the S1. Coincident S1 pulses within the skin detector are vetoed. There are two selections: a prompt cut, which removes events with any $N > 2$ PMTs coincident in the skin within 500 ns, and a delayed cut, which removes events with skin S1s within an expanded coincidence window of 1.2 ms if the skin pulses were large enough ($N > 55$). Finally, anti-accidental data quality cuts are applied (more details in Chapter 7).

Detection Threshold

The WIMP recoil spectrum is a smooth, falling exponential, which heavily incentivizes making the energy threshold as low as possible. A limiting factor here is the accidental PMT coincidence background, discussed in greater detail in Chapter 7, which sets the smallest S1 as 3 phd, detected in 3 distinct PMTs. This is more than the 2-fold coincidence requirement used in LUX. The scintillation yield decreases as deposited energy approaches zero, exacerbating the problem. The number of detected photons is a stochastic process due to the light collection efficiency, making events near threshold fluctuate downwards.

Data quality cuts are used to remove the accidental background, which also rises towards threshold, as detailed in Section 7.3. Both S1 and S2 based cuts were employed. A significant background of few-electron noise was observed, necessitating raising the S2 threshold to 600 phd. These criteria were tuned to have high acceptance at threshold for low-energy tritium and DD calibration data. In the end the S2-based cuts, which aimed to eliminate near-surface and above-anode events, end up reducing the acceptance to around 60% at $S2_{\text{raw}} = 600$ phd, and the S1 cuts maintain $> 95\%$ acceptance across the entire ROI.

Translating to reconstructed energy keV_{NR} yields Fig. 2.6. A roughly logistic curve is seen, rising from nearly 0% at 2 keV_{NR} , to 50% at 4.51 keV_{NR} , to $\approx 90\%$ across at 10-45 keV_{NR} . The WIMP search ROI removes the tail towards higher energies, as the corrected S1 pulse areas are required to be below 80 phd.

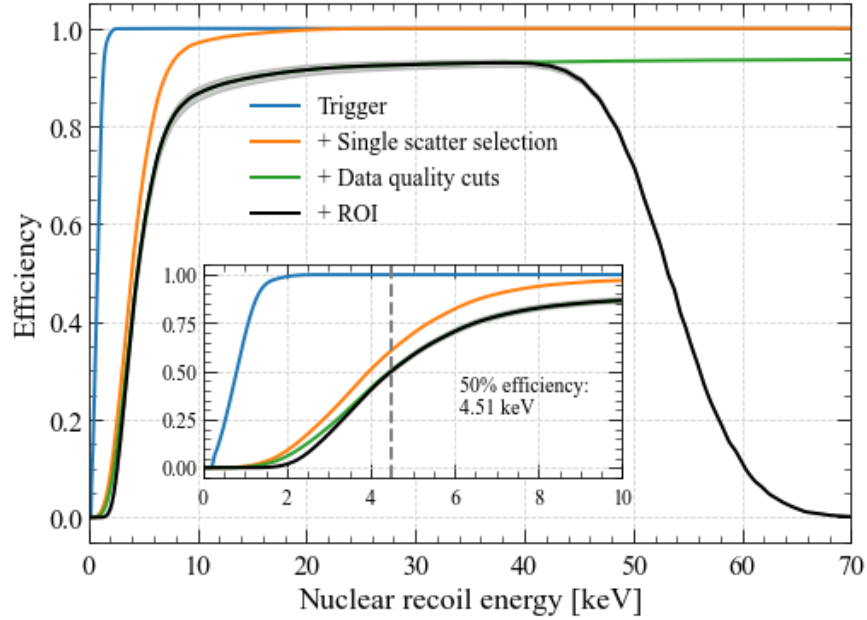


Figure 2.6: The nuclear recoil analysis acceptance curve for LZ’s first science run, taken from [53]. Note that the largest impact at threshold is the single scatter selection, while intermediate energies (10-40 keV_{NR}) are impacted most by the data quality selection, while the largest energies are removed by the region-of-interest (ROI). The inset indicates more precisely the soft energy threshold.

Profile Likelihood

LZ obtains two sided confidence intervals on the size of a physics signal based on the observed data. Rather than obtaining a ”cut and count” upper limit using the method of Feldman and Cousins[88], the analysis of the signal in the pretense of background proceeds with the *Profile Likelihood Ratio*(PLR) test statistic. In this case one estimate parameters $\boldsymbol{\mu}$ alongside some nuisance parameters $\boldsymbol{\theta}$ from data \mathbf{X} . The likelihood of \mathbf{X} is provided by the ”extended likelihood function”[89], which combines the Poisson fluctuations and probability distribution functions of the signal and background. When μ_χ is the number of signal events, θ_j are counts of background events, and $f_\mu(\mathbf{x}_i)$, $f_j(\mathbf{x}_i)$ are the distributions of each model, the extended likelihood is:

$$\mathcal{L}(\mathbf{X}; \boldsymbol{\mu}_\chi) = \frac{e^{-(\mu_\chi + \sum_j \theta_j)}}{N!} \prod_i^N [\mu_\chi f_\chi(\mathbf{x}_i) + \sum_j \theta_j f_j(\mathbf{x}_i)] . \quad (2.12)$$

The likelihood ratio generally is the ratio of two models: $q = \mathcal{L}(\mu_1)/\mathcal{L}(\mu_2)$. For the PLR the nuisance/background parameters θ_j are *profiled*, in other words the conditional probability is maximized.

$$\lambda(\boldsymbol{\mu}) = \frac{\mathcal{L}(\boldsymbol{\mu}, \hat{\boldsymbol{\theta}})}{\mathcal{L}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\theta}})}. \quad (2.13)$$

Here the double caret indicates that the likelihood has been maximized over $\boldsymbol{\theta}$ for a given $\boldsymbol{\mu}$, while the single caret over the parameters indicate the global maximum. Following the convention set in Ref. [45], the variable \tilde{t}_μ is defined:

$$\tilde{t}_\mu = \begin{cases} -2 \ln \lambda(\boldsymbol{\mu}) & \hat{\boldsymbol{\mu}} \geq 0 \\ -2 \ln \lambda(0) & \hat{\boldsymbol{\mu}} < 0 \end{cases}. \quad (2.14)$$

This handles the fluctuations below zero and is asymptotically distributed according to the χ^2 distribution with degrees of freedom equal to the number of fitted parameters less the number of data points[90]. In practice the asymptotic formulae are not used, and the distribution of $f(t_\mu|\mu)$ is found via monte-carlo sampling. The exclusion limit is then set by examining the models such that the probability of obtaining the observe \tilde{t}_μ , given the true value is μ , exceeds a predefined amount α (usually 0.1):

$$p_\mu = \int_{\tilde{t}_{\mu, \text{obs}}}^{\infty} f(t_\mu|\mu) dt_\mu > \alpha. \quad (2.15)$$

For discovery the background-only value \tilde{t}_0 is examined instead. By definition, the two sided exclusion limit "lifts off" with frequency $\alpha/2$, but this does not imply that a discovery has taken place. The coverage requirement to measure a signal a fraction α of the time leads to limits occasionally being set in regions of parameter space where the detector has little sensitivity. "Power Constrained Limits" (PCL) address this by incorporating $\pi(\mu)$, the probability to reject the background-only hypothesis in the presence of a signal, into the analysis[91].

2.8 Physics searches

WIMPs

LZ is designed to be maximally sensitive to spin-independent (SI) couplings between WIMPS and nucleons. The LZ SR1 result is shown in Figure. 2.7, setting a world-leading limit at the masses 20-40 GeV/ c^2 . The 50% acceptance threshold for NRs was found to be 5.5 keV $_{nr}$. An underfluctuation of the limit between 19-23 GeV/ c^2 resulted in the utilization of a power-constrained limit in that region[91].

Neutrinoless Double Beta Decay

Neutrino masses are currently not explained in the standard model. If they are their own antiparticle, i.e. their mass is from a Majorana mass term in the Lagrangian, certain processes become possible. In two neutrino double beta decay ($2\nu 2\beta$), a nucleus undergoes two




Figure 2.7: The LZ spin-independent WIMP-nucleon exclusion limit. Figure from [53]. Note the distinctive check mark shape, caused by the kinematic matching between the DM particle and the xenon, along with the energy threshold. The fluctuation of the exclusion limit downwards lead to the use of a power constraint between $m_\chi \in [19, 23] \text{ GeV}/c^2$. The green and yellow are “Brazil bands” which indicate the confidence interval of the background free limits.

simultaneous beta decays, resulting in a continuous distribution of total energy deposition out to the Q-value. If neutrinos are Majorana, the two outgoing neutrinos may instead connect as a virtual particle in the Feynman diagram, leading to neutrinoless double beta decay ($0\nu 2\beta$):

$$(Z, A) \rightarrow (Z + 2, A) + 2e^- . \quad (2.16)$$

The combined energy of the simultaneous β rays in this process is located entirely at the Q value of the $2\nu 2\beta$ process. Therefore, the signal appears as a sharp peak on top of the beta decay spectrum, and sensitivity is ultimately limited by energy reconstruction.

Natural xenon contains two potential candidates for $0\nu 2\beta$, ^{134}Xe and ^{136}Xe , with Q-values of 825.8 keV and 2457.83 keV [92]. At the ^{208}Tl line of 2.61 MeV, a resolution of $\sigma_E/E = 0.64\%$ is found. The projected median sensitivities are 7.3×10^{24} years and 1.06×10^{26} years for ^{134}Xe and ^{136}Xe neutrinoless double beta decay, respectively[93].

Boron-8

In the sun ${}^7\text{Be}$, part of the pp chain, occasionally captures a proton to form ${}^8\text{B}$. Boron then immediately β -decays with a 12.1 MeV endpoint. These neutrinos can coherently scatter off of the xenon nuclei. Due to the poor energy transfer, the ${}^8\text{B}$ neutrino signal appears as extremely low energy nuclear recoils, and is almost indistinguishable from a WIMP of mass 2-6 GeV/ c^2 . This signal poses a problem for the future of the WIMP search in that region, but remains a viable signal to search for in data. Lowering the coincidence threshold for S1s from 3 to 2 will increase the discovery potential by two orders of magnitude[94]. This analysis is impacted greatly by the ability to remove isolated S1s, which cause accidental coincidence background (see Chapter 7).

Transient Signals

A nearby supernova (within the galaxy) would generate neutrinos which would be observable in LZ. Approximately 350 NR events would be observed over a 10s period of time from a $27 M_{\odot}$ supernova 10 kpc away[95]. This would vastly expand the catalog of supernova events began by supernova 1987a. Cerenkov-based detectors such as Hyper-Kamiokande can probe $\bar{\nu}_e$, and the kiloton-scale DUNE far detector will be sensitive to ν_e . A clear understanding of the flux of all flavors of neutrino is valuable to the field of stellar physics. The low threshold and high N of liquid xenon TPCs affords observation of keV-scale $CE\nu NS$ nuclear recoils from $\mathcal{O}(\text{MeV})$ neutrinos. This interaction is sensitive to all flavors of neutrinos[96], unlike the other detector technologies, and can identify the timing of the core-collapse rebound.

Chapter 3

Cathode High Voltage for LZ

3.1 Introduction

An electric field must be established and maintained across the drift length of a TPC for successful operation. For dark matter searches, a large field is necessary since the electric/nuclear recoil discrimination improves with field magnitude up to a maximum between 240-290 V/cm[97]. Meter-scale TPCs are faced with the challenge of delivering large negative high voltages on their cathodes in order to establish these fields. While the drift fields are far below the minimum necessary to generate electroluminescence, localized regions of enhanced field carry additional risk. While the region around the wires have enhanced field, the cathode field is generally much lower than the field across the electroluminescence gap in the top of the detector. A volume of particular concern is the space where the cathode high voltage cable makes the connection with the electrode ring itself.

Previous LXe-TPCs have found it challenging to achieve their high voltage aspirations. Xenon-10 operated with a cathode biased to -13 kV for a drift field of 730 V/cm[98]. LUX Run-03 saw a drift field of 180 V/cm[55], corresponding to a cathode voltage of ~ 8.7 kV. Xenon-1T[99] ran at 120 V/cm in their first science run and at 80 V/cm in their second, with a cathode voltage of -12 kV and -8 kV, respectively. PandaX-4T [100] biased their cathode to values ranging from -20 kV to -16 kV (decreasing over time), which established drift fields between 93-121 V/cm. The dual phase LAr-TPC experiment Darkside-50[101] operating under similar grid biases of -12 kV on their cathode, which provided a drift field of 200 V/cm.

During LZ's design and construction, a *design* cathode voltage was specified at -100 kV, corresponding to what was thought to be an optimal drift field. A *requirement* cathode voltage was established as -50 kV, which equated to the perceived acceptable drift field for the experiment. The cathode high voltage (CHV) feedthrough would be the limiting factor in these scenarios, and therefore a novel design for TPCs was used.

In this chapter I detail the prototype testing, assembly, and deployment of the production CHV feedthrough for LZ. Following the success of the test a production version was machined and assembled in a clean room at LBNL. Finally the cable and feedthrough on-site at SURF,

making the final seals on the cryostat vessels. I explain the purpose, results, and challenges of this subsystem, along with my contributions.

This project was one of the first that I contributed to when I joined LZ in 2016. At that point in time the design and test parameters had already been selected, and the testing procedure was largely in place. I contributed to the construction of the test stand, working with project scientist Ethan Bernard. Modifications to and characterization of the photomultiplier tube (PMT) and other electronics were performed by myself. I assisted in the conduction of all tests and was responsible for the data collection and software. All data analysis was performed with code written by myself and run on local machines. The production feedthrough and cable was assembled, checked out, and packaged by myself and Ethan Bernard. The final connection was performed with a three-person team, including me, who used self-contained breathing apparatus (SCBA) to align the seals while under a nitrogen purge-induced low oxygen risk.

3.2 Design Considerations

Structure

The high voltage for LZ is generated outside the detector in high voltage power supply (HVPS) and delivered via a feedthrough. The feedthrough facilitates a continuous electrical connection between the cathode ring in the LXe space and the laboratory space where the HVPS is located. LZ made the unique decision to place the feedthrough itself on the side of the outer cryostat vessel. A popular choice for TPCs is to place the feedthrough on the top of the vessel, and extend a rod or cable vertically downwards to the cathode position. The extension is usually a coaxial design, with a high voltage wire surrounded by an insulating layer and grounded outer conductor. In order to minimize the electric field magnitude, the layers terminate in succession, with the ground layer ending first, followed by the insulating layer, leaving the bare high voltage wire to make the final connection. This technique benefits from simplicity, but comes at the cost of enhanced electric fields in the region outside the field shaping cage. Feedthroughs of this type are shown in Fig. 3.2.

The higher fields resulting from this design are problematic if one desires to instrument the region outside of the field cage. A top-down design impedes the light collection, generates high fields between the cable and the PMTs, act as an additional background source, and establishes a complicated, nonuniform electric field. The primary issue is the increased risk of electroluminescence near the connection, due to the small standoff between the high voltage and the inner cryostat vessel.

Since LZ aimed to instrument the “skin” region in order to reduce external backgrounds, these issues had to be resolved. One possible solution would be to simply increase the standoff between the field cage and the inner cryostat vessel. The increased distance at the same field would reduce the field, at the cost of additional xenon, making it uneconomical. Instead, LZ opted for a novel design of making the connection horizontally from the side, making consecutive seals to the inner and outer cryostat vessels (ICV and OCV, respectively).




Figure 3.1: A side-view cross section of the LZ Cathode High Voltage (CHV) feedthrough. Visible is the HV cap (purple) and spring connection on the right hand side. The rings (pink and gray) form the grading structure which maintains a relatively uniform electric field. The smaller and larger cones seal to the ICV and OCV, respectively. The white cylinder is the “Xe displacer” which prevents LXe from coming into contact with the regions of highest electric field. The yellow lines indicate cooling straps which help to regulate the temperature of the cone. For scale, the xenon displacer (the white cylinder and cone) structure is approximately 30 inches long. Rendering by Ethan Bernard.

The horizontal feedthrough design consists of a single high voltage cable which connects the HVPS to the cathode through a set of concentric cones extended outwards from the ICV and OCV. The cone shape allows the ground and insulation layers of the cable to end relatively close to the seal while maintaining a lower peak electric field. On the surface of the inner cone (connected to ICV) copper cooling straps were placed to help maintain the temperature at that of the rest of the ICV. Multi-layer insulation was also wrapped around the inner cone in order to suppress bubble formation, which would weaken the dielectric strength. The conductor itself is carbon-doped polyethylene, which greatly reduces the risks from mechanical stress when cooling from room temperature to liquid xenon temperatures (177 K at 2 bar). To further limit the field magnitude, a linear potential grading is established by a resistor-divider network formed from a series of doped polyethylene rings. The structure is shown in 3.1.



Figure 3.2: Photographs of examples of traditional feedthrough designs for dual-phase TPCs. From top to bottom: the designs for DUNE far-detector, Darkside, and CAPTAIN TPCs. Photograph taken from [102].

The high voltage cable makes a direct connection between the HVPS and the cathode ring located in the inner TPC volume. Shortly past the end of the inner cone, the ground layer terminates, flush with a polyethylene xenon displacer, which prevents direct contact between the LXe and the cable's insulating core. This increases the standoff distance between the LXe and the regions of highest electric field, reducing the risk of electroluminescence and breakdown. The fit between the cable and the xenon displacer is virtually gapless, achieved through a cryofitting procedure which inserted the slightly oversized cable into the bored hole. The high voltage wire itself makes contact with a stainless steel detent, which captures a conductive plastic cap. A spring is inserted into this cap, and makes the final connection to the cathode ring. Using a flexible connection makes for a simpler installation procedure, and limits the mechanical stresses from thermal contraction. It separates the concerns of aligning the inner cone seal and the high voltage connection itself, while vastly improving the tolerance of the design.

The small current flowing into the cathode can return to ground through the TPC forward field region resistor-divider network, the reverse field regions resistor-divider network, or the CHV grading structure. The grading structure is another resistor-divider network which reduces the peak field on any one part of the design. The large, round end cap is held at the cathode voltage and lacks sharp points which can enhance the field. Concentric, conducting plastic rings form the rest of the network, with $1 \text{ G}\Omega$ resistors connecting them. The large, round shape of the rings is again to avoid sharp, proud points which can enhance the electric field. This design limits the peak electric field magnitude to 35 kV/cm at a cathode voltage of 100 kV .

Because made of plastic, the cable is flexible, and when warm can be fit with an o-ring seal to separate volumes. The two spool can seals at the top of the water tank, one air-to-vacuum and the other vacuum-to-gxe, are made this way. Because the plastic loses pliability at LXe temperatures the seals could not be made there, so for the length of the water tank the cable is routed through a vacuum-jacketed (VJ) hose to the spool can. Rather than make a seal between the VJ hose and the OCV, The OCV vacuum space extends the length of the water tank through a bellows which surrounds the VJ hose. This thermally insulates the VJ hose, and adds extra protection against any ingress into the xenon space.




Figure 3.3: CAD rendering of the LZ water tank with components relevant to the CHV feedthrough annotated. This illustrates the scale of the CHV feedthrough relative to the LZ OCV and water tank. Note the turn the bellows makes, as well as the “spool can” at the top of the water tank, which forms the seals to the outside. Rendering by Ethan Bernard.

Unfortunately the polyethylene emanates radon, which is exacerbated by the large distance spanned while in xenon. This is mitigated through several methods. Firstly, there is a flow restrictor immediately inside the inner cone, which limits the exchange of LXe between the TPC and the VJ hose. An FEP sleeve is fit over the cable, acting as a guard against radon emanation by routing the gasses away from the TPC and into the circulation loop, where it passes through an inline radon reduction system. A small diameter tube runs partially down the outside of the sleeve, allowing some Xe to bypass the radon removal system and connect to the detector cooling and circulation system.

The connection of the VJ hose to the warm end is also topologically complex. In order to protect the VJ hose from the water in water tank, a protective bellows surrounds, joined with the outer cryostat vacuum. The VJ hose has a risk of softening over time, necessitating access to a pumpout port. These constraints require the installation of a CF tee connection between the reverse side of the VJ hose flange and the outer bellows flange. Because the VJ hose must also connect to the water tank feedthrough, this results in a triple-stack CF seal, with two gold gaskets on either side of the VJ flange.




Figure 3.4: The results of a finite element field simulation in Maxwell, performed by then graduate student Evan Pease. The region of peak field is seen to be the corner of the first field shaping ring after the HV cap. This simulation was performed assuming a -100 kV cathode voltage.

Electroluminescence

In LXe the threshold for electroluminescence is 412 kV/cm[103]. This exceeds the surface fields of the design but can be achieved by asperities on the conducting surfaces. The overall goal of the LAr prototype tests are to ensure that this threshold is never crossed. In XeBrA (Chapter 8) it was observed that it is possible to observe enhanced light production without breaking down, at least over a period of several minutes.

The TPC is separated from the CHV feedthrough by the PTFE wall and field cage. Significant light leakage was not observed between the two regions. As such the main area of concern is either degradation of the skin tagging efficiency, or loss of exposure, as a result of elevated light rates.

3.3 LAr Prototype Test

Purpose and overview

As explained above, the unique design of LZ's CHV feedthrough was expected to pay dividends by enabling the skin veto within LZ. However it was necessary to verify the integrity of the design as built, in order to inform possible design changes or recommendations for commissioning. A prototype of the feedthrough was machined and tested in LAr with the goal of identifying the maximum voltage which could be applied without generating a single




Figure 3.5: An annotated diagram of the spool can connection / warm feedthrough at the top of the water tank. The interior of the spool is in the GXe space, and is topologically connected to the TPC. Inside the circle, the cable makes a loop, which allows it to contract as the temperature drops and reduces mechanical stress. On the side of the can are small outlet tubes, one of which connects to a bypass tube, and the other goes to the inline radon reduction system. The cross at the top is evacuated and forms the warm o-ring seals. At the bottom, the GXe-to-vacuum seal is formed, then the cable makes a right angle and feeds through the vacuum-to-air seal. The remaining plumbing exists for pumping out the cross and for leak checking purposes.

photon background, as measured with a large area PMT. The design of the prototype was identical to the intended production design, save for those features necessary for installation (e.g. bolt holes, access ports for leak checking, seals between the inner and outer cones). Only the inner cone (which attaches to the ICV and is filled with xenon during the lifetime of LZ), was tested, as the outer cone will not experience any large fields. LAr was chosen as a proxy for LXe due to its low cost and similar dielectric strength[104, 105].

The metric for success was the maximum voltage held for at least one hour. It was intended to reach the LZ design voltage of -100 kV for this period of time. Two runs were performed in order to isolate features in case of failure. The first run tested the grading structure, while the second tested the spring connection.




Figure 3.6: A photograph of Harwin pin which connects the resistors in series between the CHV grading rings. Note the slightly canted edges of the $5\text{ G}\Omega$ resistors. These are particular high voltage risks due to the proximity of the voltage drop between the rings and the pins.

Grading Structure

The construction of the prototype grading structure took place in several steps. First, a section of cable was removed from a spool and a short section, approximately one foot long, had its outer conductive layer removed. This section was chilled in an environmental chamber to shrink and stiffen the plastic. The polypropylene Xenon displacer was heated with a heat gun to exploit a small amount of thermal expansion. The cable was quickly cryofit into the displacer in order to make a gapless connection.

Taking place inside a class 100 clean room, the grading rings were placed one by one over the exposed insulating layer on the polyethylene cable. Between each ring was placed a delrin spacer to establish the fields. Each division had four $5\text{ G}\Omega$ resistors placed between them, in two parallel sets of two. The resistors were attached to each grading ring with Harwin pins, shown in Fig. 3.6.

The cable was connected to the high voltage power supply. The supply itself was a Cockroft-Walton generator which multiplied the high voltage sourced from a Spellman power supply.

Test Stand

The prototype CHV feedthrough was tested in a 16 inch inner diameter dewar. Threaded rods suspended the entire structure from the lid. An electropolished stainless steel cone stood in for the production cone. All wiring was similarly routed through the top. In the first run the structure terminated in the HV cap, instead of completing the electrical connection. For the second run, an electropolished stainless steel mock spring extended from the plastic HV

cap. This took the place of the spring connection from the grading structure to the cathode ring itself, but due to space constraints it remains on axis.

The CHT test stand was instrumented with a Hamamatsu R5912-02-MOD PMT[106], and two Cremat CR-150[107] charge sensing circuit boards. These three channels were digitized using a CAEN 1720 DAQ with custom software. The base of the PMT circuit was modified from a previous configuration to support cathode biasing, and furthermore to saturate at a lower voltage. The latter modification was necessary due to the potential for copious amounts of light to impinge on the photocathode during the testing. The resulting circuit board is shown in Fig. 3.9.

The charge amplifiers were put in place primarily to identify the location of breakdowns if one were to occur during testing. With this knowledge, interventions could be taken as necessary if the goal voltage was not achieved. The two modes that were distinguished were breakdowns within the grading structure, and breakdowns from the grading rings to the inner cone. Each location can be seen in Fig. 3.10, where the second “cone” is the ground cone which surrounds the high voltage cap terminating the grading structure.

The circuit employed in the test is shown in Fig 3.11. The two charge amplifiers were Cremat CR-111 charge sensitive amplifiers, with a gain of 130 mV/pC. These were simulated in PartSim¹ to find the response to shorts across various locations in the system. It was observed in simulations (and in data) that there was significant cross-talk between the channels due to the capacitive connection between the cone and the grading structure. However, the simulations predicted a difference in polarity between the two modes. Charge entering the mount through the grading structure show up as negative current (and therefore, positive voltage signals due to the inverting amplifier) in that particular charge amplifier, while breakdowns to the cone showed up as negative current in the other charge amplifier. The results of the simulations are shown in Fig 3.12.

In addition to the digitizing of the PMT voltage, the signal was split and drove an analog Ortec ratemeter. This allowed for real time monitoring of small spikes in the light production, guiding the ramp rate. When the ratemeter fluctuated too much in a short amount of time the voltage was immediately reduced.

Spring connection

In order to test the ergonomics and feasibility of making the spring connection on site, a practice rig was constructed at LBL. A slide was assembled from 80-20 pieces which allowed the prototype grading structure to slide forwards and backwards. The cap and spring was placed at the other end of the rig. A steel plate was installed to mimic the clearance the operator would have while working on the seal on the OCV. This successfully demonstrated that an operator could reliably stretch the spring and make the final connection. My role on this portion of the project was assembly of the rig, and as an independent data point confirming the ease of the eventual procedure.

¹Partsim.com




Figure 3.7: A photograph of the CHV prototype test stand. *Left*: The design in Run 1, without the spring connection. The baffles which contain the cables are seen, along with the clean room tent. This entire column is lowered into the argon dewar with a crane. *Right*: The design used in Run 2, with the mock spring connection. The PMT and HV cap can be seen in finer detail.




Figure 3.8: Schematic of the charge sensing circuit for the CHV feedthrough test.

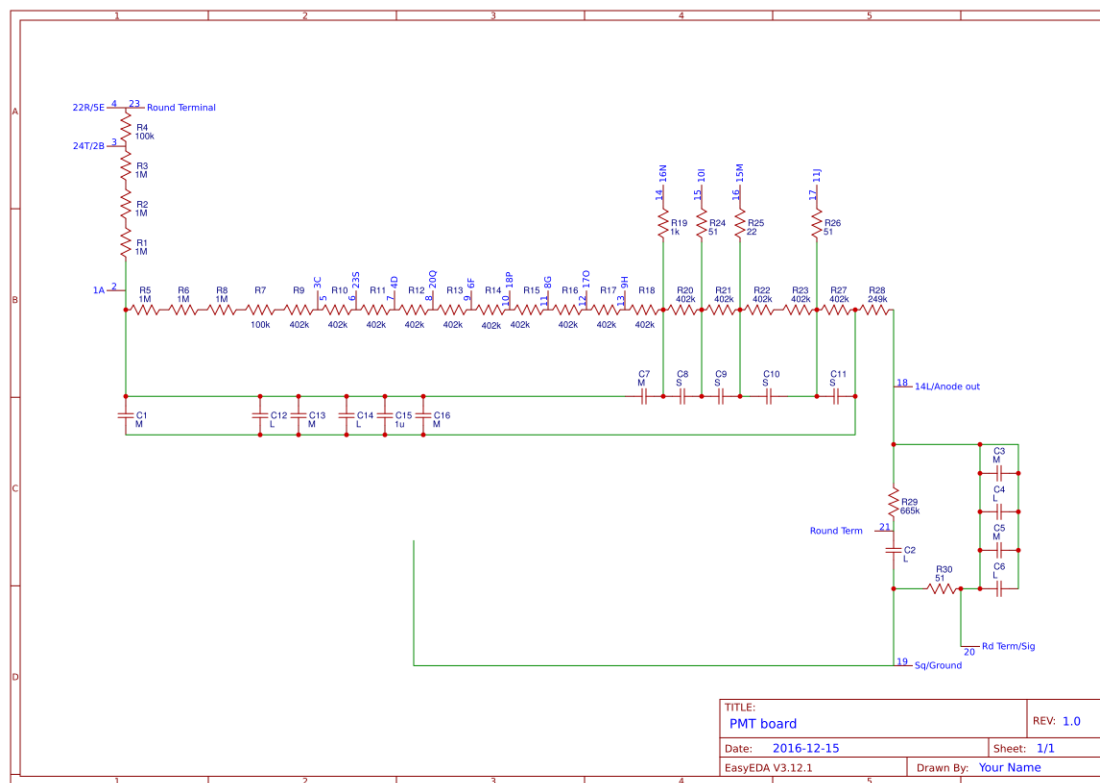


Figure 3.9: The schematic for the base for the CHV photomultiplier tube. Note the capacitors in parallel with the final stages of the voltage divider. This help to keep the PMT response linear to higher numbers of photoelectrons. The anode signal is read across the 51Ω resistor on the right side of the diagram. Further details on the base design may be found in the Hamamatsu guide[108].

Signal Processing

Strategy

The signal processing procedure consists of the following algorithm for each event window.

1. Attenuate the noise from the signal and find the pedestal from which to measure the pulses.
2. Locate the pulse boundaries for PMT signals.
3. Cluster the pulses together and merge them appropriately.
4. Calculate the reconstructed quantities for the merged pulses.

This is performed on the PMT channel. For the charge amplifier channels, the clustering and noise reduction steps are omitted. This is due to the fact that the impulse response of




Figure 3.10: CAD rendering of high voltage test structure and surrounding instrumentation. *Left:* The structure during Run 1. *Right:* The structure during Run 2. Note the addition of the “mock spring” which extends out from the cone.

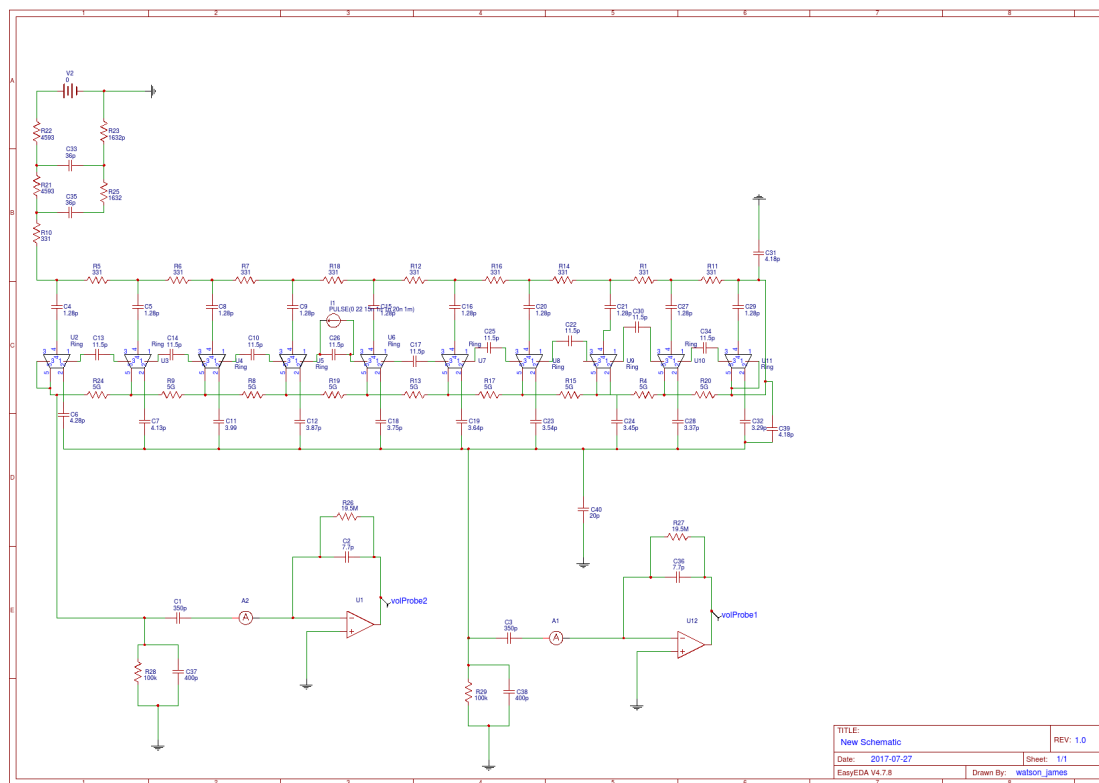


Figure 3.11: The schematic of the charge sensing circuit for the CHV prototype test. The trapezoids are custom components representing the resistance and capacitance matrices of the grading rings. The connections on the bottom of the diagram represent the charge-sensitive amplifiers. One amplifier is connected to the inner cone, while the second is connected to the lowest potential grading ring. Also visible is a current source which simulates the effect of a short between the grading rings.

the charge multipliers are much longer than that of the PMT, obviating the need to detect the merging of short pulses.

Noise filtering

The unamplified voltage signal from the PMT was digitized by a CAEN 1720 DAQ board and stored on disk using the custom software *PixeyCalc* [109]. The DAQ operated at a frequency of 250 MHz, allowing each PMT pulse to span several 4 ns samples. Data was taken using a heartbeat trigger, wherein a fixed acquisition window was written to disk whenever a square wave signal from a Rigol function generator crossed a specified threshold. The trigger rate and acquisition window was approximately tuned in order to maximize livetime. This was conducted by adjusting the frequency on the function generator until the *PixeyCalc* software failed to yield an event rate identical to the input threshold crossing rate. Final values were

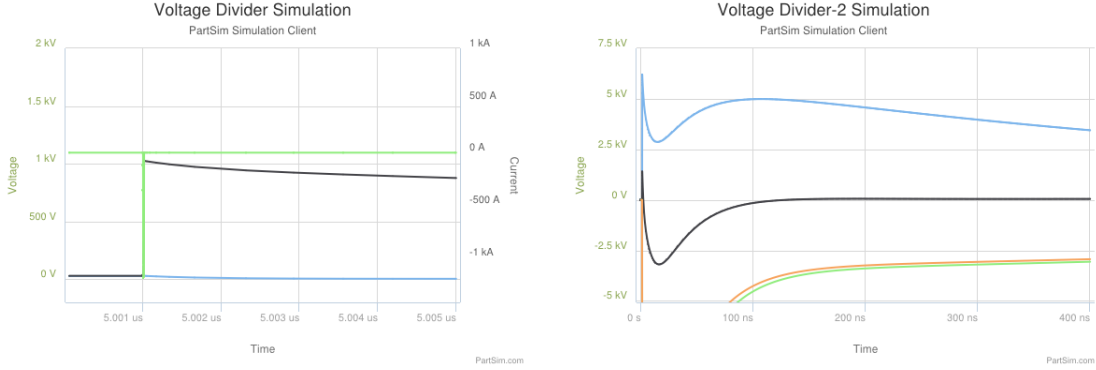


Figure 3.12: Results of SPICE² simulations using the schematic in Fig. 3.11. The first figure shows a short across the first set of 5 G Ω resistors. The second shows a short from the end of the grading structure (referred to as the “mushroom head” to the inner cone. In both cases the black voltage trace indicates the mount charge amplifier and the black the cone charge amplifier.

a 64 μ s window ($=2^{14}$ samples) with a 650 Hz signal, for a 4.16% total livetime. When additional channels were digitized, the trigger rate was reduced proportionally.

A persistent challenge when analyzing the digitized PMT signal offline was the spectrum of noise present in the signal. Ideally a power spectrum is white, i.e. constant in magnitude across the sampled frequencies. However, in the prototype testing the noise environment exhibited a combination of low and high frequency signals, as shown in Fig. 3.13. The power spectrum P_k itself is defined as

$$P_k = \frac{X_k X_k^*}{\sum_{j=0}^{N-1} X_j X_j^*} \quad (3.1)$$

$$X_k = \sum_{j=0}^{N-1} x_j e^{-i2\pi jk/N}, \quad (3.2)$$

where N is the number of samples and X_k is the discrete Fourier transform (DFT). An example of the appearance of this noise is shown in Fig. 3.14. Significant effort was undertaken to reduce this error, in order to achieve a high single photon efficiency. A bandpass filter method, where the frequencies found in the noise spectrum was tried initially, but this was found to distort the PMT single photon response in an undesirable way. Since the noise spectrum is dominated by a small number of pure tones, it was decided instead to fit a sinusoid $A \cos(\omega t - \phi)$ to the data. The amplitude returned by this method, A , differs from the Fourier analysis components in that ω is unconstrained, no longer necessarily a multiple of $2\pi/N$. This is similar to the Lomb-Scargle periodogram [110], which is used for irregularly spaced data.




Figure 3.13: The noise power spectrum for Run 1 of the CHV Prototype tests. There is a significant DC component, along with numerous sharp high frequency peaks.

Prior to the Lomb-Scargle periodogram, the DC component had to be subtracted. This was done by finding the mean and rms of the waveform, masking the outliers (more than three standard deviations from the mean), and repeating the process until no more outliers are observed. The outliers are recalculated at each step in order for large pulses to not bias the initial step.

The fitting procedure employed in the analysis consisted of two stages: one to find the optimal relative values globally, and then to actually fit the template to individual waveforms. The baseline model consisted of a sum of N_c cosine functions, each with parameters A_i , and ϕ_i . The frequencies were taken to be harmonics of a fundamental frequency, i.e. $\omega_i = (i + 1)\omega_0$. An initial pass through the data was performed, removing outlier data points so that the waveforms were free of SPE pulses. Then, an unconstrained fit was performed with the harmonics. The fitted values for each waveform were collected, and the mean values for A_i , ϕ_i , and ω_0 were found. Using these results, the relative values, A_i/A_0 and $\phi_i - \phi_0$ were frozen. This effectively created a noise template of fixed shape, with remaining free parameters A_0 and ϕ_0 . The analysis to follow fit this noise template to the outlier-removed waveforms, and then subtracted this function from the waveform so that the residuals could be searched for PMT pulses. An example of the procedure can be seen in Fig. 3.15, and

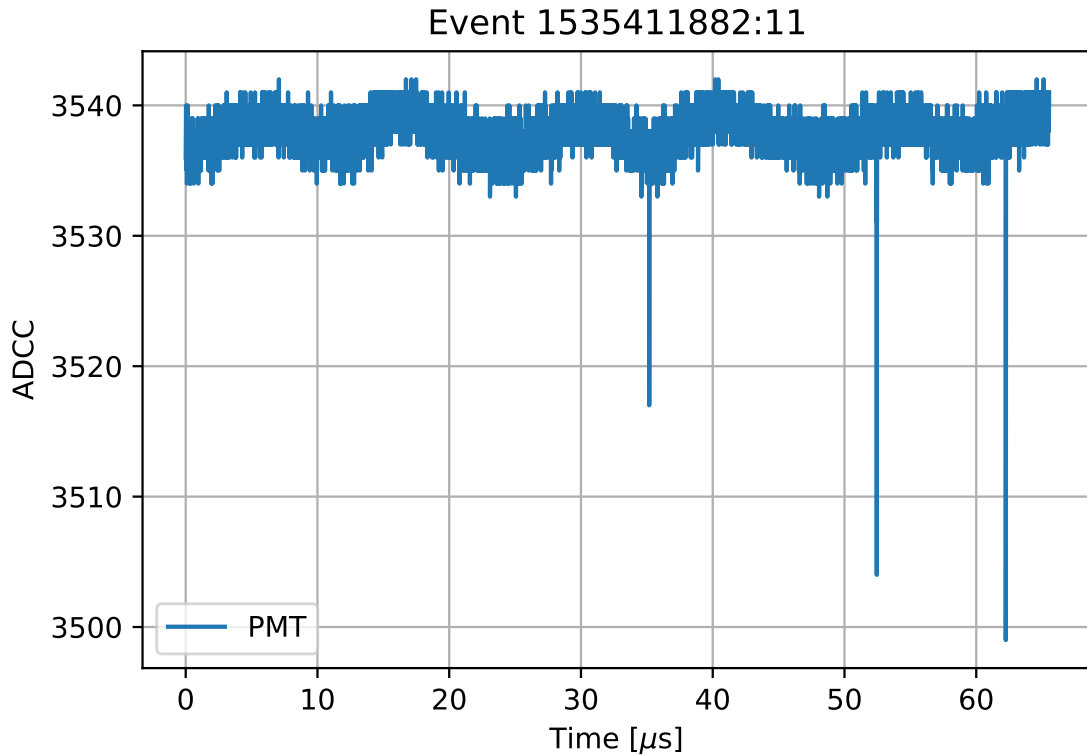


Figure 3.14: A representative waveform from a CHV prototype test stand acquisition. This particular event is from an LED calibration. The DC roll of the pedestal is visible, along with three apparent detected photons.

the effect on the noise spectrum is shown in Fig. 3.16. After some optimization, a value of $N_c = 4$ was determined to provide adequate results. This achieved excellent low frequency noise reduction with minimal distortion of the signal pulses.

A short “chirp” was observed in some, but not all, datasets. An example is shown in Fig. 3.17. These events are barely above the noise threshold. It is possible that it is some encoded cell phone traffic or similar digital signal. The carrier frequency was located in the Fourier transform at 7.95 MHz. A wavelet of with that frequency and a cosine envelope was created and convoluted with the waveform, revealing regions contaminated by the chirp signal. Pulses inside or nearby those regions discarded as being too unreliable. This was found to reduce the size of the background by a factor of ~ 2 , but it was not enough to completely eliminate it. In any case, the size of this signal does not affect the calibration greatly, though it does pollute the data.

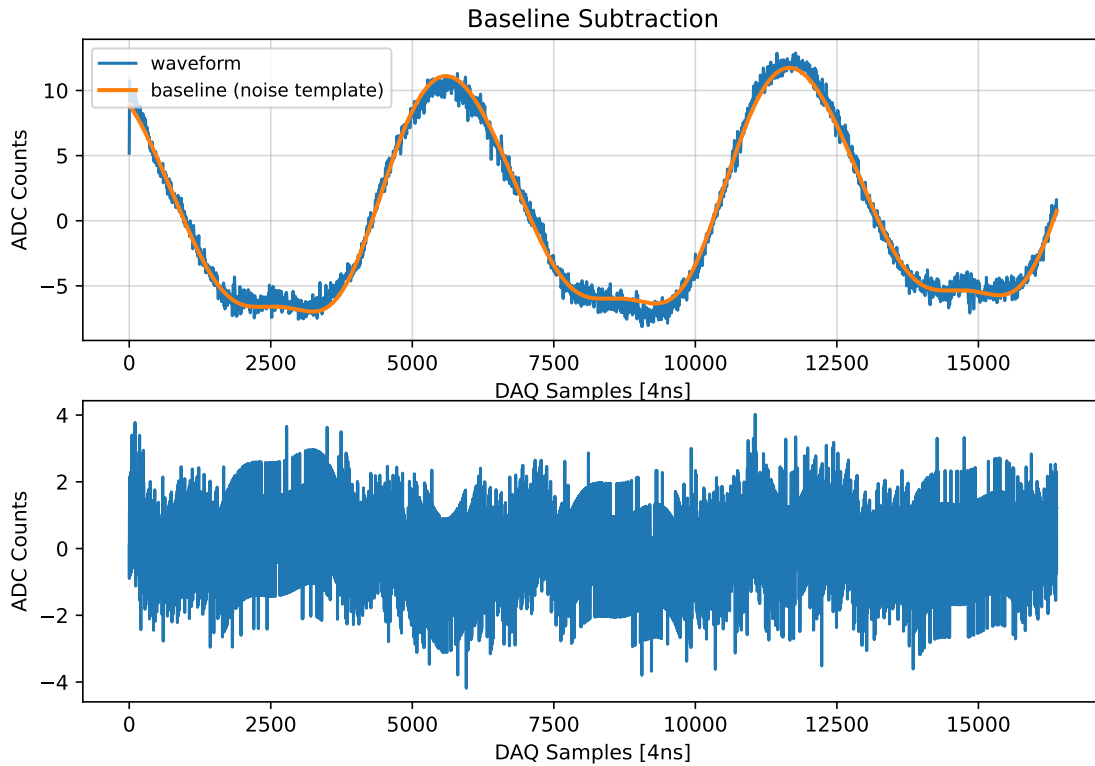


Figure 3.15: Waveform results from the baseline subtraction technique. *Top*: the raw waveform (blue), along with the fit sinusoidal harmonics (orange) which accurately follow the slow component. *Bottom*: The residuals from subtracting the Lomb-Scargle pedestal.

Pulse Finder

No matter the methodology for baseline removal, a signal model must be used in order to not also filter out the PMT response. This was obtained by examining LED calibration data and running a simple threshold-based pulsefinder. This resulted in pulses which were biased upwards in amplitude. Locations with amplitudes below a threshold were selected, and shifted in time such that the minimum of the pulses were aligned. In order to improve the spectral resolution, eliminate boundary effects, and reduce the overall length of the filter, the impulse response was multiplied with a window function. The spectral response of a finite impulse response is the convolution of the signal with the window it was observed over. For a basic truncation at some length, this results in a rectangular window, which has the frequency response of an *aliased sinc function*, $\text{asinc}(\omega M) = \sin(M\omega/2)/M \sin(\omega/2)$. The rectangular window has the narrowest main lobe, but worst frequency sidebands. Other windows attenuate the sidebands at the cost of a wider main lobe. For my purposes I found that a cosine window, where $H(t) = \cos(\pi t/2M)$, was able to adequately smooth out the impulse response boundaries. The averaged pulse converged to the true impulse response of

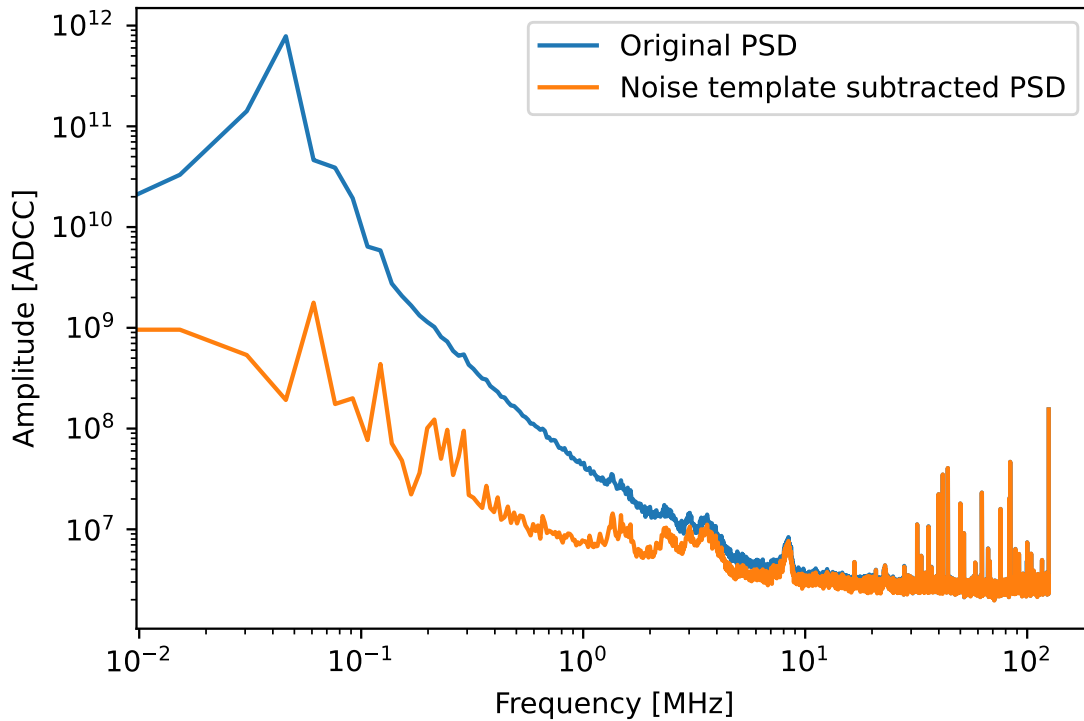


Figure 3.16: The noise power spectrum (FFT) before and after the baseline subtraction technique. The spectrum is forms from the average over many power spectrums $|F(\omega)|^2$. Note the much stronger attenuation of the DC component.

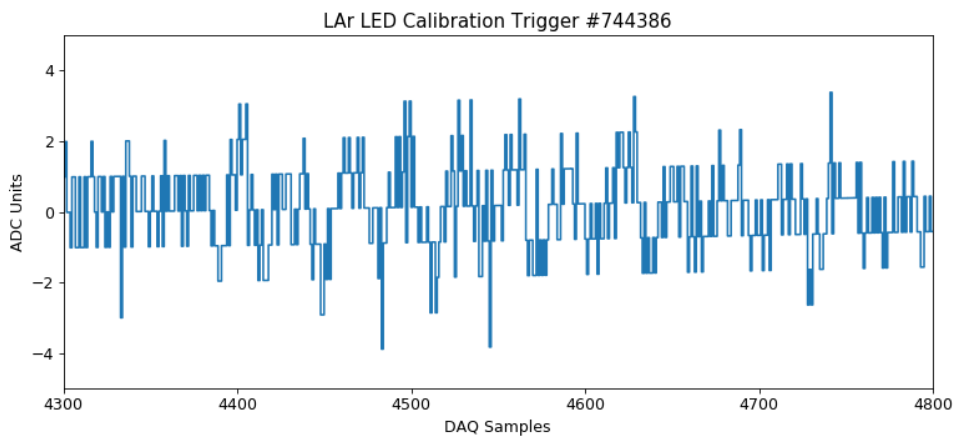


Figure 3.17: A representative “chirp” waveform which was observed in the data. These chirps require an additional technique to remove, as they are barely above the noise level, but can skew the results of the remaining analysis.




Figure 3.18: The Hamamatsu PMT impulse response, obtained from CHV prototype waveforms. Note the double peaks in the first 50 ns, following by an AC pulse some time later. This pulse is possibly due to improper termination of some cable, though this phenomena was taken into account during the design process for the electronics chain. This template is consistent with previous characterizations of the R5912-02 MOD PMT[111].

the PMT-DAQ system, shown in Fig. 3.18.

One method for pulse finding used the “optimal filter” method, similar to the method used in the Edelweiss surface run[112]. This consists of constructing a finite-impulse-response (FIR) filter using the impulse response H of the PMT to the PSD N^2 of the noise spectrum, as described in Eq. 3.3. This filter is normalized such that the convolution of the filter with itself is one. By convolving this optimal filter with the input signal, the resulting waveform indicates regions of high similarity with the template, which can then be identified as pulses. Since much of the noise is low frequency, this design has the side effect of removing much of the DC signal. Thus, the filter response is not in general suitable for reconstructing the pulse areas and amplitudes.

$$F(\omega) = \frac{H^\dagger(\omega)}{N^2(\omega)} \frac{1}{\int d\omega H^\dagger H / N^2(\omega)} . \quad (3.3)$$




Figure 3.19: The optimal FIR filter for the CHV PMT response. The left plot shows the filter and its inputs in frequency space, while the right plot shows the filter in the time domain. Note that the optimal filter mostly resembles the template response, with some distortions. The filter is shifted here so that the peak is in the same location as the template.

The result of this equation as applied to the PSD and single photon response in the CHV test stand is shown in Fig. 3.19. Following convolution with the optimal filter, regions of high amplitude are identified as pulses. An additional step is used when accurate timing information is desired, such as when the triplet lifetime was being measured. This is accomplished by minimizing the χ^2 between the optimal filter and the waveform, but with the filter shifted continuously in time through convolution with the following all-pass filter:

$$T_{\text{shift}}(t) = e^{-i2\pi kt/N} . \quad (3.4)$$

This affords sub-sample timing precision, along with more accurate amplitude reconstruction due to minimizing the time offset between the filter and the signal.

Due to the computational cost of this procedure, it was only employed for specific sub-analyses that required this level of precision (e.g. the triplet lifetime analysis). In other cases, a simpler, threshold-based algorithm is used. This algorithm consists of setting an ADCC threshold and collecting the indices where the voltage signal crossings. Upwards-going crossings are merged with nearby downwards going crossings in order to be robust against fluctuations. This proceeds iterates until no more merge candidates exist. Pairs of downgoing and upgoing crossings are then grouped together into pulses. After some testing a minimum distance of 5 samples between adjacent pulses was established, with a minimum number of samples below threshold of 2.

Pulse Merging

The pulse finding algorithm described above is adequate for finding the spikes corresponding to the PMT impulse response. However, due to the issues of reflection, combined with the

particular shape of the impulse response (seen in Fig. 3.18), it is necessary to perform an additional clustering step to determine if adjacent peaks are actually from the same scintillation signal. This is accomplished through a procedure known as *deconvolution*. If the impulse response of a system is $H(s)$, and the input is $X(s)$, and $Y(s) = X(s)H(s)$ is the response of the system, then deconvolution attempts to find the most likely value of the input, $X(s) = H^{-1}(s)Y(s)$.

With a noiseless system with a finite impulse response, this process is as simple as presented above. In the presence of noise, and when the signal response might run past the window, this becomes far more challenging. Two popular techniques which were explored for this purpose were the Weiner filter[113] and the Richardson-Lucy deconvolution[114]. Additionally, a fit-and-subtract method was attempted, using the χ^2 minimization technique described above. This involved fitting the pulse using a continuous time shift, subtracting the fitted impulse response from the signal, and then repeating the process until no prominent pulses remained. This was found to be prohibitively slow, and having a higher error rate, than the other two techniques. All three strategies were initially attempted in a global fashion, i.e. replacing the initial pulsefinder. However, this proved to be computationally expensive, and therefore these strategies were then incorporated into the *narrow phase* of the pulse finding, which checked if nearby pulses were to be merged. The *broad phase* consisted of either the threshold-based pulsefinder, or the optimal filter, depending on the analysis.

The nearby pulse merging was resolved via Richardson-Lucy deconvolution, which was found to reconstruct the delta functions of pulse arrivals better than the Weiner filter, which tended to result in more oscillatory behaviour near large pulses, the very phenomena that I intended to remove. This procedure works with the point spread function p , the observed signal y , to estimate \hat{x} , the posterior probability for the underlying source, such that $y = \hat{x} \star p$. This procedure begins with an initial prior $x^{(0)}$ and iteratively refines the posterior $x^{(i+1)} = f(x^{(i)})$, with the full equation given by:

$$x^{(i+1)} = x^{(i)} \cdot \left(\frac{y}{x^{(i)} \star p} \star p^T \right). \quad (3.5)$$

The base algorithm is further modified to exploit the expected signal. The first modification consists of specifying a noise threshold. All samples with amplitudes below the noise threshold are set to zero. This exploits the fact that the signals are only a single polarity. Setting the threshold at a nonzero value has the additional benefit of preventing the RL algorithm from inadvertently amplifying on noise and producing instabilities at later iterations. The second modification is that a smoothed waveform is periodically generated in order to capture the low frequency noise, and not force the RL algorithm to estimate it using the point spread function. This is accomplished by periodically (every 5 or so iterations) taking the residuals and applying a butterworth filter, and then adding them to the current estimation of the pedestal. This pedestal is subtracted from the signal before each iteration of the RL algorithm. The last modification is a standard tactic in deconvolution, which is the use of regularization. The second derivative of the signal is calculated, and this curvature value is used to attenuate the residuals that the RL algorithm attempts to fit. Thus, the error values at each point become $\epsilon_i = (\hat{y}_i - y_i)/(1 + \lambda \nabla^2 y|_i)$, where λ is a hyperparameter which controls




Figure 3.20: An example event waveform with the zero-suppressed Richardson-Lucy deconvolution algorithm applied. This shows the benefits of noise suppression and allows for higher confidence in identifying pulses close to one another.

the strength of regularization. The results of this algorithm is shown for a representative multi-pulse event in Fig 3.20. After the deconvolution is calculated, prominent peaks are located and identified as individual, distinct pulses.

Data selection criteria

Following the subtraction of the baseline and the pulsefinding procedure, pulses within $\Delta t = 1.6 \mu\text{s}$ in time were grouped into clusters. The clusters were then classified into three categories based on a cut on M .

1. **Single photons** / coincident single photons. These are events with $N_{\text{pulses}} < M$. They are considered to be accidental coincidences and counted as N_{pulses} SPEs.
2. **Multiple photoelectrons** (MPE). If $N_{\text{pulses}} \geq 3$, the pulses are considered to be causally related, because the probability to observe 3 pulses in the same trigger under normal circumstances is astronomically small. An SPE rate of ~ 200 Hz was typical, making the expected number of SPEs in a waveform 0.01.

3. **S1s.** These are a subset of MPEs that are suitable for triplet lifetime measurements. They have total length $> 1600 \mu s$ and $N_{\text{pulses}} > 10$. These events are selected to fail the long-timescale production, described below, in order to distinguish scintillation from electroluminescence.
4. **Long-timescale light production.** These are the events which are most likely to be caused by charge moving as a result of the high-voltage structure. Events with an center of mass/ first moment $> 2 \mu s$ and $N_{\text{pulses}} > 100$ are selected, and are considered on a waveform-by-waveform basis. This means that there are a maximum of one long-timescale light event per waveform. This is motivated by the fact that the charge amplifiers respond slowly, and in those channels only one charge event can be counted per waveform. The presence of long-timescale light events in a PMT channel are compared to the presence of charge events in order to look for correlations. These cuts were chosen to provide low false-positive rates but also to correlate with large spikes in the Ortec ratemeter.
5. **Charge amplifier events.** Similar to the long-timescale light production, these are counted on a trigger-by-trigger basis due the slow response time. The charge amplifier signals were found to have a low-frequency noise component which grew in proportion to the voltage across the grading structure. The charge amplifier response would happen at similar timescales to the noise, which frustrated attempts at a filtering approach to locating charge events. Instead, it was observed that the mean value of any particular waveform did not change in proportion to the high voltage power supply, *i.e.* there was no bias. As such, the mean value of the waveforms became the discriminating value to find charge events. Due to the low-frequency noise the estimators of the mean \bar{y} became non-gaussian, so the standard deviation underestimated the errors. The distribution was double-peaked, and changed as a function of grading structure voltage. The final selection criteria was then chosen as $|\bar{y} - \text{med}[\bar{y}]| > \min(2 * \text{IQR}[\bar{y}])$, where $\text{IQR} = Q3 - Q1$ is the interquartile range of the distribution, estimated in bunches of data approximately 5 minutes long.

An example of a waveform satisfying criteria (4) and (5) is shown in Fig. 3.21.

Calibration

LED Calibration

At two points during the testing, the PMT was calibrated, with an aim towards locating the single-photo-electron (SPE) peak. The first calibration took place between Run 1 and Run 2, while the dewar was filled with nitrogen Gas. The second took place following the conclusion of Run 2, while the dewar was filled with a combination of LAr and GAR vapor as it boiled off. An LED with a white phosphor was placed in the dewar near the top of the baffles and was triggered with a function generator. The signal from the function generator was split and triggered the DAQ. The trigger location was placed in the middle of the waveform so




Figure 3.21: An event waveform classified under criteria (4) and (5), showing all three channels: PMT and both charge amplifiers. This event illustrates activity in all channels, i.e. the PMT and both charge amplifiers. As such it is likely the result of a breakdown in the xenon.

that the background spectrum could be subtracted from the LED-on spectrum. Examples of this procedure for the two calibrations are shown in Fig. 3.3.

The event rate in nitrogen gas was much higher than the event rate in LAr, and as such, the analysis was very sensitive to the subtraction of the two spectra. An example of these spectra is shown in 3.3. After fitting the difference between the LED on and LED off histograms to a sum of several Gaussians, the single photoelectron gain was extracted. The model used was such that there were two background Gaussians with unconstrained mean, standard deviation, and amplitude. The last Gaussian is interpreted as the single-photoelectron peak, corresponding to a gain of $\sim 4.6 \times 10^8$. The second Gaussian is interpreted as some anomalous reduction of gain in the SPE peak. This skipping could occur by a photon striking the first dynode, or electrons missing a stage due to some unknown effect. Due to the uncertainty on the second peak's location, a rather poor constraint on the interstage gain is found at 11.9 ± 52 . The Hamamatsu specifications list 4-5 as the interstage gain. The χ^2/DOF of this fit is 1.29.

When statistics were sufficient at a particular voltage, a 2-PE peak was added at $\mu = 2\mu_{\text{SPE}}$, $\sigma = 2\sigma_{\text{SPE}}$, with unconstrained amplitude. This was repeated at several PMT photocathode voltages and the results are shown in 3.3.




Figure 3.22: The charge amplifier impulse response. This was produced by completing the sense circuit with a fast switch, sending a quick surge of current into the CR-111.

After the conclusion of Run 2, the PMT was again calibrated on 08/27/2018. At the time the LAr level height meter read 1312 pF. This time, the dewar was still filled with mostly liquid argon, so the characteristics of the PMT were most similar to what they were during testing. The LED off background was negligible compared to the LED on signal, in stark contrast to the room temperature data. Calibration data were taken at a single voltage, -1550V, the same voltage used in the high voltage ramps.

A previous calibration of this model of Ref. [115] demonstrated the temperature dependence of the location of the SPE peak. That calibration was made with a slightly different base design using a grounded cathode and a 1500 V anode bias. Because the base designs were different, it was decided to simply verify that the *relative* drop in temperature was the same between our test and the Nikkel data. The SPE curves in our room temperature calibration lack the large populations of smaller pulses we find in the LAr calibration. The results of that calibration are shown in Fig. 3.3. The gain obtained at this work at room temperature is clearly lower than that obtained in Ref. [115].

After fitting the room temperature gain results to a power law $y = (x/a)^b + c$ (physically motivated by the discrete number of gain stages in a PMT), the gain at $V_{\text{cathode}} = -1550$ V was interpolated. This was then compared to the calibration data taken at -1550 V cathode at LAr temperatures. A ratio of $\alpha \equiv G_{300K}/G_{85K} \sim 1.7$ was obtained. In Fig. 3.3, a vertical line for the liquid argon boiling point is shown, along with a horizontal line for the $T = 300$ K gain from [115], divided by α . The vertical, horizontal, and data lines intersect,




Figure 3.23: The Pulse area spectrum for the two PMT LED calibrations. Double Gaussians + pedestal (either constant or exponential) were fit to the histogram. *Top*: the calibration taken in nitrogen. *Bottom*: the calibration taken in liquid argon after Run 2.

demonstrating that we observe the same temperature effect as in Ref. [115].

Triplet lifetime

Argon scintillates via the formation of dimers following ionization and excitation of argon atoms. The dimer then decays into two ground state argon atoms and a VUV photon. Argon is largely transparent to its scintillation light because the decay occurs from a molecular state, which has a different spectrum than the atomic state, and the three body interaction of two argon atoms and the photon is exceedingly unlikely to occur. The lowest energy states are the singlet and triplet states, with lifetimes of ~ 1 ns and $1.6 \mu\text{s}$ respectively. Triplet excitations happen to be more likely during electron-recoils than nuclear recoils, a fact which some experiments exploit in pulse-shape discrimination analysis. Because the triplet lifetime is so long, the dimers have time to find impurities and undergo a non-radiative decay, a phenomenon referred to as quenching. The effect of this quenching can be seen in the triplet lifetime, which becomes shorter with increasing impurities. The quenching coefficient was measured by [116] to be $0.54 \text{ ppm}^{-1} \mu\text{s}^{-1}$ oxygen equivalent.




Figure 3.24: Overview of the LED background subtraction procedure. *Top Left*: The pulse start time distribution, with the LED trigger indicated with the dotted vertical orange line. *Top, right*: The area spectra of pulses found in the pre-and post trigger regions in liquid argon. *Bottom*: The results of same procedure for nitrogen gas. The LED-off spectrum is subtracted from the LED-on spectrum before the SPE area is extracted from a Gaussian+Exponential fit to the residual.

Due to mechanical failure of the purity monitor, the triplet lifetime was used to estimate the oxygen-equivalent impurities. The datasets used for the triplet lifetime estimation are the zero-field conditions taken when no voltage was applied to the grading structure. The selection cuts for this portion of the analysis are as follows:

1. Event length $> 1.6 \mu\text{s}$.
2. $N_{\text{pulses}} > 15$.
3. Event 1st moment $< 8 \mu\text{s}$.

The event traces were summed, after which a low-pass IIR filter was applied to compensate for an apparent signal overshoot only seen in large pulses, i.e. S1s. The average




Figure 3.25: The R5912-02-MOD gain as a function of applied cathode voltage. This curve was found through LED calibrations, scanning the PMT photocathode voltage in nitrogen and locating the single photon gain from the double-Gaussian fitting procedure described above. The data was taken at room temperature. The “X” indicates the gain measured at the LAr boiling point (only one data point was available).




Figure 3.26: The PMT gain as a function of PMT temperature. Data from Ref [115] is shown in blue. The purple dotted line indicates the gain from Ref [115], reduced by the temperature dependence shown in Fig. 3.3 (≈ 1.7). This coincides with the argon boiling point (orange), indicating that we observe the same temperature dependence. Ref [115] observed a large gain overall, possibly due to the change of HV bias from anode to cathode.




Figure 3.27: The measured oxygen-equivalent impurity level, obtained via the Argon triplet lifetime, over the course of Run 2. An apparent rise in impurity was observed over time, due to a possible ingress of air. Later the contamination decreases for unknown reasons. The boiling off of argon may have changed the dissolved amount of oxygen and nitrogen over time.

waveforms were then fit to an exponential for the time windows $1 - 8 \mu\text{s}$ after the first pulse in each event. The results of these fits are shown in Fig. 3.27. For Run 1, the purity was found to be 10 times worse than the value measured by the purity monitor. A possible explanation for this is the presence of an unknown quantity of nitrogen dissolved in the liquid argon. The argon is obtained at a purity better than 1ppm of any contaminant, and the purifier only reacts with oxygen, so potentially 1 ppm of N_2 could be dissolved at the end of filling. Nitrogen differs from oxygen in both the scintillation quenching factor and the electron attenuation length factor. However, an identical concentration of nitrogen will attenuate drifting electrons orders of magnitude less than oxygen [117], but quench the scintillation by a similar amount ($0.11 \text{ ppm}^{-1} \mu\text{s}^{-1}$) [118]. This could explain how the triplet lifetime implies a worse purity than the purity monitor.

Purity Monitor

The purity monitor failed to operate to specifications during these tests. While passing continuity tests, it did not demonstrate the ability to observe the electron lifetime. One troubleshooting attempt was performed near the end of the testing. The bottom heater was run at 75 W (31 V) to stimulate boiling. An inconclusive cathode signal was observed, along with what was interpreted as microphonics from the boiling of the argon. The fiber feedthrough was then bypassed to improve transmission. A healthy cathode signal was seen after unscrewing the flashlamp and barely touching it to the port, as shown in Fig. 3.28. The oscilloscope was photographed with anode voltage = +2800V and cathode voltage =




Figure 3.28: A photograph of an oscilloscope from a test of the purity monitor performed while bypassing the feedthrough. It shows the cathode charge signal, as integrated by the CR-111 charge amplifier, when a pulse of light from a xenon flash lamp is incident on the gold-plated cathode. In this particular test the threaded feedthrough coupling was bypassed, and the flash lamp was pressed directly against the fiber optic connection. This was attempted due to problems seeing a stable signal from the purity monitor. The success of this test, in that a cathode signal is evident, shows that the feedthrough contributed to these issues, and that the argon purity was sufficient for testing.

-400V. Since a clear cathode signal was observed absent an anode signal, only a weak limit on purity could be set. This is consistent with impurity levels higher than what the monitor can estimate.

3.4 Prototype Testing Results

Run 1

The first iteration of the prototype testing examined the design of the grading structure without the mock spring connection. Each day the dewar mass was examined to ensure that the liquid argon had not boiled off to the point of weakening the feedthroughs. The capacitive level meter was also examined to make sure that the LAr was at an acceptable fill. LAr loss rate was estimated as 2 kg / hour by the dewar scale. The purity monitor struggled with stability, but the presence of any cathode signal indicated purity sufficient for our initial testing. Success was defined as the ability to reach 120 kV without producing excess light for one hour.

The cameras were tested with infrared LEDs, and only one fiber achieved an acceptable focus. Following an attempt to modify the LabView program, the PLC lost the ability to communicate with the logging computer. Due to the limited time span we had before the LAr boiled off, it was decided to proceed with the test, but logging the voltages and ratemeter manually. Following a small modification to the charge amplifier circuit, the test began on May 23, 2018 and ran until May 25, 2018.

During the ramps, the rate almost immediately spiked upwards and stayed high for the duration of the test. This is believed to be due to a fault in the PMT high voltage feedthrough which flickers slightly. This feedthrough is near the top, where the gas pocket forms. Intermittent dips in the rate seem to support this hypothesis.

Though the PLC was malfunctioning, the Ortec ratemeter, coupled to the PMT, was monitored by eye, along with the current and voltage produced by the HVPS. Larger breakdown events would be observable on the HVPS readout. Transient light production events were observed at various times during the test. Whenever the Ortec ratemeter saturated the dynamic range of its scale, the scale was adjusted and the high voltage partially ramped down until the rate was extinguished. Light production could last from seconds to minutes and correlated with movement in the charge amplifier attached to the cone. Charge entering the mount was observed regularly throughout the ramp which was not coincident with observable light production.. This is likely a result of capacitive coupling between the mount and the grading structure. Occasional charge and light events were missed in the ratemeter. Sustained light production in the PMT can appear to be a single pulse, which would not show up in the ratemeter and would fail the selection criteria used in Fig. 3.29. The longest static hold on this day was terminated by a single large-timescale light event. This event is shown in 3.30.

It was discovered during these tests that the two ceramic alumina feedthroughs are partially translucent and leak light into the system. A cell phone flashlight, when shined at the feedthroughs, would produce an increase in the ratemeter by a factor of a few. Unused feedthroughs were obtained, and it was confirmed that they allow through visible light. Electrical tape was placed over the installed feedthroughs, as replacing them at this point would have spoiled the purity of the argon. Despite this intervention, the ratemeter still read 1250 Hz. The room and utility lights were then shut off, and the rate remained at ~ 1250 Hz. We concluded that the light leak was sub-dominant to whatever source was producing this light, presumably the feedthrough.

Occasionally two ramps were performed during a single day. Background data at 0V with 10Hz trigger rate were taken overnight when possible. In these ramps, the previous sustained increase in SPE rate was observed, as well as the sporadic *dips* in rate. No transient upward spikes in SPE, no non-ohmic behaviour, no cone events, and no large light events were observed during the second two days of data taking. Though, the rate during the ramps was persistently higher than the OV data. No apparent drop in LAr purity was observed in the purity monitor cathode signal. The grading structure was ramped up and no bursts in light were observed.

Following the conclusion of this run the 5W heater was turned on for two weeks, followed by a larger 80W (32 V) heater. Afterwards the scale read 427.5 kg, implying ~ 95 kg of

argon remaining. The jacket space was confirmed empty by blowing gas into the liquid fill tube. Cold gas was felt leaving, but no sound of liquid. The large and small valves were opened to atmosphere. The small supply was set to 10W and left until the argon was boiled off.

Run 2

This test examined the grading structure equipped with a mock spring connection. The mock spring was not an actual spring but rather electropolished steel which was inserted into the HV cap. As the mock spring was proud of the cone, some protection for the PMT photocathode was necessary in the event of a breakdown. An additional grounded stainless steel cone extended from the PMT mount ring. The mock spring had ridges which approximated the pitch of the extended spring connection. The dewar was completely evacuated with a turbopump out between runs, and the resistance was checked. Between Run 1 and Run 2 the PLC software was repaired, so the ratemeter and HV power supply were logged to disk, and the voltage from the HVPS was controlled in software. The scaling factor between the voltage output from the control panel and the voltage output by the power supply was later confirmed. The purity monitor anode and cathode were confirmed to be connected to their respective power supplies. At the beginning of each day the test stand was topped off with LAr. The jacket was refilled before the start of testing to a scale weight of 530 kg. The level meter capacitor was read with a Craftsman multimeter showed 1010 pF. Zero-field data was taken with PMT photocathode voltage = -1750 V while the LAr was being refilled.

Some troubleshooting of the purity monitor was done during this time with the goal of observing an anode signal. The gender changer in the purity monitor feedthrough was exchanged in order to improve the light transmission. The optical paths going into the dewar were taped up, which resulted in a loss our temp monitoring.

Zero field overnight data was taken throughout this run. The ratemeter was logged and was mostly quiet, with rate 300 to 500 Hz. Small bursts could be seen, as in Fig 3.33. The temperature in the resistive temperature devices (RTDs) increased overnight, as well. The PMT power supply seemed to fluctuated more than before. It was decided to temporarily shut off the PMT. A working hypothesis that we attributed this phenomenon to was a faulty PMT power feedthrough. The temperature read by certain RTD's were observed to gradually approach room temperature, indicating the existence of a gas layer.

For Run 2 a filter was used to clean the argon as it filled the dewar. This filter used activated copper to getter the impurities out the the argon, and relied on conditioning before the fill to use effectively. Conditioning consisted of passing hydrogen gas through the filter material at a temperatures exceeding 200C. The argon fill relied on a gas regulator between the storage dewar and the filter.

An upper limit on the background rate was calculated for the ramps, indicated by the dotted black dotted line in Fig. 3.32. Sporadic deviations can also be seen in this limit, but these seem to be statistical fluctuations resulting from inconsistent livetimes in the bins. The changes in livetime can be seen in the dotted background rate limit. Upward deviations in that limit indicate less livetime, and these fluctuations correlate with the downward fluc-

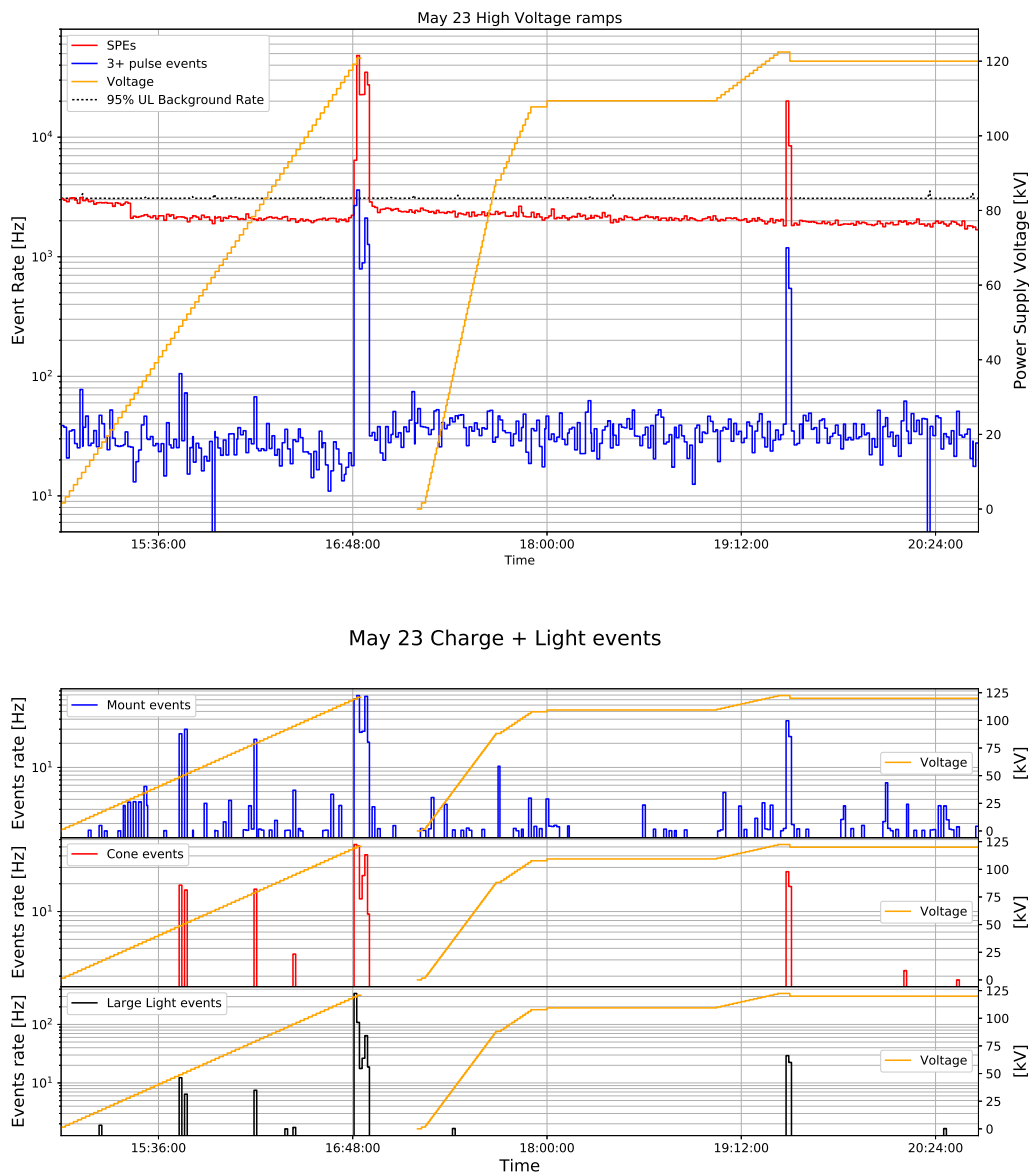


Figure 3.29: PMT (top) and charge amplifier(bottom) data from the first day of Run 1 high voltage ramps. Spikes in rate were seen in two locations, but a quiet region was observed following.




Figure 3.30: The waveforms from terminal event in the first day of Run 1's high voltage ramps. This is the activity which defined the length of the quiet region for later analysis.




Figure 3.31: PMT time series data from the May 25 high voltage ramp. This ramp was far less eventful than the first day, with the rate maintaining a consistent value overall.

tuations in SPE rate. These fluctuations also occur at regular intervals, indicating that the DAQ may have been blind during some process e.g. saving a file.

This run's data exhibits several interesting features. There were sharp upward spikes in rate which decayed away over an approximately hour-long timescale, seen in Fig. 3.32. This is in sharp contrast to the spikes observed in the first run, which had spikes which sharply turned on and off over a timescale of several minutes. These spikes correlated with charge events and large light events. Interestingly enough, even though the SPE rate remained high and slowly decayed, the charge+light "sparks" only correlate with the initial rise, not with the tail. Uncorrelated large light events occurred during the entire ramp, in stark contrast to the entirety of Run 1. Large light events in excess of the uncorrelated rate were observed at the initial onset of the test. These events do not show up in the SPE rate, though.

These spikes appeared at lower voltages than the apparently safe ranges indicated from Run 1. It was feared that this indicated that the mock spring was to blame. However, the low rates of pulses into the charge amplifier, along with general instabilities in the PMT voltage, indicated that the PMT feedthrough might be to blame. The fact that the rates of SPEs were maintained for some time, while the larger area pulses rapidly fell back to background levels was an indication that the noise came from a persistent, low-intensity source. Early into the run it was discovered that the LAr was not filled to the level expected. The regulator it seems did not maintain the desired pressure, and therefore the LAr had to be topped off before starting the tests. An additional gas layer likely weakened the HV feedthrough used for the PMT. It was discovered after the first day of testing that lowering the photocathode voltage from -1750 V to -1550 V extinguished the long-timescale light production.

Ramps following this change to PMT gain still failed to achieve the voltages of Run 1, but were far more stable, as seen in Fig. 3.34. For the purposes of maintaining the integrity of the test, if a burst of SPE was seen in the offline analysis, the "clock" was restarted, and the time period following is considered a second static hold. An example of a terminal event is shown in Fig. 3.36. Five "holds" were recorded over a period of ten calendar days. In general voltages exceeding -70 kV were attainable, which when correcting for purity and surface area translates to a success for the LZ design voltages.

The Gaseous Argon Test

The liquid argon surrounding the test structure in the second liquid argon test was slowly boiled over several weeks and the structure reached room temperature. High voltage was applied to the structure with the Glassman power supply in the same manner used in the liquid argon tests. An attempt to view light produced by the structure with the PMT failed because the PMT base was unable to sustain the voltage needed to operate the PMT without breakdown of the surrounding argon gas. The current through the Glassman supply was monitored to determine the onset of excess conduction.

Three voltage ramps were made, each at about 500 V/min. In each case the current through the Glassman supply increased Ohmically (*i.e.* $V \propto I$) until a threshold voltage was reached. At the threshold voltage the Glassman power supply current suddenly increased by




Figure 3.32: PMT data from the August 14 high voltage ramp.




Figure 3.33: PLC and PMT data from the August 14-15 overnight zero field data.




Figure 3.34: PMT time series data from the August 24 (Run 2) high voltage ramp. Here the ratemeter data is indicated in green.




Figure 3.35: Charge amplifier data from the August 24 high voltage ramp. The initial spike in rate cause the ramp to pause for the remainder of the test that day.




Figure 3.36: Waveforms from the terminal event in the August 24 high voltage ramp, defining the quiet time period for Run 2.

about a factor of ten above the Ohmic current value. The observed thresholds were 18.1 kV, 17.6 kV, and 18.0 kV.

The reproducibility of the onset voltage suggests that this test does not damage the feedthrough and that it is suitable as a quality control check for the production feedthrough that will be installed in LZ. The magnitude of the sudden increase in current suggests that it is not due to conduction between rings of the feedthrough (as would be the case if the breakdown were bypassing a high voltage resistor). The large magnitude of the current can only be generated by conduction to ground from one of the structures connected directly to the high voltage cable: the last grading ring, the hemispherical end, the spring mimic, or the hanging bean.

Summary

For each run, any period of time longer than twenty minutes where no light or charge events was logged, along with the power supply voltage during that time. The goal voltage in any case was -100 kV LZ equivalent, after corrections. Several effects were taken into consideration to translate the LZ goal voltage to an equivalent test voltage for the setup used in these tests. One such effect is that of purity. It has been observed that electronegative impurities enhance the electric field noble liquids can hold before breakdown [119]. Our setup is less pure than LZ’s goal purity, so the goal of -100 kV must be adjusted upwards. Another effect is that of “stressed electrode area,” (SEA) or the surface area of an electrode exposed to the strongest electric fields (often 90% of the maximum). As observed in Ref. [104], the breakdown voltage exhibits a scaling given by

Run	Highest Voltage [kV]	Target Voltage [kV]	Duration	Purity [ppb] (Triplet)	Stressed Area [cm ²]	PMT Voltage [kV]
5/23	109	105	1h3m	155 ±11	68	-1.75
5/23	120	105	42m	155 ±11	68	-1.75
5/24	120	105	46m	155 ±11	68	-1.75
5/25	120	105	1h7m	155 ±11	68	-1.75
8/14	58.6	70.1	39m	205 ±6	387	-1.75
8/15	63	71.8	50m32s	418 ±15	387	-1.55
8/16	72.5	72.6	28m	587 ±11	387	-1.55
8/16	72.5	72.6	33m	587 ±11	387	-1.55
8/17	77.8	73.4	1h28m	805 ±1	387	-1.55
8/24	77.8	72.06	1h4m	460 ±13	387	-1.55

Table 3.1: Test results for the CHV feedthrough Runs 1 and 2. The “highest voltage” column is the actual recorded voltage from the HVPS, while “target voltage” indicates the voltage necessary to achieve -100 kV on the LZ cathode, when correcting for surface area and purity.

$$E_B = C(\text{Area}[1 \text{ cm}^2])^{-b}, \quad (3.6)$$

where $C = 124.26 \pm 0.09 \text{ kV/cm}$, $b = 0.2214 \pm 0.0002$. In Run 2, a mockup of the spring connection was inserted that mimicked the maximum field, but severely enhanced the stressed area. Because of this, the -100 kV target is revised downwards to an equivalent target voltage, as seen in Table 3.1. The purity is a far weaker effect than the stressed area.

These tests suggest that the margin for functionality in LZ is wide, but two things can eat into the margin: First, there are differences in dielectric strength between liquid argon and liquid xenon. However, these appear to be minor, at least in terms of complete dielectric breakdown (spark discharges) [104]. Electroluminescence data were collected in liquid xenon (Chapter 8), but no direct comparison in the same configuration exists for liquid argon. Second, the ‘passing’ criteria was for the feedthrough to hold a fixed voltage of 100 kV[57] without light production or excess current conduction for one hour. For LZ to function, cathode needs to operate in a stable manner for years. The test data suggest that the feedthrough is not damaged by the events that cause the onset of light and excess conduction. The occasional (at timescales > 1 hour) occurrence of these events in LZ would probably not cause irreversible damage and could probably be eliminated by a small (a few percent) retreat in the applied cathode voltage.

The two above considerations suggest that the feedthrough may not be able to operate long-term in LZ at the highest plateau voltages reached in the LAr tests, but almost certainly can operate above the requirement voltage of 50 kV.

3.5 Deployment

Clean room build

Production Grading Structure

Following the success of the prototype design, work began on the production LZ grading structure. The prototype was not deployed itself due to the lax cleanliness standards used in the testing at LBNL. While there was a class 10,000 tent with a HEPA filter, this was insufficient to be allowed within the LZ xenon volume. Additionally, there was no way to guarantee that damage had not occurred during the stress testing of the feedthrough.

The construction of the production feedthrough proceeded much in the same way as the prototype feedthrough. One difference was that the cryofitting took place within a class 1000 cleanroom, rather than in the tent in building 77A. This required the transportation and installation of plumbing which transports liquid nitrogen from a 400 kg dewar in the antechamber into the clean room. The modified cryofitting procedure took place by dipping the cable in a small, 5L dewar filled with LN₂, before inserting it into the polypropylene xenon displacer.

The cryofitting procedure failed several times before the eventual successful fit. Initially it was believed that the displacer was not bored to a large enough diameter. Reaming plastic is challenging due to it deforming instead of chipping away, and its tendency to expand and deform due to the heat from the lathe. After the first failure to insert the cable, the displacer was removed and re-reamed with the closest size up. This again failed, with the cable becoming stuck part way.

Several concerns made the rapid resolution of this problem pressing. Firstly, this work was time sensitive for the LZ construction timeline. Secondly, if the cable became stuck partway, a new displacer would have to be used, of which there were a finite number of spares. This was particularly concerning due to the possibility for the crystalline cable to shatter inside the displacer. Lastly, each time the cable failed, an additional section of cable would have to be cut. At some point the slack built into the cable length would be exhausted, and the cable would not be able to stretch from the o-ring seals to the high voltage power supply.

Therefore, in the final, successful attempt it was decided to cool the cable for a longer amount of time. The cable was instrumented with a resistive temperature device (RTD) in order to measure the temperature in real time. Over the course of half an hour, the cable temperature was observed to drop. When the temperature reached a minimum and began to rise again, the cable was quickly and successfully inserted into the displacer.

The stress cone was then attached to the displacer, and the setup was affixed to the steel mount. Throughout this process, the displacer was wrapped with clean wrap at the end of the day. The final construction of the grading structure would take place beneath a HEPA filter within a class 100 clean room, shown in Fig. 3.37

The FEP radon guard was affixed on the underside of the steel mount. This component redirected all of the emanation from the plastic to the gas phase of the xenon, where it was




Figure 3.37: *Left*: The appearance of the xenon displacer before the grading structure was constructed. The plastic wrap was to prevent any additional dust from collecting on the surfaces. While the structure was regularly wiped down with shed-free wipes saturated with ethanol, there was still the possibility for trapping dust during construction, thus the additional precautions. *Right*: The FEP radon guard tube connection.

redirected into the inline radon reduction system. Each grading ring was then placed over the grading structure with delrin spacers establishing the gradient. Inside each grading ring is a small recess where the resistors reside, so as to not be too proud and create excess fields. The resistors hook into small plastic guards which prevents shorts to the conductive grading rings, which are nicknamed “boomerangs” for their obtuse angle. These boomerangs, along with the first two grading rings, are shown in Fig. 3.38

After all nine rings were placed and connected, the continuity of the chain was tested with a multimeter. Since the resistance of the grading structure is orders of magnitude larger than the internal resistance of the multimeter, it was necessary to test the resistance with the meter in voltage mode, rather than resistance mode. An external voltage supply provided a small voltage to one end of the grading structure, with the multimeter in series with the structure. The resistance of the chain is then given by $R_1/R_2 = V_2/V_1$, where R_i is




Figure 3.38: Photographs from the production grading structure construction process in the LBL clean room. *Left*: The placement of the first grading ring along with the resistor “boomerang.” *Right*: The placement of the first two grading rings, along with the attached resistors.

the resistance in that particular section and V_i is the voltage drop across the same section. The proper resistance was confirmed, allowing the HV cap to be placed. This HV cap will attach to the spring connection to the cathode, and makes a tight connection via a detent and wave spring. The completed structure is shown in Fig. 3.39.

Cable Routing

Following the construction of the production grading structure, the cable plumbing was built. This build took place in the class 1000 clean room due to space constraints. The main challenges for this stage were the careful planning of the topology and order of operations in order to not inadvertently trap components, and the elimination of torsion which can concentrate over the several meters on each component. It was found that spooling the cable prior to threading kept kinks from forming, which would have been challenging to work out.



Figure 3.39: A photograph of the completed production grading structure in the LBL clean room. Left: Ethan Bernard. Right: the author, James Reed Watson.

Two particularly challenging spots were the link between the inner and outer cones, which consisted of a small bellows, and the triple stack, where the VJ hose connects to the spool can. The first was a struggle due the need to carefully handle the inner cone to avoid shearing the cable. Additionally, the bellows required testing to confirm that it would not fail when placed under large positive pressure, since it had a head of several meters of LXe on one side, and vacuum on the other. The triple stack was challenging due to the weight on either side of the seals, along with the intrinsic difficulty on forming two CF seals on either side of a piece. The leverage caused the conflat to not be tightened uniformly, leading to failure to seal adequately. These failures were discovered and the gaskets replaced. Successful sealing took place after proper support was added to both sides of the seal, along with proper utilization of a tension wrench.

After the system was sealed, a Helium leak check was performed. Each CF seal was checked, along with the spool can's o-rings. Each individual seal, along with the system as a whole, met its leak specification.



Figure 3.40: A photograph of the cathode high voltage delivery crate. Aluminum panels lined the box. The plastic wrap kept the bellows relatively dust free, despite the box not being airtight. Assorted tools necessary for installation were shipped in the same box. The HV cable is seen spooled on the top, protected by a separate plastic sheath. A challenge of this method was orienting the bellows correctly to minimize strain and torsion.

Delivery and Installation

Following the successful leak test, the entire system was packed up and shipped to the Sanford Underground Research Facility (SURF). The entire structure was wound up in a double-decker steel box, along with all tools necessary for testing and installation. This is shown in Fig. 3.40. During this process the xenon volume was completely sealed, so the only areas that were exposed were regions that would be exposed to the air post installation. Additionally, the bellows was covered in plastic to prevent dust accumulation. Soft supports were placed beneath the grading structure to damp vibrations.

Following delivery, the package was unpacked and leak checked once more. The cone was lowered through a hole in the LZ water tank and placed onto a custom cart. Following extensive training and preparation, the grading structure was connected to the cathode. This commenced while the ICV was under a nitrogen purge. Due to the risk of asphyxiation in the confined space of the (then empty) water tank, the workers (including the author) wore self-contained breathing apparatus (SCBA), which provided oxygen. The final connection is shown in Fig. 3.41.

Following the final sealing, multi-layer insulation (MLI) was wrapped over the inner cone. Copper thermal links were installed prior to this, along with RTDs to monitor the cooling of the cone. The seals were leak checked one final time, and the effective resistance was confirmed to be consistent with the forward field region, reverse field region, and grading




Figure 3.41: Photographs taken during the cathode high voltage feedthrough final installation at SURF. *Top*: The preparations taking place within the water tank. While the TPC is visible at this point it was covered in a plexiglass cover. *Bottom*: Shortly before the final connection between the feedthrough and the LZ cathode. Pictured is technician Derek Lucero in SCBA.

structure resistances in parallel.

3.6 Conclusion

The LZ cathode high voltage connection was successfully built and tested at LBNL. A prototype was tested in liquid argon, with charge and light sensors to examine changes in the rate over time. At 100 kV equivalent voltage, no elevated rate of single photons was seen for over an hour. A production design was constructed in a clean room, passed all checks, and transported to SURF. It was installed in LZ under a nitrogen purge and commissioned for SR1.

For SR1 the goal was to reach a voltage of -50 kV on the cathode. A slightly elevated rate of light was observed in the skin PMTs facing the cathode connection. As the design is functionally identical, it is possible this has to do with the link to the cathode ring itself. In Chapter 7, a scan over cathode voltages is analyzed, demonstrating that the cathode does not significantly affect TPC photon rates. Rather than troubleshoot the system while it reached stability, an acceptable voltage of -32 kV was chosen for SR1. The ER/NR discrimination the CHV feedthrough enables is a major contributor to the success of the first science result out of LZ.

Chapter 4

Electric Field Simulations for LZ

4.1 Introduction

A TPC's electric field affects every aspect of the experiment. The drift field magnitude determines the drift velocity, changing the mapping of time to z-position in the detector. It also impacts the discrimination power between electron recoils and nuclear recoils via its affect on the recombination probability. Nonuniformities in the field distort the field lines and electron trajectories. While TPCs often aim for electric fields of the highest achievable uniformity and magnitude, going too high in field can have undesired results.

A particularly instructive example of a nonideal electric field is LUX Run 4. Following the completion of Run 3, a grid conditioning campaign was conducted, with the intention of burning off impurities on the electrodes and increasing the allowable applied voltages. This unfortunately had the opposite effect, distorting the electric field and reducing the average electric field magnitude[120]. This is believed to be due to the creation of a permanent negative charge on the PTFE walls of the TPC of between $-3.6 \mu\text{C}/\text{m}^2$ to $-5.5 \mu\text{C}/\text{m}^2$. Due to this nonuniformity, a complex field model had to be fit to the data, in order to properly simulate the detector response.

While it was not expected that the field nonuniformity in LZ would be as severe as in LUX Run4, it was still necessary to model the electron drifts and extraction fields. The electric fields had to be solved for a variety of possible electrode configurations, as it was not known ahead of time which voltages would be applied.

4.2 Finite Element Simulation

Maxwell's Equations

The software used to simulate LZ's electric field (FEniCS[121]) requires a more in depth knowledge of the details of finite element simulation than other packages, for example Comsol. At its core, the finite element method seeks a field u which is the solution to some differential equation, subject to boundary conditions. The field is represented as an element

of a particular vector space V . Each node of the mesh is assigned a value $u_i = u(x_i)$, and the spatial derivatives are calculated based on the nearby nodes $x_{j \neq i}$. The differential equations therefore appear as constraints specifying the relationship between nearby nodes, i.e. $\sum_k a_k u_k = 0$. The entire problem then becomes system of equations, $\mathbf{A}\mathbf{u} = \mathbf{f}$, where \mathbf{A} is the finite element matrix and \mathbf{f} is the source term. For electrostatics, this source term is physically the free charges in the solution. Constraints, such as Neumann (derivative) or Dirichlet (value), can be applied to the boundaries in the form of Lagrange multipliers.

There are many solvers capable of finding u , but it is observed to be beneficial to put the differential equations in “weak formulation” first[122]. This is a conjugate problem whereby the following transformation is performed:

$$\mathbf{A}\mathbf{u} = \mathbf{f} \rightarrow v^T \mathbf{A}\mathbf{u} = v^T \mathbf{f} , \quad (4.1)$$

where $u, v \in V$, a vector space describing the problem. This solution must hold for all v . The solutions found with the weak formulation are equivalent to the solutions found with the strong (original) formulation. The advantage of performing this technique is that the constraints are all first order constraints, rather than second order, which is more numerically stable. This is somewhat reminiscent of the equivalence between Lagrangian and Hamiltonian mechanics. In the case of electrostatics we are solving the Poisson equation where u represents the electric potential. Therefore, the weak formulation, with units of permittivity $\epsilon = 1$ becomes:

$$\nabla(\mathbf{D}) = \nabla(\epsilon\mathbf{E}) = -\nabla(\epsilon\nabla\phi) = \rho_f \quad (4.2)$$

$$-\int v \nabla(\epsilon\nabla\phi) d^3x = \int \rho_f v d^3x$$

$$\int (\nabla v \cdot \nabla\phi) d^3x - \int \epsilon \frac{\partial\phi}{\partial n} d^2x = \int \rho_f v d^3x \quad (4.3)$$

$$\int (\nabla v \cdot \nabla\phi) d^3x + \int \sigma_f d^2x - \int \rho_f v d^3x = 0 . \quad (4.4)$$

Without free charge accumulation, it is unnecessary to insert the dimensionful value of ϵ_0 . For the LZ simulation I exploit the approximate axisymmetric symmetry of LZ, which requires a change to cylindrical coordinates, given by:

$$\int (\partial_r u \partial_r v + \partial_z u \partial_z v) r dr dz + \int \sigma_f (r dz dr) - \int \rho_f v r dr dz = 0 . \quad (4.5)$$

LZ Axisymmetric model

To simulate the electric fields, a simplified, axially symmetric model of the LZ TPC was constructed. Primarily this was used to evaluate the field nonuniformities, both in magnitude and in trajectory. Fine details that would not affect the overall solution of the TPC were excluded, such as wires and assorted connectors. Insulating spacers and PMTs were included,




Figure 4.1: The LZ 2D axisymmetric finite element mesh, produced and rendered in GMSH. *Left*: The complete mesh. Colors indicate distinct volumes, but are not necessarily the same material. Note the PMT banks on the top and bottom of the detector, and the relative sizes of the facets. *Right*: A zoom-in of the field shaping rings. The mesh has a far higher resolution near the resistors because the field changes drastically in that area. The region around the rings is the PTFE wall. The fine details such as the small gaps between the pieces are not modelled. The dielectrics are mismatched by $\approx 0.25 \epsilon_0$, but the areas are relatively small, so are not expected to affect the result considerably.

but were modelled as an annulus with the same side-view cross section. The electrode grids, while in reality are woven meshes, were modelled as rings with the same diameter as the physical wires. In order to match the characteristics in the bulk, I followed the recommendation from Ref. [123], and halved the pitch of the simulated grid. This was shown in simulations to perform similarly to the full 3D simulation when at least 1 pitch distance away from the grids.

The sagitta (the greatest distance between an arc of a circle and a chord) of the grids was found via an iterative procedure done by described in Ref. [124]. There, the grids were deformed in a hyperbolic cosine profile and the fields resimulated until the electrostatic forces balanced the tension and weight of the grids. The sagitta were of order 1 mm and these values were compared against extraction region data taken during commissioning. The cathode was similarly simulated, however its sagitta has relatively diminished impact on the overall result.

Under the force of gravity, a wire of density ρ , cross-sectional area σ , and length L under a tension T , will curve to a sagitta determined by the following equation [125]:

$$s = \frac{\rho g \sigma L^2}{8T} . \quad (4.6)$$

In the presence of an external field E_0 , each element will experience a force $dF = \lambda E_0 dx$, which effectively subtracts from the differential weight $dF = \rho \sigma g dx$. Additionally, the sagitta will perturb the electric field to first order in s/d , where d is the distance between the grids. This results in an enhancement of the sagitta relative to the nominal ($d = \infty$) case. Since both the external field E_0 and the charge density λ depend linearly on the applied voltage V , then the overall sagitta scales as $s \propto V^2$.

The case of a wire grid with pitch a , radius r_0 , and wire-electrode distances b_1, b_2 , and $b = b_1 + b_2$ is worked out in [126], where the charge density is given by

$$\lambda = \frac{abV_0}{4\pi b_1 b_2 (1 + K)} \quad (4.7)$$

$$K \equiv \frac{ab}{2\pi b_1 b_2} \ln \frac{a}{2\pi r_0} , \quad (4.8)$$

with a bulk fields on either side of $E_1 = V_0/b_1$ and $E_2 = -V_0/b_2$, evidently resulting in the V_0^2 scaling.

Particular attention was paid to keeping the number of elements in the simulation large enough to ensure accuracy, and small enough to fit into memory. While mesh refinement methods were explored for the purposes of automating this process, it was decided that the resolution would be set manually. This decision was guided by the fact that process of adding the basis fields for each electrode together in superposition is faster when all fields are calculated on the same mesh. Resolutions of 1 mm were set over the PTFE walls, and the grids were forced to have the resolution necessary to model their circumference. This small facet size was unnecessary in the middle of the TPC, so certain volumes were cut out and a larger resolution (generally cm-scale) was applied to their boundaries. In order to make

Parameter	Value
Liquid Xenon dielectric	1.874 [127]
Gas Xenon dielectric	1
PTFE dielectric	2.1 [128]
PEEK relative permittivity	3.2 [129]
Cathode-Gate distance	1456 mm
Gate-Anode distance	13 mm
Gate-Liquid surface distance	5 mm
Cathode-Bottom grid distance	137.5 mm
PTFE radius	728 mm
PMT photocathode bias	-1250 kV
Bottom grid bias	-1250 kV

Table 4.1: Parameters used in the LZ electrostatics simulation. The electrode biases are a parameter which was explored as part of this work.

the increased resolution necessary to model the grid wires not “spread” to other regions, the grids were placed in holder volumes. These generally specified the locations where the resolution needed to drop to 1mm in order to keep the simulation tractable. These volumes had to have additional height to manage the sagitta of the grids. These holder volumes can be seen as a change in color in the mesh zoom-ins in Fig. 4.2. The mesh was generated in GMSH using the *Frontal-Delaunay* algorithm.

Resistor chain optimization

A preliminary model developed in Quickfield¹ was initially examined, which indicated a large amount of uniformity. It was desired to explore possible interventions to make the field more uniform. These interventions were restricted to decisions that could be made at that point in the assembly process. Since LZ could not be redesigned from the ground up, this meant that the options were limited to swapping out the resistor values in the field cage. The field cage consisted of titanium field shaping rings connected to one another with two 1 G Ω resistors in parallel. The assayed resistors included values of 500 M Ω , 1 G Ω , and 2 G Ω , allowing 7 possible values per gap. With 57 field shaping rings plus the gate, this makes 58 resistor gaps to search. While the electron trajectories are nonlocal, the overall field is only sensitive to the difference in potential between the middle of the TPC and the edge/wall. This means that the largest effects come from the fields at the edge of the field shaping ring, allowing an approximate ordering of the configurations.

The starting field exhibited a large amount of electron deflection towards the center, which I intended to mitigate. This is caused in part by the size and shape of the cathode ring. The ring extends past the cathode grid wires in both the vertical and horizontal

¹quickfield.com




Figure 4.2: The zoomed-in 2D mesh for the LZ cathode, gate, and anode rings. *Left*: The cathode at $z = 0$ in LZ simulation coordinates. The shape juts out in all directions, causing a distortion of the electric field in the vicinity. Note the holder volume for the cathode mesh. Also visible is the first reverse field region (RFR) field shaping ring. *Right*: The extraction region, consisting of the gate, anode, and PEEK spacer which separates the TPC and Skin volumes. Particular features include the fact that the final “ring” is actually an extension of the gate, and the relative proximity of the gate electrode to the liquid surface.

directions. Similarly, the gate ring also extends below the gate electrode. These geometric changes cause the distance over which the potential must vary to decrease relative to that of the center of the detector, increasing the field in that location. The effect of the grid wires exacerbates this, as the thin wires enhance the local electric field relative to a flat plane. While the net effect could not be helped, an increase in the resistor values in the final gap, below the gate, would lead to an increased ΔV between those surfaces, distorting the field lines into a more vertical direction at a critical juncture.

The optimization procedure was conducted by utilizing Quickfield’s python bindings. This allowed simulations to be controlled in software, rather than a tedious manual search. At each point in the solution space of resistors, the fields were calculated. The metric for success was the volume-weighted field uniformity as calculated in the fiducial volume. The solution was reweighted by the radius r due to the cylindrical coordinates. Then, proposed solutions were explored in the vicinity of the current configuration. This was limited to changing the resistors one value up or down (among the seven options), and was limited to the first and last two resistor values in the chain. Then, the “direction” with the largest decrease in fiducial volume E-field variance was chosen, and the procedure repeated. The previous solutions were cached to save time on recalculation.

The aforementioned procedure converged on a tentative solution, but it was later dis-




Figure 4.3: The resulting simulated electric field for the LZ forward field region from Fenics. This simulation uses the SR1 electrode configuration, and the equipotentials are steps of 1 kV. The field is highly uniform, with most of the nonuniformities located along the edges of the TPC.

covered that the model itself had been causing some of the observed field nonuniformity. The preliminary mesh modelled the wires with a diameter far exceeding that of the physical wires. Additionally, there were fewer simulated wires than in reality. This was done to avoid the computationally expensive meshing procedure which would result from accurate wire placement. However, setting the appropriate diameter and pitch is crucial for accurate simulations.

The wires were simulated with the same diameter, half-pitch configuration and the procedure repeated. The nonuniformity in the starting configuration was found to be acceptable, but there was room for improvement. At the goal cathode voltage of -100 kV, with 11.5 kV ΔV in the extraction region, the optimal solution ended up being simple: replace one of the final ≈ 1 G Ω resistors with a 2 G Ω resistor, effectively increasing the resistance in that stage from 500 M Ω to 667 M Ω . This improved the field uniformity to the point where the fiducial




Figure 4.4: Radial electric fields and equipotentials near the cathode region. *Left*: Before the modification to the resistor chain. *Right*: After the modification to the resistor chain. The change to the electric field near the cathode results from the modifications to the resistor chain grading. Each equipotential is a step of 100 V from its neighbors. The nonlinear grading results in a larger initial voltage step between the cathode and FFR1. This has the effect of pulling the field lines straighter, reducing the magnitude of E_r , the radial component of the electric field.

volume had less than 1 V/cm volume-weighted standard deviation.

The impact of the change to the field uniformity is shown in Fig. 4.5. The volume-weighted average is computed over $z \in [20, 1336]$ mm and $r < 688$ mm. The change to the resistor grading causes insignificant impact on the mean field, and reduces the standard deviation from 3.2 V/cm to 1.0 V/cm at the goal cathode voltage of -50 kV and gate voltage of -5.75 kV. In order to investigate the field uniformity outside of the strict WIMP search fiducial volume, an “engineering fiducial” volume was also used. This was defined as being 2 cm from the grids and PTFE wall. This region unsurprisingly has a slightly higher nonuniformity at 1.6 V/cm when using the nonlinear resistor grading.

Anode Correction

The primary focus of this work was finding the map of real space to reconstructed S2 space $f : (r, z) \rightarrow (S2R, DT)$ within the TPC, along with the electric field magnitude $|E|(r, z)$. While the anode voltage is an important component of the overall physics, it is not crucial




Figure 4.5: Volume-weighted histograms of simulated field magnitudes in the LZ TPC. *Left:* The impact of the resistor grading change on the fiducial volume electric field uniformity. The curves indicate the volume-weighted field magnitudes within the fiducial volume, *i.e.* a grid of points is used to sample the electric field, and the average is calculated by weighting each point by r . The blue (orange) curve indicates the unoptimized (optimized) resistor gradings. *Right:* The impact of the choice of “engineering” vs “analysis” volumes on the uniformity. The engineering fiducial is a larger volume meant to encapsulate the full nonuniformity without being sensitive to the fine details of the grids.

for these quantities. Because the anode grid would increase the facet count by nearly a half (since the facet count is dominated by the cathode and gate grids), it was simplified to being a plane. This plane is not parallel, but still maintains the predicted sagitta.

However, the electric field near a wire is not identical to a plane even in the idealized situations. The potential near an infinitely long cylinder of charge is $\phi(r) = \lambda \ln(r/a)$, and therefore on the scale of a pitch distance the infinite plane approximation must break down. Modelling the grid wires as a collection of unit cells, with a wire placed between two Dirichlet boundaries in one direction, and two Neumann conditions in the other, a correction can be calculated. The field far (multiples of the pitch) from the wires with voltage V_0 , with electrode voltages on either side V_1, V_2 , at distances b_1, b_2 , is given by [126]:

$$E_1 \approx \frac{V_1 - V_0}{b_1} + \frac{b_2}{b} \left[\frac{V_0 - V_2}{b_2} - \frac{V_1 - V_0}{b_1} \right] \frac{K}{1 + K}. \quad (4.9)$$

In order to correct the anode grid, V_1 is taken as the gate voltage, while V_2 is taken to be the PMT voltages. The factor K is given by Eq. 4.8. This is additionally complicated by the change in dielectric between the anode and gate due to the liquid-gas phase change at $z = 1461$ mm. For the purposes of this calculation both grids are temporarily modelled as having zero sagitta, *i.e.* perfectly level. Since the goal is to find the voltage to apply to the planar anode such that the electric field in the extraction region is matched, the sensitivity of the gas field to the anode voltage was calculated:

$$\frac{dV_0}{dE_1} = s_0 + s_1 \frac{\epsilon_0}{\epsilon_1}, \quad (4.10)$$

where s_0 is the liquid surface-anode distance and x_1 is the liquid surface-gate distance, with ϵ_0 and ϵ_1 being the gas and liquid dielectric constants, respectively. From these calculations, a correction factor of $\Delta V = -430$ V on the anode was found to match the gas extraction region field between the grid and plate conditions.

Cathode Voltage Simulation

It was not known at simulation time which voltages the detector would run at, particularly that of the cathode grid. Therefore, it was necessary to understand the results at a range of cathode voltages in order to make informed decisions during commissioning. Therefore, the analysis done for the resistor chain optimization was repeated at a number of cathode and voltages, and the engineering and fiducial volume averages were computed. The cathode scan was performed at a fixed gate voltage of -4 kV, while the gate scan was performed at a fixed cathode voltage of -32 kV. An ideal cathode voltage from the perspective of field uniformity was found at -45 kV. It was observed that the gate voltage has a disproportionate effect on the field nonuniformity, *i.e.* changing the gate voltage by a kilovolt had a larger impact on the nonuniformity in the fiducial volume than changing the cathode by an identical amount. This is caused by the field leakage from the extraction region, as well as the interaction between the gate grid and the upper corner of the detector. These effects are illustrated in Fig. 4.6. The ideal cathode voltage, conditional on a particular gate voltage, increased by approximately 10 kV for every 1 kV on the gate (with the same polarity).

4.3 LZElectricField simulations package

Drift simulation

While particular tools such as *Garfield* exist for drift simulations, it was decided that the particular parameters of interest were best suited with a bespoke simulation. In addition to the electric field, the features calculated across the TPC are the:

1. Drift time: the time it takes for the electrons to drift from their point of creation to the liquid surface at $z = 1461$ mm.
2. S2 radius: the radial coordinate where the electron cloud, on average, crosses the liquid-gas boundary.
3. Probability of an electron attaching to the PTFE walls. This occurs when the field lines either intersect the PTFE walls at $R = 728$ mm, or pass nearby them such that electrons have a chance to diffuse past the boundary.




Figure 4.6: The field nonuniformity (RMS magnitude) as a function of cathode and gate voltage, averaged over the analysis and engineering fiducial volumes. Note the outsized impact the gate has relative to the cathode. *Top*: Scans over the negative cathode voltage. Blue: Sensitivity paper[93] fiducial volume. Orange: less restrictive fiducial volume, 2 cm away from the grids and wall. *Bottom*: Scans over gate voltage. Note that the engineering fiducial has a larger RMS value, and the optimal value is shifted closer to zero.

4. Transverse diffusion. This is parameterized by the diffusion constant D_T . The probability distribution function (pdf) of the electron's displacement in the direction transverse to the drift direction is a Gaussian distribution with standard deviation $\sigma_T = \sqrt{2D_T t}$, where t is the drift time[130]. EXO-200 measured this quantity to be $D_T \sim 55 \text{ cm}^2/\text{s}$ [131].
5. Longitudinal diffusion: similar to transverse diffusion but in the direction of drift. For the fields used in LZ the values for D_L are approximately $25 \text{ cm}^2/\text{s}$ [132]. This diffusion is still parameterized in terms of distance, so the effective diffusion in drift time is then

$$\sigma_t = \sqrt{\frac{2D_L t^3}{L^2}}, \quad (4.11)$$

where L here is the drift length in cm and t is the drift time.

The observable boundaries of the TPC are of particular interest, as the fiducial volume is defined relative to them. As a result of the low background design of the experiment, the only features with which to compare simulations and data are the edges themselves. These edges are the cathode and gate electrodes, along with the PTFE walls. Therefore, the initial conditions of the simulation must be chosen so that the boundary is sufficiently explored. The algorithm consists of the following steps:

1. Place the starting points for the trajectories. This is a nontrivial task, since it is desired for the resulting maps to be approximately uniform in sampling density. A subset of starting points are initialized along the top and wall, with points along the top displaced a short distance beneath the gate grid. The placement is uniform in R^2 , with the caveat that points are “snapped” to the pitch of the simulated grid wires. This is to guarantee to a high degree of confidence that the resulting field lines drift vertically through the gaps in the mesh, rather than be deflected radially. The density of points is specified near the wall to be 2.5 mm, and the snapped points are placed a distance of 3 mm below the gate grid. Along the wall, points are placed in a manner which is aware of the field shaping rings. In the middle of the height (Z) of each ring, points are placed at intervals of $\Delta R = 1 \text{ mm}$ inwards, without snapping to the grid wires. Along the gaps between each field shaping ring, points are placed a small distance ($\epsilon_R = 100 \mu\text{m}$) from the wall in vertical intervals of $\Delta Z = 500 \mu\text{m}$. The choices ensure that the regions where electron trajectories intersect the wall are mapped adequately, along with the boundaries between the intersecting and non-intersecting trajectories. Points are not placed along the bottom boundary since they are adequately explored via the other points. The cathode excursion into the TPC has additional points placed on its corners.
2. Map the points bidirectionally, *i.e.* find the trajectory which begins on boundary of the TPC and which passes through this particular point (r, z) . These trajectories integrate the electric field $\vec{E}(r, z)$ using RK4 (explained in the following section 4.3). When a boundary is crossed in a particular timestep, the crossing location is linearly interpolated. The drift time between each step is logged, and after the trajectory is

found the drift time at each step is integrated from the point where it crossed the liquid surface. Similarly the local contributions to the diffusion $\Delta_{\sigma^2} = 2D\Delta T$ are logged and then integrated backwards to the bottom(right) boundaries.

3. Downsample the points. The timestep is chosen to precisely map out the shape of the trajectories, but it is unnecessary in most parts of the detector to use every single timestep as the trajectories are fairly straight. As such, a “stride” value in either direction is specified, which enforces a fixed number of sample points within each cell. Sampled points near the wall relax this requirement, instead specifying a fixed distance from their starting points within which all points are kept. This is done in an attempt to not save an excessive number of points, which would be detrimental to the resulting interpolator performance.
4. Perform corrections. Due to the funneling effect of the grids, the trajectories have deflections inwards and outwards in R . This is an artifact of the simulation not expected to appear in the actual data, since the physical grids are woven mesh, and therefore the funneling effect is averaged over 2π . Therefore, the radial coordinate where the trajectory crosses a specified boundary slightly under the gate grid is logged, and this is then written as the $S2_R$ which is expected to be observed. This standoff distance which accomplishes this task is approximately one (simulation) pitch distance away, or 2.5 mm. The other correction which must be performed deals with trajectories which do not end at the liquid surface, i.e. on the PTFE wall. Because these points do not produce S2s, they do not have values to write out. However, leaving these areas blank would result in challenges for the interpolator, as it would then be attempting to interpolate between valid and invalid points. To deal with this, the valid points are grouped in tranches of z . The invalid points on wall-intersecting trajectories select the $S2$ positions from the valid point in that tranche with the largest R coordinate. The $S2_R$ value is then extrapolated based on the ratio between the valid and invalid physical R values.
5. Conduct Monte Carlo simulations. While points near the center of the detector may have their diffusion simulated analytically with little error, this is not true for points near the wall. This is because many such trajectories experience charge loss, which then skews the distribution of the surviving cloud. Because the trajectories are somewhat oscillatory, the transverse and longitudinal components mix somewhat, with the resulting σ differing from values at the same z but closer to the center. Because of this, points at most 2 cm from the PTFE wall are selected for tracking via Monte Carlo diffusion. A set of paths are simulated from the same starting point. At each time step, the transverse and longitudinal δx are drawn from their respective Gaussian distributions. Since the simulation is axisymmetric, at each timestep the points are transformed from cylindrical to Cartesian coordinates, the azimuthal diffusion simulated, and projected back into cylindrical coordinates. The resulting value is not critical for the resulting analysis but does serve as a control to measure the aforementioned transverse-longitudinal mixing against. Points which intersect the wall or otherwise

fail to reach the liquid surface are not counted for the purposes of calculating the resulting $\sigma_R(x)$ and $\sigma_T(x)$, but do allow for the calculation of a real-valued *attachment probability*.

6. Augment the boundary. In order to prevent errors resulting from points outside the boundary of the TPC, points were duplicated. Points within a small distance ϵ from the wall and inner boundary at $r = 0$ are copied and inserted a distance ϵ away from the wall. This allows points to be queried along the boundaries without running into errors at the wall, but is particularly helpful for points on the inner $r = 0$, $z = 0$ boundaries. Because the coordinate system breaks down, one can not sample points exactly at $r = 0$. Similarly, since the cathode wires are a finite radius, and the reverse field region leaks into the forward field region, electron trajectories often do not reach $z = 0$ when drifting backwards in time. In order to make the maps even more resilient against meshing issues, a special behavior was implemented in the LZElectricField software. When the nearest facet does not actually contain the query point, the interpolation method switches from barycentric interpolation (details below) to a simple nearest-neighbor method. In other words, instead of taking a linear combination of the facet vertices, the closest vertex to the point is selected as the returned value.

Runge-Kutta Method

In order to integrate the field lines without requiring too fine of timesteps Δt , a higher order method was used. Runge-Kutta methods are a family of integrators which solve the first order differential equation $\frac{dy}{dt} = f(y(t), t)$. In the case of integrating the field lines, the velocity field is static and therefore any time dependence can be removed. The fourth order method, known colloquially as “RK4,” has a local error of $\mathcal{O}(\delta t^5)$ and a global error of $\mathcal{O}(\delta t^4)$. This small error is necessary in order to accurately handle the probability of trajectories intersecting the wall without decreasing the time steps to a prohibitively expensive degree. The RK4 update step is given by the following equations:

$$\vec{x}_{i+1} = \vec{x}_i + \Delta t \left(\frac{1}{6} (k_1 + 2k_2 + 2k_3 + k_4) \right) \quad (4.12)$$

$$k_1 = \vec{v}(\vec{x}_i) \quad (4.13)$$

$$k_2 = \vec{v}\left(\vec{x}_i + \frac{\Delta t}{2} k_1\right) \quad (4.14)$$

$$k_3 = \vec{v}\left(\vec{x}_i + \frac{\Delta t}{2} k_2\right) \quad (4.15)$$

$$k_4 = \vec{v}(\vec{x}_i + \Delta t k_3) , \quad (4.16)$$

where $\vec{v} = v(|E|(\vec{x}))\hat{E}(\vec{x})$ is the drift velocity at a given location as a function of the electric field.

Interpolation

The points which make up the drift map are not sampled on a grid, but rather the trajectories/field lines themselves are selected. As such, query points between the samples can't be interpolated using methods used on regular grids, such as bilinear interpolation. Instead, a triangulation must be created, such that the triangles enclosing the query points can be interpolated instead.

Given a set of points, there are many ways to form a triangulation, where no two triangles overlap. There is one arrangement, the *Delaunay* triangulation, which forms a graph such that each triangle's circumcircle contains only that triangle's vertices. This has the side effect of keeping the vertices of each triangle a minimal distance from their centroids. Delaunay triangulations are dual to the *Voronoi* diagram, which is the set of regions defining the nearest vertex to a particular location. A Delaunay triangulation can be calculated using the *quickhull* algorithm.

An *n-simplex* is a shape consisting of $n - 1$ points $x \in \mathcal{R}^n$. The 3-simplices in the Delaunay triangulation are also referred to as the *facets*. The $(n - 1)$ -simplices are referred to as *ridges*. For $n = 2$, the facets are triangles and the ridges are line segments. For $n = 3$, the facets are tetrahedra and the ridges are triangles.

The quickhull algorithm is a method for calculating the *convex hull* of a set of points. The convex hull is the smallest convex shape which encloses the given vertices. For instance, in 2D one can visualize the convex hull by placing nails in a board, then stretching a rubber band around the perimeter. Likewise in 3D this would look like a balloon extending around a cloud of points. The quickhull algorithm is conducted by selecting the two points farthest from one another. It then draws a line segment between those two points, and sorts the remaining points into two sets depending on which side of the line they fall. The algorithm then proceeds recursively on these subsets. For each set, the point farthest from the dividing line is selected as a boundary point and a line is drawn perpendicularly to the line from the previous iteration. This eventually results in a series of vertices and line segments defining the convex hull.

Unfortunately the quickhull algorithm does not automatically result in the Delaunay triangulation of a set of points in \mathcal{R}^n , only its convex hull. However, a particular transformation will allow the quickhull algorithm to compute the desired graph. Each point is extended to \mathcal{R}^{n+1} , with the last component being the sum of the squares of the remaining elements, *i.e.* $z_{in} = \sum_{j=0}^{n-1} x_{ij}^2$. This transformation is known as "lifting." The convex hull is computed on these "lifted" points, and the resulting $(n + 1)$ -simplices are projected downwards onto the original space. This results in the desired triangulation.

The popular open source software qhull² is used to perform these calculations. While this work was conducted in 2D, it was thought that a 3D implementation might be used at some point in the future. Qhull has the ability to compute the convex hulls for arbitrary dimension n . It also contains an implementation for a directed search of the conjugate graph. This implementation was augmented with an additional caching trick for the LZ problem. The

²qhull.org




Figure 4.7: Illustration of the principle of barycentric coordinates. Each vertex is a primary color (red,green,blue). The color at any point within the triangle is an RGB vector with values given by the barycentric coordinates λ .

previous query point was stored, which made repeated requests to nearby locations cheap. This speeds up the calculation of electron clouds initiated from a single location.

While the points in the simulation are generally highly irregular, there are regions where the repeating patterns of the system can lead to uniform spacing of vertices. The grid wires, forward field rings, and photomultiplier tubes are all evenly spaced, and therefore may result in regions of regularity. Normally this would be desirable, but the condition for the Delaunay triangulation is that no point is only contained within its own circumcircles. Therefore, a regular grid would result in a non-unique Delaunay triangulation, and in fact qhull exits with a failure code in this condition. To circumvent this in general, I “joggle” the points slightly, adding a small random component to each dimension before calculating the mesh.

After the facet enclosing a query point is located, the values of the field at each vertex are collected. These are then interpolated using *barycentric* interpolation. This type of interpolation is calculated by, in effect, selecting one vertex as an origin, and using the remaining points on the facet to define (non-orthogonal) axes. The unit vectors are then orthogonalized, and the field at the query point is calculated in this new space. The barycentric coordinates λ_i of a point \mathbf{x} of a triangle defined by $\mathbf{r}_1, \mathbf{r}_2, \mathbf{r}_3$ are given by:

$$[(\mathbf{r}_1 - \mathbf{r}_3), (\mathbf{r}_2 - \mathbf{r}_3)]\lambda = \mathbf{x} - \mathbf{r}_3 . \quad (4.17)$$

The typical use case for barycentric coordinates is the interpolation and smoothing of images. In this case, each vertex maps to a color, and the pixel value at each point within the triangle is a mixture of the vertices according to the barycentric coordinates. This is illustrated with Fig. 4.7.

Wall Attachment

Charge Loss Simulations

It was noticed early on that the fields near the wall are highly non-uniform. Field lines initiated near the edge often intersected the wall, a feature which did not abate with smaller time steps. This resulted in the particular attention paid to the simulations near the wall, detailed above.

The particular cause of this effect is the aspect ratio of the field shaping rings. Each section is 25 mm tall, with 3 mm gaps between each ring. The inner radii of the rings themselves are 5 mm from the inner radius of the PTFE wall. This results in the equipotentials needing to transition from the approximately parallel lines in the center of the detector to approximately concentric lines surrounding the rings. Field lines near the wall therefore alternately point inward and outward in a pattern following the field shaping rings. The somewhat sinusoidal pattern intersects the wall, resulting in regions where any free electron is directed towards the wall.

The fate of electrons which attach to the wall is unknown. They may eventually detach and form a source of delayed electron noise, or recombine with ions. They could perhaps conduct along the surface of the PTFE into the field shaping rings themselves. In principle these electrons could accumulate over time, causing distortions of the electric fields over the lifetime of the experiment. Xenon-1T reported a change over time of their wall position, which indicated the possibility of wall charging. They also reported regions absent of S2 due to the polygonal cross section of their wall.

These wall attachment zones extend approximately 3 mm inwards from the wall and are contained entirely outside the fiducial volume. These charge loss/wall attachment zones are illustrated in Fig. 4.8. Attenuation of S2 size contributes to loss of position resolution, which increases the likelihood that scatters are misreconstructed into the fiducial volume.

Comparison to data

While charge loss model was not used in SR1 simulations, a wall model is planned for SR2 and beyond. I examined the charge loss present in a dispersed source near the wall, ^{83m}Kr . This affords high statistics in the charge loss and non-charge loss region. The calibration injection datasets were selected for physical drift times $dT \in [60, 950] \mu\text{s}$, S1 areas $\in [100, 300] \text{ phd}$, and S2 area $< 10^{4.8} \text{ phd}$, to efficiently tag the Krypton data. Basic SR1 livetime cuts, *e.g.* e-train, muon, OD, and skin vetoes, were applied. The S2 areas were corrected based on position. Two slices of radius were chosen to compare: the first, between 68 and 70 cm, was near the wall, but not subject to significant charge loss. The second was S2 radius $> 70 \text{ cm}$, which is partially subject to charge loss.

For simulations I take the uncharged wall as a baseline. Positions are uniformly sampled with radii within 4 mm of the wall at $R = 72.8 \text{ cm}$, and the attachment probability was interpolated from the simulated drift map. The corrected S2 areas were sampled from the ^{83m}Kr events within the inner $R \in [68, 70] \text{ cm}$ slice and attenuated by the attachment probability. The boundary for the simulated events (4 mm) was chosen to match the approximate




Figure 4.8: Illustration of the “dead zones”/ charge-loss regions near the TPC wall. *Left:* Oscillatory field lines near the PTFE wall. The color indicates the field magnitude, and the black lines indicate the equipotentials. Electrons ionized between intersections will be directed towards the wall. *Right:* The simulation results near a section in the middle of the wall. Each dot indicates a point sampled along a field line. While the field line density does not indicate field strength, the apparent oscillations are real trajectories. The color represents the attachment probability, or the fraction of electron trajectories which begin at the particular point and end on the PTFE wall at 728 mm. The attachment probabilities were found through the Monte Carlo diffusion method.

power law of the charge loss “rain,” shown in Fig. 4.9. The rise, fit on the calibration data between $\log_{10} S2c \in [10^3.2, 10^4]$ is described by:

$$\frac{dR}{d \log_{10}(cS2[\text{phd}])} = R_0 \times (19.5 \pm 0.2[L]) \times (\log_{10} cS2[\text{phd}])^{0.69 \pm 0.02} \quad (4.18)$$

where R_0 is the activity per liter of the Krypton 41 keV_{ee} line, and R is the total charge loss activity across the drift times $[60, 950]\mu\text{s}$. A broad agreement was seen between simulations in data. The peak of the Gaussian in $\log_{10} S2c$ is shifted lower than in data. Otherwise, the low area slopes are largely consistent, validating the diffusion model for charge loss. In order to obtain this agreement, the slices in radius had to be tuned appropriately due to the fact that the data is being compared against $S2_R$ and real radius, and the position resolution is on the cm scale.




Figure 4.9: The volume-weighted charge loss probability distribution function, used for validation of simulations against data. The blue peak is ^{83m}Kr calibration single scatter events which pass a fiducial selection criterion. The orange distribution is selected from the same ^{83m}Kr calibration, but events with reconstructed radius near the wall. For the green histogram, points are drawn from a grid in real space near the PTFE wall, and the “attachment probability” $p(r, z)$, is interpolated. These interpolated values are then used to attenuate corrected S2 areas drawn from the ^{83m}Kr spectrum: $S2' = p(r, z)S2$. Binomial fluctuations were considered insignificant due to the large number of electrons in the Krypton peak ($10^{4.5}/g_2 \approx 672$), leading to fluctuations on the $\sigma_p = \sqrt{\hat{p}(1 - \hat{p})/n} \approx 0.02$ level.

Cancellation

It is not known on what timescale the electrons which attach to the wall eventually detach. Fluorine is highly electronegative, so it is conceivable that they remain there for some time. PTFE has a surface resistivity of $10^{15} - 10^{18} \Omega$ [133]. It is conceivable that the electrons which do not make it to the surface remain on the PTFE for an extended period of time, or are recharged by new events rapidly enough, to reach an equilibrium. The production rate of electrons, inferred from random trigger data, is 77 kHz, or 0.012 pA across the entire detector. This provides an extremely optimistic timeframe of 17 years to charge the surface of the detector to $1 \mu\text{C}/\text{m}^2$. Then, the radial component of the electric field can be made strictly nonnegative, preventing further charge loss.

The complexities of the boundary conditions near the wall create challenges for attempting to arrange charges such that the charge loss zones shrink or are removed entirely. Sev-

eral charge profiles were attempted. The profiles were strictly negative charges, as it is not expected that the positive ions will attach to the walls. Additionally, trap energies are 0.05-0.1 eV lower for electrons than for holes [134], meaning positive charges will be cleared out of the bulk slightly faster.

All attempted charge profiles were a pattern which repeated on the scale of the field shaping ring heights, $Z_0 = 24.925$ mm. Generally these patterns were smoothly varying with the exception of rectification, and placed a larger charge density in areas where the unperturbed electric field points away from the wall. These profiles were inspected:

1. Sinusoid raised to a power: $\sigma \propto \sin^4(\pi(\frac{Z}{Z_0} - \frac{1}{4}))$
2. Rectified sinusoid: $\sigma \propto \max(1 + \sin(2\pi(\frac{Z}{Z_0} - \frac{1}{2})), 0)$
3. Square wave function: $\sigma \propto \theta(\sin(2\pi(\frac{Z}{Z_0} - \frac{1}{2})) + 0.2)$
4. Rectified Sinusoid, raised to a power: $\sigma \propto (\max(1 + \sin(2\pi(\frac{z}{Z_0} - 0.52)), 0))^4$

This eclectic set of models were analyzed by evaluating the radial component of the electric field along a contour of constant r . In order to not be overly sensitive to the finite element modelling of the change in dielectric, ϵ from LXe to ptfe, the contour was chosen at $r = 727.4$ mm. The radial component oscillates in a manner reminiscent of a sinusoid, as seen in Fig. 4.10. The figures of merit for these models were the amplitude of oscillations along this contour, A and the mean around which they oscillate, μ . The phases were adjusted to to achieved maximum cancellation for each model. The maximum local charge density was similarly adjusted until the smallest peak E_r was observed.

Model (1) was observed to have the best performance, resulting in oscillations of slightly higher frequency but lower amplitude than the others. This was achieved at a charge density amplitude of $-0.675 \mu\text{C}/\text{m}^2$. As the macroscopic charge is on this scale based on the wall position, it is conceivable that the charge profile is similar to this.

However, arranging the charges in this way is predicated on the electrons having sufficient mobility to move and achieve equilibrium. From the observed event rate, it is unlikely that the charge loss itself could provide this charge. The likely source of such charge is from the VUV light from the grids and S2s, creating electron hole pairs within the surface traps of the teflon [134]. If this is the case, then the charge profile is likely smooth over the scale of individual field shaping ring panels, making the motivation of field cancellation tenuous. For these reasons, the fast-charging profile needed for charge cancellation is not included in the final simulation, only the slow-varying charges needed to match the wall profile.

Grids

A similar phenomenon to the charge loss on the walls is seen near the cathode. Because the cathode wires are the lowest potential in the TPC, and they define the boundary between the forward field region and reverse field region, there is necessarily a contour of local minima in electric potential. This contour extends outwards from the cathode wires themselves,




Figure 4.10: The radial component of the electric field near the TPC wall boundary, with and without a particular charge distribution. The blue contour is the original contour, and the orange is after the attempted cancellation. Reduced amplitude of the oscillations comes at the cost of smaller oscillation period.

and causes volumes nominally in the FFR to direct charge downwards into the RFR, and eventually to the bottom grid. These volumes extend approximately 1.1 mm upwards at their greatest extent. These “cathode excursions” contribute to the multiple-scatter-single-ionization (MSSI) background near the cathode. This topology is where a multiple scatter event loses all but one of its distinct S2s, causing an artificially low $S2/S1$ ratio, placing an electron recoil within the nuclear recoil band. The simulated electric fields surrounding these excursions are shown in Fig. 4.11.

In addition to the field at distances larger than one pitch from the electrode grids, two quantities of interest are the electron *transparency* and the shielding *inefficiency* of the gate grid. The inefficiency is defined in Ref. [135] as

$$\sigma = \frac{dE_A}{dE_C} \approx \frac{a}{2\pi d_A} \log \frac{a}{2\pi r_0}, \quad (4.19)$$

where $d_{A(C)}$ is the distance between the gate grid and either the anode(cathode), and $E_{A(C)}$ is the asymptotic electric field between the gate and anode(cathode). Thus an efficient gate will isolate the two regions effectively, while an inefficient gate will have a larger dependence between the two volumes. The LZ gate grid, ignoring the impact of the change in dielectric at the liquid-gas boundary, is 5%.

The transparency of the grids is defined as the fraction of field lines initiating in the cathode-gate region which end on the wires. Based on the density of field lines, one may




Figure 4.11: The equipotentials which trace out the cathode excursions. The saddle points between each wire define the locations where electrons are directed upwards or downwards.

presume that if the electric field in the gate-anode region is higher, then the gate would be 100% transparent. It turns out that, due to the anisotropic charge induced in the surface of the wires, it is possible for electric field lines to both initiate and terminate on the electrodes even when the field changes across the boundary. Requiring that the grid wires have strictly nonpositive induced charges along their circumference leads to a requirement for the ratio of fields on either side[126]:

$$\frac{E_{GA}}{E_{CG}} \geq 1 + \frac{4\pi r_0}{a}(1 + K) , \quad (4.20)$$

which, when applied to LZ, yields a minimum ratio of fields for full transparency of 1.53. The LZ drift field-extraction region field ratio for flat grids, taking into account the phase change, is 21, safely placing LZ into the full transparency region.

With full transparency, the field lines from the cathode side will compress away from the wires, or “funnel”, as they pass into the extraction region. This means that the electron trajectories for short drift lengths the S2 reconstructed positions will deviate from their axisymmetric predictions on the scale of the gate grid wire pitch (5 mm). At longer drift lengths the transverse diffusion cause the electron cloud to straddle multiple grid “holes.” Events with many extracted electrons (>10) can still have excellent position resolution due to the position reconstruction algorithm averaging over the electron positions at the liquid level.

In order to study the grid funneling effect, an example grid cell was modelled using the same software as the LZ axisymmetric model(gmsh and FENICS), but in 3D. This grid cell was a rectangular prism, with the central axis of the prism running through the middle of a hole in the mesh. In other words, the grid wires appeared as a square along the edge of the cell, with half of their circumference cut away. Cyclical boundary conditions were applied to the edges of the cell. While not matching the full aspect ratio of LZ, the electric field ratio on




Figure 4.12: Illustration of the principle of partially transparent grids, taken from Ref. [135]. A cathode at point A , and anode/collector at point P , establish a bulk field. The gating grid placed at G to shield it from the induced fields from the positive charges produced by the ionizing radiation at point Q . The wire radius r , pitch d , and distances GP and AG allow one to calculate the inefficiency of the grid $\sigma = \frac{dE_P}{dE_Q}$, along with the bulk fields from specific voltages.

either side of the wires matched using planar electrodes. A map of (x,y) coordinates below the gate were mapped to (x', y') positions above the gate using the same RK4 integration as before. The spatial compression ratio, along with the path length/temporal dispersion were then obtained. These are shown in Fig. 4.13.

Superposition

In order to rapidly test changes to the electrode configuration, I developed a set of user-friendly tools which can recalculate the new fields and drift maps. The field map calculations exploit the linearity of the Poisson equation. As long as the geometry which defines the Dirichlet boundary conditions do not change, solutions may be added or subtracted in a linear combination. A set of “basis fields” for LZ was created, where each electrode was biased to a unit voltage (1 V) and the remaining boundaries were grounded. The voltage basis fields which were created are the PMT arrays (biased to their actual voltages, as these are not expected to change significantly), bottom, cathode, gate, and anode. From a command line interface, each of these voltages may be specified and the correct field is output.

A limitation of this method was that the sagitta could not be easily changed, as they are a geometric feature, so they are fixed at the nominal SR1 values. Changes to the gate-anode extraction voltages will have a significant impact due to their proximity to one another, but changes to the cathode voltage has a negligible effect on the sagitta. This method can not

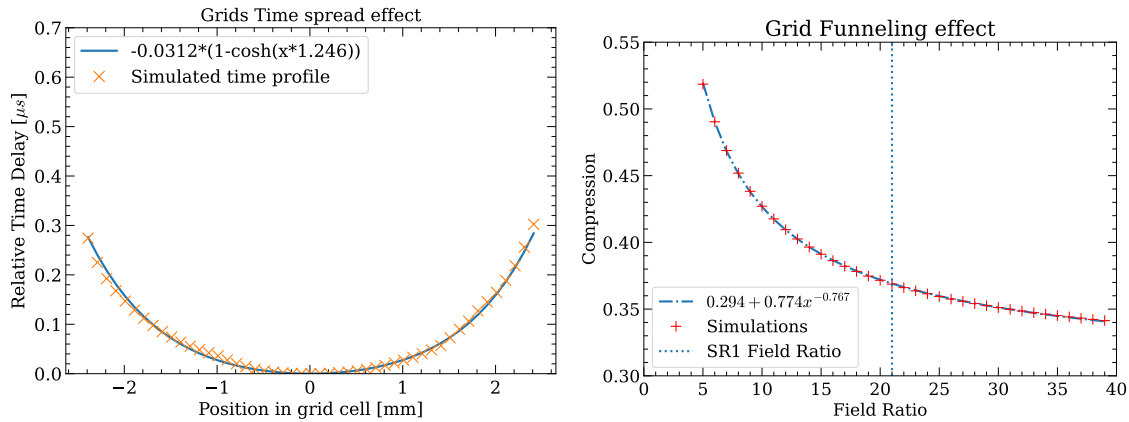


Figure 4.13: The funneling and time delay effect of the woven mesh grid as a function of displacement from the middle of a cell. *Left*: The time delay, relative to the middle of the cell, of electrons displaced from the middle. This is fit to a hyperbolic cosine. *Right*: The compression, or scaling factor, of electrons within a grid cell. The larger the ratio of bulk field magnitudes, the more the image of the cell is shrunk.

easily create new drift maps, as the integrated field lines are not local. The drift maps must be resimulated with every change to the potentials. However, for rapid testing the Monte Carlo diffusion was disabled, which allowed for quick scans over detector parameters.

4.4 Comparisons to SR1 data

Cathode Voltages

In order to quantify the effect of the wall attachment areas, and compare against SR1 data, two metrics were considered. First is the total volume which sees charge loss. This is parameterized by the minimum charge loss fraction and is compared against the charge loss events within calibration data. Second is the fraction of the wall which is invisible to S2s. This metric is compared against events known to take place on the surface of the walls, in this case ^{210}Po α decays. Between SR1 and SR2 a series of data were taken at a variety of voltages on the electrodes. This afforded an opportunity to study the impact of changes to the drift field on the wall attachment probability. The predictions for the charge loss volumes are shown in Fig. 4.14.

The wall events studied here are ^{210}Po α scatters. These are the result of ^{210}Pb atoms plating out on the PTFE during construction. The ^{210}Pb half life is 22.2 years, and it β -decays to ^{210}Bi with a Q-value of 63.5 keV. The ^{210}Bi then β -decays with a half life of 5.01 days to ^{210}Po , which itself has a half-life of 138.4 days. Somewhat uniquely, the ^{210}Po decays with a 5.3 MeV α which causes the daughter ^{206}Pb nucleus to recoil with kinetic energy




Figure 4.14: The fraction of charge loss regions near the wall ($R = 728$ mm) as a function of drift field. A charge loss region is a surface where either the radial component of the electric field points inwards, or the field line eventually reconnects with the wall radius. The “dead fraction” indicates the ratio of the total surface area of these patches to the total surface area of the wall. The data points were evaluated using isolated S1 data from ^{210}Po α s during the post-SR1 drift field scan. For the simulations (orange line) the trajectories were evaluated at a small standoff distance ($100 \mu\text{m}$) from the wall. Higher drift fields cause the electron trajectories to become more vertical, which increases the probability that an electron will eventually diffuse into the wall. The different curves indicate volumes where at least the specified amount of charge loss occurs.

106keV. Depending on the direction the α is emitted, the ^{206}Pb could be the only detected event, which is problematic as it is just barely outside the WIMP search ROI for LZ.

The method for evaluating the wall rate relied on *isolated S1s* (more detail in 7). These are S1 pulses without an associated S2. Since LZ triggers off of S2s, a special “long random” acquisition mode was utilized. Here, a square wave triggered the DAQ to write to disk at a rate which maximized the livetime. Additionally, the event window was expanded from 4.5 ms to 11.1 ms. This allows a search for S1s which did not have an associated S2. Events with near or total charge loss will show up in this dataset. Since the ^{210}Po α s are quite energetic, it is common to see photoionization on the grids from the large S1, even when no charge actually drifts to the surface from the location of the decay. As such, events with an S2 below 1000 phd were considered to be isolated S1s if their S1 area was consistent with being a ^{210}Po α decay. This is far below either the nuclear recoil band or the electron recoil band, requiring significant ionization signal loss if the S1 and S2 are causally related.

In order to select the S1s by area, a position dependent area correction had to be applied. Due to the different collection efficiencies of the top and bottom PMT arrays, the raw S1 area ends up being larger for events of the same energy near the cathode than events near the gate. This is corrected for single scatters based on drift time, but for isolated S1s this is not possible. However, the relative difference between the light collected in the top and bottom arrays, or *top-bottom-asymmetry* can be used as a proxy for drift time instead. At small S1, the binomial fluctuations between the two arrays makes the uncertainty on this estimate quite large, but with $\mathcal{O}(10^4)$ photons these fluctuations become subdominant to other sources.

After the TBA corrections are performed, a window of $S1c \in [34.25, 36.75]$ phd is searched for S1 pulses. This is taken as the rate of charge-loss ^{210}Po α -decays. This rate is divided by the rate calculated at zero drift field in order to estimate the fraction of the wall which remains S2-dead at various cathode voltages.

For the comparison against simulations, a scan of cathode voltages was performed. The wall attachment was calculated a short standoff distance (0.5 mm) from the wall. The fraction of that contour which exceeded 95% loss was considered the loss fraction and compared against the ^{210}Po - α data.

Wall Charge

Wall Positions

It was observed early into SR1 that the wall position was highly uniform. For simplicity, it was decided that simulations would use an uncharged drift map for the initial SR1 result, without any wall attachment modelling. The fiducial volume was then chosen in a data-driven manner. It is informative to validate this decision against data. Here, I compare the observed wall positions against the uncharged simulations, and perform an optimization in order to estimate the wall charge.

There are several ways to estimate the wall position. Here, I use the $^{83}\text{Kr}^m$ calibration data. The $^{83}\text{Kr}^m$ is produced via the decay of ^{83}Rb and is injected into the TPC volume. The isotope mixes with the LXe and decays with a 1.83 hour half life[136]. This provides a high-activity, monoenergetic, quasi-uniform calibration source. The decay proceeds via two internal conversion electrons at 9.4 and 32.1 keV, respectively. The combined energy of 41.5 keV_{ee} is just beyond the WIMP search ROI, which allows it to be regularly injected during the search itself, although this was not done during SR1.

The $^{83}\text{Kr}^m$ spatial distribution was not perfectly uniform, but it was sufficiently uniform on the scale of ~ 10 cm that analyses can perform analyses on spatial bins as if it was. Two methods of identifying the wall using this source were examined. The first was to divide the data into drift time bins, and find the radius where the activity within that bin dropped to half of its value at the center. The second was to look for charge loss events (events with S1 consistent with the $^{83}\text{Kr}^m$ source, but with severely diminished S2s), and calculate the mean of the S2 radii within drift time bins. The first method benefited from higher statistics, and had smaller uncertainties for the same drift time bins, so was used instead of the charge loss.

When searching for the loss of activity near the edge, it is assumed that the underlying distribution of scatter locations is an ideal step function, going from a flat rate when $r < r_{wall}$ to zero when $r > r_{wall}$. This underlying distribution is effectively convolved with a Gaussian which models the accuracy of the position reconstruction. The resulting profile at constant drift time then appears as an error function (erf), multiplied by radius due to the cylindrical volume element:

$$f_t(r; A, \mu, \sigma) = r \frac{A}{2} (1 + \operatorname{erf}[-\frac{(x - \mu)}{\sigma\sqrt{2}}]) , \quad (4.21)$$

where the μ parameter is the estimated wall position within the drift time bin, A is a normalization factor, σ is the resolution at the wall. The covariance matrix of the fit provides the uncertainty.

The proposed fields are then used to drift the individual points backwards in time using RK4, starting at the liquid level. In a perfect reconstruction, the points are all located at the wall radius at the specified drift time. Due to the complexity of the charge loss zones, instead of drifting all the way to the wall, the reconstructed $\hat{\mu}$ were scaled inwards to some standoff distance, and the locations were compared to that radius instead. For this work I used the modified wall radius of $r=72$ cm, but all plots here rescale the results to the actual radius of 72.8 cm. In order to not bias the result, the drift times are similarly corrected for the 8 mm $\sim 5 \mu\text{s}$ loss of drift distance. The loss function is the sum of the squares of the residuals of the reverse-drifted points from the modified radius.

The wall charge is parameterized as an azimuthally symmetric polynomial function. Each basis field was calculated from a charge profile of

$$\sigma_i(z) = A_i \tilde{z}^i , \quad (4.22)$$

where A_i is a scaling coefficient and $\tilde{z} \equiv z/146.1$ cm is the reduced height in the TPC. I calculate fields using polynomials up to $i = 4$. The function is not analytically differentiable, so a non-derivative based minimization algorithm had to be used. For this purpose the *Nelder-Mead* algorithm[137] was selected. Nelder-Mead uses simplices to identify a local minimum, and therefore does not require a derivative to converge. The downside is that a covariance matrix does not get calculated automatically.

The result of the fits to the various krypton injections is shown in Fig. 4.15. A general decrease in the integrated charge was observed, shown in Fig. 4.16, at a rate of approximately $12.5 \text{ nC m}^{-2}\text{day}^{-1}$ over this time period. This data period straddles both SR1 and some calibration data sets. The charging rate increases over time, possibly related to the changes in event rate during these time periods. Between the first two injections was almost the entirety of SR1, which was significantly lower activity than the calibrations data taken after the conclusion of SR1. As the grids were debiased on occasion during this time period, it is unlikely that the change in the charging rate is due to electron attachment on the walls. It is more likely that this is a function of VUV photons impinging on the PTFE over time, an effect implied by LUX RUN04 [120]. There, after a period of grid conditioning, the field was




Figure 4.15: The results of the estimated charge density on the reconstructed location of LZ PTFE walls. This was evaluated over all of the available high-statistics Kr83m injections. The total integrated charge has the largest impact on the radius of the bottom of the detector. While certain time bins exhibit larger errors, the overall profile is able to be accommodated with the polynomial charge profile. A trend of increasing estimated wall charge at later injections can be observed. The notation DXX indicates the day of the injection relative to the first (*i.e.* the first injection is D0).

observed to be highly nonuniform, requiring extensive drift corrections in order to analyze the data.

Some general features are apparent. All of the fits favor distributions with larger absolute charge density near the edges. The charge near the bottom of the detector is generally denser than that at the top. Because of the positive induced charge in the TPC field shaping rings, the charge on the This may be due to the effect of the cathode ring. The electric field is highly nonuniform due to it jutting into the TPC by approximately 5 mm.

Impact of wall charge on TPC Field

Beyond the effect that the charging of the PTFE wall has on the reconstructed wall positions, it is important to understand the overall impact on field uniformity and wall attachment. It might be tempting to assume that the addition of surface charge would result in an overall decrease in the charge loss near the wall, due to the increase in the radial component of the electric field, E_r . However, this is only true for a plane of charge without Dirichlet boundary conditions. Because the field shaping rings enforce a nearby voltage, the addition of negative charge on the walls induces positive image charges in the conducting volumes. This makes it




Figure 4.16: The estimated total surface charge accumulated on the TPC walls over time. The date is relative to the date of the first injection. The charging rate appears to increase over time, possibly due to the post-SR1 calibrations, which had an elevated event rate, along with changes to the grid voltages which produced copious light.

so that the field far from the wall falls off more similarly to that of a charged capacitor. Near the walls, the field is highly nonuniform due to the alternating pattern of induced charges, as shown in Fig. 4.18.

The overall impact on the charge loss volumes is shown in Fig. 4.19. Somewhat disappointingly, addition of negative charge has little impact on the volume of near-complete charge loss. Negative charge on the walls has a slight beneficial effect on the partial charge loss areas. This is likely due to the nonuniformities that are enhanced under the presence of surface charge: trajectories are either captured immediately (within the 25 mm distance between the field shaping ring resistors), or they are pushed further away, where diffusion is less likely to cause a loss of the electron.

Electron lifetime

Due to the small amount of nonuniformity in the field magnitude, the electron recombination probability will vary from point to point within the detector. If there is a correlation between the height where a scatter occurs in the detector and the electric field, this can mimic the effect of electron attenuation due to electronegative impurities. To investigate this effect, a monoenergetic source was simulated using the LZ simulations package LZLAMA.




Figure 4.17: The estimated surface charge density profile for the calibration data. Note the relatively larger density near the top and bottom of the TPC. Higher degree polynomials, or a spline fit may accommodate the data better.

Single scatter electron recoils with energy 40 keV (approximately $^{83}\text{Kr}^m$ -like) were placed uniformly throughout the TPC. The S1 and S2 response were calculated based on the local electric field. The electron lifetime for the simulation was set to infinity so as to remove the effect from the actual impurities. The resulting events were fiducialized away from the walls, in order to remove the nonuniformities at extreme r . Events with $r < 688$ mm were selected and binned in drift time. The mean S2 area was calculated within each bin, and a trendline was fit between 200 and 900 μs .

The intrinsic lifetime due to the electric field was found to be 289 ± 72 ms. This far exceeds the measured electron lifetime in SR1, 5 ms. Since the effective lifetime adds in inverse $\tau^{-1} = \tau_1^{-1} + \tau_2^{-1}$, I conclude that this has a negligible effect on the SR1 result. At drift times shorter than 200 μs a shorter intrinsic lifetime is seen. With real data, this would be confounded by the effects of the grids. Overall, this result is shown in Fig. 4.20.

4.5 Handoff

These simulations were incorporated into the larger LZ analysis and simulations toolchains. The handoff includes three components: the electric field map, the drift map, and the software to read and query the maps. Each field map is a mesh mapping $(r, z) \rightarrow (E_r, E_z)$, while the drift map projects $(r, z) \rightarrow (S2_R, S2_T, P_a, \sigma_R, \sigma_\phi, \sigma_t)$, where P_a is the probability

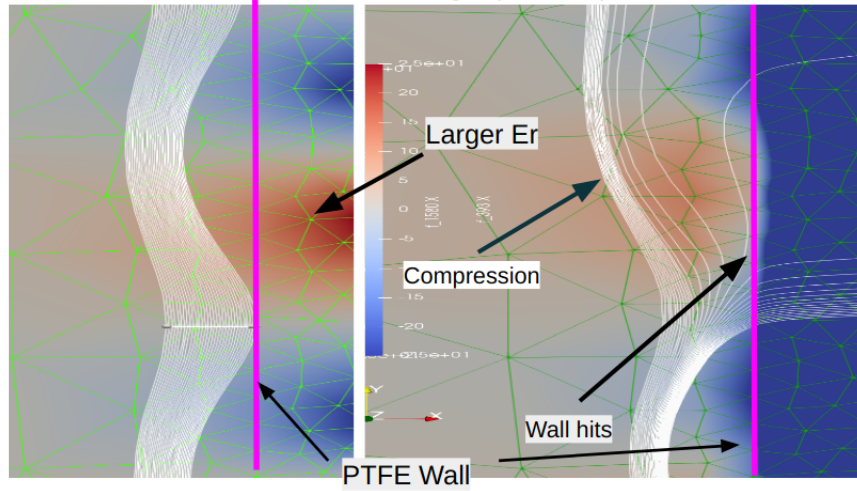


Figure 4.18: The impact of a uniform wall charge on the field lines near the walls. Both images were produced in Paraview, with the starting points of the streamlines being the same. The left image is uncharged, while the right image shows a wall charged with uniform $1 \mu\text{C}/\text{m}^2$ on the PTFE surface. The color maps the radial component of the electric field. The closest field lines to the wall begin to intersect with the addition of new charge. Those that manage to not intersect immediately are instead pushed further away, where they have in increased chance of surviving until the surface.

for the electron to become attached to the wall and the assorted σ are the relevant diffusion quantities. The points which are interpolated are not the same in real (r, z) space, because the drift map is based on integrating individual trajectories, and not the points from the field map. A method for a reverse drift map was developed. This involved a map from $(S2_R, S2_T) \rightarrow (r, z)$, and was accomplished in a rather straightforward manner. Since the maps were stored in comma-separated-value format, the real and S2 points could be transposed before running the quickhull algorithm to create the mesh. Due to the grid funneling affect, and the wall attachment, the points were downsampled prior to this step when doing the reverse map.

Together these form the *LZElectricField* package, or LZEF. LZEF is used in *BAC-CARAT*[138], the Geant4[139]-based Monte Carlo particle tracker for LZ. Within *BAC-CARAT* there are two methods of event generation: raytracing and deposit-only simulations.

Raytracing mode makes calls to NEST[85], where the quanta (photons and electrons) are generated. The electric field in raytracing mode queries the electric field at the location of each deposit, and uses that field as an input to NEST in order to calculate the recombination probabilities. The drift maps are queried in order to find the mean S2 location and drift time. Each electron within the cloud is then randomly perturbed in the transverse and longitudinal directions, with Gaussian distributions with standard deviations σ_T , σ_L , also interpolated from the maps. Note that the electrons, despite having access to the electric




Figure 4.19: The charge loss volume vs uniform charge density. These are the volume of xenon near the wall with charge loss probabilities exceeding 80% (blue), 90% (orange), and 95% (green), and were found with the stochastic diffusion model and not the deterministic integration. Due to the boundary conditions imposed by the nearby titanium field shaping rings, the impact on charge loss volume is relatively muted across this range of values. The field is distorted near the wall, which affected the charge loss *surface area* far more.

field vectors, are not transported step by step within the BACCARAT simulation.

In the deposit only BACCARAT mode, the LZEF package is not queried. Only the real locations, particle type, and magnitude of the energy deposits are stored. This is much faster, as individual photons and electrons are not tracked. The final S1 and S2 values are calculated in a post-processing simulation package called LZLAMA. Here, LZEF is queried for electric fields and S2 coordinates like in the raytracing mode. The main difference comes in the form of the effect of diffusion. With a parametric model, the electrons are not simulated, so the effect of the diffusion on the position resolution has to be simulated in a separate step. A single-electron resolution σ_1 is calculated, and the resulting S2 resolution is then $\sigma_N = \sigma_1/\sqrt{N}$. The single-electron resolution is a function of drift time and radius. In the end, the S2 reconstructed position is then sampled from a Gaussian distribution with the unperturbed location as the mean, and standard deviation σ_N . These parameters were tuned on data, and more detail is provided in Chapter 6.

Within the Raytracing mode, the work in the chapter has two additional features. The first is the ability to utilize the attachment probabilities which are calculated on a per-deposit basis. The number of electrons generated locally is calculated by NEST. This number is then used to calculate the size of the S2 based on the electron lifetime. The electron lifetime being




Figure 4.20: Uncorrected S2 areas as a function of drift time with simulated perfect xenon purity. The impact of the nonuniform electric field on S2 size is indicated with the red trend line. Monoenergetic 40 keV ^{83m}Kr events are simulated, illustrating the artificial/intrinsic lifetime due only to the electric field nonuniformities. The slope indicated in the red line is the result of the slight change of recombination probability as a result of the change in e-field magnitude over the height of the TPC.

τ_e , and the drift time being dT , the survival probability per electron is then $\lambda = \tau_e dT$. This λ becomes the expectation value for a Binomial distribution. Near the PTFE wall, the attachment probability P_A simply modifies the λ from the electron lifetime:

$$n_{\text{extracted}} \sim \text{Binom}(n_e, (1 - \tau_e dT)(1 - P_A)) . \quad (4.23)$$

The second additional feature is the handling of the electron funneling through the gaps in the gate grid. In raytracing mode each electron has its diffusion simulated between the initial point and the liquid surface. When the grid-funneling flag is set in BACCARAT, the electrons are binned into their respective pitch centers. From there, the compression ratio and temporal dispersions are calculated based on the field ratios, as shown in Fig. 4.21. The electron positions are then uniformly scaled inwards towards the center of the grid square, and the drift times are corrected. This has minimal effect within the fiducial volume, but short drift times have a larger impact.

For the SR1 simulations and analysis these additional features were not enabled. Since the events within the fiducial volume are largely insensitive to the wall attachment and the grid funneling, these effects were believed to be a complicating feature for the first science result.




Figure 4.21: Electron location scatter plots at the liquid level from grid funneling simulations. A gamma ray Compton scatters multiple times, producing distinct S2 positions. The energy deposit locations are indicated by the red x's. *Left*: The original distribution of electron positions at the liquid level. *Right*: The distribution of electron positions for the same event after the grid funneling feature was enabled.

Chapter 5

Multiple Scatter Dark Matter Search

5.1 Introduction

The typical WIMP-nucleon differential recoil rate approaches a value proportional to the ratio of the WIMP-nucleon cross section and the WIMP mass, $\sigma_{\chi p}/m_\chi$, when the WIMP mass greatly exceeds the target mass: $m_\chi \gg M_T$. Limits in this regime are determined by the exposure and acceptance of the respective experiments (see Eq. 1.31). For xenon, this approximation becomes valid beyond about $m_\chi \sim 100 \text{ GeV}/c^2$, and is observed in Fig. 2.7. Beyond the unitarity limit of $\mathcal{O}(100)$ TeV, however, the maximum allowable WIMP annihilation cross section is insufficient to explain the observed relic density[14]. This necessitates the introduction of theories which relax the assumptions of thermal relic WIMP theory, such as initial thermal equilibrium[36], weak interactions[140, 141], CP conservation[142], or pointlike particles[38]. Under such a theory, it becomes possible for an experiment such as LZ to set exclusion limits up to the Planck Mass. In this ultrahigh mass regime, the analysis techniques used for the GeV-scale WIMPs must be re-examined.

The main limiting factor for the largest masses that an experiment can probe is set by the particle flux. At $m_\chi = 40 \text{ GeV}/c^2$ (where most LXe experiments such as LZ[53], PANDAX-4T[100], and XENON1T[143, 144] set the tightest constraints on σ_p), there are frequent transits of dark matter particles through the TPC, but these transits rarely result in detectable nuclear recoils due to the large mean free path. As m_χ grows, the upper limit on σ_p increases and the flux of dark matter decreases, leading to an inverted situation where transits become infrequent, but each transit regularly results in a detectable nuclear recoil. Eventually we transition to a regime where only $\mathcal{O}(1)$ transits occur over the lifetime of an experiment, but those transits result in, on average, *many* energy deposits each. This violates one of the central assumptions of the traditional WIMP search: that dark matter events are singular nuclear recoils. Much of the background discrimination comes from the ability to distinguish single vs. multiple scatters, which necessitates new methods to achieve a low background search.

Several limits exist on ultrahigh mass, multiply scattering dark matter, in a variety of media. Recently, DEAP-3600 set limits on dark matter models exceeding the Planck

mass[145]. In this chapter I detail the extension of LZ SR1 WIMP search for ultrahigh mass dark matter with a multiple scatter event topology. My most significant contributions to this analysis was in the area of the simulations. I also made, consulted on, and developed a subset of the analysis criterion, and performed code refactoring as required.

5.2 Models

Composite Asymmetric Dark Matter

In the typical WIMP paradigm, dark matter is a fermion which annihilates with its anti-particle (which, if dark matter is Majorana, is the same particle) to achieve the correct relic density today. It is assumed that χ and $\bar{\chi}$ are produced in equal amounts. However, since the baryonic and non-baryonic matter densities are the same order of magnitude, $\Omega_\chi \sim 5 \times \Omega_b$ [3], this leads to the interesting possibility that the two densities could be related, with perhaps $n_\chi \approx n_b \rightarrow m_\chi \approx 5 \times m_b$.

The matter-antimatter asymmetry in the universe is currently unexplained. Three basic conditions must be met in order for baryogenesis to occur:

1. Baryon number (B) violation: it must be possible for an interaction vertex to change the B number.
2. C and CP violation: baryons must be distinct in the theory from anti-baryons, otherwise the B-violating processes would still alter B and \bar{B} at identical rates.
3. Out-of-equilibrium interactions: With the previous two conditions met, the comoving baryon number density n_b evolves over time. However, while in thermal equilibrium the production and destruction of baryons occur at equal rates. A departure from thermal equilibrium is necessary in order to overproduce baryons.

Together, these requirements are known as the Sakharov conditions[146]. In asymmetric models of dark matter, there exists a dark sector consisting of dark baryons, which couples to the standard model (SM) through some unknown portal. The initial asymmetry is produced in either the SM baryons, the dark sector, or both. The asymmetry is then communicated to the other sector through the portal. Afterwards, the symmetric portion of both fields annihilates away, leaving only the asymmetric portions. Because the density of the symmetric portion is lower than the total density, a slightly larger annihilation cross section is required [142], by small factor given by

$$\frac{m_\chi}{m_p} \frac{\eta(\chi)/q_\chi}{\eta(b)} = \frac{1 - r_\infty \Omega_\chi}{1 + r_\infty \Omega_b}, \quad (5.1)$$

where $\eta = (n_b - \bar{n}_b)/n_\gamma$ is the asymmetry of the species relative to the photon density, q_χ is the baryon number of the DM, m_χ and m_p are the DM and proton masses, respectively, Ω_χ and Ω_b are the fractions of the critical density made up of DM and baryons, and $r_\infty = \bar{n}_\chi/n_\chi$ is the fractional frozen out asymmetry.

Since the frozen-out portion of the dark matter is protected against further self-annihilation, the dark baryons may form bound states, held together by a mediator within the dark sector. This effectively decouples the theoretical problem of achieving the correct dark matter density from the mass of the dark matter state itself. Therefore, the unitarity bound no longer applies, and the bound state can achieve arbitrarily high masses, perhaps up to the Planck mass (above which the particle number flux becomes negligible for current generation detectors).

The interactions are generally considered as a result of scalar and vector interactions, analogous to the pions and vector mesons which carry the residual strong nuclear force. The Lagrangian is given by[147]:

$$\mathcal{L} = \bar{\chi}[i\gamma^\mu\partial_\mu - g_V\gamma^\mu V_\mu - (m_\chi - g_\phi\phi)]\chi + \frac{1}{2}[(\partial\phi)^2 - m_\phi^2\phi^2 - V(\phi)] - \frac{1}{4}V_{\mu\nu}V^{\mu\nu} + \frac{1}{2}m_V^2V_\mu^2, \quad (5.2)$$

where χ is the DM constituent fermion field, V is a dark vector field, and ϕ is a dark scalar. The attractive force from the ϕ coupling must be strong enough to efficiently fuse the dark nucleons into dark baryons, but it need not couple strongly to the SM. In general the mediator or portal need not be the same field ϕ which binds the nuclear state[34], though it is possible that the V_μ kinetically mixes with the SM photon, as in dark photon theories[35].

The formation and detection mechanisms of such bound states is explored in Refs. [34, 38]. There, the authors assume, based on nuclear physics, that bound states of sufficiently high dark baryon number enter into a state of constant density as a function of radius, known as saturation. Using a model incorporating both scalar and vector mediators (similar to nuclear physics, where the scalar mediator is the pion and the vector mediator is the photons coupled to the protons), they further assume that the surface tension on the state is negligible compared to the overall mass. The dark nuclear state can then be parameterized by the mass of the dark nucleon, m_χ , and the reduced mass when taking into account the binding energy, \bar{m}_χ . The radius of such a bound state with N dark nucleons is:

$$R_\chi = \left(\frac{9\pi}{4} \frac{N}{\bar{m}_\chi^3}\right)^{1/3}. \quad (5.3)$$

In the big bang nucleosynthesis (BBN) scenario, elements up to $A = 7$ can be synthesized efficiently. A bottleneck exists such that no stable nuclei exist with $A = 5$ or $A = 8$, suppressing further nucleosynthesis along with the Coulomb barrier which grows with Z . However, there is no *a priori* reason to suspect such a bottleneck with the dark nucleon fusion processes, and therefore dark fusion would have occurred until freeze-out when the reaction rate reaches parity with the Hubble expansion $\Gamma(t) \sim H(t)$. In Ref. [38], the bounds on the number of constituents in the bound state N_χ and m_χ were explored, finding

$$N_\chi \approx 10^{12} \left(\frac{g_\star(T_{syn})}{10}\right)^{3/5} \left(\frac{1 \text{ GeV}}{\bar{m}_\chi}\right)^{12/5} \left(\frac{\bar{m}_\chi}{n_{sat}}\right)^{4/5} \left(\frac{T_{syn}}{\bar{m}_\chi}\right)^{9/5}, \quad (5.4)$$

where T_{syn} is the temperature at which the fusion process begins. Self-interaction constraints are also relevant, since the dark nuclei have the potential for strong interactions between each

other. However, due to the dependence of N_χ on m_χ the self-interacting dark matter (SIDM) limits translate into upper bounds on N (and m_χ). These bounds depend on m_χ and n_{sat} but are restricted to $m_\chi < 10^{19}$ GeV/c² in order to avoid both excessive fine-tuning and to maintain the assumption of two-body interactions[38]. These models are variously referred to in the literature as “nuggets” or “blobs.” More recent papers adopt the model-agnostic term “multiply interacting dark matter.” One consequence of these models is that the extended radius of the dark matter leads to form factor effects, analogous to that of the target form factor. The recoil rate for MIMPs[148] is modified from the WIMP case[149] as follows:

$$\frac{dR}{dE_r} = A^2 N_\chi^2 |F_T(q)|^2 |F_\chi(q)|^2 \sigma_{1p} \frac{\rho_\chi}{2m_\chi \mu_p^2} \int_{v_{\min}(E_r)}^{v_{\text{esc}}} f(v) \frac{dv}{v}. \quad (5.5)$$

There are several possibilities for the mass distribution of the dark matter state, but it is thought that at sufficient number of constituents N_χ , a “saturation” occurs whereby the density becomes constant: $n \rightarrow n_{\text{sat}}$ [147]. For an extended object of constant density and radius R , the “top-hat” form factor is given by

$$F(q) = 3 \frac{j_1(qR)}{qR} = 3 \frac{\sin(qR) - qR \cos(qR)}{(qR)^3}, \quad (5.6)$$

where j_1 is the spherical Bessel function of the first kind. The A^2 coherence term is lost faster as a function of momentum exchange q with the addition of F_χ . When qR_χ exceeds the inverse spacing of the dark nucleons, the rate transitions to *incoherent* scattering, which scales with N_χ and not N_χ^2 [34]. Even with the Helm[150] form factor of xenon alone leads to a dependence of the collisional cross section $\sigma_{T\chi}$ on the velocity, as the maximum recoil energy is $E_{r,\text{max}} = rE_\chi \approx 2m_T v_\chi^2$, where $r \equiv 4m_\chi m_T / (m_\chi + m_T)^2$, when $m_\chi \gg M_T$. This effect is shown in Fig. 5.1. There, the attenuation of the collisional cross section is shown as a function of velocity, normalized such that $\bar{\Sigma}(v=0) = 1$. The collisional cross section is the average of the cross section over scattering angle:

$$\Sigma(E_\chi) = \sigma_{xp} \frac{\mu_{\chi T}^2}{\mu_{\chi p}^2} A^2 \int_{\sqrt{2m_T E_{\text{thres}}}}^{\sqrt{2m_T E_\chi}} |F_T(q)|^2 |F_\chi(q)|^2 \frac{q}{m_T} dq. \quad (5.7)$$

Composite states may also fuse from bosonic constituents [38], in which case the Pauli exclusion principle no longer forces the radius of the DM state to be large. In this case, the differential cross section is suppressed by the extent of the blob radius $R_\chi = \Lambda_\chi^{-1}$, but there is not a possibility of incoherent scattering as in the fermion case. Incoherent scattering refers to the case where the inverse momentum transfer q^{-1} becomes on the scale of the inter-constituent distance of the DM state $\Delta R = \Lambda_\chi^{-1} N_\chi^{-1/3}$. The finite size effects are therefore less significant for bosonic dark matter.

The mediator mass could be heavy or light. The heavy case leads to suppression of σ_T by $1/\mu^2$, where μ is the reduced mediator mass. In the case of a light mediator, *i.e.* $\mu < \min(m_T, m_\chi)$, an additional form factor is used based on long range interactions[34], given by

$$F_{\text{med}}(q) = q_0^2 / q^2, \quad (5.8)$$




Figure 5.1: Collisional cross section reduction as a result of coherence loss. The mode of the standard halo model is around 300 km/s, so this effect leads to significantly larger mean free path than the predictions from pointlike scattering.

where q_0 is an arbitrary scale factor which is absorbed into the per-nucleon cross section.

Freeze-in

Another way to circumvent the unitarity limit for thermal relics is to remove the restriction that the dark matter was initially in thermal equilibrium with the SM plasma. Instead, one may stipulate that the dark matter was produced gravitationally, possibly immediately following the inflationary era[36, 151]. Other mechanisms include inflaton decay during reheating and Higgs portal interactions[151]. This class of models is referred to as “freeze-in” models, to contrast the typical “freeze-out” WIMP paradigm. They never exist in chemical equilibrium, and their abundance increases until the temperature of the bath drops below their mass. The particles themselves are frequently referred to as “WIMPZillas,” referring to their large masses compared to the classical WIMP.

$$\mathcal{L}_\phi = \frac{\kappa_\phi}{2} \phi^2 \Phi^\dagger \Phi . \quad (5.9)$$

When the dark matter is generated through $2SM \rightarrow 2DM$ pair production from SM particles, the production primarily occurs at the temperature of the more massive particle [152], due to Boltzmann suppression $e^{-m/T}$. As opposed to freeze-out dark matter, the yield scales proportionally to the interaction cross section, rather than inversely proportionally. The Higgs-portal based model used in Ref. [151] parameterizes the density in terms of the Hubble constant at the end of inflation, H_e , and the temperature at the start of radiation

domination (reheating), T_{RH} . The number density of dark matter, n_χ , is a model-dependent quantity which generally scales with $T_{RH}^{-n_\chi}$, where $n_\chi > 0$. Both H_e and T_{RH} have well-motivated theoretical upper limits, and therefore m_χ can not be arbitrarily large in these models. In general, for both fermionic and bosonic WIMPzillas, $m_\chi < 10^{16}$ GeV/ c^2 in order to achieve the desired density[151]: The choice of model (i.e the spin of the particle) affects the comoving number density. This production mechanism can accommodate dark matter masses up to $m_{DM} \sim 10^{15}$ GeV, with the main limitation being constraints on the value of the reheating temperature T_{RH} .

In models where the WIMPzilla is produced via inflaton decay, the density is again proportional to the cross section[36]:

$$\Omega_\chi h^2 = m_\chi^2 \langle \sigma v \rangle \left(\frac{g_*}{200}\right)^{-3/2} \left(\frac{2000 T_{RH}}{m_\chi}\right)^7, \quad (5.10)$$

where g_* is the number of relativistic degrees of freedom. This mechanism also allows masses up to 10^{15} GeV. Due to flux limitations, as well as the typical assumption of feeble SM couplings, these models have not traditionally been converted to direct detection limits, but rather explored in the context of indirect detection. The decay of supermassive dark matter would produce extremely energetic cosmic rays. The Pierre Auger Observatory recently set limits on the dark matter lifetime in the range of $10^{14} - 10^{16}$ GeV resulting from their non-observation of events above $10^{11.3}$ GeV[153].

SIMPs

Multiple scattering dark matter necessarily has large SM cross sections, which may suggest a large self cross section, as well. An alternative to the typical $2 \rightarrow 2$ freezeout mechanism is a number-changing interaction $3 \rightarrow 2$, whereby a DM particle must meet two others in order to annihilate into two resulting DM states. This is known as ‘‘Strongly Interacting Dark Matter’’, or SIMPs[154, 140] as a result of the suggestive mass scale analogous to WIMP dark matter models. The Boltzmann equation in this case is slightly different:

$$\frac{\partial n}{\partial t} + 3Hn = -\langle \sigma v^2 \rangle (n^3 - n^2 n_{eq}), \quad (5.11)$$

which freezes out, just like in the WIMP case, when the reaction rate becomes on the order of the Hubble expansion $\Gamma \approx H$. This results in a mass scale[140]:

$$M_{\text{SIMP}} \approx \alpha \cdot (100 \text{ MeV}/c^2), \quad (5.12)$$

where α is the effective fine structure constant of these new DM forces. For $\alpha \sim \mathcal{O}(1)$, this implies a mass scale around the π_0 mass of 139 MeV, resulting in the SIMP naming.

This model could be a possibility for light multiply scattering dark matter, but such large values of the reduced cross section $\tilde{\sigma} = \sigma/M$ lead to significant effect from Earth shielding[155, 156], as discussed in Section 5.6. SIMPs risk losing too much velocity on their way to the detector and falling below the respective energy thresholds. Low-threshold, surface-operated detectors have the highest sensitivity to such models. An analysis of CRESST and

ν -cleus data[157] placed exclusion limits on such models up to $\sigma_{\chi p} \sim 10^{-27}$ cm². LZ struggles to observe such models due to Earth-shielding and 2 keV threshold, therefore SIMP models are not considered a viable MIMP candidate for this analysis.

Quark Nuggets

Dark sector models may include analogous versions of SM gauge interactions. Dark QCD $SU(3)_D$ leads to interesting phenomena, such as Dark Quark nuggets[158, 159]. The dark sector may contain a dark QCD with N_c “colors” and N_f “flavors” of dark quarks. This sector would experience a first-order phase transition, which in addition to dark baryons may nucleate dark quark states with large baryon number. The radius and mass of such states may span from keV-scale to PeV scale, and is inversely proportional to the temperature of the phase transition T_c . Probes of such “macro” dark matter include gravitational wave backgrounds, and X-ray bursts from nugget collisions[158]. SM couplings can result in direct detection signals but introduce model dependence.

Other models may lead to macroscopic dark matter, such as superconducting QCD states from axion domain walls[160], 6-quark states [161], or $Z(3)$ domain wall collapse[162]. These models explored in limits set using tracks in ancient mica and Skylab[163].

5.3 Backgrounds

While the WIMP search benefits from the extremely low single scatter background in the fiducial volume of LZ, the MIMP search is vulnerable to background sources which produce multiple scatters. These backgrounds require additional attention.

1. **Bismuth-Polonium coincidences.** The ^{214}Bi isotope β -decays into a ^{214}Po nucleus with a 3.27 MeV Q-value. This is followed by the 7.833 MeV α -decay of the ^{214}Po into ^{210}Pb . The half life of the ^{214}Po decay is 162.3 μs [164], placing it within the range of the TPC drift length. A MIMP search which extends down to double scatters would in principle observe such events as background. However, the high energy α far exceeds the energy of any nuclear recoil which may result from a WIMP. Additionally, the deposits are in the same physical location, which allows for analysis to remove them based on subtended length in the detector.
2. **Gammas:** Gamma ray photons may Compton scatter several times before being absorbed completely. Compton scattering is the dominant form of energy loss between 3 and 6 MeV[165], with a mean free path of ~ 12 cm at 3 MeV. While electron recoils can be distinguished from nuclear recoils based on S2 / S1 ratio, at low energy recombination fluctuations result in ERs being reconstructed as NRs more frequently (see Fig. 2.2). As a result, the summed S1 and S2 signals for gamma rays are more NR-like than single scatters of the same total energy. The angular distribution is predicted by the *Klein-Nishina* formula[166]:

$$\frac{d\sigma}{d\Omega} = \frac{\alpha^2}{2} \left(\frac{\omega'}{\omega}\right)^2 \left[\frac{\omega'}{\omega} + \frac{\omega}{\omega'} - \sin^2 \theta\right], \quad (5.13)$$

where $\omega(\omega')$ is the incoming (outgoing) photon frequency, $\alpha \approx 1/137.04$ is the fine structure constant, and θ is the scattering angle of the photon. Large deflection angles may accumulate over many scatters, making colinear tracks unlikely. As MIMPs are deflected by angles $\Delta\theta \sim m_T/m_\chi$, the signal tracks are more likely to lie along a single line subtending the detector.

3. **Muons:** A minimally ionizing muon at 226 MeV has a $\langle dE/dx \rangle = 3.58$ MeV/cm, depositing enormous energy into the detector. The resulting events have upwards of 10^8 TPC photons, and this quantity of light typically causes the detector to have an extremely high single photon rate for several seconds afterwards. Minimally ionizing muons, aside from their high energy, can be detected by coincidences with the OD veto detector.
4. **Electron Trains:** Following large S2s, an elevated rate of single electrons and single photoelectrons is observed over the course of tens of milliseconds, shown in Fig 5.2. These can mimic the appearance of MIMPs due to the fact that they contain many S1s and S2s. However, certain features make these not all that challenging to remove from the analysis. For instance, proximity to large S2s can predict the decay of e-trains, affording the use of an “e-train veto” precluding these time periods from further analysis. Occasionally, the e-train veto fails, as is the case with so-called “ghost muons” which are not tagged as a muon event for unknown reasons. Even in the case of such failures, these e-train events have prominent S2s of relatively uniform size which appear evenly distributed throughout the event window. As MIMPs have S2s which are at most separated by one full drift length, e-trains do not constitute a pernicious background. Additionally, e-train events frequently have S2s in the pre-trigger window, and have buffers which are not active at the start of the event, due to the activity of the previous event window carrying over. These qualities make them easier to remove.
5. **Neutrons:** Much like Compton scatters, Neutrons constitute a multiple scatter background, but are additionally concerning due to them being, by definition, multiple nuclear recoils. Neutron scatters usually have nearly isotropic scattering angles due to elastic scattering off xenon nuclei.

$$\frac{d\sigma}{d\cos\theta_{\text{lab}}} = \frac{d\sigma}{d\cos\theta_{\text{cm}}} \frac{d\cos\theta_{\text{cm}}}{d\cos\theta_{\text{lab}}} \approx \frac{1}{\cos\theta_{\text{lab}} - A^{-1}}. \quad (5.14)$$

The approximation is valid when $\cos(\theta) \gg A^{-1}$. Supplementary to the non-colinearity, neutrons can be captured on the Gd in the OD, which results in a series of detectable gamma rays. Tagging OD pulses allows for efficient rejection of neutrons.

6. **Pileup of single scatters:** Perhaps the most difficult to detect, this background consists of two true single scatters piling up, with one scatter occurring during the




Figure 5.2: An example e-train event waveform taken from commissioning data using the LZ event viewer. At this point in time, the pre-trigger window was 1.5 ms and the post-trigger window was 2.5 ms. Note the S2s which subtend multiple ms, along with the fact that the S2s begin not at the trigger ($t=0$), but at the beginning of the event window. S1s and SEs appear in the latter portion of the window.

drift length of the other. In principle this should occur as frequently as the accidental rate, with approximately one event in SR1. Three or more scatters should entirely eliminate this background. In this analysis this background is partially mitigated by requiring the time difference between S1s and S2s to be consistent with standard halo model velocities. In addition, since the single scatter background is concentrated along the edge of the TPC volume, fiducializing away from the boundary will further mitigate the background. We require a certain number of scatters (but not all) to occur in the fiducial volume, which reduces the change of pileup events appearing like a MIMP.

7. **Split S1s:** Cerenkov photons tend to have a longer tail than scintillation photons. This will occasionally lead to the pulses being split apart, like in Fig. 5.3. This is challenging to deal with without introducing too many restrictions on the pulse shape, which is undesirable since pulse merging must be assumed. However, it was observed that these split S1s are also frequently “high single channel” events, with more light concentrated in one PMT than expected from combinatorics. Removing events with such pulses based on cuts developed in Chapter 7 successfully mitigates this background.

5.4 Upper Bounds

One may ask how far into parameter space could a TPC even conceptually probe. Due to the low particle number fluxes involved, large integrated areas or long timescales are required. On the mass frontier the models and their limits start to transfer to MACHO-like limits,

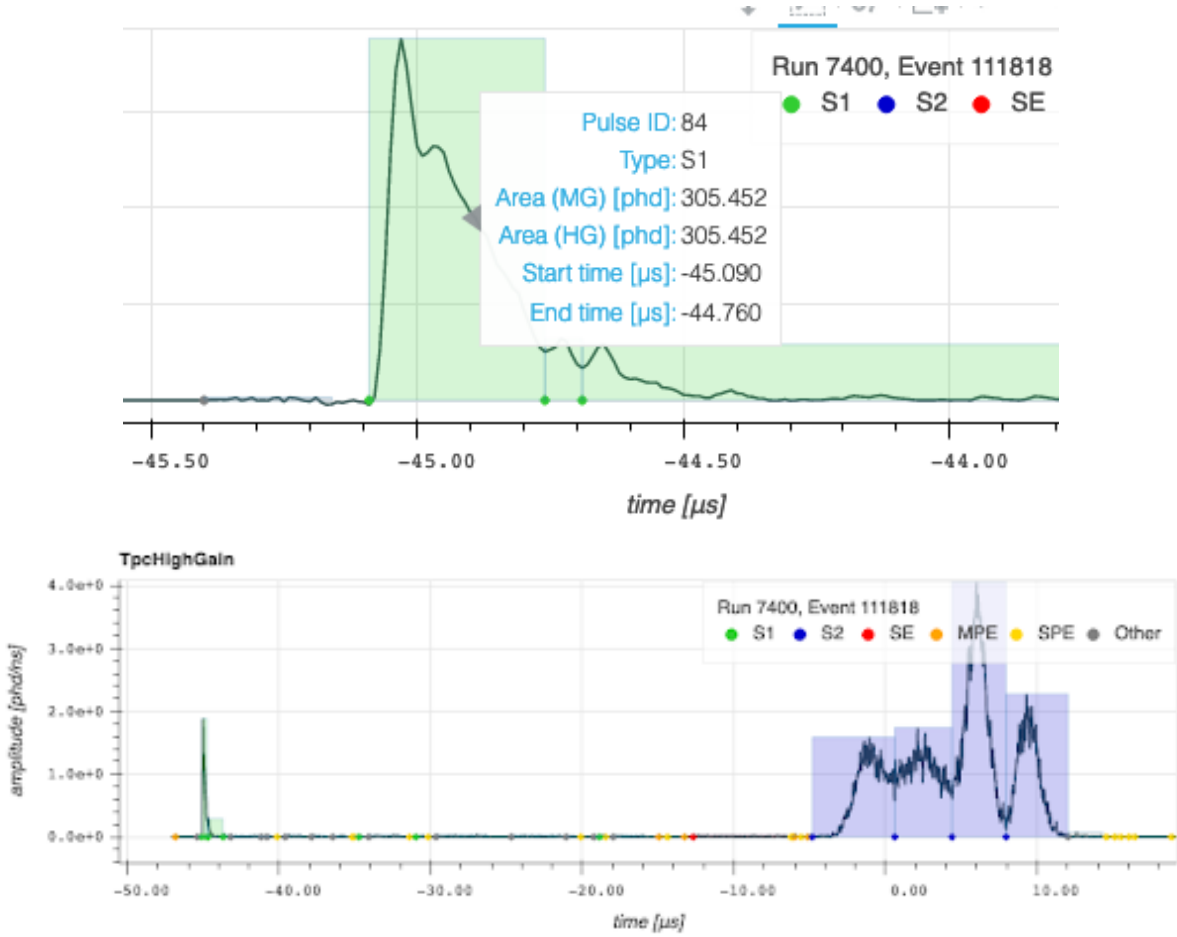


Figure 5.3: An example of a split-S1 event seen in the LZ event viewer. The S1 pulse has a longer-than-average tail, leading to three consecutive pulses identified by the pulse finding software. There were actually multiple S2s, suggesting the possibility of a γ -ray as a source.

and in any case require unrealistic integration times. For the cross section frontier, several interesting astrophysical constraints exist.

White dwarfs are stable on long time scales, kept in equilibrium by electron degeneracy pressure. If, however, they absorb dark matter at a high enough rate, a white dwarf may either heat to the point of igniting into a supernova type Ia, or collapse into a black hole. The existence of a 3 Gyr old white dwarf places a constraint of $m_\chi > 10^{15}$ GeV/ c^2 at the cross section where LZ becomes opaque to dark matter ($\sigma_{\chi p} \approx 10^{-30}$ cm 2) [167].

Interstellar gas clouds are another category that have been explored. Their immense size and mass leads to the possibility for frequent scattering with dark matter particles, if those particles have large SM couplings. The fact that these clouds have been observed with temperatures $\mathcal{O}(100)$ K allow limits to be set, as seen in Fig. 5.4 [168].

Instead of looking at large areas, one can search over long timescales instead. Ancient




Figure 5.4: Interstellar gas heating limits on contact interactions, figure from Ref. [168]. The y -axis is the DM-*nucleus* cross section, and the x -axis is the mass of the DM state. The “nucleon size” indicates a typical geometric cross section of a target nucleus ($1 \text{ fm}^2 \approx 10^{-26} \text{ cm}^2$). Assorted other limits are shown, such as those using ancient mica[169], the 2017 CRESST surface run[157], DAMA[47], and a modular liquid scintillator detector at the U. of Chicago [172].

mica crystals can be damaged by particle tracks over time, and this fact has been used to set limits on magnetic monopoles[169]. These limits can be interpreted as limits on the dark matter tracks, and the much higher energy threshold is compensated for by the GYr exposure. The upper limit on cross section was found to be[168]:

$$\sigma_{\chi p} < 10^{-15} \text{ cm}^2 \left(\frac{m_\chi}{10^{15} \text{ GeV}} \right), \quad (5.15)$$

above a threshold of $\sigma_{\chi p} > 4 \times 10^{18} \text{ cm}^2$. All of these limits, save the white dwarf supernovae limit, are far above the cross section regime that LZ is sensitive to without performing advanced pulse-shape analysis. A particular design choice which is beneficial to high-cross section searches is using modular detectors. In TPCs, large cross sections appear as an extended pulse of light, requiring a more sophisticated treatment of the waveform analysis, especially when it comes to saturation and position reconstruction. Modular experiments can detect depositions in each of their independent detectors, effectively removing the upper bound on cross section while still leveraging the colinearity of the DM tracks. Experiments such as DAMA[170], the MAJORANA Demonstrator[171], XENON1T[143], and an EJ-301 liquid scintillator experiment at the University of Chicago [172] have performed such analysis.

5.5 Projected Limits

Before selection criteria development and Monte Carlo simulation, an analytic, background-free limit was calculated as a best case scenario. A Feldman-Cousins[88] upper limit was calculated from the flux, excluding models predicting more than 2.44 events in the SR1 exposure (60 live days[53]). The limits are double valued for a given M_χ in general, due to the fact that the Poisson probability increases a maximum at $N = \lambda$ before decreasing. Because of this, I calculate the maximum excluded MIMP mass as a function of the cross section. This is done by calculating the flux of detected MIMPs at a given mass and cross section $\Phi(\sigma_i, M_0)$, through the surface area S , over the livetime T_{SR1} and scaling to find the required mass:

$$M_i \geq \frac{2.44M_0}{T_{\text{SR1}}S\Phi(\sigma_i, M_0)}. \quad (5.16)$$

The differential flux of *transiting particles* of a given velocity is given by:

$$\frac{d\phi}{dv} = \frac{\rho_\chi}{m_\chi}vf(v). \quad (5.17)$$

To calculate the flux of events, the expected number of scatters is calculated as $\lambda = \sigma(v)\bar{l}$, where \bar{l} is the average chord length. Cauchy's theorem[173] provides a value of the average chord length of a convex shape:

$$\bar{l} = 4\frac{V}{S}, \quad (5.18)$$

where V is the volume and S is the surface area of the object. The ratio of volume to surface area can be understood from dimensional analysis, while the pre-factor of 4 is a result of integration over the incident normal angles. For a square cylinder $h = 2r$ this gives $\bar{l} = 1.5r$. The probability of finding a transit with exactly N scatters is given by the Poisson probability:

$$\lambda = \mu^{-1} \equiv \bar{l}\sigma n_T \quad (5.19)$$

$$P(N; \lambda) = \frac{\lambda^N \exp(-\lambda)}{N!}, \quad (5.20)$$

where μ is the mean free path. The final step is to integrate the velocity distribution over the surface elements $\vec{v} \cdot \hat{n} = v \cos \theta$. Integrating over the incoming hemisphere gives $\frac{1}{2} \int_0^1 \cos \theta d \cos \theta = \frac{1}{4}$. The result is then

$$\Phi_N = \frac{\rho_\chi S}{4m_\chi} \int_{v_{\min}}^{v_{\max}} \frac{\lambda^N e^{-\lambda}}{N!} vf(v) dv \quad (5.21)$$

Note that this is a generalization of the usual single scatter result. Plugging in $N = 1$ to find the probability of single scatters, and the average chord length, the volume-scaling result is recovered:




Figure 5.5: The extrapolated single scatter exclusion limits for LZ (orange) and Xenon1T (blue), along with best-case scenario, background free, 100% acceptance MIMP search for LZ (green; solid is assuming sensitivity to 3-10 scatters, dotted is 2-100). Note the general features of the MIMP search: it pushes to higher masses than the single scatter search, the minimum number of scatters modifies the slope of the upper limit, and the maximum number of scatters extends the cross section higher. In this space the effect of overburden is negligible[174, 171].

$$\Phi_1 \approx \frac{\rho_\chi}{m_\chi} V n_T \int_{v_{min}}^{v_{max}} \sigma(v) v f(v) dv \quad (5.22)$$

Note that in the typical $\log \sigma_{\chi p}$ vs. $\log m_\chi$ limit plot, the upper limit from below is given by $\sigma \propto m_\chi^{1/N_{min}}$, where N_{min} is the minimum number of scatters necessary to be observed in the MIMP analysis. I also extrapolate the single scatter results from the SR1 WIMP search exclusion limit[53]. SR1 is neither background free nor a Feldman-Cousins limit, but in the high mass regime where the spectrum is independent of mass, one can calculate an effective 90% CL upper limit on the expected number of signal events. In essence this modifies the 2.44 number from FC. Taking the fiducial volume, SR1 exposure, and limit calculated at $M = 10 \text{ TeV}/c^2$, I extrapolate the SR1 single scatter limit to $10^{17} \text{ GeV}/c^2$, shown in Fig. 5.5. I do the same for Xenon1T, matching the area and average chord length of their unique fiducial volume shape[143], which is a cylinder with cutoffs near the corners.

5.6 Simulation

Monte Carlo

In order to estimate the selection criteria efficiencies, the multiply scattering events were simulated for the LZ detector. The intention was to observe the effect of certain pulse-based cuts on the spectra when the WIMP calculation is unfolded. In particular, it was necessary to see how frequently certain pulses merge together, confusing the search for multiple S1s and S2s.

Isotropic spectra are frequently simulated as initiating from a large hemisphere of radius much greater than that of the detector scale, and randomly sampled velocity vector. However, this method results in many simulated events which are not transits, and it was desired to save computing resources by only simulating true transits across the TPC. As such, the Monte Carlo sampling proceeded in the following steps:

1. Sample \vec{x}_0 , the initial starting point, from the surface of the TPC.
2. Sample the angle relative to the normal, $\cos \theta = \vec{v} \cdot \hat{n}$. This can be sampled using the inverse CDF, by sampling $x \in [0, 1]$, and transforming to $\cos \theta = \sqrt{x}$. The correct flux normalization then requires division by 4.
3. Sample the speed (equivalently, kinetic energy) of the incoming particle. This is taken from the first moment of the velocity distribution (the SHM[44] or SHM++[175]). In the LZLAMA (see Chapter 6.) simulation software, it is possible to sample the velocity using the inverse CDF method, or to simply weight the flux by the velocity probability distribution function.
4. If overburden(see Section 5.6, [174]) is being simulated, attenuate the incoming speed of the particle according to Eq 5.29. This is a function of the angle relative to the Earth's surface θ_E , and the value of σ/M . Downgoing trajectories are generally not heavily attenuated until σ/M is several orders of magnitude larger than the values which completely attenuate upgoing trajectories.
5. Track the trajectory of the particle through the detector. The deflection along the path is assumed to be negligible due to the large difference in masses. The interaction length l_i is calculated at the start. Each deposit is generated a distance $d \in \exp(-x/l_i)$ from the previous deposit, where x is sampled from a uniform distribution. The recoil energy is then sampled from the form-factor distribution in energy space:

$$P(E_r < E) = \int_0^{\sqrt{2m_T E}} |F(q)|^2 \theta(2m_T v_\chi^2 - q^2/2m_T) dq , \quad (5.23)$$

where q is the momentum transfer.

6. Simulate the S1, S2, and drift times using NEST[176] and LZLAMA (Chapter 6). The overall procedure for simulation is described in chapter 6. The length of the pulses

is simulated, and when pulses overlap they are merged with a probability which was tuned to commissioning data taken before the SR1 WIMP search exposure.

For maximal post-processing flexibility, the “uniform sampling” method was utilized, where instead of sampling the velocity distribution with the inverse CDF method, the probability distribution is used to weight the flux. Each event then has a weight attached to it, which is proportional to the frequency of transits with those properties in the detector. This is distinct from the units used in the single scatter simulations, where the weight is inversely proportional to the exposure, i.e. each weight value is in units of 1/ton/year. The flux/weight for each sampled event is given by:

$$w_i = n_\chi S v_i f(v_i) \frac{E_{\max} - E_{\min}}{N_{\text{events}}} (1 - \exp(-n_T \sigma_{\chi T} L)), \quad (5.24)$$

where S is the surface area of the detector, N_{events} is the number of simulated events, $f(v_i)$ is the velocity probability distribution function, and L is the transit length. The factor at the end corrects for the transits without any deposits, allowing the simulation to only generate transits that could possibly result in scatters inside the TPC. Even with this capability, some events do not trigger if certain conditions are met. For example, if the transit goes entirely through the RFR, it will not generate an S2, and therefore will not trigger the DAQ.

Pulse Merging

The pulse merging algorithm is critical for understanding the acceptance of the simulations. The S1 and S2 pulses have separate algorithms in the energy-only simulation which approximate the effect of the pulse finder in LZap (for more details, see Chapter 6 and Ref [138]). For S1s, a non-deterministic algorithm was developed. It calculates a merger probability based on proximity and relative ratio of the S1 areas. The map which is used to calculate this was created from ^{83m}Kr slow chain simulations. At a separation of 100ns, S1s will merge with 100% probability.

In the case of S2s, a deterministic model was used. The S2 pulse widths were calculated based on drift times, with a width of $2 \mu\text{s}$ at the center of the detector. Pulses with overlaps were merged if the edges were overlapping by an amount determined by the ratio of their areas. The minimum overlapping time was a tuneable parameter which was adjusted to match the DD neutron calibration data. In that case, the success metric was that the distribution of number of reconstructed S2s matched between the two models. When pulses are merged, the boundary of the pulse becomes the minimum and maximum of the two original pulse edges, which guarantees transitivity.

Contact Interactions

In spin-independent (SI) WIMP interactions, the WIMP-*nucleus* scattering cross section $\sigma_{\chi T}$ scales as A^4 times the WIMP-*nucleon* cross section $\sigma_{\chi p}$ when $m_\chi \gg m_A$. This is the result of the coherent sum of the contributions of the nucleons, and the scaling of the reduced mass,

both of which contribute a factor of A^2 . When considering ultraheavy dark matter, this runs into some theoretical issues. Since the 90% confidence exclusion limit in the ultraheavy limit scales as $\sigma \propto m_\chi$, the WIMP-nucleon cross sections become large as $m_\chi \rightarrow M_{\text{Planck}}$. This leads to multiple-scatter events, explored below. The cross sections that result in these events happen to coincide with the geometric cross section of xenon, which has important implications.

The normal SI scaling result comes about from the application of the first Born approximation. The scattering amplitude is given by :

$$f^{(1)} \approx -2\mu_T \int_0^{r_T} V(r') r'^2 dr' , \quad (5.25)$$

where r_T is the radius of of potential well. This result is valid when the potential is weak enough to not form bound states if $V(r)$ is attractive. The scattering cross section can be expanded into partial waves:

$$\sigma_{\chi T} = \frac{4\pi}{k^2} \sum_{l=0}^{\infty} (2l+1) \sin^2(\delta_l), \quad (5.26)$$

where δ_l are the phase shifts. Partial-wave expansion is the result of expanding the incident plane wave into spherical harmonics and solving the Schrödinger equation. Approximating the nucleus as a sphere of constant potential for $r < r_A$ leads to phase shifts above $\delta_l \approx kr_A$ which are rapidly attenuated. For s-wave scattering, it is assumed that the cross section is dominated by the $l = 0$ term, and therefore the value of δ_0 . This leads to the result that the maximum possible cross section is provided by

$$\sigma_{\chi T} = \frac{4\pi}{k^2} . \quad (5.27)$$

Solving the spherical top-hat potential results in a region of validity for the first Born approximation[177]:

$$\sigma_{\chi T}^{(1)} < \frac{16}{9} \pi r_T^2, \quad (5.28)$$

where the superindex (i) indicates the order of the Born approximation, here is the first. This result is valid in the weak coupling limit, where the first term of the Born expansion is sufficient. Stronger couplings $V(x)$ require further terms of the expansion, *i.e.* $\sigma_{\chi T} \approx \sigma_{\chi T}^{(2)} + \mathcal{O}(|V|^3)$ for the second Born approximation. The complete calculation completely saturates at the geometric cross section[177], $\sigma_{\chi T} < 4\pi r_A^2$. Due to the sharp cutoff in radius, higher partial waves do not allow this limit to be superceded greatly for repulsive potentials. Attractive potentials have the ability to form resonances, which can enhance the scattering cross section. For xenon (A=131), this saturation occurs above approximately 10^{-24} cm^2 , which translates to $\sigma_{\chi p} \sim 10^{-33} \text{ cm}^2$.

Somewhat unfortunately the maximal geometric cross section is slightly smaller than the cross section needed for multiple scatters to dominate over single scatters in LZ. Since the

xenon radius is approximately 5 fm, this occurs when $\bar{l}\sigma_{\chi T}n_T \approx 1$, where \bar{l} is the the average chord length and $n_T = 6.022 \times 10^{23}\rho_T/A_T$ is the number density of the target, which for LZ equates to $\sigma_{\chi T} = 7.87 \times 10^{-25} \text{ cm}^2$.

This result motivates the reporting of two alternate cross section limits. The spin independent A^4 scaling limits are to be shown alongside the “contact” interactions, where the per-nucleus limits are calculated. This has been done for Majorana[171] and DEAP-3600[145]. This has minimal effects for the analysis within a particular experiment as long as the detector medium is a single element with isotopes relatively close in A . The distinction between per-nucleus and per-nucleon limits come into play mostly for comparisons between experiments, where xenon’s A^4 advantage no longer applies over other media. The form factors for each target are still calculated, and lead to suppression of the rate of more energetic recoils.

Contact interactions such as these will greatly impact the overburden curve. Since there is no longer a penalty for smaller nuclei, the stopping power now only depends on the density of blocking matter. This means that, while the per-nucleon and per-nucleus limits can be approximately scaled into one another with A^4 , the overburden lower limits must be recomputed between the two cases.

Overburden

Attenuation from Earth Shielding

For DM masses above a few TeV/ c^2 and cross sections above $\approx 10^{-31}$, the mean free path through the Earth becomes low enough that transiting dark nuclei can lose a significant amount of kinetic energy before reaching the detector. Planck-scale dark matter is sufficiently massive that this is a minor effect, but for masses below $\sim 10^{12}$ GeV this overburden effect can attenuate the incident flux below detectable levels. The change in velocity when $m_\chi \gg M_T$ is given by Eq. 5.29, where the sum is over the various nuclei[174].

$$\log\left(\frac{v_f}{v_i}\right) = \frac{L}{2} \sum_i n_i \sigma_{i\chi}(v) \log\left(1 - 2\frac{m_i}{m_\chi}\right) \approx -\frac{2L}{m_\chi} \sum_i \rho_i \sigma_{i\chi}(v) \quad (5.29)$$

With the caveat that the cross section depends on v through the nuclear form factors of the nuclei that the dark nucleus encounters, the attenuation in velocity scales with $v_f \propto v_i \exp(-\sigma/m_\chi)$, which in turn results in a power law in the lower limits curve. As pointed out in Ref. [174], at sufficiently high cross sections, the number of nuclei between the surface and the detector with which to scatter saturates. After the cross section exceeds this value, there is no additional cost in terms of velocity attenuation, and therefore the exclusion limits may in principle extend arbitrarily high in $\sigma_{\chi p}$. This phenomenon is known as “saturated overburden scattering”(SOS) and is relevant for analyses sensitive to large number of scatters. In practice, the maximum number of scatters considered in an analysis typically causes the exclusion contour to turn around before the SOS threshold is reached. These analysis constraints are set by the fact that excessive amounts of energy deposition in

a detector will likely result in an interpretation other than multiply scattering dark matter, e.g. muons or grid discharge.

It should be noted that Eq. 5.29 does not take the nuclear form factor into account. A more realistic calculation utilizes the first moment of the nuclear form factor to calculate the expected energy deposition, which is no longer half the kinematic maximum, $rE_i/2 \approx 2m_T v_i^2$

$$\log\left(\frac{v_f}{v_i}\right) = \frac{L}{2} \sum_i n_i \sigma_{i\chi}(v) \log\left(1 - \frac{\langle E_{ri} \rangle}{E_i}\right) \quad (5.30)$$

$$\langle E_{ri} \rangle = \frac{1}{E_{r,\max}} \int_0^{E_{r,\max}} |F_i(q)|^2 E_r dE_r . \quad (5.31)$$

Additional perturbations occur due to the decrease in the mean free path as the DM transits the Earth, resulting from the various form factors approaching 1 as $v_\chi \rightarrow 0$. The results are shown in Fig. 5.6, where the attenuation predictably grows with the angle below the horizon, with discontinuities at the major layer transitions. Increases to the reduced cross section $\tilde{\sigma} \equiv \sigma/m_\chi$ results in stronger attenuation at a given altitude.

Density profile

The density profile used for modelling the attenuation of energy is given by the Preliminary Reference Earth Model [178]. This provides a profile of the density and atomic composition of each layer of the Earth. The densities, additionally, are modelled as piecewise quadratic. In order to find the mass subtended by the track on its way to the detector, these quadratic terms must be integrated. I use the following terms, some of which are illustrated on Fig. 5.7:

- R_\oplus = Earth's radius
- d = depth from surface
- l = length along track
- L = total length of the track.
- θ = angle between track and vertical from detector.
- $\theta' = \pi - \theta$ = interior angle.
- ϕ = angle between track and the Earth's surface normal at the entry point.
- θ_e = angle between the entry point, the center of the Earth, and the detector
- $x \equiv R_\oplus - d$




Figure 5.6: Energy attenuation as a function of incident angle, indicating The overburden or Earth shielding effect on different reduced cross sections. Here, a 300 km/s MIMP of mass 10^{17} GeV/ c^2 is simulated with various incident velocity vectors. The y -axis indicates the ratio of the incident kinetic energy at the LZ detector at the 4850 ft level of LZ to the kinetic energy at the surface. The altitude of the mantle and core layers are indicated, but there are multiple sub-layers present. The scan is performed over reduced cross section, as the effect of overburden is proportional to σ/m_χ per Eq. 5.29. Below a reduced cross section of approximately 10^{-7} barns/TeV, the upper hemisphere is almost entirely unaffected.

Assuming we know θ , d , R_\oplus , we want to solve for the remaining variables. First, we find the subtended length L . The Law of Sines states that:

$$\frac{L}{\sin \theta_e} = \frac{R_\oplus}{\sin \theta'} = \frac{R_\oplus - d}{\sin \phi} \quad (5.32)$$

$$\sin \phi = \frac{R_\oplus}{R_\oplus - d} \sin \theta \quad (5.33)$$

$$\theta_e = \pi - \theta' - \phi = \theta - \phi \quad (5.34)$$

$$L = R_\oplus \frac{\sin(\theta - \phi)}{\sin \theta} = R_\oplus \frac{\sin(\theta - \sin^{-1}(\frac{R_\oplus}{R_\oplus - d} \sin \theta))}{\sin \theta} \quad (5.35)$$




Figure 5.7: An illustration of the geometry involved in the transits of ultramassive dark matter particles through the earth. In this picture the detector is located at the top of the circle at a depth d below the surface.

The sectional density N (having units mass / area) is given by the integral of the density $n(r)$ along the track.

$$N = \int_0^L n(r) dl = \int_0^L (ar^2 + br + c) dl \quad (5.36)$$

$$r = \sqrt{R_{\oplus}^2 + l^2 - 2R_{\oplus}l \cos \phi} . \quad (5.37)$$

The constant and quadratic components are:

$$\int_0^L r^2 dl = \int_0^L (R_{\oplus}^2 + l^2 - 2R_{\oplus}l \cos \phi) dl = R_{\oplus}^2 L + L^3/3 - R_{\oplus} L^2 \cos \phi . \quad (5.38)$$

The calculation for the linear component is slightly more involved.

$$\int r dl = \int dl \sqrt{R_{\oplus}^2 + l^2 - 2R_{\oplus}l \cos \phi} = \int dl \sqrt{(l - R_{\oplus} \cos \phi)^2 + R_{\oplus}^2 (1 - \cos^2 \phi)} . \quad (5.39)$$

Making the substitutions $y \equiv l - R_{\oplus} \cos \phi$, $a \equiv R_{\oplus} \sin \phi$ this becomes:

$$\int r dl = \int dy \sqrt{y^2 + a^2} ,$$




Figure 5.8: *Left*: The density profile of the Preliminary Reference Earth Model (PREM). *Right*: The saturated overburden curves for tracks incident from above and below. These curves are calculated by finding models where the incident velocity (here $\beta = 10^{-3}$), are attenuated to a velocity where they can no longer result in a 1 keV recoil (this is approximately 20 km/s for Ixe). In addition to the additional distance when coming from below, the presence of a dense iron core $A = 55$ leads to greater stopping power.

which can be solved with a hyperbolic substitution $y \equiv a \sinh uL$

$$\int r dl = \frac{1}{2} [a^2 \sinh^{-1}(\frac{y}{a}) + ay \sqrt{1 + y^2/a^2}] , \quad (5.40)$$

providing an estimate of the projected density along the track, which gives $M = \sigma \rho_S$. I assume the composition is uniform along the track to obtain the isotopic fractions. Each element has the overburden saturate separately (for A^4 scaling, when the DM-nucleus cross sections differ between elements). As the particle slows down, the form factor suppression diminishes, which increases the stopping power. In order to simulate this second order effect, the cross sections are re-evaluated at each layer interface. Additionally, a simplified tracking procedure was implemented, whereby if the predicted attenuation exceeds a factor of two, the layer was subdivided and recalculated. The factor of two was chosen so as to only subdivide layers near the point of complete stopping. As xenon is larger than most of the elements in the Earth, the velocity dependence of the cross section from the various elements are much less than that of the detector itself. The saturated overburden curve, rounded at the corner by the number of isotopes, is shown in Fig. 5.8.

Modified Velocity Distribution

The Monte Carlo has two strategies for incorporating Earth shielding. The method described above is akin to rejection sampling, where a base distribution (the standard halo model[44]) is sampled, then the Earth shielding correction is applied to each particle. Another strategy

was considered for this purpose, which was to directly calculate the velocity distribution at the detector and sample from this modified distribution. This was ultimately not used for the simulations due to increased computation time, but it does provide an intuitive picture of the overburden effect.

The method starts by taking velocities at fixed intervals from 0 to the cutoff of the SHM at 800 km/s. Each of these “knots” are then transformed according to the overburden model at a specified zenith θ . The probability contained between the “knots” must be conserved. A modified probability distribution function is given by the following calculation:

$$f(v)dv = g(v')dv' \quad (5.41)$$

$$g(v') = f(v) \frac{dv}{dv'} = \frac{f(v)}{\left(v \frac{da}{dv} + a\right)}. \quad (5.42)$$

To generically evaluate the Earth shielding, one computes $a = \frac{E_f}{E_i}$ and takes the derivative with respect to incident velocity v (this is likely to be done numerically). Imagining a series of velocities v_n , which have their attenuations a_n evaluated. The amount of probability in each bin must be conserved under the transformation, so the transformation is given by the following equation:

$$g(v_n a_n) = f(v_n) \Delta v / (v_{n+1} a_{n+1} - v_n a_n). \quad (5.43)$$

If one keeps the above approximations, you can simply evaluate the attenuation at one mass, and safely rescale it to other masses. The mass of the DM particle only controls the average energy depletion per scatter. Some representative distributions are shown in Fig. 5.9.

At large reduced cross sections $\tilde{\sigma} = \sigma/M$ significant portions of the distribution are mapped to velocities near zero. This strains the numerical precision and leads to the possibility that the mapped velocities change orders, which would lead to negative Jacobians or divide-by-zero errors. Mapped velocities below the 1 keV_{NR} threshold in xenon of 20 km/s were placed in an overflow bin, and their probability masses were added to that value. The distribution was renormalized so that it would still integrate to 1 when performing the inverse-CDF sampling. Another technique to improve performance of the modified distributions was to take the resulting probability mass function, and then interpolate on a uniform spacing to create a new set of points to interpolate for the sampling procedure.

This was not used in the final calculations because for each point it required calculating a new distribution, with new points, each of which required a new attenuation each time.

Reduced Cross Section Scaling

To incorporate the effect of Earth shielding or overburden on the limits, I calculate several scans over σ and M in order to estimate the overall impact on the flux. For the xenon turnaround cross section of $\approx 10^{-30}$ cm², the overburden effect is negligible, so scan over

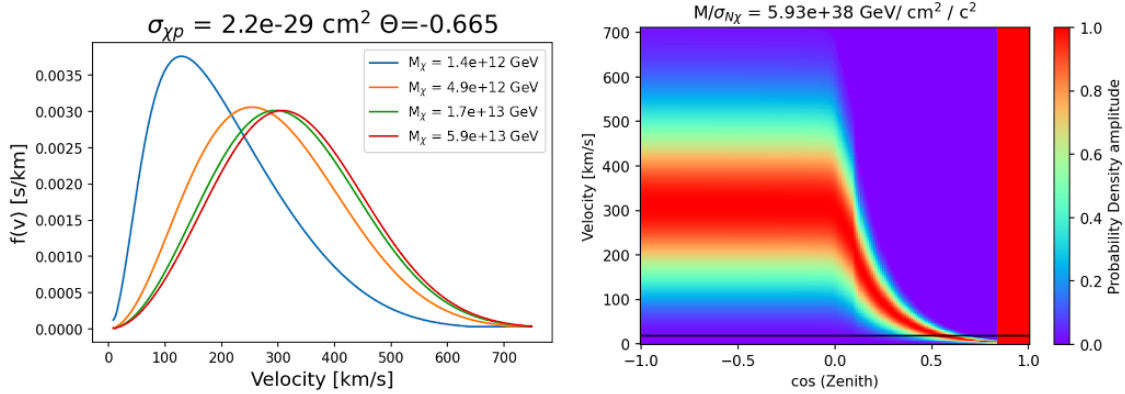


Figure 5.9: [The velocity distribution at surf after the attenuation of earth shielding. *Left*: Velocity distributions calculated from the SHM in the presence of overburden. Note that for a fixed cross section, decreasing the mass distorts the SHM towards lower velocities. *Right*: The 2D distribution of velocities, angles for another cross section. The red bar indicates complete stoppage of certain angles. Notice that the upper hemisphere is largely unaffected for this model, but the attenuation gradually increases below the horizon. The horizontal black line denotes the cutoff velocity of 20 km/s. Each horizontal scan is normalized such that the peak value in each vertical strip is 1, so this is not a true probability density function.

cross section of a fixed mass was performed. The grid scan for overburden was performed at a specific location in parameter space where multiple scatters still dominate single scatters, but the overburden effect is small enough to still allow signal through. This ends up being in the range $m_\chi \in [10^4, 10^6] \text{ GeV}/c^2$.

At each mass, successively higher cross sections were tested until the flux was too low to simulate further without prohibitive computational cost (approximately 1 MIMP event for every 10^5 transits). The maximum cross section was interpolated in log-space to find the cross section which would result in the minimum sensitive flux for SR1. A graph of mass, cross section was constructed, and the characteristic reduced cross section $\tilde{\sigma} = \sigma/M$ was determined by linear regression.

This process was hindered by the particulars of the simulations, which I had to modify in order to reduce the space requirements. The Monte-Carlo already guarantees one scatter in the TPC per transit, weighting the events properly to account for the mean free path. However, the Earth shielding is handled by first sampling from the SHM, then slowing down the particle based in its incident velocity and angle. This leads to many events being simulated which will never result in an observable scatter of 0.5 keV_{NR} (this is a loose analysis threshold, set by the threefold PMT coincidence requirement, along with the small impacts of data quality selection criterion. See Fig. 2.6 and Ref. [53]). I dealt with this first by resampling the entry point repeatedly until a surviving track was generated. This process is cheaper than writing an empty event to disk and starting the machinery over




Figure 5.10: The reduced cross section scaling of MIMP simulations.

again. These events were weighted by N_{trials}^{-1} . The incident velocity is not modified, so it is possible to become stuck in an infinite loop. I set a maximum number of iterations of 5000 to break out of the loop. When this occurs, I infer that this incident velocity is too low to be observed, and therefore should not be attempted again. The minimum velocity sampled from the velocity distribution through inverse-CDF is increased to the incident velocity + 1 km/s for all subsequent events. Events which follow are further reduced in weight by the lost probability in the velocity distribution to correct for this. Such changes drastically improved the ability of the simulations to identify the exact maximum observable cross section at a given mass. Simulation results are shown in Fig. 5.10. The lower limit for the reduced cross sections (the intercept of the overburden curves), are $\tilde{\sigma}_{A^4} = 1.91 \times 10^{-3}$, $\tilde{\sigma}_C = 325$ barns/TeV for the spin-independent A^4 scaling, and contact interaction (opaque to SM) models, respectively. In the A^4 scaling, the characteristic $\tilde{\sigma}$ is per-nucleon, while the contact interaction is reported per-nucleus.

Incident Velocity Vector

For simplicity the incident velocity vectors are sampled uniformly in angle relative to the horizon. However, the “WIMP wind” is not isotropic, but rather is biased towards coming from the constellation Cygnus at declination 42.03° and right ascension 20.62 h[45]. SURF is located in the northern hemisphere at latitude 44°N , and therefore the Cygnus constellation never actually drops below the horizon. This is highly beneficial for any dark matter search which probes high cross sections and is therefore affected by the Earth shielding effect. In the MIMP analysis, there is also a slight angular dependence to the analysis cuts due to pulse merging, which is shown in Fig. 5.11. Therefore, accurately sampling the altitude

of the incoming MIMP particles was performed in post-processing by reweighting the flux associated with each event.

I use the conventions in Ref. [45] to construct the angular distribution of transits at SURF, and I briefly restate the calculations here. In the rectangular galactic frame of reference, every halo particle has a velocity with components (v_r, v_ϕ, v_θ) , pointing inwards towards the galactic core, in the rotational direction of the galaxy, and above the galactic plane. This velocity is decomposed into the following components:

$$\vec{v}_{gal} = \vec{v}_{lab} + (\vec{v}_0 + \vec{v}_\odot + \vec{v}_\oplus(t)), \quad (5.44)$$

where $\vec{v}_0 = (0, 238, 0)$ km/s is the local standard of rest, $\vec{v}_\odot = (11.1, 12.2, 7.3)$ km/s is the peculiar velocity of the Sun, and $\vec{v}_\oplus(t)$ is the velocity of the Earth around the Sun. The Earth's velocity varies throughout the year but the mean value, taken at March 9, is $(29.2, -0.1, 5.9)$ km/s. The standard halo model [44] is a Gaussian of mean 0 and standard deviation $\sigma_0 = |\vec{v}_0|/\sqrt{2}$, with an additional cutoff of values above $v = 544$ km/s.

I sample points from the 3D Gaussian, and make the appropriate translations before calculating the laboratory frame magnitudes, which regenerates the appropriate SHM velocity distribution, peaked at ~ 300 km/s. To find the laboratory frame angular distribution, I sample points along the line of latitude containing SURF, in rectangular coordinates where the North Pole is the \hat{z} -axis. The vectors, which represent the \hat{z}' -axis of the laboratory frame, are rotated into in the galaxy frame. This is accomplished by first rotating the vectors in the y - z plane such that the galactic North Pole of declination 27° [179] is the new z -axis. The vectors are then rotated in the x - y' plane such that the right ascension of the galactic North Pole (192.86°) is the same as the celestial North Pole (122.93°). The resulting velocity magnitude distribution as a function of angle to the laboratory z -axis is displayed in Fig. 5.12, illustrating the bias towards incoming velocities. This serves as the basis to reweight the events in pos-processing.

5.7 Testing on LUX Run03 Data

The analysis was initially developed for LUX Run03 data. At the time, a mature lightly-ionizing-particle search had been performed on LUX data[180]. It was believed that adapting the analysis to multiply scattering dark matter would be viable. The main difference would be the spectrum and mean free path of the deposits.

I performed the simulations tasks for this work, using the LUX_BACCARAT software package alongside the LAMA fast chain. The MIMP physics were incorporated into LAMA, where the energy deposits were simulated coming from the surface of the TPC as discussed in Section 5.6. The particular geometry and detector response of LUX had to be implemented. It was decided that we would only use Run03 data, and not Run04, to avoid the problems of the extremely nonuniform field there[120]. Energy deposits from the WIMP NR spectrum were written out into a customized format and read into the Geant4-based LUX_BACCARAT simulations. These simulations run NEST[176] to calculate the quanta and photon gains. LUX_BACCARAT outputs a data format identical to the LUX data format, which can




Figure 5.11: The analysis efficiency as a function of altitude (angle to the horizon) for the turnaround model $\sigma_{\chi p} = 10^{-30} \text{ cm}^2$. A significant, but sharp, reduction in acceptance is seen at perfectly horizontal events.




Figure 5.12: Simulated horizon angles of SHM particles at SURF. The SHM velocity distribution in galactic coordinates was transformed into lab coordinates and averaged over the sidereal day.

then be analyzed according to the MatLab¹-based analysis frameworks to generate reduced quantities (RQs). An example waveform of a MIMP is shown in Fig. 5.13.

While the full chain generates waveforms and folds in the effect of the pulse finding code, it was found to not accurately reproduce the top-bottom-asymmetry of data, leading to significant data loss. As the software was mostly unmaintained, I decided to shift over to using fast chain, parametric simulations. Pulses were generated with NEST, and a simplified pulse merging procedure was utilized, where pulses within a fixed width were merged. Unmerged S1(S2) pulses had a width of 150 ns ($2 \mu s$).

A simplified version of a “pulse chopping” algorithm developed by LUX [181] was implemented for use in the fast chain. Each merged pulse tracked the original pulse areas and times. If a pulse was longer than a certain threshold, at the end of pulse processing it would construct the cumulative area for the merged pulse. The merged pulse would then be “chopped” into sections of 50 samples = 500 ns. The areas of the chopped pulses would be the area of the unmerged pulses (length $2 \mu s$) which overlapped with the region. The XY positions of the merged pulses were taken as the weighted centroids of the overlapping pulses, with the weights being the number of phe observed from the unmerged pulse in that region of time. This procedure helped to extend the LIP search into high charge fractions and in the MIMP search it will extend the search to higher multiplicities / cross sections.

Another algorithm which was explored was a form of “bootstrapping” for accurate predictions. A concern was that, for events which scatter in the reverse field region, S2s do not form, but S1s can. This causes the reconstructed lengths to be shorter, skewing the velocity distribution, as shown in Fig. 5.14. To mitigate this, it is possible to iteratively recalculating the positions.

The bootstrapping algorithm reconstructed the track using the initial reconstructed positions. This yields a track which may result in intersections with the reverse field region (RFR). The fraction of the reconstructed track length in the forward field region (FFR) is used to mask the S1s which are likely to have occurred in the RFR. Using these masked S1s, the S2 drift times are recalculated, and a new track intersection is generated. This procedure is repeated several times until convergence. Without head-tail information this can lead to errors, but the degeneracy may be broken with S1 TBS information, which is particularly useful in this case as the RFR S1s will have asymmetries near -1.

5.8 Analysis on LZ SR1 Data

Introduction

The dataset used for the MIMP search is identical to that used for the SR1 WIMP search[53]. This consists of 89 dark matter search days over 115 calendar days. After removing time periods due to *e.g.* the e-train veto, a final exposure of 60 live days was used for the result. The fiducial volume used was 5.5 tonnes. Most importantly for this analysis, the maximum drift time was 951 μs , with an electron lifetime in excess of 5000 μs , and a drift field of

¹www.mathworks.com/

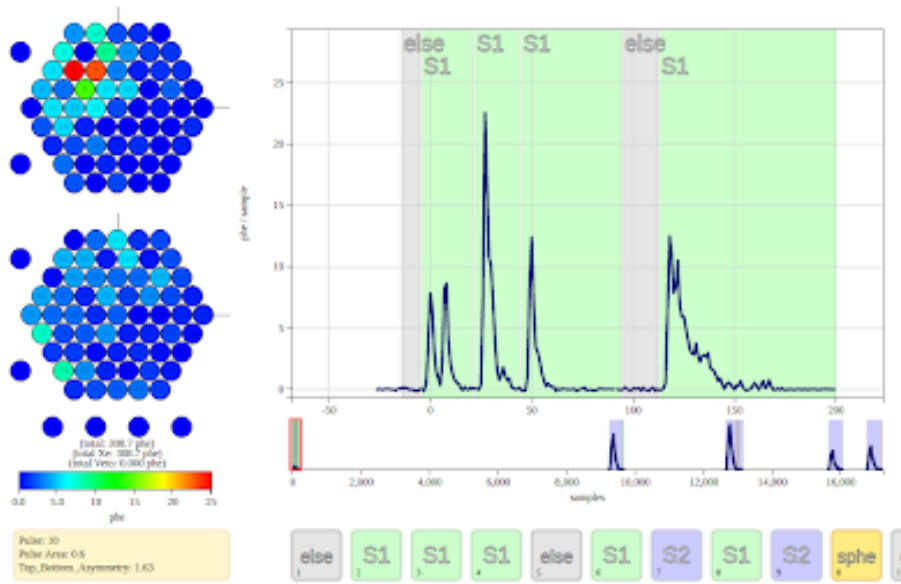


Figure 5.13: A simulated multiple scatter colinear even waveform in LUX, made using LUX-BACCARAT.




Figure 5.14: Unique reconstructed quantities for multiple scatter events, analyzed for the LUX simulations. These indicate the impact of the reverse field region on the analysis. *Left*: Reconstructed and true distances for simulated MIMPs of $\sigma_{\chi p} = 5 \times 10^{-30} \text{ cm}^2$. Red indicates reconstruction and blue is truth. Shorter track lengths are more difficult to reconstruct. *Right*: The reconstructed velocity distribution for MIMPs of $\sigma_{\chi p} = 5 \times 10^{-30} \text{ cm}^2$. Tracks which intersect the reverse field region are indicated in red, and events which do not are indicated in blue. RFR tracks are biased towards lower velocities, dominating the 0-100 km/s bin.

193 V/cm. Many of the analysis cuts were maintained from the SR1 WIMP search, with some modifications, detailed in the following sections.

Selection Criteria

1. **E-train veto, Muon veto.** These remove periods of time following large (10^4 phd or higher) S2s, which experience elevated single electron and single photoelectron noise. This constitutes a reduction of livetime of approximately 30% in SR1.
2. **Sustained Rate:** Removes events with SPE rate above 40 kHz across the entire pre-trigger region. This is an SR1 cut which is effective at reducing pileup S1s.
3. **OD Veto:** The GdLS outer detector effectively vetoes both neutrons and muons. MIMPs will not deposit significant energy into the OD due to the comparatively small nuclear recoil energies of most elements. For instance, ^{14}C will on average experience nuclear recoils of 11 keV_{NR}. Using the GdLS chemical formula[79] (C_6H_6)- $\text{C}_{(10-15)}\text{H}_{22-32}$, and density of 0.853 g/cm⁴, stopping powers of ~ 2 keV/cm are obtained for a value of $\sigma_{xp} = 10^{-27}$ cm². This is well below the muon minimally ionizing particle energy deposition of $\mathcal{O}(1)$ MeV/cm.
4. **Colinearity:** This leverages the negligible deflection in the incoming MIMP tracks. Each S2 pulse has its associated deposit position reconstructed based on the Mercury X,Y positions[86], and the Z coordinate is based on the reconstructed drift time. Because the tracks are nonrelativistic, and MIMP signals are required to have multiple S1s, this introduces some ambiguity on the drift time on a per-S2 basis. Due to the different thresholds between S2s and S1s, boundary effects of the RFR, and pulse merging, the number of S2s and S1s are in general different, preventing any systematic pairing of S1s and S2s. One may correct the drift time to either the first S1, last S1, or linearly interpolate in either the forward or backwards directions. SURF is located at 44° in the northern hemisphere, which beneficially causes most, but not all, of the WIMP wind to be downgoing. The upgoing tracks will have the first S1 be associated with the last S2 in general, as the maximum drift length of 955 μs far exceeds most transit times of 10 μs . Denoting the N S2 times T_i and the M S1 times as t_i , and the drift velocity as v_e , the corrections are given by:

$$z_{F,i} = v_e \left[T_i - t_0 - (t_{M-1} - t_0) \frac{T_i - T_0}{T_{N-1} - T_0} \right] \quad (5.45)$$

$$z_{B,i} = v_e \left[T_i - t_{M-1} + (t_{M-1} - t_0) \frac{T_i - T_0}{T_{N-1} - T_0} \right], \quad (5.46)$$

where $z_{F(B),i}$ is the forward(backward) corrected \hat{z} -coordinate of the scatter. These generally lead to percent-level corrections to the velocities. While the XY coordinates may, in principle, be corrected based on the drift maps discussed in Chapter 4, due to

the lack of head-tail information it was not performed. Because the drift map in the fiducial volume is to first order a frustrum, pushing the bottom of the detector inwards by ~ 2 cm, this leads to an error on the reconstructed direction, but is not expected to lead to improper residuals, as these are primarily sensitive to the curvature of the track. The simulations of tracks fold the nonuniform drift map, so the effect of this choice is reflected in the efficiency of the cut. The estimated uncertainty of the reconstructed altitude of the tracks, for the xenon turnaround model, for transits passing the “multiplicity”, “good S1”, and “uniformity” selections is 1.11 ± 0.08 degrees.

The colinearity requirement relies on a orthogonal-distance regression[182], where the perpendicular distances between the points and a line $\mathbf{Y} = t\hat{\mathbf{m}} + \mathbf{b}$ is minimized. This is differentiated to typical χ^2 minimization, which minimized the distance between one coordinate and the line $y = mx + b$. For two dimensions, the result is found by first calculating the covariance matrix σ_{ij} . The slope and intercept are given by[183]

$$m = \frac{\sigma_y^2 - \sigma_x^2 + \text{sgn}(\sigma_{xy})\sqrt{(\sigma_y^2 - \sigma_x^2)^2 + 4\sigma_{xy}^2}}{2\sigma_{xy}} \quad (5.47)$$

$$b = \mu_y - m\mu_x. \quad (5.48)$$

Generalizing to ND , this procedure is nearly identical to principle component analysis (PCA), where the eigenvalues of the covariance matrix are used to find the direction of maximal variance. Given centered(mean is zero along each axis) data \mathbf{x}_i , the following matrix is constructed [183]:

$$M = \sum_i^N (x_i^2 I - \mathbf{x}_i \mathbf{x}_i^T). \quad (5.49)$$

The eigenvector \mathbf{d}_0 of M with the smallest eigenvalue λ_0 minimizes the sum of the orthonogonal distances. Note that the $M_{ij} = \text{Tr}\Sigma - \Sigma_{ij}$.

The resolution is considered on a point by point basis with the following model:

$$\delta_R = \frac{1}{\sqrt{S^2}}[\sigma_0 + \sigma_1 \exp((r - R_w)/r_0)] \quad (5.50)$$

$$\delta_\phi = \frac{\sigma_2}{\sqrt{S^2}}. \quad (5.51)$$

The square root reflects the fact that the resolution is at some level based on the Poisson fluctuations in the top array. While the full covariance of each data point could be incorporated in principle, for simplicity and robustness each point was simply weighted by the inverse of the total error $\sigma_T^2 = \delta_R^2 + \delta_\phi^2$. In the z-direction the error is taken to be uniformly 1 mm, as the resolution there is set by the rising edge of the S2 and not the number of electrons. This translates to ≈ 500 ns, a conservative estimate for the uncertainty on the drift time of the cathode.

Both 2D and 3D trajectories were fit to the S2 points. For the purposes of selection criteria, only the 2D orthogonal distances were used. This was due to the mathematical challenge of propagating the error on the Z coordinate, which is calculated via the drift time rather than the Mercury position reconstruction algorithm[86]. The tails of the resulting distribution were challenging to characterize, and therefore only the XY was used for this purpose. The XYZ trajectory was used to calculate the entry and exit points in the detector, necessary to estimate the velocity. The XY colinearity requirement was selected as:

$$\tilde{\chi}_{XY}^2 = \frac{1}{N-2} \sum_i^N \left(\frac{d_{i\perp}}{\sigma_i} \right)^2 < \chi_{\text{cutoff}}^2 = 2. \quad (5.52)$$

The threshold $\chi_{\text{cutoff}}^2 = 2$ was found to have sufficient signal acceptance and background rejection. An example of an SR1 track with the reconstructed track overlaid is found in Fig. 5.17.

5. **Good S1:** In the WIMP search rejecting an S1 with a high single channel or an veto coincidence is the same as rejecting the event itself. In the MIMP search a separate decision had to be made regarding whether an excluded S1 should cause the event to be culled, rather than not counting the singular pulse. I made the conservative choice to remove events with such pulses present.
6. **Uniformity:** MIMP tracks are not only colinear, they should subtend the entire track length. The intersection of a line defined by points $(x_1, y_1), (x_2, y_2)$ with a circle is given by:

$$D \equiv x_1 y_2 - x_2 y_1 \quad (5.53)$$

$$d_r \equiv \sqrt{\Delta x^2 + \Delta y^2} \quad (5.54)$$

$$x = \frac{D \Delta y \pm \text{sgn}(\Delta y) \Delta x \sqrt{r^2 d_r^2 - D^2}}{d_r^2} \quad (5.55)$$

$$y = \frac{-D \Delta x \pm |\Delta y| \sqrt{r^2 d_r^2 - D^2}}{d_r^2} \quad (5.56)$$

The intersection points with the circle are then projected onto the cylinder in order to find the \hat{z} -coordinates of the entry and exit points. When the \hat{z} -components exceed the top and bottom of the cylinder, as is the case when the trajectory comes through the top and/or bottom, they are clipped and the missing distance is projected back onto the x, y components.

The uniformity cut's metric is the number of missing scatters along the reconstructed track trajectory. The estimated mean free path is then taken as the subtended length

of the scatters.

$$\hat{\mu} = \frac{|\mathbf{x}_{N-1} - \mathbf{x}_0|}{N - 1} \quad (5.57)$$

From the projected track length L the “missing scatters” are estimated as:

$$N_{\text{missing}} = (L - |\mathbf{x}_{N-1} - \mathbf{x}_0|)/\mu . \quad (5.58)$$

The threshold for N_{missing} was optimized by examining the receiver operating curve for the SR1 data and the simulated dataset at the turnaround cross section of $\sigma_{\chi p} = 10^{-30} \text{ cm}^2$. Due to the tails of the distribution and the reconstruction error, the threshold was chosen as $N_{\text{missing}} < 30$, a surprisingly large threshold. An alternative estimator for μ was using the estimated full length, L/N . This was found to perform worse for low numbers of scatters with the SR1 data, so was changed to the projected scatter length.

7. **Velocity:** The fact that the MIMP distribution is expected to follow the SHM [44] confers considerable leverage on the analysis. The velocity of any given track was estimated as the projected scatter length divided by the time difference of the S1s:

$$\hat{v} = \frac{|\mathbf{x}_{N-1} - \mathbf{x}_0|}{t_{M-1} - t_0} . \quad (5.59)$$

This estimator was found to approximate the input velocity distribution well, with some caveats. Because S1s are frequently lost due to fluctuations below detection threshold, and due to pulse merging, the $\delta t = t_{M-1} - t_0$ frequently underestimates the true time difference between S2s. The simulated velocity distribution has a tail which extends past the galactic escape velocity, as shown in Fig. 5.15. The boundaries for this selection were chosen as $\hat{v} \in [50, 1200] \text{ km/s}$, which was a conservative choice.

8. **Total S1 and S2 Area:** Individual deposits are expected to fall along the nuclear recoil band in liquid xenon. With pulse merging, selections based on any individual pulse become challenging. The more robust total S1 and S2 values are used instead. The selection criteria, illustrated in Fig. 5.16, are:

- $\log_{10}(S2_{\text{total}}) < \log_{10}(S1_{\text{total}}) + 3$
- $\log_{10}(S2_{\text{total}}) > \log_{10}(S1_{\text{total}}) + 1$
- $S2_{\text{total}} > 1200 \text{ phd}$

9. **Fiducial:** To mitigate the pileup and Compton scatter backgrounds a minimum number of points are required to be in a fiducial volume. Because of the multiple S2 requirement the volume can be less conservative than the WIMP search. The boundaries were chosen as $z \in [2, 135] \text{ cm}$, $r < 70 \text{ cm}$, and a minimum of 2 scatters were required to exist in this region.




Figure 5.15: A histogram of simulated MIMP velocities for a model with $\sigma_{xp} = 10^{-30} \text{ cm}^2$ using LZLAMA.

10. **Ordering:** In lieu of an XYZ colinearity requirement, a looser requirement on the z -coordinates are required. This specifies that, to the resolution of σ_z , the projected distances along the trajectory increases monotonically with z (equivalently, drift time). On a functional level this iterates through the S2 pulses, and determines the sign by the first $|\Delta d| > \sigma_z$, where d is the projected scatter length. Event with subsequent $\Delta d < -\text{sgn}|\sigma_z|$ are then rejected.
11. **Miscellaneous:** The minimum number of prominent S1s and S2s were configurable, based on the requirement that background free conditions were possible. It turned out that $N_{S1} = N_{S2} = 2$ was achievable. Using the SHM as a basis, I was able to set a hard cutoff of $5 \mu\text{s}$ for the S1 spread. I set a maximum S2 spread of 2 ms in order to accommodate the maximum drift time plus the potential for some e-train S2s to pass the event selection.

Modification of prominence algorithm

The fast chain simulation pulse merging was tuned to match simulations of deuterium-deuterium fusion simulations which typically had less than four prominent S2s. Prominence is explored more in Chapters 7 and 6, and refers to the property that a pulse is significantly larger than the noise level of pulses of its same type. A prominent pulse is meant




Figure 5.16: A heat map of simulated S1 and S2 areas for the multiple scatter analysis. The total S1 and S2 cut boundaries for the MIMP analysis are indicated. The model was the turnaround model, $\sigma_{\chi P} = 10^{-30} \text{ cm}^2$.




Figure 5.17: Example reconstructed tracks / individual S2 locations from SR1. This particular event passes the colinearity and uniformity cuts, but fails the velocity cut.

to correspond to the physical light pulse, rather than a false positive from the pulse finding code. It was noticed that events with larger number of deposits were being classified into the lower bins. This was a result of the non-causal nature of the prominence algorithm, which selects based on a pulse's height-to-width ratio and the ratio of its area to the largest pulse of its same type, regardless of the ordering. Multiple scatter events tend to classify S2s as "non-prominent" event if they occur before the largest pulse. This was determined to be a second order effect for the SR1 WIMP search, but was critical for the analysis of MIMPs, as the concern in the latter case was accurate reconstruction of the number of scatters, rather than simply single vs. multiple scatter differentiation.

I modified the prominence algorithm for this analysis to be causal, and applied it to SR1 data. The same cut boundaries were applied as in SR1, but the ratio was applied to the largest pulse of a particular type *up to that point* in the event window. In the code the raw pulses, the original non-causal prominent pulses, and the new causal prominence pulses were available, but for the reconstruction the velocities and the colinearity requirement only the causal prominent pulses were used.

Area Corrections

In order to further utilize the S1 and S2 areas, I implemented the area corrections in our analysis. These required some modification on a software level, since the versions for the WIMP search make single scatter assumptions. The S1 area corrections depend primarily on drift time, and S2 corrections are sensitive to the XY position as well as drift time due to the varying electroluminescence gain in the extraction region. The MIMP search S1c and S2c values are calculated for the three different estimates for drift time used elsewhere in the analysis: 1) the drift time from the S2 to the first S1, 2) the S1 transit corrected drift time, assuming downgoing trajectory, and 3) the S1 transit corrected drift time, assuming upgoing trajectory. The S1s have an explicit relationship between their top-bottom-asymmetry(TBA) and z -coordinate. As such I correct their areas based on their TBA so as to make fewer assumptions. Other than these small modifications, the functional dependence of S1c/S1 and S2c/S2 is identical to the SR1 WIMP search.

Count Prediction

As part of an effort to capitalize on the ER/NR discrimination power, I attempted to get accurate measurements of the mean S1 and mean S2 values for each event. The central limit theorem[184] states that these will converge at high multiplicity to gaussian distributions centered on the means of the bands. To achieve this, accurate estimates of the number of vertices for each track was required. The differing energy thresholds and widths for S1s and S2s result in a bias towards more observed S2s than S1s on average. The S1 merging depends on the number of S1s and the transit time, while the S2 merging depends on the number of S2s and the incident angle (horizontal tracks will merge into a single pulse). Various estimates for number of vertices are shown in Fig. 5.18, with the "Fit" explained below.

I chose to fit a linear regression model to the data using the python package `sklearn`², with the following features:

1. N_{S1} : number of S1s.
2. N_{S2} : number of S2s.
3. N_{S1}^2
4. N_{S2}^2
5. $N_{S1}N_{S2}$
6. Δt_{S1} : time spread of S1s.
7. Δt_{S2} : time spread of S2s.
8. Δt_{S1}^2
9. Δt_{S2}^2 .
10. $S1_{\max}$: the largest S1 area
11. $S2_{\max}$: the largest S1 area

The quadratic terms were introduced to match the nonlinearity of the data, and the time spread variables exist because pulse merging is a result of track density. Dimensionful quantities were divided by characteristic scales: $\Delta'_{S1} = 10$ ns, $\Delta'_{S2} = 100$ ns, $S1'_{\max} = 100$ phd, $S2'_{\max} = 1000$ phd. This model was trained with the `sklearn` interface, which utilized $\mathcal{L}_{\text{lasso}}(L1)$ regression.

$$L = \sum_i^N (y - \hat{y})^2 + \lambda \sum_j^M |w_j| \quad (5.60)$$

The regularization parameter λ I found using 5-fold stratified cross validation using a model of $\sigma_{xp} = 10^{-30}$ cm². The coefficients for the model weights are:

$$\hat{N} = \mathbf{X}w \quad (5.61)$$

$$w = \{.210, .899, .0209, .00315, -.0174, 2.869, \\ -1.019, 24.558, .0204, 1.575, 3.343\} . \quad (5.62)$$

The $L2$ loss, evaluated on the 30% holdout dataset, was 0.928. Vertex predictions are shown in Fig.5.19.

²scikit-learn.org




Figure 5.18: S1-S2 distributions using various estimates for the vertex counts, indicating the convergence towards the mean value.

Signal Acceptance

For the theory curves in Fig. 5.5 a known number of vertices is specified, with 100% for tracks with that number, and 0% elsewhere. This analysis does not set an explicit upper limit on the number of pulses, but rather relies on the pulse finding code, which merges S2s together if they are close in time. This effectively sets the upper boundary to whatever mean free path appears like a continuous energy deposition, which ends up being around 10^{-29} cm². I characterize the cut acceptances as a function of true vertex count in Fig. 5.20 for a model near the turnaround point (mean number of scatters in LZ = 1). The weighted acceptance for two or more scatters is $\eta_2 = 0.803$, and for three or more scatters is $\eta_3 = 0.684$. The highest acceptances are between 5 and 30 scatters, with over 95% acceptance, but as the number of vertices increases this value drops to $\sim 40\%$. These are evaluated relative to the multiplicity requirement, so as to test the quality of the analysis independent of the model. The fiducial cut is particularly costly at the higher vertex counts.

It is also instructive to show the dependence of acceptance on physical parameters, such




Figure 5.19: Pulse count predictions using the linear regression model. This indicates that using S1 or S2 along results in under-estimation of the true number of vertices (with S1 being worse in general). The linear regression model tracks the true number of vertices well.

as velocity or entry point, as shown in Fig. 5.21

Limit

The MIMP search limit was constructed in a similar way projected limit in Section 5.5. Zero background events were observed during SR1, which causes the limit to be set entirely by the surviving MIMP flux. Models were selected uniformly in mass ($M = 10^{17}$ GeV), simulated in LZLAMA without Earth shielding, then analyzed with the method described above. The surviving flux of events translates into an effective minimum mass as each point. Characteristic $\tilde{\sigma}$ from Section 5.6 were used to extrapolate the lower limit from above from the point of intersection. In Fig. 5.22 the A^4 limit for SR1 is shown, illustrating the impact of the analysis cuts. The 2+ scatter analysis has high acceptance, leading to minimal impact on the limit curve near the turnaround point. If the analysis required moving to 3+ the impact would be large, but still small at the turnaround model. To extrapolate the upper limit downwards, the minimum number of scatters, n_{\min} was used to scale the cross section:

$$\sigma \propto M^{1/n_{\min}} . \quad (5.63)$$

Extrapolation was utilized due to the computational cost of simulating multiple scatters when they are an extreme minority to single scatters, which occurs at low cross section. At 1000 live days, if zero background events are observed, LZ is projected to set a limit




Figure 5.20: MIMP search cut acceptances as a function of the number of vertices, evaluated on the turnaround model. The cumulative effect of various selection criterion are indicated.




Figure 5.21: MIMP acceptance relative to the multiplicity requirement of 2+S1s, 2+ S2s as a function of incident velocity and entry/exit points. Faster transits appear to be accepted more regularly, along with tracks which enter away from the bottom edge.




Figure 5.22: The LZ SR1 MIMP search exclusion limits for A^4 (left) and contact interactions right. The projected 1000-live days exposure for LZ, assuming a background free experimental result, is shown in dashed red on the left plot. The impacts of the analysis selection criterion appear minuscule when shown over this many decades.

of $6.27 \times 10^{18} \text{ GeV}/c^2$ at $\sigma_{\chi p} = 10^{-30} \text{ cm}^2$. Comparisons to other ultrahigh dark matter searches are shown in Fig. 5.23.

Dark Matter Radius

It was assumed in previous experimental results [171, 145] that the only models to consider were A^4 and contact interactions (opaque to SM). In the case of A^4 it was stated that one may circumvent the geometric cross section of the target by assuming a dark matter radius which is large enough to generate the desired per-nucleon cross section without needing to accommodate the model dependence of the dark matter form factor. This critical radius was supposed as $R_D \sim 1 \text{ fm}$, a typical nuclear distance scale. For completeness, I incorporate the dark matter form factor in the form of a 3D top-hat function, and show the effect of increasing the radius from zero.

The simulations proceeded in an identical fashion as the previous results. At the stage of calculating the mean free path at a particular DM velocity, $\mu(v_\chi)$, the xenon Helm[150] form factor is multiplied by the dark form factor from Eq. 5.6. These results are displayed in Fig. 5.24. As the radius of the dark “blob” or “nugget” increases, the excluded cross sections are shifted upwards.

The effect of overburden is two-fold. With the extended object, nuclear recoils are biased even more heavily towards lower energies. This means that, in addition to increasing the mean free path through the Earth, the expected recoil energy per scatter decreases, further reducing the stopping power. Additionally, the increased velocity dependence causes tracks deposit even more of their energy towards the ends of their tracks. The overburden scaling relationship between DM radius and the lower limit on cross section, when the radius exceeds

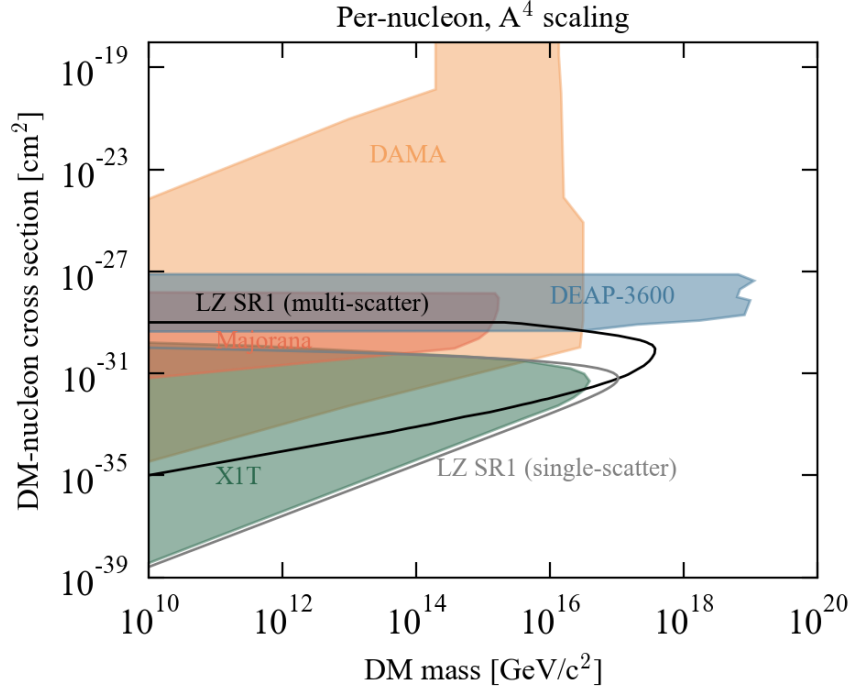


Figure 5.23: LZ SR1 MIMP search exclusion limits plotted alongside the results from competing analyses for A^4 SI scaling.

the 5 fm radius of the xenon, is approximately $\tilde{\sigma}_{LL} \propto R^{3.4}$.

5.9 Summary

The LZ SR1 MIMP search reports a background free, null result for dark matter scattering multiple times in the TPC. With A^4 scaling the experiment is world leading at $\sigma_{\chi p} = 10^{-30} \text{ cm}^2$, extending the high mass frontier to $3.80 \times 10^{17} \text{ GeV}/c^2$. The experiment is set up to further explore the high-mass frontier as exposure increases over its lifetime. With contact interactions, i.e. when the dark matter state is opaque to the SM nucleus, xenon no longer benefits from a higher atomic number, and therefore the increased flux of the LAr-based DEAP-3600[145] excluded higher masses. Contact interactions in the LZ SR1 MIMP search push lower in cross section, but into space which is already explored by the single scatter WIMP search.

The effect of dark matter form factor was explored, demonstrating the expected effect of pushing the MIMP limits higher in cross section space while maintaining the maximum mass. Future searches, if background is present, can exploit this form factor to place limits not only on cross section, but dark matter radius.

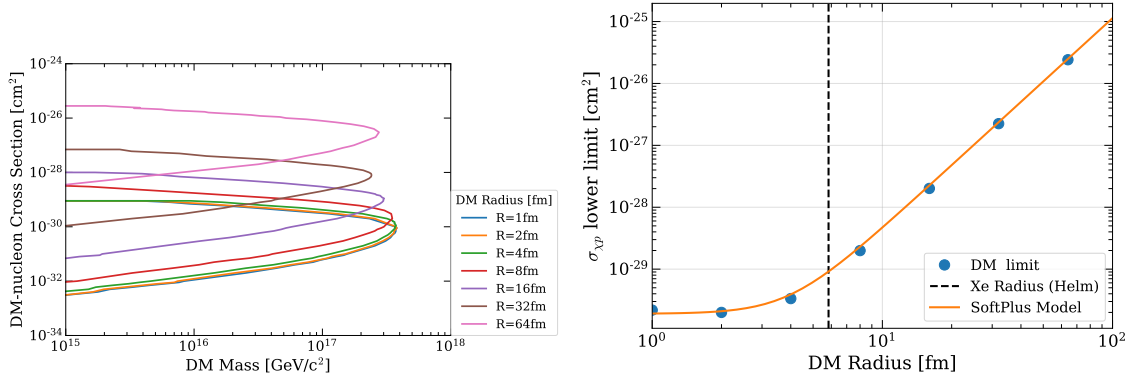


Figure 5.24: *Left*: Extended object dark matter search results. *Right*: Extended object dark matter overburden scaling for $M = 10^4 \text{ GeV}/c^2$ objects.

Extensions to the MIMP search will involve waveform-based analysis to mitigate the effect of pulse merging. Pulse classification is also uncharacterized at this point, i.e. at what point does a merger of several S1s cease to appear like an S1? It will also be beneficial to capitalize on the veto detectors. While the $\langle \frac{dE}{dx} \rangle$ in the outer detector is too small to be useful at small cross sections, the skin detector presents an opportunity to increase the effective surface area of the detector. While the skin veto is disabled, this search is agnostic to its signals. A future analysis could look for S1s which bookend the multiple scatters in the TPC, increasing the detection efficiency.

Chapter 6

Simulations

6.1 Introduction

Obtaining a result for the dark matter search requires understanding how the detector will respond to certain physics. As part of the simulations working group, I made significant contributions to the software stack in the design phase, and implemented several physics and detector response models. In this chapter I detail the design of these features. Further detail may be found in the LZ simulations paper[138] and SR1 result[53].

6.2 DMCalc

Purpose

DMCalc was build off of, and heavily modified, software developed by LZ simulations coordinator Quentin Riffard. It was created to make calculations of dark matter recoil rates. I was tasked with building off of this framework to make a general purpose, extensible physics model library which could be utilized throughout the simulations stack. Applications included making high quality plots, and creating the recoil spectra that LZLAMA and BACCARAT (below) sample from.

Design

DMCalc divides its calculations into the following concepts:

- **Spectra:** The velocity or energy distribution of incoming particles. For relativistic particles ($m \ll T$) the energy probability distribution is used, while for nonrelativistic ($m \gg T$) particles the velocity distribution is used. As of writing there are no models implemented with intermediate velocities ($m \sim T \leftrightarrow \beta > 0.01$), so the distinction is more accurately zero or nonzero mass. DMCalc takes these distributions and integrates

the differential cross section over the first moment of velocity:

$$\sigma_{\text{eff}} = \int \frac{d\sigma}{dE_R} v f(v) dv \quad (6.1)$$

In some models, this integral may be calculated analytically, and the generic numerical integral is skipped. This occurs in the dark matter case, as the *Spectrum* class is derived into the *Halo*. Typically in the literature the velocity dependence of dark matter exploits the isotropic $d\sigma/dE_R = (rm_\chi v^2/2)^{-1} \theta(rm_\chi v^2/2 - E_R)$, where $r = 4m_T m_\chi / (m_T + m_\chi)^2$ to factor the result into the “inverse velocity distribution” [44]:

$$\xi(E_R) = \int_{v_{\min}(E_R)}^{v_{\max}} \frac{f(v)}{v} dv . \quad (6.2)$$

As such, within DMCalc dark matter models redistribute the remaining factors to place all velocity dependence in $\xi(E_R)$. Neutrino models use the incident energy distribution instead.

- **Cross Sections:** Cross sections are the core of any physics model. DMCalc requires at minimum the function $\frac{d\sigma}{dE_R}(E_i, E_R)$, where E_R is the recoil energy and E_i is the incident particle energy, to be convolved with the incident energy spectrum. The interactions are defined in the laboratory frame with a target at rest. There is an implicit dependence on the target nucleus. For dark matter models, since the velocity dependence has been factored into the $\xi(E_R)$ function, the differential cross section must be independent of incident velocity. The elastic “halo cross section” is given by :

$$\left[\frac{d\sigma}{dE_R} \right]_{\text{halo}}(E_R) = \frac{2\sigma_{\chi T}}{rm_\chi} . \quad (6.3)$$

The $\sigma_{\chi T}$ term will differ between the spin-independent and spin-dependent models. In reality the recoil energy may not exceed the backscatter amount $rm_\chi v^2/2$, but this fact is folded into $v_{\min}(E_R)$ in the halo inverse velocity distribution. For neutrinos, the cross sections are calculated on a per-flavor basis. At this point in time, while scattering is not assumed to be isotropic in general, the angular dependence is not a component of the functions, as LZ and other TPCs lack directional sensitivity.

- **Target:** DMCalc, while used primarily by LZ, also aims to make comparisons between different target media. It is also necessary to calculate different scattering rates for effects like earth shielding [174]. Within DMCalc, cross sections are calculated based on the recoiling *Nucleus*, and any A, Z dependence will appear there. A detector consists of a *target* material, which is a mixture of several *nuclei*. A convenient interface allows targets to be added together and for the isotopic fractions to sum to 1, e.g. an unrealistic mixture $\text{Xe}_{\text{mix}} = 0.5\text{Xe}_{129} + 0.5\text{Xe}_{131}$. LZ uses natural xenon, but other targets are shown in Fig. 6.1




Figure 6.1: WIMP recoil rates generated with DMCalc. *Left*: WIMP dark matter differential nuclear recoil rates for different targets. *Right*: WIMP dark matter nuclear recoil rates for different model (SI, SD-p, and SD-n).

- **Rates**: This is the high level control class. The parameters of the physics models are placed in this class. For dark matter models, these parameters include the WIMP mass, WIMP-nucleon cross section, and coupling (spin-independent or spin-dependent). Targets and optional velocity distributions are specified here. Clients such as LZLAMA and BACCARAT interact through this class.

Halo Distributions

The dark matter velocity distribution is not known precisely. The standard halo model assumes a Maxwellian distribution in the galactic rest frame. Subcomponents of different distributions are conceivable, though. Notably, the results of the GAIA survey lead to the development of the SHM++[175]. For this reason I implemented several dark matter velocity distributions, along with a mechanism to combine them in a similar way to the Targets.

- **Gaussian ellipsoids**: This is the building blocks for the SHM++ model. There, a warm, slow, isotropic Gaussian distribution is combined with a cold, fast, eccentric Gaussian distribution. It is believed that this extra “sausage” is from a galaxy merger late in the development of the Milky way. In DMCalc, not only is it possible to form the SHM++, it is possible to add Gaussian subcomponents with any desired covariance matrix.
- **Debris flows**: The Milky way can tidally strip material from nearby satellite galaxies, forming flows of material which has yet to mix with the rest of the galaxy[185]. The




Figure 6.2: Velocity distributions for various dark matter halo models, produced in DM-Calcul. Shown are the SHM[44] (blue), SHM++[175](orange, jagged due to Monte-Carlo integration), debris[185] (green), stream[186](blue, line), and a custom Gaussian composite model (red, for demonstration purposes).

velocity distributions are characterized by the flow velocity v_{flow} :

$$f(v) = \begin{cases} \frac{1}{2} \frac{v}{v_{flow} v_e(t)}, & |v - v_e(t)| < v_{flow} \\ 0, & |v - v_e(t)| \geq v_{flow} \end{cases}. \quad (6.4)$$

- **Streams:** In a similar manner to debris flows, stellar streams [186] indicate that cold subhalos may travel through the main halo. These are modelled as monochromatic sources in the galaxy frame $f(v) = \delta(v - v_s)$.

Examples of these models are shown in Fig. 6.2. In addition to the probability distribution functions, the $\xi(E_R)$ functions must also be calculated.

The annual modulation due to the motion of the earth around the Sun is also modelled in DM-Calcul. One specifies the J2000 date and the velocity of the Earth in galactic coordinates v_E is calculated and used in the subsequent evaluation of the SHM velocity distribution. All dependence is rolled into v_E , and therefore any of the other velocity components benefit automatically from changing the day. Using a value of $d = 60.8$ gives the approximate mean value. The amplitude of the modulations for select SI-WIMP distributions is shown in Fig. 6.3




Figure 6.3: Relative modulation amplitudes for WIMP masses of 20 and 50 GeV/c^2 , as a function of recoil energy, calculated with DMCalc.

Bindings

To make DMCalc maximally useful, the C++ code base could be used in several ways. Every class was made serializable, such that a string could be input into a parser and the *Rate* object would be created automatically. The two main methods of supplying a configuration string were YAML[187] and via the command line. With the YAML interface additional instructions could be provided to make the driver program output the desired plots. The command line parser was somewhat more limited, as it was intended to be simple to access and extend. The internal arguments are completely encapsulated, making any change to DMCalc automatically propagate to its client programs.

The most flexible method of using DMCalc was via its python bindings, which I developed. Most of the C++ classes are accessible via the bindings, with names changed from the preferred C++ convention of camelCase to the python3 convention of snake_case.

DMCalc is compiled into a shared object `.so` library, which is then installed and linked against in the client software. The installation procedure also copies some database files which specify certain spectra. These are most critical for the Neutrino models, which calculate matter effects and have several components (e.g. pp, B8, diffuse supernovae, etc). As much of the code was made to be compile-time calculable, so that runtime calculations are as fast as possible. Compile time calculations also allow for static error analysis and unit testing, where certain computational results will trigger a compiler fault if they fail an

assertion.

Future Work

DMCalc is at the time of writing limited to the calculation of WIMP(and “blob”, see Chapter 5) dark matter and neutrino background models. Near term plans include the incorporation of sub-GeV and low-energy ER models. Neutrino magnetic moments and axions are of particular interest. Recently the Migdal effect[50] was added to the WIMP model by graduate student Andreas Biekert. This produces an ER from a low-energy NR, which requires additional convolutions over recoil energy. The intention for DMCalc was to make it simple to implement an equation and use it in simulations, and while that goal was met, some quality of life modifications will likely be made in the future to facilitate expanded use.

6.3 LZLAMA

Purpose

LZLAMA is the core of the LUX-ZEPLIN fast simulations chain. I contributed to it through consultation on the design, core code development, linkage to the DMCalc library, and module development. As with DMCalc, LZLAMA was based on a scaffolding created by Quentin Riffard for the LUX experiment. It transforms entries representing energy deposits into the LZ reduced-quantity files by simulated the detector effects in a parametric (top-down) fashion. This is different than the full chain simulations, which simulate the particle physics, photon propagation, and electronics response in a bottom-up fashion. The LZLAMA(fast) chain has the advantage of being computationally less expensive, and more straightforward to match simulation results to data. The software operates in two processing modes: *pdf-sampling*, where the recoil spectra are used to create the detector responses, and *file-processing*, where energy deposits are created using the Geant4-based particle tracker BACCARAT, and the detector response to these deposits are then simulated. Under *file-processing*, the extensive physics library of Geant4 can be leveraged without the bottlenecks involved in generating and analyzing simulated waveforms.

Design

LZLAMA takes in a configuration file which specifies the parameters of the simulation. The processing classes come in three categories: Inputs, Modules, and Writers, which are performed sequentially. Each detector effect generally becomes its own module, and these modules have certain parameters (e.g. trigger threshold) which can be tuned to match data. LZ is modelled as three separate detectors: the TPC, Skin, and OD. Each module is assigned to one or more of the detectors.

1. **Inputs:** The input classes generate the deposits. The modes *file-processing* and *pdf-sampling* control the manner in which these deposits are generated. With *pdf-sampling* enabled, LZLAMA uses DMCalc, or a hard coded histogram to generate the recoil energies. These deposits are generated uniformly throughout the cylindrical TPC volume:

$$R^2 \sim U(0, 72.8) \text{ cm} \quad (6.5)$$

$$Z \sim U(0, 146.1) \text{ cm} . \quad (6.6)$$

For certain generators the reverse field region below the cathode is included. Each generator also specifies the particle type, which is used in the NEST processing later on. The *file-processing* mode reads in BACCARAT output files, which contain the position, energy, and particle type information.

2. **Cut Modules:** The first line of modules culls the deposits which were not relevant for the performed tests. Maximum energy, time, and particle type are considered. Time cuts were implemented to assist the *file-processing* mode, where metastable, long-lifetime excited states can lead to decays seconds or even years after the primary interactions.
3. **Pulse Generation Modules:** Several modules take in the deposits and output pulses. Deposits are first *clustered*, where nearby vertices are merged together using the density-based clustering (DBSCAN[188]) algorithm. Vertices within $100\mu\text{m}$ and 10ns are merged into *clusters*. Merged pulses are then mapped with the LZElectric-Field Module, which associates with each cluster an S2 X,Y position, drift time, field magnitude, and diffusion. The mapped points are then processed using NEST[176, 85], which takes in the previous values and generates the *pulses*, i.e. S1s and S2s.
4. **Detector Triggering Modules:** Pulses are sorted according to start times, and then overlapping pulses are merged according to independent S1 and S2 pulse merging models. S1s are merged according to a stochastic model based on the ratio of pulse areas, and the time difference between them. Ratios closer to 1 result in lower merging probabilities, and closer pulses are merged with high probability. The model is based on $^{83\text{m}}\text{Kr}$ data, which has 9.4 and 32.1 keV internal conversion electrons[136], which occasionally merge into a single pulse. S2s are merged based on a drift time varying pulse widths. When the pulse boundaries overlap, they are merged based on the overlapping time and relative areas. This is tuned to match the DD calibration data. Merged S2s then have their positions perturbed using the *PositionResolution* Module. This simulates the effect of the position reconstruction algorithm (Mercury) and the transverse diffusion. This feature was implemented by myself and detailed further in Section 6.3.

The area and PMT coincidences of the merged pulses are then used to determine if they trigger the pulse finder and DAQ. S1s in LZ are identified based on a 3-fold coincident (though this is a configurable parameter). S2s have a minimum area of 20 phd necessary to trigger an event.

5. **Event Construction Modules:** In the event of a triggered event, the merged pulses are organized into an *event*. These events are then classified into different *interaction* types based on the number and time order of pulses. Interactions are one of the following: single scatter (1 S1, 1 S2), multiple scatter (1S1, 2+S2s), pileup scatter (2+ S1s, 2+ S2s), and other scatters (everything else). Based on the particular scatter, certain reconstructed RQs are calculated.
6. **Writers:** Here the events are written out into the desired .root format. The formats are:
 - **LZRQ:** a subset of the RQs in the same format as processed LZ data. The per-PMT information is not written by default.
 - **Flat Table:** A debugging format used to test the modules independently. It writes out all events, whether they triggered or not, and the format is simplified to map onto the internal LZLAMA data structures.
 - **Legacy:** A format made to be compatible with LUX events.

The written files also contain *MCTruth* information, which is the actual (i.e. not reconstructed) deposit information. Because of the multiple stages of merging, each pulse is associated with multiple vertices, and some vertices are not associated with any pulses.

Resolution

The imprecision of the position reconstruction in LZ was implemented into LZLAMA by myself using inputs obtained from data. Two independent models were used, resulting from the fundamentally different ways X,Y and Z components are estimated. For X,Y these are found with Mercury, which maximizes the likelihood of the S2 position given the observed PMT hit map. For Z this is found with the drift time and the known thermal electron velocity v_e .

Transverse

Transverse diffusion spreads the electrons in the radial and azimuthal directions according to a Gaussian $\sigma_T = \sqrt{2D_T t}$. This results in an intrinsic error on single electrons. Additionally there is an effect from the grid funneling, where the 5 mm pitch between the gate grid wires squeeze the electron trajectories inwards. Most importantly, the stochastic nature of the photon hit patterns determine the accuracy of the reconstruction. Points near the edge of

the TPC are affected by the finite extent of the top array and the reflections from nearby surfaces, reducing the accuracy. Then, the S2 position is given by $\hat{X} \sim N(X, \sigma_e(R, S2))$.

Two main scaling effects are present: reduction of the error on the estimator for the mean with increased samples, and the position dependence of the single electron resolution. The area scaling was confirmed with ^{83m}Kr data. While the total energy of these events are monoenergetic 41 keV electron recoils, between recombination fluctuations lead to a spread in S1 and S2 values. The calibration data, taken near the beginning of SR1, is shown in Fig. 6.4. For the resolution analysis, events with $S1 \in [125, 300]$ phd, $\log_1 0 S2 \in [4, 4.75]$, $\text{driftTime} \in [50, 950]\mu\text{s}$ were selected to encapsulate the ellipsoid. Additionally, events near the wall experience charge loss (QL), producing the ‘‘rain’’ below the main ellipsoid. Events with $\log_1 0 S2 \in [2.75, 4]$ are analyzed separately.

Mercury maximizes the likelihood of S2 positions based on ‘‘light response functions’’ [86]. In addition to the estimated (X, Y) of the S2 it also returns a covariance matrix Σ_{ij} . I take the elements of this matrix and rotate the matrix $R(\theta)\Sigma R^{-1}(\theta)$ such that the diagonal elements contain the $\sigma_R^2, R^2\sigma_\phi^2$ values. I use the charge loss events from ^{83m}Kr to test the area scaling relationship, as they all occur at approximately the same radius. The $1/\sqrt{S2_{\text{Top}}}$ scaling is confirmed in both the σ_R and $R\sigma_\phi$ components, as shown in Fig. 6.5.

The non-charge loss Krypton decays are also used to provide an estimate for the area scaling. To eliminate the effect of radial dependence events with reconstructed S2 radius > 71 cm were chosen. I found that the same scaling relationship extends to these events. QL events provide a third handle on the resolution near the wall. Those scatters were sorted into drift time bins of width $50 \mu\text{s}$, and the mean and standard deviation are calculated per bin. The residuals of each event are calculated relative to their respective bin, and the estimate for the error was evaluated as a function of S2 top area. This technique reveals a modification to the inverse square root scaling:

$$\sigma_R = \frac{\sigma_0}{\sqrt{S2_{\text{top}}}} + \sigma_1, \quad (6.7)$$

where the additional $\sigma_1 = 3.05$ mm term implies an intrinsic resolution that is not captured by the Mercury covariance matrix. This model is shown in Fig. 6.6. The two models coincide around 1000 phd. For simplicity the models are combined, with $\sigma_1 = 0$ but with the σ_0 scaled to pass through the mean of the charge loss data S2 values.

Radial dependence was found with the non-charge loss events, with results shown in Fig. 6.7. As before, Mercury covariance values were used. There is relatively little radial dependence over the fiducial volume $R < 68.8$ cm, but near the wall at 72.8 cm σ_r increases by a factor of approximately two, while the $R\sigma_\phi$ is consistent with a constant term. I model this effect with an exponentially falling term. The overall result which I implemented in LZLAMA is the following model:

$$\sigma_r = \frac{1}{\sqrt{S2_{\text{top}}}} [\sigma_0 + \sigma_1 \exp(\frac{(r - r_{\text{wall}})}{r_0})] \quad (6.8)$$

$$R\sigma_\phi = \frac{\sigma_0}{\sqrt{S2_{\text{top}}}}, \quad (6.9)$$




Figure 6.4: Krypton calibration single scatter S1-S2 distribution heatmap. This data follows a nominal drift time fiducial cut, along with basic anti-accidentals cuts, such as high single channel and pulse length (see Chapter 7 for more details). The main lobe of the distribution can be seen, entirely outside of the WIMP search ROI, along with a “rain” of events at lower S2s, a result of the charge loss regions near the wall.

where parameters for SR1 were estimated as $\sigma_0 = 17.5$ mm, $\sigma_1 = 10.48$ mm, $r_0 = 1.97$ mm. The smeared positions are clamped to $r \leq r_{\text{wall}} + 1$ cm.

Longitudinal

Longitudinal diffusion causes the electrons to be spread out by a Gaussian of standard deviation $\sigma_Z = \sqrt{2D_L t}$, and $\sigma_t = \sigma_Z/v_e$. Single electrons produce photons continuously as they traverse the extraction region, usually on the scale of $1 - 2\mu\text{s}$. The drift times for S2s are defined by area fraction times, but within the fast chain the details of pulse shape are simplified into start and end times. The start time is then defined as the time that the fastest electron reaches the liquid level. To find this value when $N_e > 1$, the cumulative distribution of electrons must be located. The probability that N *i.i.d* values with cdf $F_1(X)$ are greater than x is given by

$$F_N = 1 - (1 - F_1(x))^N . \quad (6.10)$$

In this case the cumulative distribution is provided by the error function, $F_1(t) = \frac{1}{2}(1 + \text{erf}(\frac{t}{\sigma_t\sqrt{2}}))$. To sample the fastest electron, a quantile is drawn from a uniform distribution $p \sim U(0, 1)$, and the inverse error function is used to obtain the appropriate value:

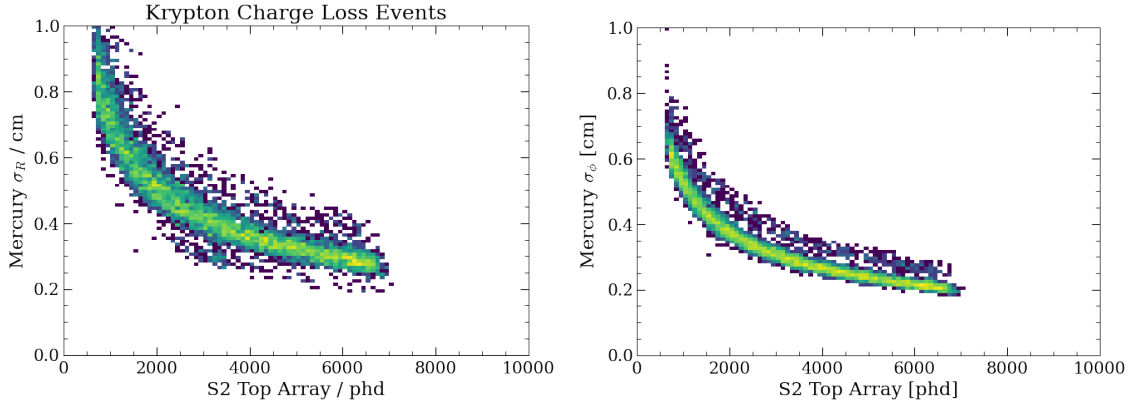


Figure 6.5: Histograms of radial uncertainty from Mercury for Krypton charge loss events as a function of S2 area in the top PMT array. S2 area scaling from charge loss events, as measured by the rotated Mercury covariance matrix.

$$t_{\min} = \sqrt{2}\sigma_t \text{erf}^{-1}\{2(1 - (1 - p)^{1/N_e}) - 1\} + \mu_t . \quad (6.11)$$

6.4 BACCARAT

Description

The full chain simulations for LZ is performed by the Geant4[139] based BACCARAT simulations package. Particles are tracked through the detector, where they generate optical photons, thermal electrons, and secondary interactions. I contributed the field models, described in detail in Chapter 4, to BACCARAT. Several field configurations were made available, with differing drift and extraction fields. The fields were utilized to calculate the recombination probabilities for the quanta. Thermal electrons were transported from the deposit sites to the liquid surface based on the drift maps. Fields within the extraction region produce S2 light based on the fields queried in that region. All effects are stored within the LZElectricField package.

PTFE Pockets

The PTFE panels covering the field cage of LZ's TPC are typically treated as if they were a single cylindrical component. In reality the cage is constructed by placing the field shaping rings between small interlocking PTFE sections. Under thermal contraction these sections fit together tightly, but small gaps exist in the vertical direction, as shown in Fig. 6.8. While still light-tight between the skin and TPC regions, this leads to a volume of Xenon which is




Figure 6.6: Independent S2 position resolution vs. S2 area scaling models. Charge loss events and ^{83m}Kr events with S2 radius > 71 cm are shown, estimating σ_R from the Mercury covariance. Charge loss (blue) and near-wall Krypton events (blue) follow the same trend. The uncertainty estimated from the charge loss reconstructed positions themselves (which are all within 3 mm of the wall, and therefore have a relatively precise “truth” location), shown in green, indicates a slightly different relationship, preferring a constant “systematic” uncertainty on top of the area-dependent relationship. The purple line indicates the implemented model, with a systematic term of zero but passing through the mean of the charge loss RMS data points.




Figure 6.7: The radial and azimuthal resolution as a function of radius within the TPC, as estimated through the rotated S2 position covariance matrix of ^{83m}Kr events. The azimuthal component does not display any dependence on radius, while the radial component shows an increase in uncertainty near the wall. Note that this shows the resolution corrected for S2 area, *i.e.* $\Upsilon = \sigma \times \sqrt{S_2}$

invisible to S2s. S1 light may reflect off the channels and lead to S1-only background (see Chapter 7).

I implemented these small pockets into the BACCARAT geometry. The generator which creates ^{210}Po decays was modified to place primaries along the new contours. LZLAMA had to be similarly modified so that it would not generate S2s from deposits in the pockets. The light collection model was simplified at this point, simply attenuating the quanta based on the distance from the wall:

$$\frac{N'_\gamma}{N_\gamma} = 1 - \frac{r - r_{\text{wall}}}{r_{\text{ring}} - r_{\text{wall}}} . \quad (6.12)$$

This new value of N'_γ propagates through to the binomial fluctuations which result from the light collection efficiency in the TPC. I simulated both the ^{206}Pb recoils and the β -decay spectra from the detector components near the wall. The charge loss model was used to find the S1-only spectrum from the decays on the regular wall. An S1-only spectrum which results from these simulations is shown in Fig. 6.9 I find that the pockets result in more S1-only background than the PTFE at $r = 72.8$ cm. With 30 mBq/m² of ^{210}Po plateout decays simulated, 31.8 mHz of S1-only is found from the normal wall with S1 < 100 phd, whereas the pockets result in 37.8 mHz of S1-only in the same region. These results depend on the light propagation in the pockets, and future work will simulate the effects of numerous reflections in this region.




Figure 6.8: *Left*: CAD rendering of the PTFE pockets. *Right*: ^{210}Po decays on the PTFE walls. The pocket profiles can be seen in the decay locations beyond the wall radius of 728 mm. The density is uniform on the surface, but appears less dense in the pockets due to the aspect ratio.




Figure 6.9: The S1-only spectrum which results from events near the wall, including the newly simulated PTFE pockets

Code Reviewer

I had the pleasure of serving as the code quality reviewer for LZ software. As part of this role, merge requests for simulations software were sent to my gitlab.com dashboard. I would issue recommendations based on the CPP core guidelines to ensure readability and performance were up to the quality demanded of LZ.

These reviews only existed for newly committed code. For old code, I judiciously utilized static analysis tools like Coverity¹ and clang-format² to enforce the core guidelines. In many cases I refactored old code within BACCARAT to be more compliant. The largest issues I observed were excessive function complexity, and shadowed variables. While it was difficult to use with the existing code base, I converted many “bare pointers” into “smart pointers” which obviate the need for manual deletion.

I implemented a feature into the gitlab repository to automatically run the memory leak checker valgrind³. Memory leaks occur when data is allocated by one section of a program, and the reference to that memory is discarded without deallocation. Depending on where these leaks occur, memory usage might grow over time until a memory error occurs. To detect these, I used valgrind-CI⁴ to render the valgrind outputs. Valgrind essentially runs the program on emulated hardware, intercepting its machine instructions to track the allocations

¹synopsis.com

²<https://clang.llvm.org/docs/ClangFormat.html>

³valgrind.org

⁴<https://pypi.org/project/ValgrindCI/>

and deallocations, which leads to slowdowns of a factor of ten or more. Because of this the tests are run in parallel to other continuous integrations, and allowed to fail. The rendered html is then placed in the artifacts folder where the developers can inspect them.

When the electric fields were incorporated into LZLAMA, I performed similar tests to inspect the impact. I used the valgrind output, alongside statistical profiling tools like GPerfTools⁵. The conclusion of these tests were that the electric field module was subdominant to the disk writes from the LZLAMA .root output. The call graph is shown in Fig. 6.10. At various times I used similar strategies in other codes which I have written, and to consult on shared analysis programs.

⁵<https://github.com/gperftools>

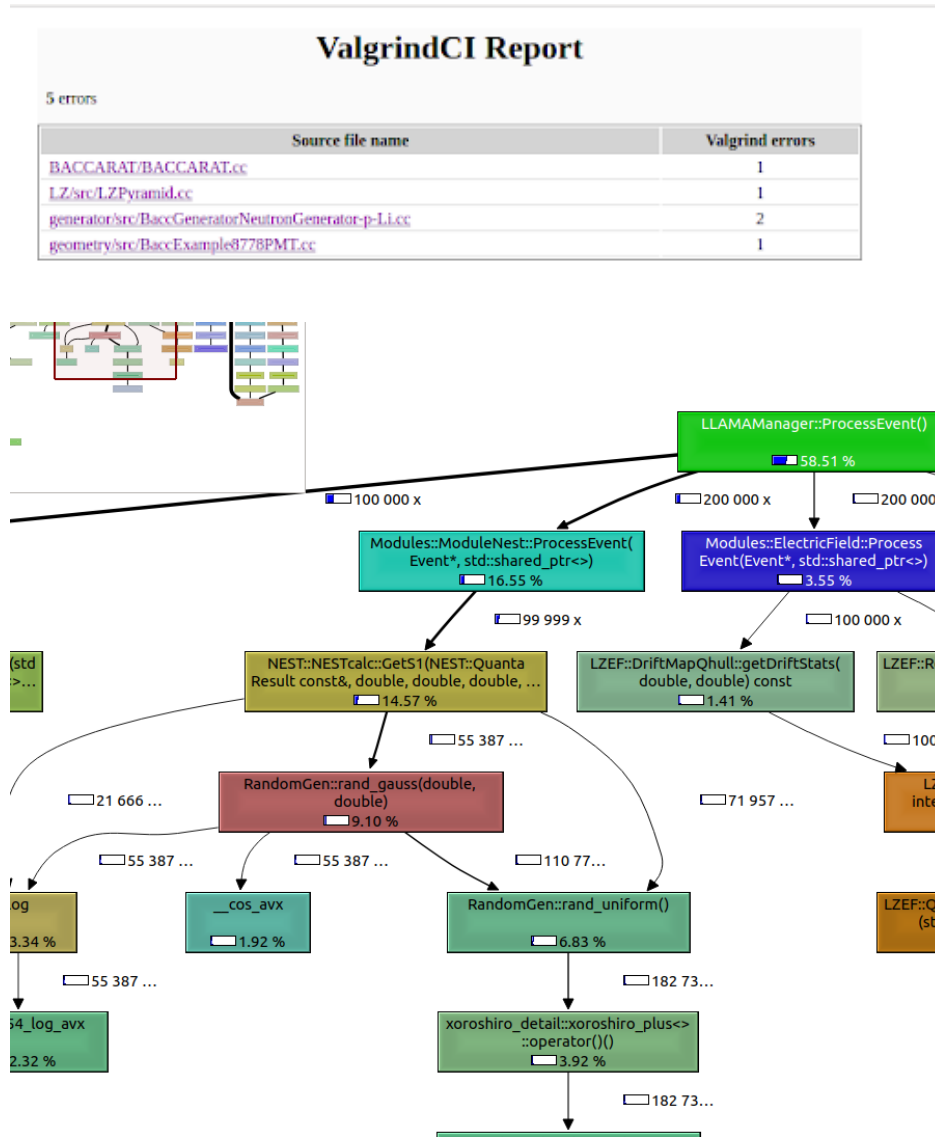


Figure 6.10: Code evaluation outputs from code profiling programs. Top: Valgrind-CI rendered for BACCARAT. The errors indicate memory leaks indentified by valgrind. Bottom: Callgrind output used to inform the optimizations for LZLAMA. Each colored node is a particular function call. The edges indicate that a function called another function, and the number indicate the number of calls which occurred. The percentages in each node represent the fraction of the parent node’s call time taken by the particular child node.

Chapter 7

Accidental Coincidence Backgrounds in LZ

7.1 Introduction

The primary LZ WIMP search analysis relies on identifying pairs of S1 and S2 pulses and analyzing joint distributions of these events. This is followed by an algorithm which identifies the category of event based on the number and order of types of pulses. When two energy deposits from different, unrelated sources occur in the same event window, this is intended to be sorted into the “pileup” event category by the interaction finder. These events are identified by having multiple S1s occurring before the multiple S2s. However, there are a variety of reasons why these pileup events could end up in the single scatter category, which would constitute a WIMP search background. The main concern is the situation in which an S1 with no associated S2 pairs up with an S2 with no associated S1 (*i.e.* that the prompt and proportional scintillation are not causally related). This event topology is known as an “accidental coincidence” event, or accidental for short. The pulses that exist outside of single scatter pairs are known as “isolated S1(2).” Isolated S2 pulses can arise from sources such as:

- **Near-liquid-surface events.** The S1 and S2 are close enough in time that the pulses appear merged together. Due to the properties of the pulse finder/classifier, this is typically identified as an S2.
- **Near-energy threshold events.** For the same energy, S1s are more likely to fluctuate below the three-fold PMT coincidence requirement than S2s are to fluctuate below the minimum number of electrons at the electroluminescence gap (5). Therefore a population of events with one- or two-fold PMT coincidence S1s paired with small S2s is expected.
- **Grid electron emission.** The cathode and gate grids will occasionally emit one or more electrons in a stochastic manner[189]. These will occasionally be large enough

to be classified as an S2. Photoionization from large S1s can also contribute to this, causing scatters in the reverse field region to be reconstructed in the TPC.

- **Electron/photon trains.** After a large S2 there is a long train of delayed electrons, which may persist over many event windows[190]. These electron / photon trains are generally easy to identify from true events, but nevertheless confuse a naïve analysis.
- **Above-anode gas events.** These events occur when an energy deposit takes place above the anode grid wires and below the top array PMT faces. The electrons experience highly distorted fields as they are drawn towards the anode wires. This results in S2s with distinct time profiles and top-bottom-asymmetry(TBA).

Some of the above sources have mitigation strategies, while others are more pernicious. Sources of isolated S1s include, but are not limited to:

- **Reverse-field region events.** Deposits below the cathode produce S1s, without the possibility of S2 production due to the field lines pointing upwards. These S1s typically can be identified by having top-bottom-asymmetry(TBA) near -1.
- **Charge loss events near the wall.** Events near the wall, where the electric field experiences significant fringing (see Chapter 4), can have missing or diminished S2 areas.
- **Electroluminescence.** The high field regions around the grid wires can generate electroluminescence in the liquid or gas[103, 109]. This can also occur near the gaps in the field shaping cage PTFE meant to accommodate the low-activity resistors.
- **Stinger S1s** following single electrons. Single electrons will frequently experience a varying electric field between the liquid level and the anode wires in locations where the gate and anode wires are anti-aligned. This occasionally results in a situation where many photons are generated at the start, and a short time later a small number of photons arrive. This can result in a separate identified “S1” pulse. These are referred to as *stingers* and are efficiently removed by requiring single electrons to be well separated from the S1 pulses. This phenomenon can also result in the SE pulse, or the SE+S1 being classified as an “Other” pulse, which is the software’s fall-through case.
- **Charge loss from impurities.** While unlikely, events with S2s near threshold may conceivably lose sufficient charge while drifting over long drift lengths. This effect is subdominant in SR1 due to the long electron lifetime (5-8 ms, where the full drift length is 0.955 ms).

Isolated S1 and isolated S2 sources have the potential to pile up, and if the false drift time is within the boundaries which define the fiducial volume, then this may appear as a WIMP search background. The naïve rate of accidental events, assuming zero correlation between pulse rates, is given by

$$R_{\text{acc}} = R_{S1}R_{S2}dT_{\text{max}} , \quad (7.1)$$

where R_{S1} is the rate of isolate S1(2) pulses and dT_{max} is the length of the fiducial drift time window which would result in a single scatter. Since the S1s and S2s are uncorrelated, they will not lie only in the nuclear recoil band in S1 vs. $\log_{10}(S2)$ space. The rate R_{acc} can be reduced by two main strategies. The first is selecting against the individual isolated pulses using pulse shape characteristics. The second is to target the joint distributions of S1 and S2s.

The primary tasks which I contributed to for the accidental coincidence background analysis are 1) simulations of isolated S1 sources, 2) estimation of isolated S1 rate in SR1, 3) development of selection criterion targeting accidental events, and 4) estimation of the accidental probability distribution function for SR1. Additional considerations were also calculated, such as isolated S1 pulses appearing in otherwise valid single scatters, reducing the effective acceptance of nuclear recoils. I show a predicted rate of isolated S1s of 0.5 – 1.0 Hz using the SR1 grid conditions in specialized datasets after applying the selection criteria, and an estimated rate of $0.653 \pm .008$ Hz using SR1 data itself. The selection criteria described in the following sections were able to bring down the isolated pulse rates by a factor of between 3 and 10, varying depending on the time period of the data. Combined with isolated S2 pulses, and unphysical drift time events in SR1, an overall accidental coincidence rate of 1.2 ± 0.3 events were anticipated for the SR1 exposure. Within the NR band a mere $0.18 \pm .04$ events are expected. This model was incorporated into the SR1 WIMP search PLR result[53].

7.2 Sources of Isolated Pulses

Dark Count Pileup

The original explanation of isolated S1s, considered from early on in the development of TPCs, is dark rate pileup. Each PMT has a characteristic “dark current” [108], which are pulses resulting from stochastic electron emission from the photocathode, without any impinging photon. The pure dark count pileup rate is the probability of finding N_c PMTs firing within a coincidence time window T_c :

$$R_{\text{pileup}} = \text{Bin}(N_c; rT_c, N_{PMT})/T_c , \quad (7.2)$$

where $\text{Bin}(k,p,n) = \binom{n}{k} p^k (1-p)^{n-k}$ is the binomial distribution, N_c is the minimum PMT coincidence requirement, r is the rate per tube, T_c is the S1 coincidence window, and N_{PMT} is the number of PMTs.

The single photoelectron (SPE) rate is estimated from the data. The Hamamatsu 11410-20 photomultiplier tubes [108] are specified to have 10 nA of dark current at room temperature, which with a single photoelectron gain of 5×10^6 , translates to approximately 12.5 kHz of dark pulses. However, dark pulses are a result of thermionic emission, and are heavily temperature dependent [115], thereby being suppressed at LXe temperature. Additionally,

there exists some evidence for PTFE florescence [191], suggesting that the SPE rate should rise following large S2s. This effect was seen in commissioning data, and is particularly strong during high activity calibrations. The data quality selection criteria described in Section 7.3 will have an impact, biasing the SPE towards the values which pass those cuts. I collect the SPE rate from the pre-S1 region of single scatter events in SR1 passing the electron/photon train holdoff. This is estimated on a per-event basis by the number of SPEs divided by the time difference between the detected S1 and the beginning of the event window(at -2 ms). A histogram of these rates is shown below in Fig 7.1, along with the calculated pileup rates for 3-, 4-, and 5-fold coincidence requirements. A lower coincidence requirement results in a lower energy threshold for the WIMP search, at the cost of higher pileup rates, resulting in more accidental coincidences and single scatter to multiple scatter contamination. The expectation value for SPE rate across the entire TPC is 20.2 kHz, or 41 Hz per tube.

There is no definitive coincidence window for S1s in LZ. Rather, the requirement is set implicitly by the width of the difference-of-Gaussians (DoG) pulsefinder. However, an effective value of 150 ns can approximate the requirement. Evaluating for $N_{PMT} = 494$, $N_c = 3$, $T_c = 150$ ns, $r = 41$ Hz into Eq. 7.2, one obtains an SPE pileup rate of 30.6 mHz. Because the pileup rate is highly nonlinear, I also calculate the expectation value. The sustained rate cut discussed in Sec. 7.3 establishes an effective upper bound of 40 kHz, which I used as the upper bound for the integral. This weighted average comes out to 45.6 mHz of 3-fold coincidences due to the dark rate.

Reverse Field Region

The reverse field region (RFR) of LZ, below the cathode grid, directs electrons downwards, away from the liquid surface. This eliminates the possibility of S2 production, resulting in a population of isolated S1s with identical spectrum to the fiducialized LZ background (following TBA corrections). This is an unavoidable feature of the design of the TPC, but only really becomes an issue at low S1 areas. By virtue of being so low in the detector, these events can typically be removed quite effectively by the top-bottom-asymmetry(TBA) selection criteria, discussed in Section 7.3. Only at S1 areas below approximately 40 phd (see Fig. 7.2) do the fluctuations of the photon partitioning become significant enough to see RFR S1s appear to come from inside the forward field region (FFR) of the TPC.

It is challenging to specifically calibrate this region. Event sources which are intended to be uniform are known to depend on the details of the LXe flow, which is not well known in the RFR. While I map out the isolated S1 TBA below, from which confidence intervals of TPC height can be constructed, obtaining the low-S1 contribution is a challenge. To obtain an estimate for this rate, I use the cathode rate in SR1 as a proxy. Selecting SR1 events with $S2 \in [600, 31000]$ phd and using the same accidentals-targeting criteria as discussed in Sec. 7.3, the rate of events from near the cathode may be estimated. An additional section of wall, equal to half the height of the RFR (13.7 cm ~ 89 μ s) was used to estimate the volumetric rate of WS ROI events. The rate calculated from this region can be doubled to provide a proxy for the RFR contribution, as this value is interpreted as half the activity of the cathode. Assuming that half the activity of the cathode is lost to the RFR, and that




Figure 7.1: A histogram of total SPE rate preceding S1s in single scatters passing the livetime-impacting data quality cuts (black line) (e.g. e-train / muon vetoes, see Sec. 7.3, Ref. [62]), with dark pileup rates indicated. The linear scale on the left indicates the probability density of SPE rates for event windows, and the logarithmic scale on the right indicates the pileup rate. The pileup rate (blue, orange, and green) grows sharply with the minimum coincidence (approximately a factor of 1000 for each additional PMT), and slightly more slowly as a function of the SPE rate itself (approximately two orders of magnitude between 10 and 100 kHz).

the entire bottom grid is identical in activity to the cathode, I take the correction factor to be 3. Using this, I find the *a priori* data driven estimate for the RFR contribution to be 0.128 mHz. The cathode events, along with TBA confidence intervals, are shown in Fig. 7.2.

Field Fringing

As discussed in Chapter 4, the electric field near the PTFE walls is heavily distorted, resulting in regions of charge loss, which may either be partial or complete. In the case of complete charge loss, this results in an isolated S1 pulse, which can go on to form an accidental coincidence. In some ways these events are similar to the RFR events, but are slightly more challenging due to their TBAs aligning with physical TPC heights.

I estimate this rate using a combination of simulations and a data-driven approach. The wall events were simulated with ^{210}Po recoils initiated on the walls of the PTFE, using a custom generator written by graduate student Eric Morrison. The 103 keV ^{206}Pb recoils were paired with the 5.4 MeV α -recoils. Depth profiles within the PTFE were also implemented, with the attenuation of the ^{206}Pb nuclei calculated. I made the additional contribution of modelling the small gaps between the PTFE panels, as discussed in Chapter 6. Generating




Figure 7.2: A 2D histogram of top-bottom asymmetry and S2 area of cathode events within the WS ROI. Events are taken between 900 and 955 μs . The binomial confidence interval is indicated with the red curves. As expected from binomial fluctuations, the TBA band widens as $S1 \rightarrow 0$.

2.3M events through *BACCARAT*, the spectrum of ^{206}Pb recoils was extracted and normalized to the rate to the number of 5.6 MeV ^{210}Po α recoils. I find that ratio to be 0.325 $S1 < 100$ phd events for every detected 5.4 MeV α -recoil.

The next step is to estimate the volume of charge loss zones near the PTFE wall. To do this, I examine a dataset of large S1 events selected from SR1, which contains the α recoils. These events are shown in Fig 7.3. Due to photoionization on the gate, this dataset actually contains the complete charge loss events. I select the charge loss as all events with $S2 < 1800$ phd and drift time $< 6 \mu\text{s}$, which is indicated in the plot. Each α peak was selected from the TBA-corrected S1 areas. Events with S1 TBA between -0.7 and -0.15 , and drift time $< 951 \mu\text{s}$ were selected in order to not be close to the grids. Because the distribution of ^{222}Rn and ^{218}Po events is not perfectly uniform in the detector due to the low-flow region in the middle of the detector, only events with $S2_R > 60$ cm were chosen. The 5.59 MeV ^{222}Rn alpha was selected as events with corrected S1 area $cS1 \in [36100, 39000]$ phd, and the ^{218}Po peak was chosen as $cS1 \in [39000, 41000]$ phd. “Good” S2s were selected as those between $\log_{10}(S2/\text{phd}) \in [5.2, 5.6]$. The complete charge loss fractions were found to be $r_{Rn} = 0.0134 \pm 0.0004$ and $r_{Po} = 0.0127 \pm 0.0004$. These are largely consistent with the ratio predicted from simulations. Assuming 50% total loss of events between 725 mm and 728 mm, a loss ratio of $r_{sim} = 0.0128$ is expected. In physical terms this is 3.1 liters over which charge is not collected.




Figure 7.3: A 2D histogram of S1 and TBA-corrected S1 areas from an SR1 α skim. Events are taken between S1 TBA -0.7 and -0.15 . The red line indicates the approximate dividing line between partial and complete charge loss events. At these energies, the complete charge loss events still result in S2s due to photoionization on the grids.

Examining SR1 data allows one to make a prediction for the fringe background. Single scatters are selected from SR1 data with $S1 < 100$ phd. A preliminary anti-accidentals cut is applied to remove accidental coincidences unlikely to be from this fringing effect, namely the high single channel (fraction of detected photons in a single PMT), anti-stinger (proximity to single electron pulses), and area fraction time $(95-5) < 350$ ns (eliminates pulses inconsistent with the xenon triplet lifetime). Identical top-bottom asymmetry requirements (between -0.7 and -0.15) and $S2_R > 60$ cm selections are applied to keep the ratios the same as the α s in the previous step. The rate in this pseudo-ROI region was found to be 2.3 mHz. When applying the r_{Po} ratio, this predicts a rate of 0.0605 mHz of isolated S1 background from field fringing effects. The spectrum of these events are shown in Fig. 7.4.

A similar method was used to predict just the ^{206}Pb nuclear recoil contribution. The ≈ 5.4 MeV ^{210}Po rate was estimated using 11.1 ms event window, randomly triggered datasets taken prior to SR1 (more details in Chapter 4) to be 0.54 Hz. This was found by applying a cut of $cS1 \in [34000, 36100]$ on specifically charge loss events ($S2 < 1800$ phd). Due to the concentration of this source on the wall due to plate out, it dominates the nearby Radon peak. Using the prediction from the simulations for the lead recoils in the ROI, I obtain a prediction of 0.176 mHz, dominating the S1-only rate predicted by the SR1 single scatters.




Figure 7.4: The field fringing background spectrum as predicted by SR1 single scatters with $S2_R > 60$ cm. The rate is scaled by the loss fraction inferred by the ^{218}Po α -charge loss S1s observed in the long random data ($r_{\text{Po}} = 0.0128$).

Source	Rate [mHz]
Dark Pulse Pileup	45.6
Reverse Field Region	0.128
Field Fringing	0.176

Table 7.1: Summary of the contributions to Isolated S1 pulses in LZ.

Summary

These predicted contributions which are on the scale of mHz are insufficient to explain the entirety of the $\mathcal{O}(1)$ Hz isolated S1 background observed in SR1. The remaining sources, which dominates all others, are likely due to the grids themselves. Electroluminescence in LXe can start at 412 kV/cm[103], which is achievable in the vicinity of the grids. Activity on the anode grid can also contribute, and likely dominates due to deployment in gas. This hypothesis is supported by the relative independence of isolated S1 rate on drift field and large dependence on the anode voltage, discussed in Section 7.6.

7.3 Analysis of SR1 Data

Rate estimation

A short recap of a typical event topology within the WIMP search ROI:

- 2 ms of pre-trigger window. A typical event has the S1 contained within this region.
- A trigger at $t = 0$. The WIMP search settings are tuned such that this should only occur for S2s, or more rarely extremely large S1s from MeV scale α decays. Single electrons can also initiate a trigger, which is not usually a problem, only resulting in the S2 pulse appearing to the right, inside the post-trigger window. There is a stochastic single electron background in LZ, in addition to occasional photoionization electrons from the grids, induced by large S1s.
- 2.5 ms of post-trigger window. This region is meant to capture additional S2s from multiple scatters, as well as to encapsulate a large part of the electron train before the next event can come online. Large S2s typically have electron trains spanning multiple event windows, though.

The isolated S1 rates were estimated with a particular mode known as “long random.” Here, instead of the normal event window of 4.5 ms length, an 11.1 ms event window was utilized. This affords a $> 90\%$ livetime (integrated duration of event windows divided by time measured by a clock on the wall), with multiple drift lengths in each event window. Correlations between pulses can be investigated, which presents an advantage for developing selection criteria.

Within the 11.1 ms window, S1 pulses which are 1 full drift length ($955 \mu\text{s}$) from any S2 are considered. Additionally, the first and last full drift length from the boundaries are removed from consideration in order to both be conservative and to maintain the necessary context for future analysis. It is assumed that S1s would not be observed if they occur coincident in time with single-electrons or Other pulses, so these pulse durations were subtracted from the livetime. If an S2 was observed, no further pulses were considered in the rest of the event window, as it is assumed to be biased. All pulses skimmed from this dataset have passed through the data quality selections, such as the e-train veto. Isolated S1s are required to be *prominent* (either the largest S1, or a sufficient height/width ratio), and to pass the OD and Skin vetoes. Because of the small S2 grid emission background, the S2s were divided into two categories, “noise” and “regular,” based in a cut of $S2_{\text{raw}} > 600$ phd, which corresponds to the SR1 analysis threshold. “Noise” S2s nearby an S1 do not have an effect on the “isolation” and are treated identically to single electrons.

The clipping of the edges of the windows limits the exposure for a fixed data taking time, but it guarantees that the context surrounding an S1 can be analyzed. I count the number and total area of all the categories of pulses in the preceding window, as well as the drift window looking forwards and backwards. This enables analysis of the correlation between features like SPE rate and isolated S1 rate.




Figure 7.5: The isolated S1 spectrum obtained from long random data. The HSC cut removes a broad band across the ROI, while the stinger cut targets primarily the iS1s near threshold. Other data quality cuts eliminate small isolated pulses individually but cumulatively lead to a more manageable isolated S1 rate.

The spectrum of isolated S1 pulses, and the effect of the various isolated S1 (iS1) cuts, is shown in Fig. 7.5. The distribution is strongly peaked towards the threshold, and the rate within the SR1 ROI of $S1c \in [3, 80]$ is 0.76 Hz. Extending the lower boundary, while keeping the 3-fold coincidence nets a rate, after all cuts, of 1.04 Hz. The effect of changing the threshold is illustrated in Fig. 7.6, which shows the integrated iS1 rate as a function of coincidence requirement and minimum S1 area.

These data were taken over a 24 hour period following grid biasing, and are therefore slightly erratic. This can be seen in Fig. 7.7. In order to test the time dependence of isolated S1s within the SR1 dataset, “pre-drift” regions were explored. These examined pulses more than one drift length away from the S2 trigger, and therefore each trigger added approximately 1 ms of exposure. While this dataset does have some bias, it does provide a high precision real time monitoring of the rate. This is shown alongside the Pre-SR1 dataset in Fig. 7.7, along with an exponential fit to the data, which yields a time constant of 790 days. The error bars are determined by Poisson fluctuations.

Selection Criteria

1. **E-train and muon veto:** after large S2s a series of electrons arrives, typically spanning several event windows. A background of single photons is also apparent, which extends further than the single electron background. The veto works by calculating a hold off region which dynamically varies based on the *Progenitor* S2 area. The length was tuned by members of the grids team to decay down to the quiescent background




Figure 7.6: Isolated S1 rate as a function of the lower bound in *Top*: S1 area, *Bottom*: coincidence threshold, with stacked analysis selection criteria. While raising the threshold does decrease the isolated S1 rate, the rate of change makes it a costly choice for the WIMP search. This is in stark contrast to the pileup predictions, which indicated that even a change of $N = 3 \rightarrow N = 4$ would almost entirely eliminate the isolated S1s.

of 40 Hz single photoelectrons, which takes times on order of tens of milliseconds. The muon veto works in a similar fashion, but sets a fixed holdoff time of 20 s after a muon is detected. The muons are determined by a skin+OD coincidence.

2. **Hot spot exclusions/High S1 Rates**: automatically detected regions of time with high S2 rates. A separate cut was tuned for regions of high S1 rates.
3. **Bad buffer**: event buffers are filled on a non-synchronized per-channel basis. Large events may not have all channels become available at the same time for the next trigger. This cut removes events where all of the buffers are not available at the start of the event window.
4. **Excess Area**: catches issues where pulses are misclassified. In a perfect event very little light is detected before the main S2 outside of the main S1. This cut removes events where too much area is found in pulses before the S1 and between the S1 and S2.

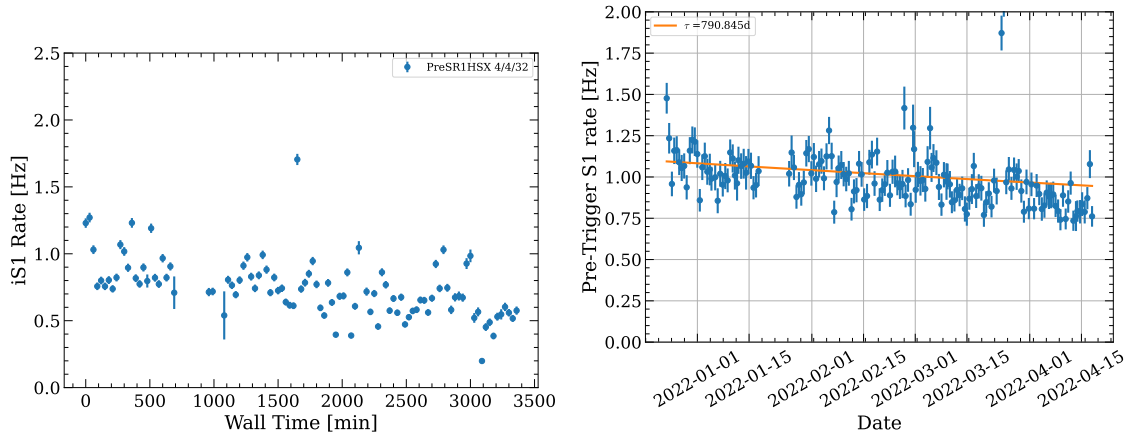


Figure 7.7: The isolated S1 rate plotted over long timescales. *Left*: The evolution of the isolated S1 rate in the Pre-SR1 dataset. Short timescale spikes are seen superimposed over a longer timescale downwards trend. *Right*: The time evolution of the SR1 pre-drift single scatter iS1 rate. The iS1 rate is following the application of the stinger, HSC, and pulse shape cuts.

5. **Sustained Rate**: catches events where the electron/photon (e^-/ph) trains did not perfectly estimate the decay. If an event has greater than 40 Hz of single photon rate in any 0.5 ms window, the event is rejected.
6. **OD Burst**: rejects events with a large amount of electronics noise.
7. **WIMP Search ROI**: Single Scatter, S1 and S2 Area, Fiducial Volume. These are the core cuts for the WS, selecting events that are WIMP-like. Internally these are divided up into separate cuts but are alike in purpose.
8. **Detector Vetoes**: OD and Skin detectors detect neutrons and Compton scatters. Coincidences between the Skin and OD detectors and the S1 will disqualify an event. The area thresholds and coincidence time window are tuned to remove backgrounds without excessive loss of livetime.
9. **S1 Prominence**: Occasional double S1 pulses are found in the pre-trigger region. Typically the smaller S1 is not classified as “prominent.” This can indicate a potential pileup event. The S1 prominence cut removes events with S1 pulses separated 750 ns in time from the prominent S1.
10. **Stinger**: Due to the nonuniform in the extraction region electric field resulting from nonaligned gate and anode grids, single electrons frequently produce electroluminescence in an initial burst, followed by a small number of photons separated by a μs or so. This cut removes S1s which occur within 2 μs of a single electron, or an “other” pulse which is larger than the S1 itself.

11. **High single channel (HSC)**: Removes events which contain too much light in a single PMT. These pulses are more likely to be the result of electroluminescence from the electrode grids.
12. **S1 Pulse Shape**: Discriminates between isolated S1s and scintillation S1, removing the isolated S1s. The time profile scintillation light is the combination of the exponential decay of the excimers, the reflection off of the PTFE walls, Rayleigh scattering, and the PMT and electronics chain impulse response. Isolated S1s result from a variety of phenomenon that generally result in a more stochastic burst of photons, generating *e.g.* longer pulses with larger first central moments. Features such as this can be exploited to discriminate against the isolated S1s.
13. **S1 Photon Timing**: Similar to the S1 Pulse shape cuts, but is based on a novel RQ calculated from the timing distribution of photon hits.
14. **S1 TBA vs. Drift Time**: Removes events based on the correlation between top-bottom-asymmetry of the S1 pulse and the height in the detector. Events which occur closer to one PMT array or the other tend to collect more light in the nearby array.
15. **S2 Width vs. Drift Time**: similar to the S1 TBA vs. drift time, diffusion of the electrons clouds as they drift towards the liquid surface lead to a correlation which can be exploited to remove accidental coincidences.
16. **S2 Shape**: narrow, early peak, rise time. These separate requirements target near-liquid-surface S2s which can be merged with S1s. These pulses have on average a much more steep rise time, along with an relatively higher prompt fraction compared to S2s which occur within the fiducial volume.
17. **S2 XY quality**: Poor χ^2 returned from the likelihood minimization program Mercury[86] are indicative of grid noise.
18. **S2 TBA**: above-anode S2s constitute a population of isolated S2s, contributing to accidental rates. This cut removes S2s with $TBA > 0.6$.

High-Single Channel Pulses

Background

It was observed in 11.1 ms event window, randomly triggered data that isolated S1 pulses tended to show an increased amount of light in one PMT, relative to low-energy calibration sources such as DD neutrons and CH_3T injection. S1 scintillation photons are generated far from the arrays, and have the opportunity to reflect off the PTFE walls and liquid surface before reaching a photocathode face. As such, most position information is almost completely lost, and the light is evenly distributed across the PMTs in a particular array (though the light is not equally partitioned between the top and bottom arrays themselves). Backgrounds which cause these pulses include Cerenkov photons within the quartz windows

of the PMT, and electroluminescence on the anode grid, which occurs close enough to the top array to have several photons impinge on a single tube.

Cut Development

Since there are 494 TPC PMTs, it is unlikely for photons to pile up in a particular tube until the pulse area becomes quite large. In statistics, this is known as the "birthday problem," *i.e.* the probability that among a group of N people, any two have the same birthday. The coincidence probability is given by[192]:

$$\begin{aligned} P(x, k, N) &= (\text{No. of ways to select } x \text{ PMTs}) \\ &\times (\text{No. of ways to assign } x \text{ photons to } k \text{ PMTs}) \\ &/ (\text{No. of possibilities for } N \text{ photons}) \\ &= \frac{\binom{N}{x} x! S(K, x)}{N^K} = \frac{(N)_x S(K, x)}{N^K}, \end{aligned} \quad (7.3)$$

where $S(k, x)$ is the Stirling number of the second time and $(N)_x = N \cdot (N-1) \cdot \dots \cdot (N-x+1)$ is the falling factorial. The probability for the maximum number of photons per PMT x is slightly simpler:

$$P(x, k, N) = 1 - \left[1 - \binom{k}{x} \left(\frac{1}{N} \right)^x \left(1 - \frac{1}{N} \right)^{k-x} \right]^N. \quad (7.4)$$

Evaluating for $N = 494$, $x = 3$, and $k = 10$, this evaluates to $P(3, 10, 494) = 4.8 \times 10^{-4}$. As such, unless a particular energy deposit is in close proximity to a PMT (which breaks the assumption that all tubes are equally likely to receive a photon), it is unlikely to produce pulses with large fractions of their pulse area in one tube. The variable in the analysis is referred to as the MaxChannelArea.

A selection criterion based on the Max Channel Area was written. Since the fluctuations are highly heteroskedatic, it was easier to work within the MaxChannelArea, $S1$ space, rather than with the ratio of the two. The decision boundary was chosen as:

$$\text{MaxChannelArea} < \max\{1 + S1_{\text{raw}}^{0.4}, 0.05 \times S1_{\text{raw}}\} \times \left(1 + 0.6 \times \exp\left[\frac{dT - dT_{\text{max}}}{30 \mu s}\right] \right) \quad (7.5)$$

The $1 + S1_{\text{raw}}^{0.4}$ power law expression encodes the low-area dependence of the max channel area on $S1$ area. Above approximately 40 phd, the fixed ratio of 5% was found to fit ^{83m}Kr contours better than the power law. The drift time independent portion of the criterion, $\text{MaxChannelArea} < S1^{0.4} + 1$, was developed based off of rejecting isolated $S1$ pulses from randomly triggered events while maintaining good acceptance(99.3% of deuterium-deuterium fusion neutrons) of low energy calibration events, and it shown in Fig. 7.8. It is clear that two lobes exist, and the HSC cut targets the upper lobe, which extends in some cases to large $S1$ areas. The boundary was chosen to allow through events at the intersection of the two lobes, and was chosen with the help of low energy calibration sources.




Figure 7.8: A 2D histogram of the maximum channel area vs. S1 area distribution of isolated S1 pulses, taken in the “long random” mode. The High Single Channel cut used in the first science run is indicated. An obvious two-band structure is observed.

The DD neutron generator creates 2.45 MeV neutrons [193], which are aimed through a conduit at the top of the detector. Using the trigger input into the DD generator, applying the NR band, and selecting the XY path of the beam, a very clean set of DD events may be curated. Tritium β decays, with an endpoint of 18.6 keV, from injected CH_3T gas, provide low energy electron recoils uniformly throughout the TPC. An advantage of the tritium calibration for this purpose was the relatively smaller event rate. Due to PTFE fluorescence, high event rate time periods, such as those during the DD calibration, have an artificially high isolated S1 rate. As a result, obtaining a clean “real” dataset to compare against is challenging for NRs. However, the tritium and DD bands do not differ significantly, as indicated in Fig. 7.9.

The drift time scaling expression similarly fits ^{83m}Kr calibration data, and encodes the moderate dependence of MaxChannelArea on drift time. As expected from a naïve solid angle calculation, deposits near the bottom of the detector, near the bottom array, observe a larger max channel area on average. This effect is more moderate near the top array due to the different collection efficiencies. These dependencies can be seen in the max channel ratios shown in Fig. 7.10. The events were selected from calibration data based on the following criteria:

- Single Scatters: *i.e.* an event containing 1 prominent S1 followed by 1 prominent S2.
- Passes E-Train Veto, OD and Skin Vetoes
- $S1 \in [120, 300]$ phd.

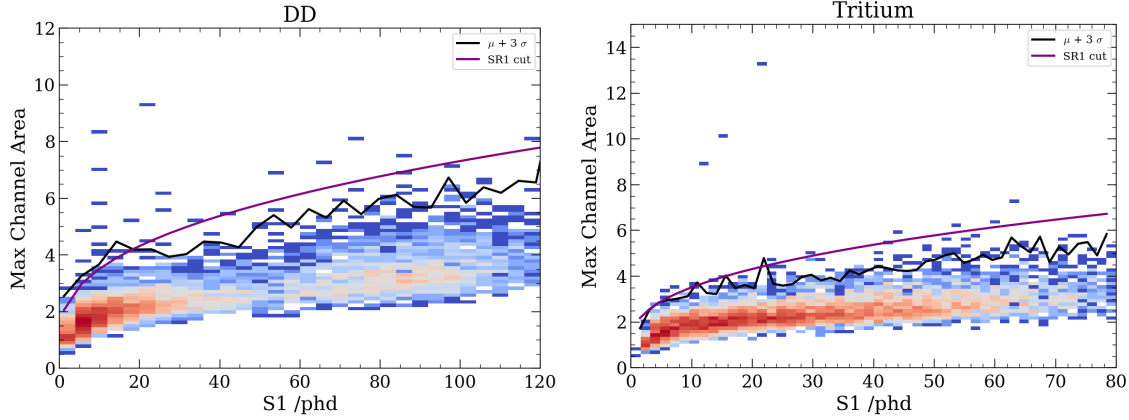


Figure 7.9: The drift time independent HSC cut development. Curves are shown for $\text{MaxChannelArea} < S1^{0.4} + 1$. DD calibration reaches lower in energy, but is located towards the top of the detector, leading to a potential bias. The tritium data subtends the entire TPC volume, allowing for a cross check of this effect. 3σ quantiles are shown, indicating the high signal acceptance of these cuts.

- $S2 \in [12000, 50000]$ phd.
- $S2$ radius < 70 cm.
- Drift Time $\in [60, 955]$ μs .

DPE Effect

The high single channel pulses are believed to be distinct from S1 scintillation light. Resulting from either electroluminescence or Cerenkov light, or a combination thereof, these pulses exhibit different characteristics. Most importantly, these pulses have a suppressed double-photoelectron (DPE effect[194]). As shown in Fig. 7.11, single scatter S1s which fail the HSC cut have a different distribution of per-channel areas than those which pass the HSC cut (i.e. real S1s).

The plot is a histogram of nonzero pulse areas for each PMT for the two pulse categories, subject to $S1_{\text{raw}} < 10$ phd in order to eliminate pileups. The events which pass the HSC cut are additionally required to have drift times within the TPC, i.e. less than $951 \mu s$. For the HSC pulses, the channel with the largest area is removed from the histogram. A double-Gaussian function is fit to the data, and the ratio of areas is computed.

Two possibilities were considered: a constrained and unconstrained fit. The constraint concerned whether the second peak's mean and variation were fixed to twice that of the SPE peak values. For an unconstrained fit, I find that the non-HSC (i.e. true scintillation pulses which pass the cut) pulses have a DPE fraction of 0.3 ± 0.04 , whereas the HSC pulses have




Figure 7.10: Maximum channel area ratio vs. drift time histogram for ^{83m}Kr calibration, with the High Single Channel selection criterion indicated in dashed red. The drift time dependence is evident at larger drift times / near the cathode, and the implemented cut is shown in red.

a DPE fraction of 1.3 ± 0.5 . This unconstrained fit is what is shown in the figure, and is not consistent with the mean μ of the second Gaussian being twice that of the first, for HSCs, but is consistent for non-HSCs. For the constrained fit, the DPE ratios were 0.275 ± 0.006 and $0.180 \pm .009$ for non-HSCs and HSCs, respectively. This corresponds to an 8.8σ result of a Z-test, indicating a statistically significant suppression of double-photoelectrons.

A suppression of the DPE effect may be due to the HSCs being a result of Cerenkov photons. Simulations of Cerenkov events in the PMT quartz windows indicate a different time signature, and higher max channel fraction than NRs in the TPC. Cerenkov photons tend to be in the visible spectrum, 100s of nm in wavelength. This lower energy can result in a suppressed DPE effect relative to the 175 nm[71] VUV scintillation light.

PMT Distribution

As HSCs are believed to result at least partially from Cerenkov light in the quartz windows of the PMTs, I examined the contribution of each channel to the HSC rate. I find that certain PMTs are over-represented in HSCs. I do not find a similarly sharp distribution in the distribution for non-HSC pulses. The elevated rate PMTs are distributed along the outer edge of the top array.




Figure 7.11: A histogram of the individual channel areas for low coincidence HSC and non-HSC pulses. This data was taken from early on in the first LZ science run in December 2021. The fitted Gaussians shown here are the result of the unconstrained fit. The shaded region indicates the data which was used for the minimization. A reduced DPE effect is evident in the distribution of pulses failing the HSC cut. Note that the first peak is not located at 1.0 phd, due to the phe-to-phd unit conversion.

HSC pulses do not exhibit a strong correlation between the identity of the maximum channel PMT and activation of nearby channels. The rate of HSCs decayed over time, similar to the non-HSC pulses. However they decayed significantly faster, being approximately six times lower at the end of SR1 than at the beginning.

Following the conclusion of SR1 a test was performed whereby one of the problematic PMTs was debiased and more long random data were acquired. Due to the aforementioned decay in HSC rate, this ended up having an inconclusive result. The rate, post-stinger cut, was 1.655 and 1.691 Hz for the post-SR1 data, PMT on/off respectively. A slightly higher rate of pulses present in the PMT-OFF data may be due to the distortion of the above-anode fields as a result of turning off the PMT. The HSC cut acceptance increases between these two configurations from 79% to 84%, indicating that this had the desired effect of attenuating this particular kind of iS1.

Top Bottom Asymmetry

Accidental events have no correlation between the S1 and S2 pulses. As such, any RQs which depend on the joint distribution will differ from a true recoil. The XY-position dependent




Figure 7.12: Histogram of channel contributions to the HSC pulses. The lower indices before(after) the break at 300 are the top(bottom) array. There are certain PMTs which are more problematic for high single channels than others, particularly PMT 239.

corrections for S1 pulse area are subdominant to the drift time, and the S1 centroids are not typically localized enough to utilize in the analysis. The primary joint RQ which can be exploited for distinguishing true and accidental nuclear recoils is the correlation between top-bottom asymmetry and drift time.

The top bottom asymmetry is defined as

$$TBA \equiv \frac{A_{\text{top}} - A_{\text{bottom}}}{A_{\text{top}} + A_{\text{bottom}}} = 2 \frac{A_{\text{top}}}{A_{\text{top}} + A_{\text{bottom}}} - 1. \quad (7.6)$$

The two arrays have different geometric packing of the circular PMT cross sections. The bottom arrays is a hexagonal packing, while the top array is a concentric series of circles, which results in more PMTs in the same solid angle on the bottom array. At the liquid-gas boundary, total internal reflection may occur for incident angles to the normal $\theta_c \geq \arcsin n_{\text{gas}}/n_{\text{liquid}} \approx 46$ degrees. These two phenomena result in a larger light collection efficiency for the bottom array, and therefore a bias towards lower TBA for all heights within the LXe.

The TBA-drift time correlation in the WIMP search ROI was found with CH₃T calibration data. The Tritium decays were selected with $S2 \in [10^{3.2}, 10^{4.4}]$ phd, and $S1 \in [10, 80]$ phd. Additional anti-accidental selection criteria were a max channel area cut : $\text{MaxChArea} < S1^{0.4} + 1$, and a drift time selection of $dT \in [50, 850]$ μs which eliminated the impact of the grid events. The data were then binned into drift time bins and a fourth-order polynomial was fit to the means of the bins. The χ^2 minimization was performed on a normalized drift time coordinate $\tilde{t} = dT/955$ μs in order to avoid floating point precision errors. The resulting relationship is:

$$TBA(\tilde{t}) = -0.0427 - 0.544 \tilde{t} - 0.826 \tilde{t}^2 + 0.906 \tilde{t}^3 - 0.263 \tilde{t}^4 \quad (7.7)$$

From this, a prediction of the top area fraction $P_T = \frac{A_{\text{top}}}{A_{\text{bottom}}}$ can be obtained for any given drift time.

The observed top area fraction then experiences binomial fluctuations. Larger S1 areas have smaller fractional fluctuations around their expectation values, while S1 near threshold undergo more severe fluctuations. A binomial distribution with observed success rate \hat{p} after n trials has a confidence interval of

$$p = \hat{p} \pm \sqrt{\frac{\hat{p}(1 - \hat{p})}{n}} \quad (7.8)$$

However, this is just the asymptotic formula, valid for when the central limit theorem may be applied ($n \geq 10$). Below that point, this equation overcovers the results[195]. The confidence interval with more accurate coverage is obtained by integrating the probability distribution function $f(k; p, n) = \binom{n}{k} (1-p)^{n-k} p^k$. The result, known as the Clopper-Pearson interval, is:

$$B(\alpha/2, k, n - k + 1) \leq p \leq B(1 - \alpha/2, k + 1, n - k) , \quad (7.9)$$

where α is the significance level and $B(c; a, b)$ is the c th quantile of the beta distribution with degrees of freedom a and b . A review of several methods of calculating Binomial confidence intervals is found in Ref. [195]. A comparison of the two methods, along with a full error propagation of the ratio of two Poisson fluctuating variables, is shown in Fig. 7.13 for an expected TBA of 0. The overcoverage of the Wald (normal approximation) interval is clear below around 20 phd. Above 20 phd the two methods become largely indistinguishable, but the Wald interval is more mathematically (and computationally) simple to calculate. The resulting cut efficiencies are shown in Fig 7.14. In the case of the isolated S1s, the cut acceptance was found by iterating over an array of drift times, evaluating the acceptance of the drift time for all pulses as if they were found at that drift time, and averaging over the imputed drift times. The Wald and Exact intervals both accept 95% or more tritium events across the entire range. For isolated S1s, the Wald interval removes more pulses between 10 and 60 phd (for the same α), but is outperformed by the Clopper-Pearson interval below around 5 phd. Due to the spectrum of isolated S1s skewing heavily towards lower areas, this makes the choice of using the exact interval easy.

S1 Pulse Shape

Isolated S1 pulses were additionally filtered with pulse shape parameters, a criterion which I developed. While the probability of 4 or more photoelectrons piling up to form a pulse is vanishingly small, other processes may lead to a pulse shape which differs from the typical scintillation signal. Because iS1 pulses are peaked near the 3-fold coincidence threshold, it is not possible in general to utilize asymptotic formulae. As a result, a generic exponential + constant parameterization which accommodates the heteroskedacity of the data is utilized:

$$t = A + B \exp(S1 \cdot C) . \quad (7.10)$$




Figure 7.13: Analytic TBA vs. S1 area calculations, with confidence intervals calculated with several methods, indicated for $\hat{p} = 0.5$. The fluctuations happen over the top area fraction, and are converted back into $TBA = 2\hat{p} - 1$. Note that the Wald confidence intervals are almost identical to the full error propagation when selecting the α appropriately.

As many features related to pulse shape are selected as possible, and generally apply an upper and lower boundary. These parameters were manually tuned rather than optimized automatically, and an example is shown for the full-width-at-half-maximum (FWHM) RQ in Fig. 7.15. Calibration DD events were used as a “signal” dataset to choose the parameters. This causes an issue for the top-bottom-asymmetry(TBA) parameter, as the DD data was skewed towards the top of the detector. This provides a high-statistics estimate of the upper TBA contour, but does not encapsulate the lower contour. This contour can be calculated analytically using the relationship between TBA and drift time at large S1s. I calibrated this contour by “calibrating” from events near the cathode, which are more likely to be real events over accidentals than events in the fiducial volume because of xenon self-shielding. In order to be conservative, and since a drift time-TBA criterion was already present, the lower boundary of TBA was not utilized for the shape cut. The parameters are tabulated in Table 7.2.

When tritiated methane CH_3T calibration data became available, this dataset became the standard by which data quality selection criteria were evaluated. The S1 shape cut, when applied to the tritium dataset, shows a similarly high acceptance of “real” S1s, and minimal dependence on drift time, as demonstrated in Fig. 7.16




Figure 7.14: The cut efficiencies as a function of S1 pulse area for 3σ confidence interval selections. The tritium calibration data is compared against the isolated S1 long random data. The isolated S1s have been pre-selected against the high single channel(HSC) and stinger (no single electrons or others in the preceding $2\mu s$). Also compared is the choice of the exact (Clopper-Pearson) and Wald intervals.

Parameter	Units	A	B	C
AFT ₇₅ -AFT ₅₀ (upper)	ns	130	70	-0.05
AFT ₇₅ -AFT ₅₀ (lower)	ns	50	-50	-0.1
AFT ₉₅ -AFT ₅	ns	350	0	0
fwhm (upper)	ns	210	110	-0.1
fwhm (lower)	ns	80	-80	-0.05
promptFrac50 (upper)	1	0.45	0.5	-0.07
promptFrac50 (lower)	1	0.2	-0.15	-0.04
promptFrac100 (upper)	1	0.8	0.3	-0.07
promptFrac100 (lower)	1	0.4	-0.15	-0.04
rms (upper)	ns	80	20	-0.05
rms (lower)	ns	30	-30	-0.04
TBA (upper)	1	0.35	1	-0.05
TBA (lower)	1	-0.75	-1.25	-0.05

Table 7.2: S1-only shape selection criteria; parameters for Eq. 7.10.

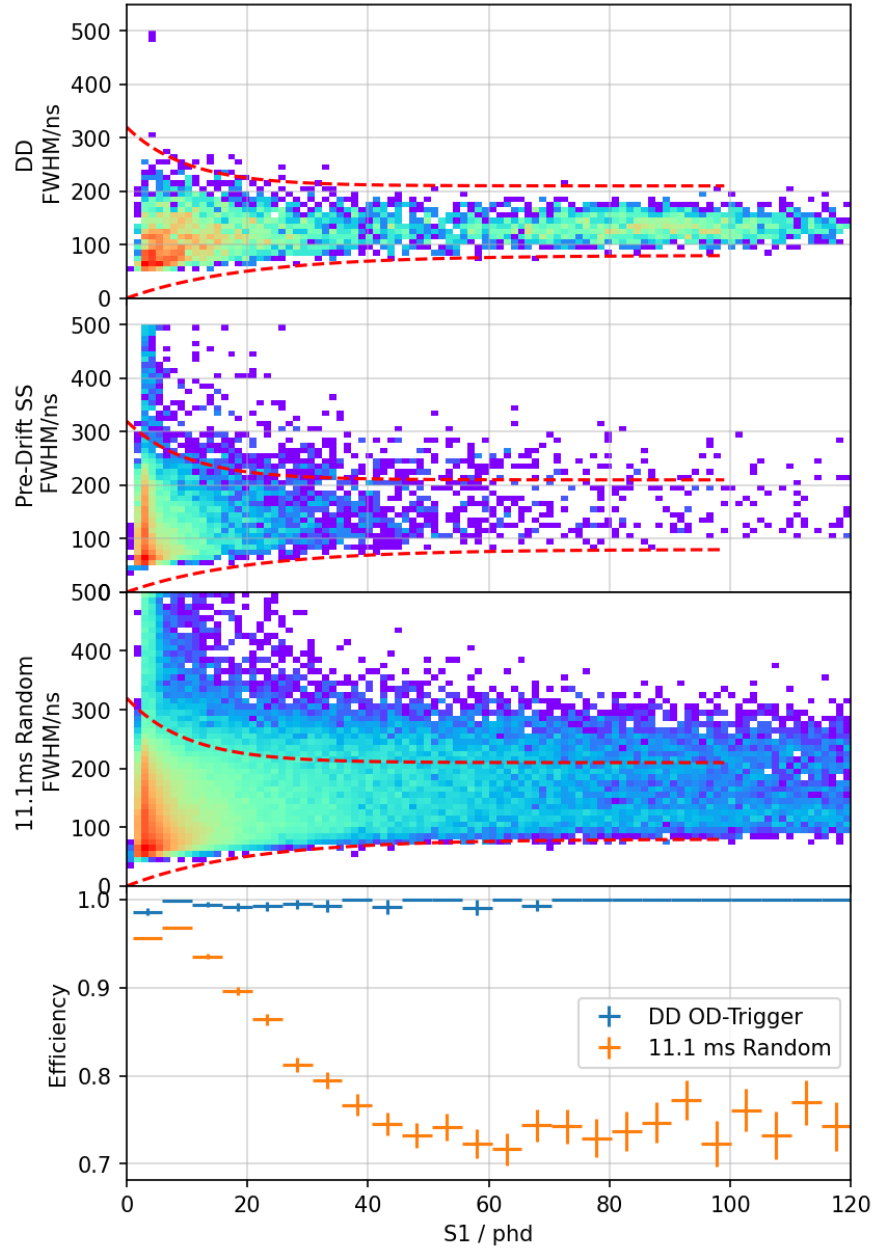


Figure 7.15: S1 shape decision boundaries for the FWHM RQ, as a function of S1 area. The top plot indicates a proxy for the “signal” which is contained. The middle two indicate iS1 backgrounds to be removed. The bottom plot shows the fraction of events in each S1 bin which fall between the indicated selection (dotted red). Isolated S1s are efficiently removed above 40 phd with $< 5\%$ loss across all bins for the DD data.




Figure 7.16: The S1 shape cut acceptance, as applied to the tritiated methane injection calibration dataset, and the iS1 pulses from Pre-SR1 data. The iS1 pulses have passed through the anti-accidentals selections such as HSC and stinger cuts. The effectiveness of the overall cut can be seen by comparing the asymptotic value here to the value for e.g. the fwhm cut along in Fig. 7.15. *Top, left:* CH₃T conditional acceptance vs S1 area. *Top, right:* CH₃T conditional acceptance vs drift. *Bottom:* DD and iS1 acceptance vs drift time.

7.4 Isolated S1 Locations

After the HSC, TBA, and stinger cuts, a population of isolated S1s still remains. These events are likely due to grid activity, based on the results of the grid electrode scan discussed in Section 7.6. While not directly exploitable for selection criterion, the sources of isolated S1s may provide clues for future mitigation strategies on the hardware side. S1s, as opposed to S2s, do not carry much positional information, besides at the upper and lower portions of the TPC, where the hit patterns of the photons are more localized. This is caused by the fact that photons typically reflect off the PTFE walls, the liquid-gas boundary, as well as Raleigh scatter within the TPC, erasing information on the deposit location. The primary handles on location are the top-bottom-asymmetry (TBA) and the centroids of the PMT hit patterns. Mercury is not run over S1s, therefore a likelihood-minimization procedure is not utilized.

In addition to TBA, another handle on the location of pulses within the detector is the “spread”, taken to be the square root of the trace of the covariance matrix of the PMT hit pattern:

$$\Sigma_{ij} = \frac{1}{N} \sum_k^N (X_{ik} - \bar{X}_i)(X_{jk} - \bar{X}_j). \quad (7.11)$$

Events occurring near the PMT arrays will have a more concentrated hit pattern, i.e. smaller “spread.” This variable is slightly correlated with TBA, as seen in Fig. 7.17. The pulses which pass and fail the HSC cut are presented separately, and all pulses have passed through the stinger selection criterion. HSC pulses appear at more extreme TBA, indicating that the grids are likely to be sourcing those pulses. Confidence intervals around particular drift times were constructed which appear to trace the TBA-S1 distributions of HSC and non-HSC pulses, as illustrated in Fig. 7.18. There, the TBA for HSCs were constructed using channels *other than* the maximum channel, to avoid bias. HSC pulses appear to be concentrated in a sharp band around drift time 500 μs , along with more diffuse distributions at larger TBA. Non-HSC pulses are more evenly distributed, but have an overdensity suggestive of a maximum drift time of 100 μs .

The top array centroid patterns are likewise differentiated between the HSCs and non-HSC pulses. With the HSC pulses, a circular pattern is present along the outer edge of the array, with a bright spot in the middle which is indicative of a diffuse pattern across the array. The non-HSCs are biased towards the top of the detector, as evidenced by the radial symmetry of the bottom array centroids. In the top array, non-HSCs show three distinct blobs, with one consistent with a diffuse pattern averaged too the center. The brightest region is apparent at (-40, 0) cm. Interestingly, the HSC pulses show a ring like pattern in the bottom array at radius ~ 30 cm (though the physical source is likely at the wall, due to the averaging process reverting towards the center). A bright spot is evident at a particular spot at the edge. Combined with the evidence from the TBA distribution, as well as the unexplained bend in the estimated wall position around that region of the TPC, signs point to some manner of short or conducting path existing in the vicinity.




Figure 7.17: Histograms of PMT hit pattern spreads vs top-bottom-asymmetry, presented for top array (left) and bottom array(right), and for HSC pulses (top) and non-HSC pulses (bottom). The y -axis is the square root of the trace of the covariance matrix $\sigma = \sqrt{\Sigma_{xx} + \Sigma_{yy}}$. All pulses have photon count >5 in the respective arrays. Note the concentration of pulses with high single channels near the top and bottom of the TPC.




Figure 7.18: TBA-Area bands for HSC cut failing (left) and passing(right) isolated S1s from the Pre-SR1 dataset. The HSC failing pulses (i.e. they have a high single channel) partially converge on a band around $500 \mu\text{s}$ drift time, approximately halfway up the TPC. The HSC passing pulses (likely scintillation) converge on locations slightly more evenly distributed throughout the detector, but appear to not exist above drift times of $100 \mu\text{s}$.

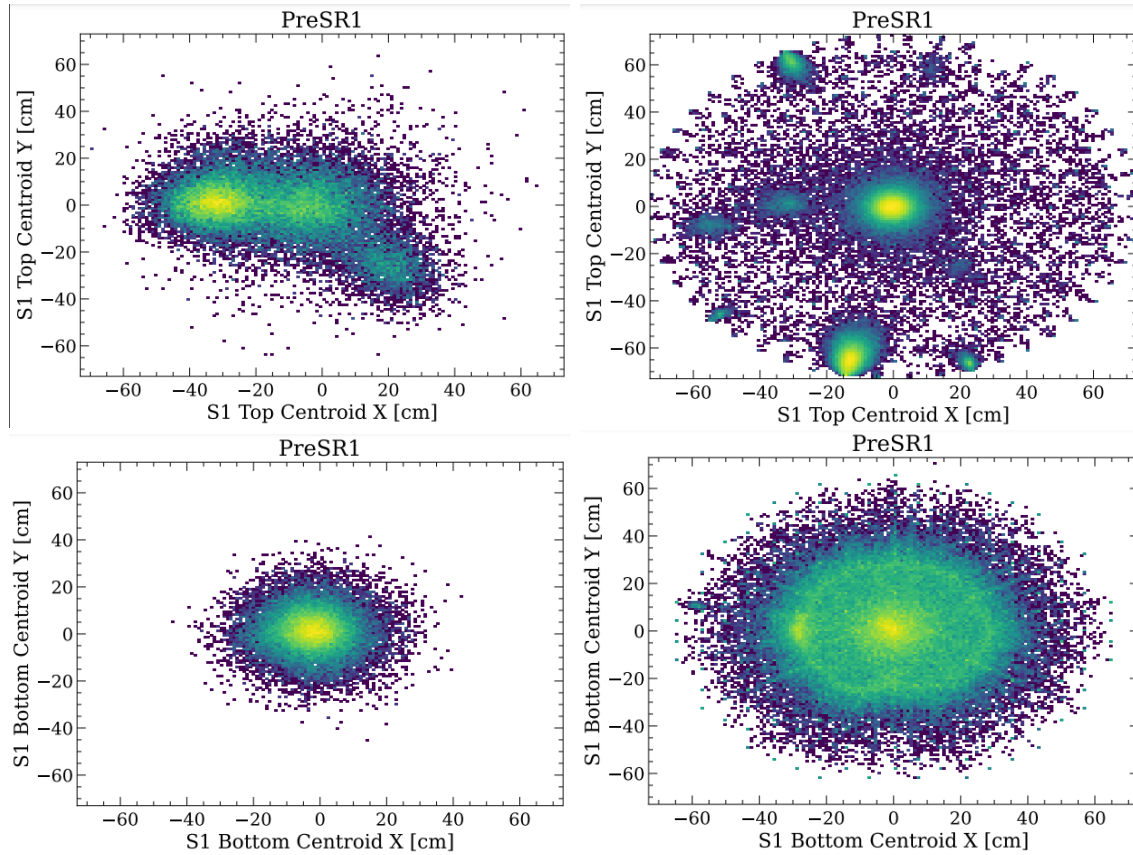


Figure 7.19: Heatmaps of the PMT array hit centroids for isolated S1s from the Pre-SR1 long random acquisition. Left column: passes HSC cuts. Right column: fails HSC cuts. Top row: top array. Bottom row: bottom array. The pulses all pass the stinger cut and have S1 area > 20 phd.

7.5 Creation of accidentals model

Introduction

While subdominant to other background sources in SR1 such as β -decays leaking into the nuclear recoil band, accidental coincidence events are particularly challenging due to their bias towards the low-energy region, making it more likely to appear not only in the WIMP search ROI, but in a region which favors a light WIMP mass. Therefore the accidentals S1/S2 spectrum must be incorporated into the complete backgrounds model. I contributed to the production and deployment of an applicable version of this model for SR1. This constituted a background in the WIMP Search ROI ($S1_c \in [3, 80]$ phd) of 1.2 ± 0.3 events[53].

Isolated Pulse Extraction

While long-random data is suitable for gathering large statistics of isolated pulses in a short amount of time, as well as examining long correlations between pulses it can obfuscate the effects of the shorter buffer windows in the SR1 data. It also blinds us to the potential time evolution of isolated rates, or particular effects that appeared after the detector settled into stability. Therefore, the pulses which are selected for construction of the accidentals probability density function are drawn from the SR1 WIMP search data itself.

Prominent pulses are identified within the event window. This is meant to be a finer selection than the initial pulse finding code, aiming to remove improper pulse boundaries. The decision is made by comparing the height/width ratio of S1(2)s against the largest S1(2) within the event, and the particular decision boundary differs between the two species. Depending on the number and ordering of prominent pulses the event is classified as one of several *interactions*:

1. Single Scatters: 1 or 2 prominent S1s preceding 1 prominent S2. If there is only one prominent S1 then the interaction is considered for the WIMP search. The double S1 events are a result of, for example, ^{83m}Kr events into this interaction type.
2. Multiple Scatters : 1 prominent S1 followed by 2 or more prominent S2s. This topology captures neutrons and Compton scatters
3. Pileup: 2 or more S1s followed by 2 or more S2s. As the name implies, these are meant to capture non-causally related energy deposits.
4. Other scatters: everything else, including events where the pulses are in the wrong order. Because of the random/heartbeat triggers, this topology included both S1-only and S2-only events.

The isolated S1 and S2 populations are selected from other scatters. Isolated S1s are selected from the pre-trigger window of random trigger types, whereas isolated S2s are selected from S2-triggered others. In both cases, the preceding (or following) drift length is scanned for the opposing pulse type to ensure that a single scatter was not misclassified due to a complicated topology. The choice of pre-trigger S1s is crucial, as it is possible for larger S1 pulse areas to trigger an event acquisition. This would end up biasing the S1 spectrum downwards in the post-trigger region, as these events would not end up in the random trigger category. Only the pre-trigger window is considered in order to mitigate this bias. No such concern exists for the S2s as they are meant to trigger an event.

Since the pulses in these dataset have already passed through both the physics and livetime cuts, namely the e-train and muon vetoes, these events are representative of pulses which would eventually survive into the subsequent data quality cuts. Since this dataset is partially meant to evaluate the cut efficacy (in particular the drift time-TBA cuts), further pre-selection is not performed. The exception is the OD and Skin vetoes. These cuts are adequately evaluated using calibration sources, so it would needlessly inflate the size of this dataset to include pulses which would be guaranteed to be removed.

Chopstitch

The *chopstitch* method was originally developed for the event salting framework for the LZ WIMP search[57]. In that context, the S1 and S2 pulses from a known single scatter source (such as from CH₃T calibration) are used to construct a new spectrum of events. These synthetic events are then injected into the data stream, the analyzers being blind to their source as a bias mitigation technique. Only after all selection criteria are frozen are the chopstitch events “unsalted” and removed from the dataset. The number of salt events, and the spectrum are meant to be unknown, but similar to an actual WIMP model in order to mitigate the most amount of bias.

It is crucial for salting that the pulses which are selected for salting are subject to certain constraints, otherwise they will be unnaturally easy to remove from the final dataset. Primarily these constraints are related to the effect of drift time on the S1 and S2 pulses. For S1s, drift time is correlated with top-bottom-asymmetry(TBA), and for S2s longer drift times result in more longitudinal diffusion, making the pulse longer and more Gaussian in shape.

For accidentals, much of these concerns evaporate, as the accidentals are by definition uncorrelated S1s and S2s. Therefore, assembling the accidentals probability density function requires *removing* the constraints of the salting framework. Taking the pulses collected in 7.5, 30030261 single scatter events were constructed from 66592 S1s and 387197 S2s, of which 28,945,817 passed the preliminary selection criteria such as sustained rate cuts. These events were passed through the data quality cuts listed in Section 7.3.

The chopstitch-accidentals events are efficiently removed with the data quality cuts. Between 99% and 99.8% of events are removed (varying as a function of S1 area), as indicated in Fig. 7.20. The acceptance rises for S1 and S2 near threshold, and is given by the number of events which pass the data quality cuts, divided by the number of events which are classified as single scatters. In total 5503 events pass all cuts, which presented some problems. This quantity of events is high enough that a probability distribution function had to be estimated for the ROI, but sparse enough that the high-energy regions would be challenging to fill. The chopstitch-accidentals dataset gradually grew in size as this problem was realized, but eventually it was realized that it would be challenging to scale the statistics by a sufficient factor. Therefore, a smoothing technique was used to interpolate the data.

The first step in this process was to verify that the S1 and S2 spectra are separable. While the inputs to the chopstitch data are uncorrelated, it was unknown whether the full analysis chain would introduce a bias. For instance, the prominence algorithm depends on all pulses in the event, which leads to some events being misclassified. The OD and Skin Vetoes are applied on the S1s before the S1 selection, so the non-coincidence requirements with the S1 should theoretically have no effect, if the chopstitch S1 is the same as the S1 which is identified for the single scatter. The OD and skin channels in chopstitch data are taken from the S1 channel alone. However, the default prominence algorithm is acausal in that it makes a determination of S1 prominence based on pulses which come later in the event window. This can alter the decision criteria, and therefore the vetoes may remove events by accident. This turns out to be a small effect, with the skin veto removing 122, and the

OD veto 518, single scatter chopstitch-accidentals events, for an inefficiency of $1.99 \pm 0.08 \times 10^{-4}$. Combined with other phenomena, effects such as these could produce unintended correlations.

To test for these possibilities, I performed two tests. First, I calculated the Pearson correlation coefficient from the covariance matrix Σ_{ij} :

$$\rho_{12} = \frac{\Sigma_{12}}{\sqrt{\Sigma_{11}\Sigma_{22}}}. \quad (7.12)$$

The S1c, logS2c estimators have correlations of $\rho = -0.001$, with a p -value of 0.462, which fails to reject the null hypothesis of no correlation.

The possibility of more complicated behaviours was explored using the G-test[196]. This consists of dividing the data into coarse bins in both dimensions, and calculating the expected counts in each bin E_{ij} assuming perfect separability $E_{ij} = \frac{S1_i S2_j}{N}$, where $S1(2)_i$ is the number of events found in $S1(2)$ bin number i . Then, the G-test statistic is given by:

$$G = \sum_i O_i \log \frac{O_i}{E_i}, \quad (7.13)$$

where O_i is the observation in bin i and E_i is the expectation for the same. This quantity is distributed like a χ^2 with degrees of freedom equal to (number of rows-1) \times (number of columns -1). Using 3 rows and 3 columns, evenly dividing the WIMP search ROI $S1_c \in [3, 80]$, $\log_{10} S2_c \in [2.75, 4.5]$, I obtain $G=1.9$, for a p -value of = 0.167. This similarly fails to reject the null hypothesis that S1 and S2 are separable in chopstitch-accidentals data.

From here, the 2D probability density distribution could be constructed. 1D Histograms were constructed in $S1_c$ and $\log_{10} S2_c$ spaces and combined. Due to the aforementioned sparsity, Poisson fluctuations in the histogram bins $\sigma_i = \sqrt{N_i}$ become significant for larger areas, which poses a problem for analysis. This is a particular worry for S1s, as there are only a few dozen surviving events above $S1_c = 60$ phd. One possible way to mitigate this effect would be to have non-uniform bins. However, this would pose a logistical issue when trying to add this pdf to the less sparse datasets that comprise the rest of the PLR. Another method would be to transport the data, calculate the histograms, then transform back, possibly followed by an interpolation step.

Rather than implement those techniques, the S1 and S2 histograms were transformed in slightly different ways. For both, a particular threshold was chosen such that below which the raw bin content would be used, and above which a smoothed value would replace it. In the case of the S1s, the smoothing technique was a power law fit. This favored a power law exponent of $x \propto S1^{-2.1}$, for a $\chi^2/NDOF = 1.3$, shown in Fig. 7.21. For the S2, since the data was significantly less sparse, and had no obvious analytic form, a kernel-density-estimator[197] was used instead. The cutoff thresholds were 16 and 2012 phd for the S1 and S2 spectra, respectively.




Figure 7.20: The total chopstitch cut acceptances as a function of corrected S2 area (left) and corrected S1 area (right). “Conditional CombinedGroup” here indicates that the acceptances are calculated according to the number of events passing the S1 and S2-based selections, divided by the number of events passing the “previous” livetime-impacting cuts such as *sustained rate* and the vetoes.




Figure 7.21: Chopstitch 1D histograms for corrected S1 and S2 areas, which form the inputs to the outer product which becomes the complete “accidentals pdf.” Note the location where the spectrum transitions from the bin content (blue) to the fit(S1) or KDE (S2) values.

UDT Normalization

While the PLR will allow the overall normalization of the chopstitch-accidentals dataset to float to best fit the SR1 data, it is advantageous to check the results against expectations, particularly with the low number of events. The main purpose of the exercise is to account for potential effects which only appear at the event level, such as unanticipated biases in the classification code. Taking the $iS1$ and $iS2$ estimated rates from Pre-SR1 long random data (1.04 Hz and 7.41×10^{-2} Hz, respectively), one obtains a rate of 7.32×10^{-7} Hz, corresponding to 3.8 events in SR1. Note that the 1.04 number is for the S1 pulse area range $[0, 80]$ phd. Then, using the SR1 range of $S1 \geq 3$ phd, one obtains 2.8 events within SR1. To constrain/initialize the values for the PLR, the rate of “unphysical drift time” or UDT events are used. These events are otherwise typical single scatters that were found to have drift lengths exceeding that of the cathode. The cathode in SR1 was identified at $951 \mu s$ as an obvious overdensity of events at all radii, as shown in Fig. 7.24.

Since UDT events are by definition known accidentals, they are a low-systematic estimator on the accidental rate in SR1. Due to these enlarged drift times, proxies for certain reduced quantities were required, such as the drift time being imputed by subtracting one full drift length. By imputing their drift times into the physical region, applicable data quality cuts may be applied, and an estimate for the true accidental coincidence background is found. The rate was found to be 1.2 ± 0.3 events within the WIMP search region of interest. The outer product, along with the UDT events, is shown in Fig. 7.22,

My contribution to this portion of the analysis was cross-checking the validity of the UDT events against the chopstitch-accidentals dataset. Similar to that case, I evaluate the correlation coefficient and G-test. The ρ of the UDT events in $S1_c$ vs. $\log S2_c$ space is $\rho = -0.0346$, $p = 0.543$. For the G-test, $G = 6.57$, $p = 0.828$, which indicates that, like the chopstitch data, the UDT events do not have statistically significant correlation between the two pulses introduced by the processing stages. I did find an inconsistency between the overall spectra of the UDT events and the chopstitch. The chopstitch events seemingly stretch lower in S2 area than the UDT events. It is possible that this is an artifact of the method of correcting the event areas using a proxy drift time. The G-test when taking the chopstitch as expected and the UDT events as observed yields $p = 0.007$.

7.6 Grid Voltage Scan

Overview and Results

After the conclusion of SR1, a campaign was conducted to determine optimal fields for SR2. The nominal voltages were lower in magnitude than the requirement voltages, and were chosen for being sufficiently high while maintaining stability. Three scenarios were considered: changing the drift field, changing the extraction region field, and shifting the voltages uniformly up and down while maintaining the fields. I was responsible for investigating the impact of these new configurations on the isolated S1 rates, which became an




Figure 7.22: The chopstitch-accidentals pdf, before smoothing. The unphysical drift time (UDT) events from SR1 are overlaid with red points.




Figure 7.23: The smoothed accidentals PDF, with SR1 data overlaid. The UDT events are largely consistent with the predictions from the accidentals PDF. NR and ER bands are overlaid to indicate the approximate impact on the WIMP search. A single UDT point in the NR, passing all selection criteria (besides fiducialization) is visible.




Figure 7.24: Single Scatter drift time spectrum in LZ, indicating the difference between physical and unphysical drift times. Events beyond 2 ms are due to events where the S2 is not the triggering pulse. This scenario occurs for especially large S1s, or when a background grid emission electron triggers an event.

input into the overall accidentals rate. The configurations explored were, in the notation Anode/Gate/Cathode in kV:

1. +4/-4/-32 : SR1 fields
2. +4/-4/0: Zero drift field
3. +4/-4/-20: Reduced drift field
4. +4/-4/-39 : Increased drift field
5. +4.5/-4.5/-32 : Increased extraction field. This also decreases the drift field marginally.
6. +3.5/-3.5/-32 : Decreased extraction field
7. +5/-3/-31 : Asymmetric voltages, increased anode voltage.
8. +2/-6/-34 : Asymmetric voltage, decreased anode voltage.

One challenge of this analysis was the overall instability of rates as a result of changing the grid configurations, particularly the anode. Increasing the anode in particular lead to erratic behavior in the system which at times took hours to calm down to a quiescent value which

could be analyzed. Also apparent was the appearance and disappearance of hot spots on increased electron emission from the grids at specific locations. Occasionally new hot spots flared up, and other times regions would quiet down for periods of time before increasing in rate.

These flare ups were dealt with by an analyzer monitoring the S1 filter crossing rate and adding regions of time to a database to filter out of the overall livetime. The filter crossing rate is defined by the online FPGA software which has the ability to serve as an event trigger. Isolated S1 rates were still elevated, even with these changes, and therefore I added additional time periods to remove elevated regions where possible. This process proceeded by analyzing the iS1 rate as described above, but binning the results into 30 minute exposure windows. This provides sufficient temporal information to see the decay of light after altering the voltages.

In some cases stability was never reached over the course of the test. While all the voltages listed above yielded definitive results, some cases were not able to acquire a quiescent value. A particular loss was a completely unbiased dataset with 0/0/0 voltages, as well as a dataset with zero extraction but still had drift fields.

The bursts themselves have interesting phenomena worth commenting on. Pulses found within the burst regions appear to occupy an intermediate region of max channel area vs total S1 area, between HSCs and the band mapped by DD and tritium. The top array centroids appear to be concentrated in a circular pattern around the edge, which is different than the patterns which appear in the other isolated S1 datasets. Bursts increase the rate across all areas, but in particular at low S1.

The general conclusion of the tests, from an isolated S1 perspective, is that the drift field is not a driver of accidental rates, as seen by the fact that the +4/-4/0 dataset did not considerably alter the iS1 rates. This can be seen in Table 7.6. It is somewhat less conclusive, but the extraction region field by itself does not appear to drive the rates, either. Rather, the anode voltage appears to be the largest determining factor for the voltages, based on the high rates associated with the 4.5/-4.5/-32 and +5/-3/-31 configurations. The associated spectra for those configurations is shown in Fig. 7.25.

Light Bursts

Following changes to the electrode configurations, the rates were elevated for a period of time. I selected regions of time for exclusion to estimate the rates above, as shown in Fig 7.26. For certain datasets no quiescent value was obtained, and therefore the overall rate was significantly higher than the others. During these burst events, pulses appear with maximum channel area between that of HSCs and non-HSCs, as shown in Fig 7.27. This new species appears in a ring surrounding the upper rim of the TPC.

Isolated S1 pulses which occur during these erratic periods have different hot spot locations, even for the non-HSC pulses, as shown in Figure 7.28. These patterns did not shift around for the SR1 field configurations. This indicates that isolated S1s could be in some way related to transient asperities on the grids [198].

Range	Anode	Gate	Cathode	Rate	Rate	Rate (Time Excl.)	Rate (Time excl.)
Units	kV	kV	kV	3-80 phd	0-80 phd	3-80 phd	0-80 phd
				Hz	Hz	Hz	Hz
PreSR1	4	-4	-32	0.760	1.040	-	-
PostSR1	4	-4	-32	0.670	0.876	-	-
20 kV Cathode	4	-4	-20	0.368	0.517	-	-
0 kV Cathode	4	-4	0	0.364	0.450	-	-
7kV Extraction (Sym)	3.5	-3.5	-32	1.305	1.511	0.501	0.647
Asymmetric Up	5	-3	-31	4.729	5.435	-	-
Asymmetric Down	2	-6	-34	1.167	1.619	0.333	0.478
Mid-Campaign SR1	4	-4	-32	3.787	4.371	0.801	1.039
9kV Extraction (Sym)	4.5	4.5	-32	3.081	3.530	-	-
39 kV Cathode	4	4	-39	0.473	0.631	-	-
End of Campaign SR1	4	4	-32	0.374	0.521	-	-

Table 7.3: Results of the electrode grid scan as applied to the isolated S1 rates. All rates are presented after the relevant cuts, and potential changes to the ROI and extra exclusion regions are tabulated. A general trend towards lower values was seen over the course of the test.




Figure 7.25: Differential iS1 rates vs corrected S1 area for various datasets over the WIMP search ROI. These datasets illustrate the impact of grid configurations consisting of changing the anode on the iS1 rate.




Figure 7.26: The decay of the isolated S1 rate following particular alterations to the grid voltages. These elevated rates were typically associated with changes to the electrodes, particularly increases to the anode grid. Periods of elevated rates (many σ from baseline) were removed where the transition was clear, in order to estimate the quiescent rate. Occasionally no stable rate was obtained.




Figure 7.27: *Left*: The max channel vs S1 distribution during a dataset with an elevated S1 rate. *Right*: The top centroid distribution of the intermediate max channel pulses. These are pulses which are above the HSC cut boundary $S1^{0.4} + 1$, but below the fixed ratio $0.7S1$. They are concentrated in a ring around the top array.




Figure 7.28: Isolated S1 hotspots in two different grid configurations. S1s with area >20 phd are shown. Changes to the gate and anode evidently activate or deactivate different features or asperities.

7.7 Machine Learning

Problem

For the SR1 WS selection criteria I hand-tuned the pulse shape cuts to provide discrimination against iS1 pulses based on a number of RQs. Likewise, a similar cut was developed by a member of the grids team to select on photon timing parameters. Looking forward, I initiated a process of automating this procedure using boosted decision trees (BDTs). These BDTs were additionally tuned on a weighted spectrum, made to look like a variety of WIMP masses. This provided more relevant information as it forces the model to place higher importance on low energy events than higher energy events. A demonstration of this procedure is illustrated in Fig. 7.29, where a combined AmLi + CH₃T calibration dataset is used to build a WIMP-like training set.

For the positive samples I use the high statistics, pre-SR1 long random dataset which I extracted above. Since this model is meant to remove events based on pulse shape parameters, those cuts were not used to generate the training set. The stinger and HSC selection criteria were applied before training so that the BDT does not learn cuts which are already close to optimal.

I used the following features for the BDT:

1. **Photon Timing:** channel photon mean, channel photon standard deviation.




Figure 7.29: The reweighted S1 spectrum which provides the negative samples for the boosted decision trees. The tritium dataset becomes similar in distribution to a 40 GeV/c² WIMP.

2. **Likelihood ratio:** An alternative method for folding the photon arrival times of an S1 into a single quantity. This tests the likelihood that the photon times, as measured by the channel peak times, are distributed according to an exponential distribution, relative to a uniform distribution. The exponential distribution of decay time τ is compared against a flat distribution of length L , both with the observed photons N . The likelihood ratio is given by:

$$-\log \frac{\mathcal{L}_{\text{exp}}}{\mathcal{L}_{\text{flat}}} = N \left(\frac{\langle t \rangle}{\tau} + \log \tau - \log L \right), \quad (7.14)$$

where $\langle t \rangle = \sum_i t_i$ is the mean of the photon arrival times. Maximizing the likelihood requires an estimator $\hat{\tau} = \langle t \rangle$. However, the best discrimination power is now found by comparing against a general exponential distribution, as we have an expectation for the decay time from the convolution of the triplet decay time (24 ns) and the electronics Gaussian shaping. As such, $\tau = 80$ ns is used, found by a receiver-operating curve analysis to be optimal.

3. **Pulse length proxies:** Area fraction times (75% - 25%), (95%-5%), pulse end-start. “Area fraction time” (AFT) refers to the time it take for a pulse to integrate to a given quantile of the total area.
4. **Prompt fractions:** the number of photoelectrons seen in a fixed window from the pulse start. Windows of (-50,50, 100,200, 500, 1000, 2000) ns were used as features.
5. **Pulse distributions:** Full-width-at-half-max (FWHM), root-mean square (RMS), peak time and skewness, based on area fraction time (AFT):

$$\text{Skew} \equiv \frac{AFT_{75} + AFT_{25} - 2 \times AFT_{50}}{AFT_{75} - AFT_{25}}. \quad (7.15)$$

6. **Pulse areas:** height, photonCount.

7. **Pulse locations:** top-bottom asymmetry (TBA), PMT spreads, PMT centroids.

Decision Tree Learning

A *Classification and Regression Tree* (CART) is fundamentally a series of yes or no questions arranged into a binary tree structure[199]. The leaf nodes of these trees provides the classification of the input features. When grown to arbitrary depth, each leaf node is associated with a particular member or members of the training set and provide a binary output. Frequently decision trees are grown to shallower depths, in which case they output the proportion of positive training values which are sorted into that node.

The training of the decision tree proceeds by recursively choosing optimal split boundaries over individual features. Any split is evaluated against the “information gain” between the

stages[200], $\Delta H = H_{k+1} - H_k$, where the entropy H is defined as:

$$H \equiv - \sum_i^C p_i \log p_i , \quad (7.16)$$

where the sum runs over the classes (in the binary case there are only two, and p_i is the probability of an element of the set belonging to class i). Every split occurs by iterating over features and sorting the data into groups based on whether that feature is larger or smaller than a certain value. The entropy, being extrinsic, is therefore $H_{k+1} = \frac{N_{>}}{N} H_{>} + \frac{N_{<}}{N} H_{<}$.

Decision trees are a low bias, high variance machine learning method, meaning that low training error may be associated with high testing/validation error. For this reason, two strategic prongs are used to reduce the variance: regularization and ensemble methods. Regularization refers to additional terms added to the information gain in order to confer a "cost" to more complex models. For DTs, the weights of the leaves are usually the input to penalty term, which could be a sum (L1) or a sum of squares (L2). Ensemble methods utilize a weighted sum of many decision trees, and two popular methods are boosted decision trees (BDT) and random forests[201]. Boosting is an iterative technique, where multiple stages of shallow trees are trained on the residuals of the previous step. The later epochs are discounted by a geometric *learning rate* and may also have a regularization term on the number of trees trained. Random forests (RF) are shallow trees which are trained in parallel, with each tree observing a randomly assigned subset of the data and features.

For this work I evaluate both RF and BDTs and evaluate them against the baseline model discussed in the previous sections. I also examine a logistic model for completeness, even though a linear decision boundary is not expected to perform well.

The hyperparameters of the models (e.g. the regularization terms) require optimization. The datasets are $O(10^5)$, and therefore dividing the events into training and testing sets would have created issues, especially considering the paucity of events near threshold. In its place, for the purpose of hyperparameter tuning, I use 5-fold stratified grid search cross-validation. This technique involves scanning over each hyperparameter, and dividing the data into 5 equal-sized *folds*, training on 4 folds, and evaluating some metric over the fifth. This process is repeated until all folds have been left out once, and the mean of the results gives an estimation of the quality of that set of hyperparameters. The *stratification* refers to the restriction that each fold should have an identical ratio of positive to negative classifications. For this work, I chose the area under the receiver operating curve (AUC) as the score. Additionally, I utilized a custom scorer, whereby the AUC was only integrated from 0 to 0.1 in false positive rate, in order not to select an optimal value that sacrificed too much signal.

The models were trained using the sklearn interface for XGBoost [202]. The hyperparameters that I scanned over were:

1. **BDT**: max_depth, learning_rate, colsample_bytree
2. **Random Forest**: max_depth, subsample, colsample_bynode
3. **Logistic**: L1 and L2 penalty terms

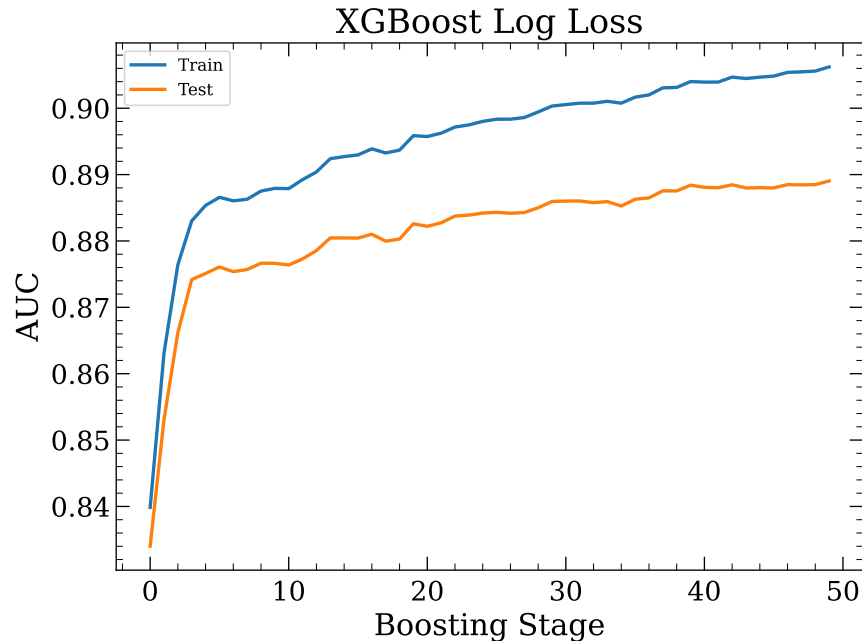


Figure 7.30: Boosting stages for the BDT training. A depth of 30-40 stages is optimal for preventing overtraining with these settings.

Results

A model was trained for each WIMP mass. Unless otherwise stated, I quote the results for the $m_\chi = 40 \text{ GeV}/c^2$ trained model. The BDT and RF models outperform the baseline RMS model, improving the iS1 identification from 18% to 60% at the same false positive rate (FPR). The random forest outperforms the BDT at extremely low FPR, but becomes comparable at higher FPR, as shown in Fig. 7.31. The acceptance as a function of S1 falls sharply for the iS1 pulses, from a peak of 0.6 to 0 at the upper end, as shown in Fig. 7.32. The Tritium+AmLi data is very flat across the range 0-60 phd, with $> 95\%$ acceptance. Weighting the negative samples according to a WIMP spectrum has a large effect at the upper end of the ROI.

Changing the weighting of the negative samples to match different WIMP spectra effects the characteristics of the surviving iS1 pulses. As shown in Fig. 7.33, the BDT automatically selects the upper bound of the data, and removes pulses above that threshold. Due to the iS1 spectrum being peaked near the threshold, which is a feature similar to WIMPs, the resulting iS1 rate does not vary by a large amount across the WIMP masses scanned. The surviving iS1 rate hovers around $0.75 \pm 0.1 \text{ Hz}$ for this model. Note that the baseline value, indicated in Fig. 7.34, is 1.27 Hz, which is higher than the value quoted above. This is due to integrating the area out to a slightly larger value for this test (115 phd vs 80) for this test.




Figure 7.31: Receiver-operating characteristic for the pulse shape cuts. “RMS” represents a flat cut on $chPhotonRMS$, which is the RQ used in the Photon Timing selection criterion. There, RMS values greater than 80ns are removed. This results in a false positive rate of 3.8%.




Figure 7.32: Acceptance vs S1 area for the BDT and RF models. The random forest and boosted decision trees perform similarly for signal (reweighted calibration ERs), while the BDT removes slightly more iS1s appear throughout a wide range of areas.




Figure 7.33: WIMP mass results. A smaller simulated WIMP mass results in more large-area iS1s being removed, as expected. However, the low-area events are not removed with the same efficiency, possibly due to the effectively reduced number of training samples.




Figure 7.34: Scan over simulated WIMP masses for the BDT selection criterion. The BDT threshold was chosen to set FPR=0.03. The dependence of WIMP mass over this range is limited, but it is unknown if this will continue to higher and lower masses.

Feature Importance

Decision trees are frequently referred to as a “black box” method, where the intuitive reason why any particular classification was chosen is not readily available. While the trees may be walked down, with dozens of boosting stages this becomes far more abstract. Interpretability is frequently separated into “local” and “global” methods. Global methods provide information about which features were most useful to the training of the overall model. Local methods provide the information about the decision boundary for a particular data point. A global method which is readily available is the “gain,” which yields the information gained (Eq. 7.16) on average with splits that use the particular feature. Another method is the *permutation score* which finds the change in the loss function as a result of shuffling the values of a feature, removing any information that it provides without altering the model.

A mixed local and global method for interpretability which is commonly used is the *Shapley values*[203]. Shapley values are a game-theoretic method for determining the share of a prize each member of a team should receive. In machine learning, the “players” are the features, and each “game” is a row in the data matrix. Essentially, various teams, or *coalitions* S are tested for their worth. Successful predictions become part of the *value* of the coalition $V(S) = \sum_{i=0}^{|S|} \hat{y}_i \log y + (1 - \hat{y}_i) \log(1 - y)$. The Shapley values are then the expected increase in value for coalitions including the feature:

$$\phi_j(V) = \sum_{S \subset \{1..p\}/\{j\}} \frac{|S|!(p - |S| - 1)!}{p!} (V(S \cup \{j\}) - V(S)). \quad (7.17)$$

The value function can be evaluated over the entire dataset, or individual points. For the XGBoost BDT, the gains and Shapley values are shown in Fig. 7.35. The broad conclusion is that the number of PMT hits and the top-bottom asymmetry are the best discriminants of scintillation / isolated S1s.

7.8 Single Scatter Contamination

The scenario of isolated S1s pairing with isolated S2s is not the only possibility for these pulses. Isolated S1s may also pile up with an otherwise good single scatter, causing it to appear like a double S1. The same situation may occur for isolated S2s, flipping the event to a multiple scatter. This causes a small loss of signal acceptance which can be accounted for.

Pulses are categorized as being “prominent” or not as a method of mitigating spurious splits done by the pulse finder. The algorithm consists of assuming that the largest S1 or S2 is prominent, and comparing the other pulses of the same type against them. Decisions are based on the height-to-width ratio of the smaller pulse, and the ratio of areas. Therefore, in order to cause a misclassification of the event topology, the accidental pulse must also pass this “prominent” algorithm.

I analyzed the isolated S1 spectrum in this space, and compared it against simulated WIMP spectra, extracting the “contamination rate.” Since WIMP spectra below $100 \text{ GeV}/c^2$




Figure 7.35: Feature importance values for the isolated S1 classifying boosted decision tree. *Left*: XGBoost outputs sorted by average information gain across nodes. *Right*: Shapley values averages across the samples of the test dataset.

vary based on the reduced mass of the target, I scan over the mass of the simulated WIMP. The simulations were performed using the LZ fast chain, LZLAMA+DMCalc.

The rate is calculated according to the following formula:

$$R_{\text{contam}} = T_{\text{pretrigger}} \int dx \int_a^b R(x) P(x, S1') dS1', \quad (7.18)$$

where $T_{\text{pretrigger}}$ is the pre-trigger length of the event (2 ms), x abstractly represents the isolated S1 variables, $R(x)$ is the isolated S1 rate, and $P(x, S1)$ indicates whether x is considered prominent. When the isolated S1 exceeds the area of the "real" S1 then it is assumed that it always passes the prominence cut (otherwise, it confounds the analysis by "stealing" the S1 from the otherwise good event). The way this works is illustrated in Fig. 7.36.

The resulting contamination rate ends up demonstrating a $\mathcal{O}(0.1)$ dependence on the WIMP mass, being worse at low masses (since the S1s are on average smaller, making the isolated S1s more likely to be the larger pulse). Over the masses considered, all contamination probabilities are less than one percent. I examined several isolated S1 datasets, including the random triggers in SR1. The Pre-SR1 dataset has a higher rate of iS1, and thereby has a worse contamination rate. The isolated S2 rate was estimated by fellow graduate student Ryan Smith, and is an input to this estimate. For a 40 GeV/c² WIMP I predict a contamination probability of 0.0037.




Figure 7.36: *Left*: The prominent cut boundaries overlaid on the isolated S1 histogram. Gray regions indicate pulses which would result in a contamination, which white regions would be deemed “non-prominent.” The red lines are the boundaries for different “largest real S1” areas. *Right*: Single Scatter contamination rate for several datasets. The elevated iS1 rate in the Pre-SR1 data (which decayed away over the course of SR1) lead to an increased likelihood of contamination, flipping single scatters to multiple.

7.9 Summary

Over the course of the commissioning, SR1 data taking, and beyond, I worked to characterize the isolated S1 spectrum, which is an input to the accidental coincidence background in LZ. I found that the rate exceeded expectations from the irreducible sources such as PMT dark pulses and the RFR. A hypothesis, supported by the data, is that the anode grid produces these isolated S1s. Isolated S1s were found to come in several flavors, one example being a “high single channel” with properties consistent with a longer wavelength light pulse. This rate slowly decayed over time from a rate of approximately 1 Hz in the WIMP search ROI to a value of < 0.4 Hz near the end of the post-SR1 / pre-SR2 calibration campaign. I developed selection criteria which reduced the isolated S1 rate by a factor of four in SR1. This aided significantly in making the accidental background subdominant in the SR1 analysis. Future improvements to the analysis can be achieved by deploying machine learning models, some examples of which were discussed in this work.

Chapter 8

The Xenon Breakdown Apparatus

8.1 Introduction

The drift fields in dark matter direct detection experiments are optimized at ~ 300 V/cm. Since the next generation of detectors is expected to be a factor of 2-3 larger in linear dimension, the applied cathode voltages must necessarily increase as well. While the bulk fields remain the same, this leads to challenges in the feedthrough design, as detailed in 3. However, in the design of these features, there is scarce LXe data available to reference. Fine features, like wires, have been studied for decades, but dielectric breakdowns in LXe at large stressed electrode area (SEA) have only recently been examined[104, 105]. In this chapter, I detail my contributions to the XeBrA experiment, which studied the risk factors for dielectric breakdown up to SEA of 33 cm². This consisted of analyzing data from initial runs, refurbishing the system following a period of inactivity, and running and analyzing several further runs.

8.2 Dielectric Breakdown Theory

Breakdowns occur when an insulating material becomes conducting, deviating from an Ohmic relationship between voltage and current. This nonlinear behavior is achieved through multiplication of charge as it moves through the medium between cathode and anode. With a sufficiently large electric field E , the kinetic energy acquired over one mean free path, μ , is $\Delta = E\mu$. When Δ exceeds the ionization energy W_0 an additional electron can be liberated, continuing the process until all electrons reach the anode.

The process of electron multiplication is described by Townsend's first coefficient $\alpha = \frac{1}{N} \frac{dN}{dx}$, where N is the number of free electrons and dx is the distance traversed through the medium. The number of electrons eventually arriving at the anode is given by

$$N(x) = N(0) \exp\left[\int_0^x \alpha(x') dx'\right] \quad (8.1)$$

The first Townsend coefficient represents the probability that an electron drifts long enough for its kinetic energy to exceed the ionization energy E_I . The mean free path through a medium is given by $1/\mu = n\sigma$, where σ is the interaction cross section. The Poisson probability for no interaction within distance x is $\mu^{-1} \exp(-x/\mu)$. When the electric field $\mathcal{E} = V/d$ is applied to an ideal gas $p = nkT$, this provides

$$\alpha = \exp(-E_I n \sigma d / eV) = \exp\left(-\frac{E_I \sigma p d}{eV kT}\right) \quad (8.2)$$

Generally speaking, this *avalanche* does not in general lead to a macroscopic discharge current. In addition to impact ionization from drifting electrons colliding with neutral atoms, secondary ionization will enhance the current over time. Secondary ionization mechanisms involve positive ion impact on the cathode, as well as photoionization from the electrons excited by the initial current. Photoionization may play a large role with liquid xenon, as the 177 nm scintillation light has an energy of ~ 7 eV, exceeding the effective steel-xenon work function of 3.8 eV [204, 205]. These processes are encoded in Townsend's second coefficient γ , where the number of electrons produced at the cathode $N'(0) = \gamma(N(x) - N(0)) = N(0)\gamma(\exp(\alpha d) - 1)$. When $\gamma > 0$, the current N becomes, after summing the infinite geometric series with $r = \gamma(\exp(\alpha d) - 1)$:

$$N(x) = N(0) \frac{\exp(\int \alpha dx)}{1 - \gamma(\exp(\alpha d) - 1)} \quad (8.3)$$

This equation diverges when $\frac{1}{\gamma} + 1 = \exp(\alpha d)$. Plugging Eq. 8.2 into Eq 8.3 yields the *Paschen-Townsend Law*, giving the breakdown voltage in terms of the physical parameters:

$$V = \frac{n\sigma d E_I}{\ln(n\sigma d) - \ln \ln(1 + \frac{1}{\gamma})} = \frac{\frac{E_I \sigma}{kT} p d}{\ln(\frac{\sigma}{kT} p d) - \ln \ln(1 + \frac{1}{\gamma})} \quad (8.4)$$

The Paschen-Townsend law is usually written in terms of empirically determined, medium specific constants $A = \sigma/kT$ and $A = E_I B$. Without *a priori* knowledge of γ_{eff} (which may in principle depend on the geometry of the apparatus), one can instead parameterize the equation with $A' = A/\ln(1 + \gamma^{-1})$:

$$V = \frac{B p d}{\ln(A' p d)} \quad (8.5)$$

Breakdown data collected in gaseous xenon are shown in Fig. 8.1.

Streamers

On centimeter scales, breakdowns typically require nonzero γ_{eff} to observe macroscopic (mA scale) breakdowns). Current also typically rise over the course of 100s of μs , owing to the need for several cycles of ionic drift. Breakdowns which occur short timescales (ns - μs) require a great deal of electron multiplication over the course of a single drift time. This phenomenon is known as *streamer* breakdown.




Figure 8.1: Paschen-Townsend law / predicted breakdown curves in xenon vs pressure-separation distance. These data were collected in gas using a variety of electrode materials and shapes. High pressure gas xenon data provides the best *a priori* prediction for the intrinsic dielectric strength of liquid xenon, when extrapolating to the appropriate density. Data from [206, 207, 208, 209, 210, 211]

In Townsend avalanche, the space charge is typically considered negligible in comparison to the bulk field. However, when αd becomes large, then the electron cloud may generate significant space charge, resulting in *streamers*. This occurs as a result of space charge, which distorts the field at the leading edge of the electron cloud. Such a distortion enhances the local electric field, increasing the effective α in a positive feedback loop. Eventually the space charge field is comparable to the bulk field. This occurs when the the *Raether-Meek* criterion[212] is met:

$$\int_0^d \alpha(x) dx \approx 18 - 20 \quad (8.6)$$

Fowler-Nordheim theory

Because the intrinsic dielectric strength of the liquid xenon is predicted to be high, it is unlikely that electrons seeded in the bulk (by, for instance, radioactive decays) lead to

breakdowns on their own. Therefore attention should be given to potential surface effects, due to their potential to generate large fields and currents which may seed streamers. Field emission is the process by which a large external field extracts electrons from a conductor into an insulating medium. Fowler and Nordheim[213] calculated the current density as a function of external field and work function, which is shown in Eq. 8.7. This equation was derived by evaluating the amplitude for quantum mechanical tunneling through a rounded triangular boundary.

$$J[\text{A/m}^2] = \frac{1.54 \times 10^{-6}}{\phi} 10^{4.52\phi^{-1/2}} (\beta E)^2 \times \exp\left(\frac{-6.53 \times 10^9 \phi^{3/2}}{\beta E}\right), \quad (8.7)$$

where E is the external electric field in V/m and ϕ is the net work function in eV. The parameter β describes the local enhancement of the electric field around an emitter. The emitter could be a small, sharp asperity on the cathodic surface. It is also possible for the Malter effect[214] to contribute. There, an insulating layer forms on the cathodic surface, which traps ions and prevents their immediate neutralization. This locally enhances the field and contributes to field emission.

The shape and size of the asperities determine the enhancement factor β . Cases with analytic solutions in the literature include a floating conducting sphere[215], a hemiellipsoid [216], and a capped cylinder. All of these arrangements were solved with the method of images, in some cases summing an infinite series of image charges and dipoles to cancel the bulk potential at the boundary. In the limit of infinite anode-cathode distance, a floating sphere of radius r in a bulk field E_0 , whose center is a distance h from the cathode, has a field of

$$E_{\text{sph}}(z) = E_0 \left(\frac{hr}{(z-H)^2} + \frac{hr^2}{2h} \frac{1}{(z-h-r^2/2h)^2} + 2 \frac{r^3}{(z-h)^3} \right) \quad (8.8)$$

Whereas an ellipsoid with eccentricity ξ , semi-major axis H , and focal point at $h = H\xi$, with infinite anode-cathode distance will have a field of

$$E_{\text{ell}}(z) = \frac{E_0}{\log(1 + \xi) - \xi} \frac{2h^3}{z(z^2 - h^2)}. \quad (8.9)$$

Taking into account the finite anode-cathode distance requires higher order terms that are suppressed by powers of h/L , where L is the electrode gap. Plugging in $z = h + r$ for the sphere and $z = H = h/\xi$ for the ellipsoid provides the field enhancement factors β_{sph} and β_{ell} , respectively. The case of a tube (cylinder capped with a sphere) can be approximated by the mean of the sphere and ellipsoid cases, $\beta_{\text{tube}} = (\beta_{\text{sph}} + \beta_{\text{ell}})/2$ [216]. All three model β can be approximated as functions of the aspect ratios H/ρ and h/L .

Reliability Analysis

The statistics related to analyzing the reliability of a system is distinct to that of other families of distributions. Typically physics experiments deal with independent, identically

distributed quantities (*i.i.d.* for short). These distributions are also usually representative of, or directly related to, the quantity of interest. When the experiment is instead testing the likelihood for a set of conditions to lead to a failure, this introduces interesting aspects to the analysis.

Two typical experimental designs in reliability analysis are the *time-to-failure* and *linear ramp* configurations. In the former, a set of conditions are established over a short time window (compared to the typical time to failure) and the results of the experiment is the amount of time it takes for the setup to reach the failure state. The experiment is then reset and repeated many times to obtain an estimate for the distribution of failure times as a function of the conditions. The latter setup chooses one condition to continuously change, and the result of each experiment is the value of the condition at the time of failure. A side effect of the linear ramp is that if the expected time to failure changes slowly enough compared to the expected ramp time, this introduces a dependence on ramp rate. Each configuration yields distinct results, but the linear ramp can be recast in terms of the time-to-failure by noting that $t = x/\dot{x}$ if $\ddot{x} = 0$.

The formalism of the time-to-failure configuration will be introduced here. Suppose a distribution of failure times $f(t)$ with accompanying cdf $F(t) = \int f(t)$. The survival function is defined as the complement of the cumulative distribution function $S(t) \equiv 1 - F(t)$ and is the probability of not seeing a failure in an experiment up to point t . One is frequently interested in the answer to the question “if I have not seen a failure at time t , how likely am I to see a failure before $t + dt$?” The pdf $f(t)dt$ does not by itself provide this answer, because of the way it was obtained. To count a failure at time t , the experiment has to have survived up to that point. Therefore, the answer to the question above comes in the form of the *hazard function*, which is the differential failure rate, conditioned on surviving up to this point:

$$h(t) \equiv \frac{f(t)}{S(t)} = -\frac{\dot{S}(t)}{S(t)}. \quad (8.10)$$

In some sense the hazard function is the more fundamental quantity, encoding how the system becomes more or less reliable over time. One can invert the equation, to find the survival function in terms of the hazard:

$$\frac{dS}{S} = -h dt \rightarrow S(t) = \exp\left(-\int_0^t h(t') dt'\right) = \exp(-H(t)), \quad (8.11)$$

where $H(t)$ is the cumulative hazard, a dimensionless, positive quantity. The interpretation of $H(t)$ is slightly less intuitive than $h(t)$. It is the number of failures expected before time t , if you were able to run the experiment after a failure without resetting. In other words, if a system could hypothetically observe a failure at time t_1 , but instead of restarting back at 0, the clock kept running until the next failure. Therefore $H(t)$ may be larger than one, and is neither a pdf or cdf.

The hazard function(s) are quite useful for answering questions related to reliability. For example, if one has a system with multiple subsystems, each of which has the potential for

failure. If the entire system fails if any one system fails (as if it were a circuit in series), the survival probability is then the probability of not seeing a failure in any subsystem, or

$$S(t) = \prod_i^N S_i(t) = \prod_i^N \exp(-H_i(t)) = \exp\left(\sum_i^N H_i(t)\right). \quad (8.12)$$

Therefore, for a composition of unreliable systems, the effective cumulative hazard is simply the sum of the cumulative hazard of the subsystems. A similar analysis can be performed for components in parallel:

$$S(t) = 1 - \prod_i^N (1 - S_i(t)) = 1 - \prod_i^N (1 - \exp(-H_i(t))) \approx \sum_i^N \exp(-H_i(t)). \quad (8.13)$$

The approximation is valid if $H_i(t) \gg 1$ for all of the subsystems. If the hazard functions of the subsystems are identical, this makes $H'(t) \approx H(t) - \ln N$. For breakdowns in XeBrA, the series configuration is more fitting, as each “subsystem” can be thought of as a surface element over the cathode.

The parameterization of the cumulative hazard function can in principle be anything, but it is frequently chosen from a general family of functions. A popular choice is the *Weibull* function, where

$$H_W(t; \lambda, \mu, k) = \begin{cases} \left(\frac{t-\mu}{\lambda}\right)^k & t \geq \mu \\ 0 & t < \mu \end{cases}. \quad (8.14)$$

The quantities λ , μ , and k are referred to as the *scale*, *location* and *shape* parameters, respectively. When the location parameter μ is set to zero (implying no minimum threshold for failure) this is referred to as the two-parameter Weibull function. Inserting the two parameter Weibull cumulative hazard function into the series equation, assuming identical elements, yields the following scaling relationship:

$$\lambda = N^{-1/k} \lambda_0. \quad (8.15)$$

The XeBrA procedure thus far does not utilize the time-to-failure experimental design, but rather the linear ramp. To do the linear ramp analysis, first we make the substitution $E = \dot{E}t$. Next is to determine the dependence of the hazard on E and t . The system may fail immediately upon any element crossing a particular threshold, which has its own particular distribution[217]. The insulating medium between the electrodes could also progressively fail in a tree like manner, growing with fractal dimension $L \propto t^r$. Either the undisturbed material between the tree and electrode may fail, or the failure could occur when the tree finally bridges the gap between electrodes and forms a conducting path. In all of these cases the cumulative hazard function can be parameterized by

$$H(t, E) = \left(\frac{t}{t_0}\right)^a \left[\frac{(E - E_1)}{E_0}\right]^b \quad (8.16)$$

Inserting $E = \dot{E}t$, with $E_1 = 0$, this becomes:

$$H(E; E_0, E_1, k) = \frac{E^{a+b}}{t_0^a \dot{E}^a E_0^b}. \quad (8.17)$$

Therefore the effective Weibull shape parameter is $k = a+b$, and the effective scale parameter is $E'_0 = (t_0^a \dot{E}^a E_0^b)^{1/k}$. It is evident that there is a dependence on the ramp speed; $E'_0 \propto \dot{E}^{a/k}$, implying that faster ramps lead to larger breakdowns, assuming $a > 0$. When this ramp speed dependence is estimated, the a and b parameters can be separated. The a parameter is interpreted as the effect of the ‘‘history.’’ When $a > 1$, this implies that a progressive weakening occurs, with hazard increasing for later times. Whereas $a < 1$ implies a progressive strengthening of the system over time, and $a = 1$ indicating no effect of history. Strengthening can be due to a sort of ‘‘conditioning’’ effect, where weak elements are eliminated during failures, leaving a stronger overall system at later times. In solid media the history effect is due to the breakdown of the crystal and chemical structures. Liquid media are not expected to show a significant effect of history, due to the ability for damaged chemical bonds to be cleared out of the regions of high electric fields. Ferrofluids showed an increased breakdown voltage with increased ramp rate[218], a trend also observed in oil-impregnated pressboard[219], borosilicate glass [220] and alumina ceramic[221]. Thin-film polymers [222], and polymer nanocomposites[219] exhibit the opposite trend, becoming weaker with increased ramp rate.

In XeBrA the testing was conducted with three sequential ramps. This was done to rapidly set the voltage to a value where breakdowns were conceivable in order to fit in more breakdowns in a given acquisition. In each run a preliminary dataset exists, where parameters for the initial two ramps were tested. When ramp rate scans were performed, the reported ramp rate was for the third and final ramp, where the breakdown occurred.

Factors Affecting Breakdown

The surface area of the electrodes exposed to high fields is referred to as the *stressed electrode area* (SEA). As SEA increases, the number of asperities that could initiate a breakdown increases as well. In fact, an inversely proportional dependence of breakdown field with stressed area has been observed in several media [223, 224, 225, 226, 227, 228]. Other dependencies have been noted as well, the main ones being:

1. **Volume:** if the medium itself is a point of failure, then the breakdown voltage would decrease in proportion to volume. This can occur in cases in which high fields cause chemical reactions to occur or ionic bonds to dissociate. In noble liquids, neither are factors. However, one might observe such an effect in cases where bubbles are present throughout the stressed volume (if the medium is very close to the boiling curve).
2. **Purity:** electrons drifting in liquid attach to electronegative impurities, such as oxygen. The reduction in charge decreases the likelihood of triggering a breakdown.

3. **Pressure:** bubble nucleation and growth are suppressed by increasing the liquid pressure. It has been observed that increasing the liquid pressure resulted in higher breakdown voltages in both liquid nitrogen (LN) [224] and liquid helium (LHe) [229, 230].
4. **Ramp speed:** if the risk of breakdown is represented as a differential breakdown rate per unit time, then different ramp speeds might result in different maximum breakdown voltages. This has been observed in borosilicate glass and alumina [220], ferrofluid [218], and oil-impregnated pressboard[219]. However, an opposite correlation was found in thin film polymers [222].
5. **Surface finish:** higher sustained electric fields have been observed in more finely polished electrodes in both LHe [231, 232, 230] and LN₂ [227]. Furthermore, acid passivation and electropolishing have been shown to be effective methods at reducing emission rates in stainless steel wires [189].

8.3 Apparatus

Terminology

The data collected from XeBrA is hierarchical, and therefore I use the following meanings to refer to various concepts:

- **Stressed Electrode Area (SEA):** the surface area of the cathode with electric field magnitude above 90% of the maximum field magnitude.
- **Run:** A cycle of xenon filling, data taking, and recovery.
- **Ramp:** A cycle of ramping the high voltage power supply on the cathode from 0V until a breakdown occurs, followed by a ramp back down to 0 V. This usually takes several minutes and constitutes several measurements.
- **Breakdown:** The discharge which causes a fault signal to be sent, and a ramp to terminate. These are almost always bright, and particularly large breakdowns can actually be heard by conducting sound through the steel. Each ramp may have multiple breakdowns in quick succession, since the slow control can not respond with a ramp down signal fast enough.
- **Dataset:** A series of ramps at with fixed settings, such as electrode separation and pressure. Changes to instrumentation do not cause a change in dataset number. Datasets are labelled according to which run they occurred, so dataset 3 from Run 5 would have a “global dataset” identifier of 0503.
- **Event:** a single acquisition of the fast DAQ. This happens continuously with no dead-time, but for each of analysis is divided into sections of 1E6 samples for each channel and analyzed as if it were a triggered event.

- **Precursor:** a transient spike in light or charge production during a ramp which does not produce a fault signal.
- **Glitch:** An intermediate size spike in current which is sufficient to distort the HVPS supply current but is insufficient to trigger a fault signal. These are typically not visible to the naked eye.
- **Tilt:** The relative angle between the cathode and anode.

Overview

The Xenon Breakdown Apparatus (XeBrA) is a 5 liter stainless steel spark chamber designed to test dielectric breakdowns in noble liquids at large, variable stressed areas. It consists in part of an inner and outer cryostat, separated by a vacuum layer and suspended by an 80-20 test stand. The outer cryostat vessel (OCV) has two viewports which enable real time monitoring of the condition of the liquid. The inner cryostat vessel (ICV) is cooled by a pulse-tube refrigerator, which connects through the lid of the outer cryostat via a cold finger. Copper cooling links maintain a consistent temperature across the ICV. High voltage is delivered to the cathode from the bottom via a conical feedthrough. Xenon (or argon) gas flows through an inlet on the lid, condenses within a heat exchanger, and is injected below the cathode. The gas outlet from the heat exchanger returns to a gas panel. When the system is continuously circulating xenon, the gas returns through an SAES getter, which purifies the gas to 1 ppb O_2 equivalent concentration. An outside view of the test stand is shown in Fig. 8.3.

The gas panel piping and instrumentation diagram (PID) is shown in Fig. 8.2. The system was designed by Lucie Tvrznikova, then a Yale graduate student. As multiple experiments share the lab space, the panel was designed for the possibility of routing the xenon to either XeBrA or another nearby noble liquid experiment. The xenon flow path runs from the high pressure storage bottles into the circulation path. A regulator (RG01) controls the flow of xenon from the high pressure area (denoted in orange) into the relatively lower pressure area (denoted in green). The MV06 valve is in parallel with a Matheson¹ moisture trap. Burst devices are placed in various locations around the panel, designed to vent to the room in the case of an overpressure event. A higher maximum pressure (9.4 bar gauge) is asserted on the the section near the circulation pump. This is to allow for higher pressures while filling and circulating, which are necessary for higher flow. The added impedance of the getter is thus crucial for the protection of the lower-pressure components closer to XeBrA itself. Purge lines run from XeBrA to the return path, which help to maintain uniformity of the xenon properties in otherwise stagnant regions like the PMT and the purity monitor.

Overpressure events are dealt with in one of several ways. Relatively slow rises are caught in software, in which case a pneumatic valve is opened and the xenon is vented into 450 L, stainless steel storage containers (adapted from propane tanks). In the event of a power loss, the burst device connecting the gas panel to the emergency storage vessels will break,

¹www.mathesongas.com/

venting automatically. Due to the desire to recover xenon back in these cases, the storage vessels must be evacuated beforehand. Due to their large size, this is a challenging prospect, and makes smaller overpressure events more costly in terms of contamination. If, for some reason, neither of these interventions works to lower the pressure, eventually burst devices will open to the room, venting the xenon. This would only occur if some kind of blockage were to prevent the xenon from flowing into the storage containers, as that burst device is designed to break first.

The interior xenon volume houses the high voltage electrodes. These are two circular, stainless steel pieces with a Rogowski[233] profile, as seen in Fig.8.4 A Kimbal sphere with six flanges comprises the core of the ICV, with features branching off in all directions. The bottom flange connects to the ceramic cathode feedthrough. On the top a rod which connects to the anode, allowing it to be moved up and down. Two of the flanges are left blank for visibility. A photomultiplier tube fits into a section at the back, intended to monitor the evolution of the single photon rate over the course of a ramp. A purity monitor occupied the remaining flange. The purity monitor is a small drift chamber which measures the attenuation of electrons generated by a xenon flash lamp incident on a gold plated cathode. This arrangement is shown in Fig. 8.5, along with the associated joints.

Rogowski electrodes

In order to test the area scaling of dielectric breakdown, large surface area electrodes are required. Ideally the electric field is uniform across as much of the surface as possible. In XeBrA, it is also required that the surface area subject to the largest electric fields is controllable. For these reasons, the Rogowski[233] profile was chosen. The Rogowski profile achieves a highly uniform electric field across the surface, and the stressed electrode area (SEA) scales with gap distance, obviating the need to replace electrodes.

Other choices of profile are possible, e.g. plane parallel, hemispheroidal, and Bruce[234]. The Rogowski profile curves away from the center according to an exponential function. The Bruce profile is a plane center and sinusoidal edge. With the plane parallel, Rogowski and Bruce profiles the edges are typically capped with a semi-circular section in order to terminate the function. Using the metric $\eta = E_{\max}/E_{\text{avg}}$ to compare the profiles, the Rogowski profile maintains an $\eta \sim 1$ over longer gap distances than any of the profiles above[235].

The reason for the Rogowski profile's excellent uniformity over its surface comes from the way it is generated. The exponential curve comes from solving for the equipotentials of a finite plane at fixed voltage a distance d away from an infinite grounded plane. This yields the following set of parametric equations[236]:

$$\begin{cases} x = \frac{d}{\pi}(\phi + e^{\phi} \cos \psi) \\ y = \frac{d}{\pi}(\psi + e^{\phi} \sin \psi) \end{cases}, \quad (8.18)$$

where ψ parameterizes the distance along field lines, while ϕ parameterizes the distance along the equipotential surfaces. In principle any value of ψ may be chosen, but $\psi = \pi/2$ is

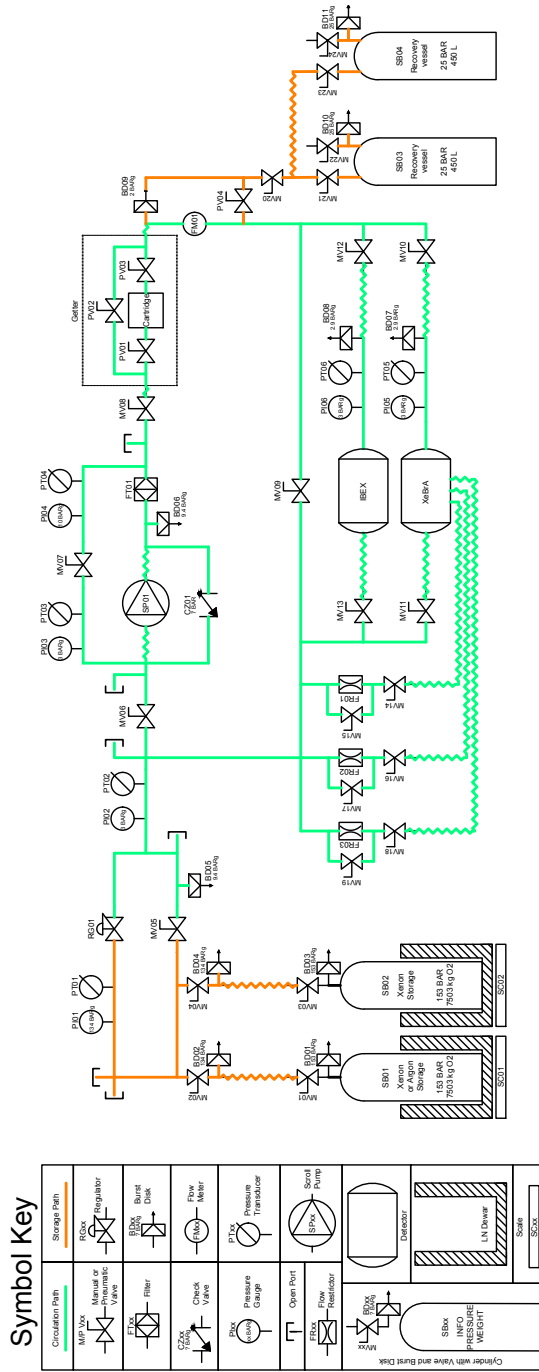


Figure 8.2: Piping and instrumentation (P&ID) of the XeBrA experiment.




Figure 8.3: The XeBrA test stand with the OCV in place. The turbomolecular pump on the top maintains the OCV vacuum. The tubes leading to the left connect to the cryocooler. The black high voltage cable is seen leading downwards from the can. The inlet and outlet hoses can be seen leading from the top to the cable tray to the right.

the most common. This choice yields:

$$\begin{cases} x = \frac{d}{\pi}\phi \\ y = \frac{d}{\pi}(\frac{\pi}{2} + e^\phi) \end{cases}, \quad (8.19)$$

where the case $\phi = 0$ yields the center of the electrode. Since the finite plane of radius R_0 is not a physical object, it is useful to work with the physical dimensions of the electrode, namely the radius and thickness of the object itself. These are given by (with the choice $\psi = \pi/2$):

$$T = \frac{d}{\pi}(2 + \sqrt{2}) \quad (8.20)$$

$$R = R_0 + \frac{d}{\pi}\sqrt{2} \quad (8.21)$$




Figure 8.4: The XeBrA Rogowski electrodes, in-situ.

The Rogowski electrodes in XeBrA are thickness 2 cm. These were simulated in FENICS, similar to the analysis done in Chapter 4. The anode radius is 56.6 mm, while the cathode is 48.4 mm. Both are constructed from 303 steel, polished to a surface roughness of $0.1 \mu\text{m}$.

Upgrades

Heat Exchanger Repair

One of the first tasks for the 2020 restart of XeBrA was to replace the heat exchanger (HX) which failed during the initial runs in 2018. A weld failed during a period of data taking, due to it not being made for positive pressure. It was essentially a brazed connection, and the 2 bar of pressure eventually broke the weld and lead to an emergency recovery, which terminated Run 3. Several kilograms of xenon were vented to the outer vacuum and could not be recovered. This heat exchanger was re-welded using a different, stronger method and replaced. Due to the Covid-19 Pandemic, this required special PPE, namely half-mask




Figure 8.5: A CAD rendering of the XeBrA ICV cross section. Note the alignment joint on the cathode. The purity monitor extends from the right hand side, and the cathode's ceramic feedthrough approaches from the bottom. The viewport is located on the left, and the anode is held by the adjustable rod from the top.




Figure 8.6: A simulated electric field profile for the XeBrA electrodes. Here, the configurations are: **Top, Left:** 1 mm separation 0 degrees tilt, **Top, Right:** 1 mm separation 1 degrees tilt, **Bottom, Left:** 3 mm separation 0 degrees tilt, **Bottom, Right:** 3 mm separation 1 degrees tilt. The black lines are the contours of constant electric field magnitude. The field is evaluated over the cathode surface.




Figure 8.7: A typical refrigeration cycle plot in pressure-enthalpy space. The xenon flows through the indicated points as it passes through the HX, XeBrA, and the circulation pump. Figure taken from [237]

respirators as I and project scientist Ethan Bernard had to be within six feet of each other to make these connections.

The operating principle of the heat exchanger is illustrated in Fig. 8.7. The XeBrA HX consists of an internal bellows surrounded by a stainless steel cylinder. Xenon enters through the outer volume, which serves as the condenser, where it is cooled on the walls via the PTR and the bellows by the outgoing cold gas. Gas exits through the bellows, which serves as an evaporator. The internal liquid xenon volume is separated from the evaporator by an expansion valve, which executes the isenthalpic expansion from points 3 to 4 in the diagram. Heat moves from the condenser to the evaporator through the walls of the bellows, causing both streams to undergo phase changes (lines 4-1 and 2-3 in the diagram). Eventually the outlet gas makes it to the circulation compressor, where it is adiabatically pumped from pressure P_1 to P_2 in the diagram.

One side effect of the loss of xenon is that the liquid level now sits lower in the apparatus when full. This leads to some amount of evaporation before the xenon reaches the evaporator, which harms the stability of the system. It was noticed that during the subcooling procedure (see 8.4) that the liquid level would sit abnormally low, roughly halfway down the viewports, too low to run the experiment. However, following raising the set point pressure (and therefore the temperature), the liquid level rose. Frequently, at this time a somewhat heavy “rain” of liquid xenon was seen on the cathode. This is interpreted as liquid xenon condensing in the relatively colder heat exchanger evaporator portion when the system pressure was low. When the circulation speed and HX temperature were raised, the liquid xenon began to boil, eventually losing surface tension and falling back into the Kimbal sphere.

Tilt and SEA Estimation

A significant improvement from the 2018 procedure is that of the improved tilt constraint. The 2018 procedure involved pressing the two electrodes together with a piece of contact paper between them. The contact paper is a graphite sheet between two pieces of white paper, which reveals a pattern of applied pressure. A circular pattern indicates that the cathode and anode tangent planes were close to parallel, whereas a smaller, elliptical pattern would indicate the opposite. The error on this method was estimated in a somewhat qualitative method. Photographs were taken of the electrodes, and the electrodes were rotated by a known, but blinded, amount in an image editing software. These pseudo-experiments were used to “calibrate” the ability of the experimenter to visually distinguish tilts. It was determined that an error of $\delta\Theta \sim 1^\circ$ was indistinguishable to the operator, and that was taken as the systematic error.

As this method drove the error in SEA, another method was developed and used for XeBrA-2020. As before, contact paper and visual confirmation was used to obtain the initial tilt angle. Additionally, a shed-free wipe was inserted between the cathode and anode, and the two electrodes were pressed together with the set screws loose, allowing the cathode to be pushed into place. Subtracting the thickness of the wipe, this allowed for calculating the tare of the linear shifter. When the electrodes were under sufficient compression, the set screws were tightened to immobilize the cathode. The wipe was then removed.

After the setting of the angle as close to zero as possible, the tilt was estimated using photogrammetry. Photos were taken from each viewport at a fixed camera perspective. The linear shifter was moved, increasing or decreasing the anode-cathode separation distance. These photos were then analyzed offline, with an aim of mapping out the contours of the anode and cathode. Graphs of separation distance (in pixels) as a function of linear shift setting were found at various points on the surface. The y-intercepts of points on either side of the center provide an estimate of the tilt angle from that perspective.

This process is performed for each viewport and the tilt estimation (and uncertainty) is calculated based on adding the angles in quadrature:

$$\theta^2 = \theta_L^2 + \theta_R^2 \quad (8.22)$$

$$\delta_\theta = \frac{1}{\theta} \sqrt{(\theta_L \delta_L)^2 + (\theta_R \delta_R)^2} \quad (8.23)$$

The results of these calculations are tabulated in Table 8.1.

While I assisted in the development of the technique, I was not involved in the execution of the fitting procedure. I propagated these errors into the SEA using the field simulations from 8.3. The simulations were performed over a fixed grid of tilt angle and unperturbed separations. Since the tilt of the cathode moves the point of minimum distance off-center, the physical separation distance differs from the untilted distance. Rather than pre-calculating this perturbation, I instead calculated the fields all at the same starting separation before applying the rotation matrix. From the resulting electric field magnitude, I located the peak electric field. As unit voltage was applied, I take the physical separation as

$$d = V_{applied}/E_{max} \quad (8.24)$$




Figure 8.8: Simulated stressed electrode areas vs. cathode-anode distances for each of the tilts which were estimated from the photogrammetric method. Notice the vast difference, almost a factor of three, between the SEA at the lowest and highest tilts, at the same separation distance.

Using the value of d from the simulations, and the estimated tilt angle, I found the SEA as a function of separation distance for all of the estimated angles. This procedure was repeated whenever the electrodes were replaced or touched. The results of the simulations are shown in Fig. 8.8.

Runs	Front [°]	Side [°]	Total [°]
4-5	0.1 ± 0.1	-0.6 ± 0.1	0.6 ± 0.1
6-7	-0.2 ± 0.1	-0.2 ± 0.1	0.3 ± 0.1
8	0.3 ± 0.1	-1.0 ± 0.1	1.0 ± 0.1
9	1.2 ± 0.2	-0.4 ± 0.1	1.3 ± 0.1
10	-0.1 ± 0.1	0.5 ± 0.1	0.5 ± 0.1

Table 8.1: Estimated cathode tilt angles for different Runs. The front and side tilt labels refer to the tilts of the electrodes in the image plane for images taken from the side and front viewports, respectively, and the total value is obtained from adding those angles in quadrature.

Electronics and Instrumentation

One important upgrade from the 2018 setup to the 2020 one was the use of a National Instruments PCIe-6376 DAQ with a 3.571 MHz sample rate, significantly slower than the 250 MHz sample rate in the prototype test. This was done so that zero-deadtime data could be acquired, giving insight into the evolution of light and charge production over the course of an entire ramp, up until the breakdown.

The slower sample rate would pose a challenge for directly digitizing the pulses from the new light sensors, the ONSem 60035C-Series 6x6 mm Silicon Photomultiplier (SiPM) and Hamamatsu model R9228 photomultiplier tube (PMT) pulses which are $\mathcal{O}(10\text{ ns})$. Because of this, the outputs were amplified with a Gaussian amplifier before digitization. The Cremat CR-200 Gaussian shaping amplifier was used to transform the (falling exponential profile) output of the CR-11X charge sensitive amplifier into a Gaussian profile of the desired width ($2.4\ \mu\text{s}$ full-width at half maximum). Separate amplifier chains were used for the SiPM and anode current. A considerable effort was taken to choose the appropriate charge-sensitive amplifier for the various instruments, as well as the coupling technique. The anode was AC coupled to a stock CR-150 test module with a CR-111 amplifier, through a surge protection circuit. For the SiPM, a CR-150 base was also used, but with a CR-113 amplifier as recommended by the Cremat specifications. The SiPM output was initially DC-coupled directly to the charge amplifier input during the first runs. Later, when the SiPM was moved from outside the outer viewport to the vacuum space, the SiPM was AC coupled. This means that the power supply is connected to the detector through the bias-in connection of the amplifier, and the signal is instead the current drawn *into* the sensor, rather than exiting it. Advantages of this coupling scheme included additional gain from lower input resistance, and only needing to route one cable into the XeBrA vacuum space.

For the PMT a CR-150 board was modified to be electrically identical to a CR-Z-PMT board, which is coupled through a larger capacitor ($1\ \mu\text{F}$ vs. $10\ \text{nF}$), and in parallel with a smaller resistor (1 vs. $200\ \text{M}\Omega$). In order to impedance match to the PMT board, a $50\ \Omega$ resistor from the PMT anode to the ground was replaced with a $1\ \text{M}\Omega$ resistor.

The XeBrA anode for both the 2018 and 2020 setups was held at a virtual ground which ran through a surge protection circuit. A spark gap to ground is connected to the anode cable near the apparatus itself, ensuring that large breakdowns short to ground. A signal splitter box provided additional protection and coupling to ground. This circuit consisted of $1\ \text{k}\Omega$ resistors in series, with TVS diodes to ground after each. A schematic of the full anode charge sensing scheme is shown in Fig. 8.9.

The Cremat CR-200 shaping module is equipped with a 10x gain switch, an inverting switch, a fine gain adjustment trimpot, and a trimpot for adjusting the pole-zero cancellation of the shaping circuit. It is necessary to tune the pole-zero resistor such that the time constant of the charge amplifier is cancelled out, leaving the Gaussian pulse of specified width instead. Under-cancellation (too much resistance) leads to the impulse response of the circuit keeping a long time constant, while over-cancellation (too little resistance) results in overshoot, with a decaying signal of opposite polarity to the initial pulse. The anode charge signal, SiPM, and PMT were all tuned using the “test-in” port of the Cremat CR-150 board.

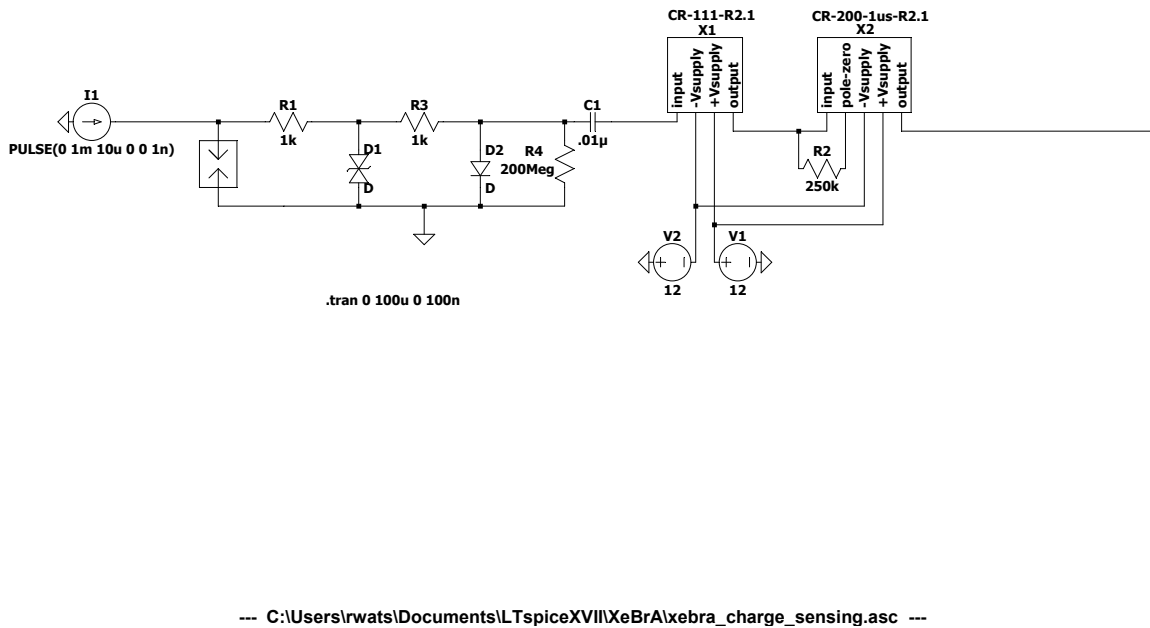


Figure 8.9: The schematic for the XeBrA charge sensing circuit, produced in LTSpice. The symbol on the left represents the spark gap.

This technique converts a square wave voltage signal (low impedance) to a delta function current signal (high impedance) by coupling the signal through a 1 pF capacitor. Results of these calibrations are shown in Table 1.

The anode pole-zero and gain were tuned and largely untouched for the duration of testing. The SiPM required more troubleshooting, and exact cancellation was hard to achieve due to parasitic capacitances and possible crosstalk. After several pole-zero tunings resulted in overshoot in-situ, the SiPM sensing scheme was tuned using an LED flash in a dark box. This results in large signals with the actual SiPM response, allowing for accurate tuning of the signal. For the PMT, an LED flash in-situ (installed in XeBrA) was used to tune the gains.

Calibrations were generally done between each run, as some experimentation was happening to maximize sensitivity. Occasionally the chain was altered slightly, splitting the anode (or SiPM) signal into a NiM differentiator circuit in order to trigger the high-speed camera

acquisition. This results in the gain being halved, and is accounted for in the conversion from DAQ to current.

Due to the high dark rate of the SiPM of 50 kHz/mm²[238] (translating to ~ 4 counts in the 2.4 μ s shaping time), a single photon threshold was not achievable during ramps. Single-photon sensitivity was achieved with the photomultiplier tube due to its lower dark rate. The PMT single-photon gain was obtained by pulsing a blue LED. A histogram of areas surrounding the LED pulse signal was obtained, and a double Gaussian plus constant function was fit to the data. The smaller Gaussian mean was identified as the single-photon gain.

Purity Monitor Repair

During the 2018 runs of XeBrA, the purity monitor became non-functional. It was believed that this may have been due to charring. One of my first tasks when reviving XeBrA for the 2020 runs was to disassemble and clean the individual components of the monitor. The monitor, built by then undergraduate student Glenn Richardson, consists internally of a series of six brass field shaping rings, a brass anode, a gold plated cathode, and two Frisch grids. These components are connected via three PEEK rods which run the length of the cylinder. The inter-ring distance is set by ceramic spacers. The spacers, along with the internal fiber optic cable which transmits the 177 nm photons to the cathode, were the main areas of concern.

Following careful disassembly, each individual ceramic spacer, and the PEEK rods, were inspected and cleaned. Some spacers exhibited charring, much of which washed away with isopropyl alcohol. Each 1 G Ω field cage resistor was similarly cleaned, and the resistances were verified individually. Following the reassembly of the internals of the purity monitor, the entire resistor-divider chain was verified with a multimeter / power supply combination. A new fiber optic cable was inserted, as well.

An area of concern which became apparent during these checks were the relative proximity of the screws which held the resistors in place. These screws all face the same direction, but are slightly proud in either direction. While the field shaping rings are 10.2 mm top face to top face, the gaps between the bottom of the screw tops and bottoms are much shorter. The shortest gap is the stage between the anode Frisch grid the anode, which creates a large, localized field enhancement.

A modification was done to the readout scheme of the purity monitor for the 2020 runs. As before, the cathode and anode signals were AC-coupled to a CR-111 charge amplifier, which measures the current running into those electrodes. The relative integrated charge of the cathode and anode signals provides an indication of the charge attenuation over the drift length. The previous iteration digitized the charge amplifier output with an oscilloscope, and fit the falling exponential signals to the wave form. In the new version, the charge amplifier output is instead fed into a CR-200 shaping amplifier, converting the impulse response into a Gaussian pulse. Further details of the analysis may be found in Section 8.7.




Figure 8.10: A photograph of the XeBrA purity monitor field cage. Note the slightly shorter inter-stage distance near the top, which creates a high voltage weak point. Plastic sheaths are present to protect the wires, as the cage has a small amount of clearance within the device.

High speed photography

High speed photography has been utilized in liquid argon breakdown experiments[239], which lent itself well to observation of the method of dielectric breakdown in those setups. Streamer-type breakdowns were observed, with an ionizing path being formed between the cathode and anode. Streamer mechanisms are a rich phenomenon, with positive (anode-initiated) streamers exhibiting a higher fractal dimension than negative (cathode-initiated) streamers [240]. Cell-phone photography was used in XeBrA-2018 to capture images of breakdown, but it was challenging to obtain detailed images of the development of breakdowns.

It was our goal with the XeBrA-2020 upgrades to look for possible warning signs of breakdowns, particularly bubbles preceding breakdowns. To that end, a pair of AOS Technologies PROMON u750 cameras were obtained, capable of capturing 2,146 frames per second. These cameras were held in place with a set of 3D printed clamps, designed by then-undergraduate student Eric Deck.

The procedure for these cameras was to acquire a rolling buffer of video prior during breakdown. When a breakdown occur, acquisition would trigger, writing a predefined duration of images before and after the image to disk. The total length was typically around 5 seconds, giving the operator ample time to activate the software trigger.

Post-processing analysis of these cameras was performed by Eric Deck and graduate student Jose Soria to obtain reconstructed breakdowns. It was hoped that the distribution of such breakdowns would reveal information about the nature of dielectric breakdowns in LXe. In particular we wanted to know if the breakdowns happened repeatedly in the same location, or had a preference for the regions of highest electric field. Examples of streamers seen in LXe are shown in Fig. 8.11, which all exhibit a mushroom like quality, and grow from the cathode to anode.




Figure 8.11: Three distinct examples of streamers observed using the high speed cameras installed in XeBrA. A bright spot is evident near the anode, while a columnar arc extends upwards from the cathode. These frames precede a frame where the camera is completely saturated.

My primary contribution to this component of XeBrA was the attempt to trigger automatically, freeing up person-power for other tasks. I attempted to do this by splitting the anode signal and routing the signal through a differentiation circuit which output a NIM signal. This NIM signal was then converted to a TTL signal, which was then routed to a linear fan-out. This fan-out signal then connected to each camera separately, which then triggered the acquisition and writing to disk.

Problems with the hardware trigger can in the form of unreliability. Frequently one camera would trigger, and not the other. Other times the cameras would trigger on larger precursors, while missing the visible breakdown a few seconds later because of the long dead time. Many attempts were made to rectify this, from replacing the NIM components, re-ordering the proprietary cables, de-straining the connections, and changing the input to the differentiation from the raw anode signal to the charge amplifier output. All of these methods proved capable of trigger on GAR breakdowns in a controlled environment, but LXe breakdowns proved finicky. Because each breakdown took between 2 and 5 minutes, losing most of the breakdown videos would severely impact the final analysis. For this reason manual triggering was the most common setup for these runs. I served numerous shifts waiting for the fault signal to arrive. Future upgrades may include coupling the camera hardware triggers directly to the fault signal emitted from the HVPS power supply.

8.4 Procedure

Main Procedure

The main procedure of XeBrA was carried over from the 2018 procedure [105], with some modifications. Each data taking run consisted of the following steps:

1. **Prepare the system.** All components which were previously opened are cleaned with isopropyl alcohol. The cathode feedthrough is affixed to the cone if it was not already. The high voltage cable is sealed. Using the method described in 8.3, the cathode is brought into contact with the anode, and the zero point of the linear shifter and the relative tilt are estimated. After the inner and outer volumes are connected, they are pumped out using the turbo pump on the top of the apparatus. Liquid nitrogen is obtained for the future recovery. All of the electronics are checked. If necessary, the gains on certain variable-gain components (such as the shaping amplifiers) are re-measured.
2. **Pre-cool the stainless steel** to liquid xenon condensation temperatures. In order to maintain a uniform temperature, and to prevent potential freezing of xenon on the pipes when it is introduced, a small amount of gas xenon is introduced. 2 bar absolute of GXe is introduced into the system while at room temperature. This is filled through the chemical scrubber to remove H₂O and O₂. After filling, the system is placed in circulation mode, where the path to the bottle is closed, and the circulation pump moves xenon through the system. In this mode, the GXe passes through the SAES getter, further purifying to ppb O₂-equivalent concentrations. At the same time, the cryocooler is activated, starting the process of cooling the stainless steel. The continuous circulation maintains the system temperature above 165 K. This process was typically started on a Friday, allowing an entire weekend (72 hours) for the steel to reach equilibrium. However, it took several attempts to find the correct circulation speed. As the pressure is low, the gas injected through the HX carries less of a heat load, which forces the resistive heaters to output more of their maximum power in order to maintain the heat balance. Initially the strategy was to lower the pressure to 1 bar absolute, which is above the xenon triple point of 0.8 bar and would prevent sublimation. However, the resistive heaters could not balance the heat load at these low pressures. This necessitated keeping the pressure higher, around 1.4 bar absolute. Eventually it was discovered that it was necessary to keep the heater tape on the gas inlet at a high temperature for this technique to work. When done correctly, when the filling procedure starts, condensation begins almost immediately.
3. **Condense the xenon.** This was usually performed on a Monday. The xenon bottle was opened, and the xenon was condensed through a path passing through either the chemical scrubber (for the initial runs) or the SAES getter (for the final run). The circulation pump was bypassed during filling. Care was taken to disallow the GXe pressure in the ICV from falling below the xenon triple point. This involved keeping

a close eye on PT05 and opening the regulator as necessary to keep the pressure sufficiently high. Though the latent heat of vaporization helps to balance the cooling power from the PTR, the pressure drop from the chemical scrubber leads to a reduced flow rate from the bottle to XeBrA. This necessitates additional heating from either the heating tape affixed to the inlet port, or the resistive heaters on the cold finger in the ICV. Over seven hours, 12 kg of xenon are transferred.

4. **Circulate and clean the xenon.** For 12-16 hours the xenon is circulated at a pressure of 1.4 bar absolute. The getter purifies the xenon during this time. Over this time the temperature of the steel reaches a new equilibrium following the changes which resulted from condensation.
5. **Take data.** Data in the most recent runs (4-10) were taken in circulation mode. Typically the first day of data taking involved opening up the xenon bottle to slowly condense the remaining 4kg of xenon, to a total weight of 16 kg. The linear shifter was moved into the desired separation, and new slow control files were started. The independent variables which were scanned over in each run were the xenon pressure, ramp rate, and separation. The surface condition was tested by swapping out the mechanically polished electrodes for passivated ones.

The high voltage was controlled with the LabView slow control. Each “ramp” consisted of three stages of progressively slower voltage ramps. The first two ramps exist simply to skip over low risk voltages faster, getting to the higher voltages we desire to test. The ramp rate reported in the data is the rate of the third and final ramp. When changing settings, the ramp schedule was tuned in order to prevent premature trips on the initial surge of current. A breakdown is detected by the HVPS, which tests if the current drawn exceeds a threshold, at which point it sends out a fault signal to the slow control. The threshold was similarly tuned to not trip on the initial surge. This typically resulted in a trip of 1-5 μA and was set as low as possible.

Some form of optical sensor was always active, barring malfunctions. The three possible optical sensors were the SiPM, the high speed cameras, and the PMT. For Runs 4-8, the SiPM and the cameras were alternated, as the SiPM was affixed on the outside of the outer viewport. In Runs 9-10 a new feedthrough was installed, enabling simultaneous acquisition of SiPM and camera data. The PMT was only used for static tests, and therefore was never active during a breakdown ramp. This data were digitized with a NI DAQ with a sample rate of 3.571 MHz with 100% livetime.

Periodically during data taking the optical sensors were disabled, in order to check the xenon for bubbles. Excess bubbling was never observed between the electrodes, only a slow wave pattern on the liquid surface.

6. **Recover the xenon** with cryopumping. The empty steel xenon bottle is immersed in liquid nitrogen, bringing the container to 77 K. The valve, at risk of becoming brittle, is kept hot with heater tape pressed tight against the steel bottle. A thermostat is used to prevent overheating. Circulation is stopped, the cryocooler is shut off, and the

return path to the xenon bottle is opened. The boil off heaters on the bottom cone are activate, providing an additional 60 W of heating power. This procedure usually only takes around 4-5 hours if done correctly. It was occasionally sped up by softening the outer vacuum with nitrogen. However, the risk of introducing condensation into the electronics was high, so this was only done on the occasions where the boil off heaters were malfunctioning. Near the end of the recovery, the pressure may become dangerously close to the triple point. Because of this, a hands on technique was employed, referred to internally as the “seesaw trick.” Essentially the valve to the bottle was opened and closed such that the pressure in the system oscillated between 1 bar and 1.2 bar absolute. This maintained the system above the triple point, making sure that no xenon was frozen. At the last cycle, if done correctly, the pressure drops from 1.2 bar to 0 rapidly when the valve is opened. If done incorrectly the pressure becomes stuck at the triple point, and the operator must wait for the ice to vaporize to continue. The vaporization is both time consuming and risky, as the ice may rapidly vaporize and damage the system or blow a burst disk.

Subcooling

An important modification to the procedure from 2018 was the introduction of a cooling cycle in between working days to mitigate bubble production. The original procedure involved filling through the getter, followed by 24 hours of circulation, then slowing circulation and continuously condensing the remaining xenon from the bottle in order to keep the system below the boiling curve. With the loss of xenon, less surplus xenon could be kept in the bottle, and therefore continuous condensation would be unsustainable. Furthermore, more system parameters were to be scanned over, so multiple days of data taking were necessary for each cycle of filling and recovery. Thus, another solution was necessary, which was discovered somewhat accidentally.

The subcooling procedure is conceptually simple. After the data were taken for the day, the system set point pressure was lowered to a value below any which we intended to examine. When this happened, the heating tape thermostat was also lowered. This resulted in a generalized cooling of the system, though the HX temperature rose due to the xenon changing phases. The circulation speed was forced to compensate for the lower pressure of the xenon in order to maintain heat balance with the PTR cooling power. Over several hours the cone steel cooled, approaching the condensation curve. Ideally this would be maintained long enough to achieve system equilibrium, but typically the cone bottom temperature was still falling the next morning.

The next morning (between 8 and 12 hours later), the system pressure was raised, along with the heater tape thermostat increasing. This boiled off xenon at the top of the detector rapidly, raising the temperature at the top of the apparatus without having it translate to the cone steel immediately. When the set point pressure was achieved, the heater tape and circulation speed was adjusted to maintain stability. The steel ICV was at this point in a temperature inversion: the hottest portions were at the top, while the coldest portions were at the bottom, near the feedthrough. Ideally, the cone bottom would be the boiling




Figure 8.12: An example time series of measurements during of a subcooling cycle in Run 7. Pressure transducers and thermometers are shown. The primary result of this method is an increase in the separation of the cone bottom and Kimbal sphere RTDs.

temperature of xenon at the overnight pressure, while the liquid level would be the boiling point at the current pressure.

With the cathode feedthrough 10s of Kelvin below the boiling point, bubbles could be suppressed for longer than an entire workday. Periodically the viewports were checked to make sure that the xenon was placid. After subcooling became part of the procedure, the xenon was always placid except for when the pressure was being altered as part of a scan (lowering the setpoint results in boiling off of xenon, carrying heat to other portions of the apparatus).

An example of the slow control data from a subcooling cycle is shown in Fig 8.12. It took some amount of trial and error to find the correct settings for subcooling. Because the PTR has a fixed cooling power, heat must be injected from the resistive heaters and gas xenon in order to achieve balance. An issue with the system is that it is naturally unstable, tending towards runaway boiling or condensation. This is because the gas xenon being injected through the heat exchanger constitutes a significant portion of the heat load. When the gas pressure raises, this leads to a higher density of room-temperature xenon being pushed into the HX, which increases the heat load, which in turn boils off more xenon, raising the gas pressure. The inverse cycle is also possible, and is actually more likely to occur due to the inability to throttle the cryocooler. The resistive heaters are controlled via the PLC's PID loop in order to achieve the correct heat load to maintain the pressure, and are typically successful in doing so if the equilibrium power input is near 50% the capacity.

8.5 Run History

Introduction

The XeBrA experiment is a research and development test stand. Many runs were taken in the 2020 iteration, partially due to some trial and error of the procedures and experimental

goals. A brief history of XeBrA is presented here, along with my contributions to each.

XeBrA-2018

Runs

XeBrA was constructed between 2016 and 2018. While I was not heavily involved in the construction process, my primary contributions in this stage were in machining of parts and the fast DAQ setup and analysis. I machined the thermal link to the bottom of the cathode cone out of oxygen-free-copper. This kept the temperature of the steel relatively uniform. Later on, I set up the CAEN digitizer to analyze anode data. While not critical to the analysis for the 2018 analysis, it did form the motivation to focus more heavily on precursor pulses for XeBrA-2020

1. Runs 0-1: liquid argon runs. These were performed in order to troubleshoot elements of the procedure in a safe environment, as in case of emergency argon could be thrown away with impunity.
2. Run 2: A complete liquid xenon run. A method of suppressing bubbles by constantly condensing was utilized. It was realized that the xenon was heavily contaminated ($\mathcal{O}(\text{ppm})$ O_2 equivalent), based on RGA measurements. This led to the saturation of the getter cartridge in XeBrA, which required replacement.
3. Run 3: A partial liquid xenon run. During data taking, a weld in the heat exchanger failed, leading to a loss of xenon. Before this, due to filling through in-line purifiers, the xenon froze instead of condensed. Rapid changes in pressure as a result most likely put additional stress on the weld.

TPC

In order to estimate the liquid purity during Run 3, XeBrA was temporarily turned into a time-projection chamber. In this mode, the electrodes were separated by the maximum possible amount (20 mm) and liquid xenon was filled up to 5 mm below the anode. The Hamamatsu PMT collected information on the electroluminescence in the chamber. The CAEN DAQ was set up to trigger on the S2s. In order to trigger events, thoriated welding rods were affixed to the outside of the outer vessel. The Compton scatters from assorted gamma rays were observed to drastically increase the trigger rate.

Due to not being designed for this mode, the light collection efficiency is unknown. Instead of attempting accurate per-event reconstruction, the intention was to obtain the S2/S1 ratio as a function of the drift time. As a result of the anode being placed far above the intended height, the field was less uniform than during the breakdown data taking. This causes the S2 width to vary as a function of position at the liquid surface.

The function $S2 = S1g \exp(-dT/\tau_e)$, where $S(i)$ is the S1(S2) pulse area, dT is the drift time, g is the average gain, and τ_e is the electron lifetime, was fit to the S2-triggered data.

An electron lifetime $\tau_e = 2.21 \pm 0.02 \mu\text{s}$ was obtained. From this, the O_2 -equivalent impurity was estimated as ~ 200 ppb. I performed the analysis of this pseudo-TPC data, extracting the lifetime and estimated the uncertainty using jackknife resampling.

XeBrA-2020

All breakdown data in XeBrA-2020 were taken in liquid xenon, though some preliminary electronics testing took place in gas xenon and gas argon.

1. **Run 4:** The first XeBrA-2020 run. Data were collected up to 6 mm separation, more than other runs. Relatively straightforward compared to future runs. No SiPM data was taken, and the high speed cameras were tested out.
2. **Run 5:** A failure in filling, resulting in an overpressure event. As it happened previously in XeBrA-2018, the xenon was filled through the in-line purifier, which resulted in some sublimation on the steel. This resulted in some frozen xenon in the system, likely the HX. While attempting to melt the ice, a piece rapidly vaporized, causing a burst disk to break and xenon to be vented into the recovery vessel. This occurred relatively early into the fill, and therefore filling continued. Some data was taken after this, though the reduced quantity of xenon lead to thermal instabilities and bubbles. The xenon was successfully recovered early, with only $\mathcal{O}(100)\text{g}$ loss of xenon. Afterwards the purity estimated after the fact with a cold trap-RGA method similar to that done in Ref. [241]. The purity was determined to be <100 ppb. This run accidentally lead to the development of the subcooling procedure detailed in Section 8.4.
3. **Run 6:** Another aborted run. This time, while beginning the subcooling procedure, which involved slowing the flow rate, the inlet pipe heated up considerably. This is believed to be a result of the reduced cooling power going into the metal, which at this point was not regulated by a thermostat. This hot metal ended up melting the plastic tube which connects the turbopump to its backing scroll pump. This spoiled the vacuum in the jacket, which while not catastrophic did result in an early, controlled recovery of the Xenon back into the bottle. No xenon was contaminated during this process. Some data were collected prior to this incident. The plastic tube was later replaced.
4. **Run 7:** A successful xenon run. This was the first run where pressure scan data were taken. SiPM and camera data were taken as well, but not simultaneously.
5. **Run 8:** A xenon run with alternate electrodes. Following the conclusion of Run 7 the electrodes were inspected. Between certain runs, the electrodes were cleaned with isopropyl alcohol in order to remove what was thought to be debris. It was later realized that pitting was occurring as a result of very large breakdowns. The electrodes were inspected under a microscope, revealing populations of large and small pits. An alternate set of mechanically polished electrodes were used for this run, which was otherwise uneventful. The original electrodes were sent off to be re-polished.

6. **Run 9:** The original electrodes were swapped back in and another run was done. This was the first run where ramp speed scan was performed, along with a pressure scan. Additionally, the SiPM was now placed inside the outer vacuum, affixed to the ICV viewport.
7. **Run 10:** A successful run, identical to Run 9, but with passivated electrodes instead.

8.6 Analysis

Breakdown Field

Selection Criteria

At each separation and run dataset, the Weibull parameters described in 8.2 were extracted. The breakdowns were initially filtered to remove confounding data. The selection criteria were:

1. Breakdown voltage must be larger than 4 kV due to occasional spurious trips on the ramp start;
2. Pressure must be within 50 mbar of the set point;
3. The first ten breakdowns of any dataset are vetoed to mitigate possible conditioning effects;
4. The breakdown must occur at least 0.1 s away from a step in voltage, which ensures that we only consider data during periods of quasi-static cathode voltage.

Criterion 1 was chosen due to the fact that the ramp control software jumps from 0 V to 800 V in one step at the beginning of a ramp. This appears to be a setting within the HVPS itself and could not be altered. An instantaneous jump in current of this size had a tendency of tripping the fault. With the assistance of the viewports, it was evident that these events were not breakdowns, as no spark was seen between the electrodes. Thus, this criterion was introduced to make sure that these breakdowns were not counted.

Criterion 2 is fairly simple. The xenon pressure in the spark chamber is an independent variable of the tests, and therefore random excursions from the set point must be excluded. This is usually only an issue on the first day of testing, where the additional 4kg xenon is condensed, which leads to challenges for the PID loop to control the pressure.

Criterion 3 exists to remove the possible conditioning effect of the electrodes. Fresh electrodes have the potential for sharp asperities on their surface. These asperities enhance the local electric field, increasing the chance for breakdowns at lower bulk fields. However, dielectric breakdowns are violent events, and can melt the metal, smoothing out the asperities. In fact, this is the process by which electropolishing[242] occurs. Due to the fact that only a small (~ 100) number of breakdowns are taken with each configuration, only 10 breakdowns are

vetoed. This is repeated for each configuration, rather than each replacement of electrodes, in case new asperities have been introduced.

Criterion 4 was introduced to be insensitive to the effect of the current surge from the discrete steps in voltage. While a ramp rate is specified in software, the voltage is not adjusted continuously, but rather in 100V increments (the smallest that the HVPS can control). The ramp rate, therefore, only controls the time between these discrete steps. In order to test the quasi-static DC breakdown conditions, rather than these AC current surges, breakdowns which occur within one slow control sample (10 Hz = 100 ms) of the step are removed. The distribution of the length of time at the breakdown voltage is given in Fig. 8.13.

Weibull fit

The breakdown fields from the linear ramps were compared against the Weibull distribution. The cumulative distribution of breakdown fields can be transformed in such a way to ease the interpretation. A two-parameter Weibull cdf $F(E) = 1 - \exp[-(E/E_0)^k]$ can be mapped in such a way as to appear linear:

$$\log(-\log(1 - F(E))) = \log(H(E)) = k \log(E) - k \log(E_0) . \quad (8.25)$$

A distribution of breakdowns will appear linear in this space, with the slope providing the shape parameter, and the y-intercept providing the scale parameter. Introducing a location parameter causes the distribution to have nonzero curvature in this space.




Figure 8.13: The distribution of the time differences between breakdowns and their associated discrete voltage steps. Events in the first bin (between 0 and 0.1s) are removed by criterion 4. This data is the combination of all runs in XeBrA-2018 and XeBrA-2020

The survival function $S(E)$ must be estimated from the finite datasets available. There exist several methods for doing so. The empirical distribution function is the most straightforward, as is given by the number of entries less than the argument:

$$\hat{F}_e(E) = \frac{1}{N} \sum_i^N \theta(E - E_i) . \quad (8.26)$$

Another popular method is the “median rank” method[243]. Given a series of breakdown fields, they are ranked in ascending order such that E_1 is the smallest and E_N is the largest. The estimated cumulative distribution is then:

$$\hat{F}_{MR}(E_i) = \frac{i - 0.3}{n + 0.4} . \quad (8.27)$$

This method has the advantage of being well-defined at the boundaries, but if values are repeated then it requires some corrections. The method used in this analysis for diagnostic purposes is the Kaplan-Meier estimator[244].

$$\hat{S}_{KM}(E) = \prod_{i:E_i \leq E} \left(1 - \frac{d_i}{n_i}\right) , \quad (8.28)$$

where the parameters d_i refers to the number of breakdowns observed at point E_i , and n_i is the number of ramps known to be active up to point E_i . The distinction is interesting, as the Kaplan-Meier estimator is frequently used in the medical context, where individuals may drop out of trials at some point. In XeBrA the distinction is less crucial, as we always observe a breakdown for a ramp.

It was noticed in XeBrA that the Weibull plots were often a straight line, indicating the validity of the model. However, this was not always the case, with some datasets demonstrating one or more “knees” in the plot. In reliability analysis this indicates the presence of several failure modes. Two ways to combine hazards are in a simultaneous or transient manner. A simultaneous combination indicates a situation where either failure mode can appear at the same time, and the effective hazard function is given by the sum of the individual hazards. On the other hand, a transient failure mode would appear as a sum of the respective probability distribution functions (pdfs). An example of the simultaneous failure modes is shown in Fig. 8.14

The simultaneous combination was not used for this analysis, but rather the linear combination of the pdfs. This choice was in part due to the increased flexibility of being able to weight the two contributions. More critically, it was motivated by the apparent auto-correlation of the breakdown voltages within datasets. Rather than having the relatively high or low breakdown voltages distributed throughout the dataset, the breakdowns tended to be clustered within their own modality. An example of this is shown for dataset 701 in Figure 8.15. These clusters are referred to as “transient dips” in breakdown field.

Transient dips are not consistently associated with any observable change in thermodynamic conditions (e.g. temperature, pressure, flow rate, etc). It is my hypothesis that the




Figure 8.14: An analytical example of combining two hazard functions. Here $H(x) = H_1(x) + H_2(x)$. The Weibull parameters are $\lambda_1 = 10$, $\mu_1 = 5$, $k_1 = .5$, and $\lambda_2 = 20$, $\mu_2 = 0$, and $k_2 = 5$.

dips are due to debris, or persistent bubbles, floating into the stressed volume and weakening the field over multiple breakdowns before dissipating. Debris could cause such weakening through the Malter effect[214]. The data is fit by maximizing the likelihood over the following function:

$$\ln \mathcal{L} = \sum_i^N \sum_j^M a_j \frac{k_j}{\lambda_j} \left(\frac{E_i - \mu_j}{\lambda_j} \right)^{k_j - 1} \exp \left[- \left(\frac{E_i - \mu_j}{\lambda_j} \right)^{k_j} \right], \quad (8.29)$$

where the i runs over the N data points and j runs over the M components. The parameters E_0 , E_1 are replaced with λ and μ , respectively in order to not overload the purpose of the superscript in this case. For this analysis, only $M = 2$ was considered.

Since not every dataset appeared to contain multiple components, $M = 1$ models were initially considered for them. A likelihood ratio test was performed, which analyzed the relative increase in the likelihood when adding the additional degrees of freedom to the function. The p-value was calculated based on the χ^2 distribution with degrees of freedom equal to the change in free function parameters. Two three component Weibull functions with a fractional contribution contains 7 d.o.f, making the likelihood ratio distributed according to

$$\ln \frac{\mathcal{L}_1}{\mathcal{L}_2} = \sum_i \ln [a f_{W1}(E_i) + (1 - a) f_{W2}(E_i)] - \ln f_{W0}(E_i) \sim \chi_4^2, \quad (8.30)$$

where the sum runs over points, $f_{W1(2)}$ are the probability distribution functions of the components of the double Weibull model, and f_{W0} is the three-parameter, single component model. When this p-value was below 0.05, the two-component model was chosen. This procedure was also used to select between the two- and three-parameter Weibull functions, i.e. with $E_1 = 0$ or floating. When the two-component model was selected based on the




Figure 8.15: High voltage power supply voltage vs. time, with the identified breakdowns for a representative dataset(701). The red stars indicate the breakdowns and serve as a diagnostic for the breakdown finder code. The lower breakdown fields are seemingly clustered together.

p-value, it was only approved in software if the size of the smaller component, i.e. $\min(a_j)$, was greater than 0.1. This was to ensure that there were sufficient breakdowns to constrain the parameters of that component. In the case that the two-component, three parameter model was selected as the superior fit, this left the question of which component to compare against the single-component models. I decided to use the component with the larger modal value, for the simple reason that a transient dip in performance is a more physically interpretable situation than a transient increase in performance. An example of a model with two components preferred is shown in Fig. 8.16

The likelihood minimization was performed with `scipy`'s `minimize` function, utilizing the sequential least squares programming(SLSQP) algorithm. The derivative of the Weibull log-likelihood was specified to ensure rapid convergence. A hessian(second-derivative) matrix was used to estimate the error on the fit parameters for the two-parameter Weibull function, as the covariance matrix is given by the inverse hessian. For the three-parameter(one- and two-component) models, instead of an analytic function the covariance matrix is estimated using jackknife resampling[245].

Generally, the uncertainties were $\mathcal{O}(0.1)$ - $\mathcal{O}(1)$ for the scale E_0 and shape E_1 parameters. This is related to the statistical fact that the moments of the Weibull function (mean, variance, skew) are not functions of only one of the Weibull parameters, but generally all three. Because of this, the off-diagonal elements of the correlation matrix are large. Instead of reporting E_0 or E_1 separately for inter-dataset comparisons, I chose to introduce a new




Figure 8.16: An example of a breakdown distribution from a dataset that strongly favors a two component, three-parameter Weibull PDF (i.e. Model 3, see text). This is for Run 4 and electrode separation 1 mm. *Top*: Histogram of breakdown fields (blue) and two-parameter Weibull PDF overlaid on top. *Bottom*: The cumulative hazard distribution (black) and the model that provided the best fit to data (orange dash line). E_2 is equal to the shifted scale parameter $E_0 + E_1$, corresponding to the 63rd quantile of the distribution, and is plotted here for reference (black dash). The cumulative hazard can be interpreted as the expected number of breakdowns that would occur before reaching a given breakdown field, if ramps were restarted at the point of failure.

parameter,

$$E_2 \equiv E_0 + E_1 \quad (8.31)$$

and perform all comparisons with it instead. Other quantities such as the mean and median of the Weibull distribution were explored, but E_2 had the advantage of relatively simple error propagation. Statistically, E_2 is the 63rd percentile of the Weibull distribution (corresponding to a cumulative probability of $1 - 1/e$).

Charge and Light Analysis

The light and charge data from the SiPM and anode current were analyzed using a similar technique to the LZ-CHV prototype data in Chapter 3. Low frequency noise is removed, followed by removal of chirp signals which were present. Then, pulses are located using two passes, one for shorter and one for longer pulses. Following large pulses in the SiPM, particularly following breakdowns, the pulsefinder is disabled until the system recovers from the shock. Pulses then have their associated charge (and photon count) reconstructed. Some selection criterion are then applied to remove spurious pulses. The current from each channel is then associated with breakdowns and given a time from the start and end of a breakdown.

Each acquisition is divided into “events” of length 10^6 samples per channel by the Polaris software[246]. As the clocks on the DAQ computer and the slow control are not perfectly synchronized, and the sampling rates were different, the breakdown locations were found from the DAQ data, agnostic of the slow control/HVPS data. The breakdowns were located by looking within each event for the overshoots from the massive surge of current from a breakdown. A positive pulse of current into the anode was observed coincident with HVPS fault trips. It is also occasionally associated with distortions of the cathode voltage which *do not* result in fault trips. These “sub-threshold breakdowns”/“glitches” or trips were disambiguated from true breakdowns(defined as a signal which *does* result in a trip) by the presence of three or more overshoots in a single event window. In Fig. 8.17 an assortment of event classifications are shown, from an empty event, to an event containing pulses, to a “glitch” event, to a “breakdown” event. Anode and SiPM pulses within two seconds of a glitch do not count towards future analysis.

“Chirp” signals in the anode, similar in nature to those found in the CHV prototype testing were filtered out using a sliding window method. The root-mean-square (RMS) value within the sliding window of length $N = 900$ samples was calculated, and the maximum amplitude of the signal within the same window was compared against it. If the RMS is > 20 analog-to-digital converter counts (ADCC) for more than $L = 250$ samples, and the ratio of amplitude to variance is < 0.9 , then the region of the waveform is replaced with a second-order lowpass filter with cutoff frequency of $\omega_c = 2\pi/4096$. A representative waveform which has been “dechirped” is shown in Fig. 8.18.

The pulsefinder is based around a difference-of-Gaussians (DoG) filter, which is an approximation of the Laplacian operator. The filter is a small standard deviation, positive polarity normalized Gaussian function, from which a wider, normalized Gaussian function is subtracted. This effectively performs a low pass filter to subtract slower trends and baseline

Model	C [kV/cm]	b	ρ_{Cb}
M. Polished (2020)	169.5 ± 3	0.262 ± 0.006	-0.995
Passivated (2020)	530 ± 28	0.667 ± 0.02	-0.920
M. Polished (2018)	171.5 ± 8	0.13 ± 0.02	-0.995

Table 8.2: XeBrA power law fit parameters to $E_2(A) = C(A/\text{cm}^2)^n$. C is a multiplicative constant, and $-b$ is the coefficient of the SEA in the power-law scaling.

drifts, then identifies regions of significant amplitude on top of these trends.

$$\log(t; \sigma_1, \sigma_2) = \frac{1}{\sigma_1\sqrt{2\pi}} \exp\left(-\frac{t^2}{2\sigma_1^2}\right) - \frac{1}{\sigma_2\sqrt{2\pi}} \exp\left(-\frac{t^2}{2\sigma_2^2}\right) \quad (8.32)$$

Local maxima of the DoG-filtered waveform were identified as pulses if they were above a given threshold. Reconstructed quantities (RQs) of the pulses were calculated, such as area, amplitude, width, full-width-at-half-max (fwhm), and rise time. In order to further remove chirps, pulses with $\text{fwhm} < 8.4 \mu\text{s}$ were vetoed. As the charge shaping circuits were frequently tuned between runs, the particular pulsefinder parameters, thresholds and gains also varied. The σ_1, σ_2 of the DoG filter were tuned by identifying large pulses in each channel, and scanning over each parameter to find a global maximum of the signal-to-noise ratio (amplitude over rms). While the values differed from dataset to dataset, the values typically were in the range $\sigma_1 \in [8, 12]$ samples, $\sigma_2 \in [30, 40]$ samples, and $H_{min} \in [100, 500]$ ADCC.

8.7 Results

Breakdown Risks

Stressed Electrode Area

The core analysis of XeBrA revolves around the stressed electrode area (SEA) scaling relationship. Each run consisted of datasets taken at several separation distances, which correspond to different SEA. Due to the varying cathode tilt between certain runs, datasets with the same nominal gap have different areas. Though this creates issues for replication of experimental results, it does lead to additional leverage for the eventual fit. The $E_2 \equiv E_0 + E_1$ values were compared between datasets. In the case of a two-component model (Model 3) being preferred, the larger component was selected. The results of this examination are shown in Fig. 8.19. A power law is fit to the mechanically polished and passivated electrode data separately, producing the parameters in Table 8.2.

As justification for the choice of selecting the larger Weibull mode, I also examined the alternative choice. For the single-component models nothing changed, but for the two-component models I calculated the E_2 values there, as well. The choices of smaller or larger

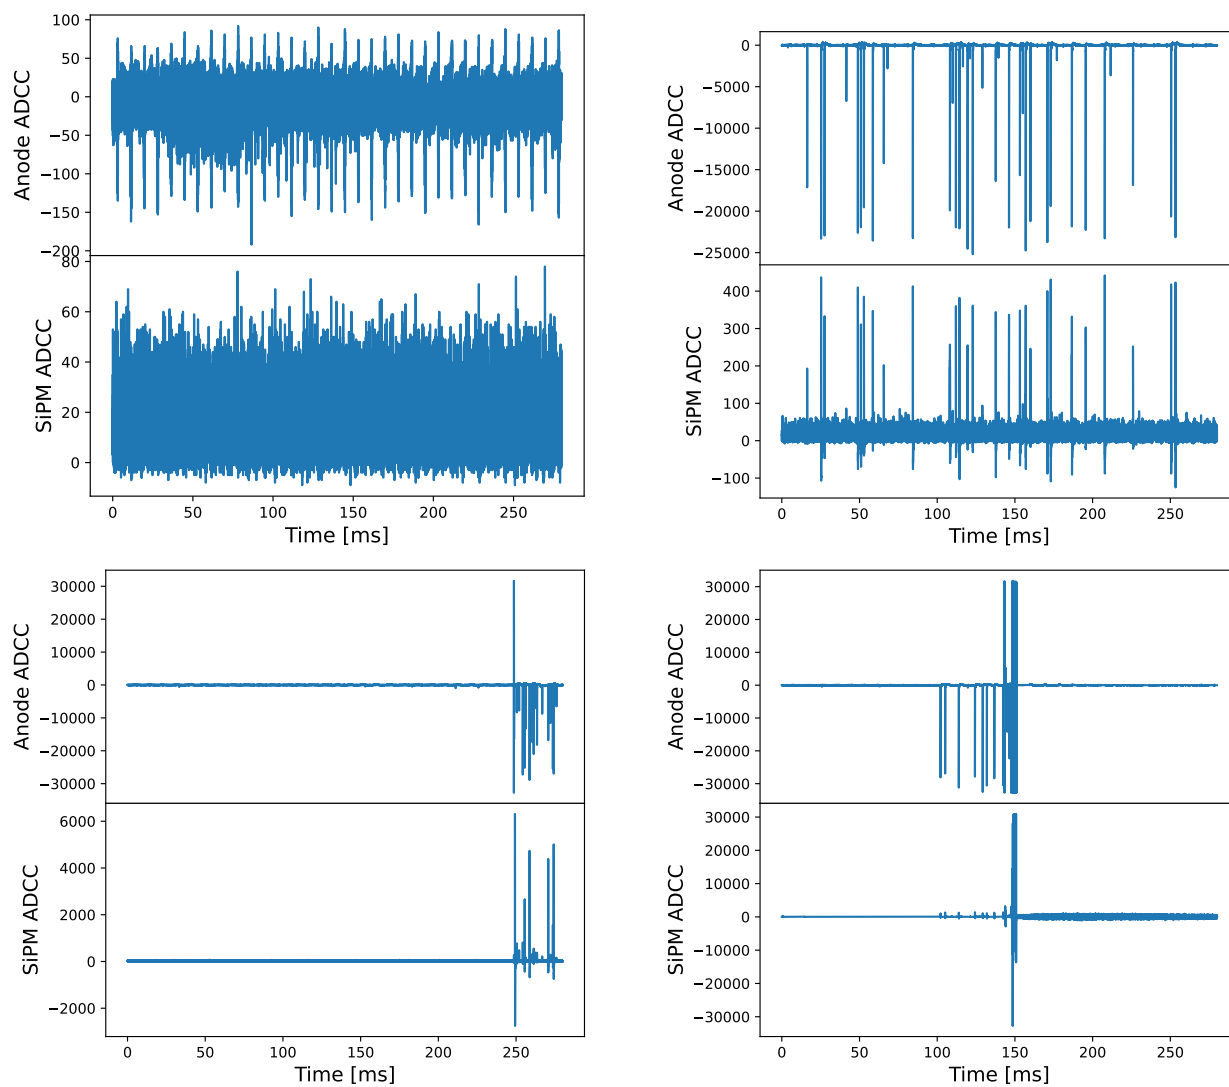


Figure 8.17: Various categories of event acquisitions waveforms. *Top left:* An example of background noise. The oscillatory behaviour is removed through a filtering technique. *Top right:* an example of the coincidence between anode and SiPM pulses. This event was recorded close to the breakdown point. *Bottom left:* A subthreshold discharge. These occasionally occur during a ramp, but did not trigger an HVPS fault signal. *Bottom right:* A true breakdown event, coincident with an HVPS fault signal.




Figure 8.18: The results of the dechirping procedure on a typical anode signal event waveform. *Top*: before the dechirping, with a low frequency component and regular bursts. *Bottom*: after the dechirping, with the chirps replaced with the results of a lowpass filter.




Figure 8.19: The SEA scaling from XeBrA. XeBra-2020 data is divided into data for passivated and non-passivated electrodes. Points are shown for the XeBrA 2018 data as well.

mode lead to R^2 values of 0.240 and 0.860, respectively. Using the F-statistic[247], a p-value of 0.01 was calculated, resulting in rejection of the null hypothesis that the choices have identical explanatory power.

The effect of passivation can also be seen in the data. Passivated configurations consistently demonstrate larger breakdown fields in XeBrA, but with a more severe SEA scaling dependence than the mechanically polished electrodes. It is unknown whether this effect will continue indefinitely, so I chose not to extrapolate this pattern to a general recommendation.

Pressure

No significant relationship between pressure as measured by PT05 and breakdown field was observed. Changing the pressure occasionally lead to wild swings in electric field strength, but this effect was seemingly unrelated to the sign of the change. That is, moving the pressure higher or lower was always associated with larger breakdown fields.

While not definitive, it is interesting that Run 8 was not observed to have as great a change in the breakdown fields for its pressure scan. This Run used a different set of mechanically polished electrodes than the rest. This makes it challenging to tell whether the change is truly from a change in procedure. However, in Run 8, the pressure scan was performed over the course of a single day, and the pressure was monotonically increased. It was the only run to do so, with the remaining runs starting the pressure scans later in the day, and usually moving up and down in pressure over two working days. It is possible that the choice of performing the pressure scans for Runs 4-7, 9-10 in this way disrupted the system thermodynamics in an unpredictable manner, negating the benefits of subcooling (see Sec 8.4).

Ramp Rate

The ramp rate was scanned over in a similar fashion to pressure. Each high voltage ramp consists of three stages of progressively slower rates. The schedule was prepared such that breakdowns at any point other than the third and final rate were extremely unlikely. Comparisons were made over scans of the terminal rates, selected as 100, 150, and 200 V/s. All other data present use ramps of 100 V/s. Only a small amount of data were taken with these tests, unfortunately resulting in inconclusive results. Passivated electrodes demonstrate a positive correlation of breakdown field with ramp speed, as expected from theory. Mechanically polished electrodes did not have a significant dependence. The results are shown in Fig. 8.21. A linear regression on the passivated and mechanically polished electrodes yields slopes of 2.1 ± 0.2 and 0 ± 4 minutes/cm, respectively.

Xenon Purity

The concentration of oxygen-equivalent impurities contained in the LXe were measured by an in-situ purity monitor. The monitor measures the electron lifetime as they drift between a gold-plated cathode and a brass anode. The electrons are generated via the photoelectric effect by a xenon flash lamp, with UV photons guided to the cathode by a fiber optic cable.




Figure 8.20: XeBrA pressure scan data. The data is presented per-run to facilitate the intra-run comparisons, without the effect of differing cathode tilt.




Figure 8.21: XeBrA estimated breakdown voltages vs. ramp rate scan results.

Run	Electron lifetime [μs]	O ₂ -equivalent impurity [ppb]
4	19.4480 ± 0.0007	23.3470 ± 0.0008
5-7	-	<100
8	34.1 ± 0.1	13.34 ± 0.05
9	77 ± 2	5.9 ± 0.1
10	125 ± 5	3.6 ± 0.2

Table 8.3: Purity measurements taken during the 2020–21 data-taking runs. A general trend of lower purities are observed between consecutive runs. For Runs 5–7 the purity was inferred from a cold trap method similar to the one described in Ref. [241]. The discrepancy in uncertainties is due to the averaging of the waveforms and inconsistent flash rates. In certain datasets for unknown reasons the purity monitor was able to maintain bias for longer than others, allowing for longer integration periods. This introduces a bias towards shorter lifetimes / higher impurities due to the capacitor in series with the anode requiring recharging.

Frisch grids allow for the cathode signal to be distinguished from the anode signal. The pulses are read out by a Cremat CR-111 charge amplifier and CR-200-1 μs shaping amplifier, generating easily distinguishable Gaussian pulses. Data are digitized using an oscilloscope and analyzed offline. Then, the electron lifetime is found by filtering the voltage signal with a boxcar filter and taking the ratio of the corresponding pulse heights. This method also yields highly accurate timing information for the pulses themselves. The ratio is converted into lifetime as

$$\tau = \frac{t}{\log(A_C/A_A)}, \quad (8.33)$$

where t is drift time, A_C is the cathode pulse area and A_A is the anode pulse area. The oxygen-equivalent impurity is calculated utilizing the attachment coefficients found in Ref. [61], which results in the conversion formula

$$O_2 = \frac{455 \text{ ppb} \cdot \mu s}{\tau}. \quad (8.34)$$

The oxygen-equivalent impurity measurements for the 2020–21 data-taking runs are shown in 8.3. During Run 5, an overpressure incident led to xenon mixing with residual gas in one of the recovery vessels. The purity of the xenon after this incident was estimated using a sampling system and it was found to be less than 100 ppb. Run 6 was aborted before a purity measurement could be taken and Run 7 suffered a data corruption issue.

The uncertainties on the purity monitor pulse areas are found by examining the pre-cathode pulse signal RMS, while the uncertainty on the drift time is taken to be twice the sampling period of the oscilloscope. The overall results are tabulated in 8.3. The discrepancy in uncertainties is primarily due to inconsistencies in the number of pulses the oscilloscope averaged over.

FN plots

Dielectric breakdown in noble liquids have been analyzed through the lens of field emission for LHe[230]. Fowler-Nordheim field emission has also been examined on wire grids for LXe[109]. Field emission may provide the Joule heating necessary for gas bubbles to grow until either superheating or percolation occurs[248]. Due to the highly variable precursor current rates, a definitive current vs. voltage could not be obtained. A picoammeter was used in XeBrA-2018 to measure DC current, but this piece of equipment was requisitioned for use in LZ cathode high voltage testing and could not be utilized for XeBrA-2020. The anode pulse rates correlated better with the time until breakdown, rather than the voltage themselves.

It was attempted to analyze the breakdown hazard data similar to the way it was done in Ref. [230]. In short, the cumulative hazard was taken to be proportional to the FN current $H(E) \propto J(E)$, with $J(E)$ given by equation 8.7. Due to the multiple modes observed in the Weibull analysis, the Fowler-Nordheim plots are also curved. FN current form straight lines of negative slope in a plot of $\log(I/E^2)$ vs. $1/E$, from which the values of emitter area and β , the field enhancement, could be obtained. Since the exact proportionality between $H(E)$ and $J(E)$ is not precisely known, the emitter area can not be extracted this way. However, knowing the work function of steel to LXe allows one to unambiguously obtain β from the slope alone.

Due to the multiple components, the FN model was fit to the quasi-linear region at high field. This region was detected automatically by calculating the $\chi^2/NDOF$ as points are progressively added in descending order of E . When additional points no longer reduce the fit quality, the data is cut off, and the slope of the line is reported. Analyzing and extracting the β over many data sets reveals a slight correlation $\rho = 0.38$ with SEA across all runs, and a slightly higher correlation $\rho = 0.51$. Significant intra-run variation of β was observed. From this I conclude that, if the FN model can be utilized, that the asperities can be conditioned away over the course of a run. Additionally, sharper asperities are more likely to be found when integrating over larger SEA. This analysis was performed for completeness, and the Weibull fit forms the core of the results.

Pre-Breakdown phenomena

Charge Data

Since the observation of precursors like those in Fig 8.24 were seen in XeBrA-2018 data, it became a secondary goal of XeBrA-2020 to determine if any warning signs of breakdown could be identified. Following the procedures laid out in Section 8.6, data were collected for the assorted runs. DAQ saturated frequently occurred in the anode channel, even for precursor events, though the SiPM channel generally only saturated during actual (fault-triggering) breakdowns. Therefore it is sometimes necessary to report rates in terms of physical charge, and compare against the rate of pulses, as if the sensor was a Geiger counter.

The goals of the analysis were to identify patterns of increasing charge and light in the lead up to breakdowns, and to identify whether pre-breakdown current is correlated with any




Figure 8.22: Fowler-Nordheim plots (cumulative hazard divided by squared field magnitude, vs. inverse field magnitude) from Run 7 data. The dotted black lines indicate the regions over which the linear regression χ^2 / NDOF decreases with additional points.

particular features. Typically, an increase in anode pulse rate of around a factor of 100 over the background rate was observed in the preceding 60 seconds to a breakdown, as shown in Fig. 8.25. The lower detection threshold of the SiPM pulses lead to the rise in pulse rate only being visible in the final few ms. The pulse height spectrum was bimodal for the anode pulses, as shown for a representative dataset 901 in Fig. 8.27.

The anode current pulse rate was generally higher in runs with larger SEA, while pressure and ramp speed had no significant effect on the time profile of precursor pulses. The precursor rate within any particular ramp did not rise continuously (at least until the last few seconds), but instead appeared to spike in rate several times preceding a breakdown (see an example in 8.23). Overall, a consistent rise in anode current before breakdown was observed. For instance, 8.25 shows the average current (as measured by the charge amplifier) over time bins of 1s for the different parameter scans in Run 10. These plots show a first surge in current around 50–60s and a rapid increase in current in the last few seconds.

Comparing the polished and the passivated electrodes, we observed that passivation only has a small effect on the preceding activity before a breakdown. As 8.26 shows, there is only a moderate suppression in both the anode current and SiPM rate in the immediate few seconds before a breakdown when comparing Run 9 (mechanically polished electrodes) to Run 10 (passivated electrodes).

The evolution of the average current in a ramp is shown in 8.28 for Runs 9 and 10. The current into the anode was averaged over a time period well separated from either the

breakdown itself or the initial ramp. The region of 45 s to 120 s from the start of the ramp was chosen to obtain a large window and avoid capturing the immediate few seconds before a breakdown. Considerable variance is shown for this quantity, with a slight downward trend for the mechanically polished electrode data. The passivated electrodes seemingly fluctuated downwards starting at earlier breakdowns, which points towards the idea that passivated electrodes are pre-conditioned.

This pre-conditioning can also be seen in the pulse rate data for the final 30s in Fig 8.29. In order to be more resilient against DAQ saturation, the data is presented in the aforementioned Geiger-like mode, and averaged over the final 30s before a breakdown. Only pulses above 1000 ADCC were counted. Terminal pulse rates appears to scale with increased stressed electrode area. The passivated electrodes (Run 10) are slightly lower in terminal pulse rates. For the pitted electrodes (Run 7), the SEA scaling is not as clear. From these patterns, I infer that the precursor discharge rate in the final 30s is in part due to asperities, and that these asperities can be conditioned away through repeated breakdowns.

In summary, small discharges on the anode were seen preceding breakdowns and grew rapidly in the few tens of seconds previous to a breakdown. Also, SiPM pulses were observed in coincidence with large enough anode charge pulses. These precursor discharges flare up considerably as the voltage increases. A small conditioning effect is also seen whereby the precursor current diminished over time. Strategies for exploiting such observations will depend on the particulars of the experiment.

Future work in this regard could add additional SiPMs to obtain position additional position reconstruction data. More careful calibration can also help to reconstruct anode current even in the case of DAQ saturation. Better thermal coupling between the SiPM and the steel can reduce the dark rate, which will grant additional sensitivity.

PMT Static Voltage Observations

Although the main focus of the XeBrA analysis was to examine the risk of breakdowns, it was also of concern whether a certain voltage produces excess single photon background. To examine this phenomenon, the linear ramps were replaced with a ramp-and-hold scheme for certain datasets. These data were always taken after the linear ramp data. The high voltage setpoint was assigned a value that was approximately 10 kV below the smallest breakdown observed during the linear ramps. The separation distance was set to 5 mm, which gave the power supply considerable room for error.

In Fig. 8.30 the results of one test are shown. Despite no breakdowns occurring over the course of the ramp-and-hold procedure, the single-photoelectron rate spikes upwards several times over the ten minute time period. An additional spike happened during the ramp down.




Figure 8.23: *Top*: Pulses found over the course of a single ramp. The anode and SiPM pulses are nearly coincident with one another. The anode pulses have a better signal-to-noise ratio, and therefore smaller peaks can be resolved on this scale. *Bottom*: Complementary information obtained from the slow control logs for the same ramp. The burst in current at time 0 is not the breakdown itself, but rather the discharge of current through the HV power supply due to the sudden drop of electric potential energy.




Figure 8.24: A precursor event window. Dotted lines indicate locations of identified pulses. Note the coincidence between the anode (top) and SiPM (bottom) waveforms. The relative scale indicates that the SiPM signal to noise ratio is considerably worse than the anode signal. The overshoot of the SiPM signal is also apparent.




Figure 8.25: Distributions of anode current preceding a breakdown for different parameter scans in Run 10, over time. The anode current is averaged over in time bins of 1 second. Pulses within 2 event windows (approximately 750 ms) of a subthreshold breakdown (defined in text) were discarded. A general rise in anode current in the few tens of seconds before breakdown is observed in all cases.




Figure 8.26: Comparison of anode current and SiPM rate in the 30 seconds preceding a breakdown between Runs 9 (orange, mechanically polished electrodes) and Run 10 (blue, passivated electrodes). The anode current is averaged over in time bins of 1s. An electrode separation of 2 mm is considered in both cases. The two sets of electrodes perform similarly, except for the last few seconds in which the polished cathode exhibits increased activity relative to the passivated cathode.




Figure 8.27: Pulse height distribution for Run 9, demonstrating frequent DAQ saturation.




Figure 8.28: Average current between 30 s and 120 s from the beginning of the HV ramps for each breakdown in Run 9. A slight downwards trend is apparent.




Figure 8.29: Anode pulse Rate vs. stressed electrode area over several runs. Run 7 used the original electrode set, and was later discovered to be pitted heavily. The remaining runs here use fresh electrodes, with Run 8 using an alternate set of mechanically polished electrodes, Run 9 using the original electrodes after re-polishing, and Run 10 using passivated electrodes.




Figure 8.30: A representative PMT ramp and hold pattern.

Qualitative Results

Index of Refraction

With the assistance of the viewports, the xenon could be visually monitored over the course of a ramp. An intriguing pattern emerged for separation distances of 3 mm or greater, where a large volume of stressed xenon could be observed. As the voltage on the cathode increased, the xenon exhibited a fascinating optical phenomenon. It would undergo a “shimmering” effect, where the background image of the apparatus would appear to ripple. The intensity of the ripples would increase in intensity up to the point of breakdown. No ripples were ever observed while the electrodes were unbiased. A close analog for this effect is the appearance of alcohol being poured into water and subsequently mixing.

A possible explanation for this effect is the change of index of refraction from the heating of liquid xenon. Using the Lorentz-Lorenz relationship:

$$\frac{n^2 - 1}{n^2 + 2} = K\rho . \quad (8.35)$$

Together with the value of $\epsilon_r \approx n^2 = 1.874$ at the triple point[127], and the thermal contraction of xenon from 2.98 g/cm³ 161K to 2.7 g/cm³ at 200K[249], the temperature dependence of n may be estimated. I find that the dependence is given over this temperature range can be approximated as

$$\frac{dn}{dT} = -9.95 \times 10^{-4} \text{ K}^{-1} . \quad (8.36)$$

The possibility of xenon heating serves as a possible explanation for the breakdown fields in liquid. Superheating in liquid was explored in Ref. [248], where two mechanisms were identified: burst production and percolation. In the former, the liquid superheats until the point of absolute stability at the spinodal line, where bubbles will spontaneously nucleate and trigger a breakdown. In the latter, small bubbles nucleate from impurities until a sufficient volume is occupied by the gas phase between the two electrodes. Then, a conducting path forms via discharges jumping between the weak gaseous regions. Pressure dependence arises in the percolation case from the Joule heating of the liquid needing to balance the surface tension of the bubbles.

Percolation is a potential explanation for some of the sub-threshold breakdowns and precursors. At bubble densities below the threshold needed for percolation, (1/3 of the total volume within the path), avalanches or streamers may occur, growing to observable size within a single void, but is not sufficient to trigger a conducting path across the entire gap.

Bubbles

In the high speed camera videos, frames are frequently visible showing bright spots prior to breakdown. These are often too small to distinguish from a background glow signally the onset of breakdown. On occasion these spots will appear, and then move around before a breakdown occurs in the spot the bubble lands.




Figure 8.31: A series of high-frame-rate photographs of a breakdown containing two apparent discharges. Four frames are shown in sequential order from top to bottom. Two light spots appear in the center of the frame before a big flash of light is emitted and two bubbles emerge from that location immediately after.

Some interesting phenomena were evident in the videos that I hand scanned. In multiple cases I observed a bubble drifting *downwards* from the anode to cathode before triggering a breakdown. Other times, two distinct bubbles were observed to grow over several frames, before breakdowns were seen at their respective locations. After any breakdown, copious bubbles were generally observed. These expand and slowly drift off the side of the anode.

It has been pointed out that bubbles may play an important role in the breakdown of insulating liquids [248, 240]. Bubbles can also explain the precursor pulses observed in XeBrA. Small gas bubbles may be sufficient to break down at some fields, but the breakdowns themselves do not open up a conducting ionizing channel from cathode to anode. This quenching mechanism causes small breakdowns to become more frequent over time, until the bubbles generated via Joule heating reach a critical density, allowing streamers to jump between voids, leading to a macroscopic breakdown.

A point in favor of this hypothesis is the visual observation of precursors at larger gap separation. At 3 mm or higher, breakdown voltages are typically around 30 kV or higher, allowing ample time to observe interesting phenomena (like the “shimmering”). Occasionally short sparks were seen in the liquid which *did not* subtend the full gap, and appeared at random distances between the electrodes. These had the appearance of a handheld sparkler, and are apparent in videos taken, but are difficult to observe in still photography.

Pitting

It was observed after Run 7 that the electrodes were sustaining damage from the hundreds of breakdowns impinging on their surface. While this may help initially, due to burning away asperities, the effect of the numerous craters is a factor that we can not easily control for. This prompted the replacement of the electrodes with duplicates for Run 8. Additional changes were that for Run 9 and onwards, separation distances > 3 mm were not performed.



Figure 8.32: The XeBrA electrodes after hundreds of breakdowns, showing damage in the form of white pits.

This was apparently helpful, as larger craters were absent following the Run 9 to Run 10 swap. Examples of pitting are shown in Fig. 8.32.

8.8 Summary

Conclusion

The XeBrA experiment, which employed a pair of Rogowski electrodes with adjustable separation distance, we confirmed that breakdown field scales inversely with stressed electrode area in LXe for stressed areas up to 33 cm^2 . No significant, reproducible scaling with LXe pressure was observed between 1.5 bar and 2.2 bar absolute. Similarly, no significant dependence was observed with the high voltage ramp speed in the range between 100 V/s and 200 V/s. By contrast, a small increase in breakdown field was observed with the passivation of the electrode surfaces.

The location of a breakdown on the electrode surface was reconstructed from high-speed videos taken by a pair of cameras affixed to the perpendicular viewports. They show moderate correlations with FN current as obtained by electric field simulations. Moreover, evidence for bubble nucleation initiating breakdown was observed with low-speed videos. We hypothesize that the bubble nucleation itself is caused by localized heating due to field emission from field enhancers, such as local asperities.

Finally, some practical recommendations can be drawn from this work. First, significant breakdown precursor activity was observed. Small discharges were detected with a SiPM and a charge-sensitive amplifier connected to the anode. These discharges slowly increase in intensity in the last 30 seconds before a breakdown and sharply accelerate in the last second. Hence, having the ability to ramp down the voltage on the scale of seconds could be

beneficial for avoiding electrical breakdown. Second, a downward trend in current collected at the anode was observed over the course of a run, which motivates conducting an initial conditioning campaign prior to the start of any science data-taking campaign. Third, the extrapolation from our measured stressed area scaling shows that the breakdown field for the next-generation, LXe TPC experiments, will only be on the order of a few tens of kV/cm. These results have important implications for the design of the high voltage delivery system of such experiments.

Recommendations

One of the current generation liquid xenon TPCs is the LUX-ZEPLIN (LZ) experiment. The LZ cathode ring has a simulated stressed area of approximately 940 cm^2 . Evaluating the model from this work for the mechanically polished electrodes gives a predicted shifted scaled parameter (E_2) of $28.2 \pm 0.5 \text{ kV/cm}$. This is a revision downwards from the prediction in Ref. [104], shown in blue in the top figure of 8.19. Using the model from XeBrA's 2018 analysis, the predicted breakdown field at the LZ cathode ring SEA is $70.4 \pm 3 \text{ kV/cm}$. We believe the discrepancy is a result of the following changes: the uncertainty on the cathode tilt angle was taken into account (instead of assuming a 0° tilt of the cathode), each run consisted of several days of operation in which the xenon was continuously being circulated (as opposed to one very long day of data-taking), and a new method to mitigate bubble formation in the bulk was introduced. To provide some reference, the maximum design voltage in LZ was set at 50 kV/cm [57], which is intermediate between these two extrapolated values. In addition, the predicted breakdown voltage for a 40 tonne active volume TPC (e.g. DARWIN [77]), assuming that SEA scales with mass to the power of $2/3$, would be $20.7 \pm 0.4 \text{ kV/cm}$. This is illustrated in the bottom panel of 8.19, where the estimated stressed area of the cathode ring for LZ and two hypothetical Generation-3 (G3) experiments are shown. Consequently, careful engineering of the electrostatics of any future, large TPC will be critical in achieving the desired fields. This analysis indicates that future LXe experiments of the scale of DARWIN should design around a maximum electric field of 20 kV/cm on their cathode rings in order to have some safety margin. However, we note that improvements on cathode design could raise this threshold.

Future Work

It became evident during the 2020 run that certain features of the system required additional controls to make quality inter-run comparisons.

The main source of error on the SEA comes from the tilt of the cathode. While the photogrammetric method constrains the cathode tilt to within $\pm 0.1^\circ$, this still leads to considerable error on the SEA, especially at lower separation distances. For this reason, future runs of XeBrA will utilize a set of molds which fits the Rogowski profiles of the cathode and anode. When squeezed together, these molds will force the cathode into a specified angle. I designed these molds, which were then 3D printed. I then confirmed that they fit the electrodes and fit within the ICV. On this front the work which remains is to

estimate, and possibly correct for, the level of the molds. Even if the molds do not force the electrodes into a parallel configuration, being able to have a consistent tilt when the electrodes are swapped will vastly reduce the systematic error.

Due to the changes in pressure over the course of a run, the circulation rate is frequently altered. The placement of PT05, which measures the gas pressure within XeBrA, is on the HX inlet port. The XeBrA interior constitutes a significant impedance, leading to a pressure drop. As a result, the system pressure is always measured in a flow state, leading to the PT05 being biased upwards from the gas pressure at the liquid surface. Future runs of XeBrA will benefit from a change in location of PT05. It has been moved from the inlet port to the bellows which connects the inner and outer volumes for pumpout. This allows the pressure to be measured, and therefore moderated by the PLC, in a neutral flow state. Care had to be taken to make sure that the emergency recovery system will still operate under this condition.

The pressure conditions were quite challenging to deal with. No definitive relationship with xenon pressure was observed, but significant variance in breakdown field was observed with changes in pressure. I, along with the rest of the XeBrA team, believe that this is due to the methodology of how the pressure was changed. There was a lack of consistency with what order the pressure data was taken (e.g. increasing, decreasing, monotonic, same or different days). The only run where the breakdown field was relatively consistent across the scan was Run 8, where the pressure was monotonically increased over the course of a single day. Future XeBrA runs will take greater care to test the exact same sequence of pressure adjustments during each run.

Finally, an additional goal of future analysis will be conducting time-to-failure analysis. While the linear ramp results in many breakdowns over a fixed window of time, it does not answer the primary concern of LXe-TPCs, which is which breakdown field can be held over year plus time scales. Ideally a test up to and including the peak fields for a given experiment would be conducted. Then, the voltage would be held until a breakdown occurs, and is then recorded. The disadvantage of this method is that each data point will take a long amount of time, potentially. This means that alternative conditions such as pressure will be difficult to examine.

PTFE is known to fluoresce[191] due to the hydrocarbon contaminants. This may constitute a background to LXe-TPCs. Future work with XeBrA may include measurements of the PTFE fluorescence over long timescales.

Conclusion

In the preceding dissertation I have described the contributions that I have made in the field of particle astrophysics over the course of my graduate studies. My primary focus was participation in the construction, operation, and analysis of the LZ dark matter direct detection experiment. The projects which I worked on were varied, allowing me to build skills in a variety of sub fields. Overall, my work revolved around the impact that electric fields have on the future of dark matter experiments. The cathode high voltage feedthrough project directly determined the magnitude of the drift electric field within the LZ TPC. Much of my simulations work involved the modelling and deployment of the spatially-varying electric fields within the xenon. With regards to the accidental coincidence backgrounds, the broad conclusion is that the electric fields in the extraction region, particularly in the gas phase, are a significant driver of uncorrelated S1 and S2 pulses. The Xenon Breakdown Apparatus collected a large amount of data which supported the theory of stressed area scaling for breakdown voltages, as well as evidence for activity preceding discharges. While not tightly coupled to the electric fields, the multiple scatter / ultraheavy dark matter search in LZ is influenced by field nonuniformities and the isolated pulse rate, both of which are strong functions of the electric field. Moving forward, both LZ and its potential successors will require accurate predictions for their electric fields. Intensive research and development is recommended in order to optimize for the effects observed over the course of this thesis.

References

- [1] Gianfranco Bertone and Dan Hooper. “History of dark matter”. In: *Rev. Mod. Phys.* 90.4 (2018). eprint: 1605.04909, p. 045002. DOI: 10.1103/RevModPhys.90.045002.
- [2] Scott Tremaine and James E. Gunn. “Dynamical Role of Light Neutral Leptons in Cosmology”. In: *Physical Review Letters* 42.6 (Feb. 1979). Publisher: American Physical Society, pp. 407–410. DOI: 10.1103/PhysRevLett.42.407. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.42.407> (visited on 11/07/2022).
- [3] Planck Collaboration et al. “Planck 2018 results. VI. Cosmological parameters”. In: *Astronomy & Astrophysics* 641 (Sept. 2020). arXiv:1807.06209 [astro-ph], A6. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201833910. URL: <http://arxiv.org/abs/1807.06209> (visited on 07/18/2022).
- [4] Pablo Villanueva-Domingo, Olga Mena, and Sergio Palomares-Ruiz. “A brief review on primordial black holes as dark matter”. In: *Frontiers in Astronomy and Space Sciences* 8 (May 2021). arXiv:2103.12087 [astro-ph], p. 681084. ISSN: 2296-987X. DOI: 10.3389/fspas.2021.681084. URL: <http://arxiv.org/abs/2103.12087> (visited on 11/07/2022).
- [5] Fritz Zwicky. “Die Rotverschiebung von extragalaktischen Nebeln”. In: *Helvetica Physica Acta* 6 (Jan. 1933). Provided by the SAO/NASA Astrophysics Data System, pp. 110–127. URL: <https://ui.adsabs.harvard.edu/abs/1933AChPh...6..110Z>.
- [6] A. M. Bykov, F. B. S. Paerels, and V. Petrosian. “Equilibration processes in the Warm-Hot Intergalactic Medium”. In: *Space Science Reviews* 134.1-4 (Feb. 2008). arXiv:0801.1008 [astro-ph], pp. 141–153. ISSN: 0038-6308, 1572-9672. DOI: 10.1007/s11214-008-9309-4. URL: <http://arxiv.org/abs/0801.1008> (visited on 07/18/2022).
- [7] Douglas Clowe et al. “A direct empirical proof of the existence of dark matter”. In: *The Astrophysical Journal* 648.2 (Sept. 2006). arXiv:astro-ph/0608407, pp. L109–L113. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/508162. URL: <http://arxiv.org/abs/astro-ph/0608407> (visited on 11/30/2022).
- [8] Scott W. Randall et al. “Constraints on the Self-Interaction Cross Section of Dark Matter from Numerical Simulations of the Merging Galaxy Cluster 1E 0657–56”. en. In: *The Astrophysical Journal* 679.2 (June 2008). Publisher: IOP Publishing, p. 1173.

- ISSN: 0004-637X. DOI: 10.1086/587859. URL: <https://iopscience.iop.org/article/10.1086/587859/meta> (visited on 07/20/2022).
- [9] Richard H. Cyburt et al. “Big bang nucleosynthesis: Present status”. In: *Reviews of Modern Physics* 88.1 (Feb. 2016). Publisher: American Physical Society, p. 015004. DOI: 10.1103/RevModPhys.88.015004. URL: <https://link.aps.org/doi/10.1103/RevModPhys.88.015004> (visited on 12/02/2022).
- [10] Brian D. Fields, Paolo Molaro, and Subir Sarkar. “Big-Bang Nucleosynthesis”. In: *Chinese Physics C* 38.9 (Aug. 2014). arXiv:1412.1408 [astro-ph], p. 090001. ISSN: 1674-1137. DOI: 10.1088/1674-1137/38/9/090001. URL: <http://arxiv.org/abs/1412.1408> (visited on 12/02/2022).
- [11] Carlos A. Bertulani, Francis W. Hall, and Benjamin I. Santoyo. *Big Bang nucleosynthesis as a probe of new physics*. arXiv:2210.04071 [astro-ph, physics:nucl-th]. Nov. 2022. DOI: 10.48550/arXiv.2210.04071. URL: <http://arxiv.org/abs/2210.04071> (visited on 12/05/2022).
- [12] Wayne Hu. “CMB Temperature and Polarization Anisotropy Fundamentals”. In: *Annals of Physics* 303.1 (Jan. 2003). arXiv:astro-ph/0210696, pp. 203–225. ISSN: 00034916. DOI: 10.1016/S0003-4916(02)00022-2. URL: <http://arxiv.org/abs/astro-ph/0210696> (visited on 11/08/2022).
- [13] Solène Chabanier, Marius Millea, and Nathalie Palanque-Delabrouille. “Matter power spectrum: from Ly α forest to CMB scales”. In: *Monthly Notices of the Royal Astronomical Society* 489.2 (Oct. 2019). arXiv:1905.08103 [astro-ph], pp. 2247–2253. ISSN: 0035-8711, 1365-2966. DOI: 10.1093/mnras/stz2310. URL: <http://arxiv.org/abs/1905.08103> (visited on 11/09/2022).
- [14] Kim Griest and Marc Kamionkowski. “Unitarity Limits on the Mass and Radius of Dark Matter Particles”. In: *Phys. Rev. Lett.* 64 (1990), p. 615. DOI: 10.1103/PhysRevLett.64.615.
- [15] Nausheen Shan. *Minimal Supersymmetric Standard Model (MSSM)*. 2003. URL: homes.psd.uchicago.edu/~sethi/Teaching/P487-S2003/MSSMnausheen.pdf.
- [16] Gerard Jungman, Marc Kamionkowski, and Kim Griest. “Supersymmetric dark matter”. en. In: *Physics Reports* 267.5 (Mar. 1996), pp. 195–373. ISSN: 0370-1573. DOI: 10.1016/0370-1573(95)00058-5. URL: <https://www.sciencedirect.com/science/article/pii/0370157395000585> (visited on 11/07/2022).
- [17] G. L. Kane et al. “Study of constrained minimal supersymmetry”. In: *Physical Review D* 49.11 (June 1994). Publisher: American Physical Society, pp. 6173–6210. DOI: 10.1103/PhysRevD.49.6173. URL: <https://link.aps.org/doi/10.1103/PhysRevD.49.6173> (visited on 11/07/2022).
- [18] John Ellis et al. *The CMSSM Survives Planck, the LHC, LUX-ZEPLIN, Fermi-LAT, H.E.S.S. and IceCube*. arXiv:2210.16337 [astro-ph, physics:hep-ex, physics:hep-ph, physics:hep-th]. Oct. 2022. DOI: 10.48550/arXiv.2210.16337. URL: <http://arxiv.org/abs/2210.16337> (visited on 11/06/2022).

- [19] Francesca Chadha-Day, John Ellis, and David J. E. Marsh. *Axion Dark Matter: What is it and Why Now?* arXiv:2105.01406 [astro-ph, physics:hep-ex, physics:hep-ph]. Oct. 2022. URL: <http://arxiv.org/abs/2105.01406> (visited on 11/07/2022).
- [20] C. Abel et al. “Measurement of the Permanent Electric Dipole Moment of the Neutron”. In: *Phys. Rev. Lett.* 124.8 (2020). eprint: 2001.11966, p. 081803. DOI: 10.1103/PhysRevLett.124.081803.
- [21] R. D. Peccei and Helen R. Quinn. “ CP Conservation in the Presence of Pseudoparticles”. In: *Physical Review Letters* 38.25 (June 1977). Publisher: American Physical Society, pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.38.1440> (visited on 11/07/2022).
- [22] Raymond T. Co, Lawrence J. Hall, and Keisuke Harigaya. “Axion Kinetic Misalignment Mechanism”. In: *Physical Review Letters* 124 (June 2020). ADS Bibcode: 2020PhRvL.124y1802C, p. 251802. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.124.251802. URL: <https://ui.adsabs.harvard.edu/abs/2020PhRvL.124y1802C> (visited on 11/07/2022).
- [23] R. Khatiwada et al. “Axion Dark Matter eXperiment: Detailed Design and Operations”. In: *Review of Scientific Instruments* 92.12 (Dec. 2021). arXiv:2010.00169 [astro-ph, physics:hep-ex, physics:physics, physics:quant-ph], p. 124502. ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0037857. URL: <http://arxiv.org/abs/2010.00169> (visited on 11/07/2022).
- [24] Maximiliano Silva-Feaver et al. *Design Overview of the DM Radio Pathfinder Experiment*. arXiv:1610.09344 [astro-ph, physics:physics]. Oct. 2016. DOI: 10.48550/arXiv.1610.09344. URL: <http://arxiv.org/abs/1610.09344> (visited on 11/07/2022).
- [25] Fuminobu Takahashi, Masaki Yamada, and Wen Yin. “XENON1T excess from anomaly-free ALP dark matter and its implications for stellar cooling anomaly”. In: *Physical Review Letters* 125.16 (Oct. 2020). arXiv:2006.10035 [astro-ph, physics:hep-ph], p. 161801. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.125.161801. URL: <http://arxiv.org/abs/2006.10035> (visited on 11/07/2022).
- [26] S T Petcov. “13. NEUTRINO MASS, MIXING, AND OSCILLATIONS”. en. In: (2012), p. 58.
- [27] Carlo Giunti and Chung Wook Kim. *Fundamentals of neutrino physics and astrophysics*. en. OCLC: ocm76935610. Oxford ; New York: Oxford University Press, 2007. ISBN: 978-0-19-850871-7.
- [28] A. Boyarsky et al. “Sterile Neutrino Dark Matter”. In: *Progress in Particle and Nuclear Physics* 104 (Jan. 2019). arXiv:1807.07938 [astro-ph, physics:hep-ex, physics:hep-ph], pp. 1–45. ISSN: 01466410. DOI: 10.1016/j.ppnp.2018.07.004. URL: <http://arxiv.org/abs/1807.07938> (visited on 11/07/2022).

- [29] Esra Bulbul et al. “Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters”. In: *The Astrophysical Journal* 789.1 (June 2014). arXiv:1402.2301 [astro-ph], p. 13. ISSN: 0004-637X, 1538-4357. DOI: 10.1088/0004-637X/789/1/13. URL: <http://arxiv.org/abs/1402.2301> (visited on 11/15/2022).
- [30] Basudeb Dasgupta and Joachim Kopp. “Sterile Neutrinos”. In: *arXiv:2106.05913 [astro-ph, physics:hep-ex, physics:hep-ph]* (June 2021). arXiv: 2106.05913. URL: <http://arxiv.org/abs/2106.05913> (visited on 06/14/2021).
- [31] Pablo Villanueva-Domingo and Kiyotomo Ichiki. “21 cm Forest Constraints on Primordial Black Holes”. In: (Apr. 2021). eprint: 2104.10695. DOI: 10.1093/pasj/psab119.
- [32] M. R. S. Hawkins. “A new look at microlensing limits on dark matter in the Galactic halo”. In: *Astronomy & Astrophysics* 575 (Mar. 2015). arXiv:1503.01935 [astro-ph], A107. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201425400. URL: <http://arxiv.org/abs/1503.01935> (visited on 11/07/2022).
- [33] Asher Berlin et al. “Dark Matter, Millicharges, Axion and Scalar Particles, Gauge Bosons, and Other New Physics with LDMX”. In: *Physical Review D* 99.7 (Apr. 2019). arXiv:1807.01730 [astro-ph, physics:hep-ex, physics:hep-ph], p. 075001. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.99.075001. URL: <http://arxiv.org/abs/1807.01730> (visited on 11/08/2022).
- [34] Ahmet Coskuner et al. “Direct Detection of Bound States of Asymmetric Dark Matter”. In: *Physical Review D* 100.3 (Aug. 2019). arXiv:1812.07573 [hep-ph], p. 035025. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.100.035025. URL: <http://arxiv.org/abs/1812.07573> (visited on 09/23/2022).
- [35] R. Essig et al. *Dark Sectors and New, Light, Weakly-Coupled Particles*. Oct. 2013. URL: <http://arxiv.org/abs/1311.0029> (visited on 11/08/2022).
- [36] Edward W. Kolb, Daniel J. H. Chung, and Antonio Riotto. *WIMPZILLAS!* arXiv:hep-ph/9810361. Oct. 1998. DOI: 10.48550/arXiv.hep-ph/9810361. URL: <http://arxiv.org/abs/hep-ph/9810361> (visited on 09/13/2022).
- [37] Kathryn M. Zurek. “Asymmetric Dark Matter: Theories, Signatures, and Constraints”. In: *Physics Reports* 537.3 (Apr. 2014). arXiv:1308.0338 [astro-ph, physics:hep-ph], pp. 91–121. ISSN: 03701573. DOI: 10.1016/j.physrep.2013.12.001. URL: <http://arxiv.org/abs/1308.0338> (visited on 10/30/2022).
- [38] Moira I. Gresham, Hou Keong Lou, and Kathryn M. Zurek. “Astrophysical Signatures of Asymmetric Dark Matter Bound States”. In: *Physical Review D* 98.9 (Nov. 2018). arXiv:1805.04512 [astro-ph, physics:hep-ph], p. 096001. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.98.096001. URL: <http://arxiv.org/abs/1805.04512> (visited on 08/22/2022).

- [39] Julien Billard et al. “Direct Detection of Dark Matter – APPEC Committee Report”. In: *Reports on Progress in Physics* 85.5 (May 2022). arXiv:2104.07634 [astro-ph, physics:hep-ex, physics:hep-ph], p. 056201. ISSN: 0034-4885, 1361-6633. DOI: 10.1088/1361-6633/ac5754. URL: <http://arxiv.org/abs/2104.07634> (visited on 11/08/2022).
- [40] The Fermi-LAT Collaboration. “The Fermi Galactic Center GeV Excess and Implications for Dark Matter”. In: *The Astrophysical Journal* 840.1 (2017). arXiv:1704.03910 [astro-ph], p. 43. ISSN: 1538-4357. DOI: 10.3847/1538-4357/aa6cab. (Visited on 09/23/2022).
- [41] Dan Hooper and Lisa Goodenough. “Dark Matter Annihilation in The Galactic Center As Seen by the Fermi Gamma Ray Space Telescope”. In: *Physics Letters B* 697.5 (Mar. 2011). arXiv:1010.2752 [astro-ph, physics:hep-ph], pp. 412–428. ISSN: 03702693. DOI: 10.1016/j.physletb.2011.02.029. URL: <http://arxiv.org/abs/1010.2752> (visited on 11/08/2022).
- [42] Man Ho Chan and Chung Hei Leung. *Ruling out dark matter interpretation of the galactic GeV excess by gamma-ray data of galaxy clusters*. arXiv:1710.08123 [astro-ph]. Oct. 2017. DOI: 10.48550/arXiv.1710.08123. URL: <http://arxiv.org/abs/1710.08123> (visited on 11/15/2022).
- [43] Christopher Dessert, Nicholas L. Rodd, and Benjamin R. Safdi. “The dark matter interpretation of the 3.5-keV line is inconsistent with blank-sky observations”. In: *Science* 367.6485 (Mar. 2020). arXiv:1812.06976 [astro-ph, physics:hep-ph], pp. 1465–1467. ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.aaw3772. URL: <http://arxiv.org/abs/1812.06976> (visited on 11/08/2022).
- [44] Christopher McCabe. “The Earth’s velocity for direct detection experiments”. en. In: *Journal of Cosmology and Astroparticle Physics* 2014.02 (Feb. 2014), p. 027. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2014/02/027. URL: <https://dx.doi.org/10.1088/1475-7516/2014/02/027> (visited on 10/31/2022).
- [45] D. Baxter et al. “Recommended conventions for reporting results from direct dark matter searches”. In: *The European Physical Journal C* 81.10 (2021). arXiv:2105.00599 [hep-ex, physics:hep-ph], p. 907. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-021-09655-y. (Visited on 06/13/2022).
- [46] SuperCDMS Collaboration et al. “Constraints on low-mass, relic dark matter candidates from a surface-operated SuperCDMS single-charge sensitive detector”. In: *Physical Review D* 102.9 (Nov. 2020). arXiv:2005.14067 [astro-ph, physics:hep-ex, physics:physics], p. 091101. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.102.091101. URL: <http://arxiv.org/abs/2005.14067> (visited on 11/08/2022).
- [47] R. Bernabei et al. *The dark matter: DAMA/LIBRA and its perspectives*. Oct. 2021. DOI: 10.48550/arXiv.2110.04734. URL: <http://arxiv.org/abs/2110.04734> (visited on 11/08/2022).

- [48] G. Angloher et al. “Results on light dark matter particles with a low-threshold CRESST-II detector”. en. In: *The European Physical Journal C* 76.1 (Jan. 2016), p. 25. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-016-3877-3. URL: <http://link.springer.com/10.1140/epjc/s10052-016-3877-3> (visited on 11/08/2022).
- [49] EDELWEISS Collaboration et al. “Search for sub-GeV dark matter via the Migdal effect with an EDELWEISS germanium detector with NbSi transition-edge sensors”. In: *Physical Review D* 106.6 (Sept. 2022). Publisher: American Physical Society, p. 062004. DOI: 10.1103/PhysRevD.106.062004. URL: <https://link.aps.org/doi/10.1103/PhysRevD.106.062004> (visited on 11/15/2022).
- [50] Masahiro Ibe et al. “Migdal Effect in Dark Matter Direct Detection Experiments”. In: *Journal of High Energy Physics* 2018.3 (Mar. 2018). arXiv:1707.07258 [hep-ex, physics:hep-ph], p. 194. ISSN: 1029-8479. DOI: 10.1007/JHEP03(2018)194. URL: <http://arxiv.org/abs/1707.07258> (visited on 11/11/2022).
- [51] J. Gascon et al. *Low-mass Dark Matter searches with EDELWEISS*. arXiv:2112.05467 [physics]. Mar. 2022. URL: <http://arxiv.org/abs/2112.05467> (visited on 11/08/2022).
- [52] Liron Barak et al. “SENSEI: Direct-Detection Results on sub-GeV Dark Matter from a New Skipper-CCD”. In: *Physical Review Letters* 125.17 (Oct. 2020). arXiv:2004.11378 [astro-ph, physics:hep-ex, physics:hep-ph, physics:physics], p. 171802. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.125.171802. URL: <http://arxiv.org/abs/2004.11378> (visited on 11/08/2022).
- [53] J. Aalbers et al. *First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment*. arXiv:2207.03764 [astro-ph, physics:hep-ex]. July 2022. URL: <http://arxiv.org/abs/2207.03764> (visited on 10/23/2022).
- [54] S. E. Vahsen et al. *CYGNUS: Feasibility of a nuclear recoil observatory with directional sensitivity to dark matter and neutrinos*. arXiv:2008.12587 [astro-ph, physics:hep-ex, physics:hep-ph, physics:physics]. Dec. 2020. DOI: 10.48550/arXiv.2008.12587. URL: <http://arxiv.org/abs/2008.12587> (visited on 11/08/2022).
- [55] The LUX Collaboration et al. “Improved Modeling of β Electronic Recoils in Liquid Xenon Using LUX Calibration Data”. In: *Journal of Instrumentation* 15.02 (Feb. 2020). arXiv:1910.04211 [hep-ex, physics:physics], T02007–T02007. ISSN: 1748-0221. DOI: 10.1088/1748-0221/15/02/T02007. URL: <http://arxiv.org/abs/1910.04211> (visited on 09/23/2022).
- [56] D. Yu Akimov et al. “The ZEPLIN-III dark matter detector: instrument design, manufacture and commissioning”. In: *Astroparticle Physics* 27.1 (Feb. 2007). arXiv:astro-ph/0605500, pp. 46–60. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2006.09.005. URL: <http://arxiv.org/abs/astro-ph/0605500> (visited on 12/01/2022).
- [57] B. J. Mount et al. “LUX-ZEPLIN (LZ) Technical Design Report”. In: *arXiv:1703.09144 [astro-ph, physics:hep-ex, physics:physics]* (Mar. 2017). arXiv: 1703.09144. URL: <http://arxiv.org/abs/1703.09144> (visited on 08/17/2021).

- [58] D. S. Akerib et al. “The LUX-ZEPLIN (LZ) radioactivity and cleanliness control programs”. en. In: *The European Physical Journal C* 80.11 (Nov. 2020), p. 1044. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-020-8420-x. URL: <https://doi.org/10.1140/epjc/s10052-020-8420-x> (visited on 12/26/2022).
- [59] *Photomultiplier tube R11410-20 — Hamamatsu Photonics*. en. 2021. URL: https://www.hamamatsu.com/us/en/product/optical-sensors/pmt/pmt_tube-alone/head-on-type/R11410-20.html (visited on 10/26/2022).
- [60] D. S. Akerib et al. “Identification of Radiopure Titanium for the LZ Dark Matter Experiment and Future Rare Event Searches”. In: *Astroparticle Physics* 96 (Nov. 2017). arXiv:1702.02646 [hep-ex, physics:physics], pp. 1–10. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2017.09.002. URL: <http://arxiv.org/abs/1702.02646> (visited on 11/08/2022).
- [61] George Bakale, Ulrich Sowada, and Werner F. Schmidt. “Effect of an electric field on electron attachment to sulfur hexafluoride, nitrous oxide, and molecular oxygen in liquid argon and xenon”. en. In: *The Journal of Physical Chemistry* 80.23 (Nov. 1976), pp. 2556–2559. ISSN: 0022-3654, 1541-5740. DOI: 10.1021/j100564a006. URL: <https://pubs.acs.org/doi/10.1021/j100564a006> (visited on 08/17/2021).
- [62] J. Aalbers et al. *Background Determination for the LUX-ZEPLIN (LZ) Dark Matter Experiment*. arXiv:2211.17120 [hep-ex, physics:physics]. Nov. 2022. URL: <http://arxiv.org/abs/2211.17120> (visited on 12/01/2022).
- [63] Carl Eric Dahl. “The physics of background discrimination in liquid xenon, and first results from Xenon10 in the hunt for WIMP dark matter”. PhD Thesis. Princeton U., 2009.
- [64] Brian Lenardo et al. “A Global Analysis of Light and Charge Yields in Liquid Xenon”. In: *IEEE Transactions on Nuclear Science* 62.6 (Dec. 2015). arXiv:1412.4417 [astro-ph, physics:physics], pp. 3387–3396. ISSN: 0018-9499, 1558-1578. DOI: 10.1109/TNS.2015.2481322. URL: <http://arxiv.org/abs/1412.4417> (visited on 08/09/2022).
- [65] Tadayoshi Doke et al. “Absolute Scintillation Yields in Liquid Argon and Xenon for Various Particles”. en. In: *Japanese Journal of Applied Physics* 41.3R (Mar. 2002). Publisher: IOP Publishing, p. 1538. ISSN: 1347-4065. DOI: 10.1143/JJAP.41.1538. URL: <https://iopscience.iop.org/article/10.1143/JJAP.41.1538/meta> (visited on 11/09/2022).
- [66] M. Miyajima et al. “Average energy expended per ion pair in liquid argon”. In: *Physical Review A* 9.3 (Mar. 1974). Publisher: American Physical Society, pp. 1438–1443. DOI: 10.1103/PhysRevA.9.1438. URL: <https://link.aps.org/doi/10.1103/PhysRevA.9.1438> (visited on 11/09/2022).

- [67] E. Aprile et al. “Observation of two-neutrino double electron capture in ^{124}Xe with XENON1T”. en. In: *Nature* 568.7753 (Apr. 2019). Number: 7753 Publisher: Nature Publishing Group, pp. 532–535. ISSN: 1476-4687. DOI: 10.1038/s41586-019-1124-4. URL: <https://www.nature.com/articles/s41586-019-1124-4> (visited on 11/09/2022).
- [68] J. B. Albert et al. “An improved measurement of the $2\nu\beta\beta$ half-life of Xe-136 with EXO-200”. In: *Physical Review C* 89.1 (Jan. 2014). arXiv:1306.6106 [hep-ex, physics:nucl-ex, physics:physics], p. 015502. ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.89.015502. URL: <http://arxiv.org/abs/1306.6106> (visited on 11/09/2022).
- [69] EXO-200 Collaboration et al. “Searches for Double Beta Decay of ^{134}Xe with EXO-200”. In: *Physical Review D* 96.9 (Nov. 2017), p. 092001. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.96.092001. (Visited on 11/09/2022).
- [70] J. Aalbers et al. “Cosmogenic production of ^{37}Ar in the context of the LUX-ZEPLIN experiment”. In: *Physical Review D* 105.8 (Apr. 2022). arXiv:2201.02858 [astro-ph, physics:hep-ex, physics:hep-ph], p. 082004. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.105.082004. URL: <http://arxiv.org/abs/2201.02858> (visited on 11/09/2022).
- [71] Keiko Fujii et al. “High-accuracy measurement of the emission spectrum of liquid xenon in the vacuum ultraviolet region”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 795 (Sept. 2015), pp. 293–297. ISSN: 0168-9002. DOI: 10.1016/j.nima.2015.05.065. URL: <https://www.sciencedirect.com/science/article/pii/S016890021500724X> (visited on 12/12/2022).
- [72] J. Mock et al. “Modeling pulse characteristics in Xenon with NEST”. en. In: *Journal of Instrumentation* 9.04 (Apr. 2014). Publisher: IOP Publishing, T04002–T04002. ISSN: 1748-0221. DOI: 10.1088/1748-0221/9/04/T04002. URL: <https://doi.org/10.1088/1748-0221/9/04/t04002> (visited on 08/09/2022).
- [73] Peter Sorensen and Carl Eric Dahl. “Nuclear recoil energy scale in liquid xenon with application to the direct detection of dark matter”. In: *Physical Review D* 83.6 (Mar. 2011). arXiv:1101.6080 [astro-ph, physics:hep-ph], p. 063501. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.83.063501. URL: <http://arxiv.org/abs/1101.6080> (visited on 08/09/2022).
- [74] J. Thomas and D. A. Imel. “Recombination of electron-ion pairs in liquid argon and liquid xenon”. In: *Physical Review A* 36.2 (July 1987). Publisher: American Physical Society, pp. 614–616. DOI: 10.1103/PhysRevA.36.614. URL: <https://link.aps.org/doi/10.1103/PhysRevA.36.614> (visited on 11/22/2022).
- [75] J Lindhard. “RANGE CONCEPTS AND HEAVY ION RANGES”. en. In: *Kgl. Danske Videnskab. Selskab. Mat. Fys. Medd.* 33.14 (1963), p. 44. URL: <https://www.osti.gov/biblio/4153115>.

- [76] S. Bruenner et al. “Radon daughter removal from PTFE surfaces and its application in liquid xenon detectors”. In: *The European Physical Journal C* 81 (Apr. 2021). DOI: 10.1140/epjc/s10052-021-09047-2.
- [77] J. Aalbers et al. “DARWIN: towards the ultimate dark matter detector”. In: *Journal of Cosmology and Astroparticle Physics* 2016.11 (Nov. 2016), pp. 017–017. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2016/11/017. URL: <https://iopscience.iop.org/article/10.1088/1475-7516/2016/11/017> (visited on 08/17/2021).
- [78] W. Turner et al. “Optical Calibration System for the LUX-ZEPLIN (LZ) Outer Detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 1010 (Sept. 2021). arXiv:2102.06281 [physics], p. 165551. ISSN: 01689002. DOI: 10.1016/j.nima.2021.165551. URL: <http://arxiv.org/abs/2102.06281> (visited on 11/07/2022).
- [79] S. J. Haselschwardt et al. “A Liquid Scintillation Detector for Radioassay of Gadolinium-Loaded Liquid Scintillator for the LZ Outer Detector”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 937 (Sept. 2019). arXiv:1808.05595 [hep-ex, physics:physics], pp. 148–163. ISSN: 01689002. DOI: 10.1016/j.nima.2019.05.055. URL: <http://arxiv.org/abs/1808.05595> (visited on 10/31/2022).
- [80] S F Mughabghab. “THERMAL NEUTRON CAPTURE CROSS SECTIONS RESONANCE INTEGRALS AND G-FACTORS”. en. In: *I N D C INTERNATIONAL NUCLEAR DATA COMMITTEE* (2003), p. 30.
- [81] D. S. Akerib et al. “Projected WIMP sensitivity of the LUX-ZEPLIN (LZ) dark matter experiment”. In: *Physical Review D* 101.5 (2020). arXiv:1802.06039 [astro-ph, physics:hep-ex, physics:physics], p. 052002. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.101.052002. URL: <http://arxiv.org/abs/1802.06039> (visited on 11/07/2022).
- [82] C. H. Faham et al. “Measurements of wavelength-dependent double photoelectron emission from single photons in VUV-sensitive photomultiplier tubes”. en. In: *Journal of Instrumentation* 10.09 (Sept. 2015), P09010. ISSN: 1748-0221. DOI: 10.1088/1748-0221/10/09/P09010. URL: <https://dx.doi.org/10.1088/1748-0221/10/09/P09010> (visited on 11/22/2022).
- [83] Laura Baudis, Patricia Sanchez-Lucas, and Kevin Thieme. “A measurement of the mean electronic excitation energy of liquid xenon”. In: *Eur. Phys. J. C* 81.12 (2021). eprint: 2109.07151, p. 1060. DOI: 10.1140/epjc/s10052-021-09834-x.
- [84] EXO-200 Collaboration et al. “Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment”. In: *Physical Review C* 101.6 (June 2020). arXiv:1908.04128 [hep-ex, physics:physics], p. 065501. ISSN: 2469-9985, 2469-9993. DOI: 10.1103/PhysRevC.101.065501. URL: <http://arxiv.org/abs/1908.04128> (visited on 11/08/2022).

- [85] Matthew Szydagis et al. “A Review of Basic Energy Reconstruction Techniques in Liquid Xenon and Argon Detectors for Dark Matter and Neutrino Physics Using NEST”. en. In: *Instruments* 5.1 (Mar. 2021). Number: 1 Publisher: Multidisciplinary Digital Publishing Institute, p. 13. ISSN: 2410-390X. DOI: 10.3390/instruments5010013. URL: <https://www.mdpi.com/2410-390X/5/1/13> (visited on 08/09/2022).
- [86] LUX Collaboration et al. “Position Reconstruction in LUX”. In: *Journal of Instrumentation* 13.02 (Feb. 2018). arXiv:1710.02752 [hep-ex, physics:physics], P02001–P02001. ISSN: 1748-0221. DOI: 10.1088/1748-0221/13/02/P02001. URL: <http://arxiv.org/abs/1710.02752> (visited on 11/01/2022).
- [87] A. G. Singh et al. “Analysis of $^{83\text{m}}\text{Kr}$ prompt scintillation signals in the PIXeY detector”. In: *JINST* 15.01 (2020). eprint: 1911.03999, P01023. DOI: 10.1088/1748-0221/15/01/P01023.
- [88] Gary J. Feldman and Robert D. Cousins. “A Unified Approach to the Classical Statistical Analysis of Small Signals”. In: *Physical Review D* 57.7 (Apr. 1998), pp. 3873–3889. ISSN: 0556-2821, 1089-4918. DOI: 10.1103/PhysRevD.57.3873. (Visited on 11/02/2022).
- [89] Steve Baker and Robert D. Cousins. “Clarification of the use of CHI-square and likelihood functions in fits to histograms”. en. In: *Nuclear Instruments and Methods in Physics Research* 221.2 (Apr. 1984), pp. 437–442. ISSN: 0167-5087. DOI: 10.1016/0167-5087(84)90016-4. URL: <https://www.sciencedirect.com/science/article/pii/0167508784900164> (visited on 11/07/2022).
- [90] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. en. In: *The Annals of Mathematical Statistics* 9.1 (Mar. 1938), pp. 60–62. ISSN: 0003-4851. DOI: 10.1214/aoms/1177732360. URL: <http://projecteuclid.org/euclid.aoms/1177732360> (visited on 11/07/2022).
- [91] Glen Cowan et al. *Power-Constrained Limits*. arXiv:1105.3166 [hep-ex, physics:physics]. May 2011. DOI: 10.48550/arXiv.1105.3166. URL: <http://arxiv.org/abs/1105.3166> (visited on 11/07/2022).
- [92] The LUX-ZEPLIN et al. “Projected sensitivity of the LUX-ZEPLIN (LZ) experiment to the two-neutrino and neutrinoless double beta decays of ^{134}Xe ”. In: *Physical Review C* 104.6 (Dec. 2021). arXiv:2104.13374 [nucl-ex, physics:physics], p. 065501. ISSN: 2469-9985, 2469-9993. DOI: 10.1103/PhysRevC.104.065501. URL: <http://arxiv.org/abs/2104.13374> (visited on 11/08/2022).
- [93] D. S. Akerib et al. “Projected sensitivity of the LUX-ZEPLIN experiment to the $0\nu\beta\beta$ decay of ^{136}Xe ”. In: *Physical Review C* 102.1 (2020), p. 014602. ISSN: 2469-9985, 2469-9993. DOI: 10.1103/PhysRevC.102.014602. URL: <http://arxiv.org/abs/1912.04248> (visited on 11/08/2022).
- [94] D. S. Akerib et al. *Enhancing the sensitivity of the LUX-ZEPLIN (LZ) dark matter experiment to low energy signals*. arXiv:2101.08753 [astro-ph, physics:physics]. Jan. 2021. URL: <http://arxiv.org/abs/2101.08753> (visited on 11/08/2022).

- [95] Dev Ashish Khaitan. “Supernova Neutrino Detection in LZ”. In: *Journal of Instrumentation* 13.02 (Feb. 2018). arXiv:1801.05651 [astro-ph], pp. C02024–C02024. ISSN: 1748-0221. DOI: 10.1088/1748-0221/13/02/C02024. URL: <http://arxiv.org/abs/1801.05651> (visited on 11/10/2022).
- [96] Rafael F. Lang et al. “Supernova neutrino physics with xenon dark matter detectors: A timely perspective”. In: *Physical Review D* 94.10 (Nov. 2016). arXiv:1606.09243 [astro-ph, physics:hep-ex, physics:hep-ph], p. 103009. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.94.103009. URL: <http://arxiv.org/abs/1606.09243> (visited on 11/10/2022).
- [97] LUX Collaboration et al. “Discrimination of electronic recoils from nuclear recoils in two-phase xenon time projection chambers”. In: *Physical Review D* 102.11 (Dec. 2020). Publisher: American Physical Society, p. 112002. DOI: 10.1103/PhysRevD.102.112002. URL: <https://link.aps.org/doi/10.1103/PhysRevD.102.112002> (visited on 09/23/2022).
- [98] E. Aprile et al. “Design and Performance of the XENON10 Dark Matter Experiment”. In: *Astroparticle Physics* 34.9 (Apr. 2011). arXiv:1001.2834 [astro-ph], pp. 679–698. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2011.01.006. URL: <http://arxiv.org/abs/1001.2834> (visited on 09/23/2022).
- [99] E. Aprile et al. “XENON1T Dark Matter Data Analysis: Signal Reconstruction, Calibration and Event Selection”. In: *Physical Review D* 100.5 (Sept. 2019), p. 052014. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.100.052014. URL: <http://arxiv.org/abs/1906.04717> (visited on 09/23/2022).
- [100] Yue Meng et al. “Dark Matter Search Results from the PandaX-4T Commissioning Run”. In: *Physical Review Letters* 127.26 (Dec. 2021). arXiv:2107.13438 [hep-ex, physics:physics], p. 261802. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.127.261802. URL: <http://arxiv.org/abs/2107.13438> (visited on 09/23/2022).
- [101] The DarkSide Collaboration et al. “DarkSide-50 532-day Dark Matter Search with Low-Radioactivity Argon”. In: *Physical Review D* 98.10 (Nov. 2018). arXiv:1802.07198 [astro-ph], p. 102006. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.98.102006. URL: <http://arxiv.org/abs/1802.07198> (visited on 09/23/2022).
- [102] B Rebel et al. “High voltage in noble liquids for high energy physics”. In: *Journal of Instrumentation* 9.08 (Aug. 2014), T08004–T08004. ISSN: 1748-0221. DOI: 10.1088/1748-0221/9/08/T08004. URL: <https://iopscience.iop.org/article/10.1088/1748-0221/9/08/T08004> (visited on 08/17/2021).
- [103] E. Aprile et al. “Measurements of proportional scintillation and electron multiplication in liquid xenon using thin wires”. In: *Journal of Instrumentation* 9.11 (Nov. 2014). arXiv: 1408.6206, P11012–P11012. ISSN: 1748-0221. DOI: 10.1088/1748-0221/9/11/P11012. URL: <http://arxiv.org/abs/1408.6206> (visited on 09/16/2021).

- [104] L. Tvrznikova et al. “Direct comparison of high voltage breakdown measurements in liquid argon and liquid xenon”. In: *Journal of Instrumentation* 14.12 (Dec. 2019). arXiv: 1908.06888, P12018–P12018. ISSN: 1748-0221. DOI: 10.1088/1748-0221/14/12/P12018. URL: <http://arxiv.org/abs/1908.06888> (visited on 08/16/2021).
- [105] Lucie Tvrznikova. “Sub-GeV Dark Matter Searches and Electric Field Studies for the LUX and LZ Experiments”. arXiv: 1904.08979. PhD thesis. Yale University, 2019. URL: <http://arxiv.org/abs/1904.08979> (visited on 08/17/2021).
- [106] *LARGE PHOTOCATHODE AREA PHOTOMULTIPLIER TUBES*. URL: https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/LARGE_AREA_PMT_TPMH1376E.pdf (visited on 09/23/2022).
- [107] *CR-150-R5 evaluation board: application guide*. URL: <https://www.cremat.com/CR-150-R5.pdf> (visited on 09/23/2022).
- [108] *PHOTOMULTIPLIER TUBES: Basics and Applications*. en. 2007.
- [109] E. Bodnia et al. “The Electric Field Dependence of Single Electron Emission in the PIXeY Two-Phase Xenon Detector”. In: *Journal of Instrumentation* 16.12 (Dec. 2021). arXiv: 2101.03686, P12015. ISSN: 1748-0221. DOI: 10.1088/1748-0221/16/12/P12015. URL: <http://arxiv.org/abs/2101.03686> (visited on 01/11/2022).
- [110] N. R. Lomb. “Least-squares frequency analysis of unequally spaced data”. en. In: *Astrophysics and Space Science* 39.2 (Feb. 1976), pp. 447–462. ISSN: 1572-946X. DOI: 10.1007/BF00648343. URL: <https://doi.org/10.1007/BF00648343> (visited on 09/06/2022).
- [111] T. Caldwell, S. Seibert, and S. Jaditz. “Characterization of the R5912-02 MOD photomultiplier tube at cryogenic temperatures”. en. In: *Journal of Instrumentation* 8.09 (Sept. 2013). Publisher: IOP Publishing, pp. C09004–C09004. ISSN: 1748-0221. DOI: 10.1088/1748-0221/8/09/C09004. URL: <https://doi.org/10.1088/1748-0221/8/09/c09004> (visited on 09/18/2022).
- [112] E. Armengaud et al. “Searching for low-mass dark matter particles with a massive Ge bolometer operated above-ground”. In: *Physical Review D* 99.8 (Apr. 2019). arXiv:1901.03588 [astro-ph, physics:physics], p. 082003. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.99.082003. URL: <http://arxiv.org/abs/1901.03588> (visited on 09/18/2022).
- [113] Norbert Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications* —. 1964. ISBN: 978-0-262-25719-0. URL: <https://ieeexplore.ieee.org/book/6267356> (visited on 09/20/2022).
- [114] William Hadley Richardson. “Bayesian-Based Iterative Method of Image Restoration*”. EN. In: *JOSA* 62.1 (Jan. 1972). Publisher: Optica Publishing Group, pp. 55–59. DOI: 10.1364/JOSA.62.000055. URL: <https://opg.optica.org/josa/abstract.cfm?uri=josa-62-1-55> (visited on 09/20/2022).

- [115] James A. Nikkel, W. Hugh Lippincott, and Daniel N. McKinsey. “Demonstration of photomultiplier tube operation at 29 K”. In: *Journal of Instrumentation* 2.11 (Nov. 2007). arXiv:astro-ph/0702202, P11004–P11004. ISSN: 1748-0221. DOI: 10.1088/1748-0221/2/11/P11004. URL: <http://arxiv.org/abs/astro-ph/0702202> (visited on 10/02/2022).
- [116] R. Acciarri et al. “Oxygen contamination in liquid Argon: combined effects on ionization electron charge and scintillation light”. en. In: *Journal of Instrumentation* 5.05 (May 2010). Publisher: IOP Publishing, P05003–P05003. ISSN: 1748-0221. DOI: 10.1088/1748-0221/5/05/P05003. URL: <https://doi.org/10.1088/1748-0221/5/05/p05003> (visited on 10/02/2022).
- [117] Steven D. Biller et al. “Effects of oxygen and nitrogen on drifting electrons in a liquid argon TPC”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 276.1 (Mar. 1989), pp. 144–150. ISSN: 0168-9002. DOI: 10.1016/0168-9002(89)90627-X. URL: <https://www.sciencedirect.com/science/article/pii/016890028990627X> (visited on 10/02/2022).
- [118] R. Acciarri et al. “Effects of Nitrogen contamination in liquid Argon”. In: *Journal of Instrumentation* 5.06 (June 2010). arXiv:0804.1217 [nucl-ex], P06003–P06003. ISSN: 1748-0221. DOI: 10.1088/1748-0221/5/06/P06003. URL: <http://arxiv.org/abs/0804.1217> (visited on 10/02/2022).
- [119] R. Acciarri et al. “Liquid Argon Dielectric Breakdown Studies with the MicroBooNE Purification System”. In: *Journal of Instrumentation* 9.11 (Nov. 2014). arXiv:1408.0264 [physics], P11001–P11001. ISSN: 1748-0221. DOI: 10.1088/1748-0221/9/11/P11001. URL: <http://arxiv.org/abs/1408.0264> (visited on 09/23/2022).
- [120] LUX Collaboration et al. “3D Modeling of Electric Fields in the LUX Detector”. In: *Journal of Instrumentation* 12.11 (Nov. 2017). arXiv:1709.00095 [hep-ex, physics:physics], P11022–P11022. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/11/P11022. URL: <http://arxiv.org/abs/1709.00095> (visited on 10/05/2022).
- [121] Martin S. Alnæs et al. “The FEniCS Project Version 1.5”. In: *Archive of Numerical Software* 3.100 (2015). DOI: 10.11588/ans.2015.100.20553.
- [122] Gagandeep Singh. “Short Introduction to Finite Element Method”. en. In: *TDT24* (2010), p. 23.
- [123] Riccardo Bevilacqua et al. “A procedure for the characterization of electron transmission through Frisch grids”. In: *Nuclear Instruments & Methods in Physics Research. Section A: Accelerators, Spectrometers, Detectors, and Associated Equipment* 770 (2015), pp. 64–67. ISSN: 0168-9002. DOI: 10.1016/j.nima.2014.10.003.
- [124] Ryan Edward Linehan et al. “High voltage electrode development and the LZ experiment’s WIMP search”. PhD thesis. Stanford, California: Stanford University, 2022.
- [125] Walter Blum, Werner Riegler, and Luigi Rolandi. *Particle Detection with Drift Chambers*. Springer, 2008. ISBN: 978-3-540-76683-4.

- [126] Kirk McDonald. (PDF) *Notes on Electrostatic Wire Grids*. URL: https://www.researchgate.net/publication/246980598_Notes_on_Electrostatic_Wire_Grids (visited on 10/02/2022).
- [127] Ralph L. Amey and Robert H. Cole. “Dielectric Constants of Liquefied Noble Gases and Methane”. en. In: *The Journal of Chemical Physics* 40.1 (Jan. 1964), pp. 146–148. ISSN: 0021-9606, 1089-7690. DOI: 10.1063/1.1724850. URL: <http://aip.scitation.org/doi/10.1063/1.1724850> (visited on 09/27/2022).
- [128] James R. Baker-Jarvis et al. “Dielectric and Conductor-Loss Characterization and Measurements on Electronic Packaging Materials”. en. In: *NIST* (2001). Last Modified: 2021-10-12T11:10-04:00 Publisher: James R. Baker-Jarvis, Michael D. Janezic, Billy F. Riddle, Christopher L. Holloway, Nicholas Paulter, J Blendell. URL: nist.gov (visited on 09/27/2022).
- [129] *Supplier Data - Polyetheretherketone (PEEK) (Goodfellow)*. en. Section: Materials Article. Mar. 2003. URL: <https://www.azom.com/article.aspx?ArticleID=1882> (visited on 09/27/2022).
- [130] A. D. McDonald et al. “Electron drift and longitudinal diffusion in high pressure xenon-helium gas mixtures”. en. In: *Journal of Instrumentation* 14.08 (Aug. 2019). Publisher: IOP Publishing, P08009–P08009. ISSN: 1748-0221. DOI: 10.1088/1748-0221/14/08/P08009. URL: <https://doi.org/10.1088/1748-0221/14/08/p08009> (visited on 10/03/2022).
- [131] EXO-200 Collaboration et al. “Measurement of the Drift Velocity and Transverse Diffusion of Electrons in Liquid Xenon with the EXO-200 Detector”. In: *Physical Review C* 95.2 (Feb. 2017). arXiv:1609.04467 [hep-ex, physics:physics], p. 025502. ISSN: 2469-9985, 2469-9993. DOI: 10.1103/PhysRevC.95.025502. URL: <http://arxiv.org/abs/1609.04467> (visited on 10/03/2022).
- [132] O. Njaya et al. “Measurements of electron transport in liquid and gas Xenon using a laser-driven photocathode”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 972 (2020), p. 163965. ISSN: 0168-9002. DOI: 10.1016/j.nima.2020.163965. (Visited on 10/03/2022).
- [133] *Overview of materials for Polytetrafluoroethylene (PTFE), Extruded*. URL: www.matweb.com (visited on 10/07/2022).
- [134] Guan-Jun Zhang et al. “On the surface trapping parameters of polytetrafluoroethylene block”. en. In: *Applied Surface Science* 253.4 (Dec. 2006), pp. 1995–1998. ISSN: 0169-4332. DOI: 10.1016/j.apsusc.2006.03.082. (Visited on 10/07/2022).
- [135] O. Bunemann, T. E. Cranshaw, and J. A. Harvey. “Design of grid ionization chambers”. In: *Canadian Journal of Research* 27a.5 (Sept. 1949). Publisher: NRC Research Press, pp. 191–206. ISSN: 1923-4287. DOI: 10.1139/cjr49a-019. URL: <https://cdnsiencepub.com/doi/abs/10.1139/cjr49a-019> (visited on 12/06/2021).

- [136] L. W. Kastens et al. “Calibration of a Liquid Xenon Detector with Kr-83m”. In: *Physical Review C* 80.4 (Oct. 2009), p. 045809. ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.80.045809. URL: <http://arxiv.org/abs/0905.1766> (visited on 10/04/2022).
- [137] J. A. Nelder and R. Mead. “A Simplex Method for Function Minimization”. In: *The Computer Journal* 7.4 (Jan. 1965), pp. 308–313. ISSN: 0010-4620. DOI: 10.1093/comjnl/7.4.308. URL: <https://doi.org/10.1093/comjnl/7.4.308> (visited on 10/04/2022).
- [138] D. S. Akerib et al. “Simulations of events for the LUX-ZEPLIN (LZ) dark matter experiment”. en. In: *Astroparticle Physics* 125 (Feb. 2021), p. 102480. ISSN: 0927-6505. DOI: 10.1016/j.astropartphys.2020.102480. URL: <https://www.sciencedirect.com/science/article/pii/S0927650520300529> (visited on 10/09/2022).
- [139] S. Agostinelli et al. “Geant4—a simulation toolkit”. en. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (July 2003), pp. 250–303. ISSN: 0168-9002. DOI: 10.1016/S0168-9002(03)01368-8. URL: <https://www.sciencedirect.com/science/article/pii/S0168900203013688> (visited on 10/09/2022).
- [140] Yonit Hochberg. “SIMP Dark Matter”. en. In: *SciPost Physics Lecture Notes* (July 2022), p. 59. ISSN: 2590-1990. DOI: 10.21468/SciPostPhysLectNotes.59. URL: <https://scipost.org/10.21468/SciPostPhysLectNotes.59> (visited on 10/29/2022).
- [141] Philip Harris, Philip Schuster, and Jure Zupan. *Snowmass White Paper: New flavors and rich structures in dark sectors*. arXiv:2207.08990 [hep-ex, physics:hep-ph]. July 2022. URL: <http://arxiv.org/abs/2207.08990> (visited on 11/08/2022).
- [142] Kalliopi Petraki and Raymond R. Volkas. “Review of asymmetric dark matter”. In: *International Journal of Modern Physics A* 28.19 (2013), p. 1330028. ISSN: 0217-751X, 1793-656X. DOI: 10.1142/S0217751X13300287. URL: <http://arxiv.org/abs/1305.4939> (visited on 08/17/2022).
- [143] XENON Collaboration 7 et al. “Dark Matter Search Results from a One Ton-Year Exposure of XENON1T”. In: *Physical Review Letters* 121.11 (Sept. 2018). Publisher: American Physical Society, p. 111302. DOI: 10.1103/PhysRevLett.121.111302. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.121.111302> (visited on 11/18/2022).
- [144] The XENON collaboration et al. “Projected WIMP Sensitivity of the XENONnT Dark Matter Experiment”. In: *Journal of Cosmology and Astroparticle Physics* 2020.11 (Nov. 2020). arXiv:2007.08796 [astro-ph, physics:hep-ex, physics:physics], pp. 031–031. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2020/11/031. URL: <http://arxiv.org/abs/2007.08796> (visited on 11/18/2022).

- [145] Michela Lai. “Planck scale multi-scattering dark matter in DEAP-3600”. en. In: *Journal of Physics: Conference Series* 2156.1 (Dec. 2021), p. 012067. ISSN: 1742-6588, 1742-6596. DOI: 10.1088/1742-6596/2156/1/012067. URL: <https://iopscience.iop.org/article/10.1088/1742-6596/2156/1/012067> (visited on 08/17/2022).
- [146] A. D. Sakharov. “Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe”. In: *J. Exper. Theor. Phys.* 5 (1967), pp. 24–27.
- [147] Moira I. Gresham, Hou Keong Lou, and Kathryn M. Zurek. “Nuclear Structure of Bound States of Asymmetric Dark Matter”. In: *Physical Review D* 96.9 (Nov. 2017). arXiv:1707.02313 [hep-ph], p. 096012. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.96.096012. URL: <http://arxiv.org/abs/1707.02313> (visited on 11/19/2022).
- [148] A. Butcher et al. “Can Tonne-Scale Direct Detection Experiments Discover Nuclear Dark Matter?” In: *JCAP* 10 (2017). eprint: 1610.01840, p. 035. DOI: 10.1088/1475-7516/2017/10/035.
- [149] David G. Cerdeno and Anne M. Green. “Direct detection of WIMPs”. In: (Feb. 2010). eprint: 1002.1912, pp. 347–369. DOI: 10.1017/CB09780511770739.018.
- [150] Richard H. Helm. “Inelastic and Elastic Scattering of 187-Mev Electrons from Selected Even-Even Nuclei”. In: *Physical Review* 104.5 (Dec. 1956). Publisher: American Physical Society, pp. 1466–1475. DOI: 10.1103/PhysRev.104.1466. URL: <https://link.aps.org/doi/10.1103/PhysRev.104.1466> (visited on 11/03/2022).
- [151] Rocky W. Kolb and Andrew J. Long. “Superheavy dark matter through Higgs portal operators”. In: *Physical Review D* 96.10 (Nov. 2017). arXiv:1708.04293 [astro-ph, physics:hep-ph], p. 103540. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.96.103540. URL: <http://arxiv.org/abs/1708.04293> (visited on 08/29/2022).
- [152] Thomas Hambye et al. “Direct Detection is testing Freeze-in”. In: *Physical Review D* 98.7 (Oct. 2018). arXiv:1807.05022 [astro-ph, physics:hep-ph], p. 075017. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.98.075017. URL: <http://arxiv.org/abs/1807.05022> (visited on 09/13/2022).
- [153] Esteban Alcantara, Luis A. Anchordoqui, and Jorge F. Soriano. “Hunting for superheavy dark matter with the highest-energy cosmic rays”. In: *Physical Review D* 99.10 (May 2019). arXiv:1903.05429 [astro-ph, physics:hep-ph], p. 103016. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.99.103016. URL: <http://arxiv.org/abs/1903.05429> (visited on 09/13/2022).
- [154] Nicolas Bernal, Camilo Garcia-Cely, and Rogerio Rosenfeld. “WIMP and SIMP Dark Matter from the Spontaneous Breaking of a Global Group”. In: *Journal of Cosmology and Astroparticle Physics* 2015.04 (Apr. 2015). arXiv:1501.01973 [hep-ph], pp. 012–012. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2015/04/012. URL: <http://arxiv.org/abs/1501.01973> (visited on 10/29/2022).

- [155] Ivone F. M. Albuquerque and Laura Baudis. “Direct Detection Constraints on Superheavy Dark Matter”. In: *Physical Review Letters* 90.22 (June 2003). arXiv:astro-ph/0301188, p. 221301. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.90.221301. URL: <http://arxiv.org/abs/astro-ph/0301188> (visited on 10/29/2022).
- [156] Bradley J. Kavanagh. “Earth-Scattering of super-heavy Dark Matter: updated constraints from detectors old and new”. In: *Physical Review D* 97.12 (June 2018). arXiv:1712.04901 [astro-ph, physics:hep-ph], p. 123013. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.97.123013. URL: <http://arxiv.org/abs/1712.04901> (visited on 10/29/2022).
- [157] Jonathan H. Davis. “Probing sub-GeV mass SIMP dark matter with a low-threshold surface experiment”. In: *Physical Review Letters* 119.21 (Nov. 2017). arXiv:1708.01484 [astro-ph, physics:hep-ph], p. 211302. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.119.211302. URL: <http://arxiv.org/abs/1708.01484> (visited on 10/29/2022).
- [158] Yang Bai, Andrew J. Long, and Sida Lu. “Dark Quark Nuggets”. In: *Physical Review D* 99.5 (Mar. 2019). arXiv:1810.04360 [astro-ph, physics:hep-ph], p. 055047. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.99.055047. URL: <http://arxiv.org/abs/1810.04360> (visited on 10/29/2022).
- [159] Edward Witten. “Cosmic separation of phases”. In: *Physical Review D* 30.2 (July 1984). Publisher: American Physical Society, pp. 272–285. DOI: 10.1103/PhysRevD.30.272. URL: <https://link.aps.org/doi/10.1103/PhysRevD.30.272> (visited on 10/29/2022).
- [160] Ariel R. Zhitnitsky. ““Nonbaryonic” Dark Matter as Baryonic Color Superconductor”. In: *Journal of Cosmology and Astroparticle Physics* 2003.10 (Oct. 2003). arXiv:hep-ph/0202161, pp. 010–010. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2003/10/010. URL: <http://arxiv.org/abs/hep-ph/0202161> (visited on 10/30/2022).
- [161] Yang Bai and Andrew J. Long. “Six Flavor Quark Matter”. In: *Journal of High Energy Physics* 2018.6 (June 2018). arXiv:1804.10249 [astro-ph, physics:hep-ph], p. 72. ISSN: 1029-8479. DOI: 10.1007/JHEP06(2018)072. URL: <http://arxiv.org/abs/1804.10249> (visited on 10/30/2022).
- [162] Abhishek Atreya, Anjishnu Sarkar, and Ajit M. Srivastava. “Reviving quark nuggets as a candidate for dark matter”. In: *Physical Review D* 90.4 (Aug. 2014), p. 045010. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.90.045010. URL: <http://arxiv.org/abs/1405.6492> (visited on 10/30/2022).
- [163] David M. Jacobs, Glenn D. Starkman, and Bryan W. Lynn. “Macro Dark Matter”. In: *Monthly Notices of the Royal Astronomical Society* 450.4 (July 2015). arXiv:1410.2236 [astro-ph, physics:hep-ph], pp. 3418–3430. ISSN: 1365-2966, 0035-8711. DOI: 10.1093/mnras/stv774. URL: <http://arxiv.org/abs/1410.2236> (visited on 10/30/2022).

- [164] M.-M. Bé et al. *Table of Radionuclides*. Vol. 8. Monographie BIPM-5. Pavillon de Breteuil, F-92310 Sèvres, France: Bureau International des Poids et Mesures, 2016. ISBN: 978-92-822-2264-5. URL: http://www.bipm.org/utills/common/pdf/monographieRI/Monographie_BIPM-5_Tables_Vol8.pdf.
- [165] M.J. Berger et al. *XCOM: Photon Cross Section Database (version 1.5)*. 2010. URL: <http://physics.nist.gov/xcom> (visited on 10/30/2022).
- [166] O. Klein and Y. Nishina. “Über die Streuung von Strahlung durch freie Elektronen nach der neuen relativistischen Quantendynamik von Dirac”. de. In: *Zeitschrift für Physik* 52.11 (Nov. 1929), pp. 853–868. ISSN: 0044-3328. DOI: 10.1007/BF01366453. URL: <https://doi.org/10.1007/BF01366453> (visited on 10/31/2022).
- [167] Javier F. Acevedo and Joseph Bramante. “Supernovae Sparked By Dark Matter in White Dwarfs”. In: *Physical Review D* 100.4 (Aug. 2019). arXiv:1904.11993 [astro-ph, physics:hep-ph], p. 043020. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.100.043020. URL: <http://arxiv.org/abs/1904.11993> (visited on 11/12/2022).
- [168] Amit Bhoonah et al. “Detecting Composite Dark Matter with Long Range and Contact Interactions in Gas Clouds”. In: *arXiv:2010.07240 [astro-ph, physics:hep-ph]* (Oct. 2020). arXiv: 2010.07240. URL: <http://arxiv.org/abs/2010.07240> (visited on 06/22/2021).
- [169] P. B. Price and M. H. Salamon. “Search for Supermassive Magnetic Monopoles Using Mica Crystals”. In: *Physical Review Letters* 56.12 (Mar. 1986). Publisher: American Physical Society, pp. 1226–1229. DOI: 10.1103/PhysRevLett.56.1226. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.56.1226> (visited on 11/12/2022).
- [170] R. Bernabei et al. “Extended Limits on Neutral Strongly Interacting Massive Particles and Nuclearites from NaI(Tl) Scintillators”. In: *Physical Review Letters* 83.24 (Dec. 1999). Publisher: American Physical Society, pp. 4918–4921. DOI: 10.1103/PhysRevLett.83.4918. URL: <https://link.aps.org/doi/10.1103/PhysRevLett.83.4918> (visited on 11/16/2022).
- [171] Michael Clark et al. “Direct Detection Limits on Heavy Dark Matter”. In: *Physical Review D* 102.12 (Dec. 2020). arXiv: 2009.07909, p. 123026. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.102.123026. URL: <http://arxiv.org/abs/2009.07909> (visited on 08/30/2021).
- [172] Christopher V. Cappiello, J. I. Collar, and John F. Beacom. “New Experimental Constraints in a New Landscape for Composite Dark Matter”. In: *Physical Review D* 103.2 (Jan. 2021). arXiv:2008.10646 [astro-ph, physics:hep-ex, physics:hep-ph], p. 023019. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.103.023019. URL: <http://arxiv.org/abs/2008.10646> (visited on 10/29/2022).
- [173] Paul Reuss. “Cauchy’s theorem and generalization”. en. In: *EPJ Nuclear Sciences & Technologies* 4 (2018), p. 50. ISSN: 2491-9292. DOI: 10.1051/epjn/2018010. URL: <https://www.epj-n.org/10.1051/epjn/2018010> (visited on 11/02/2022).

- [174] Joseph Bramante et al. “Saturated Overburden Scattering and the Multiscatter Frontier: Discovering Dark Matter at the Planck Mass and Beyond”. In: *Physical Review D* 98.8 (Oct. 2018). arXiv: 1803.08044, p. 083516. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.98.083516. URL: <http://arxiv.org/abs/1803.08044> (visited on 08/17/2021).
- [175] N. Wyn Evans, Ciaran A. J. O’Hare, and Christopher McCabe. “Refinement of the standard halo model for dark matter searches in light of the Gaia Sausage”. In: *Phys. Rev. D* 99.2 (2019). eprint: 1810.11468, p. 023012. DOI: 10.1103/PhysRevD.99.023012.
- [176] M. Szydagis et al. “NEST: A Comprehensive Model for Scintillation Yield in Liquid Xenon”. In: *Journal of Instrumentation* 6.10 (Oct. 2011). arXiv:1106.1613 [physics], P10002–P10002. ISSN: 1748-0221. DOI: 10.1088/1748-0221/6/10/P10002. URL: <http://arxiv.org/abs/1106.1613> (visited on 11/01/2022).
- [177] Matthew C. Digman et al. “(Not as) Big as a Barn: Upper Bounds on Dark Matter-Nucleus Cross Sections”. In: *Physical Review D* 100.6 (Sept. 2019). arXiv: 1907.10618, p. 063013. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.100.063013. URL: <http://arxiv.org/abs/1907.10618> (visited on 08/30/2021).
- [178] Adam M. Dziewonski and Don L. Anderson. “Preliminary reference Earth model”. en. In: *Physics of the Earth and Planetary Interiors* 25.4 (June 1981), pp. 297–356. ISSN: 0031-9201. DOI: 10.1016/0031-9201(81)90046-7. URL: <https://www.sciencedirect.com/science/article/pii/0031920181900467> (visited on 09/13/2022).
- [179] A. Blaauw et al. “The New I.A.U. System of Galactic Coordinates (1958 Revision)”. In: *Monthly Notices of the Royal Astronomical Society* 121.2 (Aug. 1960). eprint: <https://academic.oup.com/mnras/article-pdf/121/2/123/8078181/mnras121-0123.pdf>, pp. 123–131. ISSN: 0035-8711. DOI: 10.1093/mnras/121.2.123. URL: <https://doi.org/10.1093/mnras/121.2.123>.
- [180] Katayun J. Kamdin. “A Search for Lightly Ionizing Particles in the LUX Detector and Research and Development For Future Liquid Xenon Time Projection Chambers”. en. PhD thesis. UC Berkeley, 2018. URL: <https://escholarship.org/uc/item/9ck617mz> (visited on 11/01/2022).
- [181] Paul Terman. *The Search for Lightly Ionizing Particles in the Large Underground Xenon Detector - ProQuest*. en. URL: <https://www.proquest.com/openview/413c5ffa6c6e5e30256d2a651dbb68d0/1?pq-origsite=gscholar&cbl=18750&diss=y> (visited on 11/01/2022).
- [182] H.P. Gavin. *Total Least Squares*. 2017.
- [183] Nathaniel Greene. *Generalized Least-Squares Regressions I: Efficient Derivations*. 2013.

- [184] Milton Abramowitz and Irene A. Stegun. *Handbook of mathematical functions: with formulas, graphs and mathematical tables [conference under the auspices of the National science foundation and the Massachusetts institute of technology]*. eng. 1964th ed. Dover books on advanced mathematics. New York: Dover publ, 1972. ISBN: 978-0-486-61272-0.
- [185] Michael Kuhlen, Mariangela Lisanti, and David N. Spergel. “Direct Detection of Dark Matter Debris Flows”. In: *Phys. Rev. D* 86 (2012). eprint: 1202.0007, p. 063505. DOI: 10.1103/PhysRevD.86.063505.
- [186] Nilanjan Banik et al. “Probing the nature of dark matter particles with stellar streams”. In: *Journal of Cosmology and Astroparticle Physics* 2018.07 (July 2018), pp. 061–061. ISSN: 1475-7516. DOI: 10.1088/1475-7516/2018/07/061. (Visited on 11/04/2022).
- [187] *YAML Ain't Markup Language (YAML™) revision 1.2.2*. URL: <https://yaml.org/spec/1.2.2/> (visited on 11/04/2022).
- [188] Martin Ester, Hans-Peter Kriegel, and Xiaowei Xu. “A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise”. en. In: *KDD-96 Proceedings* (1996), p. 6.
- [189] A. Tomás et al. “Study and mitigation of spurious electron emission from cathodic wires in noble liquid time projection chambers”. en. In: *Astroparticle Physics* 103 (Dec. 2018), pp. 49–61. ISSN: 09276505. DOI: 10.1016/j.astropartphys.2018.07.001. URL: <https://linkinghub.elsevier.com/retrieve/pii/S0927650518300707> (visited on 08/17/2021).
- [190] D. S. Akerib et al. “Investigation of background electron emission in the LUX detector”. In: *Physical Review D* 102.9 (Nov. 2020). arXiv:2004.07791 [physics], p. 092004. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.102.092004. URL: <http://arxiv.org/abs/2004.07791> (visited on 12/11/2022).
- [191] Ping-Shine Shaw et al. “Ultraviolet characterization of integrating spheres”. EN. In: *Applied Optics* 46.22 (Aug. 2007). Publisher: Optica Publishing Group, pp. 5119–5128. ISSN: 2155-3165. DOI: 10.1364/AO.46.005119. URL: <https://opg.optica.org/ao/abstract.cfm?uri=ao-46-22-5119> (visited on 10/23/2022).
- [192] Patricia Pepple Williamson et al. “Revisiting the Classical Occupancy Problem”. en. In: *The American Statistician* 63.4 (Nov. 2009), pp. 356–360. ISSN: 0003-1305, 1537-2731. DOI: 10.1198/tast.2009.08104. URL: <http://www.tandfonline.com/doi/abs/10.1198/tast.2009.08104> (visited on 10/10/2022).
- [193] Peter E. Stott. “Advances in techniques for diagnosing fusion plasmas”. en. In: *Fusion Technology 1996*. Ed. by C. Varandas and F. Serra. Oxford: Elsevier, Jan. 1997, pp. 157–165. ISBN: 978-0-444-82762-3. DOI: 10.1016/B978-0-444-82762-3.50019-7. URL: <https://www.sciencedirect.com/science/article/pii/B9780444827623500197> (visited on 10/26/2022).

- [194] D. S. Akerib et al. “Extending light WIMP searches to single scintillation photons in LUX”. In: *Physical Review D* 101.4 (Feb. 2020). arXiv:1907.06272 [astro-ph, physics:hep-ex, physics:physics], p. 042001. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.101.042001. URL: <http://arxiv.org/abs/1907.06272> (visited on 12/12/2022).
- [195] Sakakibara. “COMPARISON OF FIVE EXACT CONFIDENCE INTERVALS FOR THE BINOMIAL PROPORTION”. en. In: *American Journal of Biostatistics* 4.1 (Jan. 2014), pp. 11–20. ISSN: 1948-9889. DOI: 10.3844/amjbsp.2014.11.20. URL: <http://thescipub.com/abstract/10.3844/amjbsp.2014.11.20> (visited on 10/10/2022).
- [196] John McDonald. *Handbook of Biological Statistics*. 2014. URL: biostathandbook.com (visited on 10/24/2022).
- [197] Trevor Hastie, Robert Tibshirani, and J. H. Friedman. *The elements of statistical learning: data mining, inference, and prediction: with 200 full-color illustrations*. Springer series in statistics. New York: Springer, 2001. ISBN: 978-0-387-95284-0.
- [198] J. Watson et al. *Study of dielectric breakdown in liquid xenon with the XeBrA experiment*. arXiv:2206.07854 [hep-ex, physics:nucl-ex, physics:physics]. Oct. 2022. DOI: 10.48550/arXiv.2206.07854. URL: <http://arxiv.org/abs/2206.07854> (visited on 12/12/2022).
- [199] Leo Breiman. *Classification and Regression Trees*. New York: Routledge, Oct. 2017. ISBN: 978-1-315-13947-0. DOI: 10.1201/9781315139470.
- [200] Andriy Burkov. *The hundred-page machine learning book*. eng. Polen: Andriy Burkov, 2019. ISBN: 978-1-9995795-0-0.
- [201] Gilles Louppe. *Understanding Random Forests: From Theory to Practice*. [stat]. June 2015. DOI: 10.48550/arXiv.1407.7502. (Visited on 10/27/2022).
- [202] Tianqi Chen and Carlos Guestrin. “XGBoost: A Scalable Tree Boosting System”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. arXiv:1603.02754 [cs]. Aug. 2016, pp. 785–794. DOI: 10.1145/2939672.2939785. URL: <http://arxiv.org/abs/1603.02754> (visited on 10/27/2022).
- [203] Erik Štrumbelj and Igor Kononenko. “Explaining prediction models and individual predictions with feature contributions”. en. In: *Knowledge and Information Systems* 41.3 (Dec. 2014), pp. 647–665. ISSN: 0219-3116. DOI: 10.1007/s10115-013-0679-x. URL: <https://doi.org/10.1007/s10115-013-0679-x> (visited on 11/12/2022).
- [204] R. G. Wilson. “Vacuum Thermionic Work Functions of Polycrystalline Be, Ti, Cr, Fe, Ni, Cu, Pt, and Type 304 Stainless Steel”. In: *Journal of Applied Physics* 37 (May 1966). ADS Bibcode: 1966JAP....37.2261W, pp. 2261–2267. ISSN: 0021-8979. DOI: 10.1063/1.1708797. URL: <https://ui.adsabs.harvard.edu/abs/1966JAP....37.2261W> (visited on 09/23/2022).

- [205] Wolfgang Tauchert, Helmut Jungblut, and Werner F. Schmidt. “Photoelectric determination of V_0 values and electron ranges in some cryogenic liquids”. In: *Canadian Journal of Chemistry* 55.11 (June 1977). Publisher: NRC Research Press, pp. 1860–1866. ISSN: 0008-4042. DOI: 10.1139/v77-260. URL: <https://cdnsciencepub.com/doi/10.1139/v77-260> (visited on 09/23/2022).
- [206] A. Bhattacharya. “Measurement of breakdown potentials and Townsend ionization coefficients for the Penning mixtures of neon and xenon”. In: (1976). DOI: 10.1103/PHYSREVA.13.1219.
- [207] A. A. Kruithof. “Townsend’s ionization coefficients for neon, argon, krypton and xenon”. In: *Physica* 7 (June 1940). ADS Bibcode: 1940Phy.....7..519K, pp. 519–540. DOI: 10.1016/S0031-8914(40)90043-X. URL: <https://ui.adsabs.harvard.edu/abs/1940Phy.....7..519K> (visited on 01/19/2022).
- [208] L. F. Berzak, S. E. Dorfman, and S. P. Smith. *Paschen’s Law in Air and Noble Gases*. Apr. 2006. URL: http://www-eng.lbl.gov/~shuman/XENON/REFERENCES&OTHER_MISC/paschen_report.pdf.
- [209] R. Massarczyk et al. “Paschen’s law studies in cold gases”. In: *Journal of Instrumentation* 12.06 (June 2017). arXiv: 1612.07170, P06019–P06019. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/06/P06019. URL: <http://arxiv.org/abs/1612.07170> (visited on 08/16/2021).
- [210] L. Norman et al. “Dielectric strength of noble and quenched gases for high pressure time projection chambers”. en. In: *The European Physical Journal C* 82.1 (Jan. 2022), p. 52. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-021-09894-z. URL: <https://doi.org/10.1140/epjc/s10052-021-09894-z> (visited on 01/18/2023).
- [211] Olivier B. Postel and Mark A. Cappelli. “Parametric study of the vacuum ultraviolet emission and electrical characteristics of a He–Xe microdischarge”. In: *Journal of Applied Physics* 89.9 (May 2001). Publisher: American Institute of Physics, pp. 4719–4726. ISSN: 0021-8979. DOI: 10.1063/1.1355696. URL: <https://aip.scitation.org/doi/10.1063/1.1355696> (visited on 01/26/2022).
- [212] J. M. Meek. “A Theory of Spark Discharge”. en. In: *Physical Review* 57.8 (Apr. 1940), pp. 722–728. ISSN: 0031-899X. DOI: 10.1103/PhysRev.57.722. URL: <https://link.aps.org/doi/10.1103/PhysRev.57.722> (visited on 12/06/2021).
- [213] *Electron emission in intense electric fields — Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*. URL: <https://royalsocietypublishing.org/doi/10.1098/rspa.1928.0091> (visited on 09/23/2022).
- [214] Louis Malter. “Thin Film Field Emission”. In: *Physical Review* 50.1 (July 1936). Publisher: American Physical Society, pp. 48–58. DOI: 10.1103/PhysRev.50.48. URL: <https://link.aps.org/doi/10.1103/PhysRev.50.48> (visited on 09/23/2022).

- [215] Evgeny G. Pogorelov et al. “Corrected field enhancement factor for the floating sphere model of carbon nanotube emitter”. In: *Journal of Applied Physics* 108.4 (Aug. 2010). Publisher: American Institute of Physics, p. 044502. ISSN: 0021-8979. DOI: 10.1063/1.3466992. URL: <https://aip.scitation.org/doi/full/10.1063/1.3466992> (visited on 02/14/2022).
- [216] E. G. Pogorelov, A. I. Zhbanov, and Y.-C. Chang. “Enhancement factor, electrostatic force and emission current in a nanoneedle emitter”. en. In: *EPL (Europhysics Letters)* 85.1 (Jan. 2009). Publisher: IOP Publishing, p. 17001. ISSN: 0295-5075. DOI: 10.1209/0295-5075/85/17001. URL: <https://doi.org/10.1209/0295-5075/85/17001> (visited on 09/24/2022).
- [217] R. M. Hill and L. A. Dissado. “Examination of the statistics of dielectric breakdown”. en. In: *Journal of Physics C: Solid State Physics* 16.22 (Aug. 1983). Publisher: IOP Publishing, pp. 4447–4468. ISSN: 0022-3719. DOI: 10.1088/0022-3719/16/22/018. URL: <https://doi.org/10.1088/0022-3719/16/22/018> (visited on 10/08/2022).
- [218] P. Bartko et al. “Effect of DC Voltage Ramp Rate on Breakdown in Ferrofluid Based on Transformer Oil”. en. In: *Acta Physica Polonica A* 137.5 (May 2020), pp. 970–972. ISSN: 1898-794X, 0587-4246. DOI: 10.12693/APhysPolA.137.970. URL: <http://przyrbwn.icm.edu.pl/APP/PDF/137/app137z5p117.pdf> (visited on 05/31/2022).
- [219] Mohammed Nedjar. “Weibull Statistics in Dielectric Strength of Oil-impregnated Pressboard under Ramped AC and DC Voltages”. In: *Journal of Energy and Power Engineering* 7 (July 2013). DOI: 10.17265/1934-8975/2013.12.021.
- [220] Pia-Kristina Fischer and Gerold A. Schneider. “Influence of the experimental setup and voltage ramp on the dielectric breakdown strength and breakdown site in borosilicate glass”. en. In: *Journal of the European Ceramic Society* 41.2 (Feb. 2021), pp. 1332–1341. ISSN: 0955-2219. DOI: 10.1016/j.jeurceramsoc.2020.09.060. URL: <https://www.sciencedirect.com/science/article/pii/S0955221920307871> (visited on 08/18/2021).
- [221] Bjoern Mieller. “Influence of test procedure on dielectric breakdown strength of alumina”. en. In: *Journal of Advanced Ceramics* 8.2 (June 2019), pp. 247–255. ISSN: 2226-4108, 2227-8508. DOI: 10.1007/s40145-018-0310-4. URL: <http://link.springer.com/10.1007/s40145-018-0310-4> (visited on 05/31/2022).
- [222] Anthomas Thomas et al. “The Effect of Voltage Ramp Rate on Dielectric Breakdown of Thin Film Polymers”. In: (Oct. 2006). Conference Name: APS Four Corners Section Meeting Abstracts ADS Bibcode: 2006APS..4CF.D1001T, p. D1.001. URL: <https://ui.adsabs.harvard.edu/abs/2006APS..4CF.D1001T> (visited on 05/31/2022).
- [223] M. Auger et al. “On the electric breakdown in liquid argon at centimeter scale”. In: *Journal of Instrumentation* 11.03 (Mar. 2016), P03017–P03017. ISSN: 1748-0221. DOI: 10.1088/1748-0221/11/03/P03017. URL: <https://iopscience.iop.org/article/10.1088/1748-0221/11/03/P03017> (visited on 08/17/2021).

- [224] N. Hayakawa et al. “Breakdown mechanism of liquid nitrogen viewed from area and volume effects”. In: *IEEE Transactions on Dielectrics and Electrical Insulation* 4.1 (Feb. 1997), pp. 127–134. ISSN: 10709878. DOI: 10.1109/94.590883. URL: <http://ieeexplore.ieee.org/document/590883/> (visited on 08/17/2021).
- [225] A Blatter et al. “Experimental study of electric breakdowns in liquid argon at centimeter scale”. In: *Journal of Instrumentation* 9.04 (Apr. 2014), P04006–P04006. ISSN: 1748-0221. DOI: 10.1088/1748-0221/9/04/P04006. URL: <https://iopscience.iop.org/article/10.1088/1748-0221/9/04/P04006> (visited on 08/17/2021).
- [226] K. H. Weber and H. S. Endicott. “Area Effect and Its Extremal Basis for the Electric Breakdown of Transformer Oil [includes discussion]”. In: *Transactions of the American Institute of Electrical Engineers. Part III: Power Apparatus and Systems* 75.3 (Jan. 1956), p. 4499314. ISSN: 0097-2460. DOI: 10.1109/AIEEPAS.1956.4499314. URL: <http://ieeexplore.ieee.org/document/4499314/> (visited on 10/22/2022).
- [227] H. Goshima et al. “Statistical analysis of area and volume effects on breakdown voltage in liquid nitrogen using Weibull distribution”. In: *Proceedings of 1994 IEEE International Symposium on Electrical Insulation*. ISSN: 1089-084X. June 1994, pp. 430–433. DOI: 10.1109/ELINSL.1994.401427.
- [228] J. Gerhold, M. Hubmann, and E. Telser. “DC-breakdown strength of liquid nitrogen under different voltage ramp conditions”. In: *Proceedings of 1999 IEEE 13th International Conference on Dielectric Liquids (ICDL'99) (Cat. No.99CH36213)*. Nara, Japan: IEEE, 1999, pp. 445–448. ISBN: 978-0-7803-4759-5. DOI: 10.1109/ICDL.1999.798968. URL: <http://ieeexplore.ieee.org/document/798968/> (visited on 08/17/2021).
- [229] J. Gerhold, M. Hubmann, and E. Telser. “About the size effect in LHe-breakdown”. In: *ICDL'96. 12th International Conference on Conduction and Breakdown in Dielectric Liquids*. Roma, Italy: IEEE, 1996, pp. 324–328. ISBN: 978-0-7803-3560-8. DOI: 10.1109/ICDL.1996.565499. URL: <http://ieeexplore.ieee.org/document/565499/> (visited on 10/22/2022).
- [230] N. S. Phan et al. “A study of DC electrical breakdown in liquid helium through analysis of the empirical breakdown field distributions”. In: *Journal of Applied Physics* 129.8 (Feb. 2021). arXiv: 2011.08844, p. 083301. ISSN: 0021-8979, 1089-7550. DOI: 10.1063/5.0037888. URL: <http://arxiv.org/abs/2011.08844> (visited on 08/17/2021).
- [231] J. Gerhold, M. Hubmann, and E. Telser. “Gap size effect on liquid helium breakdown”. In: *Cryogenics* 34.7 (Jan. 1994), pp. 579–586. ISSN: 00112275. DOI: 10.1016/0011-2275(94)90183-X. URL: <https://linkinghub.elsevier.com/retrieve/pii/S001122759490183X> (visited on 10/22/2022).

- [232] J. Gerhold. “Breakdown phenomena in liquid helium”. In: *IEEE Transactions on Electrical Insulation* 24.2 (Apr. 1989), pp. 155–166. ISSN: 00189367. DOI: 10.1109/14.90264. URL: <http://ieeexplore.ieee.org/document/90264/> (visited on 10/22/2022).
- [233] W. Rogowski. “Die elektrische Festigkeit am Rande des Plattenkondensators: Ein Beitrag zur Theorie der Funkenstrecken und Durchführungen”. de. In: *Archiv für Elektrotechnik* 12.1 (Jan. 1923), pp. 1–15. ISSN: 0003-9039, 1432-0487. DOI: 10.1007/BF01656573. URL: <http://link.springer.com/10.1007/BF01656573> (visited on 08/17/2021).
- [234] I. Brodie. “Studies of Field Emission and Electrical Breakdown Between Extended Nickel Surfaces in Vacuum”. en. In: *Journal of Applied Physics* 35.8 (Aug. 1964), pp. 2324–2332. ISSN: 0021-8979, 1089-7550. DOI: 10.1063/1.1702858. URL: <http://aip.scitation.org/doi/10.1063/1.1702858> (visited on 02/11/2022).
- [235] Vinod Kumar Gandi et al. “Effect of Electrode Profile and Polarity on Performance of Pressurized Sparkgap Switch”. en. In: *Plasma* 5.1 (Mar. 2022). Number: 1 Publisher: Multidisciplinary Digital Publishing Institute, pp. 130–145. ISSN: 2571-6182. DOI: 10.3390/plasma5010010. URL: <https://www.mdpi.com/2571-6182/5/1/10> (visited on 10/12/2022).
- [236] N. Trinh. “Electrode Design for Testing in Uniform Field Gaps”. In: *IEEE Transactions on Power Apparatus Systems* 99 (May 1980). ADS Bibcode: 1980ITPAS..99.1235T, pp. 1235–1242. DOI: 10.1109/TPAS.1980.319754. URL: <https://ui.adsabs.harvard.edu/abs/1980ITPAS..99.1235T> (visited on 10/12/2022).
- [237] Neera Jain and A.G. Alleyne. “Thermodynamics-based optimization and control of vapor-compression cycle operation: Optimization criteria”. In: vol. 1. Aug. 2011, pp. 1352–1357. DOI: 10.1115/DSCC2011-6088.
- [238] *Silicon Photomultipliers (SiPM), Low-Noise, Blue-Sensitive C-Series SiPM Sensors*. URL: <https://www.onsemi.com/pdf/datasheet/microc-series-d.pdf> (visited on 08/16/2021).
- [239] M. Auger et al. “A method to suppress dielectric breakdowns in liquid argon ionization detectors for cathode to ground distances of several millimeters”. en. In: *Journal of Instrumentation* 9.07 (July 2014). Publisher: IOP Publishing, P07023–P07023. ISSN: 1748-0221. DOI: 10.1088/1748-0221/9/07/P07023. URL: <https://doi.org/10.1088/1748-0221/9/07/p07023> (visited on 08/17/2021).
- [240] Anbang Sun, Chao Huo, and Jie Zhuang. “Formation mechanism of streamer discharges in liquids: a review”. In: *High Voltage* 1.2 (July 2016), pp. 74–80. ISSN: 2397-7264. DOI: 10.1049/hve.2016.0016. URL: <https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/hve.2016.0016> (visited on 08/16/2021).

- [241] D. S. Leonard et al. “A simple high-sensitivity technique for purity analysis of xenon gas”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 621.1-3 (Sept. 2010). arXiv: 1002.2742, pp. 678–684. ISSN: 01689002. DOI: 10.1016/j.nima.2010.04.152. URL: <http://arxiv.org/abs/1002.2742> (visited on 05/10/2022).
- [242] *Electropolishing Users Guide*. URL: <https://www.delstar.com/assets/pdf/epusersguide.pdf?r=false> (visited on 10/16/2022).
- [243] Khaled Nassar, Hordur Gunnarsson, and Mohamed Hegab. “Using Weibull Analysis for Evaluation of Cost and Schedule Performance”. In: *Journal of Construction Engineering and Management-asce - J CONSTR ENG MANAGE-ASCE* 131 (Dec. 2005). DOI: 10.1061/(ASCE)0733-9364(2005)131:12(1257).
- [244] E. L. Kaplan and Paul Meier. “Nonparametric Estimation from Incomplete Observations”. In: *Journal of the American Statistical Association* 53.282 (June 1958), pp. 457–481. ISSN: 0162-1459. DOI: 10.1080/01621459.1958.10501452. (Visited on 10/17/2022).
- [245] Robert Nisbet, Gary Miner, and Ken Yale. “Chapter 11 - Model Evaluation and Enhancement”. en. In: *Handbook of Statistical Analysis and Data Mining Applications (Second Edition)*. Ed. by Robert Nisbet, Gary Miner, and Ken Yale. Boston: Academic Press, Jan. 2018, pp. 215–233. ISBN: 978-0-12-416632-5. DOI: 10.1016/B978-0-12-416632-5.00011-6. URL: <https://www.sciencedirect.com/science/article/pii/B9780124166325000116> (visited on 10/19/2022).
- [246] B. Suerfu. “Polaris: a general-purpose, modular data acquisition framework”. In: *Journal of Instrumentation* 13.12 (Dec. 2018), T12004–T12004. ISSN: 1748-0221. DOI: 10.1088/1748-0221/13/12/T12004. URL: <https://iopscience.iop.org/article/10.1088/1748-0221/13/12/T12004> (visited on 10/20/2022).
- [247] Danielle Navarro. *Learning Statistics with R - A tutorial for Psychology students and other*. en. University of New South Wales, 2022. URL: [https://stats.libretexts.org/Bookshelves/Applied_Statistics/Book%3A_Learning_Statistics_with_R_-_A_tutorial_for_Psychology_Students_and_other_Beginners_\(Navarro\)/16%3A_Factorial_ANOVA/16.05%3A_The___F___test_as_a_model_comparison](https://stats.libretexts.org/Bookshelves/Applied_Statistics/Book%3A_Learning_Statistics_with_R_-_A_tutorial_for_Psychology_Students_and_other_Beginners_(Navarro)/16%3A_Factorial_ANOVA/16.05%3A_The___F___test_as_a_model_comparison) (visited on 10/20/2022).
- [248] V.M. Atrazhev et al. “Mechanisms of Impulse Breakdown in Liquid: The Role of Joule Heating and Formation of Gas Cavities”. In: *Plasma Science, IEEE Transactions on* 38 (Nov. 2010), pp. 2644–2651. DOI: 10.1109/TPS.2010.2046337.
- [249] M. J. Terry et al. “The densities of liquid argon, krypton xenon, oxygen, nitrogen, carbon monoxide methane, and carbon tetrafluoride along the orthobaric liquid curve”. en. In: *The Journal of Chemical Thermodynamics* 1.4 (July 1969), pp. 413–424. ISSN: 0021-9614. DOI: 10.1016/0021-9614(69)90072-X. URL: <https://www.sciencedirect.com/science/article/pii/002196146990072X> (visited on 10/16/2022).

XeBrA

Run	PMT Single-photon Gain ADCC*samples	SIPM Channel gain mV / pC	Anode Channel Gain mV / pC
4	n/a	0.836	900
7	n/a	0.836	952
8	n/a	0.992	680
9	65.3	7.9	372
10	217.3	7.6	688

Table 1: Calibrated Sensor gains for each run.

Dataset	Electrodes	Tilt [°]	Gap [mm]	Ramp Rate [V/s]	Pressure [bar]	N	Model	k	E_0 [kV/cm]	E_1 [kV/cm]	E_2 [kV/cm]
400	MP1	0.61	3.0	100	1.70	19	3	3.1 ± 2.3	2.8 ± 1.9	28.6 ± 1.8	31.4 ± 0.2
401	MP1	0.61	1.0	100	2.00	66	3	5.6 ± 3.8	37.5 ± 22.2	77.4 ± 21.3	114.9 ± 1.3
402	MP1	0.61	2.0	100	2.00	60	3	16.4 ± 12.3	78.4 ± 55.4	6.3 ± 54.9	84.7 ± 0.7
403	MP1	0.61	3.0	100	2.00	145	3	12.2 ± 0.0	77.1 ± 0.0	0.0 ± 0.0	77.1 ± 0.0
404	MP1	0.61	5.0	100	2.00	40	1	15.4 ± 1.9	71.1 ± 0.7	0.0 ± 0.0	71.1 ± 0.7
500	MP1	0.61	1.0	100	2.01	35	3	44.7 ± 8.1	83.5 ± 13.0	4.9 ± 13.0	88.4 ± 0.2
501	MP1	0.61	2.0	100	2.00	81	3	12.1 ± 2.8	82.4 ± 1.7	0.0 ± 0.0	82.4 ± 1.7
504	MP1	0.61	2.0	100	2.00	20	1	9.8 ± 1.7	61.0 ± 1.4	0.0 ± 0.0	61.0 ± 1.4
505	MP1	0.61	3.0	100	2.00	34	2	1.5 ± 0.2	9.4 ± 1.3	41.5 ± 0.3	50.9 ± 1.2
506	MP1	0.61	5.0	100	1.98	97	3	9.4 ± 12.6	29.0 ± 46.6	3.0 ± 15.6	32.0 ± 56.0
601	MP1	0.28	1.0	100	2.01	53	3	1.9 ± 0.4	20.2 ± 3.3	81.7 ± 1.6	101.8 ± 1.8
602	MP1	0.28	2.3	100	2.01	21	2	1.8 ± 0.2	14.5 ± 1.9	41.0 ± 0.5	55.5 ± 1.5
603	MP1	0.28	2.0	100	2.01	38	1	7.8 ± 1.0	77.2 ± 1.7	0.0 ± 0.0	77.2 ± 1.7
701	MP1	0.28	5.0	100	2.00	42	3	9.8 ± 16.3	53.1 ± 118.3	19.1 ± 119.9	72.1 ± 4.5
703	MP1	0.28	1.0	100	2.00	50	1	11.9 ± 1.4	111.9 ± 1.3	0.0 ± 0.0	111.9 ± 1.3
704	MP1	0.28	2.0	100	2.00	37	1	17.0 ± 2.2	91.5 ± 0.9	0.0 ± 0.0	91.5 ± 0.9
705	MP1	0.28	3.0	100	2.00	40	3	19.2 ± 2.7	85.7 ± 3.9	0.3 ± 2.9	86.0 ± 1.2
706	MP1	0.28	2.0	100	1.70	41	1	17.0 ± 2.1	90.9 ± 0.8	0.0 ± 0.0	90.9 ± 0.8
707	MP1	0.28	2.0	100	2.20	44	1	18.0 ± 2.1	97.9 ± 0.8	0.0 ± 0.0	97.9 ± 0.8
803	MP2	1.04	3.0	100	2.00	47	1	10.2 ± 1.2	79.2 ± 1.2	0.0 ± 0.0	79.2 ± 1.2
804	MP2	1.04	2.0	100	2.00	36	1	18.4 ± 2.3	75.3 ± 0.7	0.0 ± 0.0	75.3 ± 0.7
805	MP2	1.04	1.0	100	2.00	36	3	34.8 ± 3.9	78.4 ± 3.2	0.6 ± 3.2	79.0 ± 0.3
806	MP2	1.04	3.0	100	2.00	18	1	12.0 ± 2.0	67.2 ± 1.4	0.0 ± 0.0	67.2 ± 1.4
807	MP2	1.04	2.0	100	2.00	18	3	25.4 ± 22.9	80.3 ± 1.9	0.0 ± 1.7	80.3 ± 1.8
808	MP2	1.04	1.0	100	2.00	33	2	1.6 ± 0.1	3.6 ± 0.3	69.2 ± 0.1	72.8 ± 0.3
809	MP2	1.04	2.0	100	1.50	42	3	3.6 ± 0.7	7.8 ± 1.4	72.5 ± 1.0	80.4 ± 0.5
810	MP2	1.04	2.0	100	1.60	37	3	51.1 ± 23.3	66.7 ± 27.6	10.9 ± 27.5	77.6 ± 0.3
811	MP2	1.04	2.0	100	1.70	34	3	2.6 ± 0.6	6.6 ± 0.9	71.9 ± 0.5	78.5 ± 0.6
812	MP2	1.04	2.0	100	1.80	38	3	19.5 ± 9.6	59.9 ± 26.6	17.9 ± 26.4	77.8 ± 0.9
813	MP2	1.04	5.0	100	2.00	39	1	13.3 ± 1.6	83.6 ± 1.0	0.0 ± 0.0	83.6 ± 1.0
900	MP1	1.26	3.0	100	2.00	29	1	13.2 ± 2.0	86.5 ± 1.2	0.0 ± 0.0	86.5 ± 1.2
901	MP1	1.26	3.0	100	2.00	47	1	13.1 ± 1.4	96.5 ± 1.1	0.0 ± 0.0	96.5 ± 1.1
902	MP1	1.26	1.0	100	2.00	67	3	10.6 ± 4.0	134.0 ± 4.5	0.0 ± 0.0	134.0 ± 4.5
903	MP1	1.26	2.0	100	2.00	81	1	10.1 ± 0.9	96.2 ± 1.1	0.0 ± 0.0	96.2 ± 1.1
904	MP1	1.26	2.0	150	2.00	33	1	16.0 ± 2.0	106.2 ± 1.2	0.0 ± 0.0	106.2 ± 1.2
905	MP1	1.26	2.0	200	2.00	35	1	10.2 ± 1.3	98.1 ± 1.7	0.0 ± 0.0	98.1 ± 1.7
906	MP1	1.26	2.0	100	1.70	34	1	13.7 ± 1.7	107.9 ± 1.4	0.0 ± 0.0	107.9 ± 1.4
907	MP1	1.26	2.0	100	1.50	40	1	13.8 ± 1.6	104.0 ± 1.2	0.0 ± 0.0	104.0 ± 1.2
1000	Pass.	0.51	3.0	100	2.00	19	1	14.1 ± 2.4	78.9 ± 1.3	0.0 ± 0.0	78.9 ± 1.3
1001	Pass.	0.51	3.0	100	2.00	40	2	2.2 ± 0.4	28.6 ± 3.7	46.2 ± 1.9	74.8 ± 2.3
1002	Pass.	0.51	1.0	100	2.00	71	3	15.4 ± 2.5	148.6 ± 2.3	0.0 ± 0.1	148.6 ± 2.3
1003	Pass.	0.51	2.0	100	2.00	78	3	8.6 ± 3.8	103.6 ± 6.1	0.0 ± 0.0	103.6 ± 6.1
1004	Pass.	0.51	2.0	150	2.00	48	3	13.8 ± 11.3	113.5 ± 5.1	0.0 ± 0.0	113.5 ± 5.1
1005	Pass.	0.51	2.0	200	2.00	38	1	12.3 ± 1.6	117.5 ± 1.6	0.0 ± 0.0	117.5 ± 1.6
1006	Pass.	0.51	2.0	100	1.60	42	1	12.8 ± 1.5	102.0 ± 1.3	0.0 ± 0.0	102.0 ± 1.3
1007	Pass.	0.51	2.0	100	1.50	40	1	15.2 ± 1.8	117.0 ± 1.2	0.0 ± 0.0	117.0 ± 1.2
1008	Pass.	0.51	2.0	100	1.70	42	1	13.9 ± 1.7	122.7 ± 1.4	0.0 ± 0.0	122.7 ± 1.4

Table 2: XeBrA dataset results for the Weibull fits.