

ABSTRACT

Title of dissertation: **IDENTIFICATION OF TRACE KRYPTON IN
THE LUX-ZEPLIN DARK MATTER SEARCH**

John Silk, Doctor of Philosophy, 2023

Dissertation directed by: **Professor Carter Hall**
 University of Maryland

Searches for Weakly Interacting Massive Particles (WIMPs) carried out by liquid xenon time projection chambers (TPCs) require a careful accounting of all background sources. Because WIMPs are a leading dark matter candidate, their possible existence is of great interest to particle physicists, astrophysicists, and cosmologists. The LUX-ZEPLIN (LZ) detector has completed an initial science run finding no evidence for WIMP scattering events. The data excludes scattering cross sections above $6.5 \times 10^{-48} \text{ cm}^2$ for a WIMP mass of $30 \text{ GeV}/c^2$. Background contributions from the beta decay of dispersed ^{85}Kr were reduced prior to the initial science run using charcoal chromatography to remove trace krypton. Over 10 tonnes of xenon were processed, and a custom mass spectrometry system measured a final mass averaged krypton concentration of 123 ± 22 parts-per-quadrillion (ppq) $\frac{\text{gram } ^{nat}\text{Kr}}{\text{gram } ^{nat}\text{Xe}}$. A delayed coincidence $\beta - \gamma$ search was also conducted to identify rare decays from ^{85}Kr in the LZ WIMP search data. The 11.0 ± 4.0 identified events are equivalent to a concentration of 183 ± 67 ppq. The total background contribution from ^{85}Kr to the WIMP search region of interest is 30 ± 11 electron recoil events.

IDENTIFICATION OF TRACE KRYPTON IN THE LUX-ZEPLIN
DARK MATTER SEARCH

by

John Silk

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2023

Advisory Committee:

Professor Carter Hall, Chair/Advisor

Professor Xiangdong Ji

Professor Zackaria Chacko

Professor Jason Kahn

Dr. Jon Balajthy

© Copyright by

John Silk

2023

Dedication

In memory of James D. Silk.

Acknowledgments

First, I would like express my gratitude to my advisor, Carter Hall. Over the past 8 years your support and experience have been invaluable to my growth as a researcher. Many of the acknowledgements expressed here are due to the connections and community you shared with me. The research you conduct within the Experimental Nuclear Physics group and the LZ collaboration has set a quality standard that I will continue to strive for.

The Experimental Nuclear Physics group at UMD has always been and continues to be full of bright and eager scientists. Thank you to my fellow Jo(h)n's; Jon Balajthy and John Armstrong. Jon's readiness to share his thoughtful analysis and comprehensive hardware skills were foundational to my skills as an experimenter. John's ability to take on and resolve complex problems was essential in bringing many of our efforts to fruition. I also want to extend my thanks to Tim Edberg and Anwar Bhatti. Tim, the thoughtful questions, cheerfulness, and experience you brought to every project was invigorating. Anwar, your patient help navigating simulations and analyzing detector data was critical to the work presented here, thank you.

I must also thank my collaborators at SLAC. In particular, thank you to the krypton removal team: Dan Akerib, Christina Ignarra, Eric Miller, Drew Ames, and Andrew Naylor. I have learned so much from each of you. Working with you has been a highlight of my research experience. Together we tackled what was already an extremely challenging problem during a period of unimaginable uncertainty. Each of you have shared unique skills, knowledge, and smiles with me, thank you. Additionally, the LZ group at SLAC has had so many wonderful people that have offered their help; Tom, Steffen, TJ, Tommie, Tyler, Kelly, Ryan, Alden, Tomasz, Rachel, Kim, Maria Elena, Shaun, Gabe, and more, thank you all.

To the LZ team and personnel at SURF. I cannot possibly thank you enough. Everyone has contributed significantly to make this marvel of technology work.

Many thanks to my peers, Zishou Yang, Sara Nabili, and Eli Mizrachi for your friendship.

I would also like to thank Charlotte. It has been chaotic to say the least. Thank you for being calm, caring, and patient. You inspire me to be the best I can. I love you.

None of this would have been possible without the unwavering support of my family. I love you all. To my mother Susan, thank you for always being a constant source of unwavering support and strength. Your confidence in me has driven me when I needed it most. To my sister Katie, thank you for being my best friend. You have always helped me, even when I was unaware I needed it. To my Aunts, Uncles, and Cousins, your support and love has been more help than you could possibly know. Thank you all. To my father Jim, thank you for the countless hours spent together discussing physics. I will forever treasure the many nights we spent solving problems, comparing notes, and laughing at our mistakes. I could not have done this without you.

If I have failed to mention you here, know your support was appreciated. I'm lucky to have had such overwhelming support.

Table of Contents

Dedication	ii
Acknowledgements	iii
1 Identifying Dark Matter	35
1.1 Motion Within Clusters	35
1.2 Galactic Rotation Curves	37
1.2.1 A Universal Rotation Curve	39
1.3 Gravitational Lensing	40
1.3.1 Geodesics in Spacetime	40
1.3.2 The Lens Equation	43
1.4 Cosmology	53
1.4.1 General Relativity	54
1.4.2 Lambda-CDM Model	56
1.4.3 Cosmic microwave background	59
1.5 Dark Matter Candidates	62
1.5.1 Baryonic dark matter	62
1.5.2 MACHOs	62

1.5.3	Neutrinos	63
1.5.4	Axions	63
1.5.5	Weakly Interacting Massive Particles	65
1.6	Conclusion and Thesis Outline	67
2	Experimental Detection of WIMPs	69
2.1	Theory of WIMP Scattering	71
2.2	Experimental Overview	77
2.3	Crystal detectors	78
2.4	Bubble Chambers	81
2.5	Dual-Phase Time Projection Chambers	81
2.5.1	Detection Mediums	83
2.5.2	Xenon Microphysics	85
2.5.3	Energy Reconstruction	87
2.6	Conclusion	92
3	The LUX-ZEPLIN Detector	93
3.1	Active Vetoes	94
3.1.1	The Outer Detector	94
3.1.2	The Skin Detector	95
3.2	The Time Projection Chamber	97
3.3	Data Handling and Software	99
3.3.1	Simulations	101
3.4	Calibration	103
3.5	Backgrounds	104
3.5.1	Background Sources	105
3.5.2	Background	106

3.5.3	WIMP Search Selection Cuts and Backgrounds	110
3.6	WIMP Search Results	114
3.7	Conclusion	118
4	Xenon Purity Monitor	119
4.1	Cold-Trap Operating Principle	120
4.1.1	Ideal System Response	122
4.2	The Generic Xenon Purity Monitor	133
4.3	The RGA	138
4.3.1	Ionizer	138
4.3.2	Ion Filter	139
4.3.3	Charge Collector	140
4.3.4	RGA Conditioning	141
4.3.5	RGA Peak Tuning	141
4.3.6	RGA Service and Operating Parameters	142
4.4	Mass Spec Pathologies	143
4.4.1	RGA Saturation	143
4.4.2	System Contamination	146
4.5	Operating Procedures	150
4.5.1	Safe State	151
4.5.2	Initial State	151
4.5.3	Collect Xenon	152
4.5.4	System Warm-up	154
4.5.5	Form an Ice Kernel	154
4.5.6	Setting the 1x Impedance State	155
4.5.7	Set the RGA High Voltage	156

4.5.8	Setting the Yx Impedance State	158
4.5.9	Clean Run	160
4.5.10	Calibration Preparation from Reference Xenon	161
4.5.11	High-Sensitivity Measurement Procedure	164
4.5.12	Static Pump-out Procedure	166
4.6	Measurement Analysis	168
4.6.1	Low-Sensitivity Analysis	169
4.6.2	High-Sensitivity Analysis	170
4.7	Mixing Panel	173
4.7.1	Performance	175
4.7.2	Air Cocktail Mixture and Identification of Reference Xenon	179
4.8	Mobile and Stationary Sampling Systems	184
4.8.1	Parts Cleaning, Assembly, and Citric Etching	184
4.8.2	Mobile Sampling System Quality Assurance and Performance	185
4.8.3	Stationary Sampling System Quality Assurance and Performance	189
4.8.4	SSS Temperature Study	194
4.9	Conclusion	200
5	The LZ Krypton Removal System	201
5.1	Gas Charcoal Chromatography	202
5.2	System Overview	203
5.2.1	Chromatography Loop	204
5.2.2	Recovery Loop	207
5.2.3	Storage	209
5.2.4	Control	210
5.3	Sampling System	211

5.3.1	Sampler Run Control	213
5.3.2	Sampler Operation Control	217
5.3.3	Sampler Devices	219
5.3.4	Sampler Subsystems	220
5.3.4.1	Mixing Panel	220
5.3.4.2	Sniffer	223
5.3.4.3	Storage Line	226
5.3.4.4	Liquid Nitrogen Distribution	226
5.3.5	Run Viewer	228
5.3.6	Freezer Warm-up as a Distillation step	230
5.3.7	Sampler Usage	233
5.4	Results	234
6	Direct Observation of ^{85}Kr in LZ Science Run 1	237
6.1	What Will a Signal Look Like	239
6.2	Simulated Waveforms	239
6.3	Data Selection Variables	249
6.4	Determining Cut Values and Efficiencies via Simulations	251
6.4.1	Identifying Events in Simulations	253
6.4.2	Selection Criteria	256
6.4.3	Fiducial Cut	256
6.4.4	Energy Cuts	260
6.4.4.1	Validating Energy Cuts	266
6.4.5	Delay Time Window	292
6.4.6	Localization	293
6.4.7	ΔTBA	298

6.4.7.1	ΔTBA Validation	306
6.4.8	Skin Veto	309
6.4.9	Signal Efficiency	317
6.5	SR1 data	320
6.6	Side-band Analysis	322
6.7	Toy Monte Carlo Model	331
6.8	Final Event Hand Scanning	338
6.9	Conclusion	348
7	Summary	350
A	List of Publications	354

List of Figures

- 1.1 Mean rotational velocity curves of 21 spiral arm galaxies measured by Rubin and Ford. The dotted lines show velocities beyond the optical radius, calculated from the hydrogen 21 cm line, where 68% of the luminous matter is contained within. From reference [1]. 38
- 1.2 The Universal Rotation Curve surface for galaxies of arbitrary type. The radius from center of galaxy is normalized to its optical radius along x , total luminous mass along y , and resulting velocity along z . From reference [2]. 39
- 1.3 A photon approaching from $\phi = \pi$ and $r = -\infty$, with impact parameter b deviates from the straight path of $\phi = 0$ to some $\phi < 0$ at $r = \infty$. The deviation in the straight path is due to warped spacetime described by the Schwarzschild metric. From reference [3]. 44
- 1.4 Geometry for a gravitational lens system. A point mass M deflects light from source S at an angle ϕ to produce an image I as seen by observer O . The observer is a distance D_{OS} from the plane of the light source and image. The angle between the observer and the light source is β . The angle between the observer and the image is θ . The point mass lens is a distance D_{OL} from the observer and a distance D_{LS} the plane of the light source and image. From reference [4]. 46

1.5	A light source in line with a lensing mass produces an Einstein ring observed at angular radius θ_E . From reference [4].	47
1.6	Mass distribution of the Abell 1689 galaxy cluster. The red contours show the projected mass density from the outer to inner regions as $(1.6, 2.4, 4.0) \times 10^{10} M_{\odot} \text{arcsec}^{-2}$. The green dot is the mass peak. The blue circle is the location of the main dark matter clump. The smaller dark matter clump is in the top left section of the image. From reference [5].	49
1.7	The image distortion of a circular object by a gravitational lens for mass over (left) and under (right) densities. E-modes (top) denote the effect of a radially symmetric mass distribution within the gravitational lens, while the curl-like B-mode denotes the effect of complex mass distributions that result in optical shearing. Top left: $\gamma_1 > 0$. Top right: $\gamma_1 < 0$. Bottom left: $\gamma_2 > 0$. Bottom right: $\gamma_2 < 0$. From reference [6].	51
1.8	Hubble Space Telescope image of Cl 0024 + 17 with a dark mass cloud illustrated to highlight the ring and core structure. The structure is distinct compared to the galaxy and hot gas distributions. From reference [7].	52
1.9	Composite image of the Bullet cluster from the Chandra x-ray image showing the hot plasma distribution with the total reconstructed lensing mass distribution. The white contours show the location of the total surface density mass peak corresponding to the 1σ , 2σ , and 3σ confidence levels. The green contours are lines of constant surface mass density expressed in steps of 0.07κ . From reference [8].	53
1.10	Temperature variations of the cosmic microwave background with respect to the black-body temperature 2.73 kelvin. The very isotropic radiation has variations on the order of μkelvin . From reference [9].	60

1.11	The angular power spectrum of anisotropies from the CMB. The deviation from the mean temperature squared shows spikes as a function of the multipole. These features correspond to mass infall from gravity (odd peaks) and rarefaction due to radiation pressure (even peaks). Of particular prominence is the third peak ($\ell \approx 850$) suggesting a driving gravitational interaction that did not contribute a restorative pressure. From reference [10]. . . .	61
1.12	The comoving number density Y (left) and thermal relic density (right) of a 100 GeV WIMP as a function of temperature (bottom) or time (top). The dotted line shows the dashed line is the number density of particles that remain in thermal equilibrium. The solid line uses an annihilation cross section that yields the correct relic density. The yellow, green, and blue bands show cross sections that differ from the middle value by a factor of 10, 10^2 , and 10^3 respectively. From reference [11].	67
2.1	Schematic of WIMP-Standard Model interactions. The arrows indicated the direction of time for a given interaction. The “X” label denotes a WIMP and “SM” denotes a standard model particle. From reference [12].	70
2.2	The nuclear density as a function of radius for a sphere of effective radius r_n and skin thickness t . The density remains constant before falling to zero with a Gaussian edge. From reference [13].	76
2.3	The integrated, spin-independent, WIMP scattering rate for different atomic nuclei as a function of threshold energy. At low energies coherent scattering off of large nuclei, like xenon, is favorable. Event rates at higher energies are drastically suppressed in xenon due to decoherence. Germanium rates suffer modestly, but the lower mass targets are unaffected in this range. From reference [14].	77

2.4	A compilation of WIMP-nucleon spin-independent cross-section exclusion limits as a function of WIMP mass. Different theoretical dark matter models are shown in the shaded regions with no outlines. Solid curves are completed measurements and dotted lines are future projections as of 2013. The yellow region on the bottom is the neutrino fog. This fog will obscure direction detection measurements due to an irreducible neutrino background. From reference [15].	78
2.5	The reported spin-independent cross-section exclusion limits for a 60 GeV mass WIMP from direct detection experiments over time. In recent years, liquid noble detectors have consistently produced the most stringent limits. From reference [16].	82
2.6	An incoming particle scatters off of a target medium within the TPC and exits the detector. A prompt scintillation signal (S1) is produced and electrons are freed. An electric field drifts ionized electrons to the liquid surface. An extraction field pulls electrons into the gas phase. The electrons excite gaseous xenon and produce a secondary light signal (S2). Energy and position are reconstructed from signal timings and sizes. An example of the summed PMT waveforms, containing an S1 and S2, is shown on the right. Taken from reference [17].	83
2.7	Measurements of L_{eff} , referred to here as relative scintillation efficiency, for nuclear recoil energies below 100 keV. The solid dots represent measurements made by Manzur et al. using 2.8 MeV neutrons. Their theoretical model is shown as a dashed line. Included is data from Aprile (blue stars and brown circles) [18] [19] and Chepel (red ovals) [20]. From reference [21].	89

2.8	Measurements of the total quanta of light and charge produced during a nuclear recoil event from a 2.45 MeV neutron. The true recoil energy of the initial neutron recoil is given on the x axis and found using the scattering angle calculated from the location of a secondary recoil. The standard Lindhard and LUX specific Lindhard models are shown with a solid blue and dashed black line. The alternative parameterization is shown by the dotted grey line. A biexcitation quenching process is accounted for that reduces the total number of quanta at recoil energies above 10 keV due to increased exciton density [22].	91
3.1	A schematic of the LZ TPC and active veto detectors within the water tank. Starting from the outside, a cylindrical wall of PMTs view the acrylic vessels (shown in green) filled with gadolinium-loaded liquid scintillator. The inner cryostat is filled with liquid xenon. The skin detector has rings of PMTs monitoring the space between the TPC and cryostat walls. At the center is the TPC where two arrays of PMTs point towards the bulk xenon. Various conduits support infrastructure and the delivery of calibration sources. Taken from the LZ Technical Design Report [17].	94
3.2	A rendering of the lower cross section of the TPC and cryostat. The skin region is located between the cryostat and TPC walls. Lower side skin PMTs are aimed up into the skin region. Dome skin PMTs monitor light production near the bottom PMT array. Top skin PMTs are not depicted here. Taken from reference [17].	96

3.3	An external image of the LZ TPC in the above ground surface area lab following construction. The white PTFE field rings are stacked forming the vertical wall of the detector. The location of the different field regions and grids within the TPC are shown. The dark purple are regions submerged in liquid xenon and light purple regions are in the gaseous phase [23]. On the left an electric field map in the liquid region with updated voltage labels reflect the SR1 run parameters. The rainbow horizontal lines are equipotentials from 0 to to a 32 kV in steps of 3.2 kV. Taken from reference [17].	98
3.4	Block diagram for an S2 trigger event. The PMT information is stored in a buffer and actively filtered for an S2 trigger. When triggered, the DAQ saves the event window to disk. LZap searches the PMT waveforms for pulses and reconstructs the event. Finally an RQ list is produced and stored in a ROOT file.	100
3.5	Block diagram for the full chain simulation.	102
3.6	Block diagram for the fast chain simulation.	102
3.7	Calibration events in $S1$ vs $\log(S2_c)$ space from the LZ detector. The blue points are electron recoils from a tritium source. The orange events are neutron events from a deuterium-deuterium (DD) generator. Grey curves denote lines of constant energy and take different values for electron (keV_{ee}) and nuclear (keV_{nr}) recoils. The solid red (blue) line is the median of the simulated nuclear (electron) recoil band. The dotted lines indicate the 10% and 90% quantiles. Taken from reference [24].	104

3.8 Background rate as a function of energy in the innermost 1 tonne xenon volume following a DD calibration. Gamma and gamma+X-ray features sit atop a combined beta spectrum. Note the four peaks from the ^{127}Xe decays and electron cascades. Similarly the six peaks are noted from the ^{125}Xe decays and cascades. The few days that followed the DD calibration were removed from SR1 to allow the short lived isotopes, ^{133}Xe and ^{125}Xe , to decay to sub dominant rates. Taken from reference [25]. 108

3.9 Total fitted detector background spectra in the SR1 fiducial volume. The fit in red to the data in black is consistent with an average residual of approximately zero. Fluctuations at certain peaks arise from imperfect energy resolution matching between data and simulations. Energies below 80 keV_{ee} have been obscured as analysis of ^{124}Xe two-neutrino double electron capture is underway. Taken from reference [25]. 110

3.10 LZ background model before fitting to SR1 data (except for the ^{37}Ar component). The total background in the WIMP search region of interest is shown in blue with the WIMP search cut SR1 data in black. This represents an event rate of $(6.3 \pm 0.5) \times 10^{-5} \text{ events/keV}_{ee}/\text{kg}/\text{day}$. Taken from reference [25]. 113

3.11 The LZ background table with sources in the left column, expected events in the middle column, and fit results in the right column. While expected events were explicitly estimated for the various β decays, the fit uses a flat energy distribution to model all of the ^{214}Pb , ^{212}Pb , ^{85}Kr , and detector electron recoils. Taken from reference [25]. 114

3.12 Spatial distribution of WIMP search events during SR1 prior to the fiducial volume cut, OD veto, and skin veto. The detector volume is enclosed by a block dotted line and the fiducial volume is shown with a black solid line. Black dots within the fiducial volume make up the final 335 events within the WIMP search data set. Grey circles satisfy all cuts but occur outside of the fiducial volume. Red crosses are events vetoed by the liquid skin and blue circles are events vetoed by the OD. Taken from reference [24]. 115

3.13 The WIMP signal efficiency as a function of nuclear recoil energy for single scatter reconstructions in the search region of interest. The error assessed using composite nuclear recoil AmLi-tritium events. A 50% efficiency is achieved at 5.5 keV_{nr} and marked with a grey dotted line. Taken from reference [24]. 116

3.14 Low-energy recoil events after all data quality and physics cuts shown in the log₁₀(S2_c)-S1_c space. A shaded grey band, orange ellipses contours (³⁷Ar), shaded green region (⁸B solar neutrinos), and purple contours (30 GeV WIMP) are shown for 1σ and 2σ best-fit models. The solid red line shows the nuclear recoil median with the dotted lines indicating the 10% and 90% quantiles. Events that fall outside of the 2σ best-fit models are shown a pie chart with wedge size proportional to the relative weight of the background fit components corresponding to that region. Taken from reference [25]. 117

3.15	The 90% confidence limit for the spin-independent WIMP cross-section vs. WIMP mass. The dotted line is the median sensitivity projection. The green and yellow bands are the 1σ and 2σ sensitivity projections. Also included are results from the liquid xenon TPCs PandaX-4T, XENON1T, LUX, and the liquid argon scintillation only detector DEAP-3600. Argon scintillation signal shapes depend on the interaction type to discriminate electron recoils as described in section 2.5.2. Taken from reference [24].	118
4.1	A depiction of a xenon-krypton mixture flowing through a cold-trap submerged in a bath of liquid nitrogen. The xenon is frozen and the outlet is enriched with krypton. This image is taken from reference [12].	121
4.2	A depiction of the equilibrium of krypton flow throughout the system. The amount of xenon in the Z1 region is fixed, but the equilibrium quantity of krypton depends on flow rate into the cold-trap and the pumping speed at which it is depleted. Impedance values Z_1 and Z_2 set the pumping speed for each region and are controlled by IV and VV respectively. Note the S_{trap} pumping speed that drives krypton entrapment flow into ice.	124
4.3	The krypton throughput parameter α as a function of relative impedance from the 1x impedance state. Taken from reference [12].	126
4.4	Observed peak P_{Kr} , measured at increasing impedance. The data fits reasonably well with the equilibrium krypton pressure equation (4.21). Taken from reference [12].	129
4.5	The PP_{Kr} partial pressure as a function of time from flow start at a flow of 3.5 slpm. The krypton signal deviates from the expected shape at 100 seconds into the xenon flow as the ice formation reaches the outlet corner of the cold-trap. Taken from reference [12].	131

4.6	A stocking trap made from 0.5” stainless steel tubing. The inlet side of the cold-trap is on the left. The bottom of this trap is 9” wide and the toe section is 7” tall. The trap is submerged in liquid nitrogen such that the bottom is 1” above the bottom the dewar to allow maximal cooling. The liquid nitrogen level is maintained 2” above the top of the toe section. A similar trap with a 11” wide bottom span was used on the mixing panel. In this trap we attempt to form an ice sleeve in the left and bottom portions where α remains constant for the duration of the flow.	132
4.7	A schematic of a generic XPM. The system valve (SV) connects the XPM to a volume from which xenon samples are drawn. The XPM uses valves V0-V3 to deliver xenon from the source bottle to the cold trap. Xenon flow is regulated by the mass flow controller (MFC) or leak valve (LV). Xenon gas pressure is monitored by a pressure transducer (PT0). Valves V5 and V6 protect the vacuum turbo-pumps (TP1 and TP2), vacuum gauges (PT1 and PT2), and RGA. The impedance valve (IV) and vacuum valve (VV) set the impedance values Z_1 (yellow) and Z_2 (pink).	133
4.8	An cross-sectional image of a modified Swagelok 304L-HDF8-1GAL bottle. End caps have been machined to fit the 1/2” FNPT threads and welded to the 1 gallon bottle. The top cap includes a welded HVCR fitting and a female HVCR nut, the bottom cap is kept blank. A relief cut along the threads avoids trapped gas volumes. The bottle is leak tested and pressure tested following modification. Taken from the LZ controlled document database [26].	134
4.9	An SRS RGA200 probe schematic. From the RGA manual [27].	138
4.10	A cross section of the SRS RGA200 ionizer. From the RGA manual [27].	139
4.11	An SRS RGA200 quadrupole mass filter. From the RGA manual [27].	140

- 4.12 An analog scan of xenon ice. The x-axis mass is in AMU and the y-axis is partial pressure in picoTorr (pTorr). This analog scan was taken while the RGA was opened to xenon ice. Some easily identifiable peaks here are due to helium (4 AMU), water (18 AMU), nitrogen (28 AMU), CO₂ (44 AMU), and the doubly ionized xenon peaks (62 to 68 AMU). We also note that the mass channels where krypton would be (80 - 86), have no peaks. 142
- 4.13 An RGA pressure versus time scan tracks the partial pressure in pTorr of different mass channels over time. The different line colors represent different gases corresponding to the mass in AMU. The total pressure, monitored by a vacuum gauge, increased from 8.6×10^{-6} Torr to 3×10^{-5} Torr due a large concentration of helium (light blue). An accurate measurement was not possible due to the saturation. Flow became restricted due to the accumulation of helium in the trap. The clogging of the cold-trap or cryopump due a non-condensable gas is known as vapor-lock. 144
- 4.14 An RGA pressure versus time scan. The different line colors correspond to mass-to-charge ratios of a specific gas. The nitrogen (pink) signal saturated during flow, beginning at 1:38 pm. As a result the signals in all channels increased. The ⁶²Xe (blue) channel rose by approximately 30% despite being the dominant pressure and known to be constant. This increase in the ⁶²Xe remained high after flow stopped at 1:43 pm, as noted by the downturn in helium and argon signals. The pathology is observed with ⁸⁷Bl (black) and produced a ⁸⁷Bl signal that was over 30% the size of the signal observed in the ⁸⁴Kr channel (red). 145

- 4.15 Three RGA pressure vs time scans. In figure 4.15a a xenon sample containing 20 ppb of krypton causes an elevated background in the red ^{84}Kr and green ^{83}Kr channels. The measurement of this dirty sample resulted in contamination. The contamination is shown in figure 4.15b, through the continued divergence of ^{84}Kr from ^{87}Bl during a measurement of xenon that is known to contain less than 10 ppq of krypton. After two days of maintenance, the XPM baseline was recovered, as shown in figure 4.15c. The ^{84}Kr and ^{87}Bl remained stable and near zero pTorr throughout the measurement of the purified xenon. All measurements were carried out under identical settings and followed the procedure given in section 4.5.11. . . . 148
- 4.16 An analog scan of xenon ice. The RGA and vacuum components were contaminated by an unknown hydrocarbon that resulted in an elevated signals across all AMU channels. Most notably, when compared with figure 4.12, we see elevated signals between 45 to 60 AMU and 70 to 85 AMU. This contamination produces similar divergence in the ^{84}Kr and ^{87}Bl channels and impedes the XPM measurement capability. 149
- 4.17 An RGA pressure versus time (PvT) scan of the ^{62}Xe partial pressure. The impedance state was previously set to the 1x state in region 1. The average ^{62}Xe partial pressure is 1960 pTorr after setting the CEM high voltage. In region 2, IV is adjusted to increase impedance. After IV is adjusted, the average ^{62}Xe partial pressure in region 3 is 137 pTorr. The resulting impedance value is $1960/137 = 13.4x$. In region 4, VV is adjusted. This brings the ^{62}Xe partial pressure back up to the operating spec of the RGA. A final ^{62}Xe value of 2040 pTorr is achieved. 159

4.18	Helium measurements for three runs with a static pump-out after Run 1 and Run 2. The helium in Run 1 (blue) caused the RGA to saturate during this measurement giving a 2.8 ppm a lower limit. Following one static pump-out Run 2 (red) measured a helium concentration of 5.45 ppb of helium. Following a second static pump-out the helium in Run 3 (yellow) was undetectable.	167
4.19	A measurement of calibration xenon prepared with impurities from a small quantity of air. The measurement was completed at 13.4x impedance with a flow rate of 0.13 slpm, maintained between the two dotted black lines. A 5 minute background interval is highlighted in yellow and the signal window is highlighted in green. The prepared xenon mixture contained 1.43 ± 0.14 ppm of nitrogen and 24.5 ± 2.5 ppb of argon. The nitrogen and argon concentrations were reported as 1.48 ± 0.44 ppm and a 31 ± 9 ppb respectively.	169
4.20	A measurement of calibration xenon prepared to a krypton concentration of $\Phi_{cal} = 1.17$ ppt. The measurement was completed at 13.4x impedance and a flow rate of 2 slpm. The signal window begins 20 seconds after the flow starts and ends 50 seconds after the flow stops.	171
4.21	A picture of the mixing panel, sniffer, and recovery bottle. Additional components that supported operations of the SLAC XPM are installed with the mixing panel on the side of the SSS.	174
4.22	In the second phase of operation six calibration were completed with krypton concentrations around 1.1 ppt. The integrated signal for each channel is given along the x-axis and the prepared concentration, including errors, is given along the y-axis.	176

4.23 In the second phase of operations, 31 zero measurements were identified for a total of 93 data points. ^{84}Kr is shown in green, ^{83}Kr is stacked in red, and ^{87}Br is stacked in blue. The integrated signal has been converted to units of ppq using the slope value from the previous plot 4.22. The measurements taken show a distribution with a mean value 11 ± 28 ppq. The standard deviation of the mean is 2.9 ppq and the mean zero concentration is reported to be larger than zero. This could be due to insufficient cleaning prior to the zero measurement or fluctuations in the RGA baseline from environmental exposure. To account for this the full standard deviation is taken as the zero error σ_z 177

4.24 In the third phase of operations nine calibration were completed. The integrated signal for each channel is given along the x-axis and the prepared concentration, including errors, is given along the y-axis. 178

4.25 In the third phase of operation, 16 zero measurements were identified for a total of 48 data points. ^{84}Kr is shown in green, ^{83}Kr is stacked in red, and ^{87}Br is stacked in blue. The integrated signal has been converted to units of ppq using the slope value from the previous plot 4.24. The measurements taken show a distribution with a mean value 0 ± 18 ppq. The standard deviation of the mean is 4.5 ppq which is consistent with zero. Again, the full standard deviation is taken as the zero error σ_z 179

4.26 An RGA trace of the ^{84}Kr partial pressure for the xenon-air mixture followed by the reference candidate xenon. The system impedance was set to 1x and the flow rate was 1 slpm over a 90 second interval for each sample. Xenon was prepared using an air injection to a krypton concentration of $\Phi_{air} = 1.61 \pm 0.07$ ppb. The ratio of the background subtracted average signals yields a reference xenon concentration of $\Phi_{ref} = 217 \pm 12$ ppt. . . . 183

4.27	P&ID of the MSS ignition control panel implemented at SURF, courtesy of John Armstrong.	186
4.28	A Solidworks render of the MSS, courtesy of Timothy Edberg.	187
4.29	An as built picture of the MSS prior to shipping from UMD. Xenon source vessels were shipped directly to SURF and installed onsite after this image was taken.	188
4.30	A measurement of the ^{84}Kr channel from a calibration mixture on the MSS. The background interval is shown in yellow, the flow interval between the dotted black lines, and the delayed signal window in green. The signal window starts 10 seconds after flow starts and ends 200 seconds after flow stops.	189
4.31	P&ID of the SSS ignition control panel while at SLAC.	190
4.32	A Solidworks render of the SSS, courtesy of Timothy Edberg.	191
4.33	As built pictures of the SSS at SLAC. The system front panel is shown in figure 4.33a. The back panel with access to the cryostat is shown in figure 4.33b.	192
4.34	A measurement of the ^{83}Kr channel from a calibration mixture on the SSS. The background interval is shown in yellow, the flow interval between the dotted black lines, and the delayed signal window in green. For a 1577 ppq equivalent sample the implied limit of sensitivity is 23 ppq.	193
4.35	Copper cold trap and block coupled to the pulse tube refrigerator head. The block is instrumented with heaters and temperature sensors.	195

4.36	Xenon with a krypton concentration of 2933 ppq measured at different trap temperatures. The xenon ice vapor pressure is calculated using the model from equation (4.38). The vertical blue is at $x=77.15$ K, and the red dotted line is at $x=63.4$ K, where the model is expected to break down. The fit finds an RGA pumping speed of $S_{RGA} = 338$ cc/s and a trap pumping speed of $S_{trap} = 0.73$ cc/s.	197
4.37	Family of curves at varying S_{RGA} . The dotted black line is the curve from figure 4.8.4. The rainbow curves show the krypton equilibrium pressure as a function of temperature of as we change S_{RGA} by factors of 10. The red dotted line shows the temperature cutoff at 63.4 K, as P_{RGA} can no longer be held constant as P_{ICE} drops below 7.9×10^{-6} Torr.	199
5.1	A block diagram of the krypton removal system. Each of the three main subsystems are shown and connection points to the sampling system. Compressors, recirculation pumps, gas cylinders, and krypton traps (KT) are shown as well. This image is taken from reference [28].	204
5.2	An image of the krypton removal system at SLAC. The system is in a metal structure, semi-exposed to the outdoors. Not shown here are the helium re-circulation pump and fluitron compressor. This image is taken from reference [28].	205
5.3	Relative concentration data versus time from the BGA during a chromatography cycle. The output of the column is monitored from the start of xenon feed marked at the green time interval. This cycle contained process xenon with a very high krypton concentration (600 parts-per-million) to show peak separation. The krypton peak is not shown to scale as the BGA response is calibrated to xenon. This plot is taken from reference [28].	206

5.4	Lowering the krypton removal freezer into the freezer vessel. As the helium-xenon mixture passes through the freezer, xenon ice forms on the aluminum fins. Aluminum plates are thermally coupled to copper tubes flushed with liquid nitrogen. Heaters and temperature sensors on each aluminum plate control the temperature.	208
5.5	Screenshots of the Ignition user interface for run control (left) and state-based alarms (right). Users can queue multiple runs and set parameters like flow rates, slug sizes, and sniffer requests. Taken from [28].	211
5.6	An image of the SLAC sampling system taken via webcam. On the right are two 3.8 liter and one 0.5 liter stainless steel source bottles. Beneath the stocking shaped cold-trap is a styrofoam dewar that raised and lowered by a pneumatically actuated shelf. On the left is the sniffer bottle and a cylindrical dewar on a pneumatically actuated shelf. Both dewars are supplied liquid nitrogen by an insulated copper line. On the top right a manifold and solenoid valves control the system valves. Directly behind this panel are the cryodump, RGA, and pumps.	212
5.7	Block diagram of a single subsystem's automation software. A user inputs a list of queued operations and starts the software. The run queue is read by run control which then launches operation control. Operation control manages the subsystem hardware. If the operation succeeds run control is informed, the run control is updated, and if another item is in queue the process repeats. If the operation fails a safe state is enforced and an alarm is triggered.	215

5.8	A screenshot of the sampler’s run control GUI. The subsystem run queues are highlighted in blue, a system queue in black, a liquid nitrogen auto-fill interface in purple, the device request queue in green, and a sniffer-krypton removal interface in red. A text box is also available for quick notes.	216
5.9	Block diagram of an operation script. Hardware interactions can only occur in the run sequence or during an abort.	218
5.10	A screenshot of the interactive piping and instrumentation diagram from Ignition. Although all system tasks were automated, custom operations and trouble system trouble shooting require some manual control given here. Overlays show the specific subsystem: A) mixing panel in green, B) sniffer in red, and C) storage line in blue. The numbered, yellow overlays highlight shared devices: 1) calibration line, 2) cryodump, 3) RGA, 4) transfer line, and 5) utility pump.	219
5.11	A screenshot of the sampler’s liquid nitrogen auto-fill GUI.	228
5.12	A screenshot of the sampler’s run viewer GUI. The table on the left are selectable runs that can be analyzed by clicking select run. The results and raw RGA data can be exported to a csv.	229
5.13	Freezer batch krypton concentration as a function of mass spanned. The x-error bars represent the total mass contained in a single bottle, and the point corresponds to the mean value of the mass. The x-axis is in units of xenon mass out of the freezer. For example the first bottle in figure 5.13a is filled with the first 9.5 kg from freezer batch #61. The x-error bar spans from 0 to 9.5 kg and the data point is centered at 4.75 kg. The second bottle is filled with the next 11.5kg, has x-error bars from 9.5 to 20 kg, and the data point is centered at 15.25 kg.	231

- 5.14 A histogram of all of the measurements taken during the course of the krypton removal campaign. Sniffer measurements, in cyan, are used for feedback for chromatography parameters. Figure 5.15 shows a more detailed view of some of the sniffer measurement results. Bottle measurements, in blue, check the final purity xenon. KT measurements, in grey, check the purity of krypton traps that maintain helium purity and capture krypton. Test measurements, in black, correspond to specialty measurements like calibrations, zero measurements, or DVR oil space checks. 233
- 5.15 A histogram of all of the sniff measurements taken during the course of the krypton removal campaign. Measurements are grouped into three colors. The green measurements are sniffs that had krypton concentrations less than 200 ppq and measure an ideal chromatography result. Results between 200 ppq and 1000 ppq are shown in orange. These are results are not ideal, but can still produce reasonably clean batches of xenon. These intermediate purity results allowed operators to make informed decisions on whether or not to continue running. Red sniffer measurements have krypton concentrations greater than 1000 ppq and usually signal a larger issue. This information was used to halt production runs as needed to address purity issues by cleaning traps, or the DVR. 234

5.16	The cumulative mass of stored xenon and the mass averaged purity versus time. The green solid green line shows the mass-averaged krypton concentration and the dotted green line shows the target concentration. The blue line shows the cumulative mass of purified xenon and the dotted blue line shows the approximate mass of xenon required to fill the TPC. From March 2020 to July 2020, access to SLAC facilities was limited to essential personnel only due to COVID. Between July 2020 and December 2020 the system was recommissioned prior to full production that began in January 2021.	235
6.1	Decay scheme of ^{85}Kr . Decay time and energies taken from [29].	237
6.2	Full simulation of a golden event displayed in the LZ event viewer.	241
6.3	Full simulation of an event with merged $S2$ signals displayed in the LZ event viewer.	243
6.4	Full simulation of two events, one with merged $S1$ pulses and another with a delay time of $0.15\ \mu\text{s}$ displayed in the LZ event viewer.	244
6.5	Full simulation of an event with multiple gamma scatters contained within the TPC.	245
6.6	A full simulation of an event with energy deposited in the xenon skin region from a gamma scatter.	248
6.7	Analysis workflow between cuts. The section highlighted in blue is informed by the fast simulation technique and the the green is informed by the full simulation technique.	252
6.8	The MCtruth z_{truth} position (cm) versus the reconstructed drift time t_{drift} in ns. The red line is the linear fit used to convert z_{truth} to an equivalent truth drift time t_{truth} using $t_{truth} = -(6558 \pm 5)z_{truth} + (956226 \pm 387)$	255

6.9	The x and y reconstructed positions of events from the simulated ^{85}Kr excited state decay that were generated within the fiducial truth volume and satisfy the minimum selection criteria (cut0).	257
6.10	The reconstructed square radius and drift time positions of events from the simulated ^{85}Kr excited state decay that were generated within the fiducial truth volume and satisfy the minimum selection criteria (cut0).	258
6.11	The x and y reconstructed positions of events from the simulated ^{85}Kr excited state decay that were generated within the fiducial truth volume, satisfy the minimum selection criteria (cut0), and the reconstructed fiducial volume cut (cut1).	259
6.12	The reconstructed square radius and drift time positions of events from the simulated ^{85}Kr excited state decay that were generated within the fiducial truth volume, satisfy the minimum selection criteria (cut0), and the reconstructed fiducial volume cut (cut1).	260
6.13	A histogram of simulated $S1$ pulse areas, identified as either an $S1_\beta$ or $S1_\gamma$. The blue shaded area accepts any $S1_\beta$ pulse with an area between 50 and 1000 phd. The green shaded area accepts any $S1_\gamma$ pulse with an area between 1300 and 3200 phd. These events have been identified using the selection criteria and are reconstructed within the fiducial volume, no energy cuts have been applied.	262
6.14	A histogram of simulated $S2_{total}$ pulse areas for ^{85}Kr excited state decays. The green shaded area accepts all events with an $S2_{total}$ area between 0.4×10^6 and 1.28×10^6 phd. These events have been identified using the selection criteria and are reconstructed within the fiducial volume, no energy cuts have been applied.	264

6.15	The $S1_{total}$ pulse area versus the $S2_{total}$ pulse area of simulated ^{85}Kr excited state decays. The green shaded area is not explicitly a cut, but represents the combined $S1_{\beta}$, $S1_{\gamma}$, and $S2_{total}$ pulse area cuts from 1350 to 4200 phd in $S1_{total}$ area and the same $S2_{total}$ area cut from 0.4×10^6 and 1.28×10^6 phd. The two diagonal lines represent the fourth and final total energy cut. These events have been identified using the selection criteria and are reconstructed within the fiducial volume, no energy cuts have been applied.	265
6.16	A 2-D histogram of the peaks from table 6.1 identified in simulations. The left plot shows the corrected $S1_c$ and $S2_c$ pulse areas where events are identified. All identified events in simulations must pass a fiducial cut, energy cut, and are checked against particle IDs. The right plot shows the uncorrected $S1$ and $S2$ pulse areas. From bottom left to top right the events are increasing in energy from the the lowest energy 39.5 keV ^{125}I peak up to the 1460 keV ^{40}K peak.	268
6.17	A 2-D histogram of the peaks from table 6.1 identified in the SR1 and DD data. The left plot shows the corrected $S1_c$ and $S2_{c,bot}$ pulse areas where events are identified. All identified events in the detector data must pass a fiducialization cut, and energy cut. The right plot shows the uncorrected $S1$ and $S2$ pulse areas.	268
6.18	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 39.5 keV ^{125}I events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot.	269
6.19	The normalized Gaussian fits of the 39.5 keV ^{125}I events, from figure 6.18. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue).	270

6.20	The spatial distributions of the 39.5 keV ^{125}I events in simulations and SR1 data. The SR1 events, in figure 6.20a, are clustered in the top corner of the detector, near the DD conduit. The simulated events, in figure 6.20b, are uniformly distributed throughout the detector.	271
6.21	The spatial distributions of the 67.3 keV ^{125}I events in simulations and SR1 data. The SR1 events, in figure 6.21a, and the simulated events, in figure 6.21b, are uniformly distributed throughout the detector, as expected.	272
6.22	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 67.3 keV ^{125}I events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot.	273
6.23	The Gaussian fits of the 67.3 keV ^{125}I events, from figure 6.22. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue).	274
6.24	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 164 keV ^{131m}Xe events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot.	274
6.25	The Gaussian fits of the 164 keV ^{131m}Xe events, from figure 6.24. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue).	275
6.26	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 276 keV ^{125}Xe events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot.	275

6.27	The Gaussian fits of the 276 keV ^{125}Xe events, from figure 6.26. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue).	276
6.28	Spatial distributions of 1173 keV ^{60}Co events in simulations and SR1 data with the standard fiducial cut applied.	277
6.29	$S1$ and $S2$ pulse areas of simulated (left) and observed SR1 (right) 1173 keV ^{60}Co events in LZ. The $S1$ and $S2$ values have been projected onto their respective axes and a Gaussian fit is performed to identify the central peak locations. The fit data is shown on the top right of the plot. Only the standard fiducial cut is applied. The data is dominated by poorly reconstructed events along the detector edge and grids.	278
6.30	Spatial distributions of 1173 keV ^{60}Co events in simulations and SR1 data. Here the stringent volume cut has been made.	279
6.31	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1173 keV ^{60}Co events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot. The stringent volume cut has been made.	280
6.32	The Gaussian fits of the 1173 keV ^{60}Co events, from figure 6.31. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue). The stringent volume cut has been made.	280
6.33	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1332 keV ^{60}Co events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot. Only the standard fiducial cut has been made.	281

6.34	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1332 keV ^{60}Co events. The stringent volume cut has been made.	282
6.35	The Gaussian fits of the 1332 keV ^{60}Co events, from figure 6.34. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue). The stringent volume cut has been made.	282
6.36	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1460 keV ^{40}K events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot. Only the standard fiducial cut has been made.	283
6.37	The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1460 keV ^{40}K events. The stringent volume cut has been made.	284
6.38	The Gaussian fits of the 1460 keV ^{40}K events, from figure 6.37. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue). The stringent volume cut has been made.	284
6.39	The ^{85}Kr $S1_\gamma$ (514 keV) and $S2_{total}$ (514 to 687 keV) pulse areas from simulations of ^{85}Kr events. The standard fiducial volume cut is taken.	285
6.40	The ^{85}Kr $S1_\gamma$ peak value is shown in red. The black dotted line is a linear fit of the identified $S1$ peaks from table 6.1. The errors are the corresponding σ values found from the Gaussian fit of each peak.	286
6.41	The original (black) and adjusted (cyan and purple) $S1_\beta$ cuts. The “cut-” and “cut+” denote the narrower and wider cuts respectively. The number of events satisfying each cut are given below the specified ranges.	288
6.42	The original (black) and adjusted (cyan and purple) $S1_\gamma$ cuts.	289

6.43	The ^{85}Kr , $S2_{total}$ peak value is shown in red. The black dotted line is a linear fit using the identified $S2$ peaks of each source in blue from table 6.1. The errors are the corresponding σ values found from the Gaussian fit of each peak.	290
6.44	The original (black) and adjusted (cyan and purple) $S2_{total}$ cuts.	291
6.45	A histogram of events that survive up to cut5. All of the simulated events have a delay time less than $20 \mu\text{s}$ as shown of the left. The plot on the right is zoomed in and shows the first $1.5 \mu\text{s}$ to emphasize the $0.15 \mu\text{s}$ cut.	292
6.46	A histogram of events from SR1 data over the course of the 120 SR1 calendar days. The bin size is one day. The 68,916 events shown have been identified by the selection criteria, but fail cut7. The events excluded by the localization cut have a clear time structure that appears to correlate with random coincidences caused by cosmogenic and neutron activation following the initial fill (near day zero) and DD neutron calibrations (near day 30).	294
6.47	The pulse waveform and PMT hitmap of an event from SR1 day 7. Event 72773 from run 7208 occurred seven days into SR1 and satisfies cut0 through cut6 but fails localization. Note the significant separation between the two clusters of $S2$ light characterized by a localization value of 76.1 cm. The event has a delay time of $67.4 \mu\text{s}$	295
6.48	A histogram of event localization that survive up to cut6. The localization cut is shown in green.	296

6.49	A delay time histogram of 108995 events that survive the fast simulation cut7. The bin size is $0.1 \mu\text{s}$ and the scale factor is $c_{scale} = 10$ bins per μs . The red line is fit to all events between 0.15 and 100.15 μs , but the plot is zoomed in to the first 10 μs . Integration of the scaled signal component accurately accounts for all of the signal events and a zero background.	298
6.50	A histogram of $S1_{\gamma}$ pulse areas from the full simulation versus the $S2_{total}$ pulse area with Gaussian fits applied.	300
6.51	A histogram of $S1$ pulse areas from the full simulation. The blue shaded area accepts any $S1_{\beta}$ pulse with an area between 50 and 1000 phd. The green shaded area accepts any $S1_{\gamma}$ pulse with an area between 1600 and 3900 phd. The full simulation $S1_{\gamma}$ peak is centered at a pulse area 18.3% larger than the peak location identified and validated in the fast simulation. Only the selection criteria, or cut0, has been applied to emphasize events excluded by the cut overlay.	301
6.52	A histogram of $S2_{total}$ pulse areas from the full simulation. The green shaded area accepts all events with an $S2_{total}$ area between 0.4×10^6 and 1.28×10^6 phd. The full simulation $S2_{total}$ peak is centered at a pulse area 5.5% smaller than the peak location identified and validated in the fast simulation. Only the selection criteria, or cut0, has been applied to emphasize events excluded by the cut overlay.	302

6.53	The $S1_{total}$ pulse area versus the $S2_{total}$ pulse area from the full simulation. The green shaded area is not explicitly a cut, but represents the combined $S1_{\beta}$, $S1_{\gamma}$, and $S2_{total}$ pulse area cuts. The $S1_{total}$ “cut” extends from 1650 to 4900 phd, which is beyond the x-axis limit shown here. The same $S2_{total}$ area cut from 0.4×10^6 and 1.28×10^6 phd is shown along the y-axis. The two diagonal lines represent cut4. Only the selection criteria, or cut0, has been applied to emphasize events excluded by the cut overlay.	303
6.54	A histogram of full simulation events and their delay time. All of the events have a delay time less than $20 \mu s$, as shown on the left. The plot on the right is zoomed in and shows the first $1.5 \mu s$. Only the selection criteria, or cut0, has been applied to emphasize events excluded by the cut overlay.	304
6.55	A histogram of full simulation event localization values. The localization cut is shown in green. Only the selection criteria, or cut0, has been applied to emphasize events excluded by the cut overlay.	305
6.56	A histogram of ΔTBA values from full simulation events that survive cut7.	306
6.57	A 2-D histogram comparing full simulation ^{85}Kr reconstructed drift times and TBA values overlaid on SR1 ^{125}Xe reconstructed drift times and TBA values. The ^{85}Kr fit is shown in blue with vertical lines to identify binning used to calculate the difference in TBA values at a fixed drift time. The red line is the SR1 ^{125}Xe fit.	307
6.58	The original (black) and adjusted (cyan and purple) ΔTBA cuts.	309
6.59	A histogram of full simulation $S1$ pulse areas with the $S1_{\beta}$ cut shown in blue and the $S1_{\gamma}$ cut shown in green. These events have been identified using the selection criteria. The red circles highlight events that fail cut9 (skin veto cut).	311

6.60	A histogram of full simulation $S2_{total}$ pulse areas. These events have been identified using the selection criteria. The red circles highlight events that fail cut9 (skin veto cut).	312
6.61	A 2-d histogram of full simulation $S1_{total}$ versus $S2_{total}$ pulse areas. These events have been identified using the selection criteria. The red circles highlight events that fail cut9 (skin veto cut).	313
6.62	A 2-d histogram of full simulation event locations in squared radius (cm^2) versus drift time (μs) pulse areas. These events have been identified using the selection criteria. The red circles highlight events that fail cut9 (skin veto cut).	314
6.63	A delay time histogram of the 46735 full simulation events that survive cut9. The bin size is $0.1 \mu\text{s}$ and the scale factor is $c_{scale} = 10$ bins per μs . The red line is fit to all events between 0.15 and 100.15 μs , but the plot is zoomed in to the first 10 μs . Integration of the scaled signal component accurately accounts for all of the signal events and zero background events.	316
6.64	A histogram of the delay time for the remainin 15 SR1 events following cut9. The bin size is $0.1 \mu\text{s}$ and the scale factor is $c_{scale} = 10$ bins per μs . The red line is fit to all events between 0.15 and 100.15 μs , but the plot is zoomed in to the first 10 μs . The scaled and integrated fit parameters identify 8.1 ± 3.0 signal events and 6.9 ± 2.8 background events over the course of SR1 before taking efficiency into account.	322
6.65	A histogram of the radial positions and drift times of the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.	323

6.66	A stacked histogram of the drift times of the remaining 15 SR1 events following all nine cuts and side-band events. The final 15 SR1 events are shown in green for delay times $\leq 5 \mu s$ and blue for all other surviving delay times. The red corresponds to the 20 side-band events that failed only the skin veto cut. Events in both the side-band and final SR1 data appear to be cluster near the bottom of the detector.	324
6.67	A histogram of $S1$ pulse areas for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.	325
6.68	A histogram of $S2_{total}$ pulse areas for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.	326
6.69	A histogram of $S1_{total}$ versus $S2_{total}$ pulse areas for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.	327
6.70	A histogram of localization for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.	328
6.71	A histogram of TBA and ΔTBA cut for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.	329
6.72	A histogram of delay time for the remaining 15 SR1 events following all nine cuts. The red fit is to the final 15 SR1 events. The red circles correspond to the 20 events in the side-band events and failed only the skin veto cut. The side-band events appear to be consistent with a flat background. . .	330

6.73	A histogram of events over the course of the 120 SR1 calendar days. The bin size is one day. Neither the 15 surviving events, nor the 20 side-band events appear to have any structure throughout SR1.	331
6.74	A histogram 100,000 M_{sig} values taken from a Poisson distribution centered at 8.1.	332
6.75	A histogram 100,000 M_{bkg} values taken from a Poisson distribution centered at 6.9.	333
6.76	A histogram of toy delay time values from one iteration containing 5 signal and 6 background delay times. The number of events were generated using the M_{sig} and M_{bkg} Poisson distributions. The delay time values were generated using $e^{-t \ln(2)/1.015}$ and a flat distribution.	334
6.77	A histogram of all toy delay time values generated from 100,000 toy fits. The number of events were generated using the M_{sig} and M_{bkg} Poisson distributions. The delay time values were generated using $e^{-t \ln(2)/1.015}$ and a flat distribution.	335
6.78	A 2-d histogram of A and B fit parameters from 100,000 toy fits. The red box highlights the bin that contains the fit values found in the SR1 data search. The mean and standard deviations of the fit parameters are $A_{toy} = 6.13$, $\sigma_{A_{toy}} = 2.21$ and $B_{toy} = 0.069$, $\sigma_{B_{toy}} = 0.027$	337
6.79	A 2-d histogram of $A_{MC,err}$ and $B_{MC,err}$ fit parameter errors from 100,000 toy fits. The red box highlights the bin that contains the fit error values found in the SR1 data search. The mean and standard deviations of the fit parameter errors are $A_{MC,err} = 2.18$, $\sigma_{A_{MC,err}} = 0.39$ and $B_{MC,err} = 0.027$, $\sigma_{B_{MC,err}} = 0.006$	338
6.80	Final event 7051 from run 95884. The delay time is $0.19 \mu s$	341

6.81	Final event 6986 from run 87111. The delay time is 0.19 μs . The delay time is 0.38 μs	342
6.82	Final event 9834 from run 7664. The delay time is 0.50 μs	343
6.83	Final event 119504 from run 8073. The delay time is 0.77 μs	344
6.84	Final event 108440 from run 6773. The delay time is 60.3 μs	345
6.85	Final event 43340 from run 6785. The delay time is 76.8 μs	346
6.86	Final event 13919 from run 7790. The delay time is 87.8 μs	347

List of Tables

3.1	Table of activated xenon isotopes. Note that for the case of ^{131m}Xe the γ^* energy is often carried away by electron(s) and/or X-ray(s) and not a single gamma. ^{127}Xe has two characteristic gammas with relative intensities 53% and 47% as well as (K,L,M) electron cascade that produces X-rays with corresponding probabilities (83.37, 13.09, 2.88)%. Similarly ^{125}Xe has three potential characteristic gammas with relative intensity (6.8, 53.8, 30)% and X-rays from electron cascades. Taken from reference [30].	108
4.1	The mixing panel volumes.	150
4.2	The concentrations of constituent gasses in the xenon-air mixture given 98.5 grams of xenon and 54.7 mg of air. The errors, propagated from the appropriate volumes and pressure values, have a combined error of 4.3% for each of the calculated concentrations reported above. Air composition by mass percentages were taken from reference [31].	181

6.1	Sources that produce electron recoil events at a wide range of energies. Auger electrons, cascade x-rays, and intermediate short-lived excited states all contribute to these mono-energetic peaks. The 39.5 keV ^{125}I source is noted to include contributions from ^{129}Xe in the SR1 data due to neutron activation during DD calibrations. The most prominent energy peaks are selected for analysis. Total energies and half-life taken from reference [29]	267
6.2	Efficiencies from the fast simulation cuts. The fast simulation included ground state beta decays which were filtered using a particle ID truth RQ prior to identifying the initial number of events contained within the fiducial truth volume.	318
6.3	Efficiencies from the full simulation cuts. The full simulation only simulated excited state decays and was checked against a particle ID RQ. The cumulative efficiency up to cut7 is 68.4%. This reasonably close to the cumulative efficiency found from the fast simulation cuts despite a the 5.5% disagreement in the $S2_{total}$ pulse area peak and an 18.3% disagreement in the $S1_{\gamma}$ pulse area peak values.	319
6.4	Data cuts	321
6.5	Table of surviving SR1 events following all cuts.	340

Chapter 1: Identifying Dark Matter

In the last century, observational astronomy has accumulated substantial evidence that a significant portion of the universe is comprised of dark matter. This matter contributes to gravitational interactions at large scales, but fails to interact electromagnetically. Indirect measurements have quantified dark matter, yet its fundamental nature remains a mystery. A well-motivated candidate is the “weakly interacting massive particle”, or WIMP. In this chapter I will discuss the evidence of a dark mass component that permeates the Universe. Surveys of galactic motion, gravitational lensing, and the standard model of cosmology require a significant portion of the Universe to be composed of dark matter. Production mechanisms in the early Universe can explain the abundance observed today through WIMPs with a mass scale on the order of 100 GeV. Supersymmetric models of particle physics independently predict a non-standard model particle in a similar mass range.

1.1 Motion Within Clusters

The first time dark matter was suggested to describe an excess of observed mass on large scales was in 1933. Fritz Zwicky observed a discrepancy in the the motion of galaxies within the Coma cluster. The collective motion of galaxies could not be explained by their underlying luminous mass. If the cluster is in a stationary state, the virial theorem can be

used to relate the average kinetic energy $\langle T \rangle$ to the gravitational potential energy $\langle V \rangle$.

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

If the matter is uniformly distributed in a sphere of constant density ρ , the differential potential energy of each layer is

$$dV = -\frac{GM_r}{r} dm = -4\pi GM_r r \rho dr \quad (1.2)$$

where the mass within a radius r is $M_r = \frac{4}{3}\pi r^3 \rho$. The total gravitational potential for a sphere of radius R can be written as

$$V = \int_0^R -\frac{16}{3}\pi^2 G \rho^2 r^4 dr = -\frac{3GM^2}{5R} \quad (1.3)$$

where M is the total mass enclosed. Zwicky calculated an expected average velocity dispersion of 80 km/s based upon the cluster luminosity, much lower than dispersion of velocities he observed. From this Zwicky concluded that the Coma cluster contains an abundance of dark matter 400 times that of the luminous matter. Without the dark matter, the cluster would have broken apart, and eight high-velocity galaxies would have escaped the cluster long ago [32] [33].

An analysis of the relative motion of galaxies within the Virgo cluster produced a similar result. The study of 500 galaxies found an average density that was 200 times larger than expected based upon the luminous mass observed. Often dust or non-luminous dwarf galaxies were suggested as a solution to the missing mass [34]. In the Local Group, Andromeda and the Milky Way make up over 80% of the luminous mass. Approximating their dynamics as a two body system and using Kepler's law shows that the total mass of the galaxies is actually a factor of 6 larger than luminosity measurements would suggest

[35].

1.2 Galactic Rotation Curves

With clear evidence of dark matter dispersed throughout galaxy clusters, astronomers were motivated to study the internal kinematics of galaxies. Babcock, in 1939, studied the rising tail of Andromeda's rotation curve and found a mass-to-light ratio of 50 [36]. The rotation of galaxy NGC 3115 reveals a constant mass density outside of the galaxy's center. The mass to light ratio is 1.8 near the center and as the luminous matter abundance drops with radius the ratio gets as high as 250. Oort noted the total mass distribution has no relation to the luminous mass and suggested faint dwarfs and interstellar material as potential explanations [37].

Another approach uses N-body simulations to study the stability of flat spiral galaxies. Models of the Milky Way lack stability without a large spherical mass component at least as massive as the disk itself. Ostriker and Peebles found stability was only achieved using a halo to disk mass ratio between 1 and 2.5. Beyond the disk itself, the halo mass could be extremely large [38].

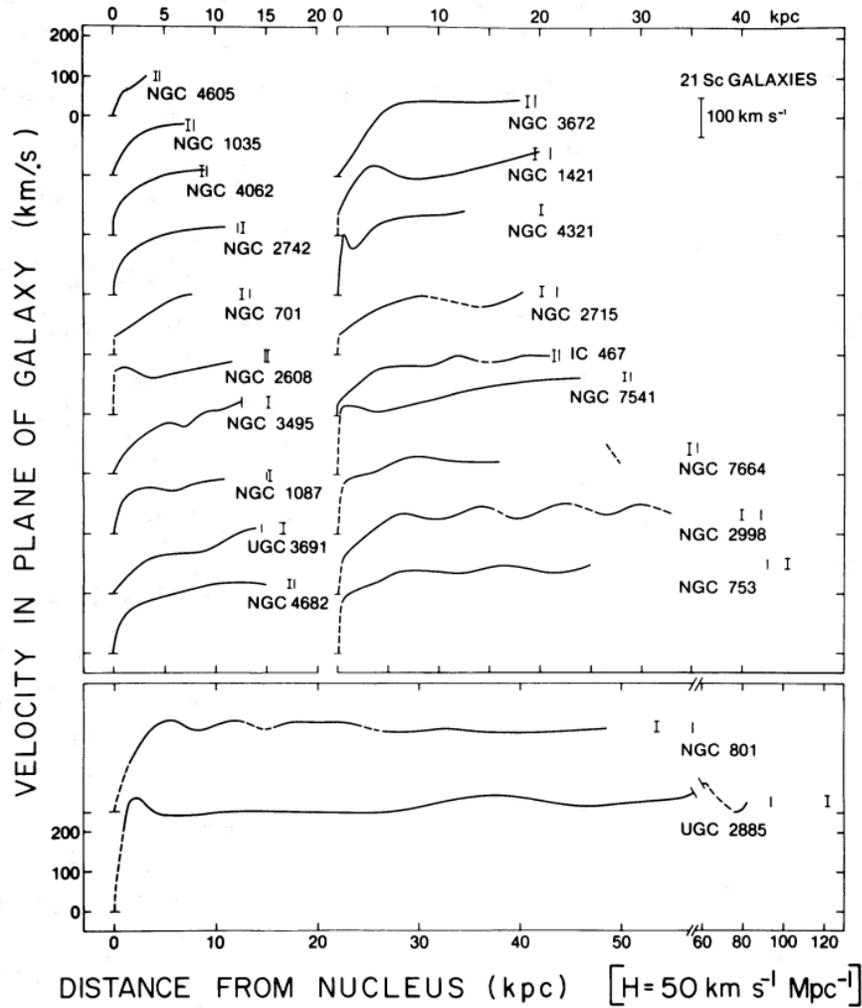


Figure 1.1: Mean rotational velocity curves of 21 spiral arm galaxies measured by Rubin and Ford. The dotted lines show velocities beyond the optical radius, calculated from the hydrogen 21 cm line, where 68% of the luminous matter is contained within. From reference [1].

In 1980 Rubin and Ford recorded rotation curves for 21 spiral arm galaxies shown in figure 1.1. Redshifted starlight and hydrogen's 21 cm line allowed velocity curves to be measured well outside the optical radius. Kepler's laws predict rotation curves that decrease far from the galactic center with $v_{rot} \propto \frac{1}{\sqrt{r}}$. Instead, they observed rotation curves that were flat or increasing at large radii. They concluded that each of these galaxies have significant mass beyond their optical radius, despite their different sizes and luminosities [1].

1.2.1 A Universal Rotation Curve

Researchers Paolo Salucci and Massimo Persic, in 1997, sought to unify different galactic rotation curves regardless of galaxy type. The universal rotation curve (URC) is composed of an exponentially decaying luminous matter component and a spherical dark halo component with a core of size a . A fit of all galaxy types is shown as a function of luminous mass and distance from the galactic center in figure 1.2.

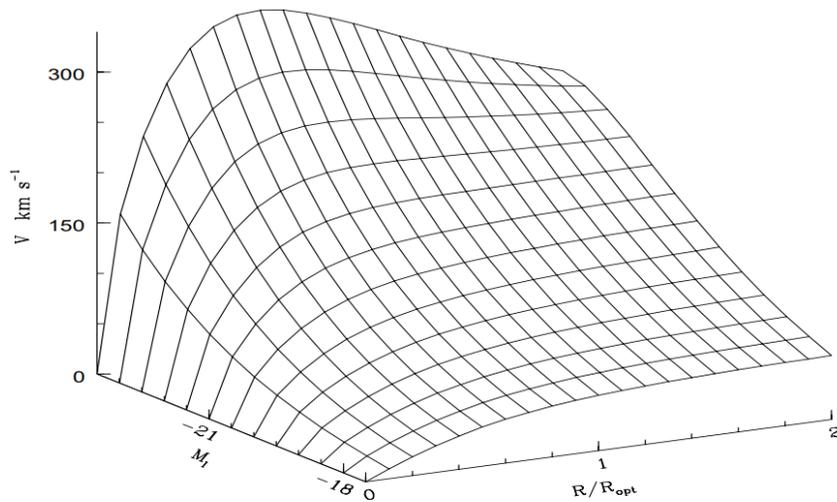


Figure 1.2: The Universal Rotation Curve surface for galaxies of arbitrary type. The radius from center of galaxy is normalized to its optical radius along x , total luminous mass along y , and resulting velocity along z . From reference [2].

The URC fits on average agree to within 1% of galaxy specific rotation curves. The general form of the dark halo mass distribution present in every galaxy regardless of type or size is modeled by

$$M_h(x) \propto (1 - \beta)(1 + a^2) \frac{x^3}{x^2 + a^2} \quad (1.4)$$

where x and a are in units of optical radius R/R_{opt} and β is the mass fraction contained within R_{opt} [2].

1.3 Gravitational Lensing

The analysis of galactic kinematics to quantify the total mass is limited to the luminous portion of the galaxy. The dark halo mass is dominant at the outermost reaches and can extend well beyond the luminous matter distribution. To measure the total mass of these objects reliably, a new tool must be used. Geodesics describe the shortest path between two points in spacetime. In general relativity, the shape of spacetime is warped and wiggled by mass and energy similar to how charges shape the electromagnetic field. The “force of gravity” is actually objects moving along geodesic paths in warped spacetime. A proper description of this effect shows that the total galactic mass can be measured through the deviation of photon paths.

1.3.1 Geodesics in Spacetime

The metric, $g_{\mu\nu}$, is a 4x4 symmetric matrix used to describe the geometry of spacetime. For now the metric will remain undefined until we choose to define one in equation (1.17). Points in spacetime are represented by four-vectors. In Cartesian coordinates, a four-vector takes the form of $x^\mu = (t, x, y, z)$ where $c = 1$. A freely falling particle between two points $x_i = x^\mu$ and $x_f = x^\mu + dx^\mu$ will follow the path that extremises the proper time τ . The proper time is the the time experienced by the particle along its world line. The differential proper time can be written in terms of the differential line element ds^2

$$d\tau^2 = -ds^2 = -g_{\mu\nu}dx^\mu dx^\nu. \quad (1.5)$$

The initial and final locations are parameterized in terms of σ , such that $x^\mu(\sigma = 0) = x_i^\mu$ and $x^\mu(\sigma = 1) = x_f^\mu$. The result can be written in terms of $L = \frac{d\tau}{d\sigma}$

$$\tau = \int_{x_i}^{x_f} \sqrt{-ds^2} = \int_0^1 d\sigma \sqrt{-g_{\mu\nu} \frac{dx^\mu}{d\sigma} \frac{dx^\nu}{d\sigma}} \equiv \int_0^1 d\sigma L(x^\mu, \dot{x}^\mu) \quad (1.6)$$

where $\dot{x}^\mu \equiv \frac{dx^\mu}{d\sigma}$. By varying the path slightly $x^\mu(\sigma) \rightarrow x^\mu(\sigma) + \delta x^\mu(\sigma)$ and fixing the endpoints, the variation of the proper time is

$$\delta\tau = \int_0^1 d\sigma \delta L = \int_0^1 d\sigma \left(\frac{\partial L}{\partial x^\mu} \delta x^\mu + \frac{\partial L}{\partial \dot{x}^\mu} \delta \dot{x}^\mu \right). \quad (1.7)$$

Integrating by parts gives

$$\delta\tau = \int_0^1 d\sigma \left(\frac{\partial L}{\partial x^\mu} - \frac{d}{d\sigma} \left(\frac{\partial L}{\partial \dot{x}^\mu} \right) \right) \delta x^\mu + \left[\frac{\partial L}{\partial \dot{x}^\mu} \delta x^\mu \right]_0^1. \quad (1.8)$$

The rightmost term is 0 as the endpoints are fixed $\delta x^\mu(0) = \delta x^\mu(1) = 0$. The path x^μ for all variations δx^μ is only minimized if the variation of the proper time $\delta\tau$ is equal to 0.

This results in the Euler-Lagrange equations

$$\frac{\partial L}{\partial x^\rho} - \frac{d}{d\sigma} \left(\frac{\partial L}{\partial \dot{x}^\rho} \right) = 0 \quad (1.9)$$

where the indices have been relabeled. Note that for an arbitrary function $f(\tau(\sigma))$ and $L = \frac{d\tau}{d\sigma}$ from equation (1.6) gives

$$\frac{df}{d\sigma} = \frac{df}{d\tau} \frac{d\tau}{d\sigma} = L \frac{df}{d\tau}. \quad (1.10)$$

Using $L = \sqrt{-g_{\mu\nu}\dot{x}^\mu\dot{x}^\nu}$ from equation (1.6), the leftmost term of equation (1.9) becomes

$$\frac{\partial L}{\partial x^\rho} = -\frac{1}{2L} \frac{\partial g_{\mu\nu}}{\partial x^\rho} \dot{x}^\mu \dot{x}^\nu = -\frac{1}{2L} \frac{\partial g_{\mu\nu}}{\partial x^\rho} \frac{dx^\mu}{d\tau} \frac{d\tau}{d\sigma} \frac{dx^\nu}{d\tau} \frac{d\tau}{d\sigma} = -\frac{L}{2} \frac{\partial g_{\mu\nu}}{\partial x^\rho} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} \quad (1.11)$$

The rightmost term of (1.9) is similarly expanded

$$-\frac{d}{d\sigma} \left(\frac{\partial L}{\partial \dot{x}^\rho} \right) = \frac{d}{d\sigma} \left(\frac{1}{L} g_{\rho\nu} \frac{dx^\nu}{d\tau} \frac{d\tau}{d\sigma} \right) = \frac{d}{d\sigma} \left(g_{\rho\nu} \frac{dx^\nu}{d\tau} \right) = L \frac{d}{d\tau} \left(g_{\rho\nu} \frac{dx^\nu}{d\tau} \right). \quad (1.12)$$

Equation (1.10) is used to remove all remaining σ parameters. Using the symmetry of the metric and changing dummy indices equation (1.12) is expressed as

$$L \left[\frac{dg_{\rho\nu}}{d\tau} \frac{dx^\nu}{d\tau} + g_{\rho\nu} \frac{d^2x^\nu}{d\tau^2} \right] = L \left[\frac{1}{2} \left(\frac{\partial g_{\rho\nu}}{\partial x^\mu} + \frac{\partial g_{\mu\rho}}{\partial x^\nu} \right) \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} + g_{\rho\nu} \frac{d^2x^\nu}{d\tau^2} \right]. \quad (1.13)$$

Combining the results from (1.11) and (1.13) in (1.9) gives

$$0 = -\frac{1}{2} \frac{\partial g_{\mu\nu}}{\partial x^\rho} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} + \left[\frac{1}{2} \left(\frac{\partial g_{\rho\nu}}{\partial x^\mu} + \frac{\partial g_{\mu\rho}}{\partial x^\nu} \right) \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} + g_{\rho\nu} \frac{d^2x^\nu}{d\tau^2} \right]. \quad (1.14)$$

Rearranging the terms results in the Geodesic Equation

$$\frac{d^2x^\nu}{d\tau^2} = -g_{\alpha\rho} \Gamma^\alpha_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau}. \quad (1.15)$$

The Christoffel symbol $\Gamma^\alpha_{\mu\nu}$ transports vectors along a manifold described by the metric and satisfies

$$g_{\alpha\rho} \Gamma^\alpha_{\mu\nu} = \frac{1}{2} \left(\frac{\partial g_{\rho\nu}}{\partial x^\mu} + \frac{\partial g_{\mu\rho}}{\partial x^\nu} - \frac{\partial g_{\mu\nu}}{\partial x^\rho} \right). \quad (1.16)$$

Due to the symmetry of the metric, the Christoffel symbol is also symmetric in its lower indices. For a diagonal metric like the Schwarzschild metric given in equation (1.17), the

The left side of equation (1.16) reduces to one quantity per equation [3, 39].

1.3.2 The Lens Equation

The Geodesic Equation completely describes the path an object will take between two points in an arbitrary spacetime geometry. Photons and other massless particles travel along “null geodesics”. These lightlike separated events are defined as having no change in proper time between them. As a photon passes massive objects, the null geodesic can be measured as a deviation from a straight line and describe the mass in a local spacetime geometry. Consider a point mass M and the nearby spacetime described by the Schwarzschild metric in spherical coordinates

$$g_{\mu\nu} = \begin{pmatrix} -(1 - \frac{2GM}{r}) & 0 & 0 & 0 \\ 0 & (1 - \frac{2GM}{r})^{-1} & 0 & 0 \\ 0 & 0 & r^2 & 0 \\ 0 & 0 & 0 & r^2 \sin^2(\theta) \end{pmatrix}. \quad (1.17)$$

For a lightlike separated event ($d\tau = 0$), in the plane $\theta = \frac{\pi}{2}$, equation (1.5) becomes

$$g_{\mu\nu} \frac{dx^\mu}{d\tau} \frac{dx^\nu}{d\tau} = - \left(1 - \frac{2GM}{r}\right) \dot{t}^2 + \left(1 - \frac{2GM}{r}\right)^{-1} \dot{r}^2 + r^2 \dot{\phi}^2 = 0 \quad (1.18)$$

where $\dot{}$ denotes $\frac{d}{d\tau}$. Since the Schwarzschild metric is symmetric and does not depend on time, the t component $\frac{\partial g_{tt}}{\partial t} = 0$ from equation (1.11). Therefore the term from equation (1.13) $g_{tt} \frac{\partial t}{\partial \tau}$ is a constant of motion

$$E = \left(1 - \frac{2GM}{r}\right) \dot{t}. \quad (1.19)$$

Similarly the metric does not depend on the coordinate ϕ resulting in another constant of motion

$$l = r^2 \dot{\phi}. \quad (1.20)$$

Rewriting the equation of motion 1.18 in terms of the energy E and angular momentum l gives

$$\dot{r}^2 + \frac{l^2}{r^2} \left(1 - \frac{2GM}{r} \right) = E^2. \quad (1.21)$$

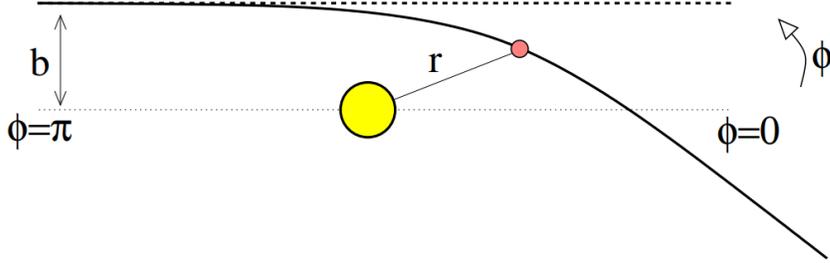


Figure 1.3: A photon approaching from $\phi = \pi$ and $r = -\infty$, with impact parameter b deviates from the straight path of $\phi = 0$ to some $\phi < 0$ at $r = \infty$. The deviation in the straight path is due to warped spacetime described by the Schwarzschild metric. From reference [3].

A photon's deflected path due to an object of mass M , as shown in figure 1.3, is described by the equation of motion (1.21). To determine the outgoing angle as $r \rightarrow \infty$, it is convenient to use the inverse distance $u(\phi) = \frac{1}{r(\phi)}$. The term \dot{r} can be written in terms of u and angular momentum l from 1.20 giving

$$\dot{r} = \frac{dr}{d\tau} = \frac{dr}{du} \frac{du}{d\phi} \frac{d\phi}{d\tau} = -l \frac{du}{d\phi}. \quad (1.22)$$

The equation of motion is now

$$\left(\frac{du}{d\phi} \right)^2 + u^2 (1 - 2GMu) = \frac{E^2}{l^2}. \quad (1.23)$$

Differentiating with respect to ϕ gives

$$\frac{d^2u}{d\phi^2} + u = 3GMu^2. \quad (1.24)$$

Ignoring the right hand side of equation 1.24, a solution for the homogeneous differential equation is found

$$\frac{d^2u_0}{d\phi^2} + u_0 = 0 \rightarrow u_0 = \frac{1}{b} \sin \phi. \quad (1.25)$$

where b is the impact parameter as shown in figure 1.3. If the coefficient $\beta = \frac{GM}{b}$ is small, a perturbative solution for u takes the form of

$$u = u_0 + \beta u_1 + \dots \quad (1.26)$$

Applying u_0 gives the non-homogeneous differential equation and corresponding solution

$$\frac{d^2u_1}{d\phi^2} + u_1 = \frac{3 \sin^2 \phi}{b} \rightarrow u_1 = A_1 \cos \phi + B_1 \sin \phi + \frac{1}{2b} (3 + \cos 2\phi). \quad (1.27)$$

To match the initial trajectory of $\phi = \pi$ coefficients are selected such that $u_1 \rightarrow 0$ as $\phi \rightarrow \pi$.

$$u_1(\phi = \pi) = 0 = -A_1 + \frac{2}{b} \rightarrow A_1 = \frac{2}{b}, B_1 = 0. \quad (1.28)$$

The solution to leading order in β is now

$$u = \frac{1}{b} \sin \phi + \frac{GM}{2b^2} (3 + 4 \cos \phi + \cos 2\phi). \quad (1.29)$$

The angle of deflection for a small ϕ as $r \rightarrow \infty$ and $u = 0$ is

$$\phi \approx -\frac{4GM}{b}. \quad (1.30)$$

This result is twice as large as the approximation in the Newtonian limit [3].

To measure the mass of an object, consider the image I produced by the deflection of light from a source S by a point mass M as described in figure 1.4.

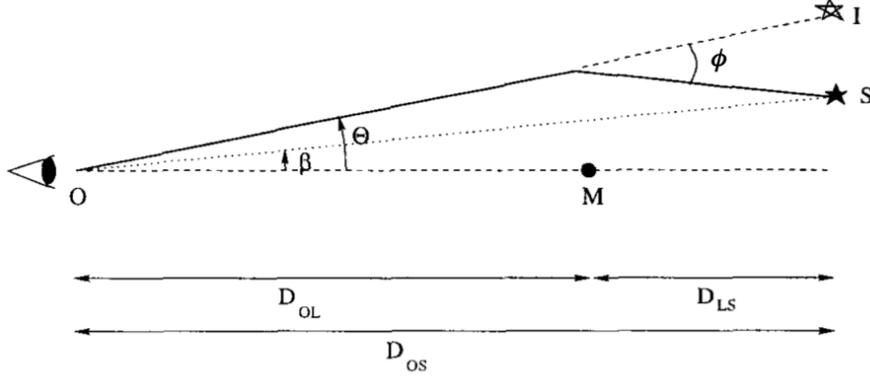


Figure 1.4: Geometry for a gravitational lens system. A point mass M deflects light from source S at an angle ϕ to produce an image I as seen by observer O . The observer is a distance D_{OS} from the plane of the light source and image. The angle between the observer and the light source is β . The angle between the observer and the image is θ . The point mass lens is a distance D_{OL} from the observer and a distance D_{LS} the plane of the light source and image. From reference [4].

For small angles β , θ , and ϕ , the perpendicular distance from the line of sight \overline{OM} to the location of the image I is

$$\theta D_{OS} = \beta D_{OS} + \phi D_{LS}. \quad (1.31)$$

Using the reduced deflection angle $\alpha \equiv \phi \frac{D_{LS}}{D_{OS}}$ between the observed source and image, equation (1.31) becomes the lens equation

$$\beta = \theta - \alpha. \quad (1.32)$$

The closest distance between the light ray and M is the impact parameter $b = \theta D_{OL}$. Plug-

ging this in to equation (1.30) and correcting the sign gives reduced deflection angle

$$\alpha = \frac{4GM}{\theta D_{OL}} \frac{D_{LS}}{D_{OS}}. \quad (1.33)$$

Using this result in the lens equation (1.32) and multiplying by a factor of θ gives a quadratic equation

$$\theta^2 - \beta\theta - \theta_E^2 = 0 \quad (1.34)$$

with $\theta_E = \sqrt{\frac{4GM D_{LS}}{D_{OL} D_{OS}}}$ and the roots

$$\theta_{\pm} = \frac{\beta \pm \sqrt{\beta^2 + 4\theta_E^2}}{2}. \quad (1.35)$$

The two solutions correspond to an image at each angular position on opposite sides of the lens. For the special case where the source, observer, and mass are inline an image of a ring is formed. The Einstein ring has an angular radius equal to the Einstein angle θ_E as show in figure 1.5.

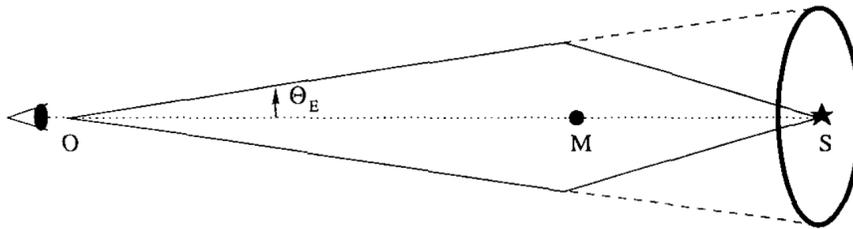


Figure 1.5: A light source in line with a lensing mass produces an Einstein ring observed at angular radius θ_E . From reference [4].

In 1936 Einstein wrote a quick response to Mandl's suggestion that gravitational lensing could be used to measure the mass of stars. Einstein responds by dismissing the cases where a ring or two images are formed by light bending around a start. He notes that resulting

image will not be large enough or bright enough to resolve from the foreground star. In the case where the source and foreground star are slightly off axis, a detectable brightness amplification of the image proportional to $\frac{\sqrt{D_{OL}}}{b}$ requires an impact parameter b so small that the image would be hidden by a star at a distance D_{OL} [40].

Zwicky followed up a year later to suggest galaxies as lensing mass candidates that could produce observable effects. The significantly longer distances for D_{OL} , larger masses, and comparably small impact parameters could result in detectable rings of sufficient size and brightness. Although telescopes at the time lacked the resolving power to test Zwicky's estimates, it motivated a new and comprehensive way to probe the total mass of galaxies [41].

Gravitational lensing that results in a ring or multiple images, are cases of “strong lensing”. The Twin Quasars (Q0957 + 561 A, B), discovered over 40 years later, was the first observed case of strong lensing. The two images of a single quasar have identical redshifts of $z = 1.41$, matching spectra, and are separated by 0.2 arcseconds [42]. The intermediate lensing galaxy (Q0957 + 561 G1) contains a mass distribution component consistent with a dark matter halo [43].

A lensing search of the galaxy cluster Abell 1689 identified a central dark matter clump merging with a smaller dark matter clump [5]. Of the 34 multiply-imaged sources, a dark matter halo model was essential in reproducing the multiply-lensed images for at least three of the source galaxies. A “weak lensing” search for distorted images was also carried out. The complimentary search methods find evidence of large dark matter distributions that permeate the cluster as well as localized distributions within the constituent galaxies.

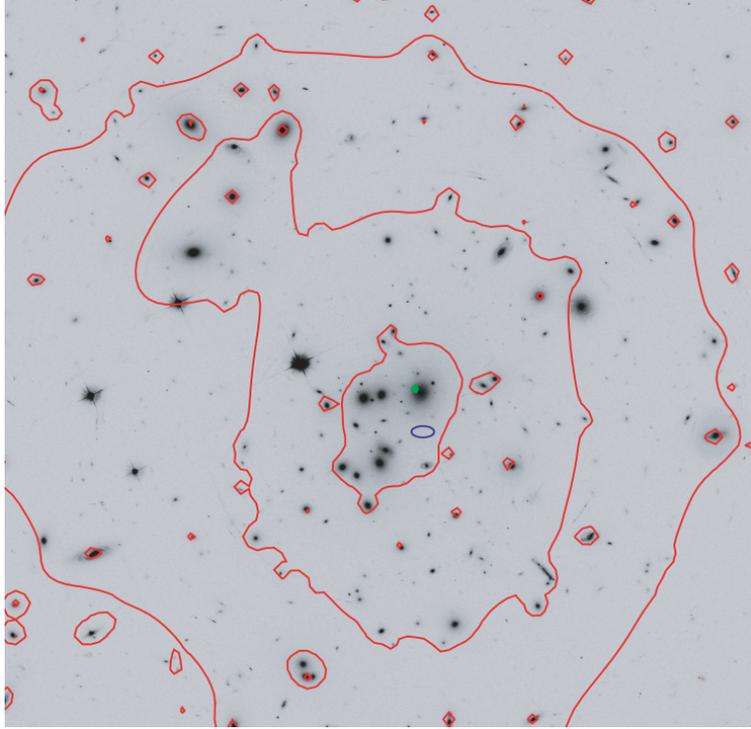


Figure 1.6: Mass distribution of the Abell 1689 galaxy cluster. The red contours show the projected mass density from the outer to inner regions as $(1.6, 2.4, 4.0) \times 10^{10} M_{\odot} \text{arcsec}^{-2}$. The green dot is the mass peak. The blue circle is the location of the main dark matter clump. The smaller dark matter clump is in the top left section of the image. From reference [5].

Instances of weak lensing are far more abundant because source objects are often not perfectly aligned with the foreground lens. This results in a slight deflection of light. With only a single image per source, an observer cannot directly measure the angle of deflection. A collection of sources all distorted by the same lens can be compared to determine the foreground mass distribution. The lens equation (1.32) in two dimensions becomes

$$\vec{\beta} = \vec{\theta} - \vec{\alpha} \quad (1.36)$$

This represents a two dimensional map between a source's true position ($\vec{\beta}$) and its image location ($\vec{\theta}$). The deflection angle α from (1.32) is now a superposition of deflection angles from a surface mass density in the lens plane. The gravitational potential can be projected

onto the lens plane to give a surface potential

$$\psi(\vec{\theta}) \equiv \int dz \phi(\vec{\theta}, z). \quad (1.37)$$

The projected surface mass density of the plane is related by taking the Laplacian of the projected potential

$$\nabla_{\theta}^2 \psi(\vec{\theta}) = 4\pi G \int dz \rho(\vec{\theta}, z) \quad (1.38)$$

where z is defined along the line of sight. The resulting deflection angle is therefore given by the gradient of the surface potential and the underlying mass surface density. The vector lens equation is now

$$\vec{\beta} = \vec{\theta} - \vec{\nabla}_{\theta} \Psi(\vec{\theta}) \quad (1.39)$$

for the re-scaled gravitational potential

$$\Psi \equiv \frac{2D_{LS}}{D_{OS}D_{OL}} \psi. \quad (1.40)$$

The Jacobian matrix can be written as a map between the angular coordinates of the source and image

$$T_{ij} = \frac{\partial \beta_i}{\partial \theta_j} = \delta_{ij} - \frac{\partial^2 \Psi}{\partial \theta_i \partial \theta_j} = \begin{pmatrix} 1 - \kappa - \gamma_1 & -\gamma_2 \\ -\gamma_2 & 1 - \kappa + \gamma_1 \end{pmatrix}. \quad (1.41)$$

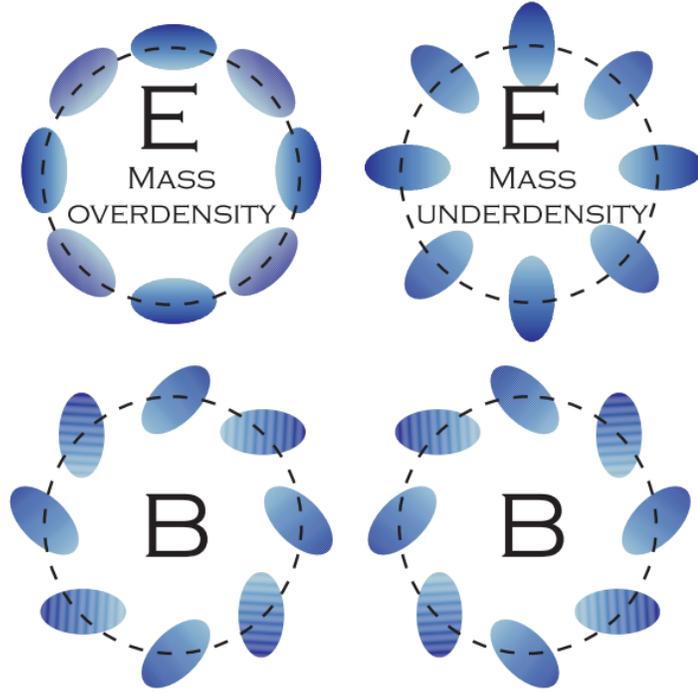


Figure 1.7: The image distortion of a circular object by a gravitational lens for mass over (left) and under (right) densities. E-modes (top) denote the effect of a radially symmetric mass distribution within the gravitational lens, while the curl-like B-mode denotes the effect of complex mass distributions that result in optical shearing. Top left: $\gamma_1 > 0$. Top right: $\gamma_1 < 0$. Bottom left: $\gamma_2 > 0$. Bottom right: $\gamma_2 < 0$. From reference [6].

The mapping matrix can be iteratively solved to find a mass distribution that properly transforms all of the sources to their corresponding image. Each of the terms in the mapping matrix modify the image while maintaining the total surface brightness. The diagonal κ term, known as convergence, is related to the surface mass density of the lens and modifies magnification of the image without changing its shape. The diagonal γ_1 term stretches the image along the radial or tangential direction with respect to the center of the lens plane. The off diagonal γ_2 shears the image diagonally. Both γ effects are shown in figure 1.7. For a circularly symmetric projected surface mass density, the γ_2 term vanishes and the image is only stretched in the tangential or radial direction [4] [44] [6].

A study of over 1300 galaxies behind the cluster Cl 0024 + 17 finds signs of strong

and weak lensing. The cluster appears to be the collision of two sub clusters along the line of sight. Lensing data from the Hubble Space Telescope and x-ray data from the Chandra X-ray Observatory are used to map the mass profile. Five multiple images of a single source are identified with an additional 2 multiply imaged source candidates. The mass distribution model relies heavily on dark matter halos and reveals a ringlike structure of dark matter, as see in figure 1.8, surrounding a dense core with a significance of 5σ . Despite the identification of the dark matter component, the center of mass offsets for the two subclusters, dark matter centroids, and x-ray clumps cannot be resolved due to the collision occurring along the line of sight [44].

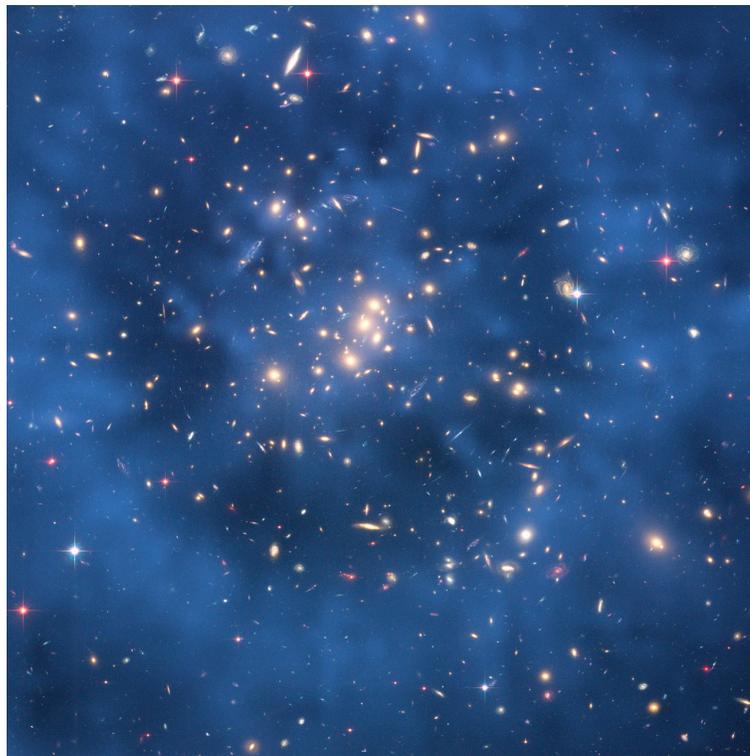


Figure 1.8: Hubble Space Telescope image of C1 0024 + 17 with a dark mass cloud illustrated to highlight the ring and core structure. The structure is distinct compared to the galaxy and hot gas distributions. From reference [7].

The Bullet Cluster is the result of two sub clusters that collided roughly 100 million years ago. The motion of the two clusters is largely in the plane of the sky with only

600 km/s relative velocity along the line of sight. A distinct bow shock, as seen on the right side of figure 1.9, indicates the sub cluster is moving away from the main cluster at 4700 km/s. Lensing data from background galaxies as seen by the Magellan and Hubble telescopes are used to determine the total mass distribution within the two sub clusters. X-rays measure the distribution of hot plasma produced from baryonic matter collisions. An offset total center of mass with respect to the baryonic center of mass is measured at an 8σ significance. This result can be explained by a collision where the baryonic components experienced significant drag and slowed to their current speeds, but the dark matter component continued without interacting during the collision [8].

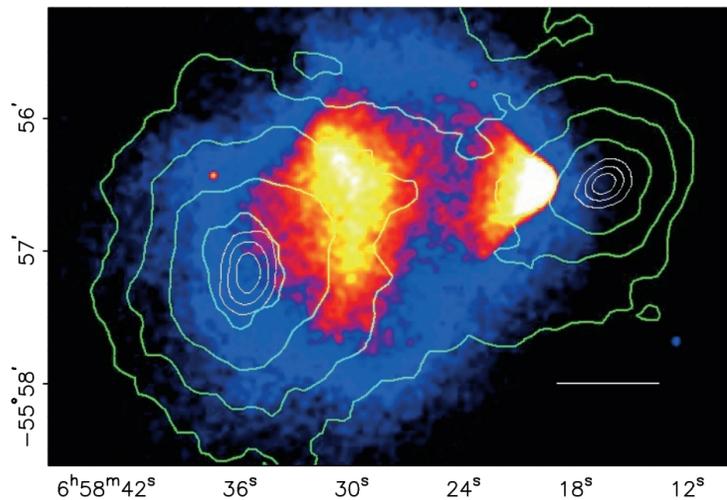


Figure 1.9: Composite image of the Bullet cluster from the Chandra x-ray image showing the hot plasma distribution with the total reconstructed lensing mass distribution. The white contours show the location of the total surface density mass peak corresponding to the 1σ , 2σ , and 3σ confidence levels. The green contours are lines of constant surface mass density expressed in steps of 0.07κ . From reference [8].

1.4 Cosmology

As the effects of dark matter are abundantly clear at large scales, a description of dark matter on the cosmological scale is motivated. The standard model of cosmology is a

description of the universe and its evolution in terms of general relativity.

1.4.1 General Relativity

Einstein's field equations (1.42) give a relationship between the geometric properties of the universe and the contents within it.

$$G_{\mu\nu} \equiv R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = 8\pi GT_{\mu\nu} \quad (1.42)$$

The Ricci tensor $R_{\mu\nu}$ is a description of the differential geometric properties given by the Christoffel symbols (1.16)

$$R_{\mu\nu} = \Gamma^{\alpha}_{\mu\nu,\alpha} - \Gamma^{\alpha}_{\mu\alpha,\nu} + \Gamma^{\alpha}_{\beta\alpha}\Gamma^{\beta}_{\mu\nu} - \Gamma^{\alpha}_{\beta\nu}\Gamma^{\beta}_{\mu\alpha} \quad (1.43)$$

where “,” denotes a derivative, for example $\Gamma^{\alpha}_{\mu\nu,\alpha} \equiv \partial\Gamma^{\alpha}_{\mu\nu}/\partial x^{\alpha}$. The Ricci scalar R is defined by

$$R \equiv g^{\mu\nu}R_{\mu\nu}. \quad (1.44)$$

The cosmological principle is the notion that at large scales, the universe is isotropic and homogeneous. All observers within the universe agree that the universe is uniform at all points and looks the same in all directions. Three geometries satisfy this principle and

are represented in the Friedmann-Robertson-Walker (FRW) metric

$$g_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & \frac{a(t)^2}{1-k^2 \frac{r^2}{R^2}} & 0 & 0 \\ 0 & 0 & a(t)^2 r^2 & 0 \\ 0 & 0 & 0 & a(t)^2 r^2 \sin^2(\theta) \end{pmatrix}. \quad (1.45)$$

A unitless scale factor $a(t)$ describes the contraction or expansion of the universe over time. The term R describes the radius of curvature and k indicates the sign of curvature corresponding to each of the geometric descriptions.

$$k = \begin{cases} +1 & \text{Spherical} \\ 0 & \text{Euclidean} \\ -1 & \text{Hyperbolic} \end{cases} \quad (1.46)$$

Solving for the components of the Ricci tensor (1.43) and scalar (1.44) gives the left hand side of (1.42). The G_{tt} component is

$$G_{tt} = 3\frac{\ddot{a}}{a} - \frac{1}{2} \left(6\frac{\ddot{a}}{a} + 6\left(\frac{\dot{a}}{a}\right)^2 - 6\frac{k^2}{R^2 a^2} \right) = 3\frac{k^2}{R^2 a^2} - 3\left(\frac{\dot{a}}{a}\right)^2. \quad (1.47)$$

The right hand side of (1.42) is the symmetric stress-energy tensor $T_{\mu\nu}$ which describes the contents of the universe. The T_{tt} describes the energy density of the universe $\varepsilon(t)$ [45] [46] [47]. The resulting relationship between the tt components of Einstein's equations is

$$3\frac{k^2}{R^2 a^2} - 3\left(\frac{\dot{a}}{a}\right)^2 = 8\pi G\varepsilon(t). \quad (1.48)$$

1.4.2 Lambda-CDM Model

This solution to Einstein's field equation is called the Friedmann equation. It describes a time evolving universe for three different geometries (1.46) in terms of the energy density. Rearranging the result from above gives

$$H(t)^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\varepsilon(t) - \frac{k}{R^2 a(t)^2} \quad (1.49)$$

where $H(t)$ is the Hubble parameter and describes the expansion rate of the scale factor. The value at the current time t_0 is Hubble's constant $H(t_0) = 70 \pm 7 \text{ km s}^{-1} \text{ Mpc}^{-1}$, where the scale factor is chosen such that $a(t_0) \equiv 1$. For a flat universe, with $k = 0$ at the current time t_0 , a critical density can be defined

$$\varepsilon_0 \equiv \frac{3H_0^2}{8\pi G}. \quad (1.50)$$

In an empty universe with zero energy density, the scale factor can either be constant in a flat universe ($k = 0$), or increasing in a universe with negative curvature ($k = -1$). A positive curvature ($k = 1$) gives a non-physical result due to the square in equation (1.49).

For a flat universe that is not empty, its contents can be modeled by a perfect isotropic fluid that fills the adiabatically expanding universe. The fluid is characterized by the energy density ε and a pressure P . For a perfectly homogeneous universe there is no bulk heat flow for a comoving volume, thus $dQ = 0$. The first law of thermodynamics can be written as

$$0 = \dot{E} + P\dot{V}. \quad (1.51)$$

The volume is proportional to the cube of the scale factor such that

$$\dot{V} = V_0 \frac{d}{dt}(a(t)^3) = V_0 3a^2 \dot{a} = 3 \frac{\dot{a}}{a} V \quad (1.52)$$

for a proper volume V_0 . The internal energy is the product of volume and energy density

$$\frac{d}{dt} E = \frac{d}{dt} V(t) \varepsilon(t) = \dot{V} \varepsilon + V \dot{\varepsilon}. \quad (1.53)$$

Combining (1.51), (1.52), and (1.53) gives the fluid equation

$$\dot{\varepsilon} + 3 \frac{\dot{a}}{a} (\varepsilon + P) = 0. \quad (1.54)$$

Differentiating the Friedmann equation (1.49) with respect to time and combining with the fluid equation gives the acceleration equation

$$\frac{\ddot{a}}{a} = \frac{4\pi G}{3} (\varepsilon + 3P). \quad (1.55)$$

The next step in describing the universe relates the pressure and energy density of the fluid using the equation of state

$$P = \sum_{\omega} \omega \varepsilon_{\omega}. \quad (1.56)$$

Each component of the fluid contributes pressure according to a unique, unitless scale ω .

The fluid equation (1.54) for a single component universe $P = \omega \varepsilon$ can be rewritten as

$$\frac{\dot{\varepsilon}}{\varepsilon} = -3 \frac{\dot{a}}{a} (1 + \omega). \quad (1.57)$$

Integrating both sides and using $a(t_0) = 1$ the time evolution of the energy density is found

$$\varepsilon(t) = \varepsilon_0 a(t)^{-3(1+\omega)} \quad (1.58)$$

where ε_0 is the critical density from (1.50).

For a flat, matter-only universe ($\omega = 0$), the result is intuitive. The number density will reduce by a factor of $a(t)^{-3}$ as the volume increases. The photon ($\omega = 1/3$) the number density is similarly reduced as the volume grows, but also includes a redshift term. This results in a reduction of total photon energy density by a factor of $a(t)^{-4}$. For the case of a single component, flat universe the Friedmann equation(1.49) can be rewritten entirely in terms of $a(t)$

$$\left(\frac{\dot{a}}{a}\right)^2 = H_0^2 a^{-3(1+\omega)}. \quad (1.59)$$

The solution for a single component universe from the Friedmann equation takes two forms

$$\omega \begin{cases} \neq 0, & a(t) = \left(\frac{t}{t_0}\right)^{2/(3+3\omega)} \\ = 0, & a(t) = e^{H_0(t-t_0)} \end{cases} \quad (1.60)$$

There are four cases of particular interest

$$\omega \begin{cases} = 1/3, & \text{radiation,} & a(t) = \left(\frac{t}{t_0}\right)^{1/2} \\ = 0, & \text{matter,} & a(t) = \left(\frac{t}{t_0}\right)^{2/3} \\ = -1/3, & \text{curvature,} & a(t) = \left(\frac{t}{t_0}\right) \\ = -1, & \Lambda, & a(t) = e^{H_0(t-t_0)} \end{cases} \quad (1.61)$$

Each component has a unique impact on the time evolution of the a flat universe. Radiation and matter both have negative second derivatives which slow down the expansion of the

universe. The curvature term is constant in its first derivative and expands the universe at a constant rate. Any observed deviations would suggest a slight negative curvature. The Λ term represents “dark energy” and has a positive second derivative. This drives the acceleration of the universe [45] [46] [47].

Current observations seek to measure these cosmological parameters and build a cohesive model called the Λ -Cold Dark Matter (Λ -CDM) model. This standard model of cosmology describes the current universe as flat and dominated by matter and dark energy. The “cold” term describes the observation that most of the matter is non-relativistic. This is well-motivated by the observation that dark matter plays an essential role in galaxy and cluster kinematics as well as the evolution of large scale structures.

1.4.3 Cosmic microwave background

The cosmic microwave background (CMB) serves as one of the best ways to constrain the energy density parameters that drive the evolution of the universe. The reported values are expressed as a fraction of the critical density

$$\Omega_\omega = \frac{\epsilon_\omega}{\epsilon_0}. \quad (1.62)$$

The early universe was a dense hot plasma in thermal equilibrium. As the universe expanded, the plasma cooled. Within the first hour of expansion, nucleons of light mass elements formed in a process called big bang nucleosynthesis (BBN). During nucleosynthesis deuterium was produced. Upon further cooling its number density was frozen out as it was no longer being produced or destroyed. Since there is no stellar production mechanism, the relative abundance of deuterium to hydrogen offers a measurement of the critical density of baryons. In 1995 a measurement of this relative abundance placed constraints on the critical energy density contributed by baryons, Ω_b , to be between between 1% and

15%. Any mass contributions to the energy density larger than 15% would require a significant non-baryonic dark matter component comprised of elementary particles produced during the earliest moments of the universe [48]. Current measurements find a baryon mass density $\Omega_b \approx 5\%$ [10].

About 380,000 years after the big bang, the average temperature of the plasma dropped below 13.6 eV and hydrogen atoms formed. As light atoms formed, the number density of free electrons dropped and photons were less likely to scatter within the plasma. The photon mean free path grew until the universe became transparent to photons. This moment of decoupling is captured in the extremely uniform, red shifted, black-body spectrum of photons in the cosmic microwave background (CMB).

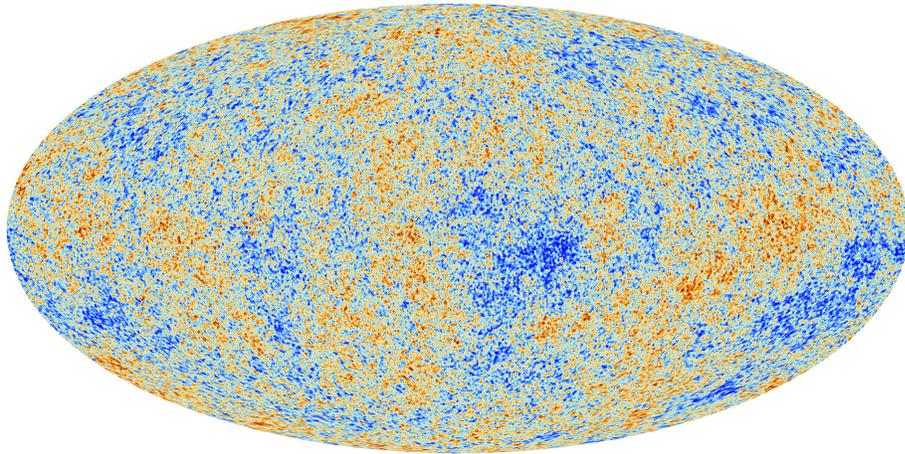


Figure 1.10: Temperature variations of the cosmic microwave background with respect to the black-body temperature 2.73 kelvin. The very isotropic radiation has variations on the order of μ kelvin. From reference [9].

The most recent measurements of the CMB constrain the cosmological parameters. The Planck satellite measures a curvature component $\Omega_K = 0.001 \pm 0.002$, consistent with flat universe, containing a matter component $\Omega_m = 0.315 \pm 0.007$, and a dark energy component $\Omega_\Lambda = 0.685 \pm 0.007$. The matter component is more than twice the upper limit allowed by a baryonic only matter universe described by BBN studies. These measurements show

that 84.4% of the universe’s matter is dark and the remaining 15.6% is baryonic [10].

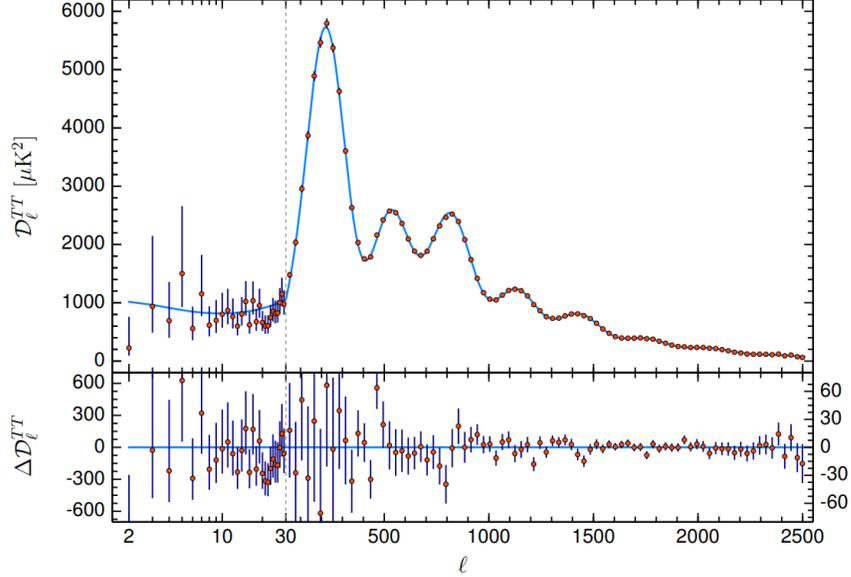


Figure 1.11: The angular power spectrum of anisotropies from the CMB. The deviation from the mean temperature squared shows spikes as a function of the multipole. These features correspond to mass infall from gravity (odd peaks) and rarefaction due to radiation pressure (even peaks). Of particular prominence is the third peak ($\ell \approx 850$) suggesting a driving gravitational interaction that did not contribute a restorative pressure. From reference [10].

A further study searches for angular structures within the CMB anisotropies. A decomposition of the 2D surface into spherical harmonics reveals the angular power spectrum in figure 1.11. This plot describes the dense hot plasma of the early universe. Prior to photon decoupling, the plasma had regions of slight over and under-densities. Over-dense regions gravitationally attracted matter. As matter fell in and heated, a counter pressure due to radiation pushed baryonic matter back, resulting in oscillations. As photons decoupled from the plasma the radiation pressure ceased. The temperature variations in the CMB are due to these so called “baryonic acoustic oscillations” (BAO). The odd peaks of figure 1.11 are caused by matter infall, and the even are from matter rarefaction. If only baryonic matter contributed to the mass over-densities, the resulting angular power spectrum would fall on an exponentially decaying envelope as the oscillations damped. Instead a prominent

third peak suggests a gravitational driving force. This can be explained by a dark matter component that interacts gravitationally but does not couple strongly to itself, photons, or baryonic matter. This dark matter component would collect in regions of over densities, driving odd peaks in the BAO, but would not contribute to rarefaction [49] [10].

1.5 Dark Matter Candidates

The term cold dark matter is a catch-all for the unaccounted mass that contributes to the observed mass density of the universe. There are many candidates that fit the description of dark matter but few are as well motivated as the WIMP candidate.

1.5.1 Baryonic dark matter

Many of the initial theories regarding dark matter suggested baryonic matter in the form of brown dwarfs, dust, or gas clouds as a potential candidate. Analysis of the CMB and residual deuterium from BBN constrain the total baryon mass contribution. As discussed in section 1.4.3 the total mass density of the universe is over six times larger than the mass density of baryons and requires an additional dark matter component [10] [48].

1.5.2 MACHOs

Massive compact halo objects (MACHOs) are another candidate. These objects can include primordial black holes, neutron stars, and brown dwarfs. While neutron stars and brown dwarfs are baryons, primordial black holes are not. Searches using gravitational lensing show insufficient MACHO mass to account for all of the cold dark matter. A search for MACHOs within the Milky Way halo was conducted over 5.7 years. Lensing data from objects that passed between the Earth and the Large Magellanic cloud showed that only 20% of the Milky Way's dark matter halo was comprised of MACHOs [50]. Also, a study

of Eridanus II and other ultra-faint dwarf galaxies combines results from microlensing and disk kinematics to rule out MACHOs as a primary source of dark matter [51].

1.5.3 Neutrinos

Another potential candidate is the neutrino. The existence of at least three flavours of neutrinos is well established, and their mass sum is constrained to be $\sum m_\nu < 0.12 \text{ eV}$ [10]. These ultra-relativistic particles cannot account for the significant portion of dark matter required by the observed structure in galaxies and galaxy clusters as they are not gravitationally bound. N-body simulations of a neutrino-dominated universe produce results inconsistent with the structure of the universe observed today [52]. Non relativistic neutrinos from a cosmic neutrino background are also constrained as the total neutrino density parameter observed is $\Omega_\nu = 3.4 \times 10^{-5}$ [45].

1.5.4 Axions

Axions are proposed as a solution to the strong charge-parity (CP) problem. Their properties also make them a promising candidate for cold dark matter. In field theory, the symmetry of charge, parity, and time reversal (CPT) is fundamental. If one symmetry is broken, one of the two remaining must be broken as well. CP symmetry violation is observed in weak interactions, but has not been observed in strong interactions. If a CP violating vacuum angle, θ , is included in the quantum chromodynamics (QCD) vacuum state, symmetry is broken for any non-zero θ . This can be measured directly through the neutron electric dipole moment which is currently constrained to $|d_n| < 6.3 \times 10^{-26} \text{ e cm}$. In terms of the vacuum angle this is $|\theta| < 10^{-9}$. Despite θ being unconstrained by QCD gauge symmetry, measurements find it to be nearly zero. If θ is treated as a dynamic variable, then an underlying mechanism could minimize it. The Peccei-Quinn solution to

the strong CP problem introduces a new symmetry with a corresponding scalar field and pseudoscalar particle. When this new symmetry is broken, θ relaxes to a minimum value. This relaxation excites the axion scalar field and could produce axion dark matter. The axion would be electrically neutral, stable, and rarely interacting. If the axion mass is found to be in the range of $6 \times 10^{-6} < m_a < 6 \times 10^{-3}$ eV, it could account for a significant portion of the dark matter.

In the early universe axions could have been produced in sufficient quantities to explain the current abundance of dark matter. While the universe was above the symmetry breaking temperature T_{PQ} , θ could take a misalignment angle θ_0 . As the universe cooled below T_{PQ} the Peccei-Quinn symmetry would break, relaxing the vacuum angle θ to 0. This vacuum realignment would have produced an axion energy density

$$\Omega_a \approx 0.15 \left(\frac{f_a}{10^{12} \text{GeV}} \right)^{7/6} \theta_0^2. \quad (1.63)$$

If the axion decay constant is in the range $10^9 \text{ GeV} < f_a < 10^{12} \text{ GeV}$, the energy scale of spontaneous symmetry breaking would correspond to the production of axions with the appropriate mass to account for the relic dark matter [53].

The existence of axions would modify Maxwell's equations to include the scalar axion field. The Axion Dark Matter eXperiment (ADMX) searches for axions through a two-photon coupling $g_{a\gamma\gamma}$ via the resonant conversion of axions to RF photons in a microwave cavity. The frequency of the induced microwave photons is proportional to the mass of the axion. Results from ADMX have excluded galactic halo axion masses in the 2.66-4.2 μeV mass range [54] [55] [56]. The ADMX search for axion dark matter is ongoing, with intentions to search for masses up to 100 μeV .

1.5.5 Weakly Interacting Massive Particles

The final compelling candidate for dark matter is the weakly interacting massive particle. This massive, stable particle can interact with itself and baryonic matter via the weak nuclear force. The WIMPs observed today as cold dark matter could be a thermal relic from the early universe. In a dense, hot early universe, all matter was in thermal equilibrium. WIMPs produced standard model particles at the same rate that standard model particles produced WIMPs. As the universe expanded, the temperature fell below the WIMP mass M_χ . This caused the dark matter production mechanism to cease, but WIMP annihilation remained. Consequently, the WIMP number density dropped exponentially as a function of the temperature $e^{-M_\chi/T}$. As expansion continued, the distribution of dark matter was diluted to the point that dark matter particles could no longer find an annihilation partner. This resulted in a “freeze-out” as the number of dark matter particles became constant. The dark matter mass observed today could be explained by a thermal relic density of WIMPs.

The evolution of the number density n is described by the Boltzmann equation

$$\frac{dn}{dt} = -3Hn - \langle\sigma_{AV}\rangle (n^2 - n_{eq}^2) \quad (1.64)$$

where H is the Hubble parameter, $\langle\sigma_{AV}\rangle$ is the average thermal annihilation cross, and n_{eq} is the WIMP number density in thermal equilibrium. The Hubble term dilutes the WIMP density through expansion, the n^2 term is due to WIMP self-annihilation, and the n_{eq}^2 term is from WIMP production. Freeze out occurs at a temperature T_f when the expansion term is comparable to the annihilation term $H \approx \langle\sigma_{AV}\rangle n$. Note that for a radiation dominated universe $H \propto T^2/M_{PL}$. At the onset of freeze-out the total WIMP density can be written by equating the now dominant dilution term to the Boltzmann suppressed thermal equilibrium

number density

$$n_f \approx (M_\chi T_f)^{3/2} e^{-M_\chi/T_f} \approx \frac{T_f^2}{M_{PL} \langle \sigma_{Av} \rangle} \quad (1.65)$$

where M_{PL} is the Planck Mass. The ratio $x_f \equiv m_f/T_f$ can be taken as a constant with a typical value of 20. The thermal relic density n_0 of the current universe at a temperature T_0 is

$$\Omega_\chi = \frac{M_\chi n_0}{\epsilon_0} = \frac{M_\chi T_0^3}{\epsilon_0} \frac{n_0}{T_0^3} \approx \frac{M_\chi T_0^3}{\epsilon_0} \frac{n_f}{T_f^3}, \quad (1.66)$$

where the approximation $\frac{n_0}{T_0^3} \approx \frac{n_f}{T_f^3}$ can be made after freeze out since the entropy of a comoving volume is fixed. This implies $T \propto a^{-1}$ and the ratio of the number density to a comoving volume is a constant. From equation (1.65) the dark matter density observed today is given in terms of average thermal cross section

$$\Omega_\chi \approx \frac{x_f T_0^3}{\epsilon_0 M_{PL}} \langle \sigma_{Av} \rangle^{-1}. \quad (1.67)$$

The evolution of the WIMP density during the early universe is shown in figure 1.12. Plugging in the appropriate constants and a dark matter density of $\Omega_\chi = 0.3$, gives an interaction cross section of $\mathcal{O}(10^{-27} \text{cm}^3 \text{s}^{-1})$. This result, known as the ‘‘WIMP miracle’’, shows that all of the dark matter observed today can be explained by thermal relic WIMPs that have interaction cross sections on the weak scale [46] [11] [57].

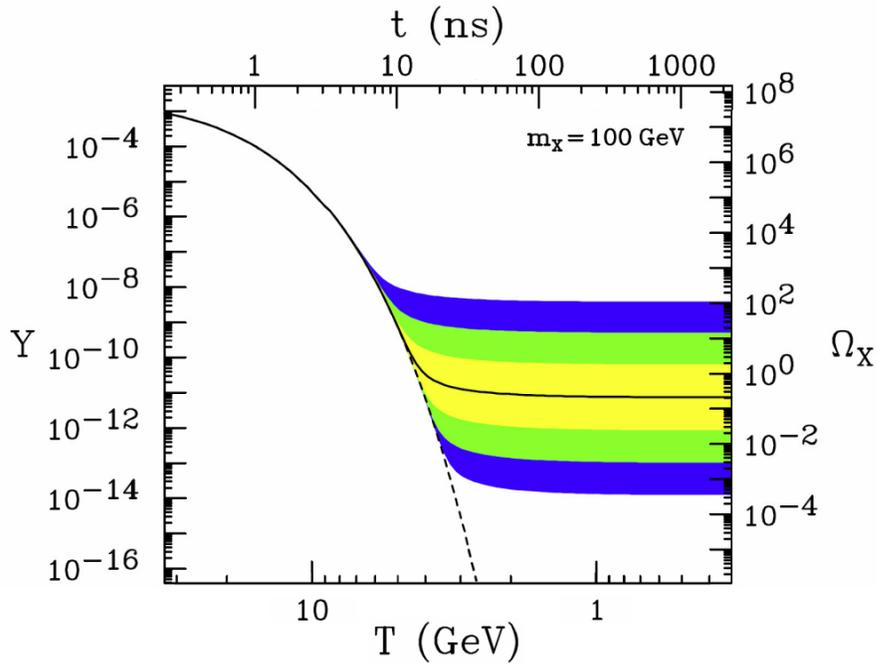


Figure 1.12: The comoving number density Y (left) and thermal relic density (right) of a 100 GeV WIMP as a function of temperature (bottom) or time (top). The dotted line shows the dashed line is the number density of particles that remain in thermal equilibrium. The solid line uses an annihilation cross section that yields the correct relic density. The yellow, green, and blue bands show cross sections that differ from the middle value by a factor of 10, 10^2 , and 10^3 respectively. From reference [11].

1.6 Conclusion and Thesis Outline

This chapter has reviewed the significant evidence that shows a large component of the universe is comprised as dark matter. The evidence spans multiple scales through its gravitational interactions with regular matter. Weak gravitational lensing searches continue to identify new sources of dispersed dark matter. The structure of anisotropies in the CMB also supports the existence of this primordial non-luminous matter and suggests dark matter may be particles with an interaction cross-section at the weak scale also known as WIMPs.

Chapter 2 outlines the kinematics of WIMP interactions that motivate direct detection searches. Chapter 2 describes the direct detection results and detector technologies ex-

plored over the last few decades. Chapter 3 gives an overview of the LZ detector and its first experimental result. Chapter 4 reviews the xenon purity monitors that LZ relies on to report krypton concentrations. Chapter 5 highlights the krypton removal campaign that purified the 10.4 tonnes of xenon used in LZ. Finally, chapter 6 describes a data driven search of the LZ data for excited state ^{85}Kr decays.

Chapter 2: Experimental Detection of WIMPs

Three complementary search techniques use independent interaction mechanisms, shown in figure 2.1. Together, they set exclusion limits on WIMP interaction cross-sections as they probe a wide range of WIMP masses between the MeV and TeV scale. Production searches try to produce WIMPs, on Earth, with accelerators. Direct detection searches look for energy deposited within a detector due to WIMP scattering consistent with signals predicted from the relative motion of the Earth and the WIMP halo. Indirect searches look for WIMP annihilation products from astrophysical sources.

This chapter reviews the three types of WIMP detection and the kinematics of WIMP scattering. A method of direct via WIMP-nucleon scattering is developed. The types of direct detection experiments motivated by WIMP-nucleon scattering are discussed. Finally, the current leading detector technology and its response to WIMP scattering events is reviewed.

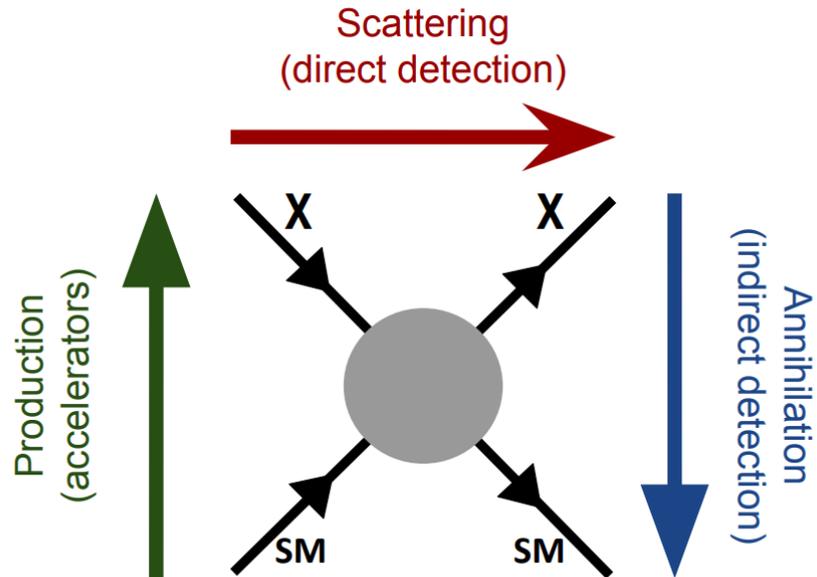


Figure 2.1: Schematic of WIMP-Standard Model interactions. The arrows indicated the direction of time for a given interaction. The “X” label denotes a WIMP and “SM” denotes a standard model particle. From reference [12].

Particle accelerators, like the Large Hadron Collider (LHC), attempt to produce WIMPs from standard model particles through proton-proton collisions. High energy collisions produce showers of standard model particles. The initial and final transverse momenta are compared. An excess of missing transverse momentum could be a potential WIMP signal [58]. Results from ATLAS and CMS, from LHC Run 2, find no significant excess of events outside of the Standard Model predictions [59] [60].

Indirect searches look for an excess of high energy standard model particles from astrophysical sources. One of the annihilation products from dark matter could be high energy neutrinos. Observatories like IceCube and Super-K have measured the neutrino flux and energy spectrum to set limits on WIMP masses from the MeV to TeV scale [61]. The non-interacting nature of neutrinos make them a challenge to measure, but they offer a direct probe of potentially WIMP-rich galactic centers. Super-Kamiokande’s

2018 publication of neutrino measurements have excluded WIMP masses ranging from 20 to 200 MeV, for self-annihilation cross-sections above $10^{-24}\text{cm}^3/\text{s}$ [62]. High energy gamma rays could be another product of WIMP annihilation. A combined analysis of high energy gammas from dwarf spheroidal galaxies used measurements from Fermi-LAT, HAWC, H.E.S.S., MAGIC, and VERITAS to set limits on self-annihilation cross-sections for WIMPs. The combined analysis, completed in 2021, excludes cross-sections from $10^{-27}\text{cm}^3/\text{s}$ to $10^{-23}\text{cm}^3/\text{s}$, for WIMP masses between 5 GeV and 100 TeV [63]. Cosmic rays could also be produced by WIMP annihilation. Cosmic ray experiments offer insight into high energy physics not accessible to colliders on Earth. An anomalous increase in the positron-electron fraction for energies between 10 GeV and 1000 GeV was reported by AMS in 2014 [64], PAMELA in 2009 [65], and DAMPE in 2017 [66]. This positron rate anomaly could be from self-annihilating WIMPs. Limits, calculated from the positron flux, exclude WIMP masses from 2 GeV to 100 GeV for self-annihilation cross-sections above $2 \times 10^{-27}\text{cm}^3/\text{s}$ [61].

2.1 Theory of WIMP Scattering

For elastic scattering the recoil energy, E_R , of a nucleus can be written as a function of the WIMP's incident energy $E_i = M_\chi v_i^2/2$

$$E_R = E_i r \frac{(1 - \cos \theta)}{2} \quad (2.1)$$

where the WIMP scatters through the nucleus with an angle θ and the dimensionless kinematic factor is

$$r \equiv \frac{4\mu_A^2}{M_\chi M_A} = \frac{4M_\chi M_A}{(M_\chi + M_A)^2} \leq 1. \quad (2.2)$$

The reduced mass is $\mu_A = M_\chi M_A / (M_\chi + M_A)$ for a WIMP mass M_χ and target nuclear mass M_A .

The total differential nuclear recoil rate, $dR(E_R)/dE_R$, is the integral of the uniform recoil distribution, $1/(E_i r)$, weighted by the differential interaction rate, $dR(E_i)$, of an incident WIMP with energy E_i . For isotropic scattering the nuclear recoil energy is uniformly distributed from 0 to $E_i r$. To produce nuclear recoil energies up to $E_i r$ an incident WIMP must have a minimum energy of at least $E_{min} = E_R/r$. This can also be expressed as the minimum WIMP velocity $v_{min} = \sqrt{2E_R/(rM_\chi)}$. The maximum incident WIMP energy, E_{max} , that contributes to a given nuclear recoil energy is determined by the galaxy's escape velocity of 540 km/s.

$$\frac{dR(E_R)}{dE_R} = \int_{E_{min}}^{E_{max}} \frac{dR(E_i)}{E_i r} dE_i. \quad (2.3)$$

In this approximation setting $E_{max} = \infty$ is appropriate and yields reasonable results as the velocity distribution of the WIMP halo model is taken to be an ideal, non-interacting gas, with a Maxwellian distribution given by

$$f(\mathbf{v}_i, \mathbf{v}_E) = \frac{1}{k} e^{-(\mathbf{v}_i + \mathbf{v}_E)^2 / v_0^2}. \quad (2.4)$$

Here $k = \pi^{3/2} v_0^3$ is a normalization factor and is calculated by integrating over all possible halo velocities. The WIMP velocity in the lab frame is given by the vector sum of the Earth's velocity and the WIMP halo velocity $\mathbf{v}_i = \mathbf{v}_E + \mathbf{v}_{halo}$. The $-1/v_0^2$ term, where $v_0 \approx 220$ km/s, is the local circular velocity in the galaxy and significantly suppresses any rate contributions from velocities greater than 540 km/s.

The differential interaction rate per kg of detector material with an atomic mass A for an incident WIMP is given by

$$dR(\mathbf{v}_i) = \frac{N_0}{A} n_0 f(\mathbf{v}_i, \mathbf{v}_E) \sigma v_i d^3 \mathbf{v}_i \quad (2.5)$$

where N_0 is Avogadro's number, σ is the WIMP-nucleon scattering cross-section, and n_0 is the local WIMP number density.

The local WIMP number density for a given WIMP mass is $n_0 = \rho_\chi/M_\chi$ [13]. Studies, completed in 2013 by Zhang [67] and Bovy [68], use lensing data from the Sloan Digital Sky Survey to calculate local WIMP mass densities of $\rho_\chi = 0.25 \pm 0.09 \text{ GeV}/c^3$ and $\rho_\chi = 0.3 \pm 0.094 \text{ GeV}/c^3$ [69]. Assuming $\rho_\chi = 0.3 \text{ GeV}/c^3$ and a 100 GeV WIMP mass this would amount to a local number density of 3 WIMPs per liter. From the average local circular velocity, the WIMP flux is approximately 700 million WIMPs/m²/s.

The WIMP-nucleon scattering cross-section describes how likely a WIMP is to scatter. The total cross-section is the product of a constant zero momentum-transfer cross-section and a momentum dependent form factor $\sigma(q) = \sigma_0 F^2(q)$. The zero momentum-transfer cross-section is the sum of spin-independent and spin-dependent terms $\sigma_0 = \sigma_{0,SI} + \sigma_{0,SD}$. The spin-independent term is

$$\sigma_{0,SI} = \frac{4\mu_A^2}{\pi} [Zf_p + (A-Z)f_n]^2 \approx \frac{4\mu_A^2}{\pi} f_n^2 A^2, \quad (2.6)$$

where A is the total number of nucleons. The coupling constants for neutrons and protons, $f_{p,n}$, are approximately equal for most WIMP models. Note that the spin-independent cross-section is proportional to A^2 , motivating the selection of a large nuclei target. The spin-dependent term is

$$\sigma_{0,SD} = \frac{32G_F^2 \mu_A^2}{\pi} \left(\frac{J+1}{J} \right) (a_p \langle S_p \rangle + a_n \langle S_n \rangle)^2 \quad (2.7)$$

where G_F is Fermi's constant, J is the total nuclear spin. The coefficients $a_{p,n}$, are spin coupling constants for the expectation values of the proton and neutron spins $\langle S_{p,n} \rangle$.

To estimate the scattering rate, the cross-section can be taken to be constant $\sigma = \sigma_0$,

with $v_E = 0$, and $v_{esc} = \infty$. Equation (2.3) can be written as

$$\begin{aligned}\frac{dR(E_R)}{dE_R} &= \int_{E_R/r}^{\infty} \left(\frac{1}{r^{\frac{1}{2}} M_{\chi} v^2} \right) \frac{R_0}{2\pi v_0^4} v e^{-v^2/v_0^2} (4\pi v^2) dv \\ &= \frac{R_0}{r^{\frac{1}{2}} M_{\chi} v_0^2} \int_{v_{min}}^{\infty} \frac{2}{v_0^2} e^{-v^2/v_0^2} v dv.\end{aligned}\quad (2.8)$$

with the total scattering rate defined to be

$$R_0 \equiv \frac{2}{\sqrt{\pi}} \frac{N_0}{A} n_0 \sigma_0 v_0. \quad (2.9)$$

Changing variables to energy and recalling that $E_{min} = E_R/r = M_{\chi} v_{min}^2/2$ gives

$$\frac{dR(E_R)}{dE_R} = \frac{R_0}{E_0 r} \int_{E_R/r}^{\infty} \frac{2}{M_{\chi} v_0^2} e^{-2E/(M_{\chi} v_0^2)} dE = \frac{R_0}{E_0 r} e^{-E_R/(E_0 r)} \quad (2.10)$$

where the most probable WIMP energy is defined as $E_0 \equiv M_{\chi} v_0^2/2$.

The total scattering rate R_0 and the mean recoil energy $\langle E_R \rangle = E_0 r$ can be found [13] [70] using the values from the flux calculation. For a xenon target, the total scattering rate for a 100 GeV WIMP is

$$R_0 \approx \frac{2}{\pi} \frac{N_0}{131 \text{ g/mole}} \frac{0.3 \text{ GeV/cm}^3}{100 \text{ GeV}} (220 \text{ km/s}) \left(\frac{\sigma_0}{1 \text{ pb}} \right) = \left(\frac{\sigma_0}{1 \text{ pb}} \right) (0.0296) \frac{\text{events}}{\text{kg} \cdot \text{day}}. \quad (2.11)$$

Assuming a 1 pb interaction cross-section there will be one WIMP event every 34 days per kg of xenon. The mean recoil energy is

$$\langle E_R \rangle = r E_0 = 0.99 \times \frac{1}{2} 100 \text{ GeV} \left(\frac{220000}{3 \times 10^8} \right)^2 \approx 27 \text{ keV}. \quad (2.12)$$

Many assumptions were made to approximate the energy and total scattering rate. In practice, the energy of a detectable event is set by the minimum energy threshold of a de-

detector. The fraction of events above the energy threshold E_{th} is $e^{-E_{th}/(E_0 r)}$. Since the event rate falls exponentially with energy it is critical that experiments minimize their energy thresholds to maximize the rate of detectable events.

WIMPs with velocities above $v_{esc} = 540$ km/s are not bound to the galaxy and are not part of the local WIMP dark matter halo. A 100 GeV WIMP with velocity v_{esc} would produce a maximum recoil energy of a 198 keV. The fraction of escaped WIMPs has a negligible impact on the detectable event rate as the maximum energy is a factor of 7 larger than the mean energy $\langle E_R \rangle = 27$ keV. For low mass WIMPs the detectable rate is suppressed when using a heavy target despite an increased WIMP number density. This is due to a smaller mean recoil energy that results in a larger fraction of events falling below a fixed detector threshold energy.

Another simplification was the assumption of a stationary Earth. The actual circular velocity of the Earth as it orbits the sun is

$$v_e(t) = 232 + 15 \cos\left(2\pi \left(\frac{t - 152.5}{365.25}\right)\right) \text{ km/s} \quad (2.13)$$

where t is the time in days from January 1st. This annual modulation with respect to the WIMP “rain” periodically shifts the recoil spectrum and causes the detectable event rate above experimental threshold to fluctuate as much as 7% [13].

A final layer of complexity is added when the nuclear form factor is considered and σ can no longer be treated as a constant. As the momentum transfer increases the scattering interaction begins to probe the internal structure of the target nucleus. The nucleus can be modeled as a sphere of constant density up to an effective radius r_n . At the edge of the sphere there is a soft skin region of thickness t where the density falls to zero with a Gaussian edge. This model is shown in figure 2.2.

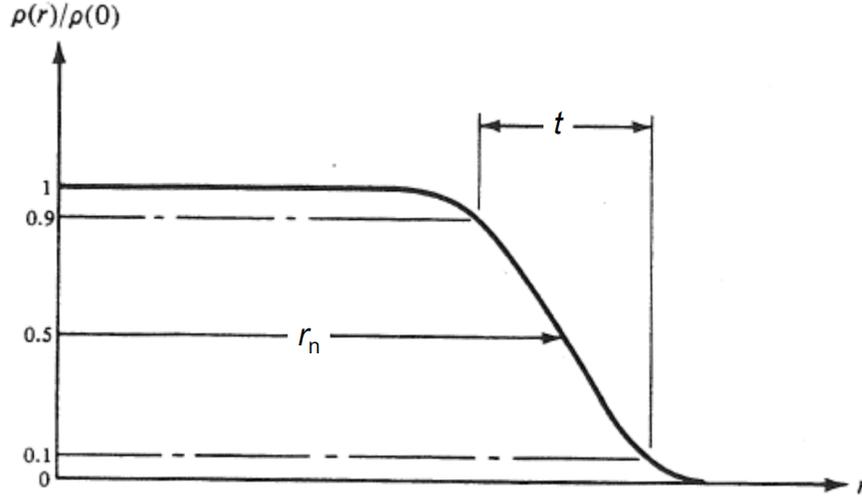


Figure 2.2: The nuclear density as a function of radius for a sphere of effective radius r_n and skin thickness t . The density remains constant before falling to zero with a Gaussian edge. From reference [13].

For the spin-independent case, a good approximation is the Woods-Saxon form factor

$$F(q) = \frac{3[\sin(qr_n) - qr_n \cos(qr_n)]}{(qr_n)^3} e^{-(qs)^2/2} \quad (2.14)$$

where $s = 0.9$ fm is related to the skin thickness and

$$r_n = \sqrt{(1.23 \text{ fm } A^{1/3} - 0.6 \text{ fm})^2 + \frac{7}{3}(0.52\pi \text{ fm})^2 - 5s^2}. \quad (2.15)$$

At low momentum transfer, the de Broglie wavelength λ is larger than the r_n , a WIMP will scatter coherently as if the nucleus was a solid sphere. At high momentum transfers, where $\lambda < r_n \propto A^{1/3}$, coherence is lost and the scattering cross-section is suppressed [13]. Despite smaller targets being less susceptible to suppression from their form factor, a high mass target with a low energy threshold is still favorable as seen in figure 2.3.

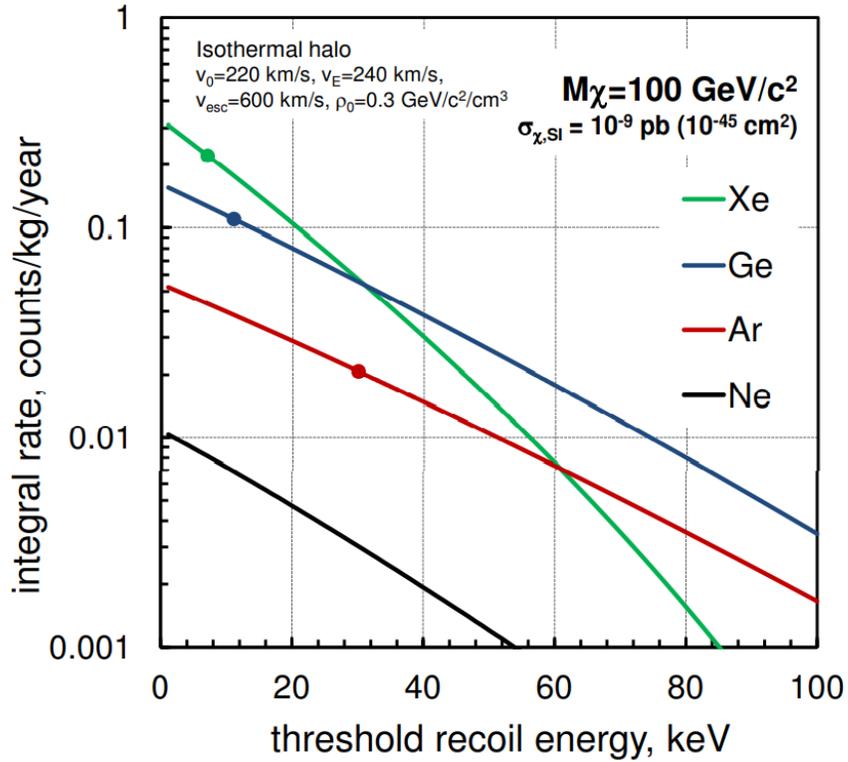


Figure 2.3: The integrated, spin-independent, WIMP scattering rate for different atomic nuclei as a function of threshold energy. At low energies coherent scattering off of large nuclei, like xenon, is favorable. Event rates at higher energies are drastically suppressed in xenon due to decoherence. Germanium rates suffer modestly, but the lower mass targets are unaffected in this range. From reference [14].

2.2 Experimental Overview

The recoil energy deposited by WIMPs can be dissipated in three distinct channels: heat, scintillation, and ionization. Heat is measured as a temperature increase in low temperature detectors. Scintillation light is produced by electron excitation and relaxation in atomic orbitals. Ionization occurs when electrons gain enough energy to be liberated from the parent atom. Detectors use one or two of these energy channels to reconstruct scattering events and probe potential WIMP signals. The detectors are usually compared by their sensitivities to the WIMP-nucleon cross-section for a specific WIMP mass. Sensitivities

to cross-sections vary at different WIMP masses due to detector thresholds and designs. Results from different experiments are shown in figure 2.4.

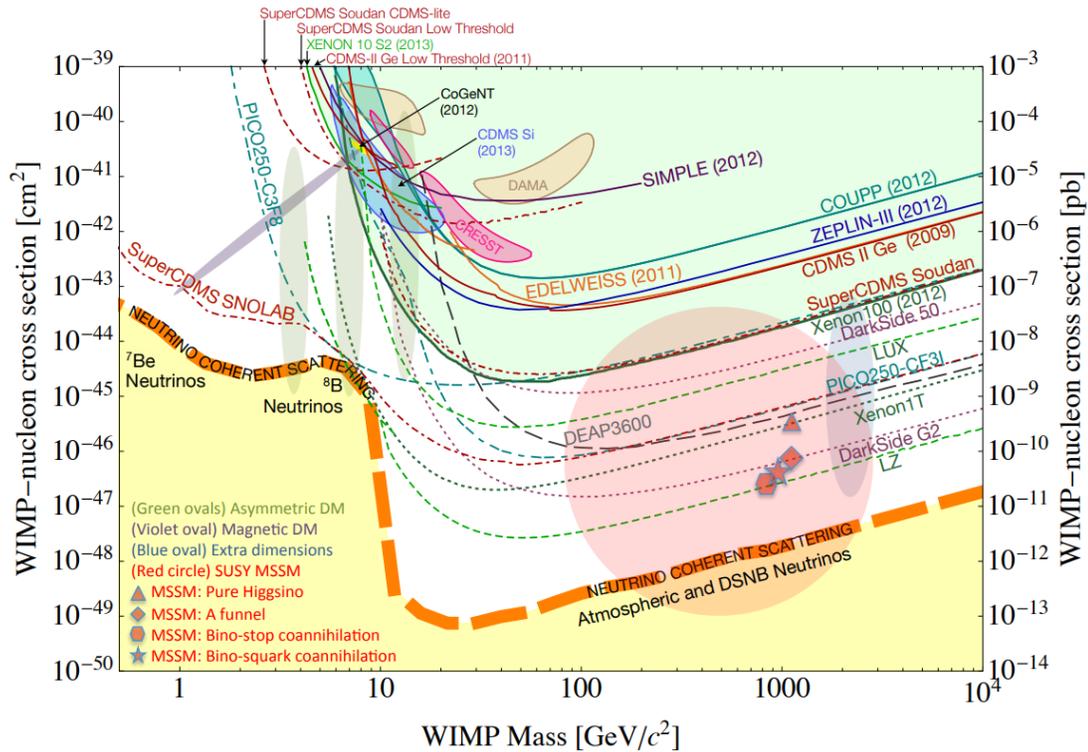


Figure 2.4: A compilation of WIMP-nucleon spin-independent cross-section exclusion limits as a function of WIMP mass. Different theoretical dark matter models are shown in the shaded regions with no outlines. Solid curves are completed measurements and dotted lines are future projections as of 2013. The yellow region on the bottom is the neutrino fog. This fog will obscure direction detection measurements due to an irreducible neutrino background. From reference [15].

2.3 Crystal detectors

Crystal charge detectors are designed to have low energy thresholds and fine energy resolution. They take advantage of the low energy required to produce electron hole pairs in reverse biased semiconductors. Upon scattering, charge is freed and measured as a current signal. Crystal charge detectors suffer from high noise as they cannot discriminate electron recoils from nuclear recoils. They can reconstruct event position, but contain many arrays

or strings of targets. As a result these detectors have poor surface-area-to-volume scaling when compared to detectors with a single large target volume. This results in backgrounds due to surface components permeating most of the detector volume. Experiments of this type are usually sensitive to WIMP masses below 10 GeV [71].

CDEX is a 10 kg germanium charge detector submerged in liquid nitrogen located at the Jinping Underground Laboratory in China. An exposure of 102.8 kg·day yielded results with an exclusion limit for spin-independent cross-sections above $8 \times 10^{-42} \text{ cm}^2$ and spin-dependent cross-sections above $3 \times 10^{-36} \text{ cm}^2$ for 5 GeV mass WIMPs [72]. CoGeNT, located at the Soudan Underground Laboratory, is another germanium charge detector. They reported a limit excluding spin-independent cross-sections above $6.7 \times 10^{-41} \text{ cm}^2$ for a WIMP mass of 9 GeV. Although a slight excess of events are noted it cannot be distinguished from backgrounds [73]. A silicon detector from DAMIC located at the SNOLAB underground laboratory excluded spin-independent cross-sections greater than 10^{-39} cm^2 for WIMP masses greater than 3 GeV over a 0.6 kg day exposure [74]. An above ground 0.019 g·day exposure test of the SENSEI detector at Fermilab reported an exclusion of WIMP-electron spin-independent cross-sections that range from 10^{-25} cm^2 to 10^{-28} cm^2 for WIMP masses between 0.5 keV and 4 MeV [75] [71].

Crystal scintillators are another type of single channel detector. They benefit from high mass targets like NaI ($A=127$) and simple designs that allow for long run times. Photomultiplier tubes (PMTs), sensitive to individual photons, monitor the crystal targets for scintillation. These detectors suffer from high backgrounds due to embedded impurities, but some recoil discrimination is possible from pulse shape discrimination. The inability to reconstruct position makes identification of a low background detector volume impossible. Detectors that can reconstruct even positions can take advantage of a “fiducial volume” or a volume with low backgrounds, usually displaced from detector surfaces. Rather than search for individual WIMP-like events, crystal scintillators often search for annual mod-

ulations consistent with equation (2.13). DAMA/LIBRA is a 250 kg high purity NaI(Tl) detector that claims to have detected a modulation consistent with a WIMP halo model to a significance of 12.9σ [76]. Despite this, COSINE-100, which also uses NaI(Tl) crystals, has been unable to reproduce the effect claimed by DAMA/LIBRA. In an attempt to reduce the backgrounds of COSINE-100, which are 2 to 3 times higher than DAMA/LIBRA, an upgrade is planned. COSINE-200 will increase the amount of target crystal mass from 106 kg to 1000 kg while reducing the background rate per kg to less than that of DAMA/LIBRA [77] [71].

Cryogenic crystal detectors can measure heat deposited from recoils through a temperature increase. To maximize sensitivity, they must be operated at cryogenic temperatures (≤ 50 mK). Transition edge detectors are current-carrying wires held at the transition temperature of their superconducting state. A small temperature increase causes a spike in resistance and a measured drop in current. These detectors can also measure charge collection. With two energy collection channels, cryogenic crystal detectors can discriminate electron and nuclear recoils to reduce backgrounds. Due to the ultra low temperatures required, they are difficult to maintain and scale. They also suffer from the same fiducialization and surface-area-to-volume issues as other crystal charge detectors. They benefit from the low energy threshold and good energy resolution. A 60g germanium detector cooled to 25 mK was able to measure the heat deposit of recoil events [78]. The large phonon energy produced in nuclear recoils was used to discriminate electron recoil events with comparable charge yields. In the few keV range, a discrimination of 10:1 between nuclear and electron recoils was achieved. This can be used to significantly reduce backgrounds in the WIMP search region of interest. SuperCDMS, another dual channel cryogenic crystal detector, reported a spin-independent cross-section exclusion limit of 1.4×10^{-44} cm² for a 46 GeV WIMP mass [79] [71].

2.4 Bubble Chambers

Bubble chambers use superheated liquids held at or just above their boiling point without actually boiling. When a nuclear recoil occurs, heat causes a local phase change and produces a bubble. Electron recoil events do not produce enough heat energy to cause a phase change and do not contribute any backgrounds. When a bubble forms cameras and acoustic sensors record the event. The cameras allow for fiducialization while acoustic sensors are used to tag alpha events. Immediately following an event, the chamber is mechanically compressed to restore the superheated liquid. During this reset, the bubble chamber incurs a long dead-time. These detectors cannot reconstruct energies, because a bubble forms for any event above the detector's threshold energy. The PICO experiment used 52 kg of superheated C_3F_8 to produce an exclusion limit on the WIMP-proton spin-dependent cross-section above $3.4 \times 10^{-41} \text{ cm}^2$ for a 30 GeV WIMP mass over a 1167 kg day exposure [80] [71].

2.5 Dual-Phase Time Projection Chambers

The noble elements, argon and xenon, make convenient detector materials that can be liquefied above liquid nitrogen temperature (77 K). Both have relatively low ionization energies and both produce scintillation light. As shown in figure 2.5, dual-phase time projection chambers (TPCs) have produced the most stringent WIMP-nucleon cross-section limits. Xenon detectors have recently dominated the field. From XENON10 in 2007 to LUX-ZEPLIN (LZ) in 2022, this rapidly maturing detector technology has scaled from 5.4 kg to over 5.5 tonnes of fiducial liquid xenon target. The WIMP-nucleon scattering cross-section exclusion limit has been pushed from $4.5 \times 10^{-44} \text{ cm}^2$ [81] down to $6.5 \times 10^{-48} \text{ cm}^2$ [24]. Although LZ is still taking data, its first science run has produced the most

stringent exclusion limits to date, shown in figure 3.15.

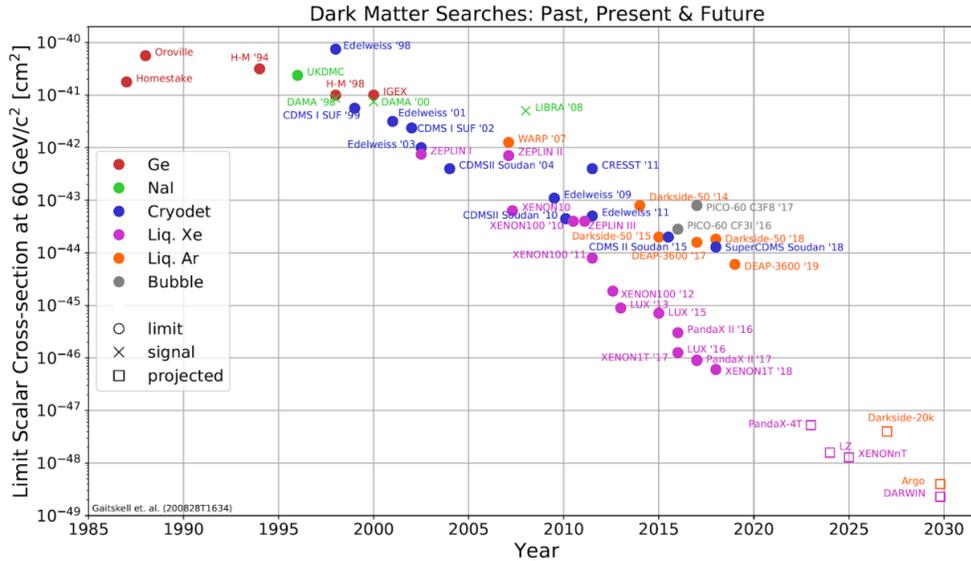


Figure 2.5: The reported spin-independent cross-section exclusion limits for a 60 GeV mass WIMP from direct detection experiments over time. In recent years, liquid noble detectors have consistently produced the most stringent limits. From reference [16].

A dual-phase TPC contains a liquid target with a gas pocket at the top. Two PMT arrays face each other from the top and bottom of the detector. During a scattering event, the recoiling target ionizes and excites nearby atoms. The free electrons drift vertically under the influence of an electric field maintained by high voltage grids. The remaining positively charged atoms and excited atoms form excited molecules with nearby neutral atoms. The excited molecule then relaxes, releasing the prompt scintillation light that forms the S1 signal as described in section 2.5.2. The free electrons continue to drift vertically towards the liquid surface where an extraction field pulls the electrons out of the liquid and into the gas phase. As the electrons traverse the gas phase they excite nearby xenon atoms and produce electroluminescence light. This S2 signal, also collected by the PMT arrays, depends on the extraction field and the amount of charge produced during the initial recoil event.

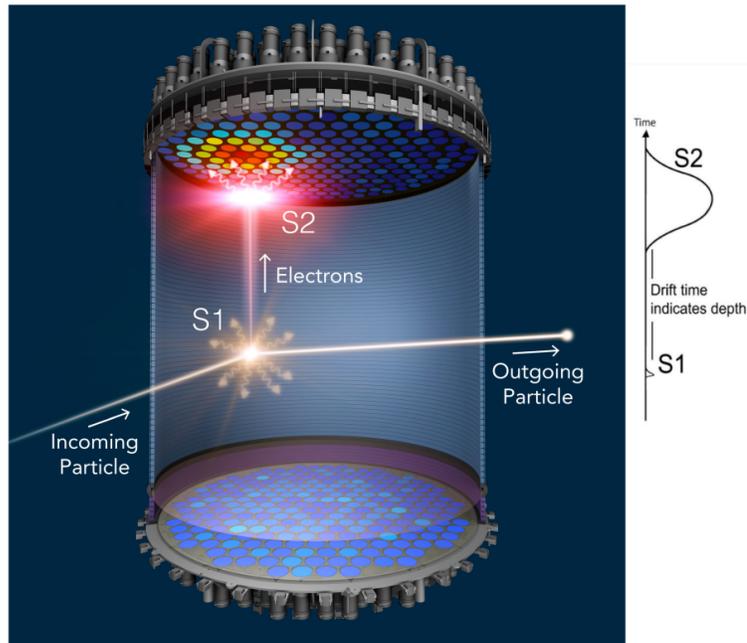


Figure 2.6: An incoming particle scatters off of a target medium within the TPC and exits the detector. A prompt scintillation signal (S1) is produced and electrons are freed. An electric field drifts ionized electrons to the liquid surface. An extraction field pulls electrons into the gas phase. The electrons excite gaseous xenon and produce a secondary light signal (S2). Energy and position are reconstructed from signal timings and sizes. An example of the summed PMT waveforms, containing an S1 and S2, is shown on the right. Taken from reference [17].

The sum of the S1 and S2 signals determines the total energy of the event. The ratio of the S1 and S2 signals can be used to identify nuclear and electron recoils. The S2 signal is produced at the top of the detector and the PMT array hitmap is used to reconstruct the event's XY position. The Z position is measured using the time it takes for the charge to drift from the interaction site to the liquid surface. The drift time and the drift velocity are used to calculate the event's depth. An illustration of this process is shown in figure 2.6.

2.5.1 Detection Mediums

While there are subtle differences between dual-phase argon and xenon detectors, their operating principles are the same. In both mediums event reconstruction and discrimination

depend on charge and light measurements. Both benefit from a low background fiducial volume. Both are inert and can be actively purified via commercially available chemical getters.

One key difference is the intrinsic ^{39}Ar background. This cosmogenically-activated isotope beta decays with a 269 year half-life. In large detectors, argon extracted from air contains enough ^{39}Ar that the event rate is dominated by this background, and may overwhelm any potential WIMP signal [82]. DarkSide-20k plans to extract and distill over 23 tonnes of argon from underground sources to ensure a manageable ^{39}Ar rate [83].

Although xenon does have cosmogenically-activated isotopes, they are relatively short-lived. The longest lived cosmogenically activated isotope, ^{127}Xe , has a half-life of only 36.4 days. This is of little concern since it decays away over the lifetime of an underground experiment. Xenon detectors contain naturally occurring ^{124}Xe , a two-neutrino double-electron capture background, and ^{136}Xe , a two-neutrino double-beta decay background [17]. These isotopes are extremely long lived with half-lives of 1.1×10^{22} [84] and 2.165×10^{21} years [85] respectively. In LZ's first science run the combined xenon backgrounds from ^{124}Xe , ^{127}Xe , and ^{136}Xe account for less than 9% of the total backgrounds [25]. With the exception of each medium's isotopic background contributions, many of the other backgrounds are similar and come from the detector itself.

Another difference between these noble elements is their atomic mass. Argon has an atomic mass $A = 40$ and xenon $A = 131$. From equation (2.6), the spin-independent cross-section is proportional to A^2 . To be comparably sensitive to WIMP interactions, argon detectors must compensate by using a larger target mass. In general a larger detector is more complicated to design, build, and maintain. To collect charge, an electric field is applied across the liquid. An optimal field strength is ≈ 200 V/cm [83] [17]. As detectors become larger, especially in the case of argon detectors, a higher voltage difference must be achieved between the grids to maintain the electric field. These grids experience an

attractive force and can deflect, impacting detector performance.

Liquid xenon is also more than twice as dense as liquid argon. As a result, the interaction length of background radiation is short. In liquid xenon, backgrounds due to trace radioactivity in TPC components, deposit most of their energy within a few cm. This property, known as “self shielding”, maximizes the potential fiducial volume. In LZ, a space between the TPC and cryostat walls is filled with liquid xenon. This xenon skin region shields the TPC from background signals produced by the cryostat vessel. This skin region is also instrumented as an independent coincidence veto detector [17]. In addition to the stronger self shielding offered by xenon, PMT photocathodes are available that efficiently convert the 178 nm xenon scintillation light from photons to photoelectrons that are amplified and measured by detector electronics. This is not the case for argon, instead, a wavelength shifter is required to re-transmit the 128 nm scintillation light to 425 nm for which there are PMTs with appropriate photocathodes. These wavelength shifters introduce losses in light collection efficiency [71]. Argon remains a promising target medium, but currently xenon has shown to be more effective in producing stringent WIMP-nucleon cross-section exclusion limits.

2.5.2 Xenon Microphysics

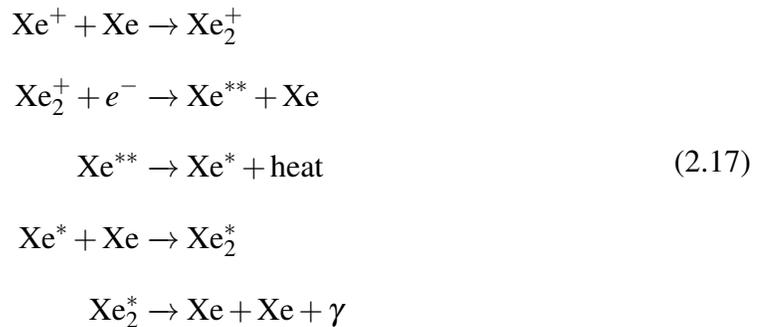
The previous description of scattering events was sufficient for a general detector overview, but the interactions that allow for signal discrimination require a deeper analysis. During a scattering event in xenon, an incoming particle imparts energy onto a nucleus or atomic electron. As the electron or nuclei recoils, it tracks through the bulk xenon exciting (Xe^*) and ionizing (Xe^+) nearby atoms.

An excited xenon atom can temporarily bond with a nearby xenon atom and produce an excimer molecule (Xe_2^*). As the newly formed excimer relaxes, the molecule breaks apart

and the excess energy is carried away in a 178 nm vacuum ultraviolet (VUV) photon. The xenon is transparent to the VUV photon with a scattering length of 36.4 cm in liquid xenon [86]. This process is described in equation (2.16).



Ionized xenon atoms also produce prompt scintillation through electron recombination. Following ionization a xenon ion forms a molecule (Xe_2^+) with another xenon atom. Many ionized electrons escape the interaction region and are collected at the anode, but some recombine with nearby ionized molecules. The amount of recombination depends on the energy of the recoil, the strength of the drift field, and whether the event is a nuclear or electron recoil. Upon recombination, the Xe_2^+ molecule breaks up, and one of the xenon atoms becomes doubly-excited Xe^{**} . The Xe^{**} dissipates some energy as heat producing the singly excited Xe^* state. This excited xenon atom undergoes the same excimer formation and relaxation to produce scintillation light. This process is described in equation (2.17).



The excimer contains both a singlet and a slightly longer lived triplet state. The half-

lives have been observed to be in the range of 2-4 ns for the singlet state and 21-28 ns for the triplet state [87]. These short timescales give the S1 signals a fast rise time. A similar excitation and ionization process occurs in argon, where the singlet state has a 6.5 ns half-life and the triplet state has a 1 μ s half-life [14]. The population of the singlet and triplet states in both argon and xenon depend on the the recoil type. The resulting pulse shape is used to discriminate electron and nuclear recoils in argon detectors [82]. A similar technique can be employed in xenon, but due to the similar half-lives, the recoil discrimination achieved is modest in comparison to argon.

The electrons that do not recombine move vertically through the liquid at a constant velocity set by the strength of the drift field. The survival fraction of these electrons depends on the amount of electronegative impurities and the depth at which the event occurs. Upon reaching the liquid surface the electrons are exposed the extraction field. This pulls the electrons from the liquid into an electroluminescence region. The electrons are accelerated through the gas until they are energetic enough to excite nearby xenon atoms. These excited xenon atoms produce secondary VUV scintillation light via excimer relaxation. This process repeats as electrons accelerate and collide with the gaseous xenon atoms in their path. The number of photons produced is proportional to the number of electrons, so this known as proportional scintillation. The amount of light produced per electron depends on the electric field in the gas phase and the distance from the liquid surface to the anode. This powerful gain mechanism allows for the detection of single electrons extracted from the liquid surface.

2.5.3 Energy Reconstruction

When energy is deposited in the detector from a scattering event, light from the scintillation and charge channels are measured by the PMT arrays. The quantities S1 and S2

are the integrated waveforms corresponding to these channels and are measured in units of photons detected (phd). Converting these values to quanta of electrons (n_e) or photons (n_γ) requires knowledge of the detector-specific gain factors g_1 and g_2 .

$$S1 = g_1 n_\gamma \tag{2.18}$$

$$S2 = g_2 n_e$$

These gains encapsulates the light collection efficiency due to photocathode efficiency, detector reflectivity, and detector geometry. In the case of charge collection, the gas gain, extraction efficiency, and electron lifetime must also be considered. While more detailed corrections can be applied by treating each detector subvolume as having its own g_1 and g_2 , here they are treated as uniform.

The total recoil energy from the scattering event is the sum of the two quanta times a conversion factor. In terms of S1 and S2 this is

$$E = \frac{W}{L}(n_\gamma + n_e) = \frac{W}{L} \left(\frac{S1}{g_1} + \frac{S2}{g_2} \right). \tag{2.19}$$

W is the average energy required to produce a single ionized or excited atom and measured to be $13.7 \pm 0.2 eV$ [88] [89]. L is the energy and recoil type dependent Lindhard factor. It compensates for charge quenching due to energy dissipated as heat during nuclear recoils. In electron recoils, very little energy lost to heat and L is defined as 1. In nuclear recoils a significant amount of energy is carried away as heat in the form of motion. Measurements of L have been recorded at a variety of nuclear recoil energies ranging from a $E_n = 2.8 \text{ MeV}$ neutron beam. Measurements were taken as neutrons were scattered off of a xenon target exposed to different electric fields. The scattered neutron angle θ determines the nuclear

recoil energy E_r .

$$E_r = E_n = \frac{2m_n m_{Xe}}{(m_n + m_{Xe})^2} (1 - \cos \theta) \quad (2.20)$$

The effective Lindhard factor is the ratio of the equivalent electron energy (E_{ee}) to the true nuclear recoil energy (E_r), corrected for scintillation yields at zero electric field. The equivalent electron energy is defined as the S1-only extrapolated energy, fit against a 122 keV gamma recoil produced by a ^{57}Co source. The scintillation fractions, $S_{e,n}$, correct for light yields under different electric fields for electron and nuclear recoils [21].

$$L_{eff} = \frac{E_{ee}}{E_r} \cdot \frac{S_e}{S_n} \quad (2.21)$$

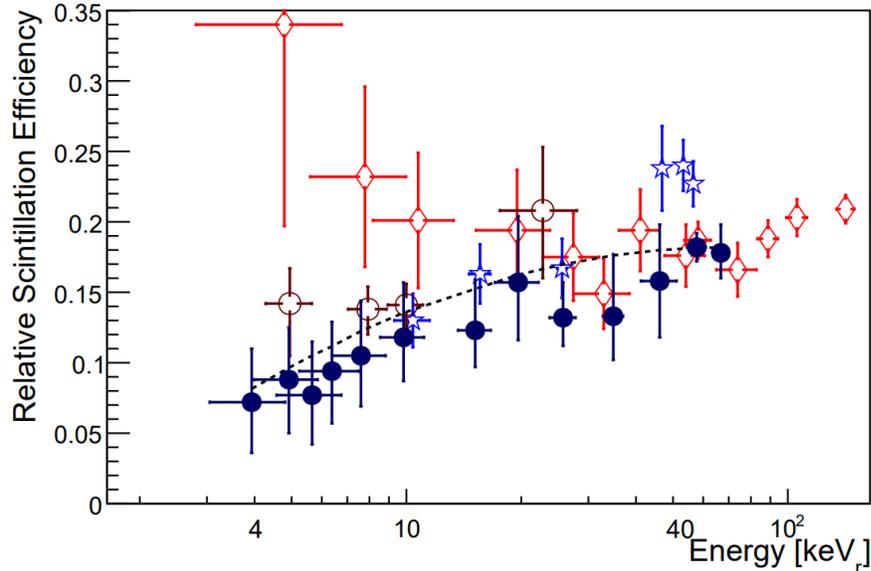


Figure 2.7: Measurements of L_{eff} , referred to here as relative scintillation efficiency, for nuclear recoil energies below 100 keV. The solid dots represent measurements made by Manzur et al. using 2.8 MeV neutrons. Their theoretical model is shown as a dashed line. Included is data from Aprile (blue stars and brown circles) [18] [19] and Chepel (red ovals) [20]. From reference [21].

Measurements of the Lindhard factor show how much additional energy is lost to heat in a nuclear recoil event when compared to an electron recoil event at the same energy under

zero electric field. The Lindhard factor is found to depend on energy and takes values from 0.07 to 0.18 for nuclear recoil energies ranging from 4 keV to 67 keV, as shown in figure 2.7 [90] [91] [89] [21] [92].

A more recent measurement of low energy nuclear recoils was conducted in the Large Underground Xenon (LUX) detector using a collimated beam of mono-energetic 2.45 MeV neutrons produced by a deuterium-deuterium (DD) generator. The scattering angle of multi-site neutron interactions was used to reconstruct the true energy of the initial nuclear recoil site. The LUX experiment measured the absolute charge and light yields of nuclear recoil events from 1.1 keV to 74 keV at an average average electric field of 180 V/cm. From these measurements LUX presented a parameterization of recoils that relies on the total light and charge yields of electron and nuclear recoils to discriminate one another as opposed to the ratio use in Lindhard models. This new parameterization is a better fit to data across the range of nuclear recoils measured, as shown in figure 2.8 [22].

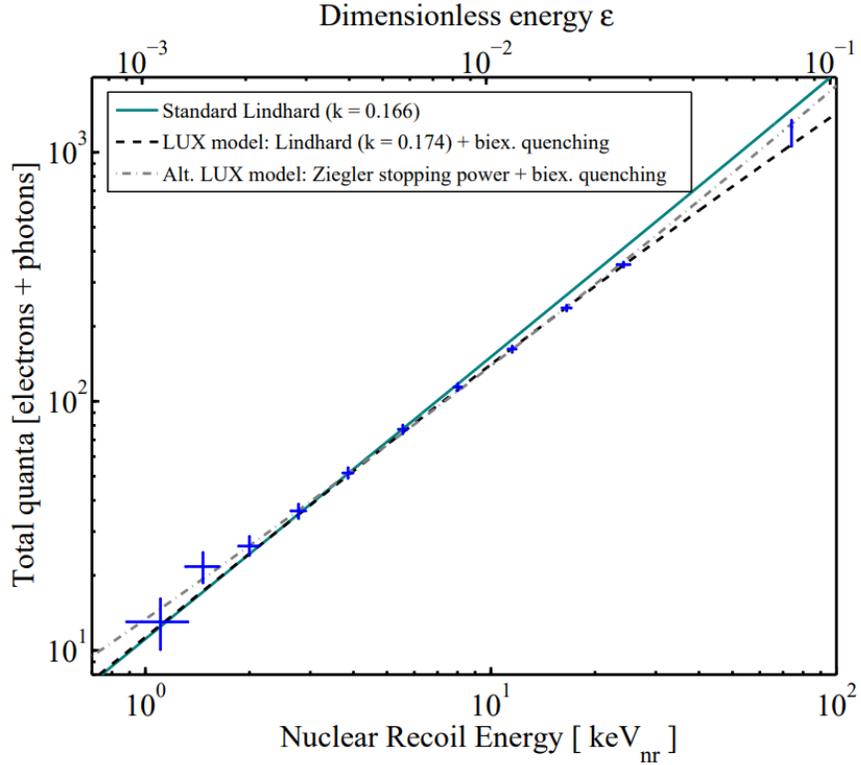


Figure 2.8: Measurements of the total quanta of light and charge produced during a nuclear recoil event from a 2.45 MeV neutron. The true recoil energy of the initial neutron recoil is given on the x axis and found using the scattering angle calculated from the location of a secondary recoil. The standard Lindhard and LUX specific Lindhard models are shown with a solid blue and dashed black line. The alternative parameterization is shown by the dotted grey line. A biexcitation quenching process is accounted for that reduces the total number of quanta at recoil energies above 10 keV due to increased exciton density [22].

From the description of charge recombination in section 2.5.2 the amount of electrons and photons produced is a mixture of the number of excited (N_{ex}) and ionized (N_{ion}) xenon atoms. This mixing is described by the recombination fraction r and depends on the recoil type as well as the local electric field strength.

$$\begin{aligned}
 n_\gamma &= N_{ex} + rN_{ion} \\
 n_e &= (1 - r)N_{ion}
 \end{aligned}
 \tag{2.22}$$

Measurements of electron and nuclear recoils in liquid xenon show that the ratio of excited

xenon atoms to xenon ions N_{ex}/N_{ion} is ≈ 1 [90] [89] [21] and 0.06 [93] respectively. The difference between N_{ex}/N_{ion} for nuclear and electron recoils manifests as two distinct bands when plotting events in the S1 vs S2 space. The separation of these bands allow for electron recoil discrimination in dual-phase xenon TPCs. A calibration plot from LZ shows the separation of these bands in figure 3.7.

2.6 Conclusion

This chapter has focused on the framework and progress of WIMP searches. A brief review of collider and indirect searches was presented. A method of direct detection through observations of recoiling nuclei is developed. The spin-independent WIMP-nucleon scattering cross-section is shown to be proportional to the number of nucleons squared in equation (2.6) and figure 2.3. Detector technologies and their recent results are discussed. Dual-phase time projection chambers, as shown in figure 2.5, have produced the most stringent exclusion limits to date. The end of this chapter gives an overview of xenon microphysics that lead to light and charge signals which can be used to reconstruct and differentiate recoil events.

Chapter 3: The LUX-ZEPLIN Detector

My description of WIMP dark matter searches has slowly narrowed to the subject of a dual phase liquid xenon TPCs. Over a century of astrophysical evidence supports the existence of dark matter. Modern theoretical models point towards a WIMP candidate. Various detector technologies have been developed in an effort to extend the search for WIMPs. Finally we arrive at the LUX-ZEPLIN (LZ) dark matter detector, the current leading WIMP search experiment. The LZ experiment is the combined effort of the LUX and ZEPLIN-III collaborations. This chapter will discuss the various components that make up the LZ dark matter detector, the data handling infrastructure, and the detector simulations used to understand the data. A review of the calibration and background sources is given followed by a WIMP search result.

The following description of the LZ experiment uses detector parameters and results from "Science Run 1" (SR1). Data for SR1 was taken from December 23rd, 2021 until May 11th, 2022. The LZ detector, seen in figure 3.1, is located 4850 ft underground in the Davis Cavern at the Sanford Underground Research Facility (SURF) in Lead, South Dakota. The detector includes a TPC containing 7 tonnes of liquid xenon nested within two active veto detectors. A WIMP signal would interact once, vetoes exclude events with multiple interactions across the independent detectors. The "skin detector" is an instrumented liquid xenon space between the TPC and the inner cryostat walls. The cryostat is 2 m tall and has a 1.85 m diameter containing 10 tonnes of liquid xenon. The "Outer Detector" (OD) surrounds the cryostat and includes 10 acrylic tanks filled with 17 tonnes of

gadolinium-loaded (0.1% by mass) liquid scintillator. The acrylic vessels and cryostat sit within a water tank filled with 228 tonnes of ultra-pure water. Conduits are used to deliver calibration sources, liquid xenon, and electrical connections between the inner cryostat and external support infrastructure.

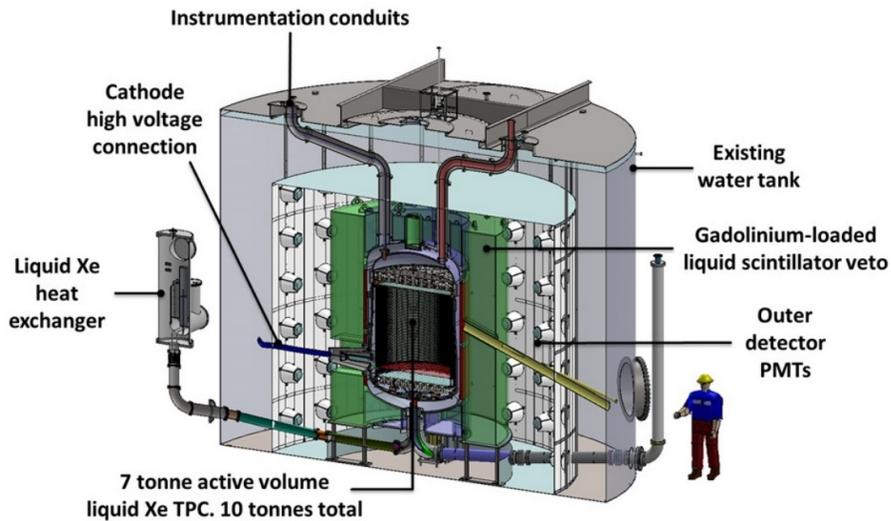


Figure 3.1: A schematic of the LZ TPC and active veto detectors within the water tank. Starting from the outside, a cylindrical wall of PMTs view the acrylic vessels (shown in green) filled with gadolinium-loaded liquid scintillator. The inner cryostat is filled with liquid xenon. The skin detector has rings of PMTs monitoring the space between the TPC and cryostat walls. At the center is the TPC where two arrays of PMTs point towards the bulk xenon. Various conduits support infrastructure and the delivery of calibration sources. Taken from the LZ Technical Design Report [17].

3.1 Active Vetoes

3.1.1 The Outer Detector

The Outer Detector uses gadolinium-loaded liquid scintillator to veto neutrons that have interacted within the TPC. A neutron with energy between 0.5 and 5 MeV that only scatters once in the TPC will be virtually indistinguishable from a potential WIMP signal. After the scatter, the neutron will exit the detector and be captured on a gadolinium nucleus. The

time scale of this capture process is approximately $30 \mu\text{s}$. Upon capture, a gamma ray cascade of 3 or 4 gammas will emit a total energy of 8 MeV. Some of these gammas will interact in the liquid scintillator and others will penetrate the cryostat and interact in the liquid skin. These gamma splashes, detected in either the OD or skin detector, are used to veto coincident interactions within the TPC.

An array of 120 eight-inch Hamamatsu R5912 PMTs monitor the scintillation light produced in the OD acrylic tanks and Cherenkov radiation produced in the water. Cylindrical scaffolding holds the PMTs 80 cm away from the scintillator tanks. The ultra-pure water attenuates any background gamma rays emitted by the PMT housings. Both the cryostat and the PMT scaffolding are covered in $600 \mu\text{m}$ multilayer Tyvek® with a reflectivity greater than 95% for the 390 to 440 nm light emitted by the liquid scintillator [17].

Signals from the OD are used to perform an OD prompt and delayed veto cut. The prompt veto excludes gammas that scatter into, or out of the TPC. An event is vetoed if a TPC S1 occurs within 300 ns of 6 OD PMTs triggering in coincidence. The delayed OD veto excludes neutron exiting the TPC after producing an S1. If the OD detects an event with an energy greater than 200 keV up to $1200 \mu\text{s}$ after the S1, the event is vetoed. The neutron tagging efficiency is $89 \pm 3\%$ calibrated using an AmLi source. The false veto rate is 5% due to accidental activity in the OD during the coincidence window [94] [95] [96] [24].

3.1.2 The Skin Detector

Between the cryostat and TPC walls there is a radial space that allows for instrumentation housing and clearance during assembly. This cylindrical skin is 4 cm thick at the top and 8 cm near the bottom. The cryostat wall is covered in PTFE (polytetrafluoroethylene or Teflon®) and the TPC wall is made of stacked PTFE rings. PTFE is over 99% reflective

to 178 nm scintillation light. The liquid xenon that fills the skin region is monitored by a total of 131 PMTs. Just below the liquid surface, a ring of 93 one-inch Hamamatsu R8520 PMTs view down into the skin region. At the bottom, a ring of 20 two-inch R8778 PMTs point up towards the skin region. A final group of 18 R8778 PMTs monitor the dome region beneath the bottom PMT array, see figure 3.2.

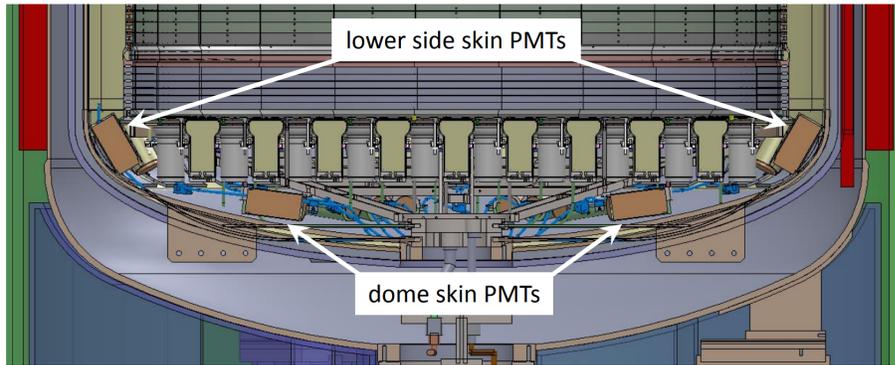


Figure 3.2: A rendering of the lower cross section of the TPC and cryostat. The skin region is located between the cryostat and TPC walls. Lower side skin PMTs are aimed up into the skin region. Dome skin PMTs monitor light production near the bottom PMT array. Top skin PMTs are not depicted here. Taken from reference [17].

Charge is not collected in the skin, but scintillation signals are used to veto coincident TPC events. Signals from the skin detector are used to implement prompt and delayed veto cuts on TPC signals. The prompt veto flags events that occur in the skin that could potentially leak light into the TPC region or gammas that scatter between the TPC and skin. This veto excludes events that produce a signal in the skin that occur within 500 ns of a TPC S1 and requires at least 3 skin PMTs to trigger, measuring a total of pulse area of at least 2.5 phd. The delayed skin veto rejects neutron events by searching for gamma rays that penetrate the cryostat following neutron captures in the OD. The resulting gamma splash is vetoed if at least 56 PMTs measure a total pulse area of at least 50 phd. The delayed veto window extends up to 1200 μ s after an S1 is detected in the TPC. The combined skin vetoes reduce the detector live time by less than a 0.05% [97] [17].

3.2 The Time Projection Chamber

At the heart of the nested veto detectors is the TPC. It is cylindrical in shape with a 1.45 meter diameter and a 1.69 meter height. The ends are capped with arrays of Hamamatsu R11410-22 PMTs, 253 on the top and 241 on the bottom. Stacked electric field shaping rings embedded within PTFE make up the TPC wall. These rings are chained together using 2 G Ω high voltage resistors.

Four stainless steel woven wire mesh grids, held at set voltages, produce an electric field graded by the voltage steps between the field shaping rings. The field regions and grid voltages within the TPC are shown in figure 3.3. The voltages applied to the “reverse field region” between the bottom and cathode grids reduces the electric field strength at the bottom of the detector. No charge is collected from events in this region but scintillation light is still produced and contributes to detector backgrounds. The “drift region” has an electric field of 193 V/cm and extends from the cathode grid to the gate grid. This 145.6 cm region contains the main detector and fiducial volume from which scintillation and charge signals are collected. Electrostatic simulations found electric field variations of 4% over the 5.5 tonne fiducial volume [23].

The gate grid sits just below the liquid surface and the anode just above it. Together they formed the “extraction region”. The 7.3 kV/cm extraction field pulls electrons, produced during ionization, into the gas phase from the liquid surface. The field in this region determines the size of the S2 signals due to the gas gain from proportional scintillation, as described in section 2.5.2. The wire pitch and wire diameters are selected to maximize light collection while maintaining the electric field shape and supporting the electrostatic forces. Optical simulations of each grid found a transparency greater than 90% [23] [24]. The maximum drift time, 951 μ s, is the amount of time it takes an electron to drift from the cathode to the top.

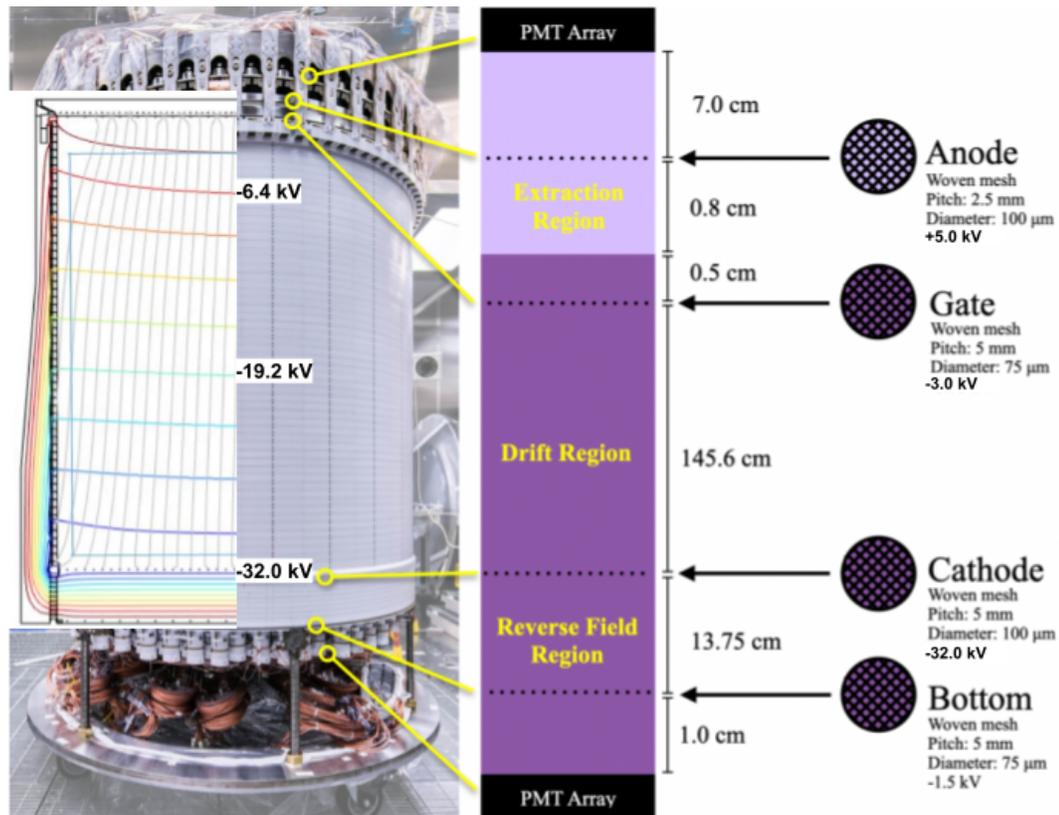


Figure 3.3: An external image of the LZ TPC in the above ground surface area lab following construction. The white PTFE field rings are stacked forming the vertical wall of the detector. The location of the different field regions and grids within the TPC are shown. The dark purple are regions submerged in liquid xenon and light purple regions are in the gaseous phase [23]. On the left an electric field map in the liquid region with updated voltage labels reflect the SR1 run parameters. The rainbow horizontal lines are equipotentials from 0 to a 32 kV in steps of 3.2 kV. Taken from reference [17].

The observed electron lifetime increased throughout SR1 from 5000 μs to 8000 μs due to the inline purification system. The electron lifetime is the amount of time a free electron can remain in the detector before it binds to an electronegative impurity and is lost. The inline purification system includes a hot zirconium getter. Gas flow is driven by two metal seal Fluitron compressors at 500 standard liters per minute (slpm) with a xenon turnover rate of 3.3 tonnes per day.

The temperature and pressure within the TPC is stable to within 0.2%, at 173.1 K and

1.791 bara (bar absolute). The liquid level is maintained by weir overflows to within 10 μm , as measured by precision capacitance sensors. Liquid xenon is transferred outside of the water tank to a heat exchanger and evaporated. The gas is fed into the purification system and supported auxiliary systems used for calibration source injection, xenon gas purity monitoring, and inline radon reduction. The purified liquid is then condensed and fed back into the bottom of the TPC [24].

3.3 Data Handling and Software

Events with energy deposition above 1 keV_{ee} produce $S1$ and $S2$ signals within the 7 tonne active region of the TPC. The data acquisition system (DAQ) uses a digital filter sensitive to $S2$ pulses to trigger the processing chain. All of the PMT data from 2 ms prior to and 2.5 ms after the trigger are saved to disk [24].

The LZap (LZ Analysis Package) software package uses the data from the detector electronics to identify pulses and reconstruct each event. The pulse shape and PMT hit patterns are used to identify pulses as either $S1$ s or $S2$ s. $S1$ pulses have a sharp rise time and have smaller pulse areas than their $S2$ counterpart. $S2$ signals have a slower rise time and a wider Gaussian profile.

Once the pulses are identified, the event is classified as either a single scatter, multi-scatter, pileup, or other event. Single scatter events have one $S1$ that precedes one $S2$. Single scatter events are where a WIMP event would be found. Multi-scatter events have a single $S1$ followed by multiple $S2$ pulses. This is usually due to a gamma Compton scattering within the detector. This event type produces a single $S1$ as all the scatters occur at nearly the same instant, but the ionization cluster occur at different z positions and produce distinct $S2$ signals. Pile-up scatters contain more than one $S1$, at least one of which is happens before one or more $S2$ s. These are usually separate events that occurred within

the same event window. An exception to this case are ^{83m}Kr events and ^{85}Kr excited state decays. A ^{83m}Kr event has two $S1$ signals prior to a single $S2$. These events are used for mapping the electric field within the detector and are identified as a special case of single scatters. “Other” scatters are a miscellaneous category all for any remaining events that fail to be classified. These events can have an $S2$ that precedes $S1$, an $S1$ only, an $S2$ only, or multiple $S1$ s and multiple $S2$ s.

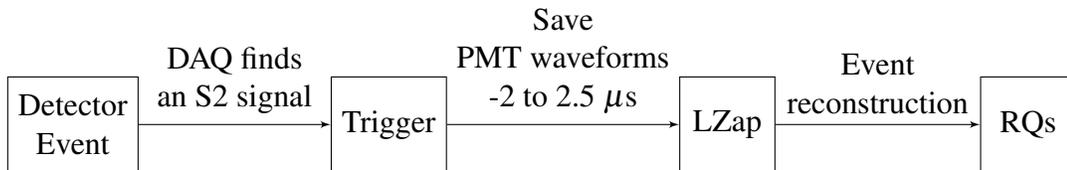


Figure 3.4: Block diagram for an $S2$ trigger event. The PMT information is stored in a buffer and actively filtered for an $S2$ trigger. When triggered, the DAQ saves the event window to disk. LZap searches the PMT waveforms for pulses and reconstructs the event. Finally an RQ list is produced and stored in a ROOT file.

Once LZap has classified the event, a reduced quantities (RQs) list is produced and stored in a ROOT file. The RQs are populated in the appropriate branch pertaining to the event classification. All event types have RQs according to raw event parameters that identify a given pulses type ($S1$ or $S2$), integrated pulse area (phd), time from trigger (ns), and xy position for $S2$ pulses. If the event is identified as either a single or multi scatter, LZap will produce additional corrected RQs and store them in the appropriate branch of the ROOT file. The corrections are determined from maps based on the event location. These maps are measured with calibration sources such as ^{83m}Kr (see section 3.4). These corrections cannot be applied to pile-up or other event types, as the event correlation between $S1$ s and $S2$ s are ambiguous. LZap also processes information from the OD and skin detectors and populates detector specific branches with appropriate RQs.

Once processed by LZap, this information (including raw waveforms and individual PMT channels) is made available on the LZ database. The “LZ offline event viewer” can

load individual events in a convenient graphical interface. From here users can inspect pulses from specific PMTs in the TPC, OD, and skin detector. The events can also be loaded from the LZ database and analyzed in depth through ALPACA (Analysis LZ PACKage). This analysis package implements standardized cuts and modules that can be shared among researchers to streamline analysis.

3.3.1 Simulations

To simulate events within the the detector LZ uses BACCARAT (Basically, A Component-Centric Analog Response to AnyThing). The BACCARAT simulation software contains the LZ detector geometry and uses GEANT4 to track particles and the energy they deposit in the detector. The energy deposited is passed to a NEST (Noble Element Simulation Technique) module within BACCARAT. NEST uses a collection of semi-empirical models to determine the physics response of the noble target medium. This converts the energy deposited into light and charge yields. NEST can either pass these yields back to BACCARAT for a “full” simulation or use an averaged detector response to convert the yields into S1 and S2 signals for a “fast” simulation.

Full NEST tells BACCARAT how much light and charge are produced in a recoil event. BACCARAT then tracks the individual electrons and photons generated by each recoil. The electron paths follow the simulated field map and are extracted into the gas phase where the electroluminescence process is modeled. All of the photon paths are modeled throughout the detector geometry and the PMT hits are then recorded. The “detector electronics response” (DER), simulates the PMT efficiencies and the amplifier response from the PMT hits. The DER then provides individual waveforms for each PMT which can be processed by LZap and studied in the event viewer. Much like real events, the full chain simulation replicates the entire detector response including software processing, but comes at the cost

of computing power and time. For SR1, well-tuned full-chain simulations were reserved for backgrounds in the WIMP search region of interest. In chapter 6 the full chain simulation is used to produce simulated excited state ^{85}Kr decays in the event viewer, inform the veto detector response to their gamma emission, and evaluate how $S1$ light disperses between the top and bottom PMT arrays.

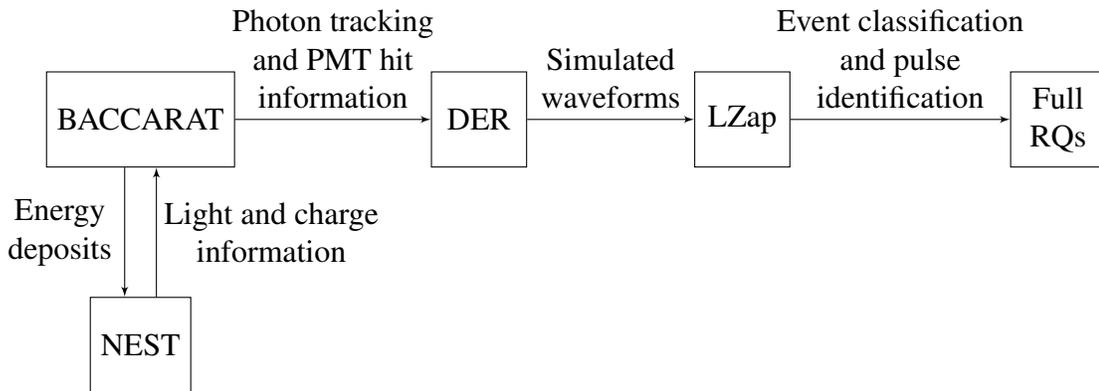


Figure 3.5: Block diagram for the full chain simulation.

To reduce the computational load, fast NEST directly translates energy deposits from BACCARAT into $S1$ and $S2$ signals based on a parameterized detector response. The fast simulation chain skips photon and electron tracking. This means that channel-specific PMT information is lost. The benefit is that fast NEST has a quick turnaround, with large statistics tuned to match the detector response at a wide range of energies. It compliments the full chain simulations [25].

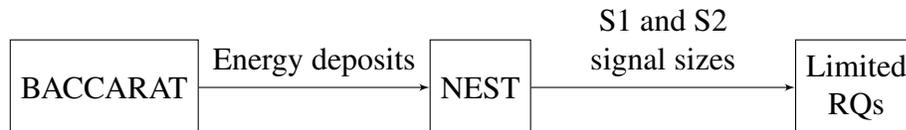


Figure 3.6: Block diagram for the fast chain simulation.

3.4 Calibration

LZ uses radioactive sources to calibrate the detector response. Nuclear recoils are calibrated using a deuterium-deuterium (DD) generator that emits monoenergetic 2.45 MeV neutrons and AmLi neutron sources deployed outside of the cryostat walls. Electron recoils are calibrated using electrons and gammas emitted from ^{83m}Kr (41.5 keV) and ^{131m}Xe (164 keV) before and during the WIMP search. Tritiated methane (CH_3T , 18.6 keV Q value) is a beta source crucial to characterizing low energy electron recoils in the WIMP search region of interest. Dispersed xenon sources were used to correct the $S1$ ($S1_c$) and $S2$ ($S2_c$) signals, as the position of these sources is uniform throughout the detector. An average $S1$ correction of 9% accounts for variations in light collection efficiencies at different locations. An average $S2$ correction of 11% accounts for variations in light collection efficiencies in the xy plane due to non-operable PMTs and extraction region electric field variations. The average size of $S2$ corrections in z is 7% due to electron attachment to impurities. The corrected parameters result in uniform response across the TPC to within 3%.

The detector response is tuned in NEST by fitting the tritium beta spectrum and monoenergetic electron and gamma peaks from ^{83m}Kr , ^{129m}Xe , and ^{131m}Xe . Calibration data finds a photon detection efficiency of $g_1 = 0.114 \pm 0.002$ phd/photon and an ionization gain of $g_2 = 47.1 \pm 1.1$ phd/electron. The NEST electron recoil model is propagated to produce a nuclear recoil model and found to be in good agreement with the DD calibration data. The bands in $\log(S2_c)$ have matching means to within 1% and widths to within 4% [24]. The two calibration recoil bands compared to the NEST fits are shown in figure 3.7.

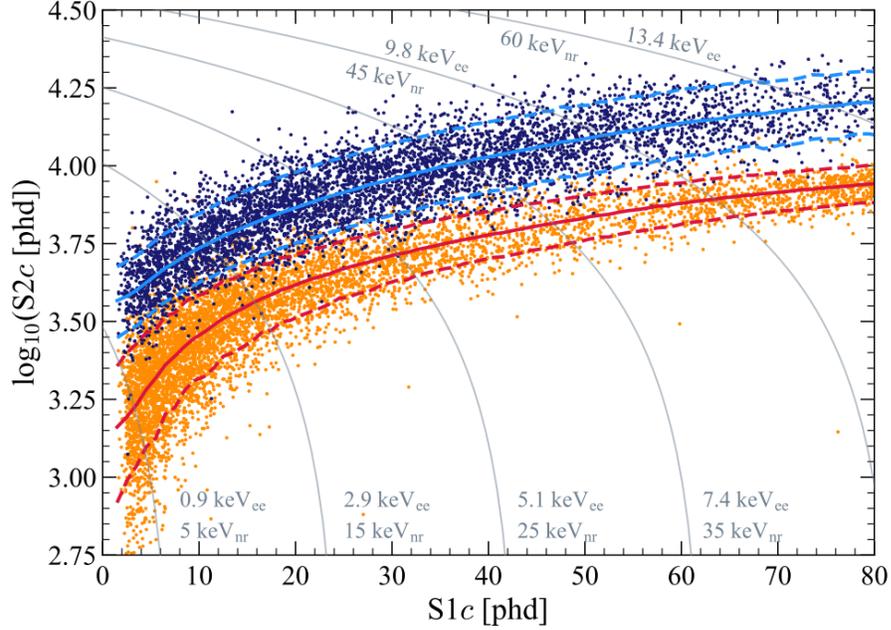


Figure 3.7: Calibration events in $S1$ vs $\log(S2_c)$ space from the LZ detector. The blue points are electron recoils from a tritium source. The orange events are neutron events from a deuterium-deuterium (DD) generator. Grey curves denote lines of constant energy and take different values for electron (keV_{ee}) and nuclear (keV_{nr}) recoils. The solid red (blue) line is the median of the simulated nuclear (electron) recoil band. The dotted lines indicate the 10% and 90% quantiles. Taken from reference [24].

3.5 Backgrounds

Detector backgrounds are determined by the measured and anticipated rates from five classes of background sources: detector components, surface contaminants, dispersed contaminants, external, and physics. The identification and reduction of each of these background sources were essential in the planning and construction of LZ. The majority of SR1 background fits were conducted on single scatter events to inform WIMP search backgrounds. The identified sources were simulated based on their expected rate and included in a global background model. WIMP search cuts were applied to the simulated backgrounds and the surviving events from the global background model informed the expected backgrounds for SR1. A total background fit is then made to the SR1 data. WIMP search

cuts were applied to the low energy portion of the total background fit and determined the final WIMP search backgrounds.

3.5.1 Background Sources

Detector components like the cryostat, grids, and PMTs contain the naturally-occurring radioactive materials ^{40}K , ^{137}Cs , ^{60}Co , ^{238}U , ^{235}U , and ^{232}Th . These sources and their progeny emit a combination of gammas, alphas, and betas. Component screening programs were carried out to quantify and reduce contributions from these sources. Additionally liquid xenon shielding minimizes backgrounds from these sources within the detector fiducial volume.

Surface contaminants like dust and ^{222}Rn daughters can accumulate on detector surfaces exposed to air during detector assembly. Of particular concern are neutrons produced from the (α, n) process and ^{210}Pb nuclear recoils that are poorly reconstructed due to their proximity to the detector wall. Both of these sources contribute to the nuclear recoil backgrounds. A rigorous cleanliness management program was implemented during assembly to minimize the exposure of dust, and construction was completed in a reduced-radon clean room. A radon assay program was also completed to identify and account the radon emanation from detector components.

Dispersed xenon contaminants are radioisotopes dispersed throughout the liquid itself and cannot be mitigated through xenon's self shielding properties. Radon emanation from materials is the largest contributor to the total LZ background. This is due to naked beta emissions that lack accompanying gammas from decays of ^{214}Pb (^{212}Pb) in the ^{222}Rn (^{220}Rn) sub-chain. Natural xenon also includes trace levels of dispersed ^{85}Kr and ^{39}Ar . Both of these beta emitters were reduced to an acceptable level through a devoted chromatography system, discussed in chapter 5.

External backgrounds are due to laboratory sources and cosmogenic interactions. Neutrons produced from muon cascades can cause nuclear recoil backgrounds. At SURF shielding from the rock overburden, equivalent to 4300 meters of water, reduces muon flux by 6 orders of magnitude. Any surviving muons are tagged via Cherenkov emission and scintillation in the OD [17]. The gamma flux from the walls in the Davis Cavern was measured using NaI detectors [98]. Neutrons from outside of the water tank are attenuated by more than 6 orders of magnitude, resulting in a negligible contribution to the LZ backgrounds [99]. Cosmogenic activation of xenon produces ^{127}Xe which has a half-life of 36.4 days and is present for the duration of SR1. Additionally DD, calibrations activate ^{131m}Xe , ^{129m}Xe , and ^{133}Xe with half-lives of 11.8, 8.9, and 5.2 days respectively.

Physics backgrounds are interesting processes in their own right but contribute background events to the WIMP search region of interest. In LZ there is a sufficient quantity of xenon to observe the two-neutrino double beta decay ($2\nu\beta\beta$) of ^{136}Xe and ^{134}Xe . Given a natural abundance of 9%, a half life of 2.2×10^{21} years, and a Q-value of 2,456 keV, LZ should observe approximately 3 million events from ^{136}Xe over a 1000 day exposure. LZ should also be able to measure the full spectrum of the ^{134}Xe $2\nu\beta\beta$ decay down to 1 keV [17]. Astrophysical neutrinos, *hep* solar neutrinos, and ^8B neutrinos should contribute to nuclear recoils while solar *pp* neutrinos should contribute to electron recoils. LZ could also potentially observe a burst of neutrinos from a nearby supernova [25] [100].

3.5.2 Background

To model the background, studies above the WIMP search energy range were conducted from 80-2700 keV_{ee}. Sources in this range were identified and characterized according to their position within the detector. A fit, accounting for all identified sources, is made in this higher energy range and extrapolated to lower energies where the WIMP search is carried

out. Alphas from the radon decay chains with energies greater than 3 MeV are readily identified from their large S1 signals. Within the fiducial volume, a fit for all single scatter alpha events is completed. Decay rates are identified for ^{222}Rn , ^{218}Po , ^{216}Po , ^{214}Po , and ^{212}Po . A spatial distribution of events associated with radon daughters is observed with higher rates near the cathode and walls.

From 200-450 keV, cosmogenically-activated xenon contributes gamma backgrounds. The rock overburden makes the equilibrium rate for these isotopes negligible, but because the start of SR1 was only 4 months after the xenon was brought underground, cosmogenic activation peaks were expected. In addition, neutron calibrations were conducted during SR1 and contributed to the xenon activation background. Two of the activated isotopes undergo electron capture and emit auger electrons or X-rays from electron cascades. The different combinations of these processes results in a variety of energy peaks with different relative intensities in this range. The xenon activation isotopes and corresponding decays are outlined in the following table 3.1 and figure 3.8. Dispersed homogeneous background rates were measured for the innermost 1 tonne cylindrical volume of LZ and extended to the remaining detector volumes.

Isotope	Decay Process	Half-life	Energy (keV)
^{133}Xe	β^-	5.2 (d)	346
^{131m}Xe	γ^*	11.8 (d)	164
^{129m}Xe	$\gamma\gamma$	8.88 (d)	196.6 + 39.6
^{127}Xe	EC + (K,L,M)	36.4 (d)	(203, 375) + (33.2, 5.2, 1.1)
^{125}Xe	EC + (K,L)	16.9 (h)	(55, 188, 243) + (33.2, 5.2)

Table 3.1: Table of activated xenon isotopes. Note that for the case of ^{131m}Xe the γ^* energy is often carried away by electron(s) and/or X-ray(s) and not a single gamma. ^{127}Xe has two characteristic gammas with relative intensities 53% and 47% as well as (K,L,M) electron cascade that produces X-rays with corresponding probabilities (83.37, 13.09, 2.88)%. Similarly ^{125}Xe has three potential characteristic gammas with relative intensity (6.8, 53.8, 30)% and X-rays from electron cascades. Taken from reference [30].

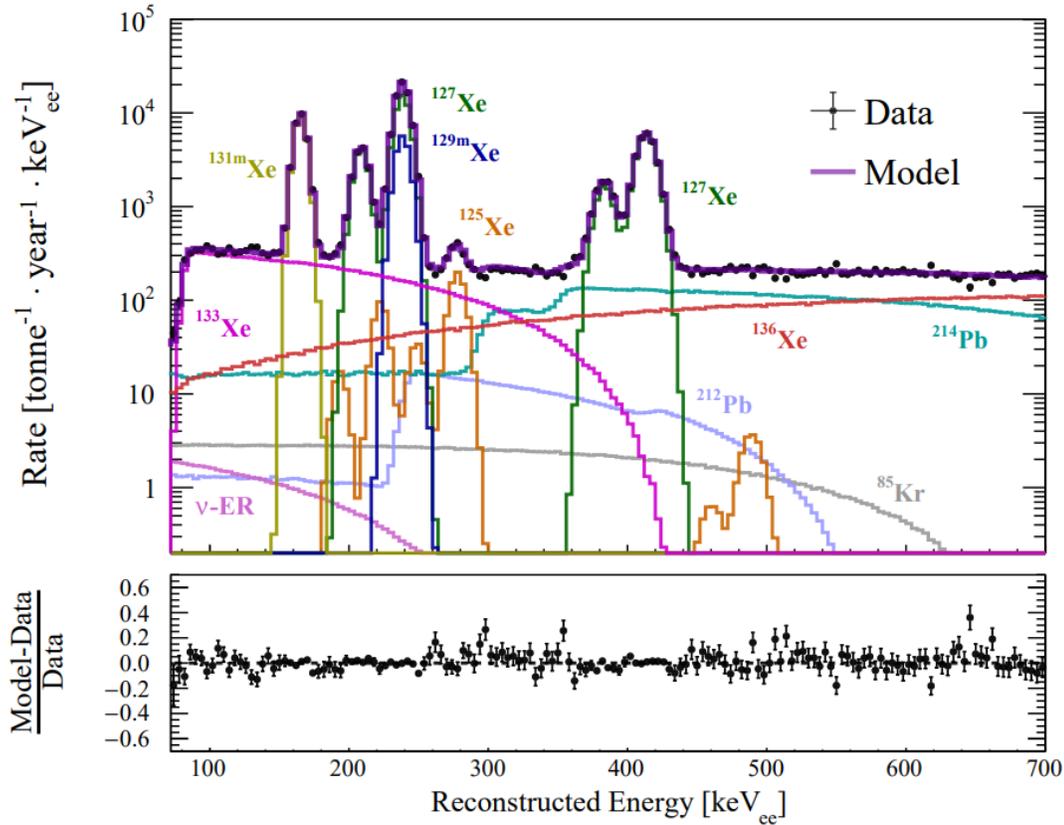


Figure 3.8: Background rate as a function of energy in the innermost 1 tonne xenon volume following a DD calibration. Gamma and gamma+X-ray features sit atop a combined beta spectrum. Note the four peaks from the ^{127}Xe decays and electron cascades. Similarly the six peaks are noted from the ^{125}Xe decays and cascades. The few days that followed the DD calibration were removed from SR1 to allow the short lived isotopes, ^{133}Xe and ^{136}Xe , to decay to sub dominant rates. Taken from reference [25].

The beta background region is defined between 80-700 keV and measured in the innermost 1 tonne volume. Extra care was given to the spatial distributions of ^{214}Pb and ^{212}Pb as they were known to be inhomogeneously distributed from the spatial distribution of alpha events from the radon sub-chain. Event rate contributions from neutrino recoils, ^{136}Xe , ^{212}Pb , ^{214}Pb , and ^{85}Kr are not impacted by DD calibrations. Backgrounds following a DD calibration are shown in figure 3.8.

A high energy gamma region is defined between 700-3000 keV. Contributions to this region come from cavern rock gammas and detector components. Again ^{238}U , ^{232}Th , and their progeny as well as ^{40}K and ^{60}Co are most prevalent. Gamma contributions were determined through a dedicated analysis during the LZ detector's early commissioning. At this time the TPC was filled with gaseous xenon while the OD and water tank were empty resulting in a cavern-rock gamma-dominant signal. The internal rate was determined using SR1 data localized to specific volume subgroups. The cumulative source activity for each subgroup was fit and compared to be in reasonable agreement with screening estimates, although ^{238}U rates were found to be a factor of 2 higher than projected. The total background fit against the SR1 data is shown in figure 3.9 [25].

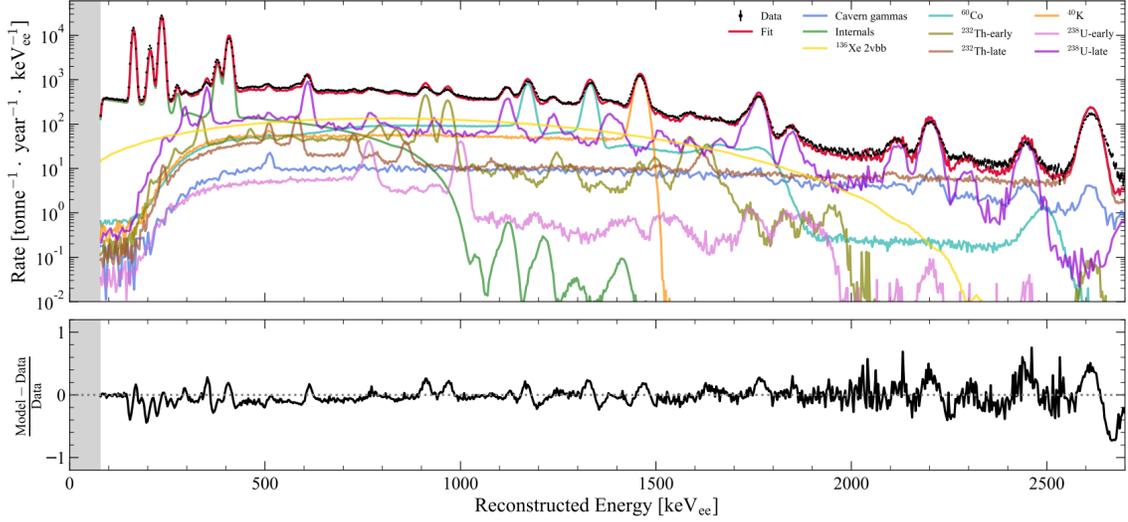


Figure 3.9: Total fitted detector background spectra in the SR1 fiducial volume. The fit in red to the data in black is consistent with an average residual of approximately zero. Fluctuations at certain peaks arise from imperfect energy resolution matching between data and simulations. Energies below 80 keV_{ee} have been obscured as analysis of ^{124}Xe two-neutrino double electron capture is underway. Taken from reference [25].

3.5.3 WIMP Search Selection Cuts and Backgrounds

The WIMP search is performed on single scatter events in a region of interest $3 < S1c < 80$ phd. A three-fold $S1$ PMT signal coincidence and $S2_{raw} > 600$ phd (uncorrected) are required. This corresponds to a WIMP interaction with an energy deposit greater than about 5 keV.

Anti-coincidence vetoes cuts were applied, as described in sections 3.1.1 and 3.1.2. Pulse cuts were used to reduce pathologies that result in accidental coincidence events. Accidental coincidences occur when an unrelated $S1$ and $S2$ are paired during reconstruction and falsely identified as a single scatter event. A live-time veto was developed to exclude periods of high event rates within the detector and discard periods of time where the probability of accidentals increased. This includes intervals of afterglow from high energy interactions within the detector. Accidental backgrounds were modeled from events

with unphysical drift times. These events have drift times that exceed the measured maximum drift time and represent an accidental population where an S1 has been paired to an S2 that a single event could not produce.

A fiducial volume cut was implemented to reduce surface backgrounds, poor detector edge reconstruction, and external backgrounds. The fiducial volume cut accepts events with drift times between 86 and 936.5 μs and a radial offset at least 4.0 cm from the TPC wall. Additional cut outs were implemented at the top and bottom of the detector. Events with drift times less than 200 μs were required to occur with radial offset of at least 5.2 cm from the wall. At the bottom of the detector events with drift times greater than 800 μs were required to have a radial offset of at least 5 cm from the TPC wall.

Ion mobility and xenon flow models were used to estimate the position distributions of ^{222}Rn and ^{222}Rn daughters. Alpha events associated with the radon decay chain were tagged and tracked to model these distributions. The decay rate was informed by ex-situ radon measurements. These simulations informed the expected beta background rates and distribution from ^{214}Pb and ^{212}Pb .

Direct measurements of the trace impurities ^{nat}Kr and ^{nat}Ar were made during the krypton removal campaign via a xenon gas purity monitor, as discussed in chapter 5. These measurements were used to determine beta backgrounds from ^{85}Kr and ^{39}Ar . An SR1-data driven search for the delayed coincidence gamma-beta decay of ^{85}Kr is described in chapter 6.

The expected contributions from detector and cavern gammas is evaluated from rates measured during construction and initial commissioning. A simulation using these rates was used to infer background contributions within the fiducial volume following all SR1 WIMP search cuts and veto rejections.

Cosmogenic spallation of xenon at the earth's surface results in ^{37}Ar production which undergoes electron capture with a half life of 35 days. The resulting 2.8 keV X-rays fall

within the WIMP search energy window. The rate of ^{37}Ar was estimated by calculating the exposure of xenon to cosmic rays prior to being brought underground and then correcting for the decay time [101]. This background was implemented as a flat distribution ranging from 0 to 288 events.

Expected backgrounds from ^{124}Xe 2ν double electron capture, ^{136}Xe $2\nu\beta\beta$, and solar neutrinos were estimated from half lives measured by XENON1T [84], EXO200 [85], and from calculated neutrino scattering rates.

The expected rate from ^{127}Xe was determined from the simulated activation due to cosmogenic and neutron activation. The semi-naked ^{127}Xe electron capture decay was simulated. In this case, a gamma ray leaves the detector and the low energy X-rays from the L or M shell cascades contribute to the WIMP search backgrounds. Other activated xenon peaks were not considered in the expected fit as they have short half-lives and SR1 excludes data following DD neutron calibrations.

The expected neutron backgrounds were derived using auxiliary fits to events tagged in the OD veto cut, but otherwise pass all WIMP search cuts. These events were then compared to AmLi and DD neutron data to determine the survival fraction of single scatter events in the WIMP search region of interest and the expected neutron rate.

A final fit was completed for all events that survived the final WIMP search SR1 cuts and was consistent with the expected event rates. This is shown in figure 3.10 and table 3.11. A final WIMP search background rate was found to be $(6.3 \pm 0.5) \times 10^{-5}$ events/ $\text{keV}_{ee}/\text{kg}/\text{day}$. This amounts to 333 ± 17 total events over the 330 ± 12 tonne-day exposure.

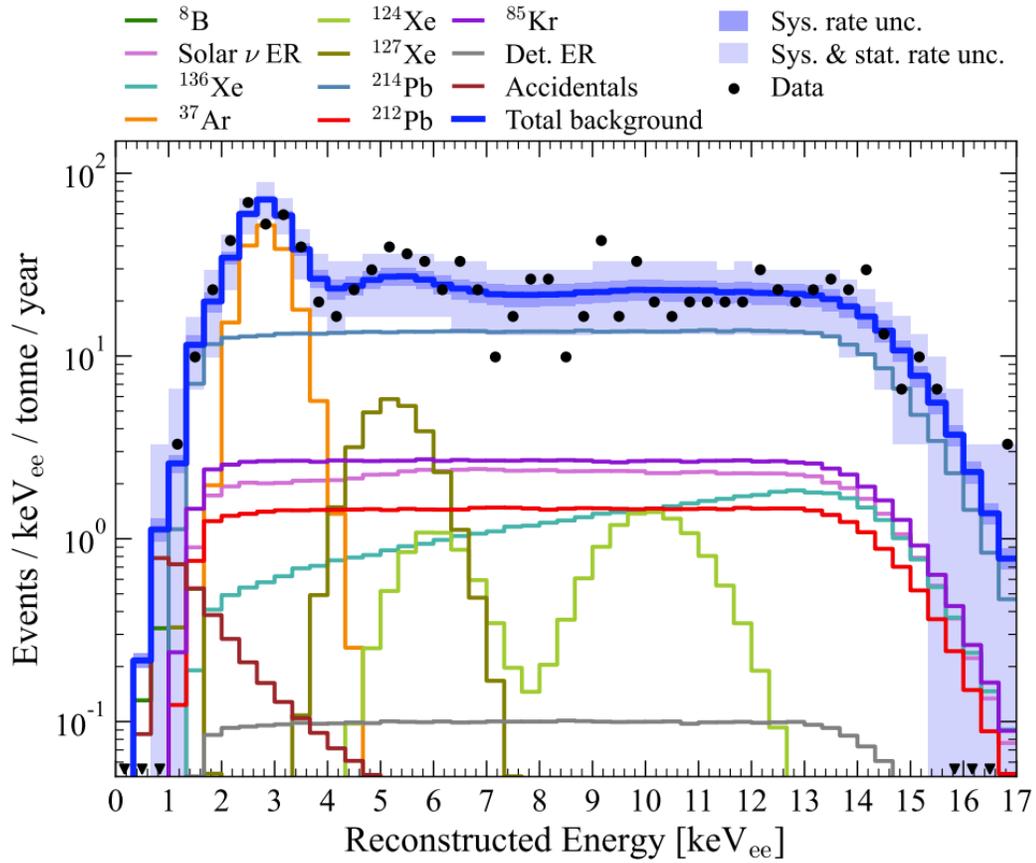


Figure 3.10: LZ background model before fitting to SR1 data (except for the ^{37}Ar component). The total background in the WIMP search region of interest is shown in blue with the WIMP search cut SR1 data in black. This represents an event rate of $(6.3 \pm 0.5) \times 10^{-5}$ events/keV $_{ee}$ /kg/day. Taken from reference [25].

Source	Expected Events	Fit Result
^{214}Pb	164 ± 35	-
^{212}Pb	18 ± 5	-
^{85}Kr	32 ± 5	-
Det. ER	1.4 ± 0.4	-
β decays + Det. ER	215 ± 36	222 ± 16
ν ER	27.1 ± 1.6	27.2 ± 1.6
^{127}Xe	9.2 ± 0.8	9.3 ± 0.8
^{124}Xe	5.0 ± 1.4	5.2 ± 1.4
^{136}Xe	15.1 ± 2.4	15.2 ± 2.4
^8B CE ν NS	0.14 ± 0.01	0.15 ± 0.01
Accidentals	1.2 ± 0.3	1.2 ± 0.3
Subtotal	273 ± 36	280 ± 16
^{37}Ar	[0, 288]	$52.5^{+9.6}_{-8.9}$
Detector neutrons	$0.0^{+0.2}$	$0.0^{+0.2}$
30 GeV/ c^2 WIMP	-	$0.0^{+0.6}$
Total	-	333 ± 17

Figure 3.11: The LZ background table with sources in the left column, expected events in the middle column, and fit results in the right column. While expected events were explicitly estimated for the various β decays, the fit uses a flat energy distribution to model all of the ^{214}Pb , ^{212}Pb , ^{85}Kr , and detector electron recoils. Taken from reference [25].

3.6 WIMP Search Results

During the 100.377 days of SR1 exposure over 110 million events were recorded. The full calendar time is 139 calendar days due to detector maintenance and calibrations. After accounting for DAQ downtime and hold offs due to anomalous trigger rates or high activity, the final WIMP search exposure was 60 ± 0.1 days containing 60 million events. Following the analysis cuts of single scatter events within the region of interest, only 31,000 events remain and are shown below in figure 3.12.

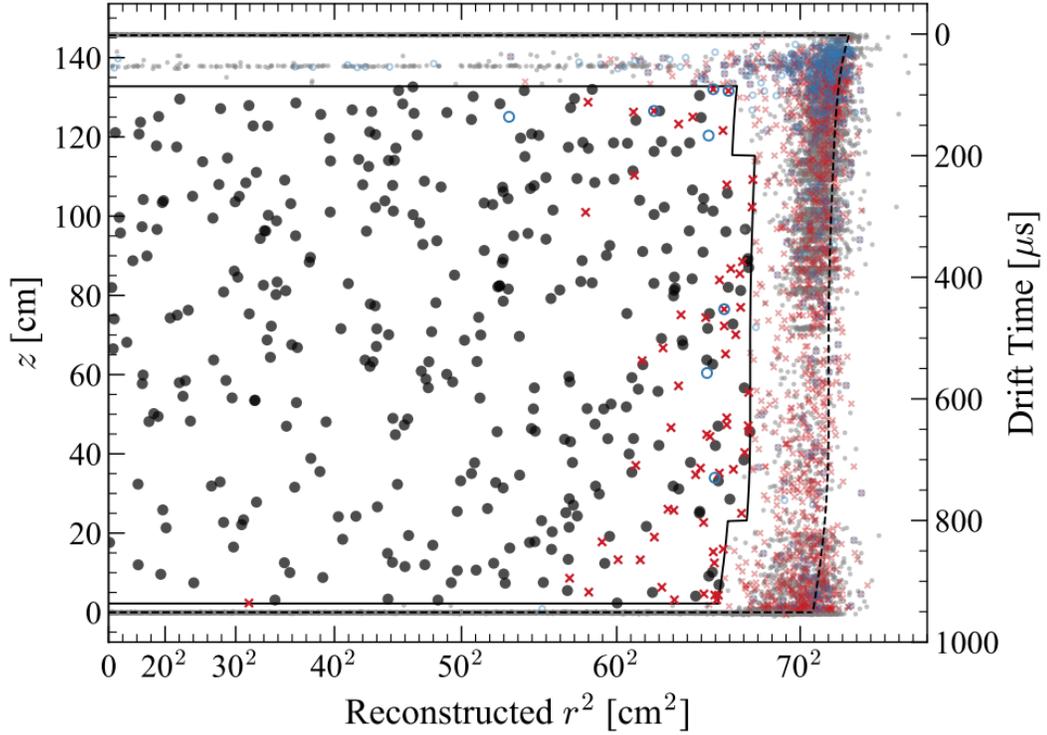


Figure 3.12: Spatial distribution of WIMP search events during SR1 prior to the fiducial volume cut, OD veto, and skin veto. The detector volume is enclosed by a block dotted line and the fiducial volume is shown with a black solid line. Black dots within the fiducial volume make up the final 335 events within the WIMP search data set. Grey circles satisfy all cuts but occur outside of the fiducial volume. Red crosses are events vetoed by the liquid skin and blue circles are events vetoed by the OD. Taken from reference [24].

The 335 events remaining after the fiducial cut, OD veto, and skin veto are shown in the $\log_{10}(S2_c) - S1_c$ space in figure 3.14. The cut efficiency for WIMPs was evaluated using composite nuclear recoil like waveforms generated to mimic an expected WIMP signal. A software package was used scan tritium data for $S1$ pulses and stitched them to small $S2$ pulses from AmLi or lower energy tritium events. The full cut efficiency is shown in figure 3.13.

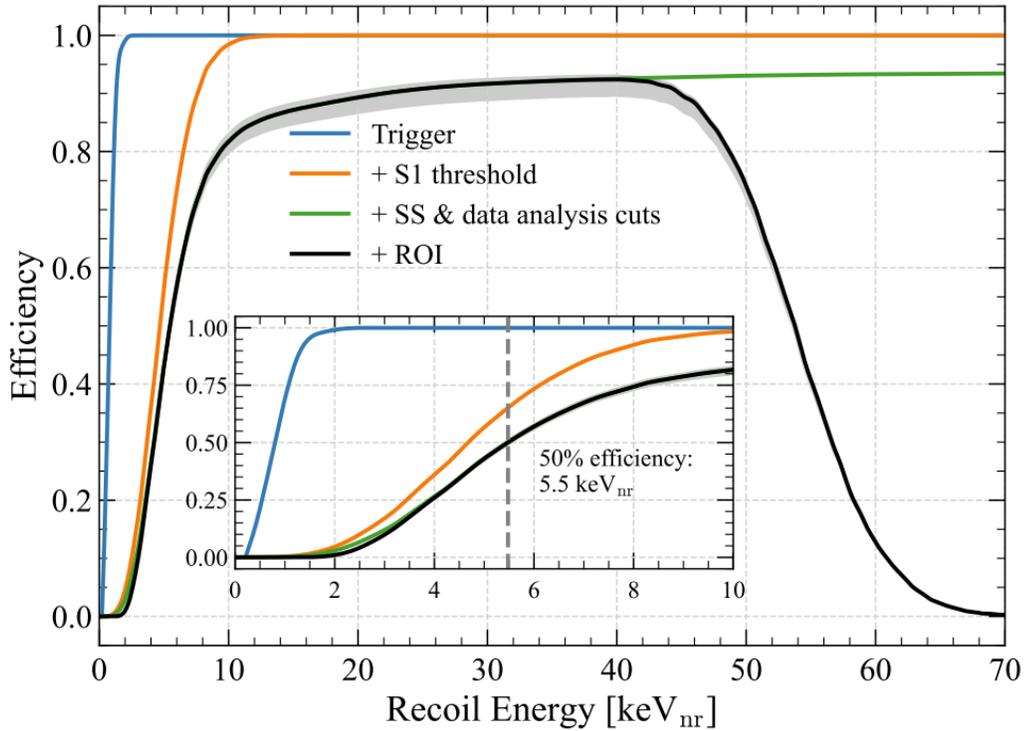


Figure 3.13: The WIMP signal efficiency as a function of nuclear recoil energy for single scatter reconstructions in the search region of interest. The error assessed using composite nuclear recoil AmLi-tritium events. A 50% efficiency is achieved at 5.5 keV_{nr} and marked with a grey dotted line. Taken from reference [24].

The final result is consistent with a background-only hypothesis for WIMP spin-independent and spin-dependent interactions for WIMP masses above 9 GeV. The most stringent limits are set for a 30 GeV WIMP mass with a spin-independent WIMP-nucleon cross-section of $6.5 \times 10^{-48} \text{ cm}^2$ as shown in figure 3.15. A strong limit is set at this WIMP mass due to a statistical under fluctuation of events within the interior purple contour of figure 3.14.

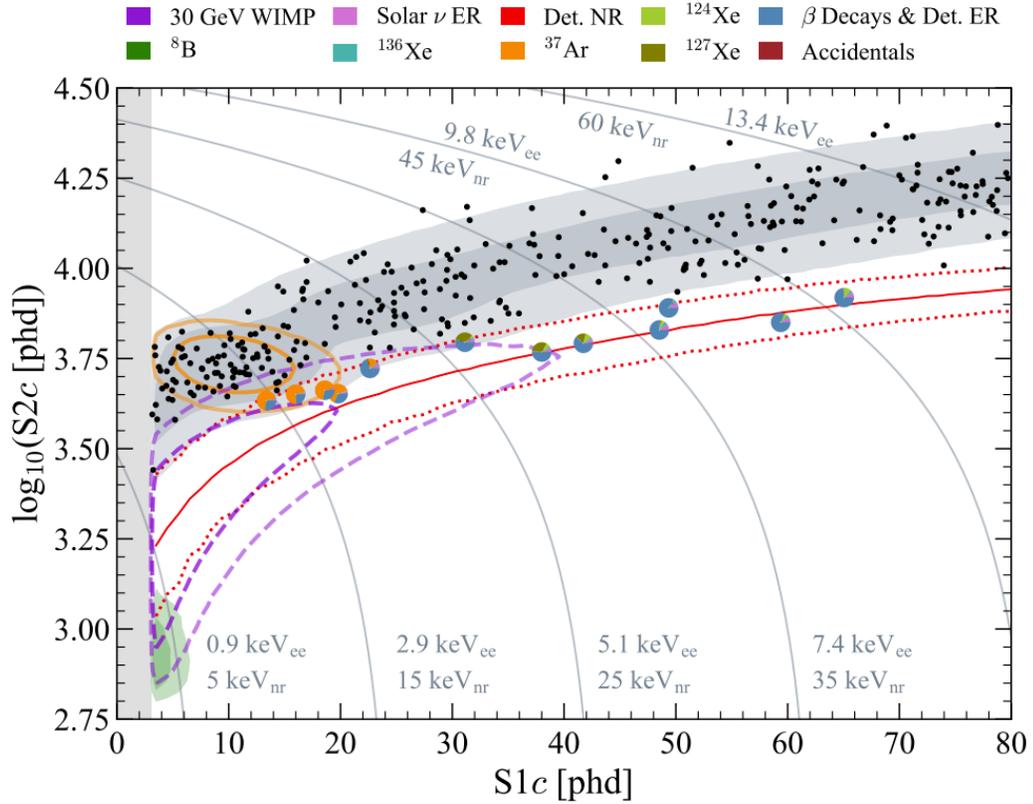


Figure 3.14: Low-energy recoil events after all data quality and physics cuts shown in the $\log_{10}(S2c)$ - $S1c$ space. A shaded grey band, orange ellipses contours (^{37}Ar), shaded green region (^8B solar neutrinos), and purple contours (30 GeV WIMP) are shown for 1σ and 2σ best-fit models. The solid red line shows the nuclear recoil median with the dotted lines indicating the 10% and 90% quantiles. Events that fall outside of the 2σ best-fit models are shown a pie chart with wedge size proportional to the relative weight of the background fit components corresponding to that region. Taken from reference [25].

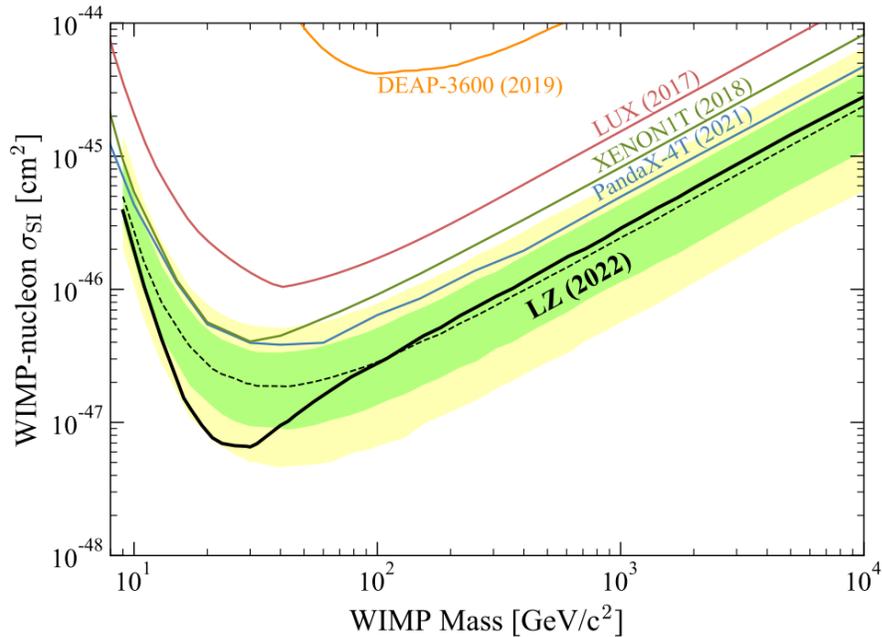


Figure 3.15: The 90% confidence limit for the spin-independent WIMP cross-section vs. WIMP mass. The dotted line is the median sensitivity projection. The green and yellow bands are the 1σ and 2σ sensitivity projections. Also included are results from the liquid xenon TPCs PandaX-4T, XENON1T, LUX, and the liquid argon scintillation only detector DEAP-3600. Argon scintillation signal shapes depend on the interaction type to discriminate electron recoils as described in section 2.5.2. Taken from reference [24].

3.7 Conclusion

In this chapter the LZ dark matter detector was presented. The function of each sub-detector as it pertains to the WIMP search was reviewed. The simulations and supporting software were discussed along with a robust calibration campaign to tune detector response for the WIMP search. The discrimination between nuclear and electron recoils was established and a determination of backgrounds was made. A final WIMP search result is reported for SR1 and shows a result consistent with a null hypothesis as well as a background model consistent with expected rates.

Chapter 4: Xenon Purity Monitor

Impurities in the LZ xenon are monitored throughout the lifetime of the experiment. Specific backgrounds of concern are caused by ^{85}Kr and ^{39}Ar . Both of these contaminants are reduced during the krypton removal campaign outlined in chapter 5. Any remaining traces of these impurities are present for the lifetime of LZ due to their long half-lives.

^{85}Kr is anthropogenically produced through nuclear fuel reprocessing. The atmospheric isotopic abundance is $2 \times 10^{-11} \text{ }^{85}\text{Kr}/^{\text{nat}}\text{Kr}$ [102]. ^{85}Kr is a beta background source with a half-life of 10.76 years and a total decay energy of 678 keV. Approximately 1% of these beta decays produce backgrounds in the WIMP search region of interest. LZ design specifications require that ^{85}Kr contributes a background comparable to the pp solar neutrino background. This amounts to an upper limit of 300 parts-per-quadrillion (ppq or 10^{-15}) by mass¹.

Cosmogenic production of ^{39}Ar results in an isotopic abundance of $8 \times 10^{-16} \text{ }^{39}\text{Ar}/^{\text{nat}}\text{Ar}$ [102]. ^{39}Ar has a half-life of 269 years and undergoes a beta decay with a 565 keV endpoint energy. The LZ design specifications require that ^{39}Ar contributes no more than 10% of the ^{85}Kr rate to detector backgrounds. This amounts to an upper limit of 9 parts-per-billion (ppb) $^{\text{nat}}\text{Ar}/^{\text{nat}}\text{Xe}$ [17] [28].

In addition to removing these noble contaminants, LZ constructed three gaseous xenon purity monitors (XPMs) to track xenon purity throughout purification and operation of LZ. The measurements from these XPMs inform the expected backgrounds shown in table

¹any concentrations reported are in mass/mass units unless otherwise stated

3.11. The three XPMs are the mobile sampling system (MSS), stationary sampling system (SSS), and mixing panel. The MSS was constructed at the University of Maryland and shipped to SURF in January 2019. The stationary sampling system (SSS) was partially constructed at UMD, completed at SLAC, and finally shipped to SURF in September 2021. The mixing panel was fully constructed at SLAC in May 2019 and shared hardware with the SSS. It served as a supplementary system to support operations during the krypton removal campaign.

This chapter will review the underlying principles of a generic XPM, a description of the system hardware, key identifiers for proper operations, operational procedures, and a review of each system's performance. A temperature-dependent study completed on the SSS is discussed as well as its implications for increasing XPM sensitivities. The first portion of this chapter draws heavily from the research and conclusions of chapter 3 from Jon Balajthy's thesis [12]. A description of the hardware and operational procedures will be given as they pertain to values taken from the mixing panel. The intention of this chapter is to offer guidance to future system operators through explicit procedures and accompanying examples of both poor, and ideal, system response.

4.1 Cold-Trap Operating Principle

The XPM relies on a freeze-distillation technique. The technique feeds xenon gas at a constant rate into a section of plumbing held at 77 Kelvin. A bath of liquid nitrogen or a cryogenic pulse tube refrigerator maintains a constant temperature. The bulk xenon is frozen to the walls of the cool plumbing within the "cold-trap". The flow of volatile impurities within the xenon remains largely unaffected by the xenon ice formation. The resulting gas mixture at the outlet of the cold-trap is the static xenon ice vapor pressure enriched with these volatile impurities. This process is depicted in figure 4.1. The krypton concen-

tration at the trap outlet is enriched by up to a factor of 10^9 . A commercially-available mass spectrometer (residual gas analyzer or RGA) monitors xenon ice background prior to flow and the accumulation of impurities during flow. The RGA is sensitive to impurities down to concentrations of one part-per-million (ppm). This combination of technique and hardware within the XPM has been demonstrated to measure xenon impurities down to approximately one part-per-quadrillion (ppq).

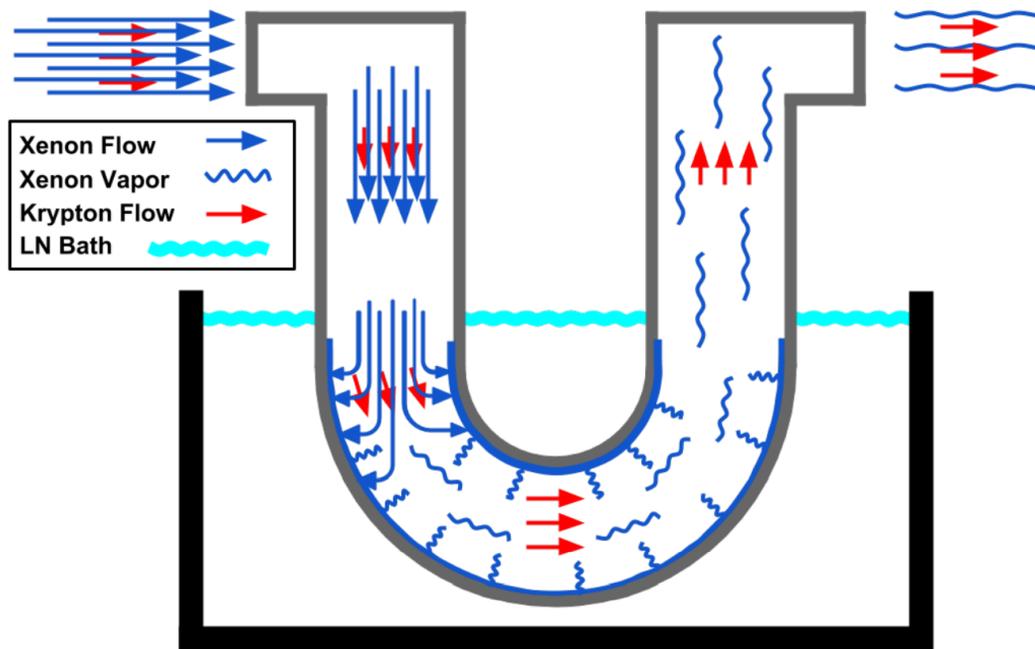


Figure 4.1: A depiction of a xenon-krypton mixture flowing through a cold-trap submerged in a bath of liquid nitrogen. The xenon is frozen and the outlet is enriched with krypton. This image is taken from reference [12].

Cold-trap mass spec techniques were initially developed to measure the volatile impurities oxygen, nitrogen, methane, argon, and helium in xenon down to the ppb scale [103]. Less volatile impurities, like water and hydrocarbons, do not pass through the trap and cannot be measured with this technique. The technique was extended to krypton and a sensitivity of 0.3 ppt g/g was achieved for the purpose of monitoring krypton concentrations in EXO [104] [105]. The LZ purity requirements motivated a further improvement

in krypton sensitivity. Jon Balajthy developed a vacuum impedance tuning technique and demonstrated a 7.7 ± 2.0 ppq limit of sensitivity [12].

4.1.1 Ideal System Response

A quick review of the detailed studies from reference [12] is given here. Vacuum equations are applied to the cold-trap and RGA volumes to inform a toy model of the system. A basic vacuum system can be treated as a simple electrical circuit where pressure, flow, and impedance act as vacuum analogues to voltage, current, and resistance. When applied to a gas system we can calculate the change in pressure ΔP (Torr), caused by a gas flowing at a rate Q (Torr·liter/sec), through a section of plumbing with impedance Z (sec/liter)

$$\Delta P = Q \cdot Z. \quad (4.1)$$

The volumetric flow rate, or pumping speed S (liter/sec), at any point within the vacuum system is defined by the rate at which a gas volume moves through it:

$$S \equiv \frac{dV}{dt} = \frac{Q}{P} \quad (4.2)$$

where P is the absolute pressure and Q is the flow at that point. At equilibrium, Q can be used to relate steady state flow between connected volumes, here equation (4.2) acts as a continuity equation. For a system evacuated by a pump with pumping speed S_p , the effective pumping speed S_{eff} at any point within it is determined by the total impedance Z_{total} between that point and the pump:

$$S_{eff} = \left(\frac{1}{S_p} + Z_{total} \right)^{-1}. \quad (4.3)$$

Then the steady state pressure at any point within the system that experiences an equilibrium gas flow Q is given by

$$P = \frac{Q}{S_{eff}}. \quad (4.4)$$

In the XPM there are two relevant volumes, the RGA volume and cold-trap outlet. Each volume has a unique pumping speed, adjusted by valves identified in figure 4.2. The RGA pumping speed is set by a turbo-pump and the adjustable vacuum valve that modifies Z_2

$$S_{RGA} = \left(Z_2 + \frac{1}{S_{TP}} \right)^{-1}. \quad (4.5)$$

The cold-trap outlet is connected to the RGA via the adjustable impedance valve that modifies Z_1 and has a pumping speed

$$S_{CT} = \left(Z_1 + \frac{1}{S_{RGA}} \right)^{-1}. \quad (4.6)$$

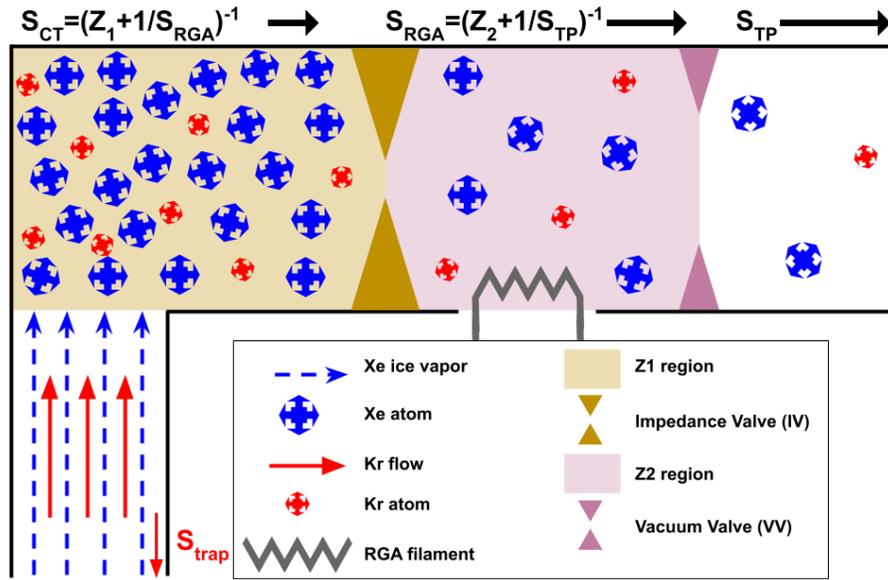


Figure 4.2: A depiction of the equilibrium of krypton flow throughout the system. The amount of xenon in the Z1 region is fixed, but the equilibrium quantity of krypton depends on flow rate into the cold-trap and the pumping speed at which it is depleted. Impedance values Z_1 and Z_2 set the pumping speed for each region and are controlled by IV and VV respectively. Note the S_{trap} pumping speed that drives krypton entrapment flow into ice.

The krypton equilibrium pressure throughout these volumes is determined by these pumping speeds and the flow rate of krypton into the system. The krypton flow into the cold-trap is set by a xenon flow with a krypton concentration of Φ_{Kr} .

$$Q_{Kr,IN} = \Phi_{Kr} Q_{Xe,IN} \quad (4.7)$$

In practice, not all of the krypton is transmitted perfectly to the RGA. Some of the krypton remains trapped in the xenon ice during formation. To describe the equilibrium flow observed at the RGA we define a throughput parameter.

$$\alpha = \frac{Q_{Kr,RGA}}{Q_{Kr,IN}} \quad (4.8)$$

Combining equations (4.7) and (4.8) gives

$$Q_{Kr,RGA} = \alpha \Phi_{Kr} Q_{Xe,IN}. \quad (4.9)$$

The corresponding pressure is found by applying the steady state pressure from equation (4.4) to the equilibrium flow at the RGA,

$$P_{Kr,RGA} = \frac{1}{S_{RGA}} \alpha \Phi_{Kr} Q_{Xe,IN}. \quad (4.10)$$

Any krypton observed at the RGA is removed from the xenon and pumped to atmosphere. By repeating measurements of the same xenon sample, under identical conditions, we measure how much krypton was trapped in the ice of the previous sample. The reduction factor is equal to the ratio of the krypton equilibrium pressures observed during consecutive runs: $P_{Kr,1}/P_{Kr,0} = (1 - \alpha)$. If we repeat this at various impedance states, we see that higher impedance states trap more krypton in the xenon ice, as shown in figure 4.3. This entrapment can be approximated as the flow of krypton into the ice within cold-trap Q_{trap} .

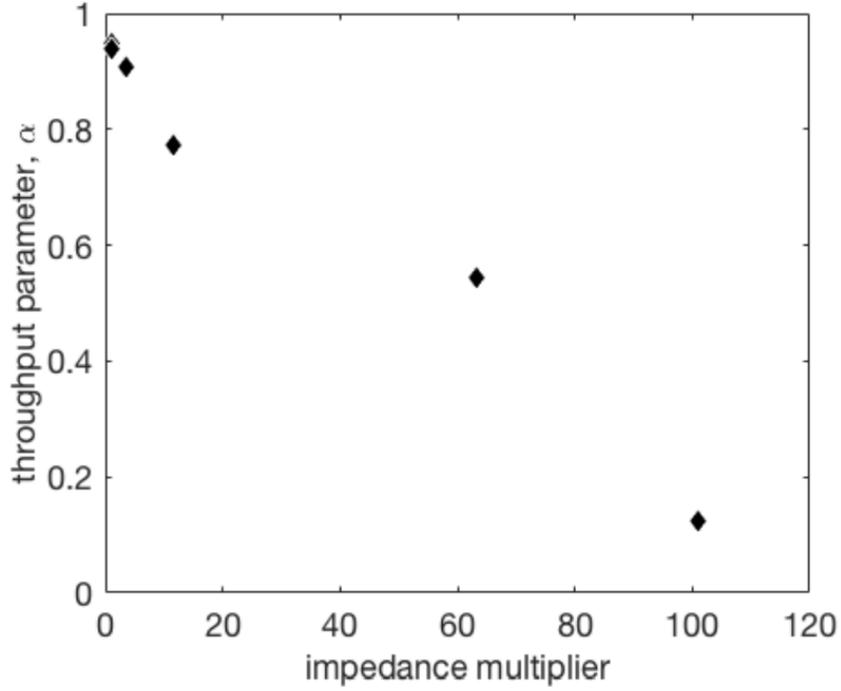


Figure 4.3: The krypton throughput parameter α as a function of relative impedance from the 1x impedance state. Taken from reference [12].

The krypton pressure time response in a volume is modeled by a differential equation that describes the rate at which krypton is supplied to and removed from that volume.

$$V \frac{dP_{Kr}}{dt} = Q(t) - SP_{Kr} \quad (4.11)$$

The response of P_{Kr} to an impulse flow, $Q(t) = P_0V\delta(t)$ is

$$P_{Kr}(t > 0) = P_0e^{-t/\tau} \quad (4.12)$$

where the time constant is given by $\tau = VS$

At the cold-trap outlet, the krypton pressure is sourced by the xenon flow and depleted by the flow of krypton into the ice as well as the flow into the RGA. Equation (4.11)

becomes

$$V_{CT} \frac{dP_{Kr,CT}}{dt} = \Phi_{Kr} Q_{Xe,IN} - Q_{trap} - Q_{RGA}. \quad (4.13)$$

At equilibrium, where $dP_{Kr,CT}/dt = 0$, the RGA flow is given by (4.10). Then equation (4.1.1) can be rewritten as

$$Q_{RGA} = \Phi_{Kr} Q_{Xe,IN} - Q_{trap} = \alpha \Phi_{Kr} Q_{Xe,IN}. \quad (4.14)$$

We note that the total amount of krypton supplied to the cold-trap is $\Phi_{Kr} Q_{Xe,IN} = S_{CT} P_{Kr,CT} / \alpha$ and express Q_{trap} in terms of the throughput parameter

$$Q_{trap} = (1 - \alpha) \Phi_{Kr} Q_{Xe,IN} = \left(\frac{1}{\alpha} - 1 \right) S_{CT} P_{Kr,CT} \quad (4.15)$$

The rightmost term shows the flow of krypton into the cold-trap is equivalent to the pumping speed $Q_{trap} = S_{trap} P_{Kr,CT} = (1/\alpha - 1) S_{CT} P_{Kr,CT}$. This allows us to rewrite α in terms of pumping speeds:

$$\alpha = \frac{S_{CT}}{S_{CT} + S_{trap}}. \quad (4.16)$$

Applying this to the steady state pressure of krypton at the RGA from equation (4.10) gives

$$P_{Kr,RGA} = \frac{1}{S_{RGA}} \left(\frac{S_{CT}}{S_{CT} + S_{trap}} \right) \Phi_{Kr} Q_{Xe,IN}. \quad (4.17)$$

We must now consider the maximum operating pressure of the RGA which is $P_{RGA} = 10^{-5}$ Torr. At 77.15 K, xenon ice has a vapor pressure of $P_{ICE} = 1.8 \times 10^{-3}$ Torr. By relating the equilibrium flow between the volumes, applying the steady state pressure given by equation (4.4), and imposing the RGA pressure condition to equation (4.1), we find

$$Q_{Xe,RGA} = P_{RGA} S_{RGA} = \frac{P_{ICE} - P_{RGA}}{Z_1}. \quad (4.18)$$

This defines our pressure drop requirement in terms of the impedance value Z_1 , matched to a pumping speed S_{RGA} . The RGA pressure condition can be satisfied by adjusting IV and VV, at a fixed temperature of 77.15 K. We can rewrite this constraint, using $P_{RGA} \ll P_{ICE}$, to give a fixed ice pressure impedance condition of

$$Z_1 S_{RGA} = \frac{P_{ICE}}{P_{RGA}} = 180. \quad (4.19)$$

As an additional consequence note that $Z_1 \gg 1/S_{RGA}$ and approximate equation (4.6) as $S_{CT} \approx 1/Z_1$. Applying (4.19) and our approximate S_{CT} we express α/S_{RGA} as

$$\frac{\alpha}{S_{RGA}} \approx \frac{Z_1}{180} \cdot \frac{1}{1 + Z_1 S_{trap}} \quad (4.20)$$

Replacing the α/S_{RGA} term in (4.17) with the approximations made above, gives a krypton equilibrium pressure at the RGA as a function of the impedance Z_1 .

$$P_{Kr,RGA} = \frac{Z_1}{180(1 + Z_1 S_{trap})} Q_{Xe,IN} \Phi_{Kr} \quad (4.21)$$

This equation relates the equilibrium krypton pressure measured at the RGA to the krypton concentration of the xenon gas fed into the cold-trap. Furthermore there are two parameters, $Q_{Xe,IN}$ and Z_1 , that an operator can adjust to maximize the krypton signal. The flow allows the operator to increase the rate at which the ice vapor at cold-trap outlet is enriched with krypton while reducing the rate at which it is depleted by pumping across IV. At very high Z_1 values, the entrapment term in the denominator $1/(1 + Z_1 S_{trap})$, begins to compete with the gains made by increasing Z_1 leading to a plateau.

We now define a “default state” for which the valves VV and IV are maximally opened. From the default state we set a “1x state” for xenon ice held at a constant 77.15 K by minimally adjusting IV or VV to satisfy the fixed ice pressure impedance condition from

equation (4.19). The 1x state is set according to the procedure in section 4.5.6. Once the 1x impedance state is set, a relative impedance state is set that raises Z_1 and lowers S_{RGA} , by a factor of Y, while maintaining the constant pressure at the RGA. This new “Yx state” is set according to the procedure in section 4.5.8.

Measurements, shown in figure 4.4, taken at a variety of impedance states, under otherwise identical conditions, fits the model well.

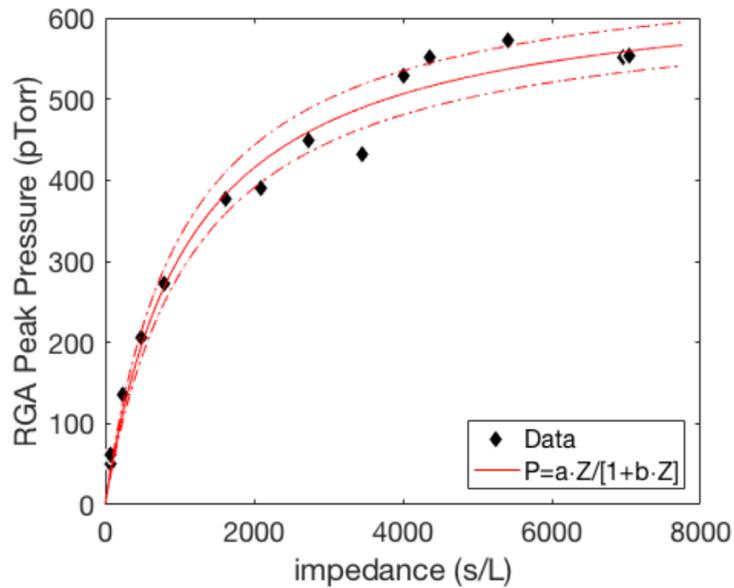


Figure 4.4: Observed peak P_{Kr} , measured at increasing impedance. The data fits reasonably well with the equilibrium krypton pressure equation (4.21). Taken from reference [12].

One behavior not captured by this model is the loss of a linear relationship between the equilibrium signal size and flow rate, $Q_{Xe,IN}$, at high flows. At flows below 1 slpm, xenon ice forms as a collar at the cold-trap inlet and quickly clogs, limiting the total xenon capacity. At flows above 1 slpm an insulating layer of xenon ice forms throughout the trap, limiting heat transfer to the cold-trap walls. This causes the ice to grow along the length of the cold-trap as a sleeve. By taking advantage of the ice sleeve, we increase the effective capacity of the trap. In doing so, the ice penetrates further into the trap until it eventually

reaches a corner bend near the cold-trap outlet. At this moment, the equilibrium signal observed at the RGA is reduced due to a change in α .

A corrected equilibrium value can be found using the krypton pressure's rising edge prior to the turnover and fitting it to the expected signal response as shown in figure 4.5. The ideal shape of the expected krypton signal as a function of time without taking into consideration the corner effect is

$$P_{Kr,RGA}(t) = \Phi_{Kr} \frac{\alpha Q_0}{S_{CT}} \begin{cases} 0 & \text{for } t < 0 \\ 1 - e^{-\frac{t}{\tau_{CT}}} & \text{for } 0 \leq t \leq T \\ e^{-\frac{t-T}{\tau_{CT}}} (1 - e^{-\frac{T}{\tau_{CT}}}) & \text{for } t > T \end{cases} \quad (4.22)$$

where $t = 0$ is xenon supply start time and $t = T$ is the time at which the xenon supply is stopped. For a xenon flow given by

$$Q_{Xe}(t) = \begin{cases} Q_0 & \text{for } 0 \leq t \leq T \\ 0 & \text{otherwise} \end{cases} \quad (4.23)$$

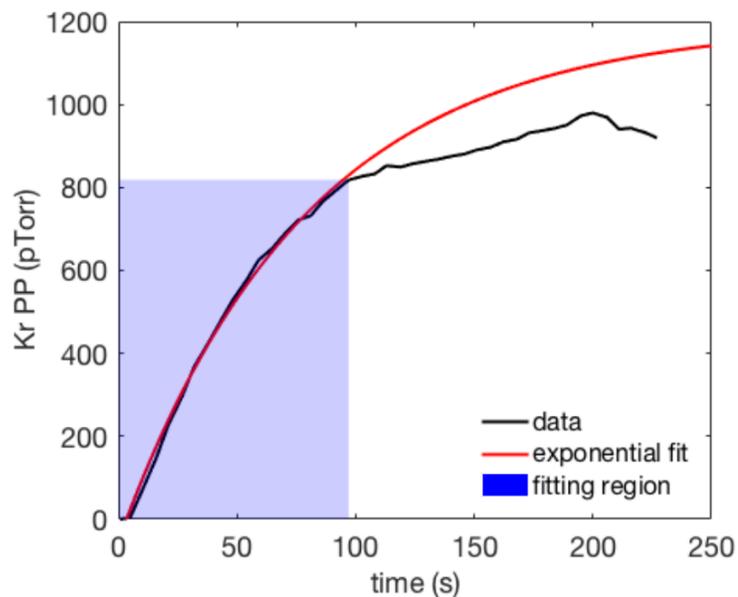


Figure 4.5: The PP_{Kr} partial pressure as a function of time from flow start at a flow of 3.5 slpm. The krypton signal deviates from the expected shape at 100 seconds into the xenon flow as the ice formation reaches the outlet corner of the cold-trap. Taken from reference [12].

From the above studies the ideal operating parameters are motivated:

- A cold trap geometry with a long straight section of tubing to host an ice sleeve, as shown in figure 4.6. We select a flow rate and duration to restrict ice formation to the straight bottom portion of the stocking trap while maximizing capacity. This ensures that α is kept constant for the duration of the measurement.



Figure 4.6: A stocking trap made from 0.5” stainless steel tubing. The inlet side of the cold-trap is on the left. The bottom of this trap is 9” wide and the toe section is 7” tall. The trap is submerged in liquid nitrogen such that the bottom is 1” above the bottom the dewar to allow maximal cooling. The liquid nitrogen level is maintained 2” above the top of the toe section. A similar trap with a 11” wide bottom span was used on the mixing panel. In this trap we attempt to form an ice sleeve in the left and bottom portions where α remains constant for the duration of the flow.

- A flow rate between 1 and 3 slpm to maximize the trap’s capacity for xenon ice while avoiding ice formation beyond the bottom straight section. The maximal flow value will ultimately depend on the specific geometry of the cold-trap.
- A relative impedance value between 5x and 30x. Above 100x impedance gains are minimal, due to the $1/(1 + Z_1 S_{trap})$ term in equation (4.21). In practice, an impedance state near 15x is suitable for LZ. In most systems this is amounts to an ultimate S_{RGA} value around 250 cc/s and Z_1 around 400 s/cc. This is a compromise between the desire to maximize any potential krypton signal and reduce chances of

signal saturation by other gas species like helium, nitrogen, and argon.

4.2 The Generic Xenon Purity Monitor

A general purpose XPM is shown in figure 4.7. This XPM is a simplification of the mixing panel described in section 4.7. I will use the mixing panel volumes, hardware, and operations to describe the generic XPM.

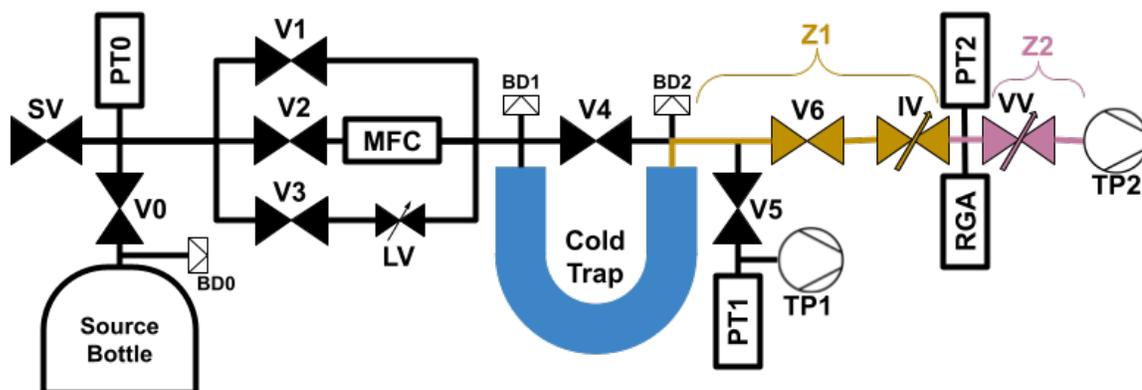


Figure 4.7: A schematic of a generic XPM. The system valve (SV) connects the XPM to a volume from which xenon samples are drawn. The XPM uses valves V0-V3 to deliver xenon from the source bottle to the cold trap. Xenon flow is regulated by the mass flow controller (MFC) or leak valve (LV). Xenon gas pressure is monitored by a pressure transducer (PT0). Valves V5 and V6 protect the vacuum turbo-pumps (TP1 and TP2), vacuum gauges (PT1 and PT2), and RGA. The impedance valve (IV) and vacuum valve (VV) set the impedance values Z_1 (yellow) and Z_2 (pink).

A description of the hardware is given below.

Tubing and fittings:

- Vacuum lines are made of 2.75" Con-Flat (CF) plumbing with copper gaskets. Vacuum spaces are located beyond valves V5 and V6.
- The cold-trap is 0.5" diameter 316L, stainless steel tubing, bent into a U-shape or stocking shape as shown in figure 4.6.
- The remaining lines are made of 3/8" diameter 316L UHP stainless steel tubing.

- All tubing junctions and valve interfaces, excluding vacuum regions, are either welded or use VCR (vacuum coupling radiation) face seal couplings. While 1/4" VCR fittings are compatible with the 3/8" HVCR fittings, 3/8" is preferred for its lower impedance.

Source Bottle

- The source bottle is a modified stainless steel Swagelok 304L-HDF8-1GAL bottle. The ends of the bottle are threaded for 1/2" female NPT fittings, but have been replaced with modified threaded caps. The caps are welded in place, to reduce potential leaks from cryogenic temperature cycling as these bottles are often used as a cryopump. A cryopump uses condensation on a cold surface to pump gasses.

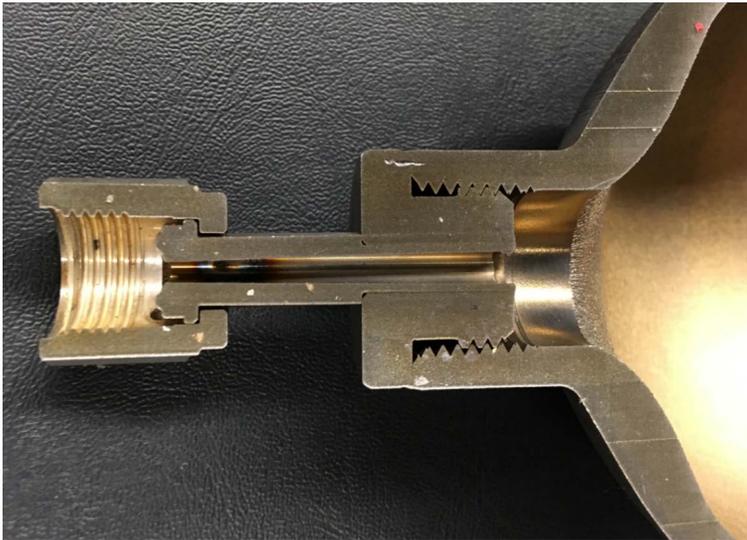


Figure 4.8: An cross-sectional image of a modified Swagelok 304L-HDF8-1GAL bottle. End caps have been machined to fit the 1/2" FNPT threads and welded to the 1 gallon bottle. The top cap includes a welded HVCR fitting and a female HVCR nut, the bottom cap is kept blank. A relief cut along the threads avoids trapped gas volumes. The bottle is leak tested and pressure tested following modification. Taken from the LZ controlled document database [26].

- In more complex systems the source bottle may be replaced an array of instrumented volumes to support the storage, mixing, or collection of xenon.

Burst Disks:

- Burst disks rupture at a set pressure to protect system components from over-pressure. The rupture pressure is selected to be less than 75% of the lowest rated component's maximum allowable working pressure. They must be included anywhere xenon ice is formed and at the output of high pressure regulators.
- For the generic XPM a value of 10 barg (bar gauge) is used. All-metal burst disks and holders, rated from 3 to 80,000 psig, are available from Continental Disk. The 10 barg disks used have a part number CD30775A0012.
- BD0 protects against over pressure in the source bottle.
- BD1 and BD2 protect against over pressure in the cold-trap.

Valves:

- Valves V0 to V6, and SV are Swagelok DF series valves. The polymer seats in these valves are susceptible to contamination from high concentrations of krypton. Valve V6 is particularly vulnerable, as it is exposed to the krypton enriched ice vapor. Avoid exposing V6 to xenon with a krypton concentration exceeding 5 ppb. The Swagelok part number for these valves is 6LV-DF###-#-#. The additional # specifies actuator type, sizing, and connection type.
- If pneumatically actuated valves are used, select normally open (NO) for valves V0, V1, V2, V3, v4 and normally closed (NC) for valves SV, V5 and V6. In the event of pneumatic pressure loss, valves will default into a safe state defined in section 4.5.1.
- SV isolates the XPM from a larger system.
- V0 isolates the source bottle from PT1.

- V1 bypasses the MFC.
- V2 selects the MFC for flow into the cold-trap.
- V3 selects the LV for flow into the cold-trap.
- V4 bypasses the cold-trap.
- V5 isolates the utility pump from the system.
- V6 isolates the vacuum RGA components from the system.
- The leak valve (LV) is a metering valve used for low flow measurements. If the MFC can regulate flows near 0.1 slpm, then the MFC can be used instead.
- The impedance valve (IV) adjusts the impedance value Z_1 . IV is a B series Swagelok valve with a metal spherical stem tip and HVCR connections, part number SS-4BG-V51. Regulating stem tip B series valves have also been used.
- The vacuum valve (VV) adjusts the impedance value Z_2 . VV is a 2.75" Con-Flat, all-metal, right angle valve from Ideal Vacuum, part number P107281.

Pressure Transducers:

- The pressure transducer (PT0) is a capacitance manometer used to monitor xenon pressure. PT0 is a Baratron absolute capacitance manometers from MKS Instruments with a range of 0-5000 Torr. The part number is 722C53TCD2FA.
- The vacuum pressure gauges PT1 and PT2 are inverted magnetron gauges. They measure vacuum from 5×10^{-9} Torr up to 1×10^{-3} Torr. These gauges must be turned off at pressures above 1×10^{-3} Torr to avoid damage. IMG-100 Agilent Varian gauges are available from Ideal Vacuum Products, part number P104146.

- PT1 monitors the utility vacuum pump-out through V4.
- PT2 monitors the RGA vacuum. PT2 should be turned off during high-sensitivity measurements as it raises RGA backgrounds.

Mass Flow Controller:

- The MFC is a Teledyne Hastings mass flow controller, part number HFC-D-302B (H). It has all metal internals and a flow range of 0-5 slpm. It is capable of stable flow at 0.1 slpm. A minimum backing pressure of 300 Torr is required for stable flow at 2 slpm.

Pumps:

- Both TP1 and TP2 are turbo molecular TwisTorr 84 FS pumps with 4.5" CF flanges. TwisTorr 84 FS turbo pumps are available from Agilent. Each turbo is backed with a dry scroll pump, historically an Agilent SH-110 or Edwards XDS10 which are no longer available. An equivalent dry scroll pump would be an Agilent IDP-7 or IDP-10.
- TP1 is a utility pump.
- TP2 is the RGA pump.

Residual Gas Analyzer:

- The SRS RGA200 is a mass spectrometer from Stanford Research Systems. The RGA operates at pressures less than 10^{-4} Torr and measures the charge-to-mass ratio of gasses. A more detailed discussion of the RGA is found in section 4.3.

4.3 The RGA

The SRS RGA200 is a mass spectrometer that monitors the partial pressures of gasses in a vacuum system. The entire package includes a gas probe and electronics control unit (ECU). The probe contains an ionizer, ion filter, and ion detector [27].

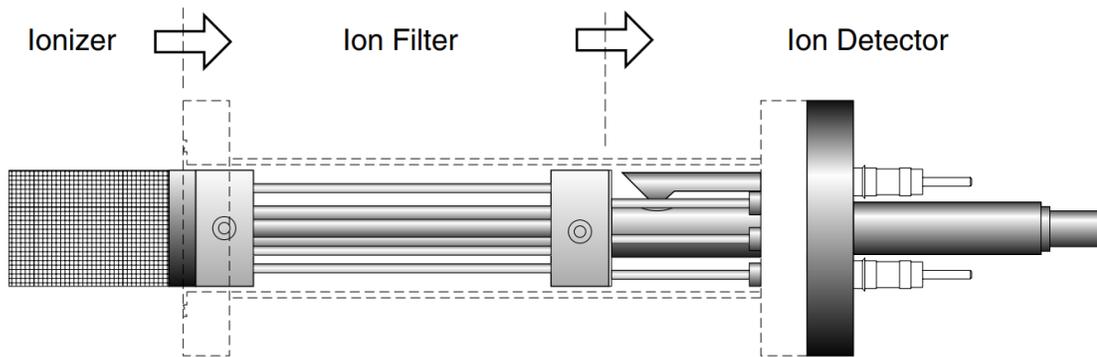


Figure 4.9: An SRS RGA200 probe schematic. From the RGA manual [27].

4.3.1 Ionizer

The first section of the probe is the ionizer. The outermost section of the ionizer is a negatively biased, cylindrical repeller cage. Within the repeller cage, a filament encircles a cylindrical anode cage. Electrons, emitted by the filament, accelerate away from the repeller cage and towards the anode. This causes electrons to recirculate about the probe's central z-axis until striking and ionizing a gas molecule. The ions collect along the central axis due to the high electron density. A negatively biased focus plate, with a hole along the z-axis, attracts the ions into the ion filter [27].

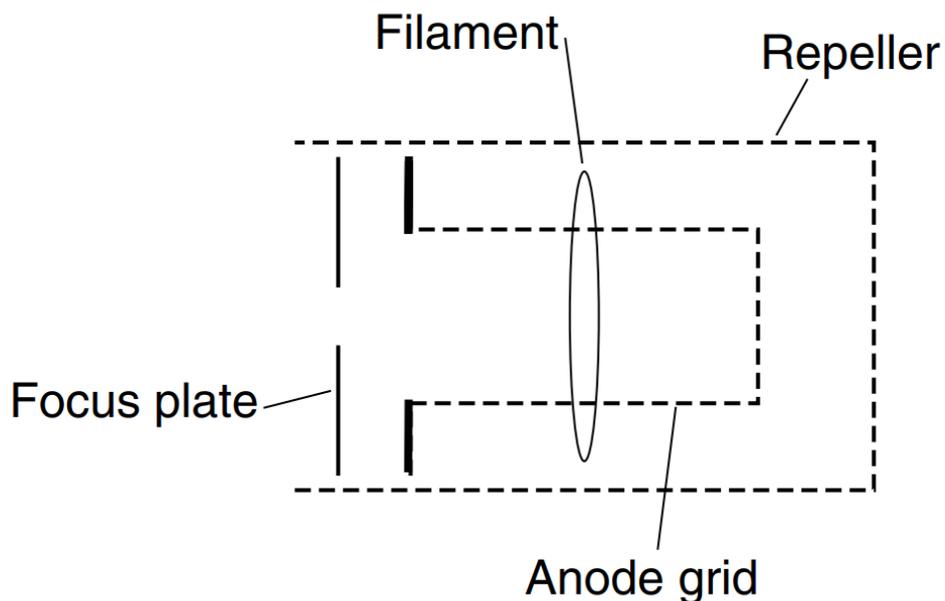


Figure 4.10: A cross section of the SRS RGA200 ionizer. From the RGA manual [27].

4.3.2 Ion Filter

The ion filter is a quadrupole mass filter that varies DC and RF voltages at a fixed frequency. As ions enter the filter along the z -axis, they start to oscillate in the x and y directions. The voltage of the rods are

$$V_{x,y} = \pm(U + V_0 \cos(\omega t)), \quad (4.24)$$

where U is the magnitude of the constant voltage component and V_0 is the maximum voltage of the oscillating RF component. Light mass ions oscillate in phase with the RF component, gaining amplitude until striking a rod. Heavier ions are not influenced by the RF oscillations, but instead drift off the central axis due to the constant voltage component. Together these effects form a band-pass and the ratio of the voltages, V_0/U , selects which charge-mass ratios have a stable path along the probe's central axis to the charge collector.

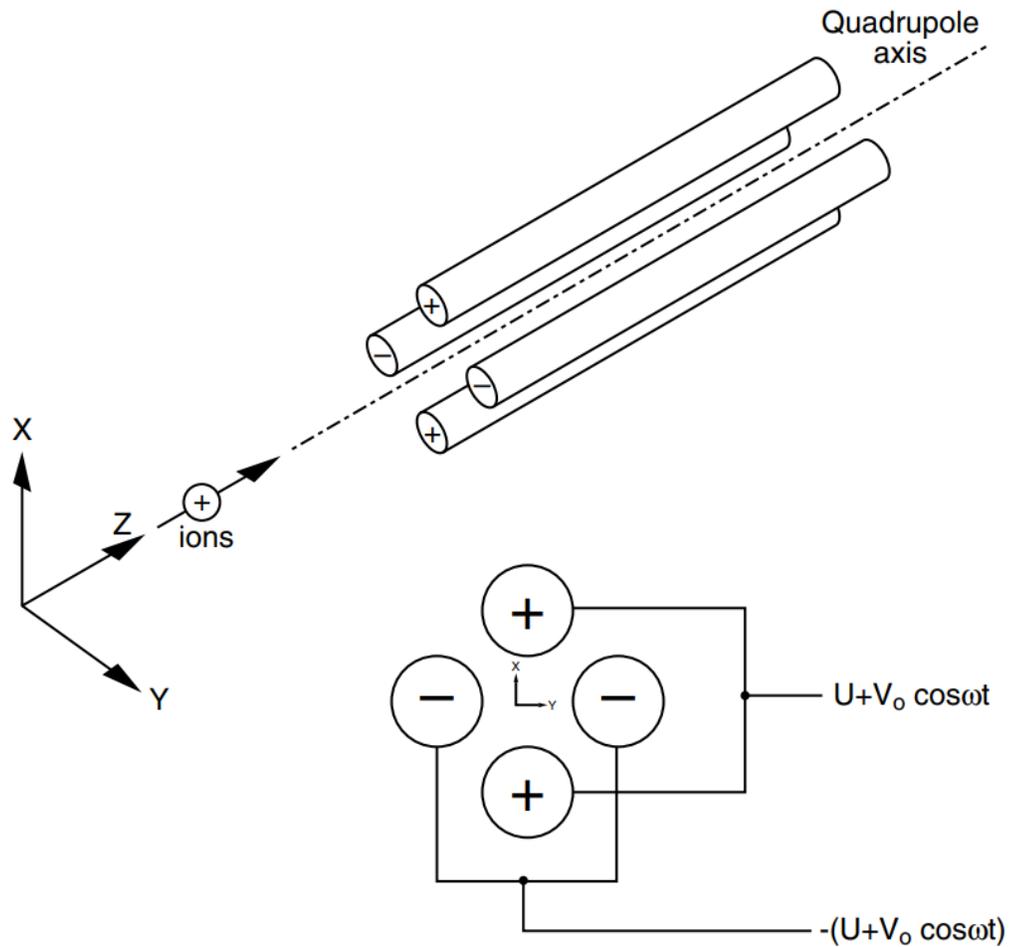


Figure 4.11: An SRS RGA200 quadrupole mass filter. From the RGA manual [27].

4.3.3 Charge Collector

Ions that survive the filter are collected by the Faraday Cup or the Continuous Dynode Electron Multiplier (CEM). Ions collected by the Faraday Cup strike a grounded wall. This induces an electron current to neutralize the ion which is measured by an electrometer. The output voltage of the electrometer is equal to the logarithm of the ion current.

When a negative high voltage is applied to the CEM, between -1000 to -2500 volts, ions are pulled off the central axis and towards the CEM. The ions accelerate and collide

with the CEM causing showers of secondary electrons. The number of electrons produced in a shower depends on the ion mass and the CEM high voltage. The electrons then strike a grounded plate at the output of the CEM and produce a positive current across the same electrometer. The output voltage equals the logarithm of the electron current which is proportional to the ion current striking the CEM. It is important to note that CEM measurements are mass-dependent due to the ion-electron conversion efficiencies.

The voltage output of the electrometer in both cases is digitized by a 16 bit analog-to-digital converter. The digitized signal is converted to a value in picoTorr through internally stored coefficients in the ECU memory [27].

4.3.4 RGA Conditioning

The RGA must be conditioned after any exposure to air. Before conditioning, check for electrical shorts between the probe's electrical pins and its Con-Flat housing. It is best practice to fully open VV to maximize the pumping speed at the RGA. The system should be pumped to less than 5×10^{-8} Torr for at least four hours before the filament is turned on. Once the filament is on the pressure will spike as the filament burns off any surface contaminants. The filament should be left to pump for another four hours. Finally the CEM can be turned on to the last operating voltage used. If a new CEM has been installed use 1000 volts. Avoid setting the CEM voltage higher than needed as this will reduce its lifetime. The RGA should be left to pump with the filament and CEM on for at least 12 hours.

4.3.5 RGA Peak Tuning

Over time the response of the RF filter can vary. This is usually caused by degrading electronics or ambient temperature swings. To check that peak values appear where

expected, an “Analog Scan”, as shown in figure 4.12, should be performed.

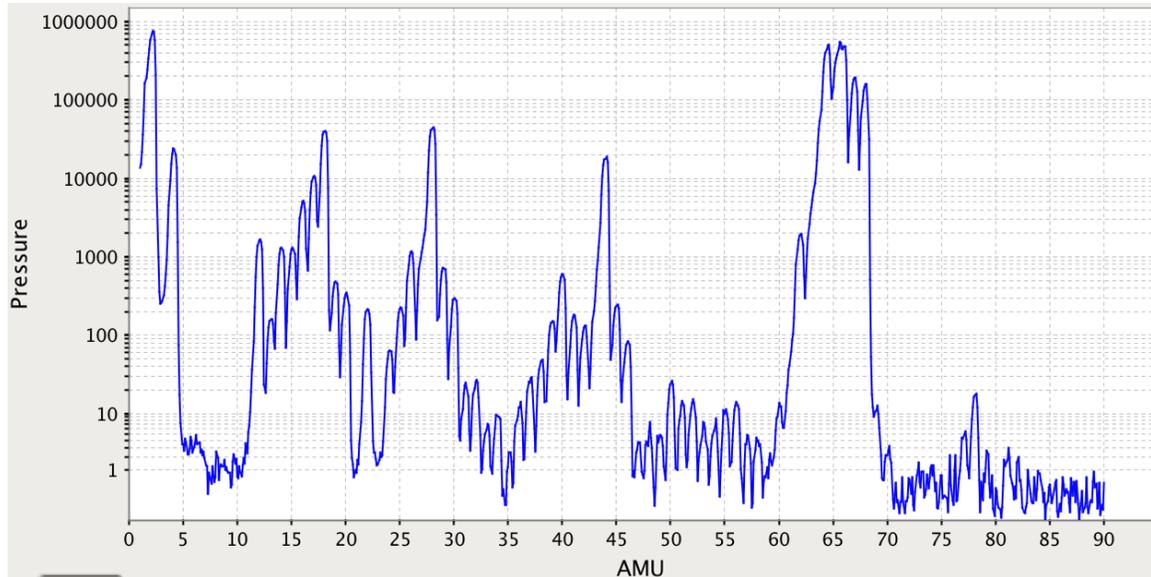


Figure 4.12: An analog scan of xenon ice. The x-axis mass is in AMU and the y-axis is partial pressure in picoTorr (pTorr). This analog scan was taken while the RGA was opened to xenon ice. Some easily identifiable peaks here are due to helium (4 AMU), water (18 AMU), nitrogen (28 AMU), CO₂ (44 AMU), and the doubly ionized xenon peaks (62 to 68 AMU). We also note that the mass channels where krypton would be (80 - 86), have no peaks.

Analog scans measure the partial pressure of each the constituent mass-charge ratios, as shown in figure 4.12. If the peak position is more that 0.3 AMU off of the expected value, the peak position should be tuned as specified in the RGA manual [27].

4.3.6 RGA Service and Operating Parameters

RGA probes and ECUs are matched by the vendor. Mismatching equipment should be avoided as it can lead to unexpected behavior. It is recommended to keep a spare RGA setup on hand in case any unit requires vendor servicing. Multiple filaments and CEMs should be kept in reserve as well. The CEM can be reconditioned and the Faraday cup cleaned according to procedures in the RGA manual [27]. Some RGA units are more susceptible than others to swings in ambient temperature and humidity, a controlled environment is

recommended.

The RGA has a user settable noise floor parameter between 0-7 (also called scan speed and with equivalent values 1-8). At lower noise floors, the RGA will spend more time averaging measurements of a single mass channel, reducing the statistical noise of a measurement. A noise floor of zero is typically used and averages the signal of each mass channel over two second intervals.

The ideal CEM voltage setting is achieved when the doubly ionized ^{124}Xe partial pressure reads 2000 pTorr while the total xenon ice pressure at the RGA is 10^{-5} Torr. The CEM will degrade over time requiring an increased voltage to maintain the gain setting needed satisfy the ^{124}Xe signal requirement. Once the voltage exceeds 1900 volts the CEM should be replaced, however it is operational up to 2200 volts. The RGA becomes less sensitive over time as the CEM degrades. XPM calibrations should be completed as the required voltage setting increases 100 volts.

4.4 Mass Spec Pathologies

4.4.1 RGA Saturation

If the pressure at the outlet of the XPM is not kept constant, either due to significant quantities of impurities or improper xenon ice temperatures, the RGA signals will be distorted. This type of “pressure saturation” can render data useless and potentially damage the vacuum equipment. Pressure increases due to improper xenon ice temperature are mitigated by the formation of an ice kernel as described in section 4.5.5. This kernel provides nucleation points for xenon ice along the length of the cold-trap. An extreme example of pressure saturation, due to excessive helium impurities, is given in figure 4.13. The increased pressure at the RGA not only induced false signals near the electronic baseline, but also diminished signals away from baseline due to gas accumulation in the ion filter.

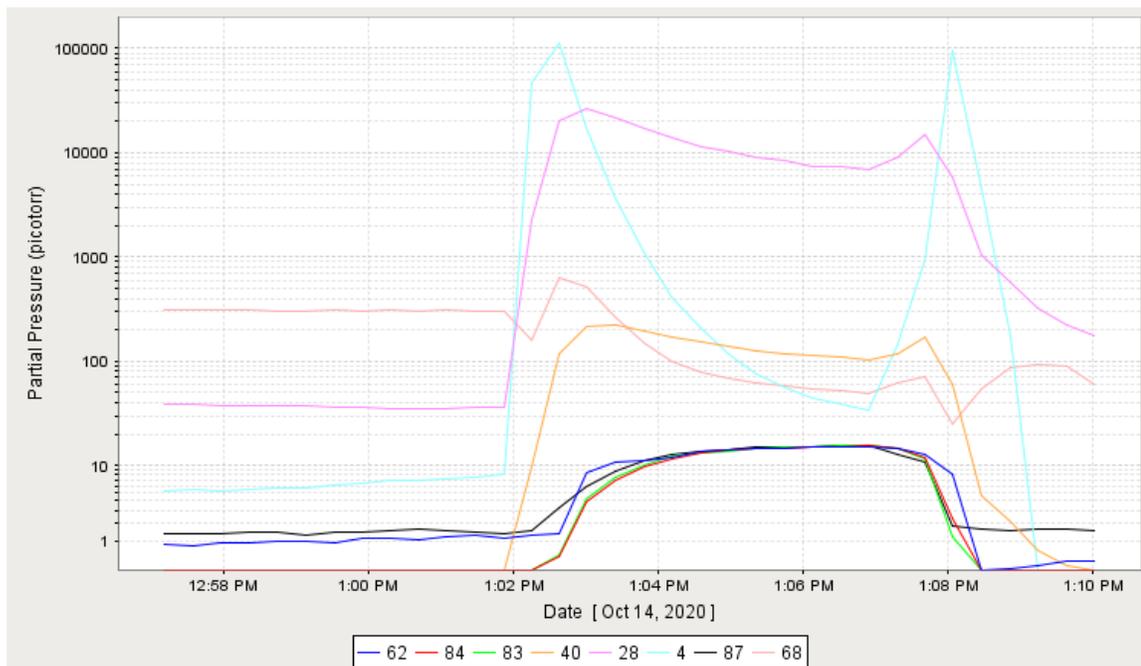


Figure 4.13: An RGA pressure versus time scan tracks the partial pressure in pTorr of different mass channels over time. The different line colors represent different gases corresponding to the mass in AMU. The total pressure, monitored by a vacuum gauge, increased from 8.6×10^{-6} Torr to 3×10^{-5} Torr due a large concentration of helium (light blue). An accurate measurement was not possible due to the saturation. Flow became restricted due to the accumulation of helium in the trap. The clogging of the cold-trap or cryopump due a non-condensable gas is known as vapor-lock.

Another type of saturation is known as “electronic saturation”. This occurs when the total pressure remains stable at the RGA, but a single mass channel produces an ion current large enough to saturate the RGA ADC. This signal impacts other channels and causes all signals to trend with the saturating channel. In general, the singly ionized xenon channel, ^{124}Xe , may cause electronic saturation and noise in the RGA’s electronic baseline. We avoid this by monitoring the doubly ionized ^{124}Xe channel at 62 AMU, referred to as “ ^{62}Xe ”. We can further track the electronic baseline behavior using the 87 AMU channel, which we refer to as “ ^{87}Bi ”. This channel is known to have no physical ions in a typical xenon purity measurements and is used to signal poor measurement conditions. The krypton partial pressure is tracked primarily through the ^{84}Kr , but the ^{83}Kr channel is monitored

as well.

An example of electronic saturation due to a nitrogen signal in the 28 AMU channel is shown in figure 4.14. The initial size of the nitrogen signal was a factor of 10 times larger than the usual ice kernel background. In future runs, this initial signal was used to flag the xenon sample as requiring additional preparation prior to measurement.

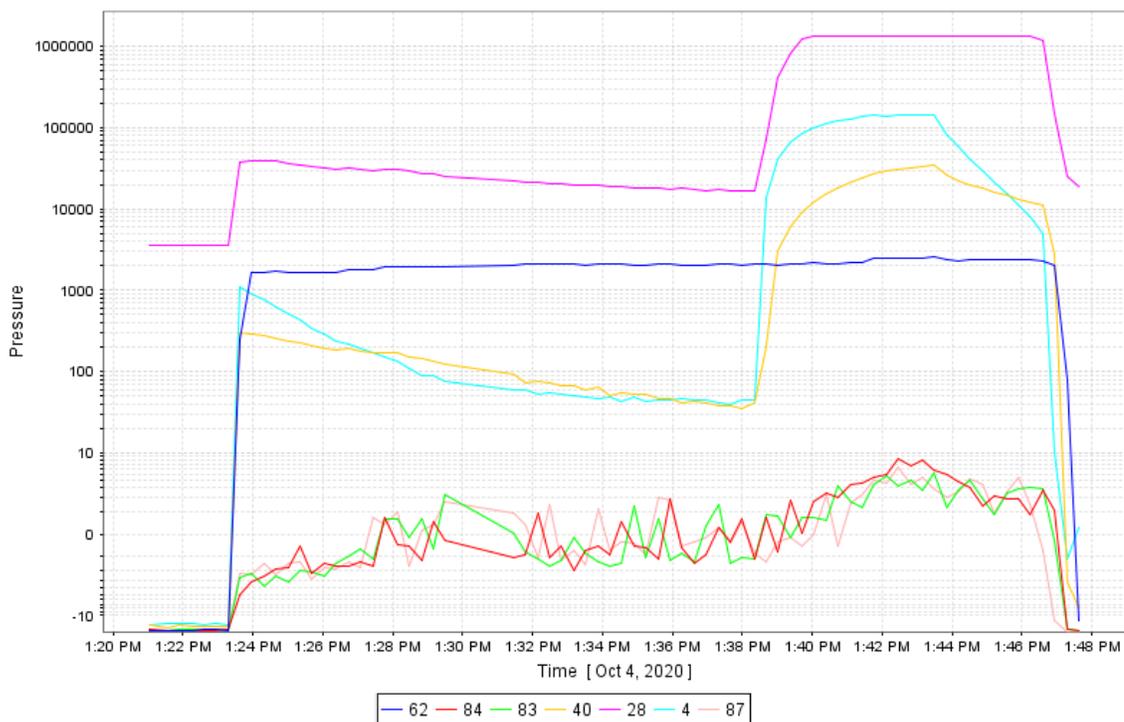


Figure 4.14: An RGA pressure versus time scan. The different line colors correspond to mass-to-charge ratios of a specific gas. The nitrogen (pink) signal saturated during flow, beginning at 1:38 pm. As a result the signals in all channels increased. The ^{62}Xe (blue) channel rose by approximately 30% despite being the dominant pressure and known to be constant. This increase in the ^{62}Xe remained high after flow stopped at 1:43 pm, as noted by the downturn in helium and argon signals. The pathology is observed with ^{87}Br (black) and produced a ^{87}Br signal that was over 30% the size of the signal observed in the ^{84}Kr channel (red).

We can avoid both types of saturation by noting the initial signal jump when exposing the RGA to the gasses accumulated during formation of the xenon ice kernel. If the signal is excessively large, additional preparation steps can be taken to either dilute the sample or remove the problematic impurity. Ideally we select a flow and impedance state that offers

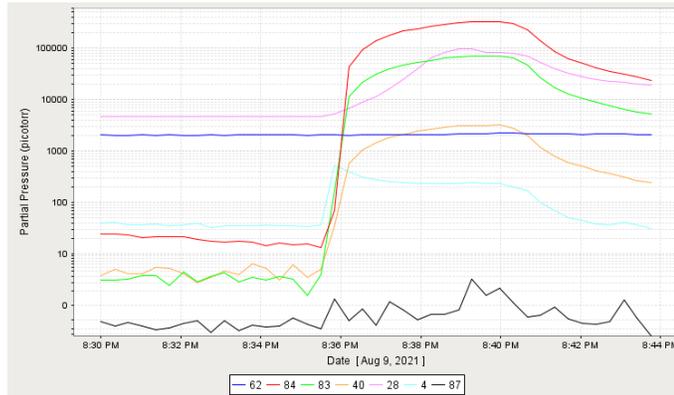
sufficient sensitivity to krypton while avoiding saturation from other gasses.

4.4.2 System Contamination

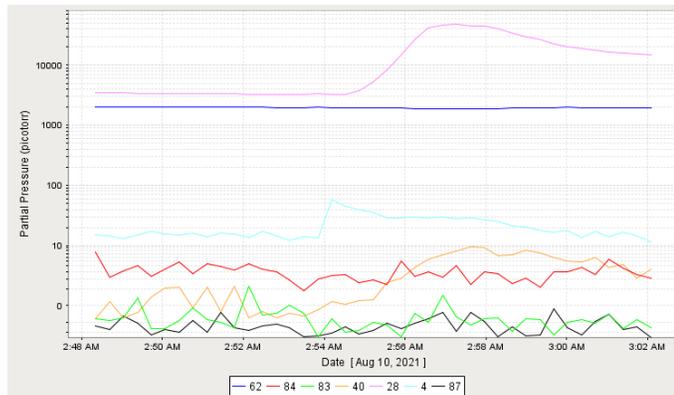
When the system or RGA is exposed to air, it is crucial to maintain a clean work environment and wear gloves. Oil or hydrocarbon contamination, particularly near the RGA, present a danger to high-sensitivity krypton measurements. If an irreducible background is identified and the system does not have a leak to air this is likely due to such contamination. Contamination manifests as an irreducible signal when measuring xenon that is known to be clean or when the concentration of a contaminant does not reduce after consecutive xenon clean runs, as specified in the “Clean Run” procedure, section 4.5.9. While some gas baselines will be higher than others, a key observation that contamination is an issue is the divergence of ^{83}Kr and ^{84}Kr from the baseline ^{87}Br signal. An example of signal divergence is shown in figures 4.15a and 4.15b, where the ^{84}Kr signal is higher than the ^{87}Br channel prior to flow. Figure 4.15a shows the contamination event, where a xenon sample with 20 ppb of krypton, contained enough krypton that it was observable even without flow. An operator had bypassed ^{84}Kr saturation check resulting in contamination of the RGA and valve V6. Figure 4.15b shows a measurement of xenon known to be clean, but with an elevated ^{84}Kr channel. The observed krypton is due to out-gassing following the contamination event.

Hydrocarbon contamination produces a similar signal divergence as a krypton contamination event, but in multiple mass channels. When hydrocarbon molecules approach the RGA they are bombarded by electrons emitted in the ionizer. These electrons can break up the hydrocarbon and produce fragmentation patterns. These patterns produce elevated signals across almost every channel and in particular a signal in the ^{84}Kr and ^{83}Kr channels. An analog scan showing in figure 4.16 shows elevated signals across the entire scan and in

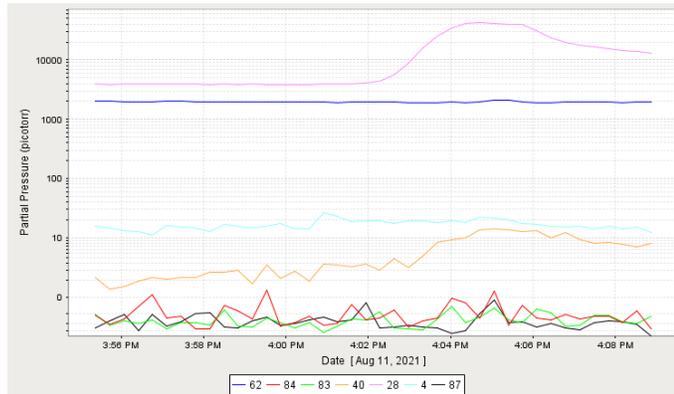
particular between 70 and 85 AMU when compared with figure 4.12.



(a)



(b)



(c)

Figure 4.15: Three RGA pressure vs time scans. In figure 4.15a a xenon sample containing 20 ppb of krypton causes an elevated background in the red ^{84}Kr and green ^{83}Kr channels. The measurement of this dirty sample resulted in contamination. The contamination is shown in figure 4.15b, through the continued divergence of ^{84}Kr from ^{87}Bl during a measurement of xenon that is known to contain less than 10 ppq of krypton. After two days of maintenance, the XPM baseline was recovered, as shown in figure 4.15c. The ^{84}Kr and ^{87}Bl remained stable and near zero pTorr throughout the measurement of the purified xenon. All measurements were carried out under identical settings and followed the procedure given in section 4.5.11.

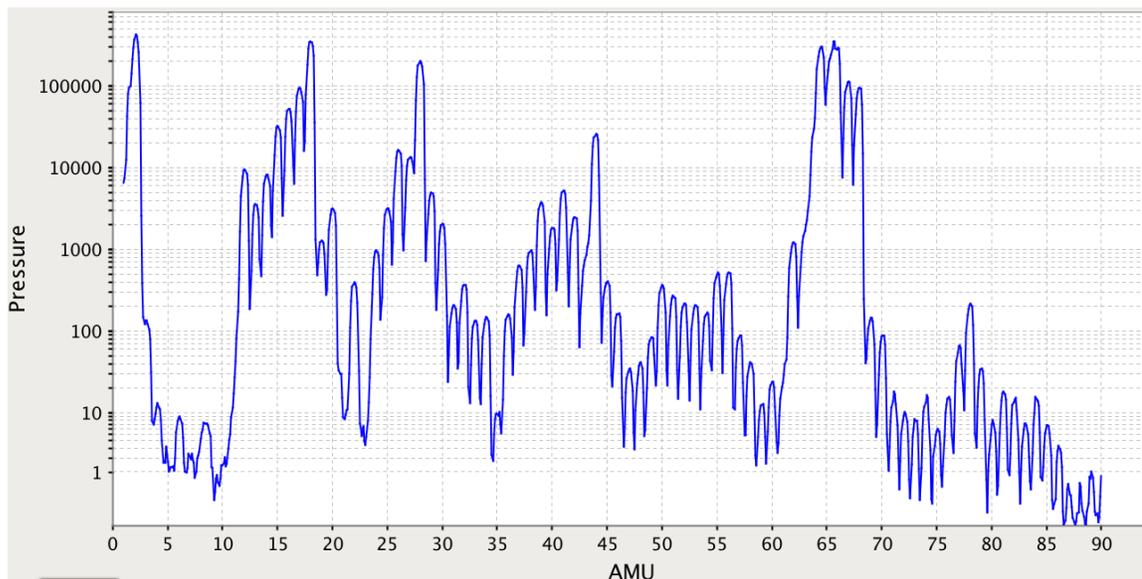


Figure 4.16: An analog scan of xenon ice. The RGA and vacuum components were contaminated by an unknown hydrocarbon that resulted in an elevated signals across all AMU channels. Most notably, when compared with figure 4.12, we see elevated signals between 45 to 60 AMU and 70 to 85 AMU. This contamination produces similar divergence in the ^{84}Kr and $^{87}\text{B1}$ channels and impedes the XPM measurement capability.

Bake-outs of plumbing and the RGA during a prolonged pump-out with heater tape can help reduce background signals from out-gassing components or residual contaminants. An operator must not exceed the max temperature of any component and should never leave a bake-out unattended. If V7 has been contaminated due to krypton exposure, it should be left open and exposed to the utility pump and RGA, while carefully baked to a low temperature. Overheating V7 can cause deformation of the valve seat and put the RGA at risk.

It can be difficult to achieve and maintain sufficient cleanliness at the RGA. In general if a system is being used for multiple measurements a week, it is best to maintain vacuum at the RGA. The filament can be left on and CEM voltage set at high impedance settings assuming the RGA was previously conditioned. This will help reduce the backgrounds from any minor contaminants introduced by unexpectedly dirty measurements. Under extreme

cases the filament can be degassed using a serial command to the ECU. Sending DG##, where ## is some value in minutes, will run extra current through the filament in attempt to burn off any extra contaminants. Degas cycles should be run while the CEM is off and can shorten the lifetime of the filament.

4.5 Operating Procedures

The following procedures outline basic operations required to commission and operate an XPM. The volumes and values used are from the mixing panel. In any XPM, equivalent spaces and procedures should be defined.

Name	Volume (cc)	description
V_{SB}	3810 ± 110	source bottle
V_{PT}	80 ± 3	PTO space
V_{CT}	200 ± 10	cold-trap

Table 4.1: The mixing panel volumes.

The initial source bottle volume is known to within 3% and introduces a systematic uncertainty to subsequent volume measurements. Other volumes are found by expanding V_{BTL} to an evacuated V_{PT} and V_{CT} , then applying $P_i V_i = P_f V_f$. This measurement procedure is repeated to account for statistical uncertainties. The error from each volume measurement is propagated to the next. Each pressure measurement includes a systematic 0.5% uncertainty. Moving forward these volumes represent a source of systematic uncertainty when mixing calibration and measuring xenon samples.

4.5.1 Safe State

First we define a safe state to ensure that all vacuum components are isolated and any xenon ice formed within the system can safely expand to a sufficiently sized buffer volume. The system should be left in this state between measurements. Normally-open and closed pneumatic valves should be chosen to match the safe state. If a control or pneumatic supply issue arises, the pneumatic control manifold can be isolated and relieved of pressure to force a safe state.

1. Close SV, V5 and V6. Isolate the XPM and vacuum components.
2. Open V0, V1, V2, V3, and V4. Set a path from each side of the cold trap directly to a buffer volume.

4.5.2 Initial State

If the system just been installed or opened to air, it must be pumped to ultimate vacuum of $< 10^{-7}$ Torr and leak checked. If the system is between xenon measurements it should be evacuated to < 10 Torr. Usually a cryopump is used to recover xenon in a devoted vessel. If the xenon previously measured is dirtier than a calibration mixture, the system should be pumped to $< 10^{-4}$ Torr after cryopumping to avoid cross contamination.

1. The system is in the safe state 4.5.1.
2. The MFC set point is 0 slpm.
3. The RGA has been previously conditioned as specified in 4.3.4.
4. The RGA filament and CEM are off.
5. TP1 and TP2 are on.

6. PT1 and PT2 are on and read $< 5 \times 10^{-8}$ Torr.

4.5.3 Collect Xenon

Two fill methods are described below. The XPM will be filled with an 80 gram xenon sample, equivalent to 2500 Torr throughout all volumes. The first method simply fills the XPM from a regulated xenon source beyond SV. The line beyond SV should be evacuated between transfers to limit xenon cross contamination.

1. The XPM is in the initial state 4.5.2.
2. Close V0, V1, V2, and V3 to isolate the PT0 space.
3. Adjust the xenon source regulator to 2500 Torr.
4. Open SV and confirm the pressure is 2500 Torr with PT0.
5. Open V0, V1, V2, and V3 to fill the system.
6. Once PT0 reads 2500 Torr the system is full and contains approximately 80 grams of xenon.
7. Close SV.
8. The system is now in the safe state 4.5.1.

If the xenon source pressure is too low to fill the system directly, the cold-trap can be used to cryo-pump xenon from the source while the MFC regulates the flow. For our purposes we will assume the liquid nitrogen dispensing is handled by an automated fill system capable of maintaining the liquid level to within $\pm 1''$ of the initial fill height so the stocking portion of the cold trap remains fully submerged at all times. This procedure

requires the supply xenon is at least 100 Torr greater than the minimum MFC backing pressure.

1. The XPM is in the initial state 4.5.2.
2. Close V0, V1, V2, V3, and V4.
3. Open SV.
4. PT0 must read at least 100 Torr greater than the MFC minimum operating pressure. In our case the MFC fails to regulate flows higher than 2 slpm around 300 Torr back pressure, so we require that PT0 read at least 400 Torr.
5. Calculate the flow interval from the amount of xenon needed $t_{flow} = \frac{M_{Xe}}{\rho_{Xe}Q}$, where $\rho_{Xe} = 5.89$ grams/liter.
6. We will use a flow rate of 2 slpm over a 6 minute 45 second interval to deliver 80 grams of xenon. If the trap forms an ice plug try increasing the flow rate and reducing the flow interval accordingly during the next attempt.
7. Submerge the cold-trap in liquid nitrogen and wait 2 minutes for boiling to stop.
8. Open V2.
9. Set the MFC to 2 SLPM.
10. Close SV after t_{flow} time has passed.
11. Set the MFC to 0 SLPM.
12. Set the safe state 4.5.1.
13. Warm-up to cold trap 4.5.4.

4.5.4 System Warm-up

After any instance of xenon ice formation the cold-trap must be safely warmed to room temperature. This applies no matter the quantity of xenon in the cold-trap. Failure to properly warm up presents a pressure hazard that can damage equipment or harm operators.

1. Set the safe state 4.5.1.
2. Lower the liquid nitrogen dewar from the cold-trap.
3. As the cold-trap warms, monitor the rise of pressure in PT0. The pressure will spike around 800 Torr as the xenon undergoes a phase change. This may cause audible rattling of ice in the cold-trap.
4. A water bath, heater, or fan can be used to speed up warm-up if needed.
5. Once warm the pressure, as read by PT0, will stable at 2500 Torr.

With a xenon sample now in the XPM a measurement can be completed. In our case the system has not yet been configured, so we must set the desired impedance state and prepare a calibration.

4.5.5 Form an Ice Kernel

We start by forming a xenon ice kernel that is later exposed to the RGA to monitor the xenon ice vapor backgrounds.

1. The sampling system is in the safe state 4.5.1, warm, and contains 80 grams of xenon.
2. Close V0.

3. Submerge the cold-trap in liquid nitrogen and wait 2 minutes for boiling to stop and PT0 reads 0 Torr.
4. Close valves V1, V2, V3, and V4.

The cold-trap now contains 5.5 grams of xenon ice.

4.5.6 Setting the 1x Impedance State

To set the 1x impedance we start with both VV and IV maximally opened with the space between IV and V6 previously evacuated. We then open the RGA to the ice kernel and minimally adjust IV or VV to achieve the impedance condition described by equation (4.19).

1. Turn on PT2,. It should read $< 5 \times 10^{-8}$ Torr. It may need up to 30 minutes warm-up depending on the model.
2. Form an ice kernel 4.5.5.
3. Open V6 to expose the RGA volume to the xenon ice.
4. Allow PT2 to settle.
5. If PT2 reads a pressure $> 1.0 \times 10^{-5}$ adjust IV in the closed direction until PT2 is between 5.0×10^{-6} and 1.0×10^{-5} Torr.
6. If PT2 reads a pressure $< 5.0 \times 10^{-6}$ adjust VV in the closed direction until PT2 is between 5.0×10^{-6} and 1.0×10^{-5} Torr.
7. Turn on the RGA filament and let the pressure settle. Check that the pressure remains within the appropriate range. Fine tune the previously adjusted valve as needed.

8. Only one valve should be adjusted when setting the 1x state from the default state.
9. Close V6 to isolate the RGA.

The current state is equivalent to the final state after forming an ice kernel 4.5.5, but with the 1x impedance state set. The system can be used as is or set to a higher Yx impedance state. At 1x impedance the system will have sufficient sensitivity to measure ppq concentrations of krypton.

4.5.7 Set the RGA High Voltage

The RGA CEM high voltage should be set prior to each measurement and prior to setting the Yx impedance state. To set the CEM high voltage we ^{62}Xe channel. The ideal CEM high voltage achieves a ^{62}Xe partial pressure of 2000 ± 100 pTorr. Over time the high voltage required to achieve the proper xenon partial pressure will increase as the CEM degrades. Once the required high voltage increases by 100 volts the system should be re-calibrated at the new high voltage value.

1. Form an ice kernel 4.5.5.
2. Set the 1x 4.5.6 or Yx 4.5.8 impedance state.
3. Turn on the filament and set the CEM high voltage to the most recently used value.
If no CEM high value has been established or the CEM was recently serviced start at 1000 volts.
4. Start a pressure versus time scan monitoring at least channels ^{62}Xe , ^{84}Kr , and ^{87}Bl .
Set the noise floor to 0.
5. Open V6 to expose the RGA to the xenon ice.

6. Take an analog scan between 60 and 90 AMU with the noise floor set to 0 and a precision of 10 points per AMU.
7. Once the scan completes wait for the signal to settle then take an average of the next 3 PvT data points the mass channels of interest.
8. Skip the high voltage setting procedure if the average partial pressures are in the acceptable ranges:
 - ^{62}Xe 2000 ± 100 pTorr.
 - ^{87}Br 0 ± 5 pTorr.
 - ^{84}Kr 0 ± 5 pTorr.
9. If not adjust the high voltage as follows:
 - (a) Use increments of ± 11 volts to adjust the high voltage until the ^{62}Xe partial pressure is 2000 ± 100 pTorr.
 - (b) Take an analog scan between 60 and 90 AMU with the noise floor set to 0 at the new high voltage.
 - (c) Once the scan completes allow the signal to settle and average the next 3 PvT data points for ^{62}Xe , ^{87}Br , and ^{84}Kr .
 - (d) Repeat if they are not in the appropriate range.
 - If the partial pressure conditions noted above cannot be achieved there is likely a contamination issue, the CEM is degraded, or peaks have shifted. When possible re-condition the RGA, bakeout, degas, or tune the RGA.
 - If the xenon pressure is unstable or increasing, isolate the RGA and check that the liquid nitrogen level is properly maintained.

- If the problem persists try a different RGA assembly. The current RGA may need to be serviced or is simply a dud. Bad RGAs can still prove useful elsewhere, but not in this system. It is not worth the headache.
- As long as the ^{62}Xe partial pressure is stable, a measurement can be taken despite elevated ^{84}Kr and ^{87}Bl backgrounds. The result will be less sensitive and the errors will be poorly quantified. If this is a new behaviour and you must operate under these conditions, calibrate immediately after the measurement and retroactively apply that calibration.

10. If stable ^{62}Xe , ^{84}Kr , and ^{87}Bl values are achieved, the RGA high voltage is now set. The final values and high voltage should be noted.

11. Close valve V6 to isolate the RGA.

The current state is equivalent to the final state after forming an ice kernel 4.5.5, but with the appropriate high voltage set.

4.5.8 Setting the Yx Impedance State

Higher impedance settings will be more sensitive to impurities while also making the system more susceptible to electronic saturation. An impedance setting near 15x offers sufficient sensitivity to krypton concentrations around 100 ppq, while minimizing the amount of preparation steps required to manage saturation.

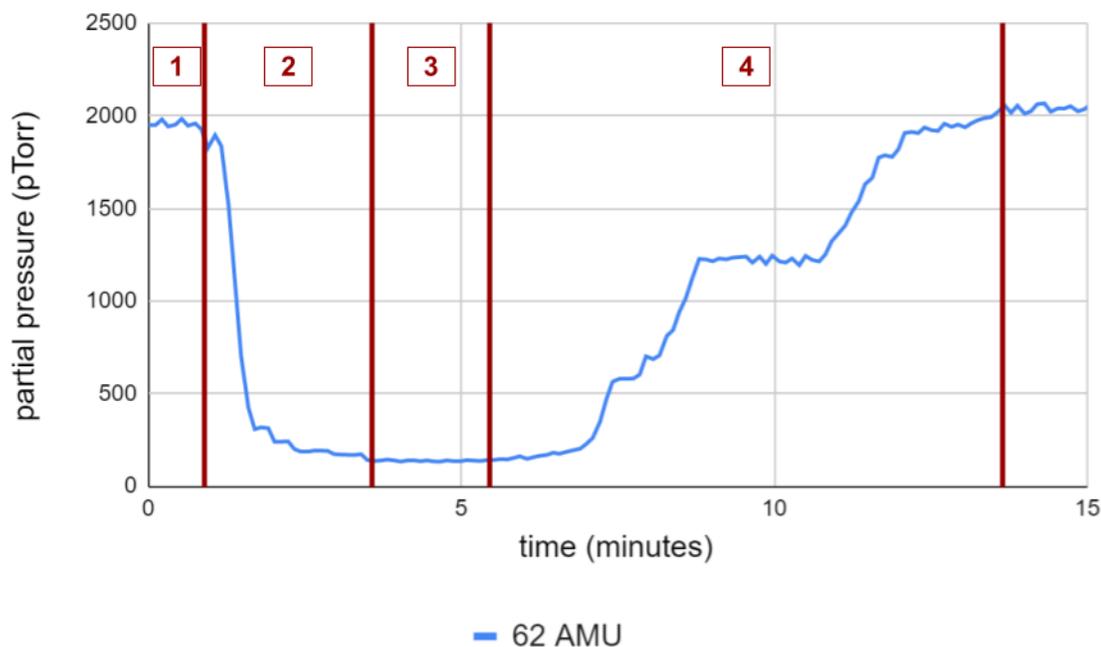


Figure 4.17: An RGA pressure versus time (PvT) scan of the ^{62}Xe partial pressure. The impedance state was previously set to the 1x state in region 1. The average ^{62}Xe partial pressure is 1960 pTorr after setting the CEM high voltage. In region 2, IV is adjusted to increase impedance. After IV is adjusted, the average ^{62}Xe partial pressure in region 3 is 137 pTorr. The resulting impedance value is $1960/137 = 13.4x$. In region 4, VV is adjusted. This brings the ^{62}Xe partial pressure back up to the operating spec of the RGA. A final ^{62}Xe value of 2040 pTorr is achieved.

Figure 4.17 shows the ^{62}Xe partial pressure during IV and VV adjustment.

1. Form an ice kernel 4.5.5.
2. Set the 1x state 4.5.6.
3. Open V6 to expose the RGA to the xenon ice.
4. Set the RGA high voltage 4.5.7. During this procedure, only ^{62}Xe needs to be monitored.
5. Take an analog scan between 60-90 amu. Use a noise floor of 0 and 10 points per amu.

6. Use a PvT scan to monitor ^{62}Xe partial pressure. Use a noise floor of 0.
7. Record the average ^{62}Xe partial pressure as shown in region 1 of figure 4.17.
8. In region 2, adjust IV the closed direction until $PP_{^{62}\text{Xe}}$ is reduced by a factor of Yx. This can be very challenging at high impedance settings, as small adjustments result in large pressure swings.
9. Record the average ^{62}Xe partial pressure as shown in region 3 of figure 4.17.
10. In region 4, adjust VV in the closed direction until ^{62}Xe partial pressure is restored to the initial value.
11. Close V6.

The current state is equivalent to the final state after forming an ice kernel 4.5.5, but with the Yx impedance state set. If VV or IV are altered, the system should be reset to 1x impedance from the default state before setting Yx again. With the final Yx impedance set, the system should be calibrated. To do so we must prepare ultra-pure xenon and identify a krypton reference source.

4.5.9 Clean Run

As noted in the α parameter study 4.3, the gasses detected in the RGA are effectively removed from the xenon. This allows the XPM system to produce ultra-clean xenon, with krypton concentrations < 1 ppq. This xenon can be mixed with other gases to produce dilution mixture. For consecutive measurements taken at 13.4x impedance the krypton concentration was reduced by a factor of 0.7 or $\alpha = 0.3$. To increase the cleaning factor, we rely on the utility pump (TP1) behind valve V5 to increase the pumping speed at the cold-trap outlet. This increases the cleaning factor to 0.3 or $\alpha = 0.7$.

1. Form an ice kernel 4.5.5.
2. Open V5 to the utility pump.
3. Set the MFC to 0 SLPM.
4. Open V0 and V2.
5. Set the MFC to 2 SLPM.
6. When $PT0 < 300$ Torr or when the MFC is no longer able to maintain constant flow, open V1.
7. Wait 4 minutes while the source bottle is exhausted.
8. If the pressure at PT0 does not go to 0 Torr, it is likely that an ice plug formed. Modify future procedures to use a higher flow rate to avoid ice plug formation.
9. Close V5.
10. Set the safe state 4.5.1.
11. Warm-up the cold-trap 4.5.4.

The cleaning procedure should be completed as needed until the desired cleaning factor is achieved. The clean xenon is now ready to be mixed.

4.5.10 Calibration Preparation from Reference Xenon

In reference [12], section 3.6, an auxiliary system is used to produce xenon mixtures from a pure krypton source and purified xenon. This is done on devoted hardware separate from the XPM as it poses a high contamination risk. The procedure involves expanding

krypton from an initial volume to an evacuated volume to control the total amount of krypton added. This method is used to produce xenon “cocktail” mixtures that range in purity from 1 ppb to 1 ppm. Small amounts of this cocktail mixture are further diluted with pure xenon to produce “reference” xenon, that contain krypton concentrations between 100 ppt to 1 ppb.

Reference xenon is kept readily available in the XPM. When a calibration is required, a small slug of reference is delivered to the XPM and diluted with ultra-pure xenon one final time. The resulting krypton concentration of the calibration mixture is usually between 1 and 10 ppt. At each stage of dilution, the xenon used to dilute the preceding mixture should be purified to have less than 1% of the subsequent target mixture’s concentration. For example if a calibration mixture of 1 ppt is made, xenon should be purified to a krypton concentration less than 10 ppq before introducing the reference xenon for dilution. While cocktail and reference xenon can be prepared and stored long term, calibration xenon is susceptible to krypton ingress from air and should be used within a few days of preparation.

Our procedure assumes that a reference source was identified as described in section 4.7.2 and figure 4.26. The reference concentration is known to be $\Phi_{ref} = 217 \pm 12$ ppt and readily sourced to the XPM through SV by an adjustable gas regulator.

To produce a calibration mixture we start by recording the initial xenon pressure $P_{Xe} = 2500$ Torr and volume $V_{Xe} = V_{SB} + V_{PT} + V_{CT}$. We form an ice kernel and fill $V_{ref} = V_{PT}$ with $P_{ref} = 600$ Torr of reference xenon. The resulting calibration concentration is

$$\Phi_{cal} = \Phi_{ref} \frac{P_{ref} V_{ref}}{P_{Xe} V_{Xe} + P_{ref} V_{ref}} \approx \Phi_{ref} \frac{P_{ref} V_{ref}}{P_{Xe} V_{Xe}} = 217 \text{ ppt} \frac{600 \text{ Torr} \cdot 80 \text{ cc}}{2500 \text{ Torr} \cdot 4090 \text{ cc}} = 1.02 \text{ ppt} \quad (4.25)$$

with a corresponding systematic percent uncertainty of

$$\frac{\sigma_{\Phi_{cal}}}{\Phi_{cal}} = \sqrt{\left(\frac{\sigma_{\Phi_{ref}}}{\Phi_{ref}}\right)^2 + \left(\frac{\sigma_{P_{ref}}}{P_{ref}}\right)^2 + \left(\frac{\sigma_{V_{ref}}}{V_{ref}}\right)^2 + \left(\frac{\sigma_{P_{Xe}}}{P_{Xe}}\right)^2 + \left(\frac{\sigma_{V_{Xe}}}{V_{Xe}}\right)^2} = 7.1\% \quad (4.26)$$

This procedure can also be used to produce dilution mixtures by simply replacing the reference xenon injected to V_{PT} with the xenon sample of interest. The dilution factor is given by $(\rho_{gas}P_{sam}V_{sam})/(\rho_{Xe}P_{Xe}V_{Xe})$, where ρ_{gas} is the density of the diluted gas. This is useful when measuring xenon of unknown origin or producing a low-sensitivity calibration mixture by injecting dry air into V_{PT} , as described in section 4.7.2.

Calibration Preparation Procedure

1. Record P_{Xe} and V_{Xe} . This can also be recorded after the final warm-up step.
2. Form an ice kernel 4.5.5.
3. Close V1, V2, and V3.
4. Set the reference xenon regulator to desired P_{ref} . The line from the reference xenon to SV should have been previously evacuated.
5. Open SV to fill V_{PT} to P_{ref} .
6. Close SV when PT0 reads the desired pressure.
7. Record P_{ref} .
8. Open V0, wait 30 seconds and close V0.
9. Open V1, V3 and wait 30 seconds.
10. Close V1, open V0.
11. Close V0, open V1, wait 30s. This sequential sweep out of V_{PT} ensures the xenon is properly mixed.
12. Set the safe state and exhaust any remaining xenon into the cold-trap 4.5.1.

13. Warm-up the cold trap 4.5.4.

The calibration xenon is now prepared and ready to be measured.

4.5.11 High-Sensitivity Measurement Procedure

With the desired xenon sample in the system and the system fully configured we can finally take a high-sensitivity xenon purity measurement.

1. Form an ice kernel 4.5.5.
2. Set the MFC to 0.
3. Open V0.
4. Turn on the RGA filament and set the CEM high voltage to the most recently used value. This is not the final high voltage setting used in measurement but will be used identify potential saturation signals when the RGA first opens to the xenon ice.
5. Configure the PvT scan to monitor species of interest. This usually includes
 - Helium: 4 AMU
 - Methane: 16 AMU
 - Nitrogen: 28 AMU
 - Argon: 40 AMU
 - Xenon: 62 AMU
 - Krypton: 84 AMU
 - Krypton: 83 AMU
 - Baseline: 87 AMU

6. Set the noise floor to 0.
7. Begin the pressure versus time scan.
8. Wait 30s
9. Open V6 to expose the xenon ice kernel to the RGA space.
10. Wait for at least 3-4 data points to be collected (about 90 seconds) while monitoring saturation:
 - Helium: 2,000 pTorr
 - Nitrogen: 100,000 pTorr
 - Argon: 2,000 pTorr
 - Xenon: 4,000 pTorr
 - Krypton: 15 pTorr
 - Baseline: 15 pTorr
11. If an initial signal is likely to cause saturation, stop the measurement procedure, isolate the RGA, and switch to the static pump-out procedure starting after the ice formation step 4.5.12.
12. Set the RGA CEM high voltage 4.5.7 from steps 6 through 10.
13. Collect backgrounds for 5 minutes.
14. Wait for 1 data point to pass.
15. Open V2.
16. Set the MFC to 2 slpm. The pressure monitored by PT0 will start to drop linearly.

17. Flow for 5 minutes.
18. Close V2 and set the MFC to 0 SLPM.
19. Wait for 3 minutes as RGA backgrounds are re-established.
20. Close V6 to isolate the RGA
21. Set the safe state 4.5.1.
22. Warm-up the cold trap 4.5.4.

A low-sensitivity procedure is used to measure helium, nitrogen, and argon signals that would otherwise cause electronic saturation during a high-sensitivity measurement. The low-sensitivity measurement procedure follows the same steps with the exception that the RGA high voltage is set to 0, the HV tuning is skipped, and the flow is adjusted. Saturation is instead monitored by PT2 and the pressure is required to remain below 1.5×10^{-5} Torr. The flow rate is reduced to deliver less than 6 grams of xenon over a 5 minute interval by either the LV or the MFC if it is capable of stable flow near 0.1 slpm. If a full suite of measurements is required, an operator should first perform a low-sensitivity measurement, followed by a static pump-out, and finally a high-sensitivity measurement.

4.5.12 Static Pump-out Procedure

Static pump-outs are used to reduce helium, nitrogen, and argon concentrations to non-saturating signal sizes prior to a high-sensitivity krypton measurement. The krypton concentration during a static pump-out in the XPM cold-trap is reduced by less than 10% while helium concentration is reduced by a factor of at least 500, nitrogen a factor of at least 10, and argon at least a factor of approximately 5. This is achieved by flowing the entire xenon sample into the cold-trap without any pump exposed until all of the ice is formed. This ef-

fectively minimizes the krypton throughput parameter α by setting an infinite impedance, as implied by figure 4.3.

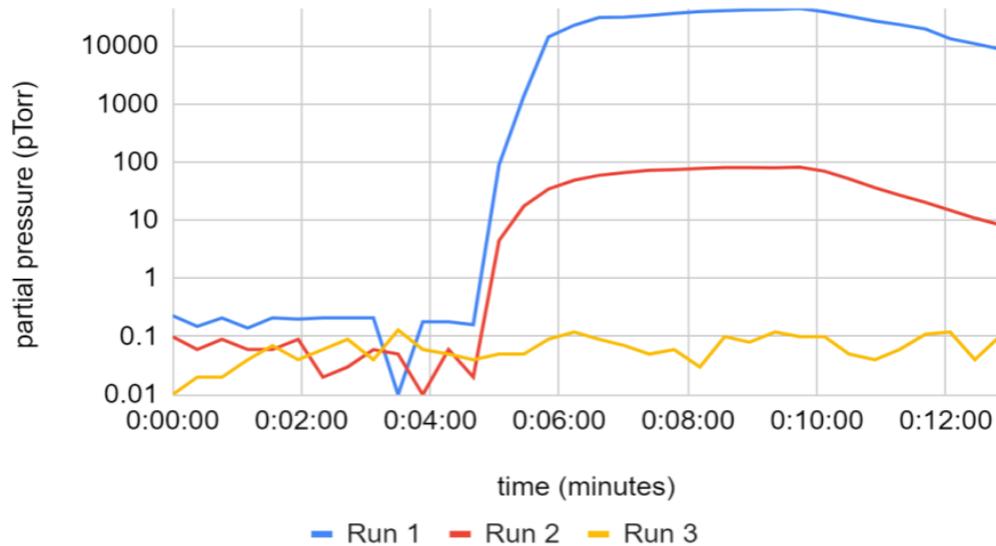


Figure 4.18: Helium measurements for three runs with a static pump-out after Run 1 and Run 2. The helium in Run 1 (blue) caused the RGA to saturate during this measurement giving a 2.8 ppm a lower limit. Following one static pump-out Run 2 (red) measured a helium concentration of 5.45 ppb of helium. Following a second static pump-out the helium in Run 3 (yellow) was undetectable.

1. Form an ice kernel 4.5.5.
2. Set the MFC to 0 SLPM.
3. Open V0 and V2.
4. Set the MFC to 2 slpm.
5. When $PT0 < 300$ Torr or when the MFC is no longer able to maintain constant flow, open V2.
6. wait 4 minutes. If the pressure at PT0 does not go to the expected baseline value, it is likely that an ice plug formed. To avoid this in the future, use a higher flow rate.

7. Open V5 to pump on the cold-trap outlet with the utility pump.
8. Wait 10 minutes.
9. Close V5.
10. Open V4 and wait 90 seconds.
11. Open V5 to pump on the cold-trap inlet and source bottle
12. Wait 5 minutes.
13. Close V5.
14. Set the safe state 4.5.1.
15. Warm-up the cold-trap 4.5.4.

4.6 Measurement Analysis

Two analyses methods corresponding to the low and high-sensitivity measurement procedures are used to observe xenon impurities. Both methods monitor xenon ice backgrounds prior to flow and define a signal interval relative to the flow start and stop times. The low-sensitivity analysis method uses a background-subtracted average signal, while the high sensitivity method uses a background subtracted integrated signal. The exact background, flow, and signal intervals vary between systems because the specific response times τ depends on the impedance setting and the cold-trap volume. We will continue to draw on the values used in the mixing panel for our description of an XPM analysis.

4.6.1 Low-Sensitivity Analysis

The background interval is defined to be five minutes of stable RGA traces prior to flow start. A signal interval is defined to be the last two minutes of the five minute flow interval. The signal of a mass channel is given by the average pressure during the signal interval minus the average background pressure multiplied by a gas specific, low-sensitivity calibration parameter C_{low} . The calibration parameter is calculated for a run completed under identical conditions where the concentration of the calibration gas Φ_{cal} , is known.

$$\Phi_{cal} = C_{low}(\overline{PP}_{sig} - \overline{PP}_{bkg}) \quad (4.27)$$

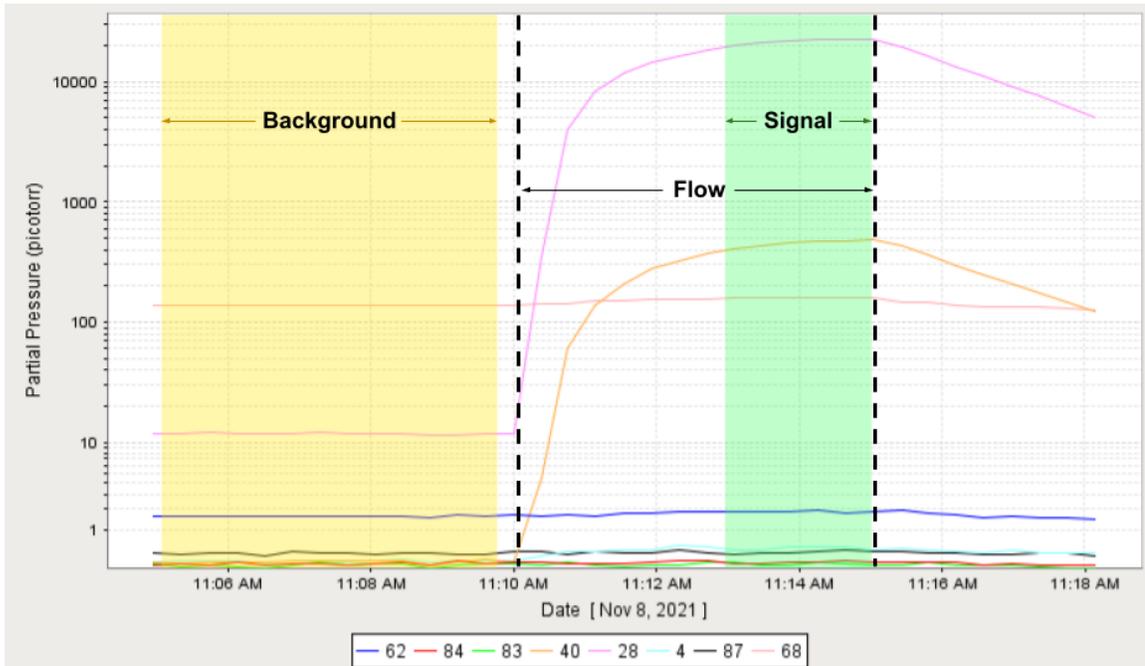


Figure 4.19: A measurement of calibration xenon prepared with impurities from a small quantity of air. The measurement was completed at 13.4x impedance with a flow rate of 0.13 slpm, maintained between the two dotted black lines. A 5 minute background interval is highlighted in yellow and the signal window is highlighted in green. The prepared xenon mixture contained 1.43 ± 0.14 ppm of nitrogen and 24.5 ± 2.5 ppb of argon. The nitrogen and argon concentrations were reported as 1.48 ± 0.44 ppm and a 31 ± 9 ppb respectively.

The error of a single concentration measurement is given by the standard deviation of the signal combined with the error from the calibration parameter. The percent error of the calibration parameter is equal to the percent error of the prepared calibration mixture, $\sigma_C/C = \sigma_{\Phi_{cal}}/\Phi_{cal}$, because the mixing error is usually large compared to the random error during the calibration measurement.

$$\sigma_{\Phi_{low}} = \sqrt{(C_{low}\sigma_{sig})^2 + (\Phi_{low}\sigma_{C_{low}})^2} \quad (4.28)$$

Ideally calibrations are completed frequently, as the system response will change and these changes are not captured in a single calibration. If many calibrations are carried out over a regular interval the system response can be classified into periods of similar behavior where the calibration parameter remains fixed and a fit can be made against these parameters to better characterize the systems response and errors. In practice, calibrations hinder system throughput. For LZ, the helium, nitrogen, and argon impurities were reduced to sufficiently small concentrations that extensive characterization of low-sensitivity system measurements was not necessary. Only three low-sensitivity calibrations were taken over a multi-year period spanning different system configurations. The reported concentrations agreed to within 30% of the calibrated values across all configurations and a 30% total error was taken for all reported low-sensitivity measurements.

4.6.2 High-Sensitivity Analysis

During a high sensitivity analysis, the background and flow intervals remain the same, but the signal interval is delayed from the flow interval and delayed. The delayed signal interval is included to account for the rising and falling edges of the low concentration krypton signal which is unique to each impedance configuration. An integrated analysis method is used instead of an averaged analysis as the equilibrium pressure of low concen-

tration samples is not achieved. For the mixing panel, the signal start is delayed by 20 seconds from flow start and signal stop is delayed by 50 seconds from flow stop. The concentration of krypton is given by integrating the background-subtracted ^{84}Kr partial pressure of N discrete signal data points $PP_{sig,i}$, multiplied by a gas-specific, high-sensitivity calibration parameter C_{high} . The calibration parameter is calculated for a run completed under identical conditions where the concentration of the calibration gas Φ_{cal} , is known.

$$\Phi = C_{high} \cdot S_{int} = C_{high} \sum_{i=1}^N (PP_{sig,i} - \overline{PP}_{bkg}) \Delta t_i. \quad (4.29)$$

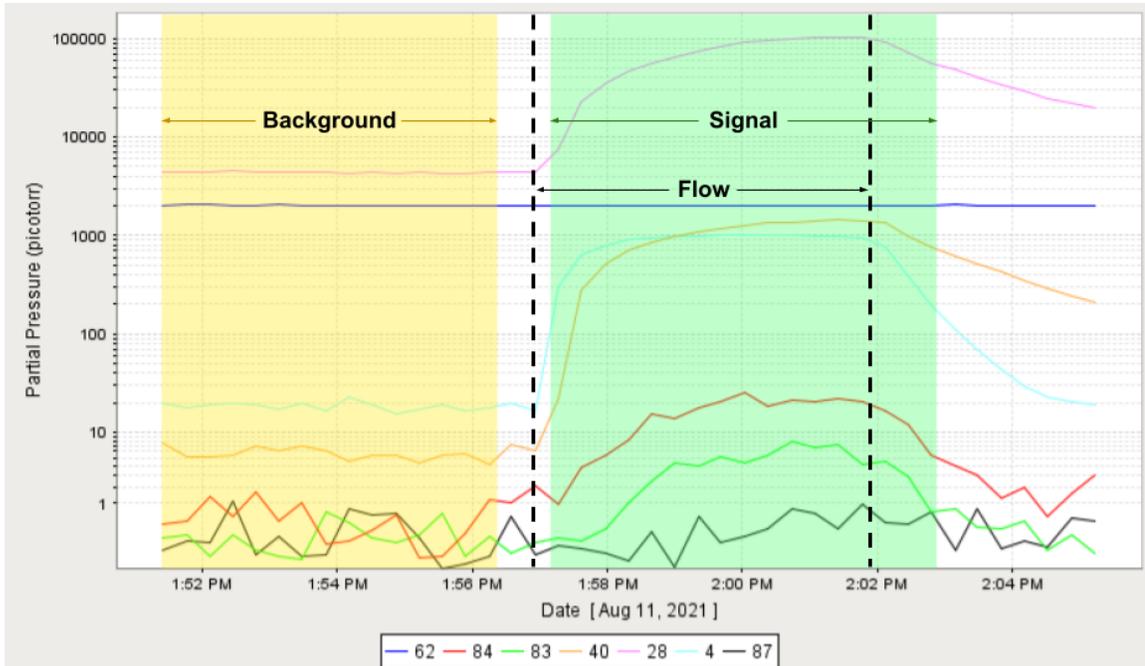


Figure 4.20: A measurement of calibration xenon prepared to a krypton concentration of $\Phi_{cal} = 1.17$ ppt. The measurement was completed at 13.4x impedance and a flow rate of 2 slpm. The signal window begins 20 seconds after the flow starts and ends 50 seconds after the flow stops.

The error for a single measurement is given by the estimated random error of the discrete partial pressure values combined with the percent error from the calibration parameter. Again, the percent error of the calibration parameter is equal to the percent error of the pre-

pared calibration mixture, $\sigma_C/C = \sigma_{\Phi_{cal}}/\Phi_{cal}$, as the mixing error is large compared to the estimated random error during the calibration measurement at concentrations around a ppt. An estimated random error is used since an equilibrium partial pressure is not achieved in small concentration samples. The random error of each discrete signal data point is scaled from the standard deviation observed in the background interval, as done in reference [12].

$$\sigma_i = \sigma_{bkg}(0.009 \text{ pTorr}^{-1} \cdot PP_{sig,i}) \quad (4.30)$$

The error σ_{int} from the integrated signal S_{int} as a percent is

$$\frac{\sigma_{int}}{S_{int}} = \frac{\sqrt{\sum_{i=1}^N (\sigma_i^2 + \sigma_{bkg}^2/N_{bkg}) \Delta t_i^2}}{\sum_{i=1}^N (PP_{sig,i} - \overline{PP}_{bkg}) \Delta t_i} \quad (4.31)$$

where N_{bkg} is the number of data points in the background interval. The percent signal error is then combined with the percent calibration parameter error from giving a final measurement concentration error of

$$\sigma_{\Phi_{high}} = \Phi_{high} \sqrt{\left(\frac{\sigma_{int}}{S_{int}}\right)^2 + \left(\frac{\sigma_C}{C_{high}}\right)^2}. \quad (4.32)$$

A detailed account of the mixing panel's high-sensitivity performance over an extended interval is covered in the next section.

An analysis of a high-sensitivity calibration can also be used to calculate an extrapolated limit of sensitivity ϕ_{lim} specific to the measurement configuration. Given a known sample concentration Φ_{cal} the extrapolated limit is given by:

$$\phi_{lim} = \Phi_{cal} \frac{\sigma_{int}}{S_{int}} \quad (4.33)$$

For the calibration measurement shown in figure 4.20, the krypton high-sensitivity mea-

surement limit of sensitivity is 15 ppq.

4.7 Mixing Panel

The mixing panel was constructed at the same time as the SSS at SLAC. Due to the 3 hour temperature cycle time of the SSS cold-trap, the SSS was unable to keep up with the demands of the krypton removal system. The mixing panel implemented a liquid nitrogen dispensing system and dewar shelf to automate the temperature cycling of a separate cold-trap. This reduced the cold-trap temperature cycle time to less than 20 minutes.

Construction of the mixing panel began in April of 2019 and was completed in May of 2019. Equivalent hardware used during construction of mixing panel is given in the description of generic XPM of section 4.2. If needed, components were washed with acetone followed by an ethanol in an ultrasonic bath and purged with dry nitrogen before use. The low and high pressure components were tested at SLAC to a pressure of 140 psig and 1300 psig respectively. Each test required the pressure hold for 10 minutes with no component failure or observable pressure drop. The finished panel is shown in figure 4.21



Figure 4.21: A picture of the mixing panel, sniffer, and recovery bottle. Additional components that supported operations of the SLAC XPM are installed with the mixing panel on the side of the SSS.

Between May of 2019 and September of 2020 the mixing panel configuration and procedures were changed to meet the needs of krypton removal commissioning. During this first phase of use, the mixing panel's behavior was poorly characterized, because ultimate measurement sensitivity was not a priority. The krypton removal system was still being optimized and automation control had not yet been implemented. The measurements taken during this initial development phase were critical in establishing the krypton removal run parameters and the final mixing panel procedures outlined in section 4.5.

In September 2020 the mixing panel's operation was stabilized largely from the implementation of automated control. The krypton removal system also began reliably producing xenon batches with sub-ppt krypton concentrations. This required accurate and

precise high-sensitivity measurement results to characterize the purification system's behavior. This marked the second phase of operations (September 2020 and April 2021). During this time, calibrations were completed every month to ensure accuracy.

A distinct change in system behavior was noted when an unknown event caused significant degradation of the RGA CEM. This event separates the second and third phase of operation as the mixing panel experienced a 20% loss in high-sensitivity measurement sensitivity. The third phase encompasses measurements taken after April 15th 2021.

4.7.1 Performance

The high-sensitivity measurements that span the second and third phases are characterized by a collection of calibrations performed during their respective intervals. All of the calibration measurements follow the same calibration preparation procedures, specified in section 4.5.10, and high-sensitivity measurement procedures specified in 4.5.11.

During a calibration measurement, both ^{84}Kr and ^{83}Kr are monitored. Each isotope offers a known concentration which the reported signal is compared to determine the system's calibration parameter C_{high} . We use the ^{84}Kr channel as a direct measurement of the total krypton concentration equal to the prepared Φ_{Kr} . Applying the relative isotopic abundance to the ^{83}Kr channel, we have an expected concentration equal to 20.175% the prepared Φ_{Kr} . The measured integrated signal, S_{int} , for each channel is then plotted against the known sample concentrations. A single-parameter chi squared linear fit is performed to determine $m_{fit} = C_{high}$ with a fit error $\sigma_{m_{fit}} = \sigma_{C_{high}}$. The estimated random measurement error, as described by equation (4.30), is on the order of 1% or less and neglected. The calibration mixing errors due to uncertainties in the volumes and pressures are included. The calibration concentration varies as the amount of clean xenon prepared and mixed to produce the calibration is not always the same. The variations in prepared calibration con-

centration and amount do not impact the flow interval and flow rate which remain identical across all measurements. Figure 4.22 shows six krypton calibration measurements versus the integrated signal taken throughout the second phase of operations.

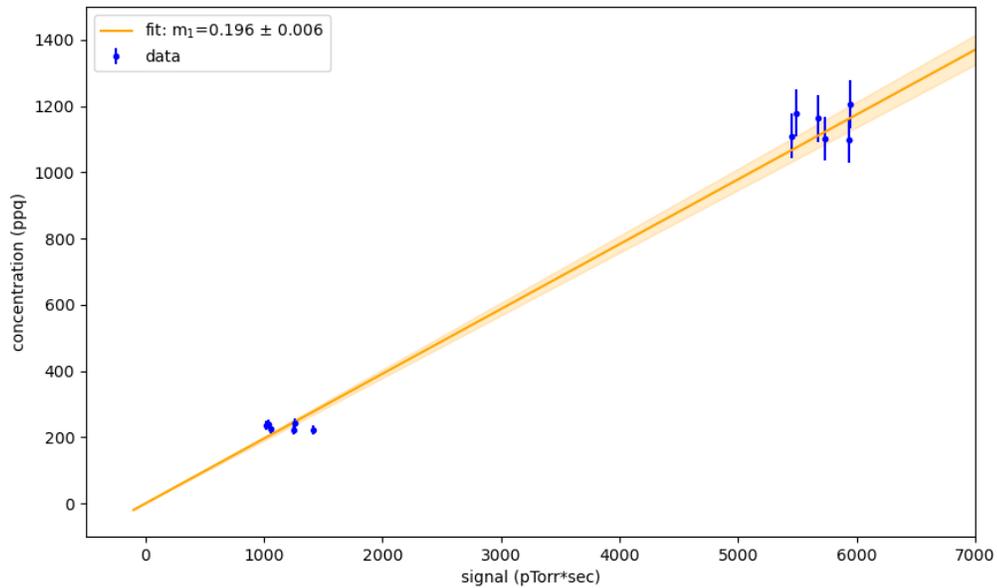


Figure 4.22: In the second phase of operation six calibration were completed with krypton concentrations around 1.1 ppt. The integrated signal for each channel is given along the x-axis and the prepared concentration, including errors, is given along the y-axis.

The variation of zero measurements, given by σ_z , account for statistical fluctuations of a null signal. The zero error is measured by preparing xenon, as described in the clean run procedure 4.5.9, to a krypton concentration below 10 ppq. Repeated measurements of this ultra-clean xenon are fit to a normal distribution to determine the zero error over the course of each interval. The measurements monitor the ^{84}Kr , ^{83}Kr , and ^{87}Kr channels, each known to contain a null signal. (Ultra-pure xenon is prepared for a variety of custom measurements including calibration mixture preparation. Other procedures requiring ultra-pure xenon are discussed in chapter 5.)

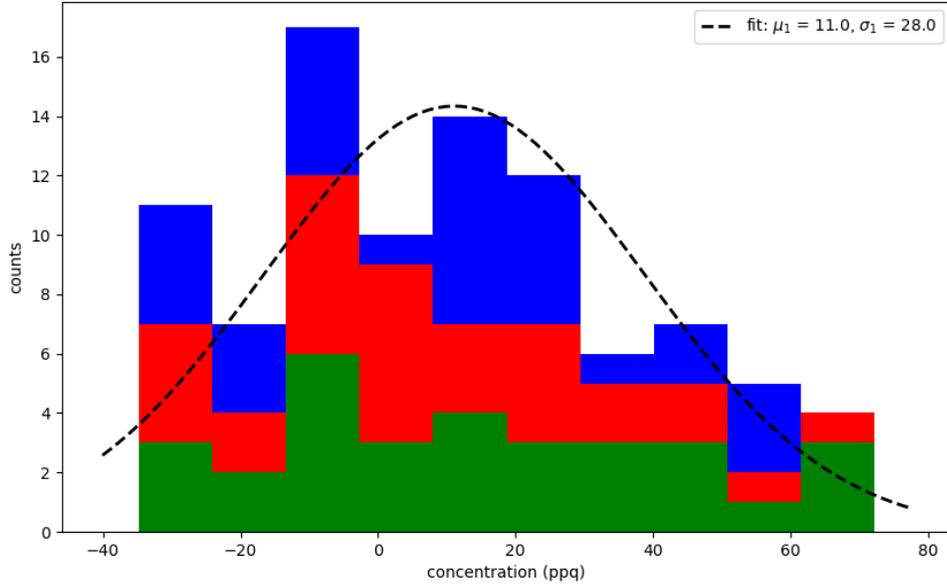


Figure 4.23: In the second phase of operations, 31 zero measurements were identified for a total of 93 data points. ^{84}Kr is shown in green, ^{83}Kr is stacked in red, and ^{87}Br is stacked in blue. The integrated signal has been converted to units of ppq using the slope value from the previous plot 4.22. The measurements taken show a distribution with a mean value 11 ± 28 ppq. The standard deviation of the mean is 2.9 ppq and the mean zero concentration is reported to be larger than zero. This could be due to insufficient cleaning prior to the zero measurement or fluctuations in the RGA baseline from environmental exposure. To account for this the full standard deviation is taken as the zero error σ_z .

For measurements taken in the second phase of operations, between September 2020 and April 15th 2021, the krypton concentration of high-sensitivity samples with an integrated signal S_{int} in pTorr·s is

$$\Phi_{Kr} = [S_{int} (0.196) \pm (7\% + 28)] \text{ ppq} \quad (4.34)$$

where the 7% is the systematic uncertainty due to the preparation of calibration xenon from system volumes, pressures, and the underlying reference xenon purity. The systematic uncertainty is more than a factor of two larger than the fit uncertainty which is neglected

here. The 28 ppq is the statistical uncertainty.

The same approach is applied to the third phase of measurements. This phase is defined by a jump in the required CEM voltage to maintain the appropriate ^{62}Xe value of 2000 ± 100 pTorr. The voltage setting changed from approximately 1400 volts to 1600 volts in April 2021. The calibration parameter is informed by the nine calibrations shown in figure 4.24 and the zero error is determined from the sixteen zero measurements shown in figure 4.25. For measurements taken after April 15th 2021, the krypton concentration of high-sensitivity samples with an integrated signal S_{int} in pTorr·s is

$$\Phi_{Kr} = [S_{int} (0.236) \pm (7\% + 18)] \text{ ppq} \quad (4.35)$$

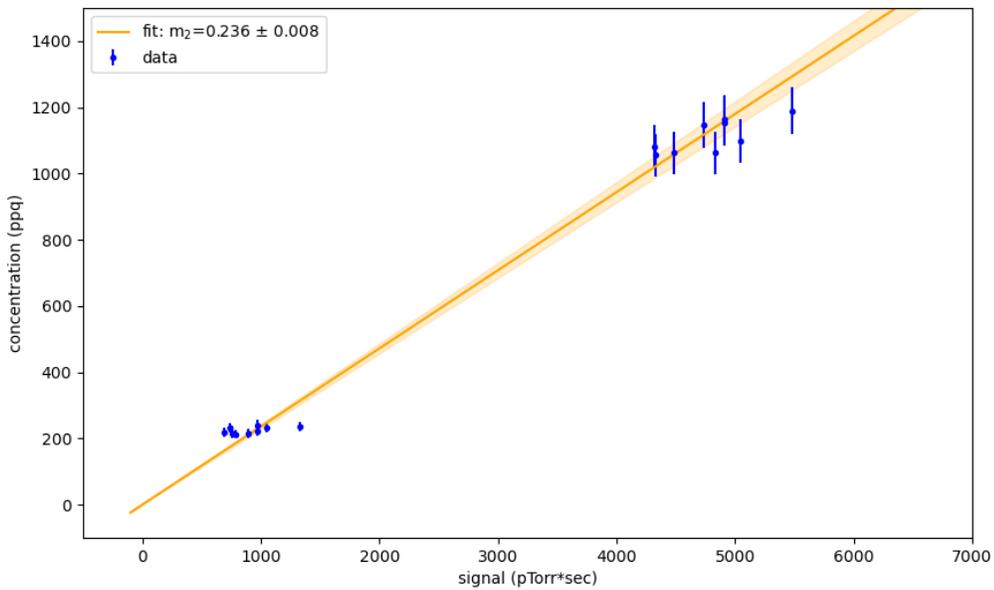


Figure 4.24: In the third phase of operations nine calibration were completed. The integrated signal for each channel is given along the x-axis and the prepared concentration, including errors, is given along the y-axis.

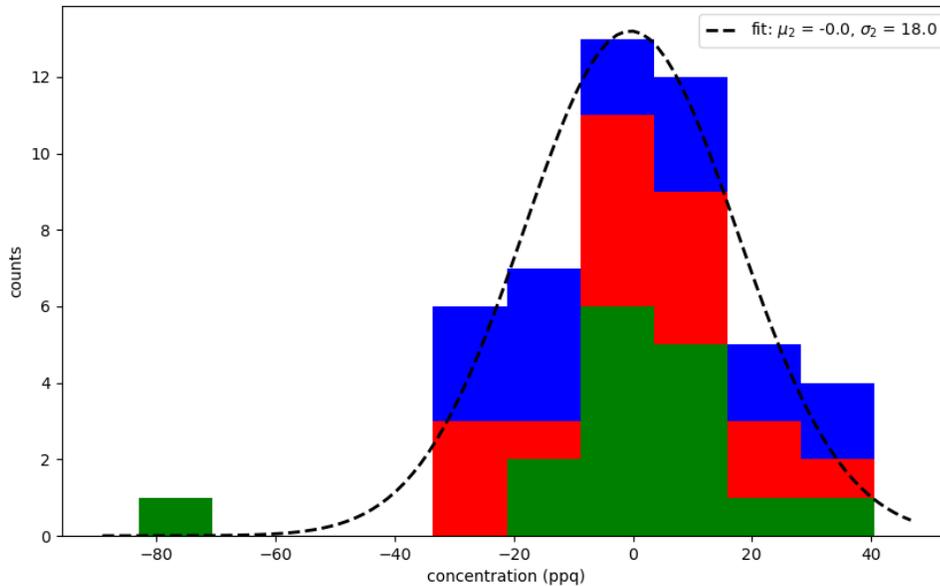


Figure 4.25: In the third phase of operation, 16 zero measurements were identified for a total of 48 data points. ^{84}Kr is shown in green, ^{83}Kr is stacked in red, and ^{87}Br is stacked in blue. The integrated signal has been converted to units of ppq using the slope value from the previous plot 4.24. The measurements taken show a distribution with a mean value 0 ± 18 ppq. The standard deviation of the mean is 4.5 ppq which is consistent with zero. Again, the full standard deviation is taken as the zero error σ_z .

4.7.2 Air Cocktail Mixture and Identification of Reference Xenon

To avoid contamination of the krypton removal system at SLAC, pure krypton sources were kept out of the facility. To calibrate the mixing panel, a krypton calibration cocktail mixture was instead prepared from air. This is similar to the cocktail mixture prepared from a pure krypton source, as described in section 4.5.10, but can be done without volume expansion steps and devoted hardware to achieve a ppb krypton concentration.. This batch of cocktail xenon was used to identify a larger quantity of xenon, suitable for use as a reference source for the lifetime of the krypton removal campaign. A candidate cylinder was identified containing over 4 kg of xenon, previously measured in 2015, to a concentration

of approximately 190 ppt. Documentation of the previous measurement was limited, so a cross-check measurement was performed.

The mixing panel's volumes, as specified in table 4.1, and an additional 510 ± 15 cc stainless steel transfer bottle, were filled with the candidate xenon to a pressure of 2762 ± 14 Torr amounting to a total xenon mass of 98.5 ± 2.5 grams. The xenon was purified with five clean runs, reducing the krypton concentration by a 99.8%.

A procedure mirroring to the calibration mixing procedure from section 4.5.10 was followed, but with air in the place of reference xenon, and volumes updated to include the transfer bottle. Air was introduced into an evacuated V_{PT} by cracking open, and immediately closing, a female VCR nut. The air pressure recorded was 424 ± 4 Torr or an equivalent mass of 54.7 ± 1.9 mg. The air was then mixed with the purified xenon by sweeping out V_{PT} with the xenon from the source and transfer bottle before being left to warm-up while exposed to both bottles. The xenon-air mixture was prepared to the following concentrations

Gas	Composition of Air by Mass	Xenon-Air Concentration (ppb)
N ₂	75.52%	4.19×10^5
O ₂	23.14%	1.28×10^5
Ar	1.29%	7.16×10^3
CO ₂	$5.1 \times 10^{-2} \%$	2.83×10^2
Ne	$1.3 \times 10^{-3} \%$	7.22
Kr	$2.9 \times 10^{-4} \%$	1.61
CH ₄	$1.0 \times 10^{-4} \%$	5.55×10^{-1}
He	$7.0 \times 10^{-5} \%$	3.89×10^{-1}
H ₂	$3.0 \times 10^{-6} \%$	1.67×10^{-2}

Table 4.2: The concentrations of constituent gasses in the xenon-air mixture given 98.5 grams of xenon and 54.7 mg of air. The errors, propagated from the appropriate volumes and pressure values, have a combined error of 4.3% for each of the calculated concentrations reported above. Air composition by mass percentages were taken from reference [31].

The contribution from water is neglected. The sample was prepared on May 8th 2019 with an ambient temperature of 70°F and the relative humidity of 70%. This amounts to a maximum water content of 0.012 mg/cc. When compared to the density of dry air at 1.225 mg/cc this amounts to 1% of the air by mass [106]. The same table was used to produce diluted air mixtures that provided a calibration for the low-sensitivity measurement technique. Air was diluted through consecutive expansions between V_{PT} and V_{CT} to reliably achieve air injection pressures near 1 Torr and mixed with xenon to produce nitrogen

calibrations in the 100 ppb to 1 ppm range and measured as shown in figure 4.19.

The resulting xenon-air mixture, prepared to identify the candidate cylinder, contained a krypton concentration of $\Phi_{Kr} = 1.61 \pm 0.07$ ppb mass. The nitrogen concentration was $\Phi_{N_2} = 419 \pm 18$ ppm, presented a significant electronic saturation hazard. Saturation was avoided by using the transfer bottle to bring 10.7 grams to a system with a hot zirconium getter to remove the nitrogen. The mixture was passed through the getter to remove the nitrogen, oxygen, carbon dioxide, and methane. The argon, krypton, neon, and helium content is unaffected during this removal process. The xenon was then recovered and reinstalled on the mixing panel.

The mixing panel was then evacuated and filled with the candidate reference xenon. A measurement is completed using a flat top averaging method with a flow rate of 1 slpm at an impedance of 1x. The xenon-air mixture was measured first over a 90 second interval. The signal was then left to settle before flowing the reference candidate over another 90 second interval.

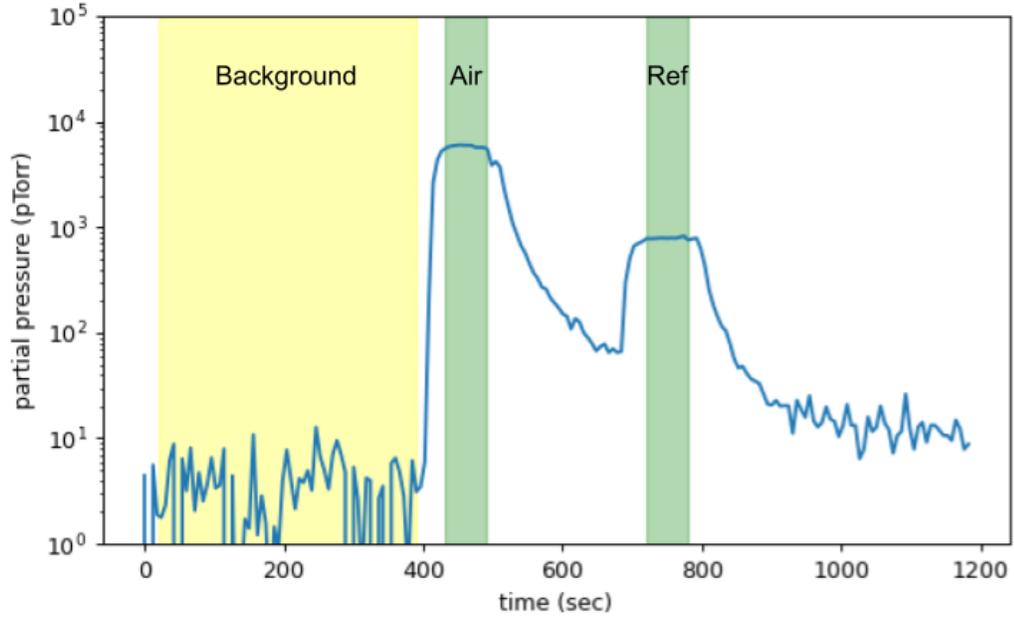


Figure 4.26: An RGA trace of the ^{84}Kr partial pressure for the xenon-air mixture followed by the reference candidate xenon. The system impedance was set to 1x and the flow rate was 1 slpm over a 90 second interval for each sample. Xenon was prepared using an air injection to a krypton concentration of $\Phi_{air} = 1.61 \pm 0.07$ ppb. The ratio of the background subtracted average signals yields a reference xenon concentration of $\Phi_{ref} = 217 \pm 12$ ppt.

The ratio of the background subtracted, averaged ^{84}Kr signals was used to calculate the reference concentration Φ_{ref}

$$\Phi_{ref} = \frac{\overline{PP}_{sig,ref} - \overline{PP}_{bkg}}{\overline{PP}_{sig,air} - \overline{PP}_{bkg}} \Phi_{air} \quad (4.36)$$

The error of the measured reference concentration is a combination of the statistical uncertainties of the average signal interval and the error from preparation of the air mixture. This error is given by

$$\sigma_{ref}^2 = \left(\frac{1}{\mu_{air}} \rho_{air} \right)^2 \sigma_{\mu_{ref}}^2 + \left(\frac{\mu_{ref}}{\mu_{air}^2} \rho_{air} \right)^2 \sigma_{\mu_{air}}^2 + \left(\frac{\mu_{ref}}{\mu_{air}} \rho_{air} \right)^2 \sigma_{\rho_{air}}^2. \quad (4.37)$$

The resulting krypton concentration of the reference xenon is $\Phi_{ref} = 217 \pm 12$ ppt.

4.8 Mobile and Stationary Sampling Systems

The MSS was constructed at the University of Maryland and shipped to SURF in January 2019. It was designed as a utility system that samples xenon throughout LZ. Of particular importance was its ability to be brought to a portion of the Davis Cavern where xenon cylinders were stored and measured after arriving from SLAC, for this reason it was built on casters.

The SSS was partially constructed at UMD and finished at SLAC. The SSS was shipped to SURF in September 2021. The SSS was integrated with the SLAC krypton removal system, but was not extensively used due to the low sample throughput. Design goals specified a measurement throughput of at least one measurement every four hours which was demonstrated, but significantly more throughput was required by the krypton removal system. The SSS has since been integrated at SURF in the xenon circulation system to monitor xenon purity in LZ.

4.8.1 Parts Cleaning, Assembly, and Citric Etching

Equivalent equipment described in the generic XPM hardware of section 4.2 was used to construct both the MSS and SSS. Much higher cleanliness standards were maintained during the construction of both systems compared to the mixing panel to limit any contaminant contributions to LZ.

Parts for both systems, with the exception of valves and sensors, were stringently cleaned prior to welding. All parts were purchased according to a ultra-high purity specification or the highest cleanliness standard available from the vendor. All tubing, fittings, and gaskets were cleaned for 15 minutes in ultrasonic baths starting with N-propyl bromide, followed by acetone, and finally ethanol. All of the ultrasonic cleaning operations were done under a fume hood with appropriate personal protective equipment, including

a respirator with 3M P100 cartridge filters and Viton chemical resistant gloves. The parts were cataloged and transferred to a class 1000 clean room. The parts were dried via nitrogen purge within the clean room.

Any component with a xenon-exposed surface was fabricated and welded within the clean room. The only exception was a copper cold-trap constructed for the SSS. The copper trap was cleaned, brazed to stainless steel, and cleaned again prior to assembly. All welding was done with a Swagelok M200 Orbital Welding System [107] and a Series 5 Weld Head [108]. The welding unit is capable of welding stainless steel tubing diameters from 1/8" to 5/8". A cerriated tungsten welding tip was used instead of a thoriated one to limit radon emanation from weld seams. At the start and end of every welding session a weld ticket was produced. Each weld ticket checked the weld bead width, thickness, coloration, and skew to ensure weld standards specified by LZ were met.

After the sub assemblies were fabricated, a fit check was performed before being citric etched. Citric etching was done to further reduce radon emanation from welded surfaces. Reverse osmosis (RO) water was mixed with citric acid to produce a 10% citric bath. The bath was heated to 60 °C and welded components were left to etch for 30 minutes before being rinsed with RO water until the conductivity of the rinse water was below 0.1 μS (Siemens).

4.8.2 Mobile Sampling System Quality Assurance and Performance

Following the complete assembly of the MSS, a series of quality assurance (QA) tests were carried out. This included electronic checks, pressure tests, and a high-sensitivity measurement demonstration. Sensor and valve control was also demonstrated prior to shipping.

The MSS does not interact with the SURF PLC, instead ADAM modules interact with

system and are controlled by the Ignition server via Ethernet. The interactive Ignition piping and instrumentation diagram (P&ID) in figure 4.27 was developed to control the ADAM modules at SURF. All of the ADAM modules were first configured and tested using the ADAM web control interface at UMD, as were all of the valves, pumps, and sensors.

Liquid nitrogen fills on the MSS are done manually. An operator is required to fill and lift the cold-trap dewar as needed. Shelves can be raised or lowered via hand cranks to submerge the cold-trap and recovery bottle volumes.

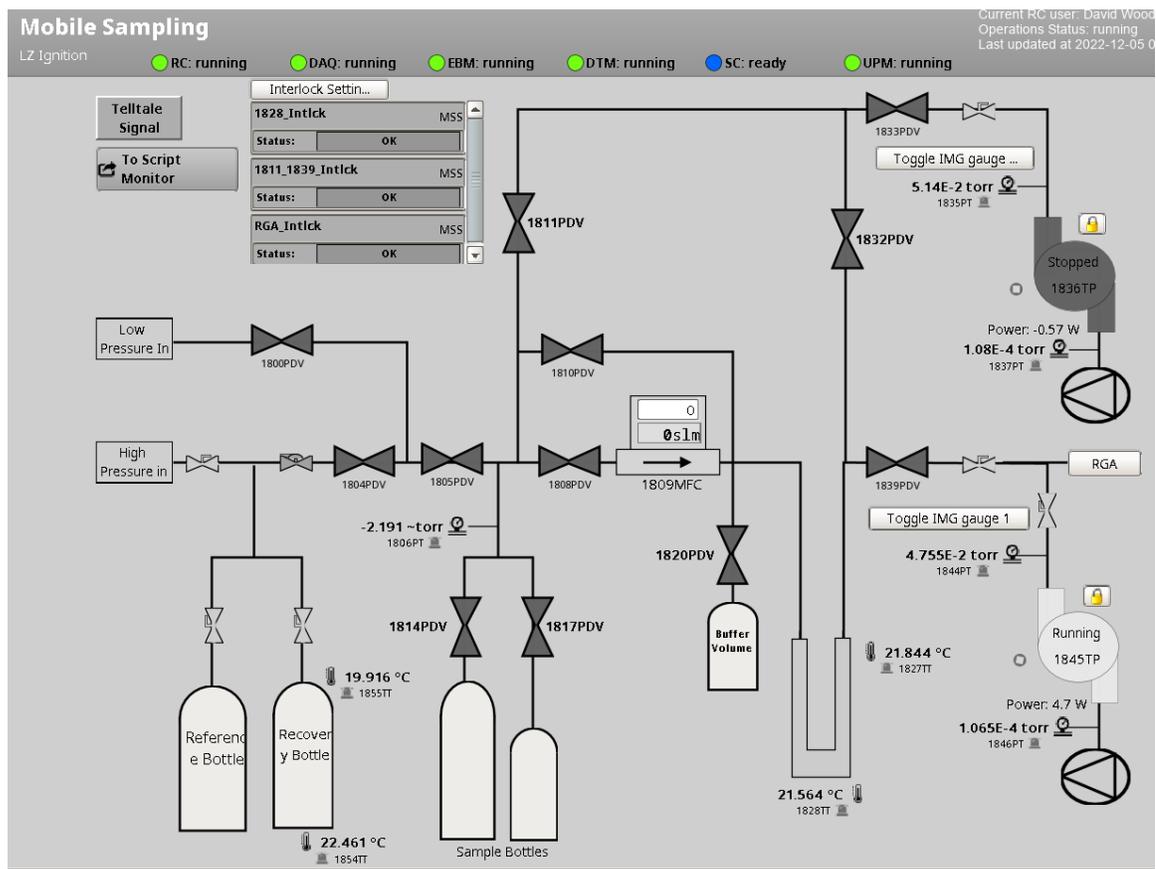


Figure 4.27: P&ID of the MSS ignition control panel implemented at SURF, courtesy of John Armstrong.

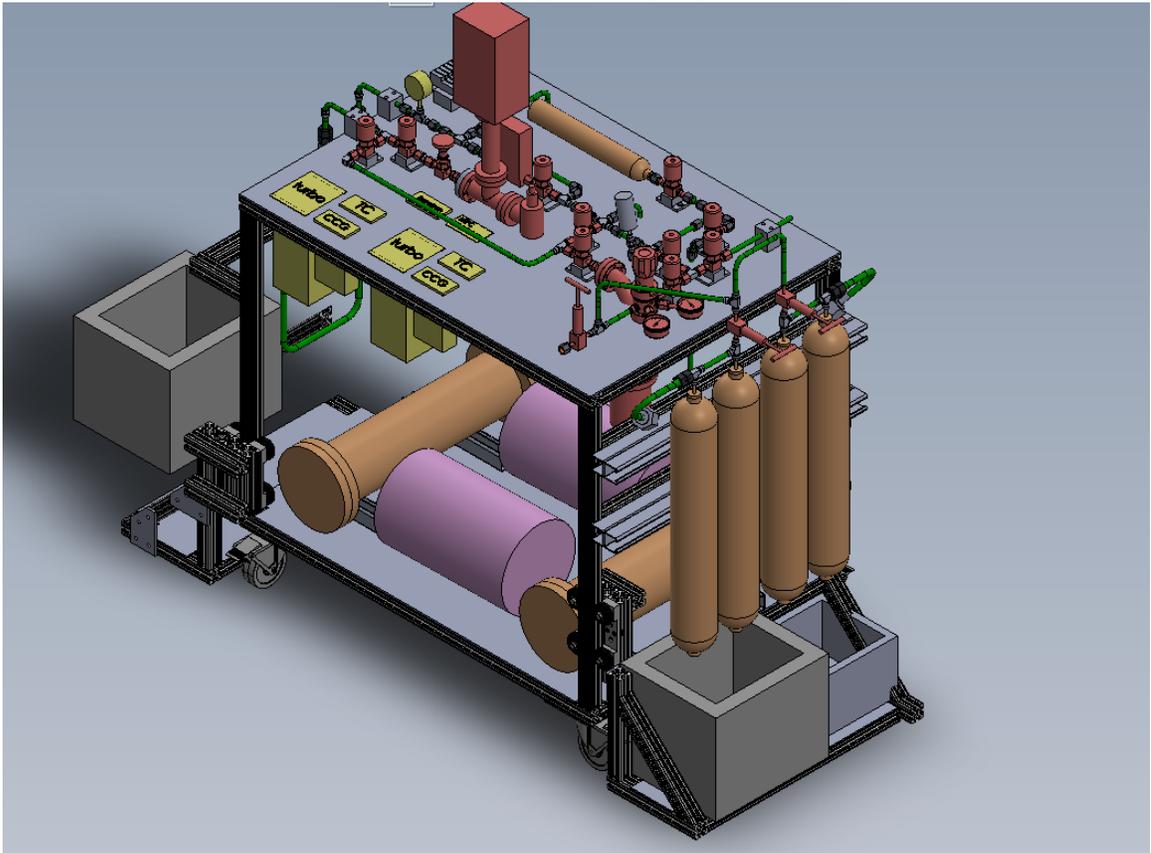


Figure 4.28: A Solidworks render of the MSS, courtesy of Timothy Edberg.

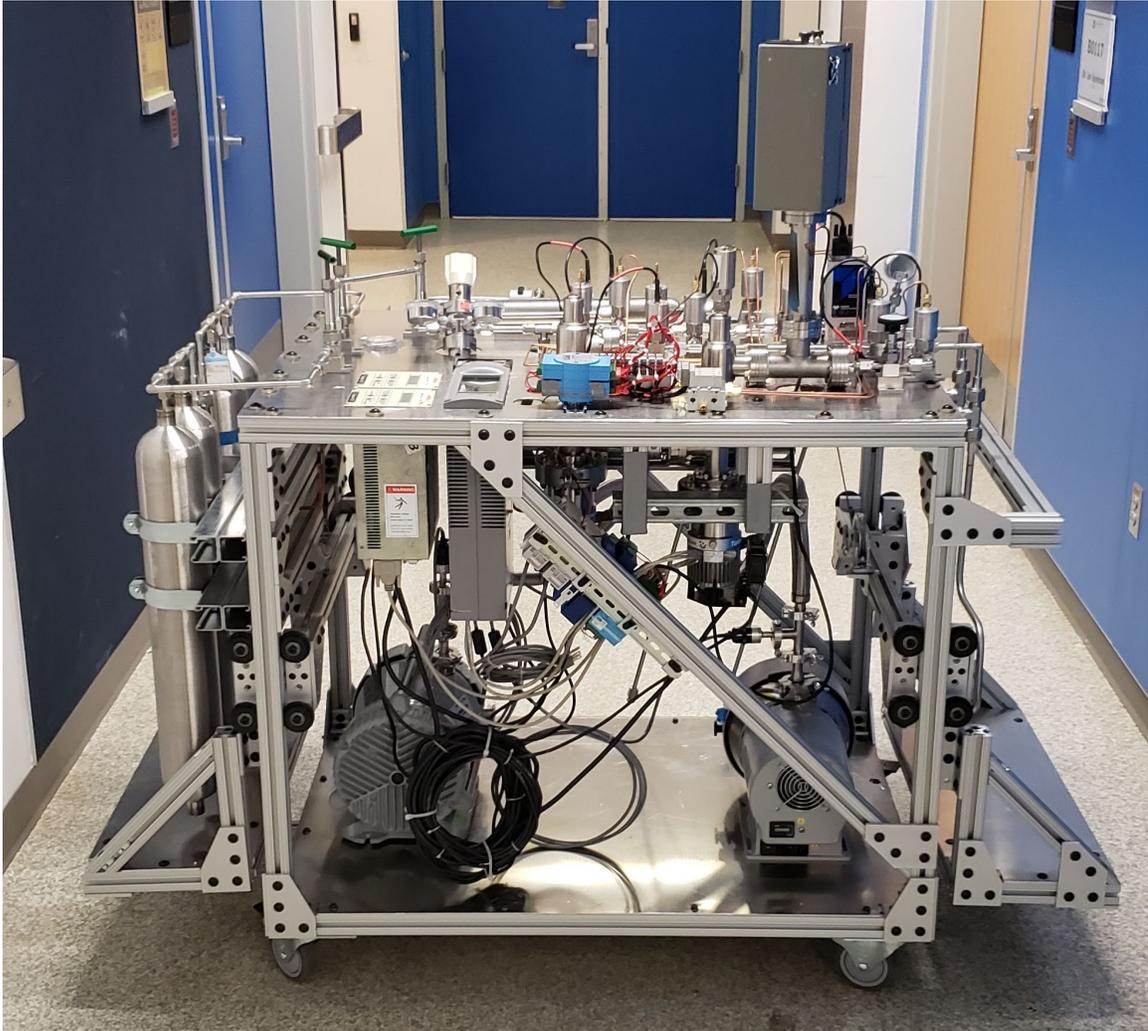


Figure 4.29: An as built picture of the MSS prior to shipping from UMD. Xenon source vessels were shipped directly to SURF and installed onsite after this image was taken.

The system was leak checked to a leak rate less than 10^{-8} mbar liters per second. The low pressure components were tested at UMD to a pressure of 180 psig for 10 minutes with no component failure or observable pressure drop. The system's radon activity was also measured at UMD to less than 0.5 mBq. The high pressure components were tested at Fermilab to a pressure of 1320 psig for 10 minutes with no component failure or observable pressure drop. Relevant quality assurance documents are kept in the LZ controlled document database, reference [109].

A final system test was carried out at UMD that demonstrated a sufficient limit of sensitivity for LZ. Ultra-clean xenon was prepared on an R&D system at UMD and mixed with reference xenon to prepare a calibration mixture containing $\Phi_{cal} = 403 \pm 40$ ppq of krypton. The calibration mixture was transferred to the MSS and measured at 1.5 slpm with an impedance configuration of 12.4x. The extrapolated limit of sensitivity, as calculated by equation (4.33), from the signal shown in figure 4.30, was found to be $\phi_{lim} = 10$ ppq.

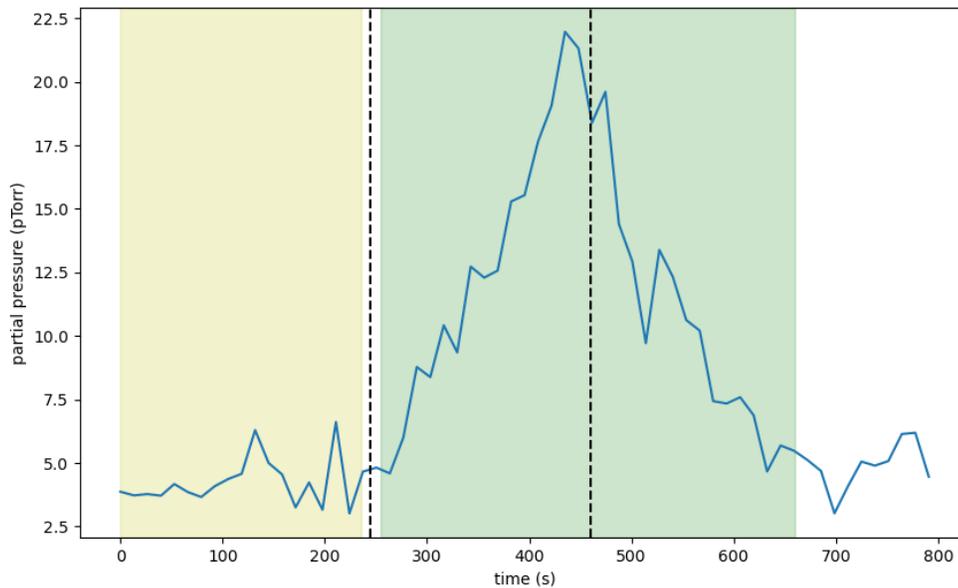


Figure 4.30: A measurement of the ^{84}Kr channel from a calibration mixture on the MSS. The background interval is shown in yellow, the flow interval between the dotted black lines, and the delayed signal window in green. The signal window starts 10 seconds after flow starts and ends 200 seconds after flow stops.

4.8.3 Stationary Sampling System Quality Assurance and Performance

The SSS was built, tested, and integrated with the PLC at SLAC during the krypton removal campaign. The interactive Ignition P&ID in figure 4.31 was developed to test the SSS valves and sensors.

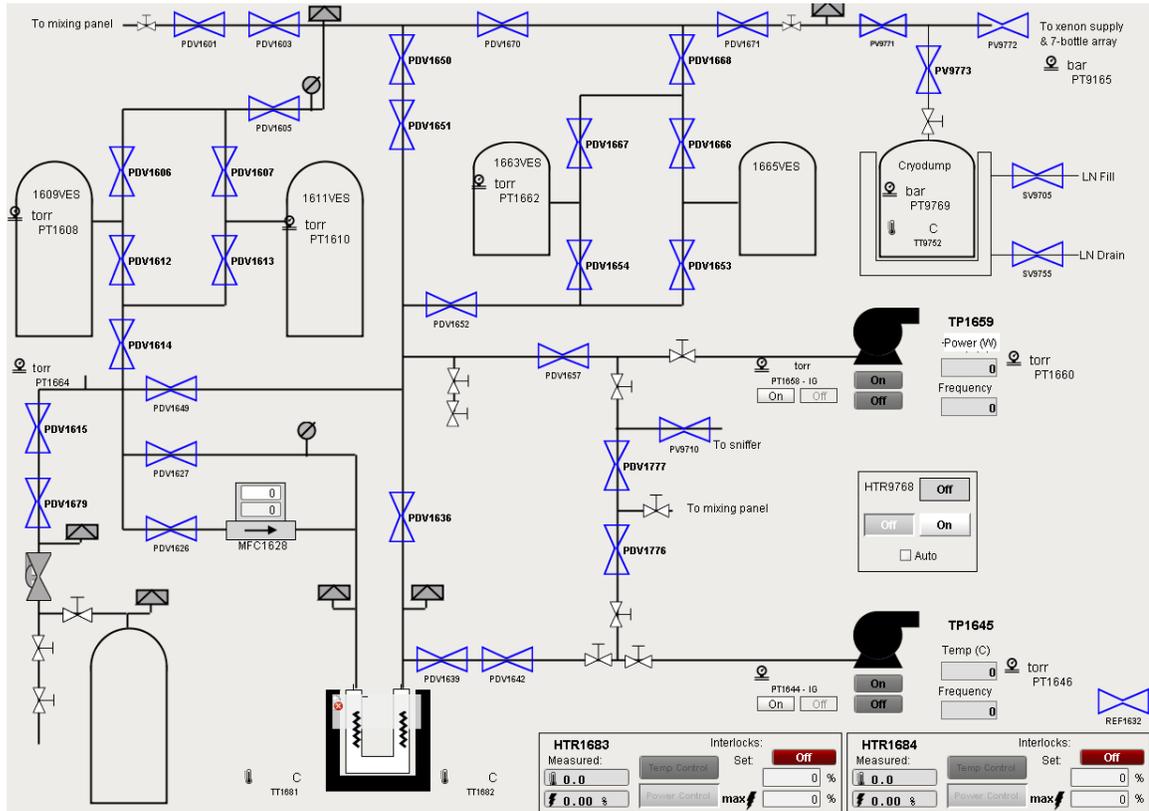


Figure 4.31: P&ID of the SSS ignition control panel while at SLAC.

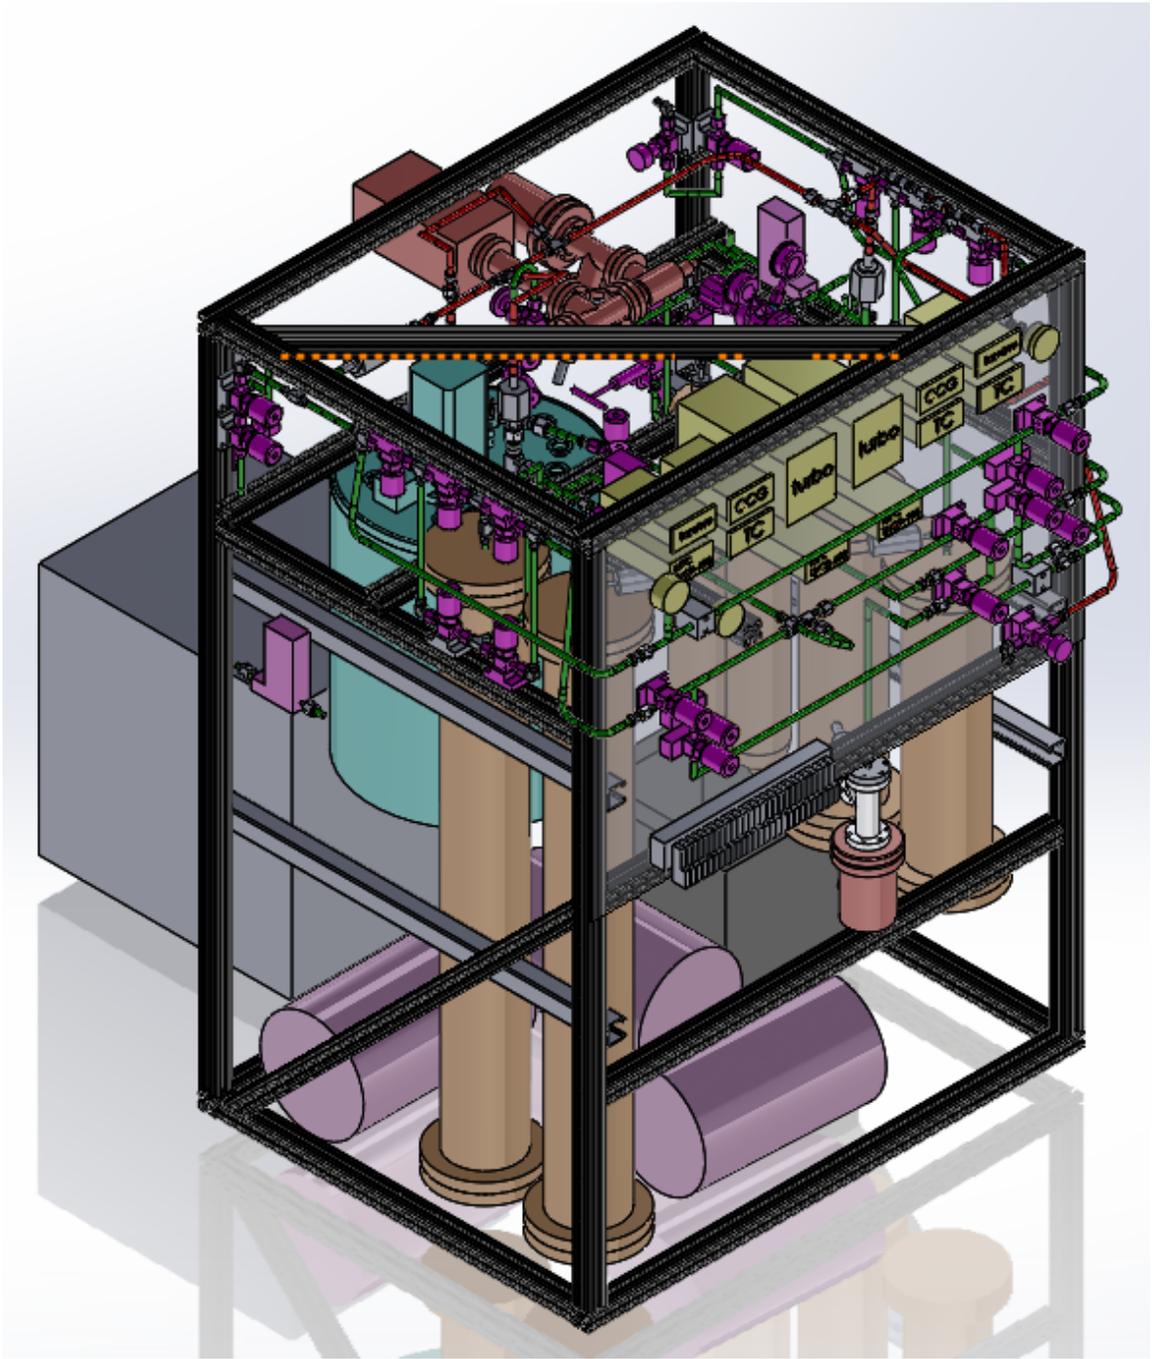
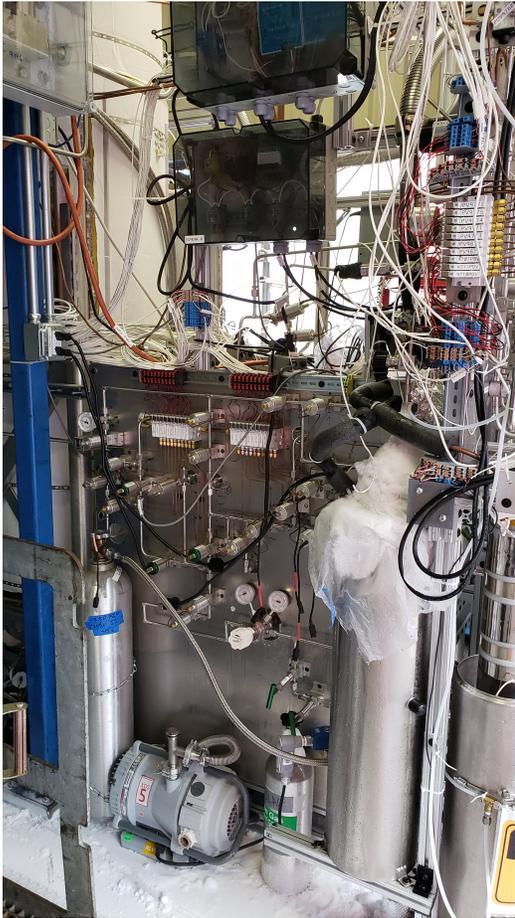
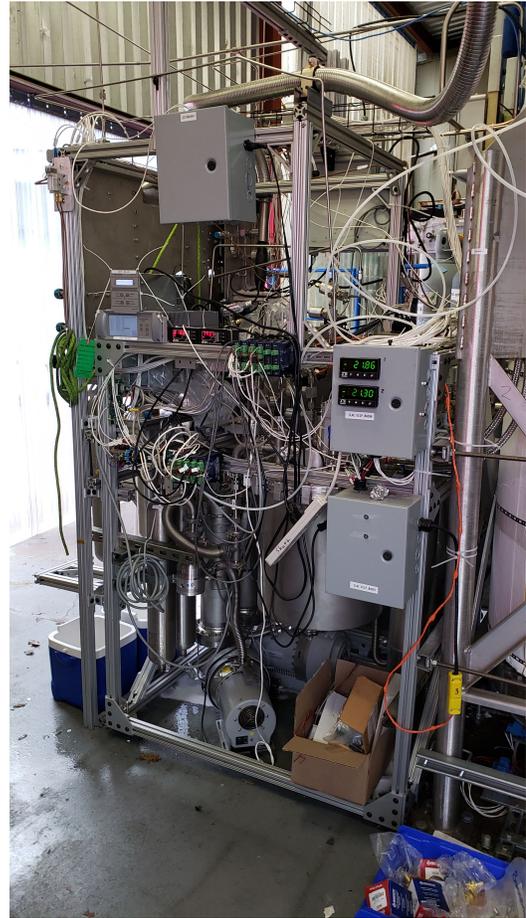


Figure 4.32: A Solidworks render of the SSS, courtesy of Timothy Edberg.



(a)



(b)

Figure 4.33: As built pictures of the SSS at SLAC. The system front panel is shown in figure 4.33a. The back panel with access to the cryostat is shown in figure 4.33b.

The SSS was leak checked to a leak rate less than 10^{-8} mbar liters per second. The low and high pressure components were tested at SLAC to a pressure of 145 psig and 1320 psig respectively. Each test required the pressure hold for 10 minutes with no component failure or observable pressure drop. A radon emanation test was not carried out on the SSS. Relevant quality assurance documents are kept in the LZ controlled document database, reference [110].

A final system test was carried out at SLAC that demonstrated a sufficient limit of sensitivity for LZ. A xenon mixture, taken from the krypton removal system during phase 1, was

measured by the mixing panel to have a krypton concentration of $\Phi_{Kr} = 7816 \pm 702$ ppq. The 9% error was calculated using the systematic uncertainty from the mixing procedure for two calibrations completed within a week of the measurement under identical configurations. Each calibration offers two known signals in the ^{83}Kr and ^{84}Kr channels for a total of four calibration data points. The average calibration parameter was calculated and their statistical uncertainty was combined with systematic uncertainty from each calibration.

The identified xenon was transferred to the SSS and measured at 2 slpm with an impedance configuration of 11.3x. The ^{83}Kr signal with an equivalent concentration of $\Phi_{Kr} = 1577 \pm 142$ ppq was used to determine the limit of sensitivity. The extrapolated limit of sensitivity, as calculated by equation (4.33) from the signal shown in figure 4.34, is found to be $\phi_{lim} = 23$ ppq.

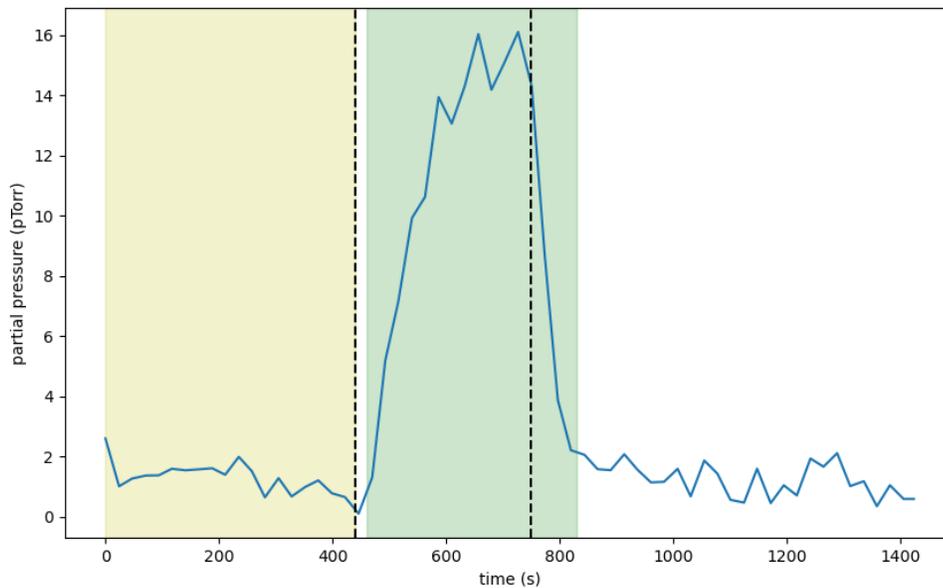


Figure 4.34: A measurement of the ^{83}Kr channel from a calibration mixture on the SSS. The background interval is shown in yellow, the flow interval between the dotted black lines, and the delayed signal window in green. For a 1577 ppq equivalent sample the implied limit of sensitivity is 23 ppq.

4.8.4 SSS Temperature Study

The SSS is cooled by a pulse tube refrigerator, model PT60 from Cryomech, capable of achieving temperatures as low as 30 K and a cooling power 60 watts at 80 K. A 3/4" diameter, copper U-shaped cold-trap, shown in figure 4.35, was brazed to 316L stainless steel reducers to 1/2" stainless steel tubing. The trap was clamped to a copper block that thermally coupled the cold-trap to the PT60 within a cryostat. Two proportional, integral, and derivative controllers (PID), part number CN8EPt from Omega Engineering [111], feedback on temperature sensors from the inlet and outlet of the cold-trap to independently regulate resistive heaters and maintain a constant temperature to within 0.1 K.

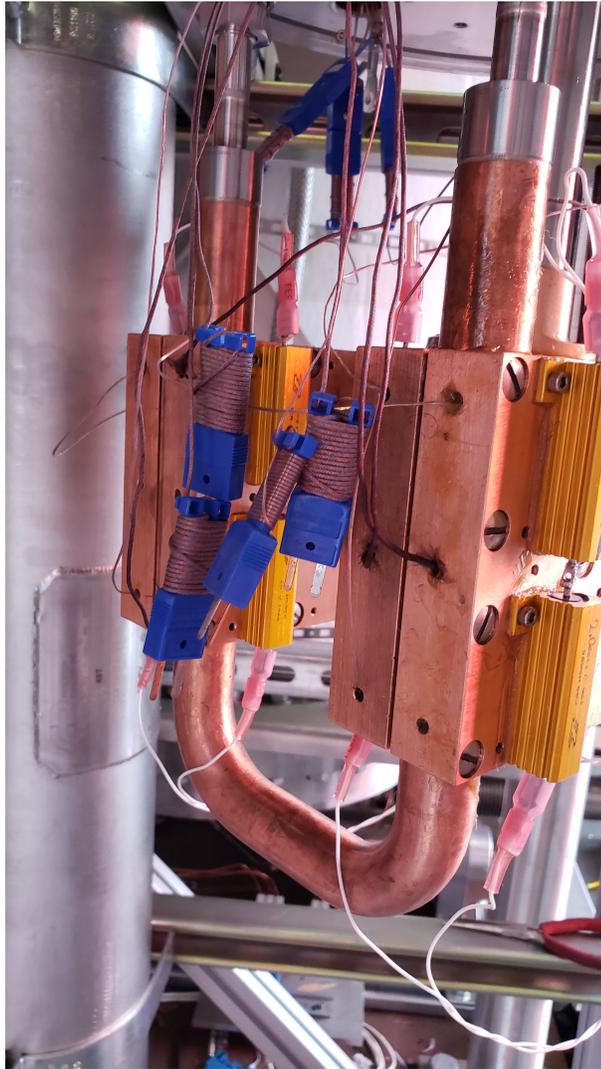


Figure 4.35: Copper cold trap and block coupled to the pulse tube refrigerator head. The block is instrumented with heaters and temperature sensors.

The ability to control the temperature allows us to set the vapor pressure of xenon ice at the cold-trap outlet. The vapor pressure of xenon as a function of temperature between 70 and 90 K is given by [112]

$$\log_{10}(P_{ICE}) = 8.044 - \frac{833.33}{T} \quad (4.38)$$

From this equation we see that reducing the temperature by as little as 2 Kelvin from

77 reduces the pressure of xenon ice by a factor of 2. This allows us to satisfy the RGA pressure condition described in equation (4.18) at a fixed pumping speed S_{RGA} and a lower impedance setting Z_1

$$S_{RGA}Z_1 = \frac{P_{ICE}}{P_{Xe}} - 1. \quad (4.39)$$

With this new impedance condition we no longer take the approximations $P_{Xe} \ll P_{ICE}$ and $Z_1 \gg 1/S_{RGA}$.

We return to our steady state krypton pressure as described in equation 4.17. Instead of increasing the impedance Z_1 and reducing S_{RGA} , we want to understand how the krypton equilibrium pressure will change as we reduce Z_1 and P_{ICE} while keeping both S_{RGA} and P_{RGA} constant. Rewriting equation 4.17 in terms of the pressures and constant pumping speeds gives

$$P_{Kr,RGA} = \left(\frac{1}{S_{RGA} + S_{trap} \left(\frac{P_{ICE}}{P_{RGA}} \right)} \right) \Phi_{Kr} Q_{Xe,CT}. \quad (4.40)$$

To test this, the SSS impedance state was configured relative to the mixing panel impedance state of 13.4x. The RGA space was shared between the two systems and at the time, krypton removal was still underway. While we are forced to use a fixed S_{RGA} to avoid perturbing the mixing panel, the SSS has an independent IV allowing us to adjust the SSS Z_1 as needed to maintain a constant pressure at the RGA. The RGA pressure requirement was selected to be $P_{RGA} = 7.9 \times 10^{-6}$ Torr to match the 13.4x mixing panel configuration.

With the target P_{RGA} and fixed S_{RGA} values selected, we would like to inspect the equilibrium krypton pressure's response to a changing temperature. A target temperature range from 73 to 81 K was selected, but a temperature sensor calibration carried out after the measurements shows the tested range was actually between 75.3 to 82.6 K.

The measurements started at the highest temperature's 13.4x "equivalent" impedance

state. This equivalent impedance state used the SSS $Z_1(T)$ configuration required to achieve the fixed target, P_{RGA} , at each $P_{ICE}(T)$. As the temperature is reduced, $P_{ICE}(T)$ falls to a new pressure, and Z_1 is reduced to compensate until P_{RGA} is achieved. Every measurement was taken at a 1 slpm flow over a 5 minute flow interval and a 5 minute background interval. The background-subtracted average signal and error, as defined in equations (4.27) and (4.28) were used as a measurement of the equilibrium partial pressure. A three parameter, chi squared fit was carried out using equation (4.40). The fit floats the S_{RGA} , S_{trap} , and $\Phi_{Kr}Q_{Xe,CT}$ which in our model and procedures remain constant.

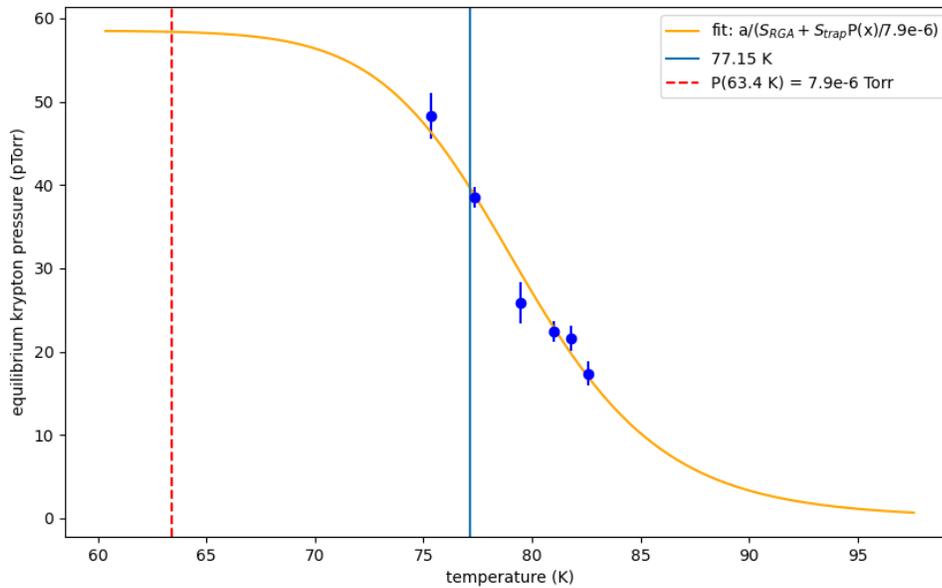


Figure 4.36: Xenon with a krypton concentration of 2933 ppq measured at different trap temperatures. The xenon ice vapor pressure is calculated using the model from equation (4.38). The vertical blue is at $x=77.15$ K, and the red dotted line is at $x=63.4$ K, where the model is expected to break down. The fit finds an RGA pumping speed of $S_{RGA} = 338$ cc/s and a trap pumping speed of $S_{trap} = 0.73$ cc/s.

The fit in figure 4.8.4 agrees reasonably well with the data. The fit values returned for $S_{RGA} = 338$ cc/s and $S_{trap} = 0.73$ cc/s are also in agreement with expected values. A

standard turbomolecular pump has a pumping speed of approximately 50 l/s for xenon gas. In the SSS, the turbo is separated from the RGA by approximately 2 feet of 1.5" diameter Con-Flat tubing which has an impedance of 1/320 s/l/cm. The equivalent RGA pumping speed for this configuration is 4.8 l/s. The mixing panel impedance was set at 13.4x impedance which corresponds to the same reduction factor applied to the RGA pumping speed giving an approximate S_{RGA} of 360 cc/s, similar to the value returned in the fit of 338 cc/s. Additionally, reference [12] studied the the relationship between impedance and krypton signal size, shown in figure 4.4, and found $S_{trap} = 0.89$ cc/s for a 1/2" diameter, stocking shaped cold-trap, similar to our fit value of 0.73 cc/s [12].

As we continue to reduce impedance Z_1 , the gains made will reach constant value as IV is fully opened and S_{RGA} remains fixed. We can further reduce S_{RGA} by increasing the impedance Z_2 via minimization of VV between the RGA and turbopump. This allows us to continue maintaining a constant P_{RGA} down to 63.4 K at which point the xenon ice vapor pressure will fall below 7.9×10^{-6} Torr. When the trap temperature remains above this limit, we can study the impact of further reducing S_{RGA} while maintaining the constant pressure condition. Figure 4.37 shows a family of curves using the same parameters, but with S_{RGA} varied by factors of 10.

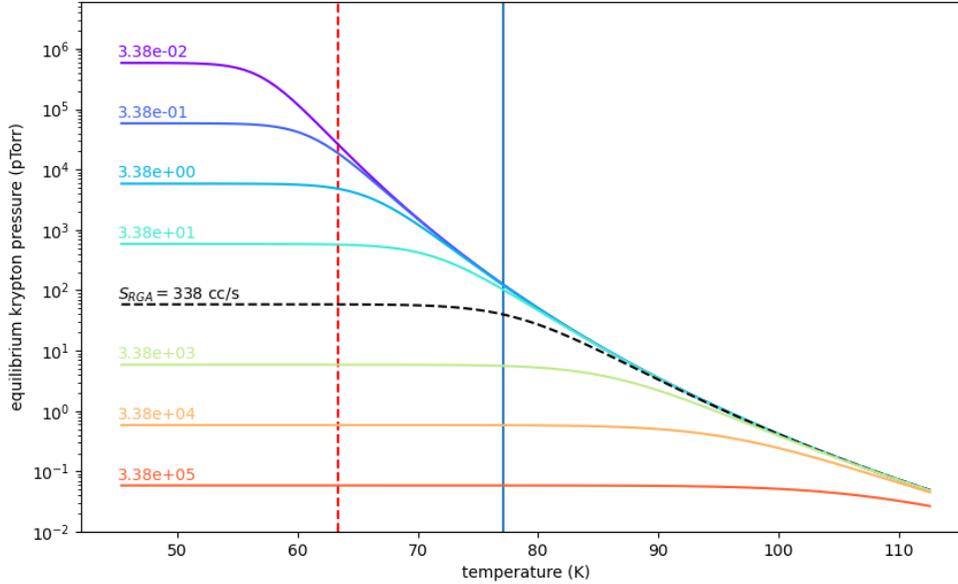


Figure 4.37: Family of curves at varying S_{RGA} . The dotted black line is the curve from figure 4.8.4. The rainbow curves show the krypton equilibrium pressure as a function of temperature of as we change S_{RGA} by factors of 10. The red dotted line shows the temperature cutoff at 63.4 K, as P_{RGA} can no longer be held constant as P_{ICE} drops below 7.9×10^{-6} Torr.

At 77.15 K these curves show that reducing S_{RGA} by a factor of 10 from 338 cc/s will increase the krypton equilibrium signal by a factor of 2.6, but reducing S_{RGA} by a factor of 100 will only increase the signal by 15% relative to the 10x case, for a total gain factor of 3. If instead we cool to a temperature just above the limiting case, say 64 K, a 10x reduction in S_{RGA} nets a factor of 9.7 increased signal, while a 100x reduction in S_{RGA} increases the signal by a factor of 78.

The results of this temperature-dependent xenon trapping method are promising and offer another avenue for significant increases in krypton signal gains. The results shown here motivate a further study of the low temperature, low impedance, and low RGA pumping speed parameter space. Some practical limitations, such as maintaining thermal equilibrium to avoid pressure saturation, may prove challenging with limited cooling power at

lower temperatures. It may prove worthwhile to investigate the use of a liquid nitrogen pre-cooler trap with a small section of ultra-cold copper tubing between the cold-trap outlet and RGA. This would allow the bulk of the xenon flow to be frozen in the pre-cooler and reduce the heat load on the pulse tube refrigerator. This would also offer a quick cycling trap that could be used to perform static pump-outs to remove potentially saturating gases prior to this new potential ultra-high-sensitivity measurement technique.

One assumption we have made is that the pumping speed of the entrapment term remains fixed at various temperatures, when in reality its behavior is unknown. The S_{trap} term was the limiting factor for impedance gains in our previous model for krypton equilibrium flow at fixed temperatures, as shown by equation (4.21) and figure 4.4. It again presents us with the limiting condition for our new model, but we have not yet characterized the effect at various temperatures and low S_{RGA} .

4.9 Conclusion

This chapter has offered a complete review of XPM theory, hardware, procedures, and measurement results. Examples of proper and poor system behavior have been presented to provide context to future operators. Sufficient sensitivity has been demonstrated for argon concentrations on the order of a ppt, a factor of 1000 lower than design specifications, and krypton concentrations on the order of 10 ppq, a factor of 30 lower than design specifications. Our theoretical model has been improved to include a variable temperature dependent ice pressure term at a fixed RGA pressure. This model was shown to agree reasonably well with data taken across 7.3 K at a fixed RGA pumping speed. A family of plots was presented from this data suggesting further sensitivity gains can be made by probing lower temperatures at lower RGA pumping speeds.

Chapter 5: The LZ Krypton Removal System

LZ purified 10.4 tonnes of xenon using a gas charcoal chromatography system. This system was scaled from the LUX chromatography system outlined in reference [113]. The purification program, often referred to as the “krypton removal campaign”, was carried out over 3 years. Construction of the chromatography system took place from 2018 to 2019, with commissioning between 2019 and 2020. Full production began in January 2021 and was completed in August of the same year. The 10.4 tonnes of processed xenon were stored in 144 custom gas cylinders, each assayed prior to shipping and upon receipt at SURF. The LZ TPC was filled with the purified xenon in September 2021.

This chapter describes how and why the noble impurities krypton and argon are removed by the krypton removal system. An overview of the component systems, how they interact, and the control systems in place is given. Some of the challenges presented by operating a complex system and how they were overcome are highlighted. Particular focus is placed on the xenon purity monitor, or sampling system, which provided essential feedback for operating and maintaining the krypton removal system. Finally the results are presented and compared to measurements taken at SURF.

Most of the xenon was purchased from commercial sources that distill air into component gasses. Industry standard research-grade xenon has a purity of 99.999%. Any remaining non-noble impurities can be effectively removed via heated zirconium getters, while noble impurities must be removed via mass separation techniques. Argon and krypton impurities contain the trace beta emitting isotopes ^{39}Ar and ^{85}Kr . Without further processing,

commercially available xenon contains enough of these isotopes produce electron recoil backgrounds dangerous to the WIMP search region of interest [17]. Without removing the krypton prior to SR1, ^{85}Kr would be the dominant background in LZ and contribute approximately 33,000 times more events than the 32 ± 5 events expected post purification.

The xenon source cylinders initially contained a 10.4 tonne mass averaged ^{nat}Kr concentration of 5 ppb and an ^{nat}Ar concentration of 20 ppb. To meet the LZ purity standards for a ^{nat}Kr concentration no more than 300 ppq and ^{nat}Ar no more than 9 ppb, the concentrations must be reduced by a factor of 17,000 and 2.2 respectively.

5.1 Gas Charcoal Chromatography

Gas charcoal chromatography is a separation technique that relies on the interactions between gas atoms and an adsorption medium. A carrier gas (mobile phase), is used to flow a process gas (adsorbate) through an adsorption medium (stationary phase). The stationary phase is a porous material that maximizes surface area for interaction sites with the adsorbate. As the adsorbate is carried over the stationary phase interacts with the adsorption media. Atoms that more strongly interact with the stationary phase will transit the adsorption media more slowly. The carrier gas maintains bulk flow through the stationary phase as the process gas stratifies into its components.

In the krypton removal system, a helium mobile phase carries the xenon gas through a charcoal stationary phase. The interactions between the noble gas atoms and the charcoal are a type of the London dispersion force. The slightly polarizable noble gas atoms form dipoles that interact with the conductive charcoal surface. Larger noble atoms are more polarizable and will more strongly interact with the stationary phase. As a result the larger atoms will move more slowly through the charcoal than the lighter mass impurities, like krypton, which elute first [114]. Once the charcoal is saturated with xenon, the leading

cloud of impurities exits the stationary phase and is discarded. The processed xenon is then pulled from the charcoal by a gas pump and the stationary is now reset.

5.2 System Overview

The krypton removal system is comprised of three primary subsystems and an auxiliary sampling system. The three primary subsystems are the chromatography loop, recovery loop, and storage branch as shown in figure 5.1. A xenon feed adds the process gas into chromatography loop where an actively purified helium stream is recirculated through a charcoal column. Once the column is saturated with xenon the feed is stopped and the recovery loop is triggered. The xenon-helium mixture is pumped from the column and passed through a freezer that collects the xenon as ice. The empty column can be used again. The krypton removal system is designed with two charcoal columns running in cadence to maximize throughput. As one column operates in the chromatography loop the other column is recovered. Once the freezer is full, the excess helium is pumped away. The freezer is then warmed and the purified xenon is fed, via compressor, into specialized storage cylinders. The auxiliary xenon gas purity monitor, also known as the "sampling system", samples gas mixtures at each step in the krypton removal system to ensure the proper operation of components and purity of each batch [28].

The krypton removal system schematic, in figure 5.1, includes the gas path of each subsystem. The helium re-circulation loops are driven by pumps that can introduce contaminants. Krypton traps (KT) are placed at the output of these pumps to maintain helium purity and capture any krypton. Helium is fed into the system from source cylinders and initially purified by KT2, a plumbing U-trap filled with 11.2 L of charcoal and submerged in liquid nitrogen. The KT2s can saturate with contaminants and if not periodically cleaned they will release krypton back into the system. This timely cleaning procedure requires the

traps to be warmed and evacuated to high vacuum via turbo molecular pump. Specialized gas sampling procedures were developed to identify when the trap was sufficiently clean and established a regular maintenance schedule based on xenon throughput. These procedures reduced unnecessary cleaning, maximized system up-time, and reduced potential batch contamination [28].

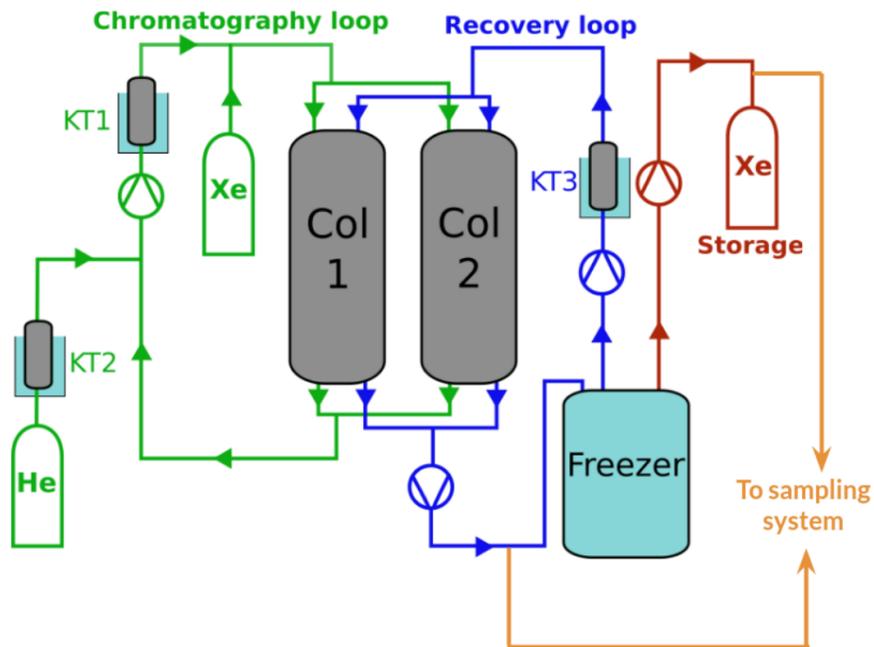


Figure 5.1: A block diagram of the krypton removal system. Each of the three main subsystems are shown and connection points to the sampling system. Compressors, recirculation pumps, gas cylinders, and krypton traps (KT) are shown as well. This image is taken from reference [28].

5.2.1 Chromatography Loop

Each of the two 4.3 m tall and 1 m diameter columns are filled with roughly 400 kg of activated charcoal. Temperature fluctuations affect xenon and krypton transit speeds through the charcoal. Extra care is taken to account for these temperature fluctuations as the krypton removal system is in an open air work area, semi-exposed to the outdoors. To compensate the helium flow rate is adjusted between 600-1200 SLPM to maintain a

constant column pressure of 1.35 bara and a 3 hour process time. Once helium flow is stable, a 16 kg xenon slug is fed into the top of a chromatography column at a flow rate of 145 SLPM.

KT1 is a small chromatography column submerged in liquid nitrogen that maintains the purity of the recirculating helium. It contains 9.7 L of charcoal topped with 6.7 L of zeolite molecular sieve. Zeolite removes any trace air contaminants introduced by the re-circulation compressor and the charcoal adsorbs any krypton separated from the xenon slug during chromatography. [28].



Figure 5.2: An image of the krypton removal system at SLAC. The system is in a metal structure, semi-exposed to the outdoors. Not shown here are the helium re-circulation pump and fluitron compressor. This image is taken from reference [28].

The helium stream is monitored by a binary gas analyzer (BGA) [115] at the output of the column. The BGA measures the speed of sound in a two component gas mixture to determine their relative concentrations, a plot of the relative xenon fraction as measured by the BGA is shown in figure 5.3. As the xenon exits the column the speed of sound

in the helium mixture decreases. Once the xenon content of the helium mixture exceeds 2.5% by mass, the end of chromatography is signaled and a column swap is triggered. A backup residual gas analyzer (RGA) [27] monitors the xenon concentration to ensure the end of chromatography was properly identified. If the end of the chromatography cycle is missed, xenon will accumulate in KT1. This causes krypton to be released back into the column and will contaminate the processed slug. Following xenon detection, a column swap begins. The helium is moved from the saturated xenon column into the empty column using a KNF diaphragm pump. Once the pressure in the saturated column drops below 15 mbar the columns are isolated from one another and the the column swap is complete. The empty column now has helium in it and is ready to start a new chromatography cycle while the saturated column is ready to start recovery.

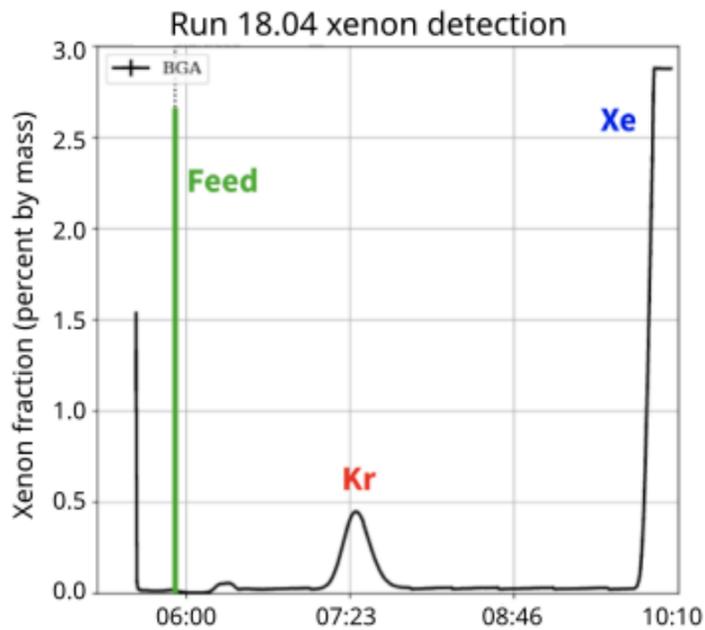


Figure 5.3: Relative concentration data versus time from the BGA during a chromatography cycle. The output of the column is monitored from the start of xenon feed marked at the green time interval. This cycle contained process xenon with a very high krypton concentration (600 parts-per-million) to show peak separation. The krypton peak is not shown to scale as the BGA response is calibrated to xenon. This plot is taken from reference [28].

5.2.2 Recovery Loop

A Leybold DryVac DVR 5000 vacuum pump drives the remaining helium through the recovery loop at a high volumetric flow. The purified xenon is pulled from the charcoal column and into the freezer. Plates with aluminum fins are thermally coupled to copper plumbing flushed with liquid nitrogen collect the purified xenon as ice. At the freezer outlet a mass flow controller (MFC) bypasses the KNF pump and maintains the recovery pressure as helium flushes the column. Between the MFC and the column is KT3. KT3 is filled with 11.2 L of charcoal and submerged in liquid nitrogen to maintain helium purity. Recovery takes between 2.5 and 3 hours to complete. Once the 16 kg xenon slug is fully collected in the freezer, the column is empty and ready to be swapped. The freezer, shown in figure 5.4, has a maximum capacity of 250 kg, or 15 xenon slugs. After capacity is reached ice formation begins to restrict flow and recovery is inhibited [28]. At various times during recovery, a sniffer sample of the helium-xenon mixture could be taken from the outlet of the DVR. These measurements were used to track and tune the krypton removal system's performance as explained in section 5.3.



Figure 5.4: Lowering the krypton removal freezer into the freezer vessel. As the helium-xenon mixture passes through the freezer, xenon ice forms on the aluminum fins. Aluminum plates are thermally coupled to copper tubes flushed with liquid nitrogen. Heaters and temperature sensors on each aluminum plate control the temperature.

The DVR vacuum pump consists of a RUVAC WH 2500 roots blower backed by two DV650 screw pumps in series. Air leaks to rubber seals throughout the DVR were managed by an acrylic nitrogen purge box. An oil space within the pump's gearbox is exposed to the gas process space. As a result the oil is susceptible to contamination and will release impurities back into the purified xenon stream. Contamination can occur from ending chromatography too early or from KT3 releasing krypton into the loop due to xenon breakthrough from the freezer. To clean the DVR oil, first KT3 must be brought up to room temperature and pumped to vacuum to remove any contaminants. With KT3 clean and cold, the DVR is run in a helium only. This allows the gearbox to agitate the contaminated oil and mix it with the purified helium stream. The high volumetric flow of the helium removes contaminants from the oil and captures them in KT3 [28].

5.2.3 Storage

Once the freezer is at capacity, every 2-3 days, the remaining helium is removed from the xenon. PMTs in the TPC are sensitive to helium ingress that results in after pulsing. LZ requires a final helium concentration of less than 200 ppb. A helium concentration of 1 ppm is achieved by directly pumping on the xenon ice with a turbo-pump for 90 minutes. The remaining helium is embedded in the xenon ice during recovery. To remove the remaining entrained helium, the freezer is isolated and warmed to melt the xenon ice. With the entrained helium now free, the xenon is re-frozen and pumped for another 90 minutes. The final helium concentration achieved is roughly 10 ppb, well below the 200 ppb requirement [28].

With the helium removed, the freezer undergoes a final warm-up. Heaters warm-up the xenon to 1.5 bara, at which point the pressure is balanced by a Fluitron D1-20/120 two-stage metal diaphragm compressor which moves the xenon into a storage pack. The

storage packs are made up of 12, 49.1 liter, industrial gas storage cylinders. Each cylinder is evacuated prior to filling and when full, holds 69 kg of xenon at 65 bara. The neck of the cylinder is continuously purged by nitrogen to limit air ingress. A total of 12 storage packs were used to transport the 10.4 tonnes of purified xenon from SLAC to SURF. Xenon can be sampled directly from the freezer during warm-up or from individual cylinders once storage is complete. Every storage pack was assayed prior to shipping.

5.2.4 Control

Automation controls were used in the krypton removal system to maximize xenon throughput while minimizing operator burden. All operations with the exception of storage were fully automated. Semi-automated storage operations presented the highest potential for xenon loss and required an operator to be on site.

A programmable logic controller (PLC) was used to handle low level operations that interacted with valves and sensors. Interlocks controlled by the PLC helped protect operators, equipment, and xenon loss. Most discrete processes were handled by the PLC for example xenon feed, column swaps, liquid nitrogen dispensing, and maintaining the chromatography and recovery loops. In the event of an emergency or component failure, interlocks enforced system safe states. Higher level functionality was implemented through Ignition, a supervisory control and data acquisition (SCADA) software from Inductive Automation. This software package includes a graphical user interface designer, an alarm system, a data historian, and a scripting framework.

The krypton removal run control software, built from Ignition's scripting framework, monitored and queued PLC processes that managed xenon feed, chromatography, recovery, and storage. The run control and alarm interfaces are shown in figure 5.5. Timing between the queued operations operations were coordinated based on system conditions and sensor

data from each subsystem. The sampling system was independently automated through Ignition. This allowed the two systems to operate independently and coordinate as needed.

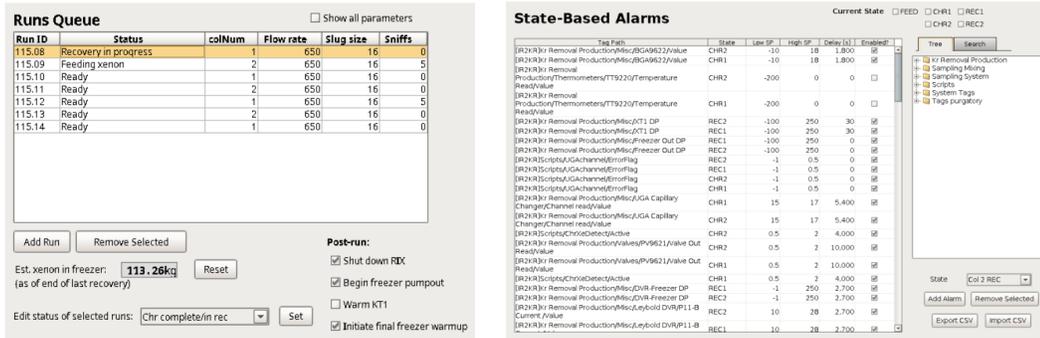


Figure 5.5: Screenshots of the Ignition user interface for run control (left) and state-based alarms (right). Users can queue multiple runs and set parameters like flow rates, slug sizes, and sniffer requests. Taken from [28].

5.3 Sampling System

The xenon gas purity monitor at SLAC, referred to here as the sampling system or sampler, was used to assay xenon purity throughout the krypton removal system. A picture of the system is shown in figure 5.6. In addition to determining the final mass averaged xenon purity, these measurements provided essential feedback used to identify sources of impurities and optimal run parameters. This section will review sampler automation and measurement results as they pertain to the krypton removal system. A thorough analysis of the underlying measurement technique and operational procedures for a generalized xenon gas purity monitor were reviewed in chapter 4 and section 4.5.

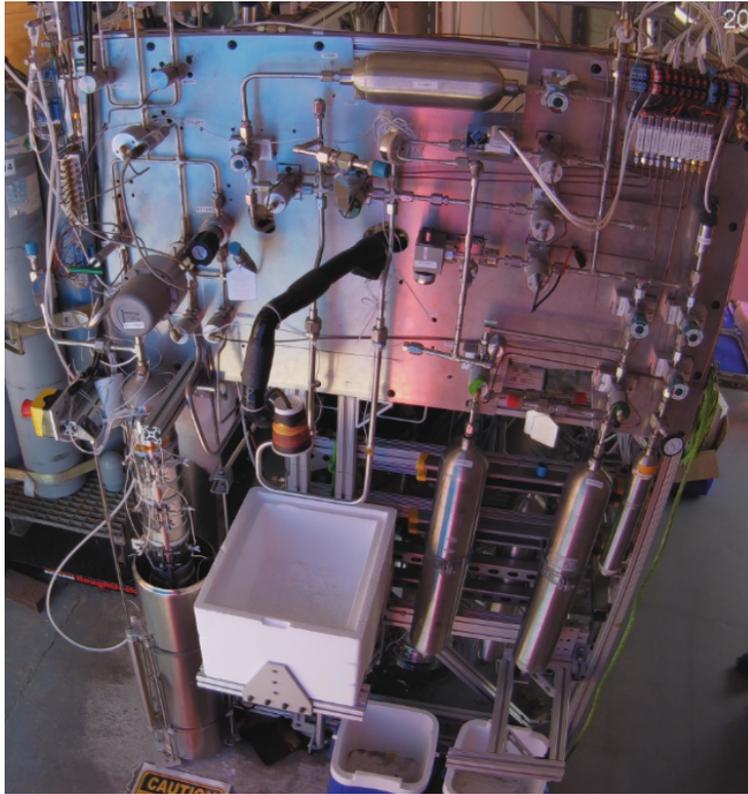


Figure 5.6: An image of the SLAC sampling system taken via webcam. On the right are two 3.8 liter and one 0.5 liter stainless steel source bottles. Beneath the stocking shaped cold-trap is a styrofoam dewar that raised and lowered by a pneumatically actuated shelf. On the left is the sniffer bottle and a cylindrical dewar on a pneumatically actuated shelf. Both dewars are supplied liquid nitrogen by an insulated copper line. On the top right a manifold and solenoid valves control the system valves. Directly behind this panel are the cryodump, RGA, and pumps.

Two measurement procedures were employed for the krypton removal system, a high-sensitivity and a low-sensitivity measurement. For a full analysis, a low-sensitivity measurement was performed to identify concentrations of nitrogen, argon, and helium, followed by a high-sensitivity measurement focused on krypton detection. Ultra-pure xenon, produced by the sampling system, was essential in commissioning and maintaining the krypton removal system. Ultra-clean xenon was used to sweep out or steep in evacuated volumes to identify sources of potential contamination. These operations are referred to as back-flow tests. These less standard procedures identified out-gassing components, small

air leaks, and confirmed the cleanliness of KT's and DVR oil after after cleaning procedures.

Xenon was collected during recovery through a devoted sniffer subsystem. A stainless steel bottle was frozen using liquid nitrogen and exposed to the helium-rich xenon stream at the outlet of the DVR pump. Once the xenon was frozen to the walls of the bottle, the helium was pumped away. This sniffing process was repeated until sufficient xenon was collected to complete a measurements. The xenon then undergoes a sniffer static pump-out before being warmed and transferred to the xenon purity monitor portion of the sampler and measured. The timing of these sniffs with respect to the start of recovery was essential in tuning the chromatography parameters. Early sniffs that were contaminated with krypton indicated incomplete separation from the xenon. A sufficiently high concentration of krypton due to poor separation can contaminate the DVR oil space and the xenon collected in the freezer from prior clean xenon slugs. Late sniffs were taken a few minutes into recovery and represent the ultimate purity achievable for that column. Late sniffs that produced dirty results between both columns indicated a contaminated oil space in the DVR. An increasing krypton concentration in late sniffs was also an indicator of krypton accumulation in the DVR oil. A late sniff that was dirty, but only seen in one column was helped identify a small air leak from a burst disk on that column.

5.3.1 Sampler Run Control

Tags are values within Ignition that correspond to things like a pressure sensor reading or valve state. Users and scripts can interact with tags by both reading and writings values, for example by checking if a valve is opened and closing it if it is. Tags can also be made to represent derived values, for example a 10 second moving average of a pressure value or a logical comparison that checks if a collection of valves are closed. Tags that represent more abstract values corresponding to system states, are especially useful when automating

system processes.

Interacting with tags through Ignition's scripting framework depends on tracking system states and automating the transition between them. Scripts are written to interact with these tags and launched by scripting slots. The slots that trigger the scripts also have tag values that can be managed in an automated fashion.

To automate the sampler, three distinct subsystems are defined to control three unique groups of tags. These subsystems are the mixing panel, sniffer, and storage line. The automation software is designed such that each subsystem is controlled independently by two devoted scripting slots. The two scripting slots are called the run control and the operation control. Only operation control scripts can interact directly with hardware tags. The run control manages which script is run in operation control. To communicate between the two, a subsystem state variable is monitored by run control and updated operation control. This tells run control when an operation is complete and so run control can feed operation control the next operation script. A block diagram of this process is shown in figure 5.7.

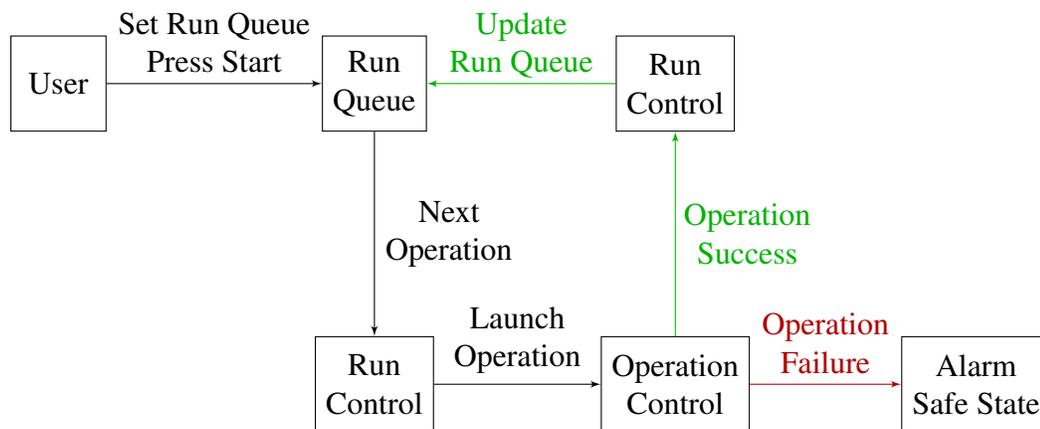


Figure 5.7: Block diagram of a single subsystem’s automation software. A user inputs a list of queued operations and starts the software. The run queue is read by run control which then launches operation control. Operation control manages the subsystem hardware. If the operation succeeds run control is informed, the run control is updated, and if another item is in queue the process repeats. If the operation fails a safe state is enforced and an alarm is triggered.

First the user populates the subsystem run queue tag with the desired operations and launches run control. Run control then checks the run queue and loads the appropriate script into operation control. Operation control updates the subsystem state variable to communicate the start of the operation to run control. Once an operation is completed, the subsystem state variable is updated. Run control removes the completed operation from run queue and repeats the process for the next item until the queue is exhausted.

In addition to the three main subsystems, there are five shared devices. The shared devices include tags that represent sensors, valves, and pumps shared between each subsystem to carry out specific operations. These devices are the calibration line, cryodump, RGA, transfer line, and utility pump.

The sampler’s run control GUI is shown below, in figure 5.8. In the blue block, each of the subsystems have a devoted run queue than a user populates with desired operations. A system queue located in the center of the GUI is used to queue whole operation cycles that chain operations between subsystems to process samples and reset the system. To start the

operations a user simply clicks start once the queues are populated. Once the operations start the “ABORT” button can be used to stop the current operation, force a safe state, and trigger a subsystem error. The subsystem error fixed by the user and acknowledge using the “Error Reset” to regain use of the subsystem’s run control. Operation queues can be edited while run control is in use, as long as it is not the top entry. Top entries can only be edited when run control is not in use. A small liquid nitrogen auto-fill interface is shown in purple. A device request queue is highlighted in green and shows which device is requested by which subsystem. The red box encloses the sniffer-krypton removal interface which is cloned in the krypton removal run control. This way krypton removal users can quickly check the sniffer status during coordinated sniff operations.

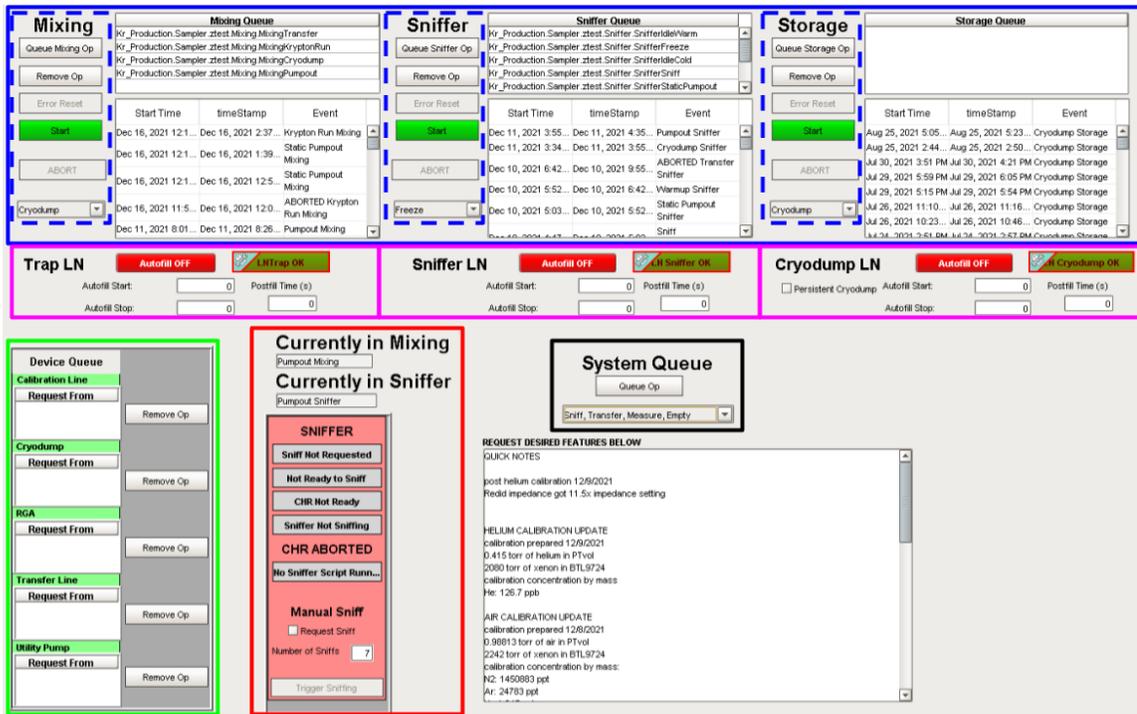


Figure 5.8: A screenshot of the sampler’s run control GUI. The subsystem run queues are highlighted in blue, a system queue in black, a liquid nitrogen auto-fill interface in purple, the device request queue in green, and a sniffer-krypton removal interface in red. A text box is also available for quick notes.

5.3.2 Sampler Operation Control

Operation scripts use three main code blocks. A requirements check, run sequence, and abort sequence. The initial requirements check valve states, pressure sensors, and temperatures to ensure the subsystem is configured as expected. If any parameters do not meet the expected values, the operation does not begin and the run control software alerts the user. This keeps the system from trying to do things like measuring an empty system or pumping a large quantity of xenon to air.

Once the requirements check is passed the run sequence begins. At the start of every run sequence the subsystem state variable is updated to inform run control. Any devices required for the upcoming operation are then requested. This is handled by a device queue tag. The script will wait for up to four hours to use the requested device. Once the request is at the top of the device queue, the run sequence checks the device configuration and uses it as needed. When the device is no longer needed, it is reset and the device queue is updated. Once the operation is complete the run sequence updates the state variable.

In the event that two subsystems must coordinate controls, a control hand-off is implemented. This type of operation is common when moving xenon between subsystems. This is implemented using an idle operation in the xenon source subsystem, which signals to the xenon receiving subsystem that it has permission to control the source subsystem's devices. This reduces the number of coordinated events to only two steps: acknowledgement of control hand-off and acknowledgement of operation completion.

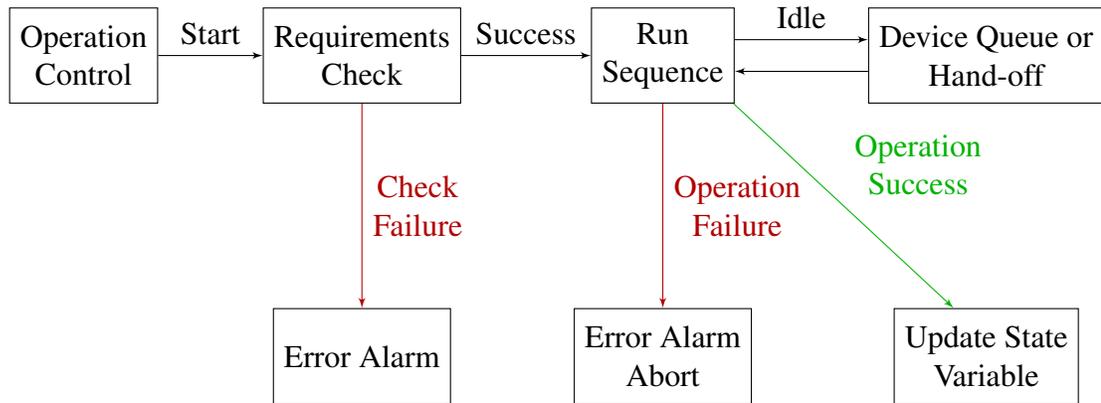


Figure 5.9: Block diagram of an operation script. Hardware interactions can only occur in the run sequence or during an abort.

If an error occurs in run control or during an operation, the operation is aborted. The abort sequence enforces a subsystem safe state and triggers subsystem alarm. Any devices and queues submitted by the aborting subsystem are reset. Alarms are set to contact users by email, cellphone, and Slack. The automation software for the subsystem remains locked until a user acknowledges the error. Detailed logs from script outputs are used to track down the error which must be fixed manually. Any configuration issues are managed through the sampler GUI shown in figure 5.10. The specific devices and subsystems are also labeled.

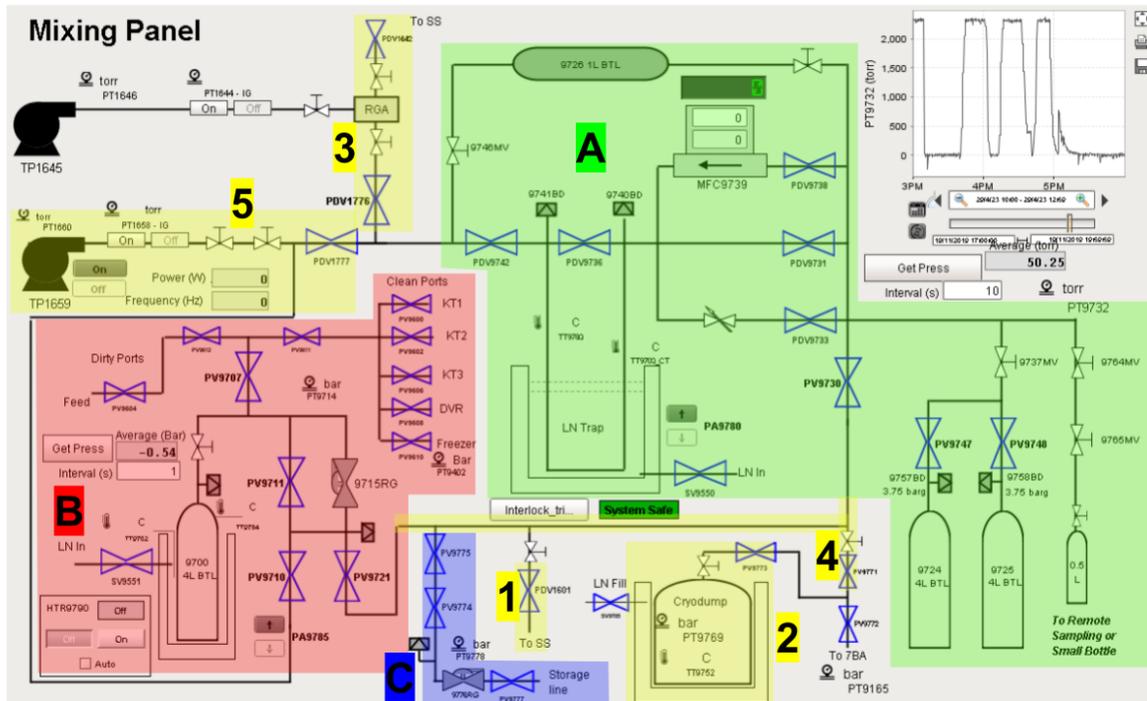


Figure 5.10: A screenshot of the interactive piping and instrumentation diagram from Ignition. Although all system tasks were automated, custom operations and trouble system trouble shooting require some manual control given here. Overlays show the specific subsystem: A) mixing panel in green, B) sniffer in red, and C) storage line in blue. The numbered, yellow overlays highlight shared devices: 1) calibration line, 2) cryodump, 3) RGA, 4) transfer line, and 5) utility pump.

5.3.3 Sampler Devices

While each subsystem is self contained, they rely on shared devices. The automation software handles how the devices are shared based on request queues. This ensures no two subsystems are using the same device at the same moment. The operation scripts support idle functions to wait in queue when a device is unavailable at the time of request.

1) Calibration line is to feed reference xenon with a known amount of krypton into the transfer line. A small amount of reference xenon is injected into the mixing panel and mixed with ultra-clean xenon to produce a calibration mixture

2) **Cryodump** is a 3.8 liter stainless steel bottle submerged in a dewar of liquid nitrogen.

When exposed to xenon gas, ice forms on the interior walls acting as a pump. This convenient method of moving xenon around is known as “cryo-pumping”. The cryodump is used to cryo-pump xenon from the rest of the sampler. Before maximum capacity is reached the cryodump is warmed and the recovered xenon is placed in a cylinder for reprocessing in the krypton removal system. The ultimate pressure achieved by cryo-pumping for a short period of time is less than 1 Torr.

3) **RGA** is used in high and low-sensitivity measurements of the xenon purity.

4) **Transfer line** connects each of the subsystems and devices to one another.

5) **Utility pump** is used to evacuate volumes to high vacuum in the event that a xenon sample was dirty enough that the residual 1 Torr left over from cryo-pumping to the cryodump could contaminate an upcoming sample. The utility pump is also used when sniffs are taken or when static pumpouts are performed.

5.3.4 Sampler Subsystems

The sampler subsystems shown are color coded in figure 5.10. This section gives a brief description of each subsystem and the automated operations it can complete.

5.3.4.1 Mixing Panel

The mixing panel subsystem is highlighted in the green A) portion of figure 5.10. This subsystem contains the XPM and automated liquid nitrogen cold-trap. Xenon from throughout the krypton removal system can be transferred into the mixing panel and measured.

Automated Mixing Panel Operations

Cryodump Between samples, the mixing panel and transfer line must be evacuated. Xenon is recovered from the mixing panel and transfer line in the cryodump. *Requires use of the transfer line and cryodump.*

Pump Out If a sample was dirtier than expected, the mixing panel and the storage line need to be pumped down to $< 10^{-4}$ Torr after the sample has been recovered by the cryodump. *Requires use of the transfer line and utility pump.*

Freezer Transfer To draw a sample from the freezer, a path must be set from the mixing panel through the transfer line, and the sniffer. The cold-trap is frozen to use as a cryo-pump and the mass flow controller regulates the total amount of xenon delivered. The cold-trap warm-up is monitored after the transfer is completed. *Requires use of the transfer line and a control hand-off from the sniffer, as well as manual coordination with the rest of the krypton removal system prior to starting the automated transfer.*

Storage Transfer To draw a sample from the storage line, a path must be set from the mixing panel, through the transfer line, to the storage line. The cold-trap is frozen to use as a cryo-pump and the mass flow controller regulates the total amount of xenon delivered. The cold-trap warm-up is monitored after the transfer is completed. While the cold-trap warms up, the cryodump is used to evacuate the storage line and transfer line. *Requires use of the transfer line, cryodump, and a control hand-off from the storage line, as well as manual coordination with the rest of the krypton removal system prior to starting the automated transfer.*

Sniffer Transfer To draw a sample from the sniffer, a path must be set from the mixing panel, through the transfer line, to the sniffer. The cold-trap is frozen to use as a cryo-

pump and the mass flow controller regulates the total amount of xenon delivered. The utility pump is used to perform a static pump-out of the xenon ice, which removes any excess helium after the transfer. The cold-trap warm-up is monitored after the transfer is completed. *Requires use of the transfer line, utility pump, and a control hand-off from the sniffer.*

Low-Sensitivity Run A low-sensitivity run is used to measure the helium, argon, and nitrogen concentrations. The cold-trap is frozen and exposed to the RGA. A leak valve delivers 5 grams of xenon to the cold trap over a 5 minute interval. The signal from the RGA during flow determines the concentrations. The cold-trap warm-up is monitored after the measurement. *Requires use of the RGA.*

High-Sensitivity Run A high-sensitivity run is used to measure the krypton concentration. The cold-trap is frozen and exposed to the RGA. If a large helium, argon, or nitrogen signal is detected prior to the flow start, the utility pump is requested, a static pump-out is completed, the trap is warmed, and the procedure starts over. If saturation is unlikely, the mass flow controller delivers 60 grams of xenon to the cold trap over a 5 minute interval. The signal from the RGA during flow determines the concentration. The cold-trap warm-up is monitored after the measurement. *Requires use of the RGA and, if needed, the utility pump.*

Static Pump Out A static pump-out is used to reduce the amount of helium, argon, and nitrogen present in a xenon sample so a high-sensitivity measurement is not saturated. The xenon is fed by the mass flow controller into a frozen cold-trap. Once all the xenon is exhausted, the ice is pumped on using the utility pump for 15 minutes. The cold-trap warm-up is monitored after the measurement. *Requires the use of the utility pump.*

Clean Run A clean run significantly reduces the amount of impurities in a xenon sample.

Clean runs are used to produce ultra-clean batches of xenon for preparing custom sample mixtures or for use in back-flow tests. The xenon is fed by the mass flow controller into a frozen cold-trap while the outlet of the cold-trap is actively pumped on by the utility pump. The cold-trap warm-up is monitored after the measurement.

Requires the use of the utility pump.

Mix High-Sensitivity Calibration To mix high sensitivity calibration, ultra-clean xenon

is prepared and mixed with reference xenon of a known krypton concentration. The calibration line and transfer line are pumped out with the utility pump. The cold-trap is frozen and the internal pressure transducer volume is emptied into the cold trap.

A path is set from the mixing panel, through the transfer line, up to the calibration line. A small amount of reference xenon is injected into the mixing panel and the pressure is recorded. The sample is mixed and the cold-trap warm-up is monitored.

The calibration and transfer lines are evacuated during warm-up to the cryodump followed by the utility pump to avoid contamination of other samples. *Requires use of the calibration line, transfer line, cryodump, and utility pump.*

5.3.4.2 Sniffer

The sniffer subsystem is highlighted in the red B) portion of figure 5.10. This subsystem contains a 3.8 liter stainless steel bottle fitted with heaters and a pneumatically actuated liquid nitrogen dewar shelf. This subsystem is used to sniff xenon from the krypton removal recovery loop. There is also a distribution manifold that can draw xenon from the freezer, KTs, and xenon chromatography feed.

A sniff requires coordination between the krypton removal system and sampler. Proper timing of the sniff is essential to accurately compare chromatography runs. To coordinate

a sniff, the krypton removal operator populates a column of the krypton removal run queue with the number of sniff cycles the storage sample will require. This depends on the helium to xenon ratio of the chromatography feed; usually seven sniffs are sufficient. The operator then prepares the sniffer by evacuating it if needed. The sniff operations can be launched up to six hours before the chromatography run begins. The operations are chained in a manner so the sniffer and sampler are ready again as quickly as possible, often with the sniffer collecting the next batch of xenon while the mixing panel measures the previous. The procedures queued for the sniffer after being pumped out are as follows: idle warm, freeze, idle cold, sniff, static pump-out, warm-up, transfer, cryodump, and pump-out. The corresponding procedures queued in the mixing panel starting after being pumped out are: transfer sniffer, high-sensitivity measurement, cryodump, and pump-out. These operation recipes can be chained as needed to maximize measurement throughput of column sniffs without operator intervention.

Automated Sniffer Operations

Cryodump Between fills from the krypton removal system, the sniffer must be evacuated.

Xenon is recovered from the sniffer and transfer line in the cryodump. *Requires use of the transfer line and cryodump.*

Pump out If xenon was dirtier than expected, the sniffer needs to be pumped down to $< 10^{-4}$ Torr after the sample has been recovered by the cryodump. *Requires use of the utility pump.*

Idle Warm The idle warm is used in conjunction with the krypton removal system to properly time sniffs during recovery. This operation ends when a chromatography run begins with a sniff request tag set.

Freeze Once the idle warm operation has ended, the sniffer is submerged in liquid nitrogen and frozen in preparation of a sniff.

Idle Cold Once the sniffer is frozen, the utility pump is requested and the turbo molecular pump is spun down. The frozen sniffer and readied utility pump signal to the krypton removal system that the sniffer is prepared. Once the charcoal column enters the recovery loop, the krypton removal system triggers the sniff operation. *This operation requires the utility pump.*

Sniff When the sniff is triggered, a preset amount of sniffs, set by the krypton removal run control, are completed to collect xenon from the helium-rich recovery stream. This is done by opening a path from the sniffer to the DVR outlet. After 60 seconds, the path is closed and the sniffer is exposed to the utility pump for 30 seconds. This removes bulk helium while the xenon collects as ice in the sniffer. These sniffs are repeated a pre-determined amount of times based on the relative quantities of xenon and helium used for that run. Usually 7 sniffs are sufficient to collect approximately 100 grams. Once all the sniffs are completed, the sniff operation is done. *This operation requires the utility pump.*

Static Pump-out The sniffed xenon ice is pumped on with the utility pump as turbo pump is brought back up to speed. The high vacuum pump-out continues for 45 minutes. *This operation requires the utility pump.*

Warm-up After a static pump-out, the sniffer can be warmed. The liquid nitrogen dewar is lowered using the pneumatic shelf and electric heaters bring the bottle back up to room temperature.

Transfer The transfer operation is an idle state used in conjunction with the mixing panel's sniffer transfer or freezer transfer operation. This operation uses a transfer tag to

signal to the mixing panel that the sniffer is ready for a control hand-off. Once the hand-off is underway the sniffer waits for the transfer tag to signal the completion of the transfer operation.

5.3.4.3 Storage Line

The storage line is a section of plumbing that leads to the storage packs where the processed xenon is kept following the freezer warm-up. After the processed xenon is moved to the storage pack via the compressor, samples from individual storage bottles can be drawn through the storage line.

Automated Storage Line Operations

Cryodump Between fills from the storage bottles, the storage line must be evacuated.

Xenon is recovered from the storage line and transfer line in the cryodump. Requires use of the transfer line and cryodump.

Transfer The transfer operation is an idle state used in conjunction with the mixing panel's storage line transfer operation. This operation uses a transfer tag to signal to the mixing panel subsystem the storage line subsystem is ready for a control hand-off. Once the hand-off is underway the storage line waits for the transfer tag to signal the completion of the transfer operation.

5.3.4.4 Liquid Nitrogen Distribution

Automated liquid nitrogen distribution is required for proper use of the cold-trap, sniffer operation, and cryodump. Due to the hazardous nature and system-wide dependence on liquid nitrogen, dispensation was handled with PLC scripts outside of Ignition. This reduces potential loss of cooling to critical systems in the event of an Ignition server outage.

In order for sampler scripts that use liquid nitrogen to run, a check requires that krypton removal system's liquid nitrogen reserves be at least 30% full.

Liquid nitrogen is dispensed through insulated copper lines and electronically-actuated, cryogenic valves. Boolean Ignition tags launch the PLC scripts used for the sampling system. The PLC script monitors a temperature sensor as the dewars are filled. The shelves for the cold-trap and sniffer dewars must be raised in order for the fill to trigger. For the sniffer, the heater must be off for the fill to trigger as well. If the feedbacking temperature sensor does not report liquid nitrogen temperatures within a set time interval the operation is aborted and the user is notified.

In the liquid nitrogen auto-fill GUI, each component has an auto-fill on/off button above a small interactive P&ID. When auto-fill is on, liquid nitrogen is dispensed until the temperature sensor reads the auto-fill stop temperature. If the post-fill timer is populated, the fill will continue for a set amount of time in seconds after the stop temperature is achieved. Once the temperature falls below the auto-fill start temperature, the dewar will begin to fill again. A plot of the relevant temperatures and pressure of each device is shown above the auto-fill control button.

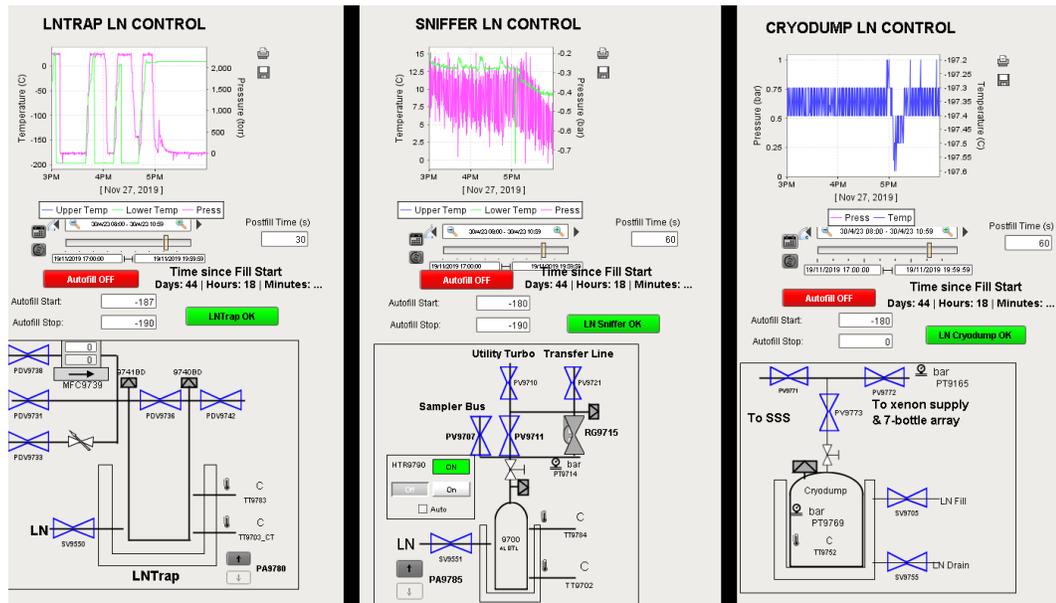


Figure 5.11: A screenshot of the sampler's liquid nitrogen auto-fill GUI.

5.3.5 Run Viewer

The sampler's run viewer is an Ignition GUI, shown in figure 5.12. Through the run viewer, any user can access all of the historical run data acquired by the sampler since implementation in August of 2020. Runs can be added retroactively or edited by directly modifying the run history dataset tag and filling out the appropriate parameters, in the Ignition Designer portal.

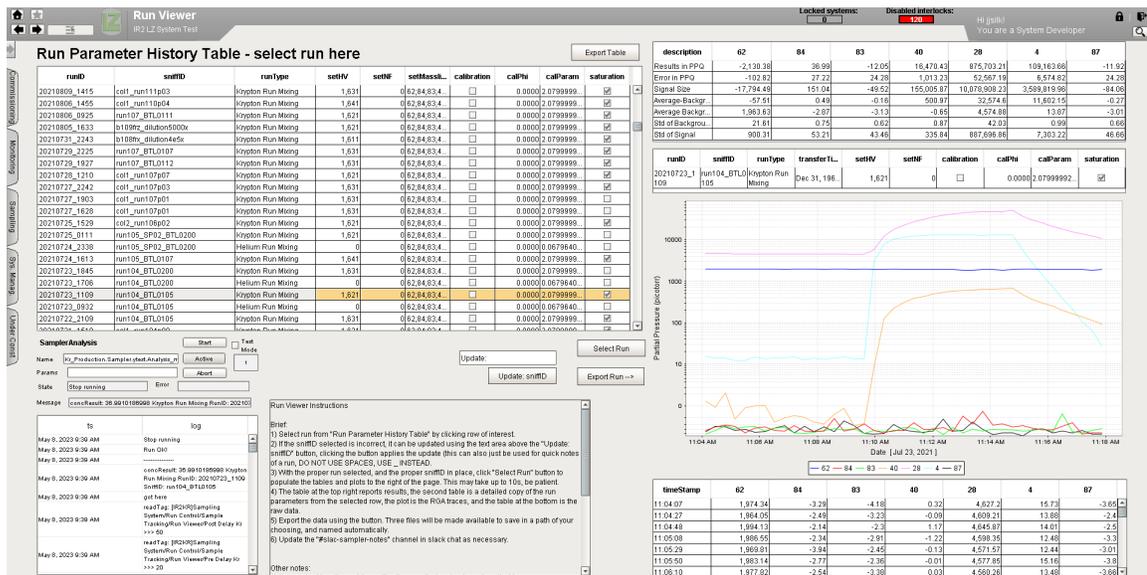


Figure 5.12: A screenshot of the sampler’s run viewer GUI. The table on the left are selectable runs that can be analyzed by clicking select run. The results and raw RGA data can be exported to a csv.

During a measurement operation, the script labels the measurements with a unique runID. A run type variable labels the measurement as low or high-sensitivity. If the xenon was transferred from the sniffer to the mixing panel, a sniffID parameter identifies what chromatography run and column the sniff was taken from. The sniffID can also be manually set after a measurement run to described where the xenon sample was drawn from. Timestamps identify the flow and background intervals used to fetch data from the data historian. If the measurement immediately follows a calibration mixing operation, a calPhi variable equal to the xenon mixture’s true krypton concentration is saved. A calParam variable is used in analysis to convert the raw RGA signal into a concentration value for the measurement. The calParam variable is derived from measurements of xenon with a known calPhi. The calParam for low and high-sensitivity measurements is different. Additional RGA parameters are tracked to monitor the RGA EM high voltage over time.

5.3.6 Freezer Warm-up as a Distillation step

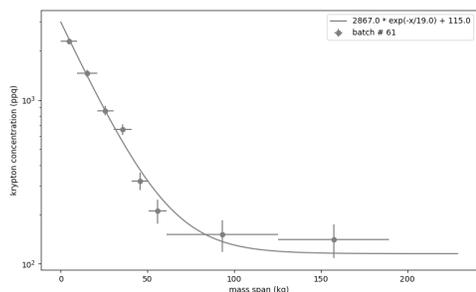
During early commissioning, measurements of the xenon purity were taken from the freezer during warm-up. An early sample and late sample were taken to compare the xenon purity with in the freezer with xenon collected during recovery sniffs. Xenon from early freezer measurements consistently contained more krypton than expected based on the sniff measurements of each xenon slug that composed the freezer batch. The late freezer measurements were noted as being consistent with or cleaner than the preceding sniff measurements. The first record of this behavior was in November of 2019 when consecutive freezer measurements were taken and the krypton concentration reduced from 56 to 11 ppt during a single warm-up. This behavior continued into production with reduction factors ranging from 2 to 100 and depended heavily on how early the first sample was taken during the freezer warm-up.

While warming, the xenon forms a liquid puddle at the bottom of the freezer. At -98 °C the volatility of krypton compared to xenon is a factor of 10.8 larger [116]. As the liquid is heated to maintain constant pressure in the freezer, the more volatile krypton is more likely to migrate from the liquid into the gas phase. The continuous removal and replenishment of the gas volume leads to krypton depletion in the liquid phase. As a result the krypton concentration drops as a function of the total xenon mass removed from the freezer. Essentially the freezer is acting as a simple xenon distillation column. Distillation of xenon alone is a viable technology as demonstrated by the XENONnT experiment. Their in situ distillation system has purified xenon to a krypton concentration of 56 ± 36 ppq (mol/mol) [117] [118].

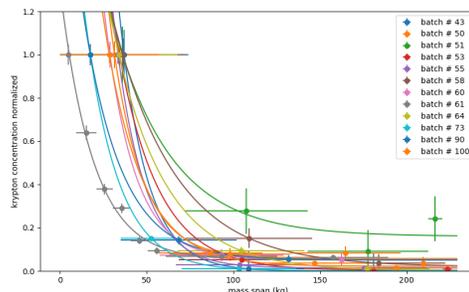
During the krypton removal campaign, an unidentified krypton source was introduced. Sniffer measurements at the outlet of the DVR noted an increased krypton concentration in April 2021. Despite attempts to recondition components upstream of the DVR outlet and

the DVR itself, the source was not reduced. A downstream purification step was required for the krypton removal system in order to meet LZ's xenon purity standards. The freezer's distillation process was an obvious candidate, and so a more thorough study of the freezer behavior was conducted.

Freezer storage batches were identified for analysis and one additional batch was used as a devoted study. Each of the batches considered starts with initially empty storage cylinders and at least three must be filled sequentially during storage from the same freezer batch. The mass deposited in each bottle was tracked and the krypton purity of each bottle was recorded. This information is used to produce plots, shown in figure 5.3.6, tracking the freezer purity as a function of the accumulated xenon storage mass. Batches ranged in initial impurity from 19 ± 0.7 ppt down to 290 ± 40 ppq achieving an ultimate purity from 190 ± 35 ppq to measurements consistent with $0 + 28$ ppq.



(a) Freezer batch 61. The y-axis is krypton concentration in units of ppq.



(b) Twelve freezer batches. The y-axis is krypton concentration normalized to the first concentration measurement.

Figure 5.13: Freezer batch krypton concentration as a function of mass spanned. The x-error bars represent the total mass contained in a single bottle, and the point corresponds to the mean value of the mass. The x-axis is in units of xenon mass out of the freezer. For example the first bottle in figure 5.13a is filled with the first 9.5 kg from freezer batch #61. The x-error bar spans from 0 to 9.5 kg and the data point is centered at 4.75 kg. The second bottle is filled with the next 11.5kg, has x-error bars from 9.5 to 20 kg, and the data point is centered at 15.25 kg.

The devoted study, shown in figure 5.13a, sliced the freezer batch into eight initially empty storage cylinders. The first six bottles were filled with about 10 kg of xenon each

and the remaining two bottles were filled with 65 kg. Each run is fit using

$$\Phi_{Kr} = a \cdot e^{-x/b} + c \quad (5.1)$$

to find a “halving mass”. This is the cumulative mass removed from the freezer that reduces the krypton concentration by half. The batch 61 study finds a halving mass of 13.2 ± 1.2 kg and an ultimate purity floor of 115 ± 35 ppq. For all twelve runs an average halving mass of 15.6 ± 4.9 kg is found with an ultimate purity floor of 50 ± 47 ppq.

This halving mass was used to selectively separate the first 10-30 kg during freezer storage and preserve the purity of the remaining 90%. Once storage was finished the dirtiest xenon is reprocessed as needed. This technique kept xenon production moving at a pace that ultimately met both the purity goals and time requirements for LZ.

5.3.7 Sampler Usage

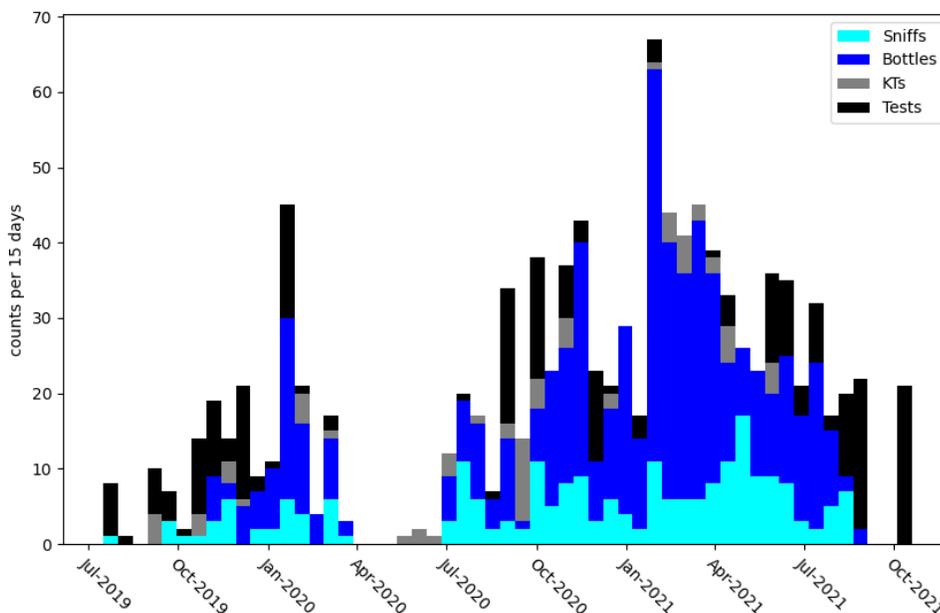


Figure 5.14: A histogram of all of the measurements taken during the course of the krypton removal campaign. Sniffer measurements, in cyan, are used for feedback for chromatography parameters. Figure 5.15 shows a more detailed view of some of the sniffer measurement results. Bottle measurements, in blue, check the final purity xenon. KT measurements, in grey, check the purity of krypton traps that maintain helium purity and capture krypton. Test measurements, in black, correspond to specialty measurements like calibrations, zero measurements, or DVR oil space checks.

Between July 2019 and September 2021 the sampler performed over 1000 measurements, as shown in figure 5.14. By the end of March 2020 high purity chromatography parameters had been identified, this critical group of sniff measurements are shown in 5.15. Between April 2020 and July 2020, SLAC access was restricted to essential personnel only and the krypton removal campaign was placed on hold. As more personnel were allowed back on site, the KT's were cleaned and measured to prepare for recommissioning. In July 2021 the krypton removal system was brought back online and by December 2021, fully optimized for both purity and xenon throughput. There was a dramatic increase in

sampler throughput to keep up with the high processing rate of the krypton removal system. In April 2021 an unidentified krypton source caused an irreducible increase in krypton as shown sniffer measurements shown in 5.15. Despite this, the distillation effect described in the 5.3.6 was used to maintain xenon purity standards and ultimately meet the purity goal set by LZ.

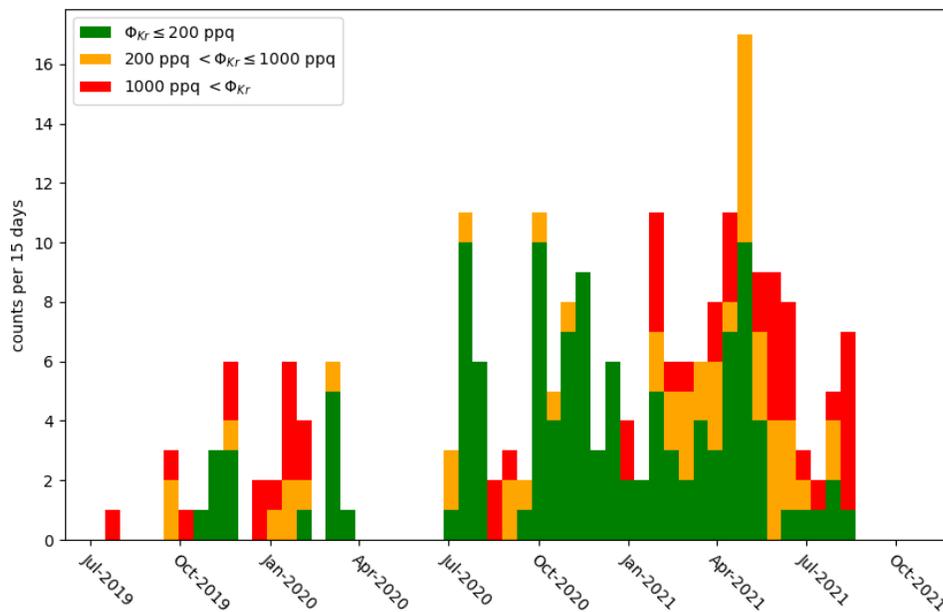


Figure 5.15: A histogram of all of the sniff measurements taken during the course of the krypton removal campaign. Measurements are grouped into three colors. The green measurements are sniffs that had krypton concentrations less than 200 ppq and measure an ideal chromatography result. Results between 200 ppq and 1000 ppq are shown in orange. These are results are not ideal, but can still produce reasonably clean batches of xenon. These intermediate purity results allowed operators to make informed decisions on whether or not to continue running. Red sniffer measurements have krypton concentrations greater than 1000 ppq and usually signal a larger issue. This information was used to halt production runs as needed to address purity issues by cleaning traps, or the DVR.

5.4 Results

Between January 2021 and August 2021, 10.4 tonnes of xenon were processed and in September 2021 condensed in the LZ detector. All of the 144 cylinders were individually

assayed at SLAC and again at SURF upon arrival. The accumulation of purified xenon at SLAC as a function of time is shown in figure 5.16.

The final mass-averaged concentration measured at SLAC prior to shipping was 123 ± 10 ppq g/g $^{nat}\text{Kr}/^{nat}\text{Xe}$. The mass-averaged concentration measured at SURF after arrival was 115 ± 17 ppq g/g $^{nat}\text{Kr}/^{nat}\text{Xe}$. Once the xenon was condensed and combined with residual xenon used for circulation commissioning the overall purity was measured to be 144 ± 22 ppq. The final mass-averaged argon concentration was measured to be 890 ± 134 ppt. The krypton removal campaign has successfully met all of the xenon purification requirements for LZ.

Xenon Purity and Mass Stored

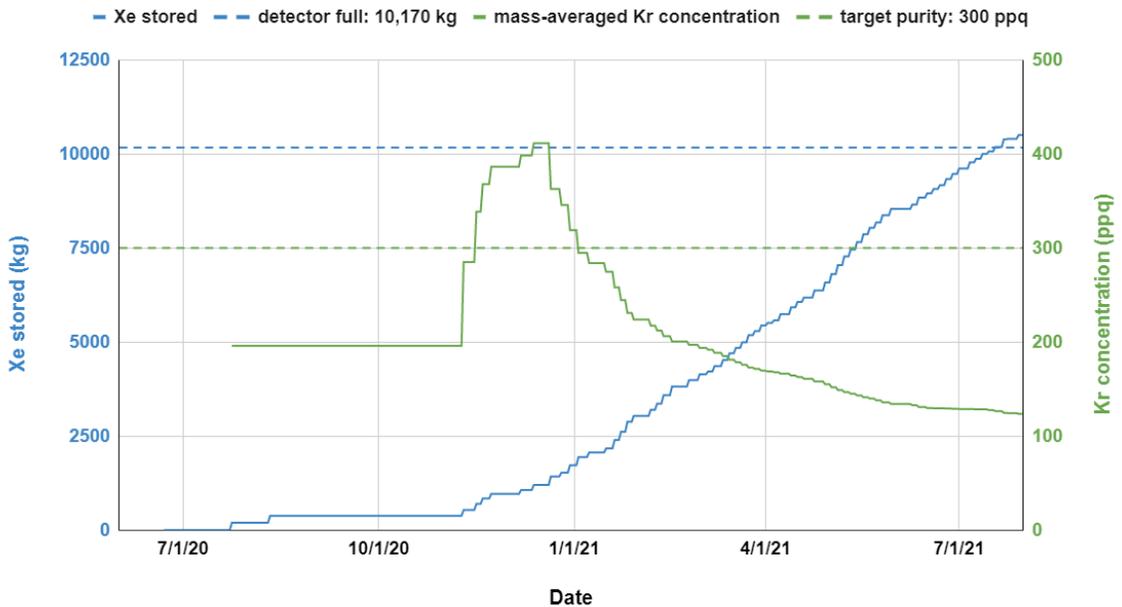


Figure 5.16: The cumulative mass of stored xenon and the mass averaged purity versus time. The green solid green line shows the mass-averaged krypton concentration and the dotted green line shows the target concentration. The blue line shows the cumulative mass of purified xenon and the dotted blue line shows the approximate mass of xenon required to fill the TPC. From March 2020 to July 2020, access to SLAC facilities was limited to essential personnel only due to COVID. Between July 2020 and December 2020 the system was recommissioned prior to full production that began in January 2021.

The operation and automation of the krypton removal campaign was a success. All of the xenon was processed to LZ specifications and arrived at SURF. The measurements upon receipt and after filling the LZ detector agree with the measurements taken at SLAC. Throughout the campaign, the sampler provided essential and reliable feedback. The automation of the sampler subsystem maintained high throughput keeping pace with the krypton removal system. During the sampler's operation, only two xenon loss events occurred due to operator error. Following sampler automation, no xenon loss events occurred. The automation made system operation accessible to users with minimal training. The run viewer made the data accessible to any user, trained or not, and maintained an organized history of results.

Chapter 6: Direct Observation of ^{85}Kr in LZ Science Run 1

In addition to in situ xenon purity measurements of krypton concentrations, ^{85}Kr is also identified in LZ Science Run 1 (SR1) data via a unique delayed coincidence signature. The relevant decay is characterized by a beta decay from ^{85}Kr to meta-stable ^{85m}Rb . The meta-stable excited state has a half-life of $1.015\ \mu\text{s}$, and upon relaxing emits a gamma. The beta decay has an endpoint energy of 173 keV and the delayed gamma has an energy of 514 keV. The branching ratio is 0.434%. The decay scheme is shown in 6.1.

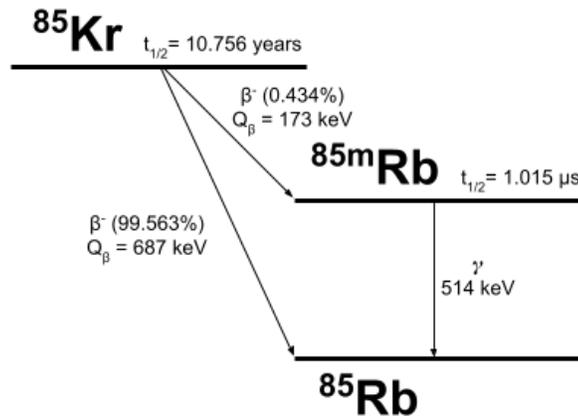


Figure 6.1: Decay scheme of ^{85}Kr . Decay time and energies taken from [29].

The mass-averaged krypton concentration over the course of SR1 was measured via the XPM to be 144 ± 22 ppq. This concentration can be used to estimate the number of excited

state ^{85}Kr decays in LZ's 5.5 tonne fiducial volume during the 100.3 day SR1 exposure. The total decay rate of ^{85}Kr is given by

$$\frac{dN}{dt} = -\frac{\ln(2)}{t_{1/2}}N \quad (6.1)$$

where N is the number of ^{85}Kr atoms. The concentration (Φ_{Kr}) is reported in units of $^{nat}\text{Kr}/^{nat}\text{Xe}$ g/g and must be converted to N using the total xenon mass m_{Xe} , isotopic abundance of ^{85}Kr to ^{nat}Kr ($m_{85\text{Kr}}/m_{nat\text{Kr}}$), the molar mass $M_{85\text{Kr}}$, and Avogadro's number N_A

$$N = \Phi_{\text{Kr}} m_{\text{Xe}} \frac{m_{85\text{Kr}}}{m_{nat\text{Kr}}} \frac{N_A}{M_{85\text{Kr}}}. \quad (6.2)$$

Putting all of this together, we find a total ^{85}Kr decay rate of

$$\frac{dN}{dt} = -\frac{\ln(2)}{10.856 \text{ years}} \cdot 144 \text{ ppq} \cdot 5.5 \text{ ton} \cdot 2 \times 10^{-11} \frac{6.022 \times 10^{23} \text{ atom}}{85 \text{ gram}} = 19.8 \frac{\text{decays}}{\text{day}}. \quad (6.3)$$

In 100.3 days we expect a total of 1986 decays. Of these, only 0.434% are excited state decays, for a total of $N_{excited} = 8.6$ events. Despite this modest event rate, the decay signature offers a powerful selection criteria unique to ^{85}Kr in this energy range.

This chapter describes the data-driven search for ^{85}Kr excited state decays in SR1. The 1.5 tonne-year exposure takes place from December 23rd, 2021 to April 18th, 2022. Two simulation techniques are used to inform cuts, efficiencies, and the detector response in terms of available RQs. These cuts are validated against mono-energetic events, readily identified in both simulations and SR1 data at energies ranging from 39.5 keV to 1.46 MeV. After validating the cuts, they are applied to the SR1 data. The final search looks for an excess of delayed coincidence signals around 1 μs and assumes a constant background.

6.1 What Will a Signal Look Like

From the decay scheme in figure 6.1 we can expect a ^{85}Kr signal to have at least two $S1$ signals, a smaller $S1_\beta$ that precedes a larger $S1_\gamma$, followed by one or more $S2$ s. The delay time between the $S1_\beta$ and $S1_\gamma$ corresponds to the ^{85m}Rb half-life of $1.015 \mu\text{s}$.

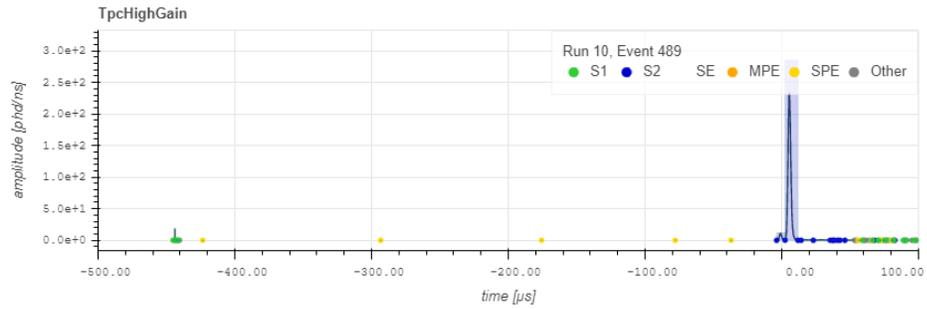
Naively we may expect two $S2$ s, one each from the beta and gamma. But there is the possibility that the gamma interaction occurs very close to the beta interaction. In this case the two electron clouds could combine, producing a merged $S2$ signal. The gamma may also Compton scatter, producing multiple interaction sites within the TPC. If this occurs, only a single $S1_\gamma$ is produced, because the time between the interactions is small, but multiple $S2$ s will be detected at different locations due to the unique drift times and xy locations. A 514 keV gamma has a total attenuation length of 3.6 cm in liquid xenon [119]. The LZ detector has an $S2$ position resolution better than 10 mm in the xy plane and 1 mm in z .

6.2 Simulated Waveforms

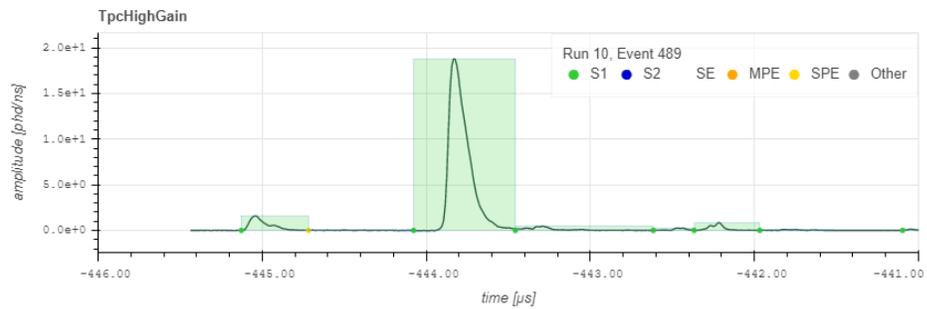
Our knowledge of the detector response is informed through a full simulation of excited state ^{85}Kr decays. With the full simulation technique, LZap produces waveforms that are studied with the LZ event viewer.

Pulse waveforms and the PMT hitmap from a simulated “golden” event are shown in the figures below. Event windows are defined from $-2000 \mu\text{s}$ to $2500 \mu\text{s}$ with respect to a global $S1$ or $S2$ pulse trigger. In figure 6.2a the window is resized from -500 to $100 \mu\text{s}$, to show the relevant $S1$ and $S2$ pulses areas. In figure 6.2b the window is resized to show the $S1_\beta$ ($445 \mu\text{s}$) and $S1_\gamma$ pulses ($444 \mu\text{s}$). The $S1$ pulse start times are separated by a delay time of $1.05 \mu\text{s}$. Note that the simulation includes a model of PMT after pulsing, which can

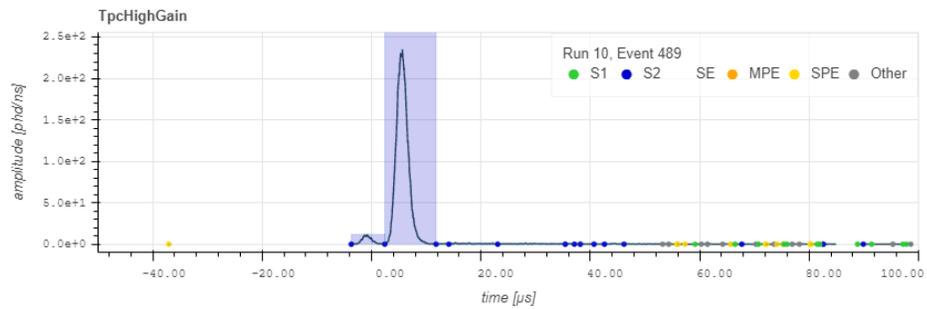
produce pulses comparable in size to the $S1_\beta$ pulse following the $S1_{\text{gamma}}$. In figure 6.2c the window is resized to show the $S2$ pulses. In this event, the $S2$ pulse produced by the beta ($0 \mu\text{s}$) is distinct from the $S2_{\text{max}}$ pulse ($4 \mu\text{s}$) produced by the gamma. This is a unique case and usually the $S2$ produced by the beta interaction cannot be readily identified due to pulse merging or gamma scatters. The time difference between $S1_\gamma$ and $S2_{\text{max}}$ is the reconstructed drift time and taken to be the gamma interaction's first z position. For this golden event, the drift time is $446.5 \mu\text{s}$, placing the event near the center of LZ. The position of the beta is neglected, only the $S1_{\text{gamma}}$ and $S2_{\text{max}}$ are used for position reconstruction. Figure 6.2d shows the light collection hit map of the top (left) and bottom (right) PMT arrays. Most of the light produced in the event is produced during charge extraction, contributing to the $S2$ signal and is concentrated around one PMT in the top array. This is used to reconstruct the gamma interaction position in x and y . For this golden event, the (x, y) position is $(-46.3, 22.1)$ cm.



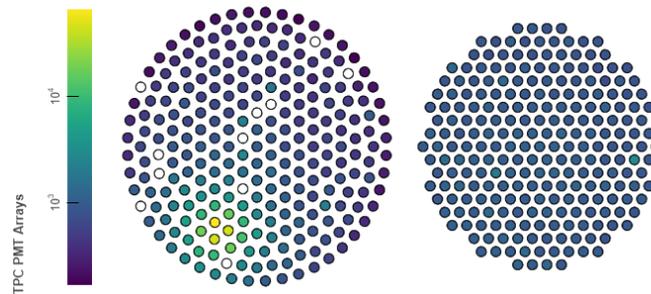
(a) Simulated event pulses between $-500 \mu\text{s}$ and $100 \mu\text{s}$. The green dots are $S1$ pulses ($-440 \mu\text{s}$) and the blue are $S2$ pulses ($0 \mu\text{s}$).



(b) A zoom in of the $S1$ pulse areas. The $S1_\beta$ pulse occurs before the larger $S1_\gamma$ pulse.



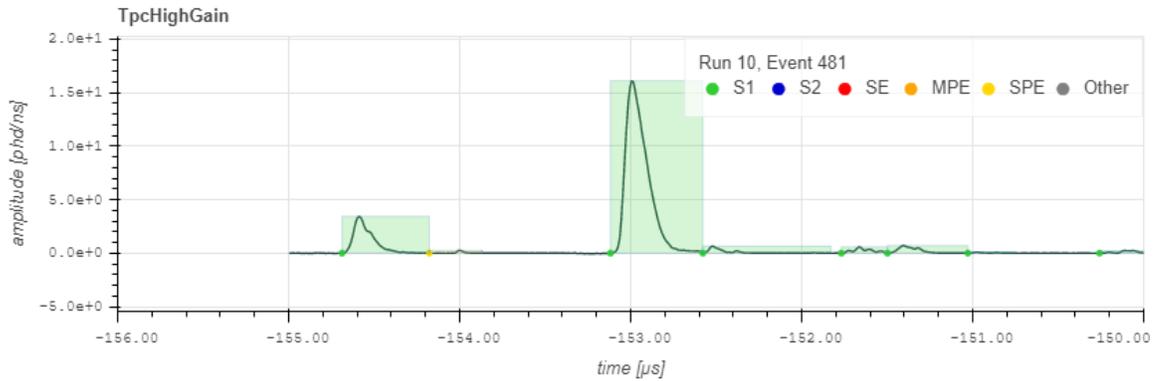
(c) A zoom in of the $S2$ pulse areas. Some PMT after pulsing and single photoelectrons are observed following the $S2_{max}$ pulse.



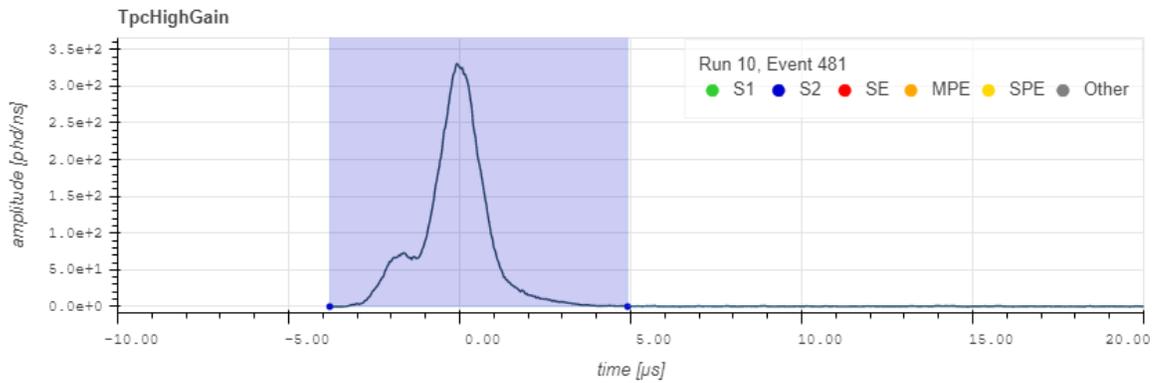
(d) The top and bottom PMT array hitmaps.

Figure 6.2: Full simulation of a golden event displayed in the LZ event viewer.

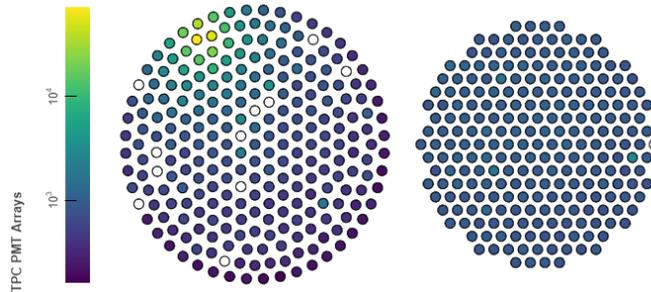
In events where the gamma interaction occurs very close to the beta interaction, the $S2$ pulses can merge. This is observed in the simulated waveforms shown in figure 6.3b. The event still produces identifiable $S1_\beta$ and $S1_\gamma$ pulses. In this example the delay time is found to be $1.57 \mu\text{s}$. The combined $S2$ pulse is treated as the $S2_{max}$ pulse. The drift time is $149.3 \mu\text{s}$ and the (x,y) position is (58.5, 29.7) cm. The $S2$ pulse contains most of the photons which are concentrated around one PMT in the top array, as shown in figure 6.3c.



(a) The $S1$ pulses are still identifiable as $S1_\beta$ and $S1_\gamma$ despite merged $S2$ pulses.



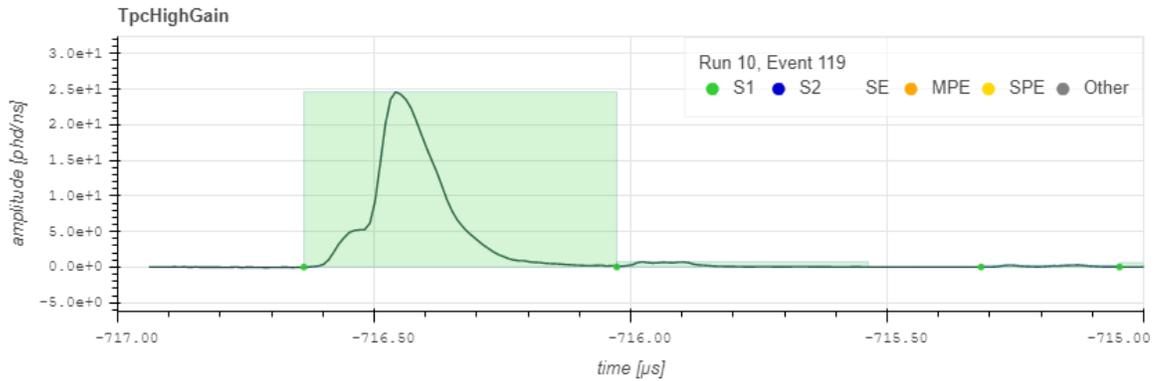
(b) Simulated $S2$ pulses which have merged.



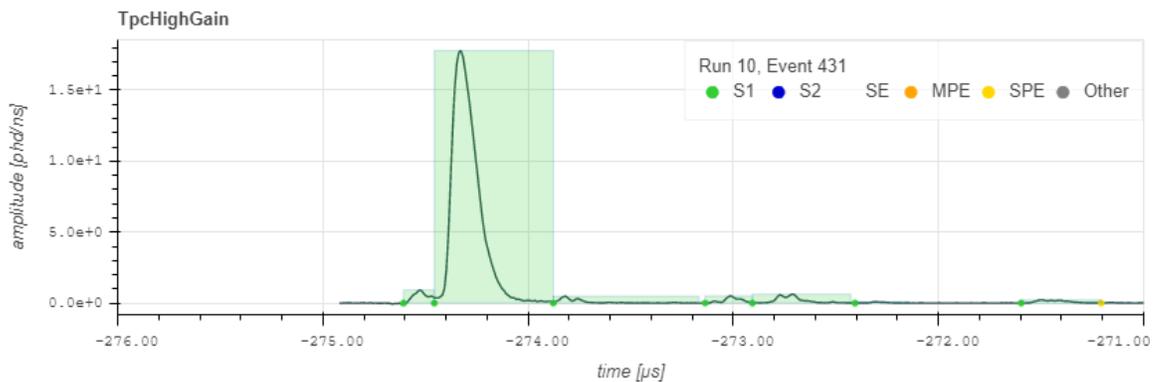
(c) The top and bottom PMT array hitmaps.

Figure 6.3: Full simulation of an event with merged $S2$ signals displayed in the LZ event viewer.

$S1$ pulses can also merge if the time between decays is less than $0.15 \mu\text{s}$, as shown in 6.4a. In this case, the event can not be fully reconstructed, because LZap has not distinguished the $S1_\beta$ and $S1_\gamma$ pulses required to calculate a delay time. Figure 6.4b shows an event with a delay time of $0.15 \mu\text{s}$ in which LZap has identified each pulse.



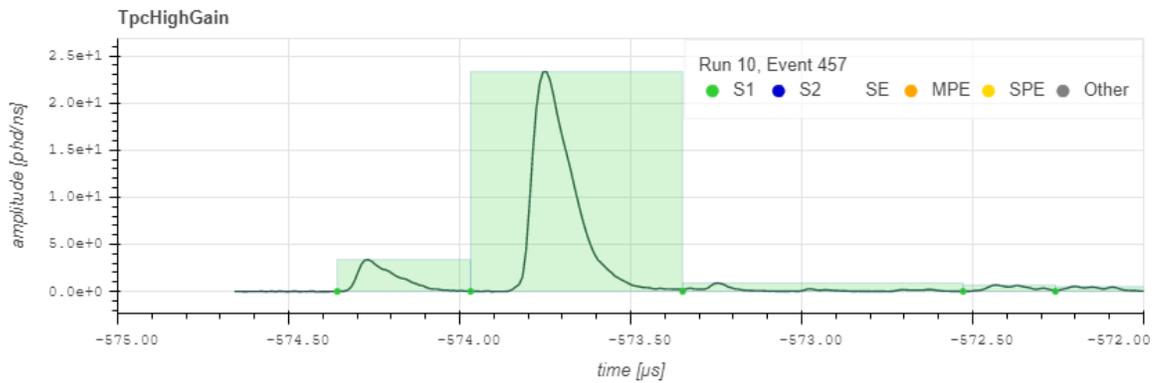
(a) A simulated event with merged $S1$ pulses. Without distinct $S1_\beta$ and $S1_\gamma$ pulses, the event delay time cannot be calculated.



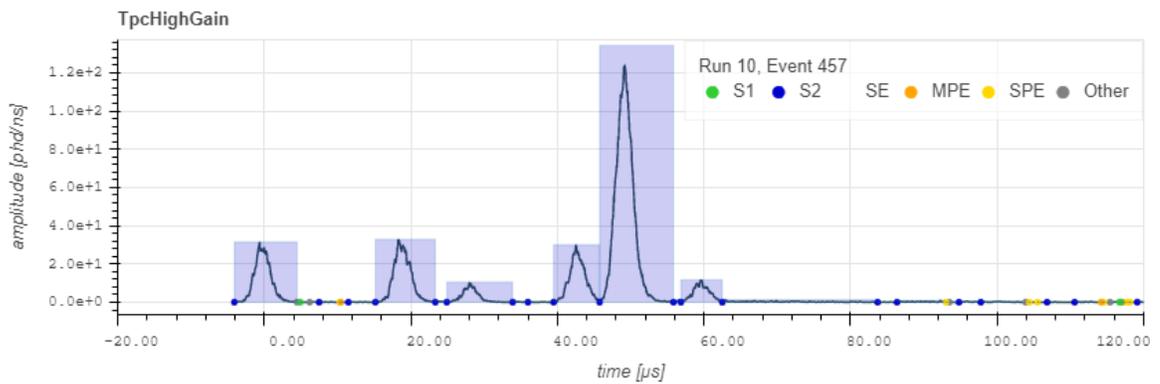
(b) A different simulated event with a delay time of $0.15 \mu\text{s}$.

Figure 6.4: Full simulation of two events, one with merged $S1$ pulses and another with a delay time of $0.15 \mu\text{s}$ displayed in the LZ event viewer.

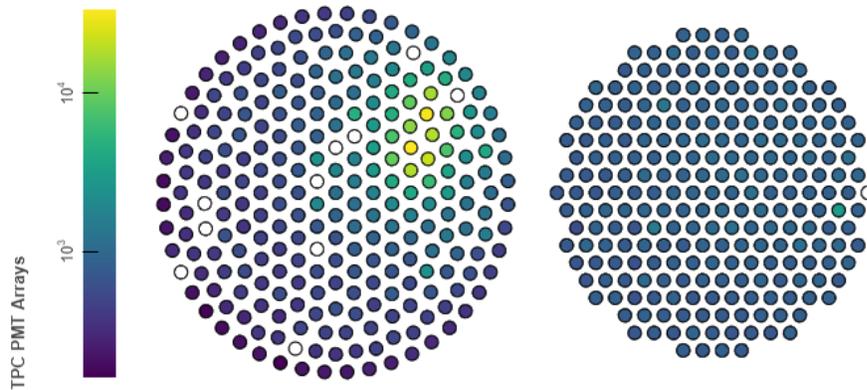
Another type of event observed in the simulated waveforms, shown in figure 6.5b, contains many $S2$ pulses due to Compton scattering of the gamma. The $S1_\gamma$ pulse, shown in figure 6.5a, is unchanged as the scatters occur simultaneously and the gamma remains within the TPC. The $S2_{max}$ pulse is still used to reconstruct the event position, despite each $S2$ having a unique position. Figure 6.5c shows that the $S2$ light is dispersed due to the multiple scatters and is no longer concentrated in a single PMT.



(a) The $S1_\beta$ and $S1_\gamma$ pulses of a simulated event with multiple gamma scatters. The $S1_\gamma$ event is indistinguishable from other $S1_\gamma$ s despite scattering.



(b) Multiple $S2$ signals observed within the TPC during a simulated event with gamma scattering.

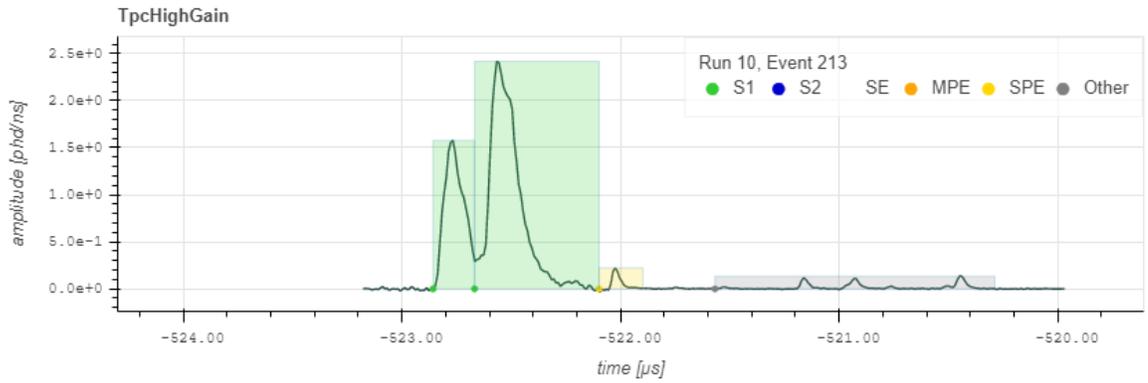


(c) The PMT hit map of the top array (left) is not concentrated in a single PMT, but spread over multiple PMTs with two containing most of the light.

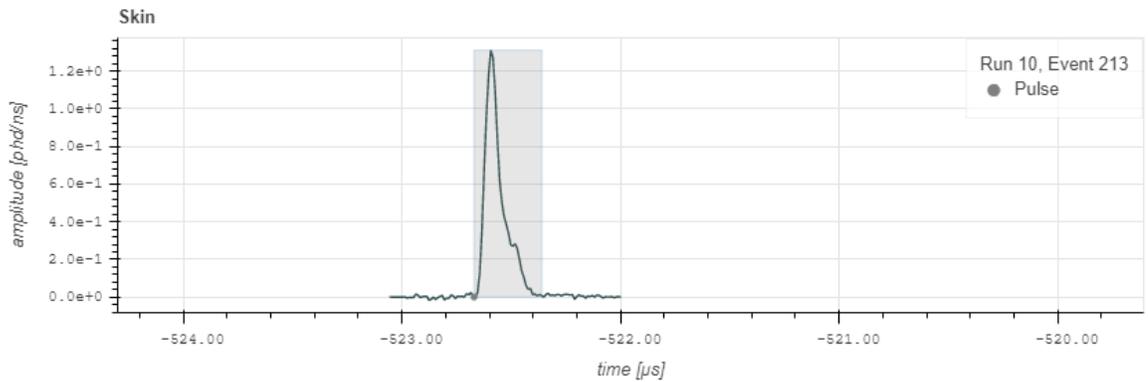
Figure 6.5: Full simulation of an event with multiple gamma scatters contained within the TPC.

The 514 keV gamma can travel an appreciable distance in liquid xenon and even leave

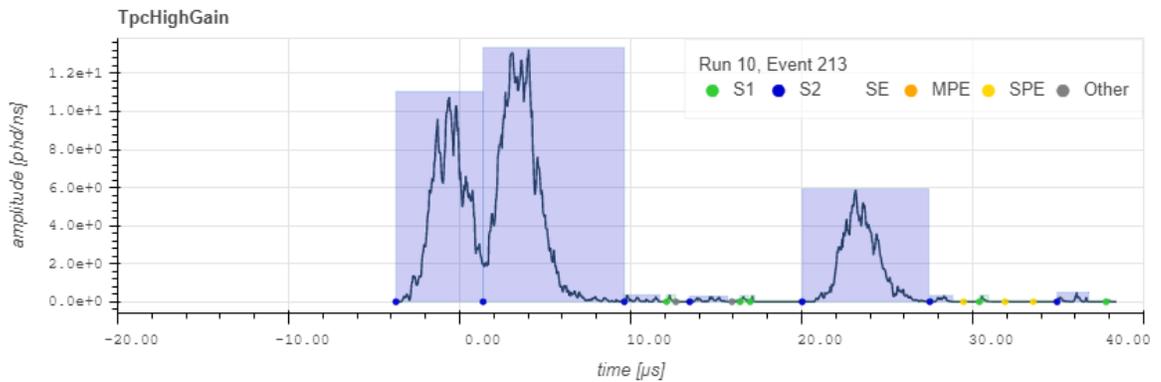
the TPC. The simulated waveforms include the skin veto detector response and capture events in which the gamma deposits energy in the skin and the OD. Figures 6.6a and 6.6b show a gamma that produces an $S1$ signal in both the TPC and xenon skin. The total $S1_\gamma$ light is split between the two regions due to the PTFE wall separating them. This results in less $S1_\gamma$ light observed in the TPC. The skin and TPC PMTs record light at the same time and in similar locations, near the edge of the detector, as shown in figures 6.6e and 6.6d.



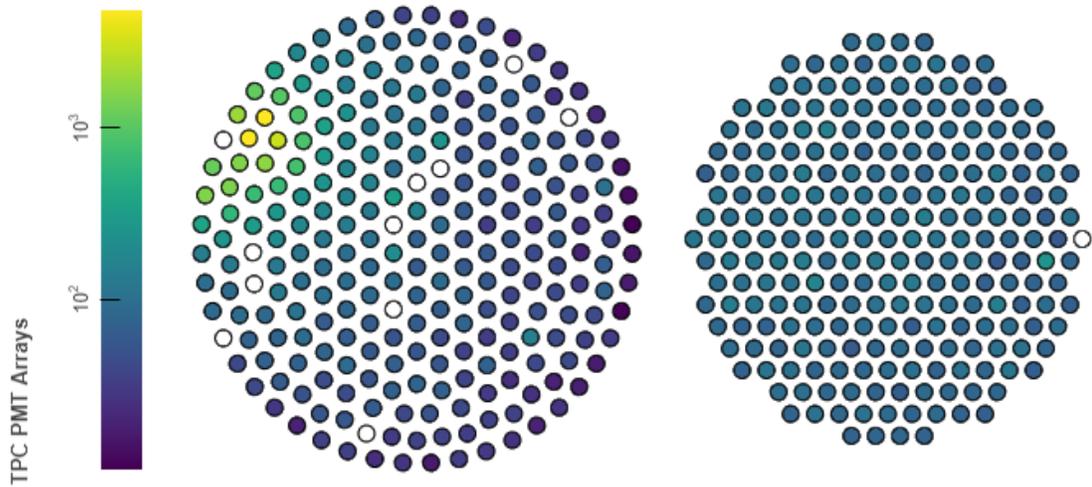
(a) The $S1_\beta$ and $S1_\gamma$ pulses of a simulated event that deposited energy in the xenon skin region. The $S1_\gamma$ observed in the TPC is smaller than those observed in the other cases.



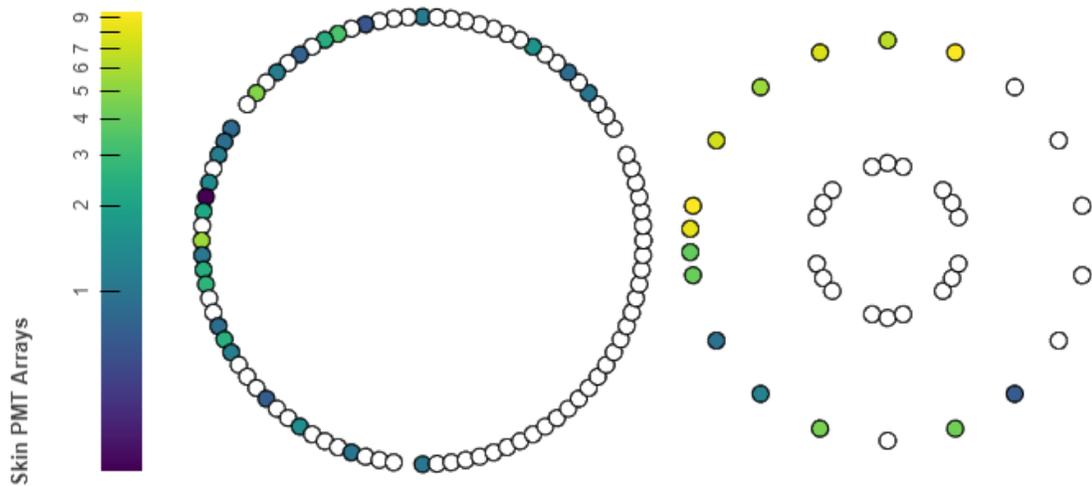
(b) The pulse detected in the skin at the same time as the TPC $S1_\gamma$.



(c) The $S2$ pulses detected within the TPC during an event that deposits gamma energy in the skin. The event has many $S2$ s but the total pulse area is less than previous cases because the charge freed in the skin region is not collected.



(d) The PMT hit map of the top array (left) is dispersed due to multiple scatters near the edge of the detector. The S2 light is spread out among multiple PMTs.



(e) The skin PMT hit map. The skin top (left) and bottom (right) PMT rings are located outside of the TPC, between field rings and cryostat wall. The light distribution in the skin PMT hitmap is similar to that of the top PMT array from figure 6.6d.

Figure 6.6: A full simulation of an event with energy deposited in the xenon skin region from a gamma scatter.

Due to the double S1s produced in an excited state ⁸⁵Kr decay, LZap identifies these event types as “pile-up” (multiple S2s) or “other” events (single S2), not “multi-scatter” or “single-scatter” events. When an event is classified as a single or multi-scatter there is a

clear association between the lone $S1$ pulse and one or more $S2$ pulse(s). This means that each $S2$ has a well defined drift time and xy position from which LZap can position based corrections. In the case where two distinct $S1$ s are present, LZap is not yet equipped to resolve the ambiguities between which $S1$ produced which $S2$ and therefore cannot correct the position. This limits the RQs used in this analysis to uncorrected or raw values. Even so, we can still pursue the analysis using uncorrected variables. However, those variables need to be validated by comparing data to Monte Carlo simulations.

6.3 Data Selection Variables

Given the various forms an event of interest can take, a cut structure must be developed to identify them. This is informed by the minimum number of pulses expected, the order of the pulses, the event location, pulse area equivalent energies, and the xenon skin veto response.

0) Selection Criteria - In order to reconstruct an event, we require that at least three signals are identified in a specific time order. An $S1_\beta$ pulse must occur before a larger $S1_\gamma$ pulse. Both $S1$ s occur before the event's largest $S2_{max}$ pulse. From these pulses the position of the interaction and the energy can be reconstructed for further analysis.

1) Fiducial Volume - We require an event to occur within the fiducial volume of the TPC. This reduces backgrounds from wall events and poor event reconstruction. The SR1 fiducial z position is enclosed by a drift time cut between 67 and 945 μ s. A radial cut is made through a standoff distance from detector walls of 5.2 cm for drift times between 67 and 200 μ s, 4.0 cm for drift times between 200 and 800 μ s, and 5.0 cm for drift times between 800 and 945 μ s. Two additional cylindrical cuts extend across all drift times with a diameter of 6 cm centered at (-71.2, 4.4) cm and (-69.2 -14.6) cm (x,y) locations to exclude light emission from field cage resistors.

2) Energy $S1_\beta$ - This cut is defined by the range of $S1$ pulse areas that correspond to a beta with a 173 keV endpoint energy. This includes a minimum pulse area requirement above zero phd to reduce contributions from small $S1$ backgrounds produced by grid emission.

3) Energy $S1_\gamma$ - This cut is defined by the range of $S1$ pulse areas that correspond to a gamma of 514 keV.

4) Energy $S2_{total}$ - This cut is defined by the range of total $S2$ pulse areas that correspond to an energy between 514 keV and 687 keV.

5) Energy $S1_{total}$ versus $S2_{total}$ - This final energy cut is defined by the diagonal edges of the $S1_{total}$ versus $S2_{total}$ pulse area and is equivalent to a total energy cut.

6) Delay time window - The delay time variable is where the final signal search is carried out. An excess of events with a delay time around $1 \mu s$ would correspond to a detection of ^{85}Kr excited state decays. This cut is selected to exclude merged $S1$ pulses and include random coincidence at delay times larger than the ^{85}Kr half-life to inform a background prediction. Random coincidence events are two unrelated events that happen within an event window. This could be due to a random beta or gamma with energy less than 173 keV occurring before a higher energy beta or gamma with energy around 514 keV.

7) Localization - “Localization” characterizes the spatial clustering of $S2$ pulses that corresponds to gamma scatters. This cut uses the $S2$ weighted pulse area position deviation from the largest $S2_{max}$ pulse location.

$$L = \frac{\sum_{n=1}^m S2_{area,n} \cdot |\vec{r}_1 - \vec{r}_n|}{\sum_{n=1}^m S2_{area,n}}. \quad (6.4)$$

Here m is the total number of $S2$ pulses and $n = 1$ is the $S2_{max}$ pulse

8) Delta TBA - The top bottom asymmetry (TBA) is a detector RQ based on the total

amount of light collected in the top and bottom PMT arrays.

$$TBA = \frac{Top - Bottom}{Top + Bottom} \quad (6.5)$$

Events that occur within the liquid xenon produce prompt scintillation photons that are reflected back into the TPC by the liquid surface. Therefore, photons produced within the liquid are more efficiently collected by the bottom PMT array and have a negative TBA value. The amount of light collected at the bottom array relative to the top array increases for S1s closer to the bottom. The difference between TBA values of the identified $S1_\beta$ and $S1_\gamma$ pulses can be used to ensure the beta and gamma events occur close to one another in z .

9) Skin Veto - This cut uses the skin veto to exclude events that would otherwise produce random coincidences from activity in the skin region or from high energy gammas scattering into or out of the TPC.

6.4 Determining Cut Values and Efficiencies via Simulations

With the variables defined in the previous section, we can determine what numerical values each data selection cut should take, as well as the resulting efficiencies. Two simulation tools are used to inform the cut values and efficiencies: fast simulation and full simulation. Fast simulation determines the selection criteria, fiducialization, energy cuts, delay time window, and localization cut. The full simulation determines the TBA cut and skin veto response. It is necessary to use both simulations because the fast simulation is more accurately tuned to the detector response, but does not contain RQs corresponding to the TBA or skin veto. The full simulation technique models the TBA and skin veto, but the energy response is not as well tuned to match data. We restrict the full simulation to only

inform the ΔTBA cut and validate the skin veto response.

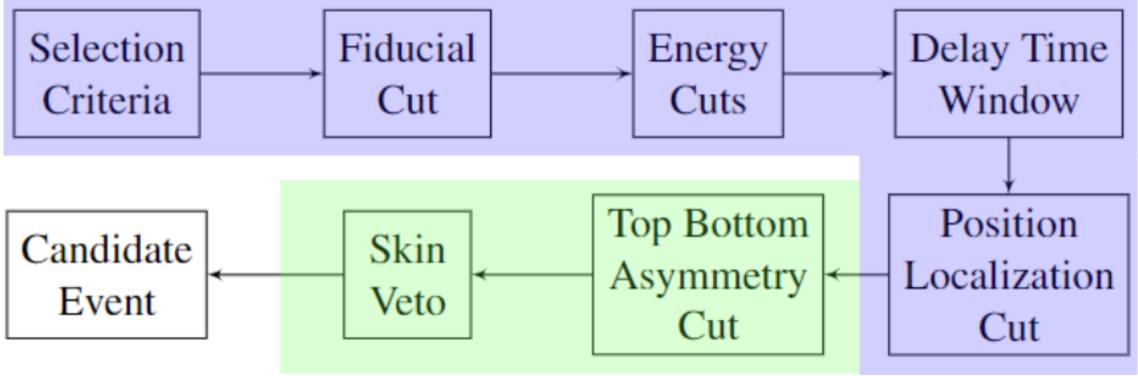


Figure 6.7: Analysis workflow between cuts. The section highlighted in blue is informed by the fast simulation technique and the the green is informed by the full simulation technique.

The workflow for this analysis is shown in figure 6.7. To determine the efficiency of these cuts we identify an initial number of events of events generated within the fiducial volume N_i . From this we count the number of events N_0 that pass our zeroth cut, the selection criteria. We repeat this process through each cut to yield a final set of N_f candidate events. Since the simulation makes use of two different techniques, separated between cut7 and cut8, the final efficiency ϵ_f is

$$\epsilon_f = \frac{N_{7,fast} N_{9,full}}{N_{i,fast} N_{7,full}}. \quad (6.6)$$

where the subscript denotes the simulation technique used to inform that cut.

The uncertainty of a cut efficiency is determined by modifying the cut values to produce a wider and narrower set of cuts. The modified cuts are applied to the simulated data and the average difference between the modified and original surviving number of events is taken to be the uncertainty. For example, a cut on an RQ with a simulated peak value C only accepts events between the bounds $[C_{lower}, C_{upper}]$. The number of events that survive this cut in simulations N_C . A comparison between SR1 and simulations show that the

simulated peak value C is off by σ_C from the observed SR1 value. The fractional error on the simulated peak value C is then taken to be $|\sigma_C|/C$. From this percent disagreement between simulations and SR1, two modified cuts are applied. A modified narrow cut is made with bounds $[C_{lower}(1 + |\sigma_C|/C), C_{upper}(1 - |\sigma_C|/C)]$ and produces a set of $N_{C,narrow}$ surviving events in the simulation. The modified wide cut is made with bounds $[C_{lower}(1 - |\sigma_C|/C), C_{upper}(1 + |\sigma_C|/C)]$ and produces a set of $N_{C,wide}$ surviving events in the simulation. The average fractional efficiency uncertainty for the cut on value C is then

$$\frac{\sigma_{N_C}}{N_C} = \frac{(N_{C,wide} - N_C) + (N_C - N_{C,narrow})}{2N_C} = \frac{N_{C,wide} - N_{C,narrow}}{2N_C}. \quad (6.7)$$

The final efficiency uncertainty is taken to be the combined percent errors of the initial number of events and each cut

$$\sigma_f^2 = \frac{\sigma_i^2}{N_i^2} + \frac{\sigma_2^2}{N_2^2} + \frac{\sigma_3^2}{N_3^2} + \frac{\sigma_4^2}{N_4^2} + \frac{\sigma_8^2}{N_8^2}. \quad (6.8)$$

The uncertainties were calculated from dedicated studies that compare simulations to real data and give confidence to values informed by the simulations.

6.4.1 Identifying Events in Simulations

In the fast simulation, ground state and excited state decays were generated, but we prefer that the ground state decays are excluded. Each simulated event has a set of MCtruth RQs that contain information about how the event was generated. In this case, a truth RQ identifies a particle at each interaction vertex as a beta, gamma, or ^{85}Rb . Excited state decays contain two ^{85}Rb vertices, one from the decay of ^{85}Kr to ^{85m}Rb and a second as ^{85m}Rb relaxes to its ground state. Of the 49,984,647 ^{85}Kr events generated in the fast simulation, 214,860 follow the excited state decay branch (0.43%), which is in agreement

with the expected branching fraction of 0.434%. A summary table of the Monte Carlo event counting corresponding to each fast simulation cut is provided in table 6.2.

The final search is restricted to fiducial volume. Events generated outside of this volume should not contribute to the initial number of events used to calculate the efficiency. These events are excluded by finding the truth x, y, and z positions and imposing the SR1 fiducial cut. This initial parent particle position is given in terms mms. To maintain consistency with SR1 analysis methods, an ALPACA function is used to apply the fiducial volume cut. This function accepts inputs for both x and y in cm, but z must be converted to an equivalent truth drift time value.

To convert the z_{truth} position to a truth-equivalent drift time, a linear fit is made between z_{truth} and the reconstructed drift time $t_{drift} = S2_{max,time} - S1_{\gamma,time}$. From this fit an equivalent t_{truth} drift time is found for each event. This allows the fiducial cut to be applied to every simulated event, even if that event lacks sufficient information for position reconstruction. The number of events remaining within the fiducial truth volume is $N_i = 147,175$.

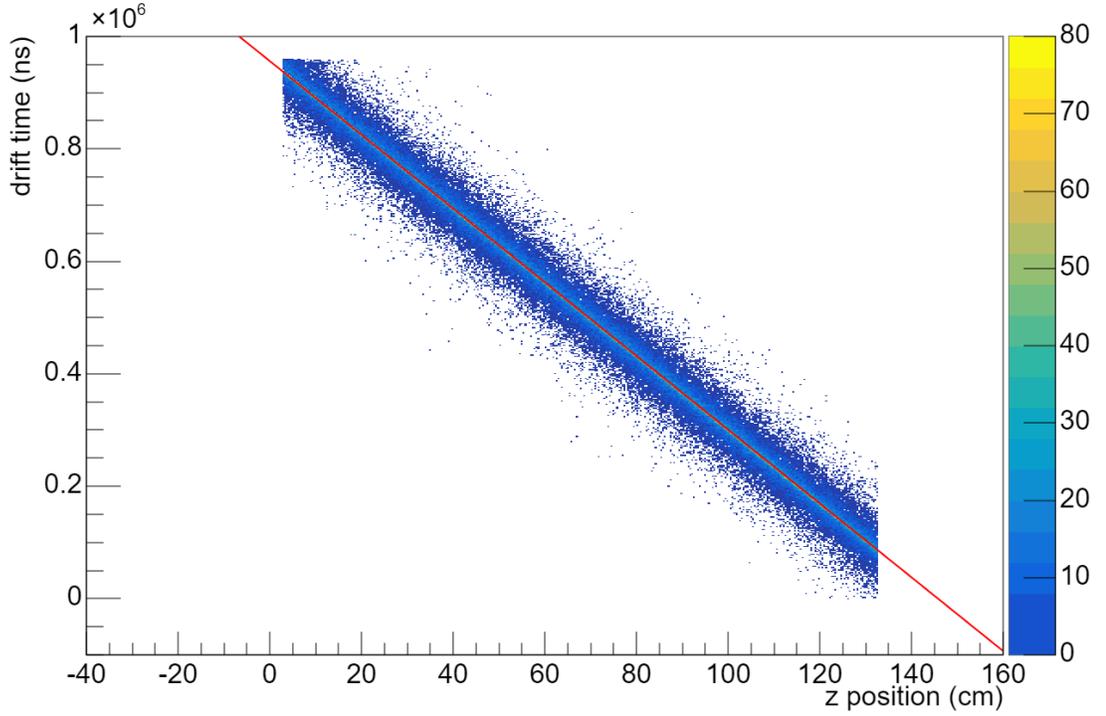


Figure 6.8: The MCtruth z_{truth} position (cm) versus the reconstructed drift time t_{drift} in ns. The red line is the linear fit used to convert z_{truth} to an equivalent truth drift time $t_{truth} = -(6558 \pm 5)z_{truth} + (956226 \pm 387)$.

The fractional efficiency error is determined by modifying the slope -6558 by its error ± 5 to produce two additional truth drift time conversion equations. This is slightly different than the analysis method described by equation (6.7). Instead of modifying the cut value we are modifying a conversion equation. One equation, with a smaller slope, results in $N_{i,more}$ events surviving the truth fiducial volume cut. The other has a slightly larger slope, and will exclude more events when the truth fiducial volume cut is applied resulting in $N_{i,less}$ events surviving. The average difference between the number of events in the two modified datasets from N_i gives the fractional error

$$\frac{\sigma_i}{N_i} = \frac{(N_{i,more} - N_i) + (N_i - N_{i,less})}{2N_i} = \frac{147290 - 147055}{2 \cdot 147175} = 0.08\% \quad (6.9)$$

6.4.2 Selection Criteria

The selection criteria requires events to have the proper amount of $S1$ and $S2$ pulses in the appropriate order. This is effectively a zeroth cut (cut0) that excludes events that could not be reconstructed. This chooses the best candidate $S1$ pulse to associate with the β and γ interaction.

The requirement finds the largest $S1$ pulse in the event and identifies it as $S1_\gamma$. The largest $S2$ pulse ($S2_{max}$) in the event is required to occur after the $S1_\gamma$ pulse. If these conditions are met, the largest $S1$ pulse to occur before $S1_\gamma$ is identified as $S1_{beta}$.

With the three minimum pulse requirements met, the position is reconstructed. The x and y location is given by the largest $S2$ pulse, $S2_{max}$ and the drift time is the difference between the $S2_{max}$ and $S1_\gamma$ pulse start times. All other quantities relevant to each the remaining cuts are also calculated at this step. Of the 147,175 events generated within the fiducial truth volume, 137,793 met the minimum selection criteria.

6.4.3 Fiducial Cut

A fiducial cut is applied to the reconstructed event positions. This is referred to as “cut1” and determines how many of the events generated in within the fiducial volume were identified within it. In some cases the event is generated within the fiducial volume, but the gamma travels outside of it and is rejected. Figures 6.9 and 6.10 show the reconstructed positions of events generated within the fiducial volume. Figures 6.11 and 6.12 show events which were both generated within and reconstructed within the fiducial volume. Of the 137,793 events identified by the selection criteria, 131,845 were reconstructed within the detector fiducial volume.

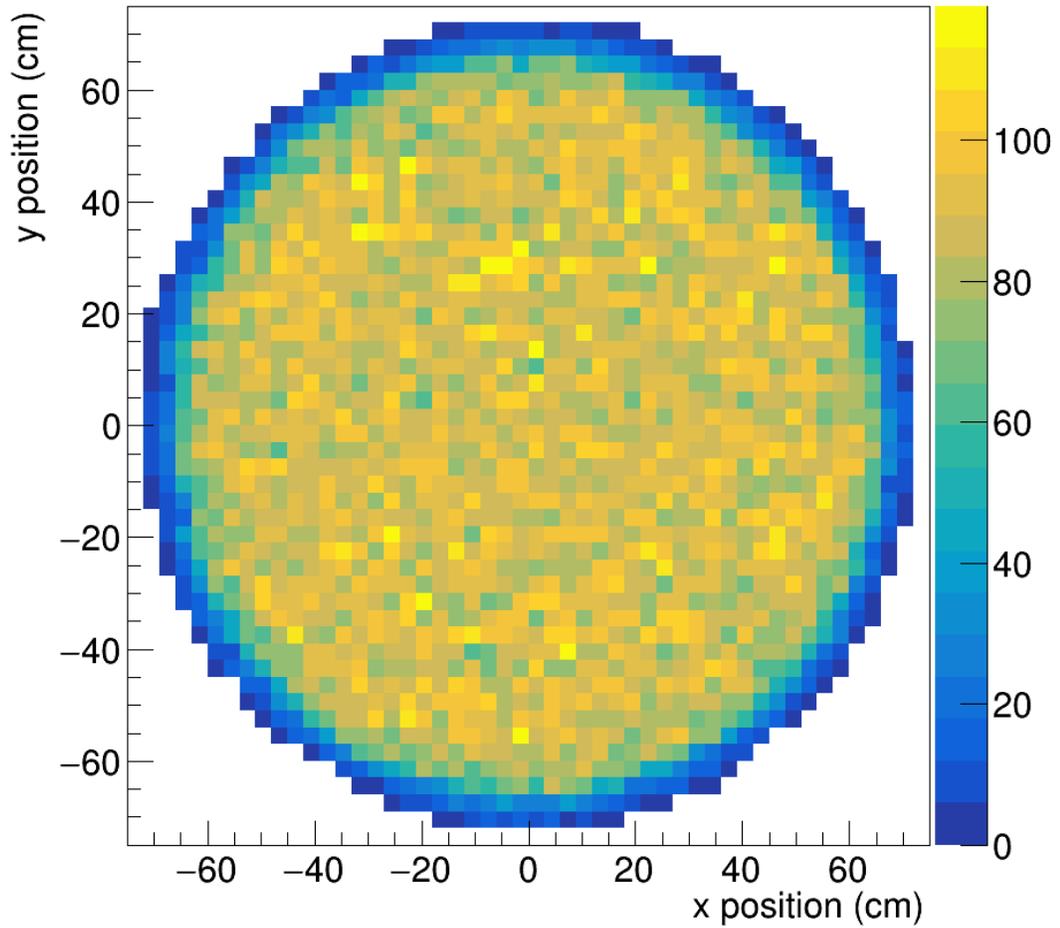


Figure 6.9: The x and y reconstructed positions of events from the simulated ^{85}Kr excited state decay that were generated within the fiducial truth volume and satisfy the minimum selection criteria (cut0).

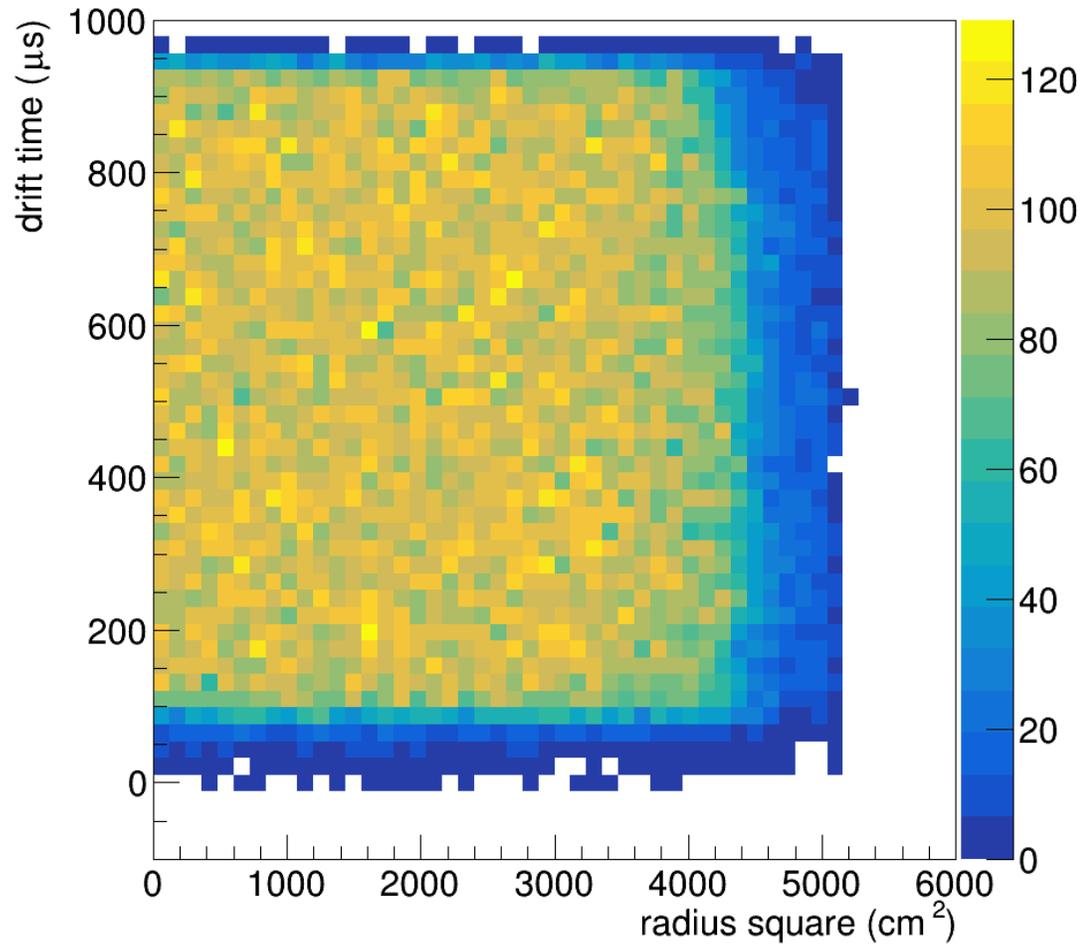


Figure 6.10: The reconstructed square radius and drift time positions of events from the simulated ^{85}Kr excited state decay that were generated within the fiducial truth volume and satisfy the minimum selection criteria (cut0).

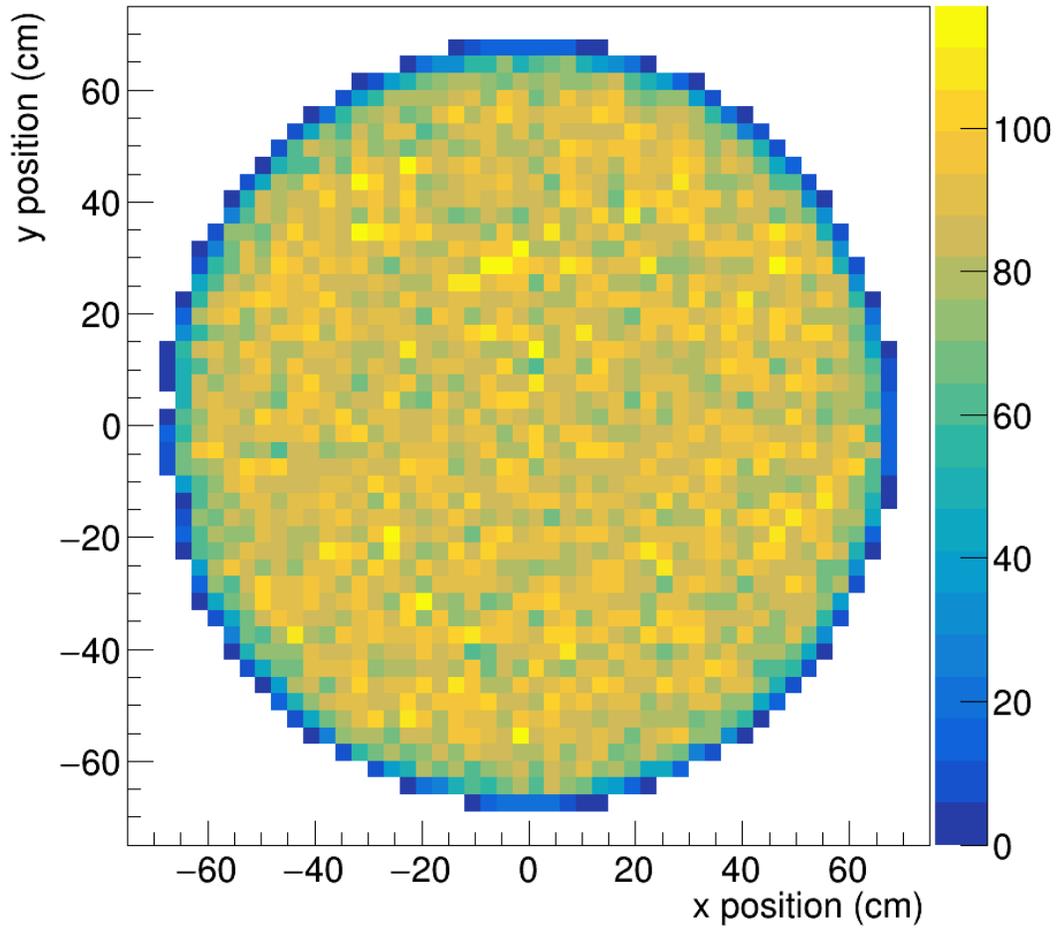


Figure 6.11: The x and y reconstructed positions of events from the simulated ^{85}Kr excited state decay that were generated within the fiducial truth volume, satisfy the minimum selection criteria (cut0), and the reconstructed fiducial volume cut (cut1).

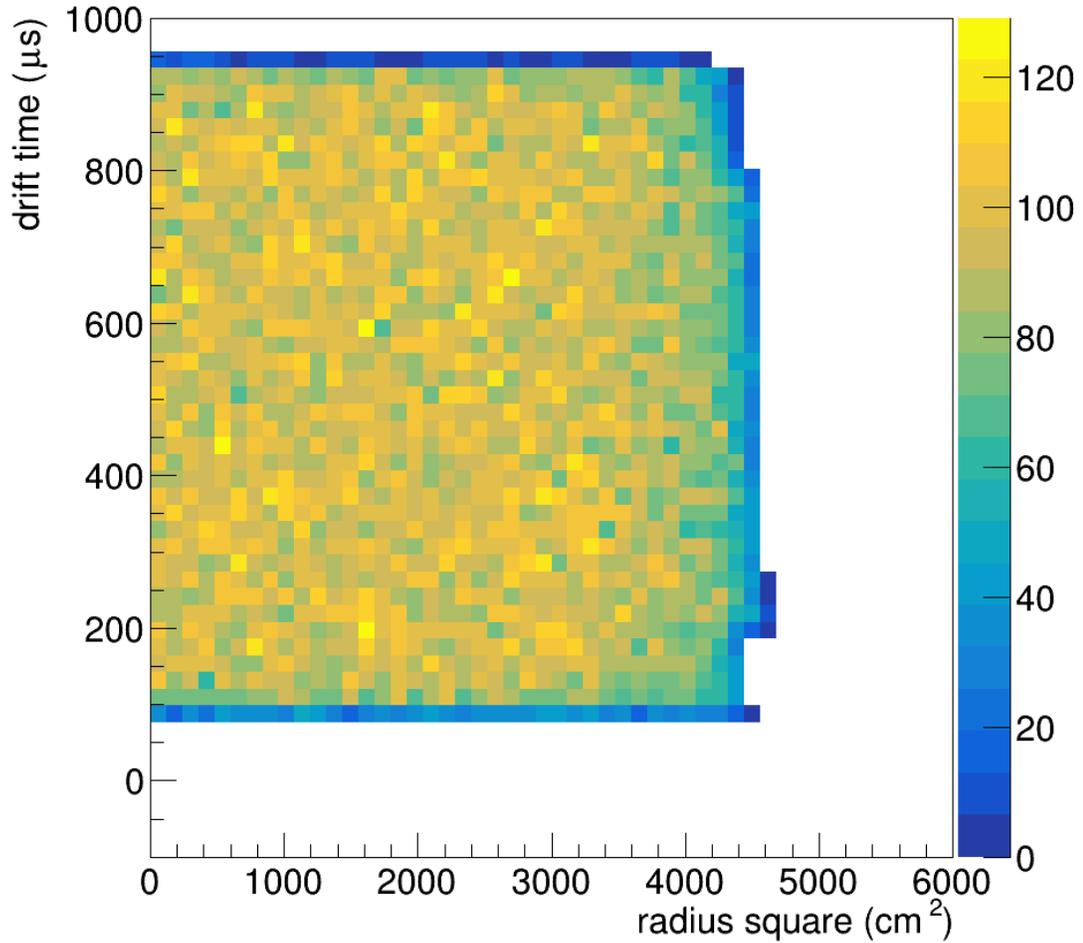


Figure 6.12: The reconstructed square radius and drift time positions of events from the simulated ^{85}Kr excited state decay that were generated within the fiducial truth volume, satisfy the minimum selection criteria (cut0), and the reconstructed fiducial volume cut (cut1).

6.4.4 Energy Cuts

Events with the same energy appear along anti-correlated lines in the $S1_{total}$ versus $S2_{total}$ space. This is due to the relationship between light quanta, charge quanta, and energy, described by equation (2.19). Since the excited state decay search is not recognized

as a single scatter or multi-scatter event by LZap, detector corrections are not applied. Due to this, the excited state decay search of ^{85}Kr is carried out using uncorrected RQs, and spatial corrections that account for charge and light collection efficiencies are not available. This means that the $S1$ and $S2$ signals will be smeared depending on the event location. To characterize the energy of an excited state decay, cuts are applied to the individual $S1_{\beta}$, $S1_{\gamma}$, and $S2_{total}$ RQs, before applying a final selection cut in the $S1_{total}$ versus $S2_{total}$ space. This amounts to four unique cuts that characterize the energies associated with each pulse and their correlation in the $S1_{total}$ versus $S2_{total}$ space.

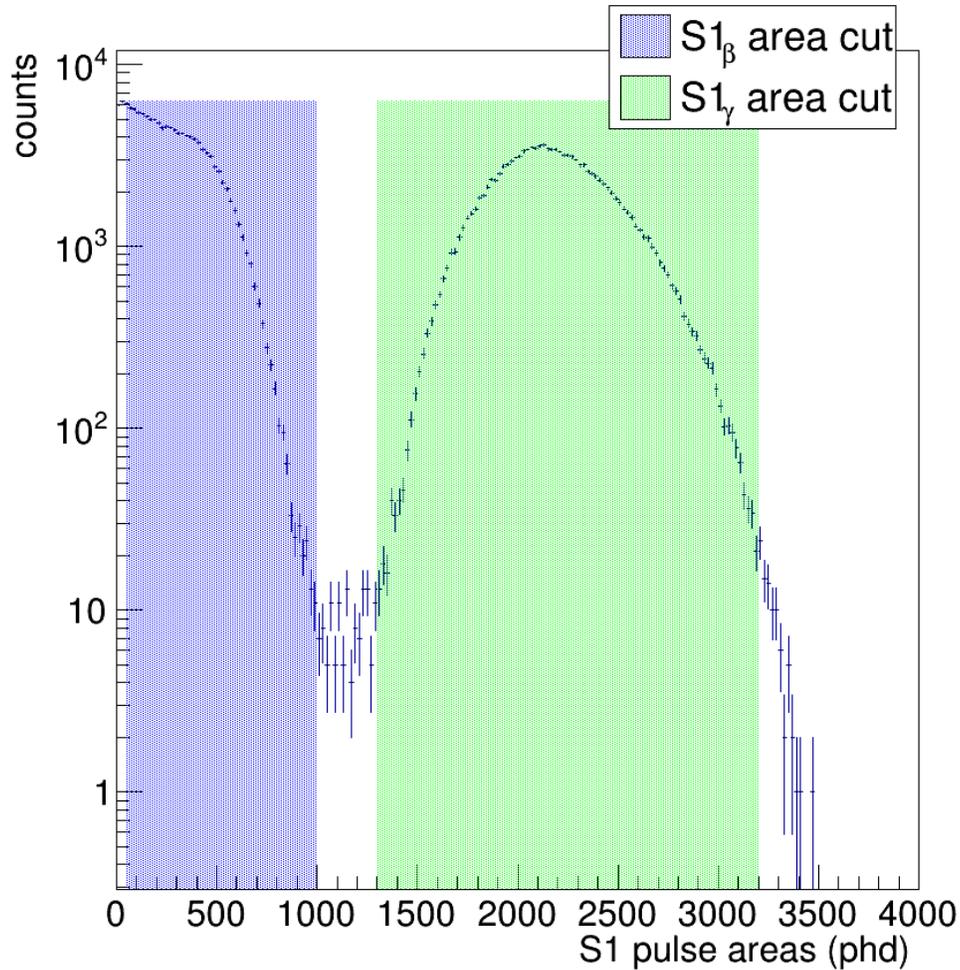


Figure 6.13: A histogram of simulated $S1$ pulse areas, identified as either an $S1_\beta$ or $S1_\gamma$. The blue shaded area accepts any $S1_\beta$ pulse with an area between 50 and 1000 phd. The green shaded area accepts any $S1_\gamma$ pulse with an area between 1300 and 3200 phd. These events have been identified using the selection criteria and are reconstructed within the fiducial volume, no energy cuts have been applied.

The cut in $S1_\beta$, or “cut2”, is from 50 to 1000 phd and shown in 6.13. Of the 131,845 events that survive the cut1, 116,477 survive cut2. The majority of the events excluded in cut2 are due to the lower $S1_\beta$ pulse area cut of 50 phd. This is chosen to reduce contributions from random coincidence backgrounds during periods of high detector activity, not

captured in simulations. These periods of high activity come from a variety of sources that cause afterglow in the detector or $S1$ flare ups. These events contribute $S1$ pulse areas up to 10 phd. If these periods of high activity are excluded with an exposure time cut, the exposure is reduced by 35% [24]. By enforcing a lower $S1_\beta$ threshold at 50 phd, excess downtime is avoided and efficiency is reduced by only 11.7%.

After removing events that fail cut2, the $S1_\gamma$ cut, or “cut3”, shown in figure 6.13, requires all events have $S1_\gamma$ pulse areas between 1300 and 3200 phd. Of the 116477 events that survive cut2, 116,215 events survive cut3.

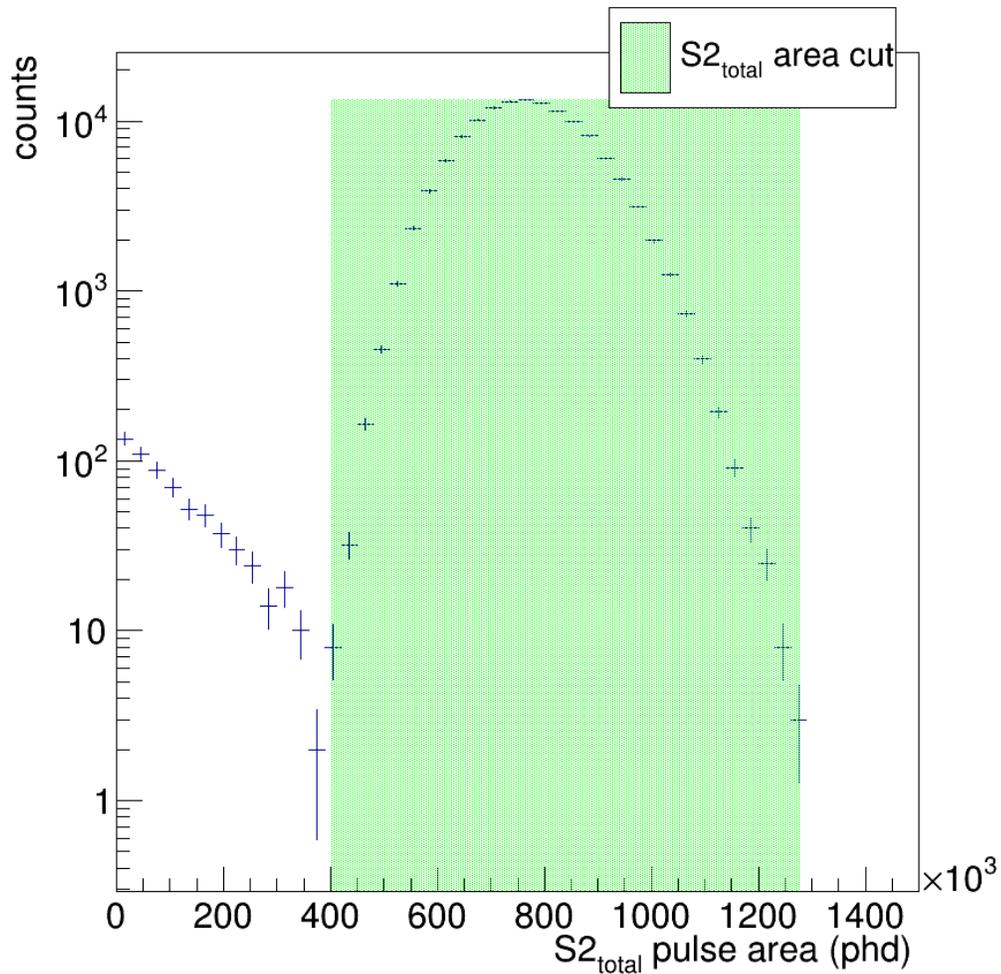


Figure 6.14: A histogram of simulated $S2_{total}$ pulse areas for ^{85}Kr excited state decays. The green shaded area accepts all events with an $S2_{total}$ area between 0.4×10^6 and 1.28×10^6 phd. These events have been identified using the selection criteria and are reconstructed within the fiducial volume, no energy cuts have been applied.

After removing any events that fail cut3, the $S2_{total}$ cut or “cut4” requires each event to have $S2_{total}$ pulse area between 0.4×10^6 and 1.28×10^6 phd, as shown in figure 6.14. Of the 116,215 events that survive cut3, 115,814 events survive cut4.

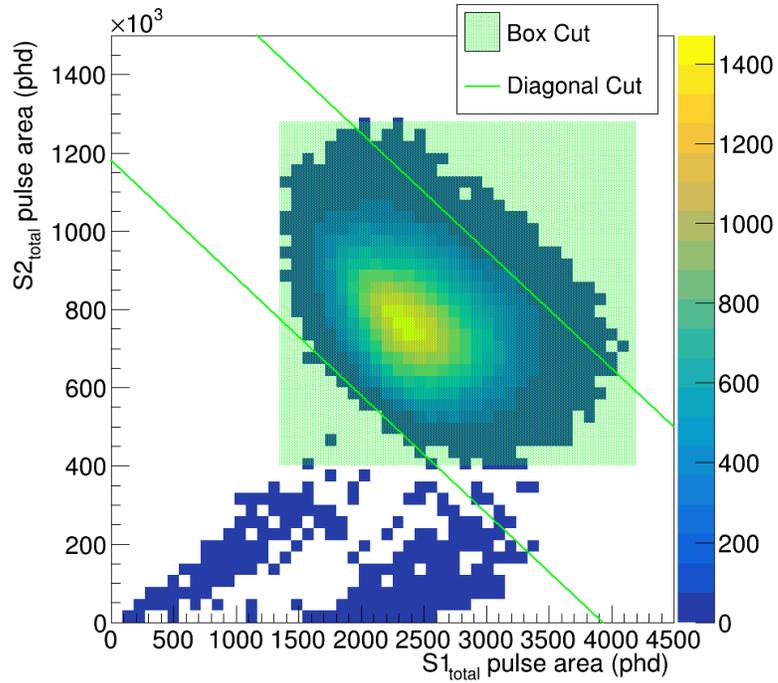


Figure 6.15: The $S1_{total}$ pulse area versus the $S2_{total}$ pulse area of simulated ^{85}Kr excited state decays. The green shaded area is not explicitly a cut, but represents the combined $S1_{\beta}$, $S1_{\gamma}$, and $S2_{total}$ pulse area cuts from 1350 to 4200 phd in $S1_{total}$ area and the same $S2_{total}$ area cut from 0.4×10^6 and 1.28×10^6 phd. The two diagonal lines represent the fourth and final total energy cut. These events have been identified using the selection criteria and are reconstructed within the fiducial volume, no energy cuts have been applied.

After removing any events that fail cut4, the diagonal total energy cut or “cut5” requires each event fall between the diagonal lines shown in figure 6.15. The lower bound is defined by

$$S2_{total} \geq -300 \cdot S1_{total} + 1.18 \times 10^6 \quad (6.10)$$

and the upper bound is defined by

$$S2_{total} < -300 \cdot S1_{total} + 1.85 \times 10^6. \quad (6.11)$$

Of the 115,814 events that survive cut4, 114,964 events survive cut5.

6.4.4.1 Validating Energy Cuts

To validate the simulated ^{85}Kr uncorrected detector response, seven calibration sources were identified in both simulations and SR1 data. These sources are listed in table 6.1 and have mono-energetic peaks from 39.5 keV to 1.46 MeV. These peaks are readily identifiable in the single-scatter corrected $S1_c$ versus $S2_c$ space of simulations and the $S1_c$ versus $S2_{c,bot}$ (bottom PMT array only) in detector data, as shown in figures 6.16 and 6.17.

The $S2_{c,bot}$ RQ is the corrected $S2$ pulse area recorded by only the bottom PMT array. This is used for identification purposes in the SR1 data only, as the $S2$ signal at high energies saturates the top PMT array. The uncorrected $S2$ pulse areas, which do saturate, are included in this comparison, because we expect the 514 keV gamma to cause saturation of top PMTs. In fast simulations the $S2_{bot}$ RQ is unavailable and saturation is not accounted for.

Isotope	Energy (keV)	Half-life	Process	Location
$^{125}\text{I}^*$	39.5	59.5 days	Electron capture + x-ray	Dispersed
^{125}I	67.3	59.5 days	Electron capture + x-ray	Dispersed
^{131m}Xe	164	11.9 days	Gamma	Dispersed
^{125}Xe	276	16.9 hours	Electron capture + x-ray	Dispersed
^{60}Co	1173	5.27 years	Gamma	Surface
^{60}Co	1332	5.27 years	Gamma	Surface
^{40}K	1460	1.25e9 years	Electron capture	Surface

Table 6.1: Sources that produce electron recoil events at a wide range of energies. Auger electrons, cascade x-rays, and intermediate short-lived excited states all contribute to these mono-energetic peaks. The 39.5 keV ^{125}I source is noted to include contributions from ^{129}Xe in the SR1 data due to neutron activation during DD calibrations. The most prominent energy peaks are selected for analysis. Total energies and half-life taken from reference [29]

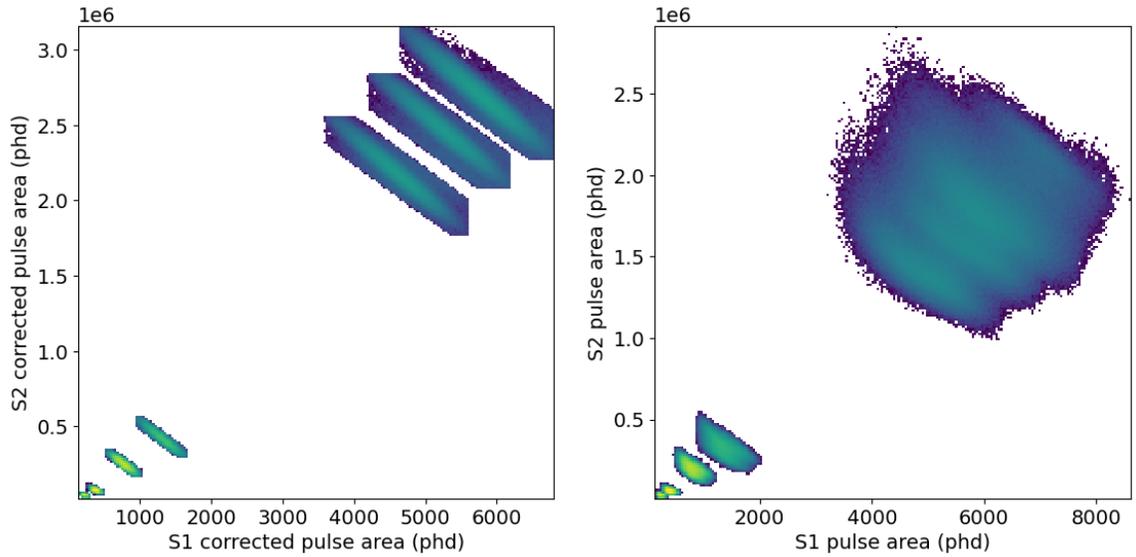


Figure 6.16: A 2-D histogram of the peaks from table 6.1 identified in simulations. The left plot shows the corrected $S1_c$ and $S2_c$ pulse areas where events are identified. All identified events in simulations must pass a fiducial cut, energy cut, and are checked against particle IDs. The right plot shows the uncorrected $S1$ and $S2$ pulse areas. From bottom left to top right the events are increasing in energy from the the lowest energy 39.5 keV ^{125}I peak up to the 1460 keV ^{40}K peak.

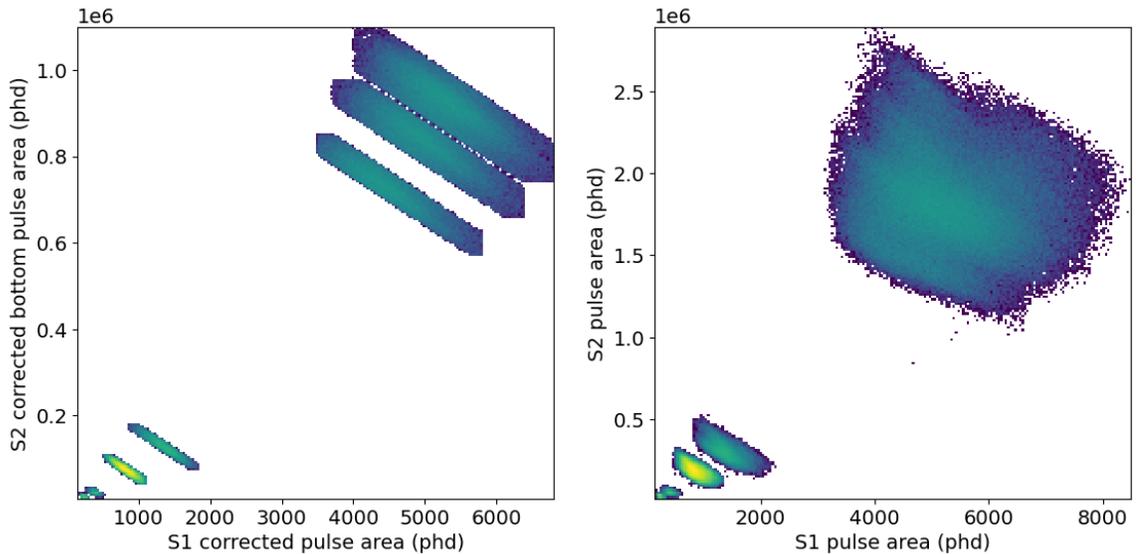


Figure 6.17: A 2-D histogram of the peaks from table 6.1 identified in the SR1 and DD data. The left plot shows the corrected $S1_c$ and $S2_{c,bot}$ pulse areas where events are identified. All identified events in the detector data must pass a fiducialization cut, and energy cut. The right plot shows the uncorrected $S1$ and $S2$ pulse areas.

With the peaks of interest identified in both simulations and data, the uncorrected pulse area distributions can be compared. Each distribution is fit to a Gaussian to determine the peak value of $S1_{SIM}$, $S2_{SIM}$, $S1_{SR1}$, and $S2_{SR1}$. Plots of $S1_{SIM}$ versus $S1_{SR1}$ and $S2_{SIM}$ versus $S2_{SR1}$ peak values are used to make a linear fit between simulations and SR1. From these fits the simulated $S1_{\gamma,SIM}$ and $S2_{total,SIM}$ peak pulse areas are mapped to an equivalent SR1 value. The percent difference between the simulated and SR1 peak values are used to modify $S1_{\beta,SIM}$, $S1_{\gamma,SIM}$, and $S2_{total,SIM}$ pulse area cuts and inform the efficiency uncertainties for each cut.

The 39.5 keV peak from ^{125}I is caused by the L-shell electron capture as it decays to ^{125}Te . ^{125}I has a 59.5 day half-life. The electron capture generates a 35.5 keV gamma and the electron cascade produces a 4 keV x-ray. ^{125}I is produced by ^{125}Xe electron captures.

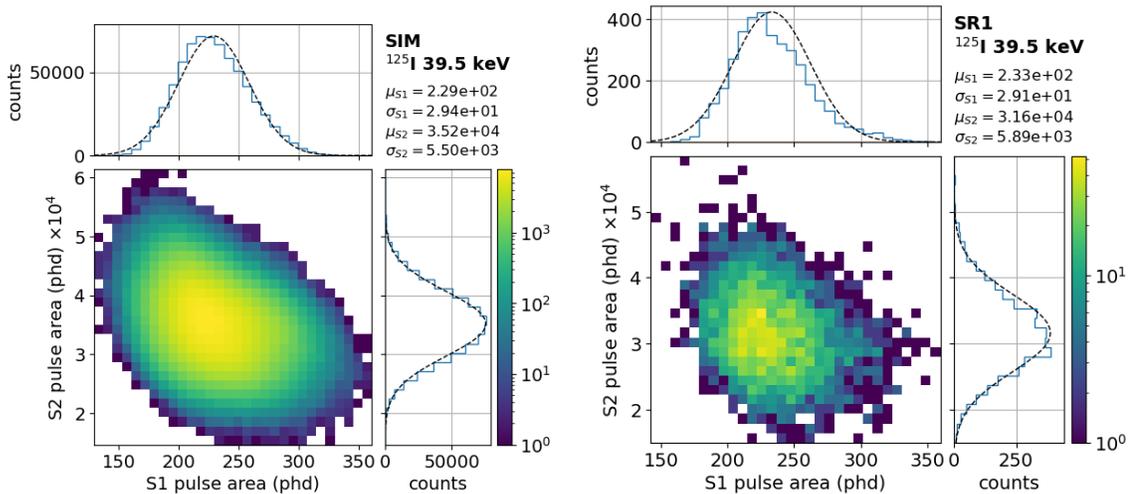


Figure 6.18: The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 39.5 keV ^{125}I events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot.

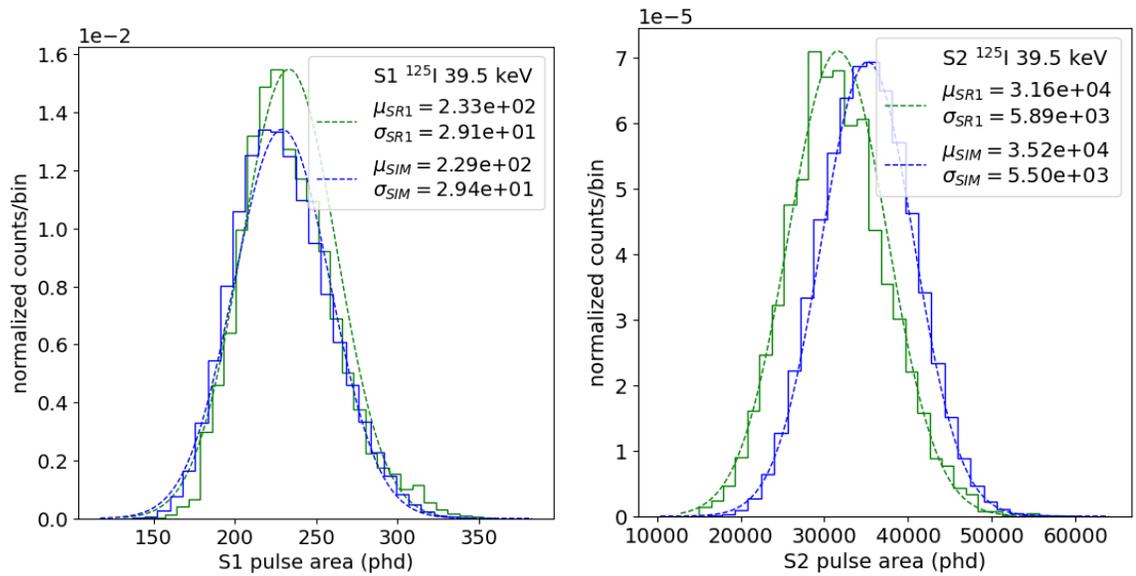
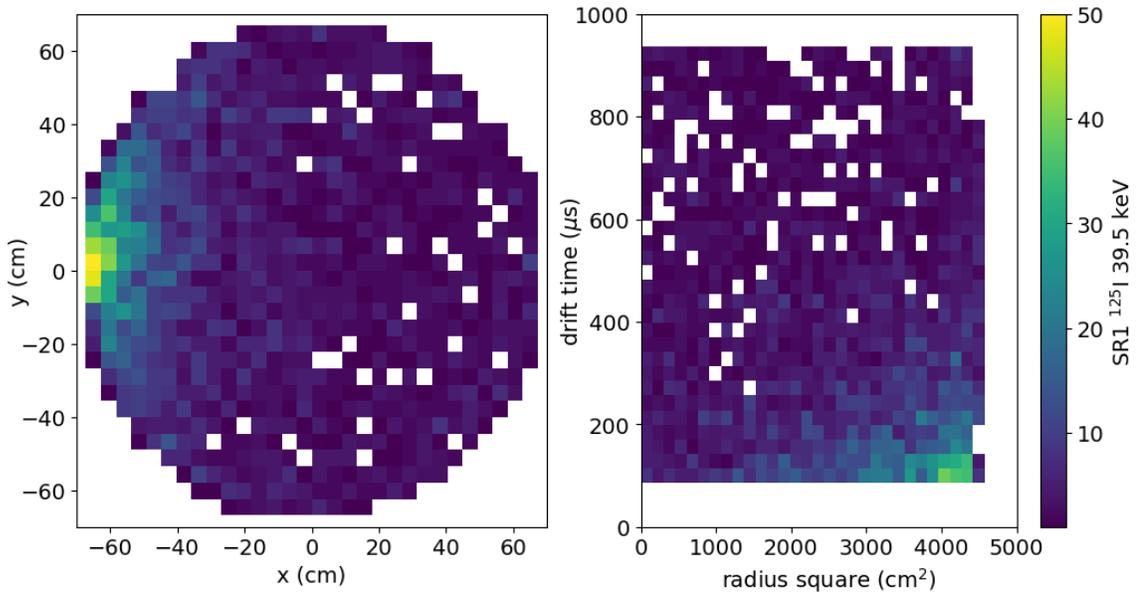
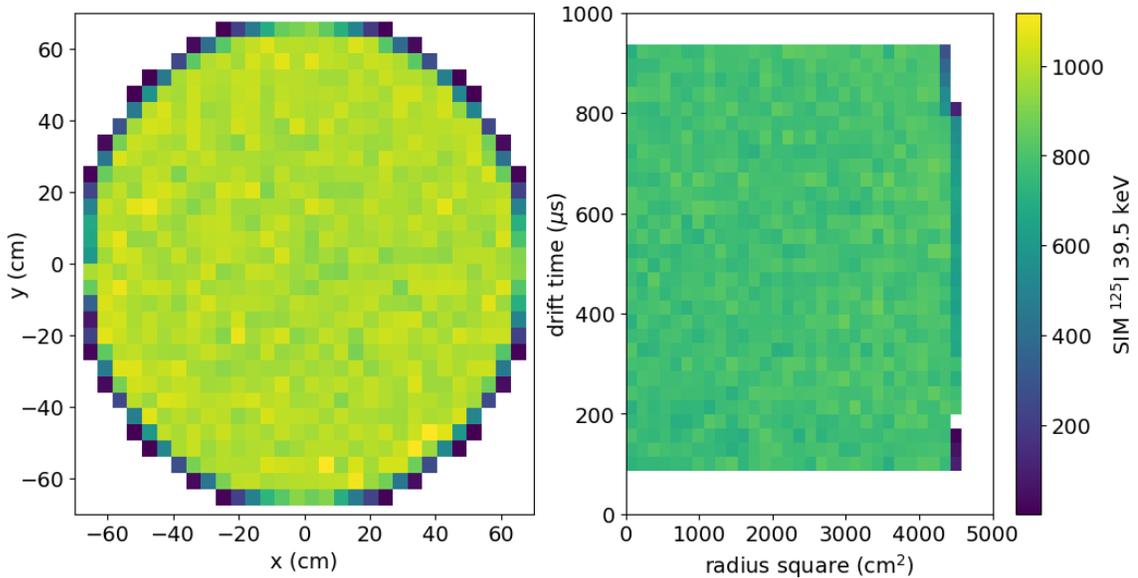


Figure 6.19: The normalized Gaussian fits of the 39.5 keV ^{125}I events, from figure 6.18. Fits show the S1 (left) and S2 (right) pulse area peaks from the SR1 data (green) and simulations (blue).

The simulated events are confirmed via the MCtruth pulse ID RQ to be the expected ^{125}I L-shell electron capture, but the spatial distribution from the SR1 data, shown in figure 6.20a, is inconsistent with the higher energy 67.3 keV K-shell electron capture events, shown in figure 6.21a. Due to the 59.5 day half-life, ^{125}I is expected to mix uniformly within the detector, as shown in both of the simulation figures 6.20b, 6.21b, and the higher energy ^{125}I SR1 data in figure 6.21a.

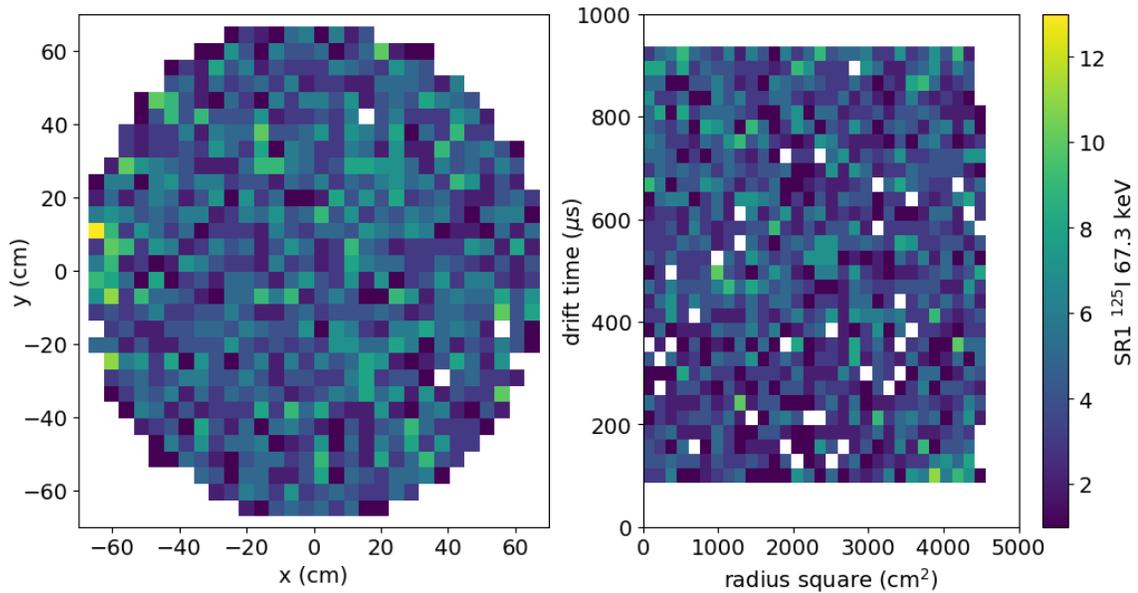


(a) The SR1 spatial distribution of events identified as ^{125}I 39.5 keV with energy cuts.

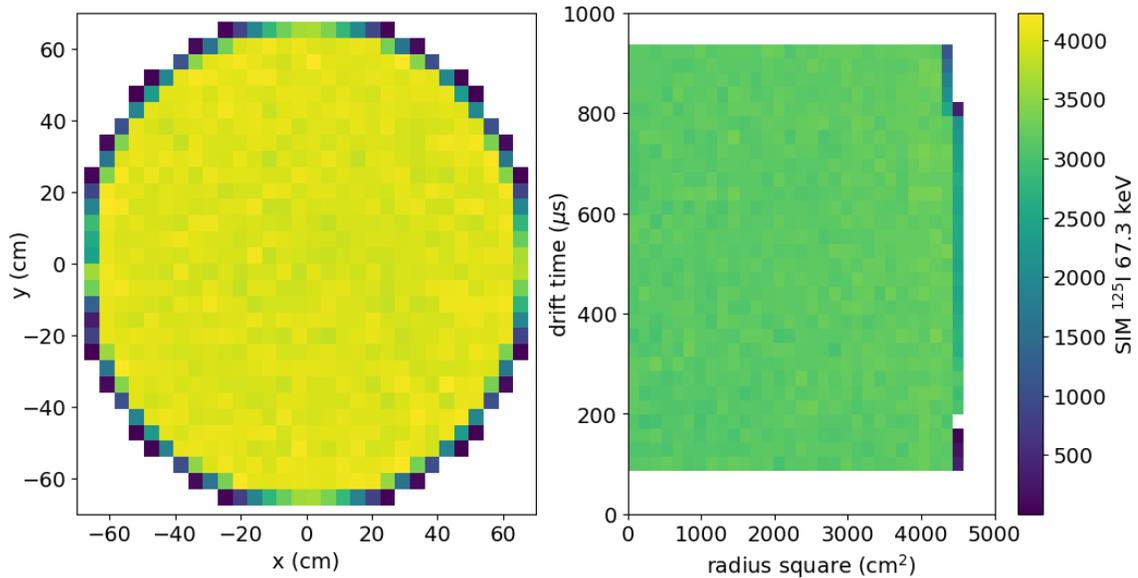


(b) The simulated spatial distribution of events identified as ^{125}I 39.5 keV with energy cuts and the MCtruth particle ID RQ.

Figure 6.20: The spatial distributions of the 39.5 keV ^{125}I events in simulations and SR1 data. The SR1 events, in figure 6.20a, are clustered in the top corner of the detector, near the DD conduit. The simulated events, in figure 6.20b, are uniformly distributed throughout the detector.



(a) The SR1 spatial distribution of events identified as ^{125}I 67.3 keV with energy cuts.



(b) The SIM spatial distribution of events identified as ^{125}I 67.3 keV with energy cuts and the MCtruth particle ID RQ.

Figure 6.21: The spatial distributions of the 67.3 keV ^{125}I events in simulations and SR1 data. The SR1 events, in figure 6.21a, and the simulated events, in figure 6.21b, are uniformly distributed throughout the detector, as expected.

The inconsistent spatial distribution of the 39.5 keV ^{125}I events are due to contributions

from the neutron activated, 3/2+ angular momentum state of ^{129}Xe . The excited state has a half-life of 97 ns and is only present during DD calibrations which explains the tight clustering of events near the DD's neutron conduit. Upon relaxing to the 1/2+ ground state, a 39.6 keV gamma is emitted which appears nearly identical to the 39.5 keV ^{125}I peak. For this analysis we choose to keep these events. Ideally a proper ^{129}Xe simulation would be used along with the ^{125}I simulation, but no fast simulation was available.

The 67.3 keV ^{125}I peak is caused by the K-shell electron capture as it decays to ^{125}Te . The electron capture generates a 35.5 keV gamma and the electron cascade produces a 31 keV x-ray.

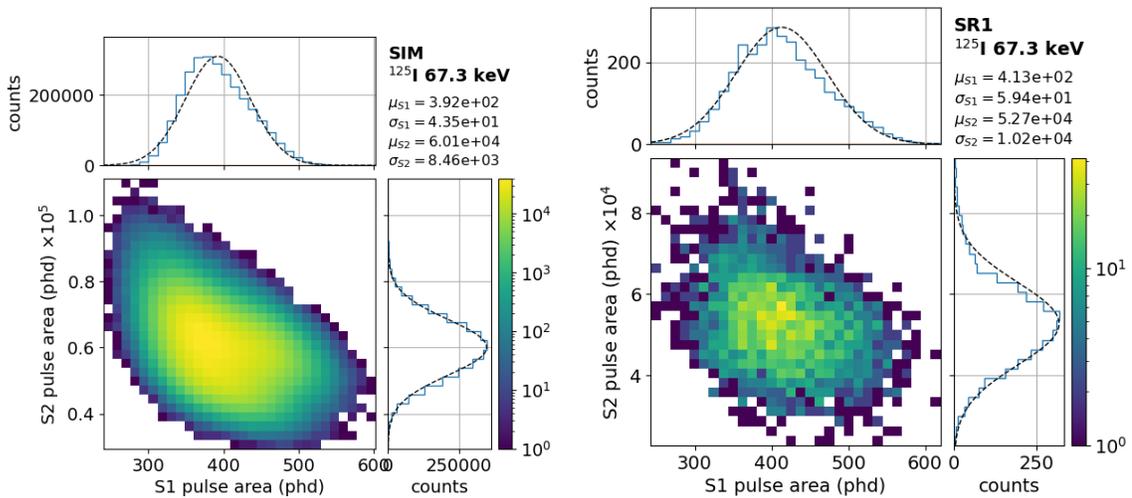


Figure 6.22: The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 67.3 keV ^{125}I events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot.

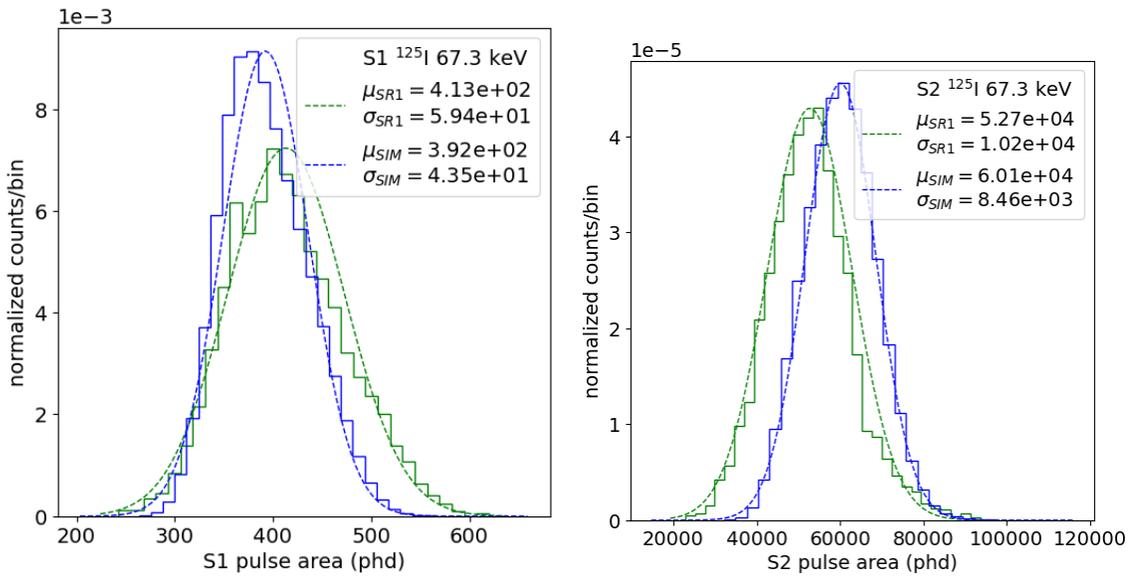


Figure 6.23: The Gaussian fits of the 67.3 keV ¹²⁵I events, from figure 6.22. Fits show the S1 (left) and S2 (right) pulse area peaks from the SR1 data (green) and simulations (blue).

The 164 keV ^{131m}Xe peak is caused by the relaxation of meta-stable ^{131m}Xe following activation caused by DD calibrations. The meta-stable state has an 11.9 day half-life and releases a 164 keV gamma upon relaxing to the stable ground state.

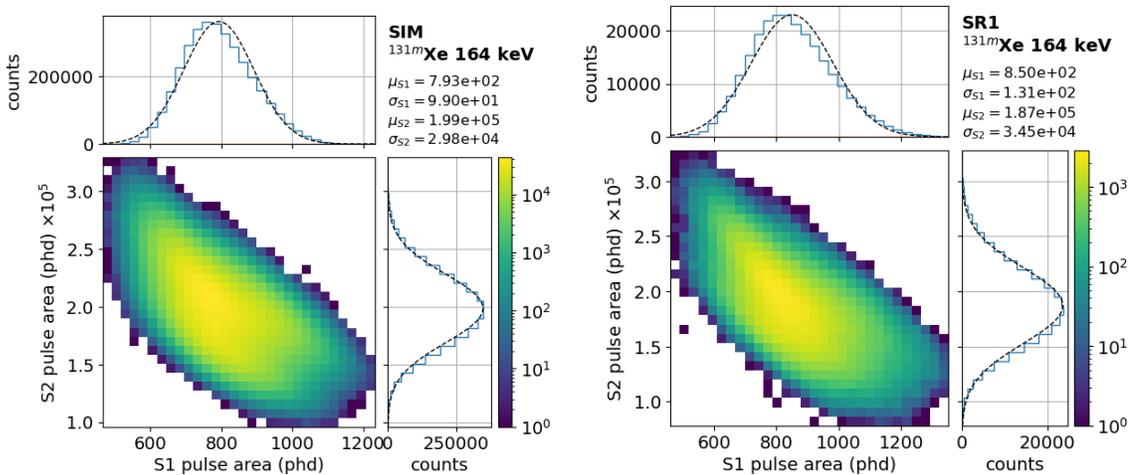


Figure 6.24: The S1 and S2 pulse areas from simulations (left) and SR1 (right) of the 164 keV ^{131m}Xe events. The S1 and S2 values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot.

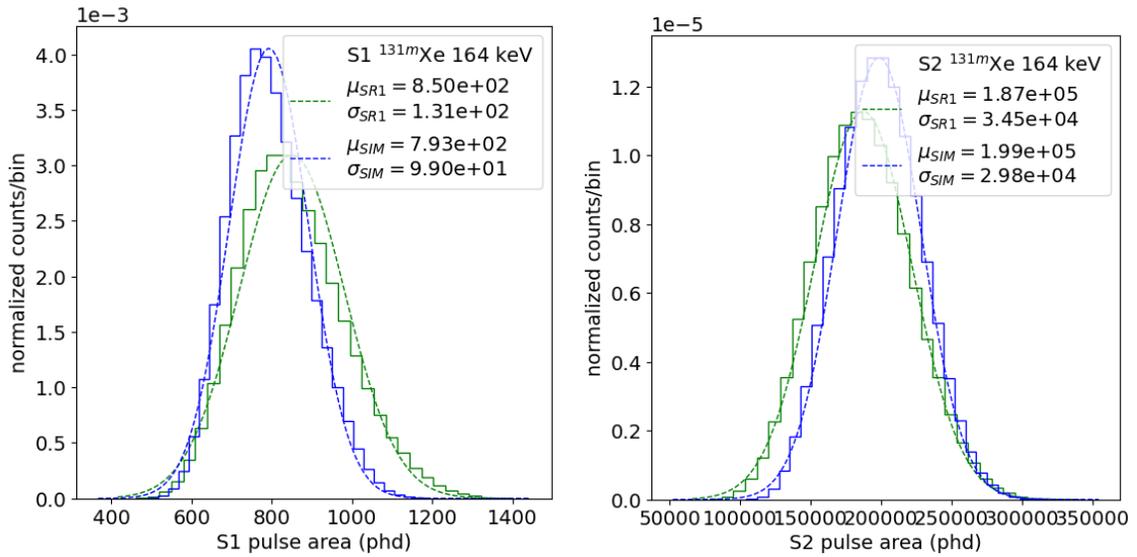


Figure 6.25: The Gaussian fits of the 164 keV ^{131m}Xe events, from figure 6.24. Fits show the S1 (left) and S2 (right) pulse area peaks from the SR1 data (green) and simulations (blue).

The 276 keV ¹²⁵Xe peak is caused by the K-shell electron capture of ¹²⁵Xe as it decays to ¹²⁵I. The electron capture generates a 243 keV gamma and the electron cascade produces a 33 keV x-ray. ¹²⁵Xe is produced via neutron capture onto ¹²⁴Xe during DD calibrations and has a 16.9 hour half-life.

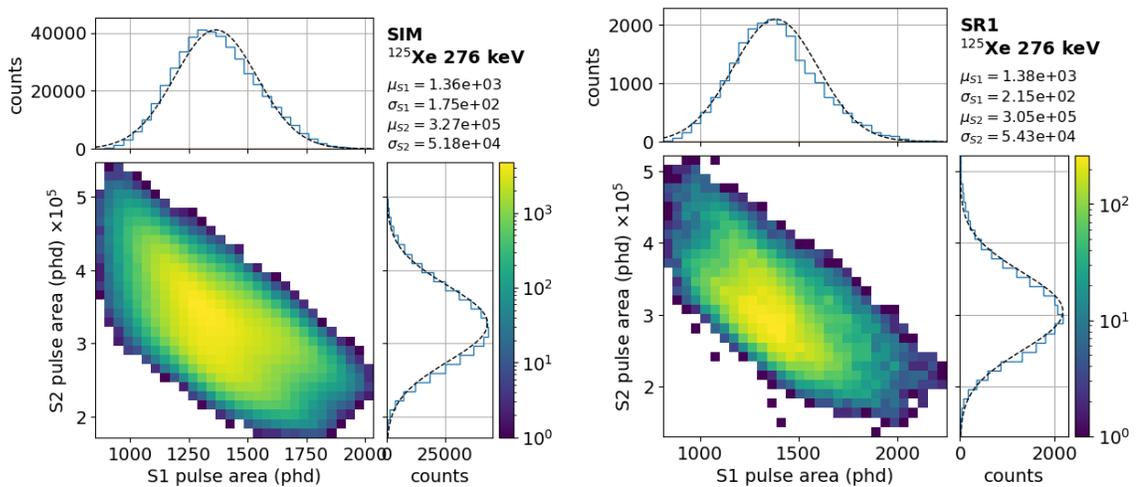


Figure 6.26: The S1 and S2 pulse areas from simulations (left) and SR1 (right) of the 276 keV ¹²⁵Xe events. The S1 and S2 values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot.

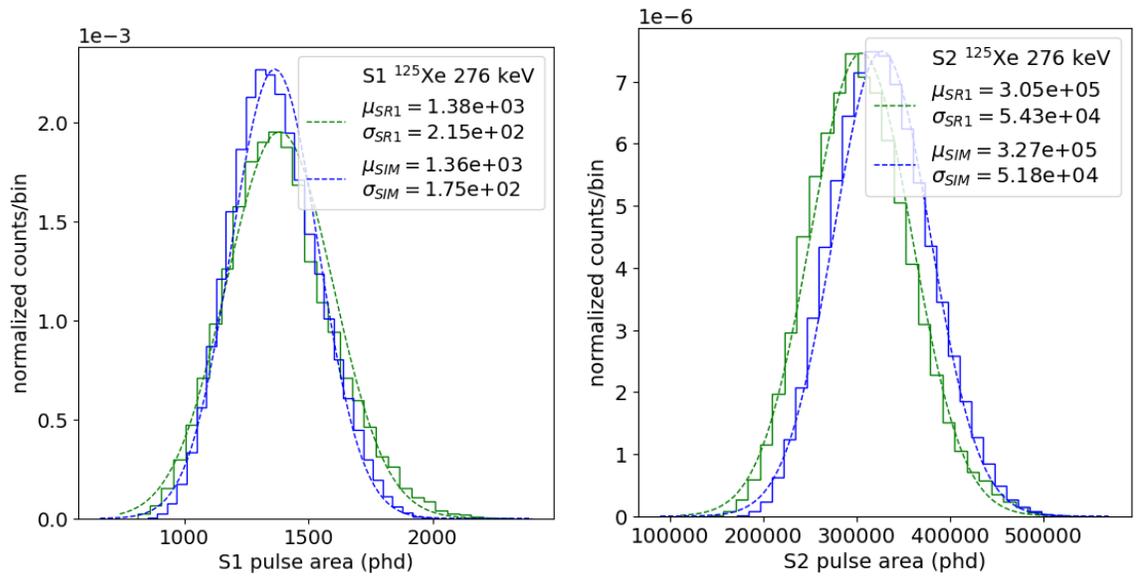
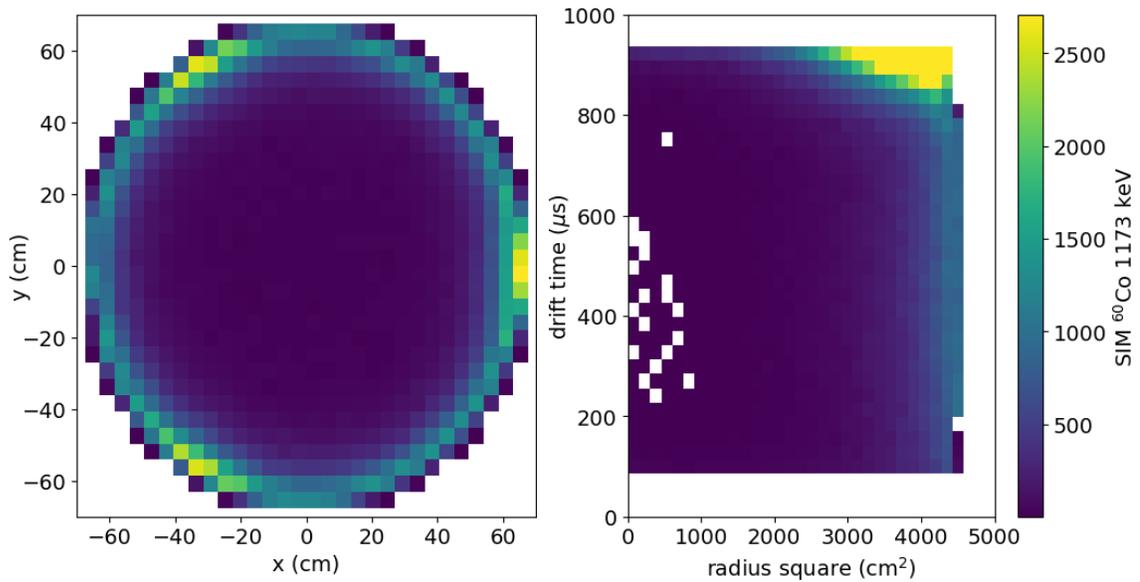


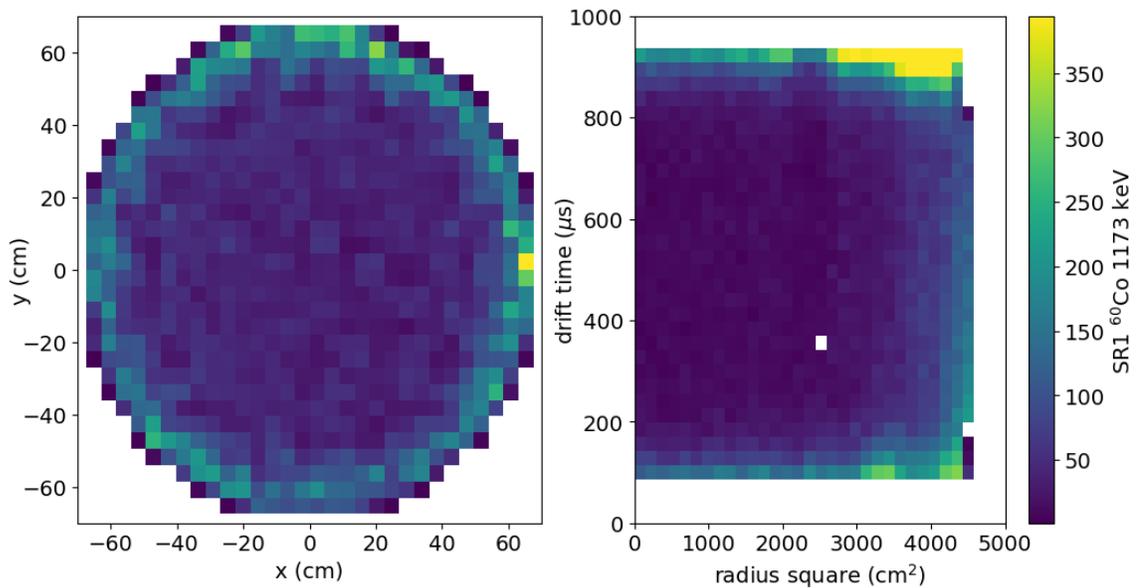
Figure 6.27: The Gaussian fits of the 276 keV ¹²⁵Xe events, from figure 6.26. Fits show the S1 (left) and S2 (right) pulse area peaks from the SR1 data (green) and simulations (blue).

The final three energy peaks considered are concentrated near detector surfaces like the walls and grids. Detector reconstruction suffers near surfaces so we must be careful to avoid unfairly weighting poorly reconstructed events in our map between simulations and SR1 data. Figure 6.28 shows the spatial distributions of events following the standard fiducial cut. Figure 6.29 shows the poorly reconstructed S1 and S2 signals due to the surplus of surface events.

The 1173 keV ⁶⁰Co peak is caused by the beta decay of ⁶⁰Co to the 4+ angular momentum state of ⁶⁰Ni, which then relaxes to the 2+ angular momentum state, releasing a 1173 keV gamma. The ⁶⁰Co is present in detector materials due to cosmogenic activation and has a half-life of 5.3 years.



(a) The simulated spatial distribution of events for ⁶⁰Co 1173 keV.



(b) The SR1 spatial distribution of events for ⁶⁰Co 1173 keV.

Figure 6.28: Spatial distributions of 1173 keV ⁶⁰Co events in simulations and SR1 data with the standard fiducial cut applied.

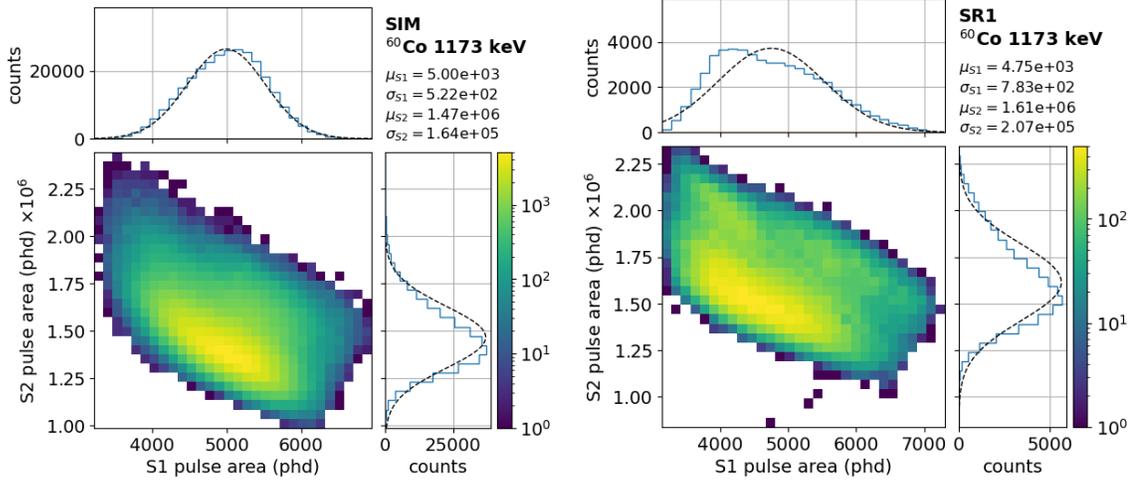
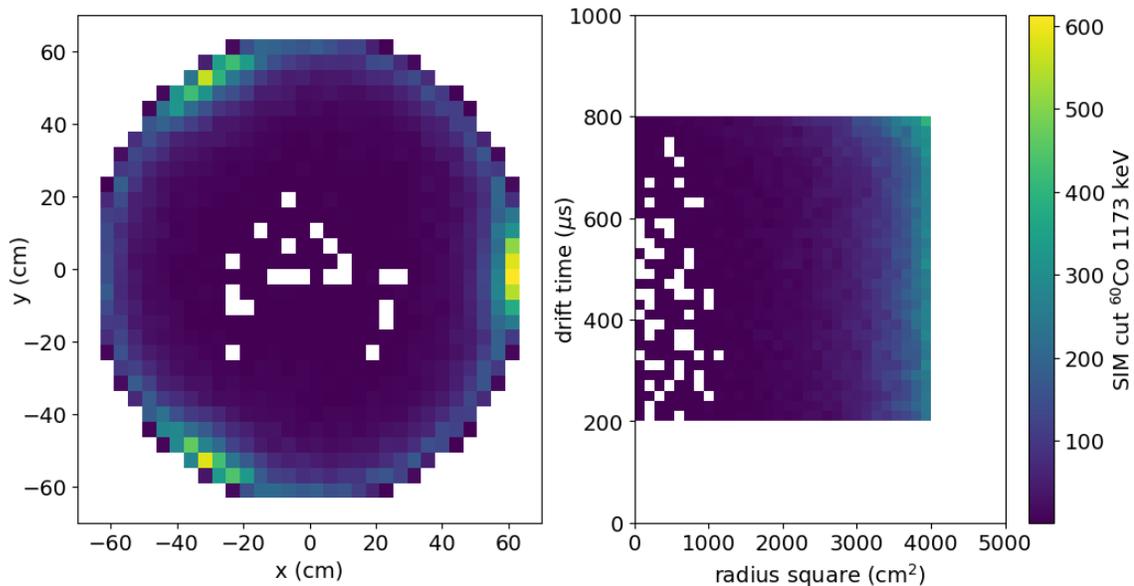
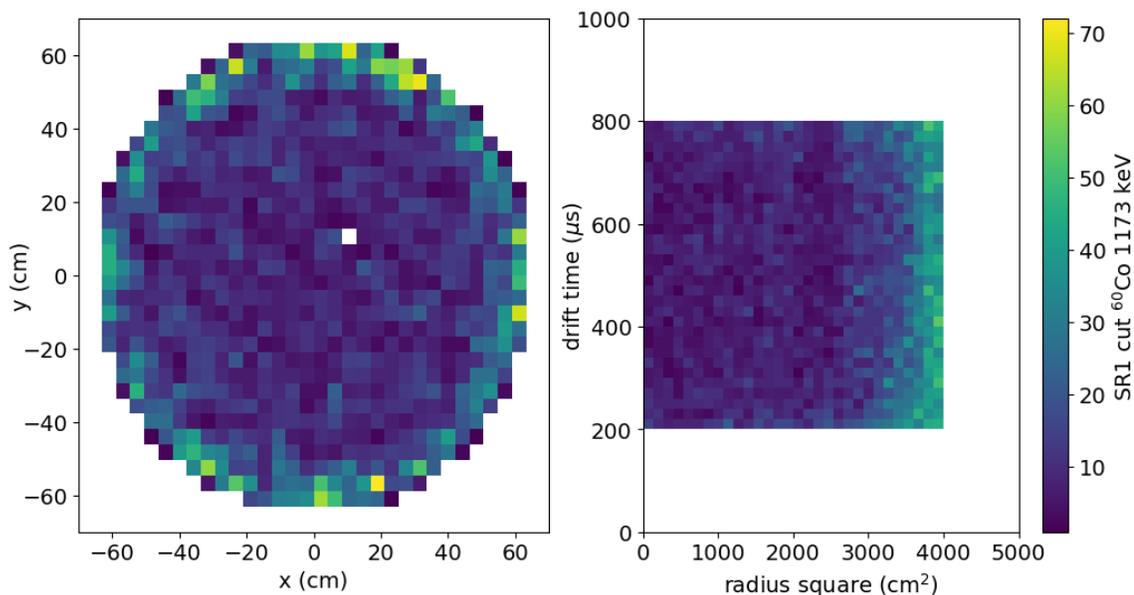


Figure 6.29: S1 and S2 pulse areas of simulated (left) and observed SR1 (right) 1173 keV ^{60}Co events in LZ. The S1 and S2 values have been projected onto their respective axes and a Gaussian fit is performed to identify the central peak locations. The fit data is shown on the top right of the plot. Only the standard fiducial cut is applied. The data is dominated by poorly reconstructed events along the detector edge and grids.

If we instead apply a more stringent volume cut, the contributions from the most poorly reconstructed populations are avoided. The stringent cut only considers events within square radius of 4000 cm^2 and a drift time between 200 and $800 \mu\text{s}$. The 1173 keV peak values in simulation changes from $\mu_{S1} = 5 \times 10^3 \text{ phd}$ and $\mu_{S2} = 1.47 \times 10^6$ to $\mu_{S1} = 4.5 \times 10^3 \text{ phd}$ and $\mu_{S2} = 1.61 \times 10^6$. After the stringent cut is applied, the SR1 peak values change from $\mu_{S1} = 4.75 \times 10^3 \text{ phd}$ and $\mu_{S2} = 1.61 \times 10^6$ to $\mu_{S1} = 4.23 \times 10^3 \text{ phd}$ and $\mu_{S2} = 1.71 \times 10^6$.



(a) The simulated spatial distribution of events for ^{60}Co 1173 keV with the stringent volume cut.



(b) The SR1 spatial distribution of events for ^{60}Co 1173 keV with the stringent volume cut.

Figure 6.30: Spatial distributions of 1173 keV ^{60}Co events in simulations and SR1 data. Here the stringent volume cut has been made.

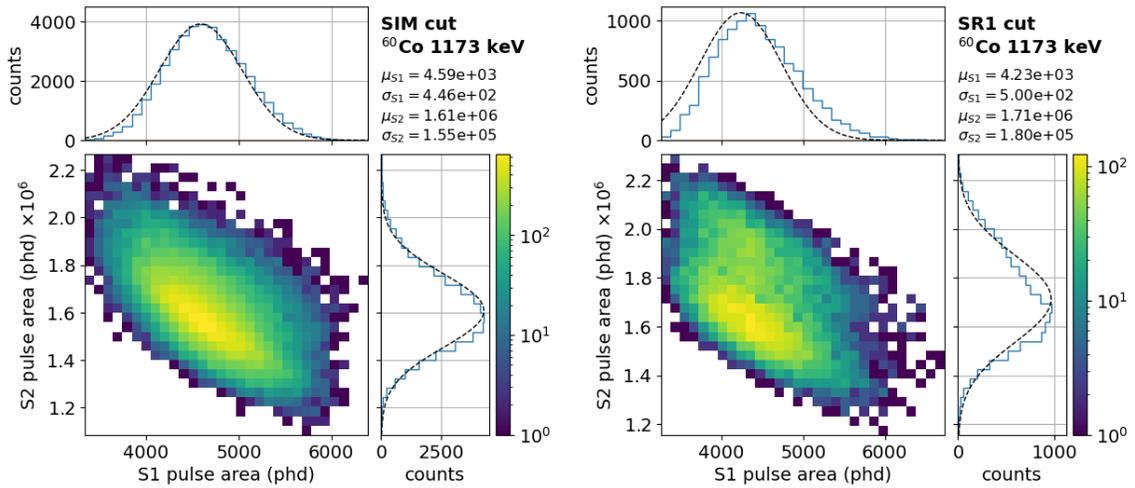


Figure 6.31: The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1173 keV ^{60}Co events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot. The stringent volume cut has been made.

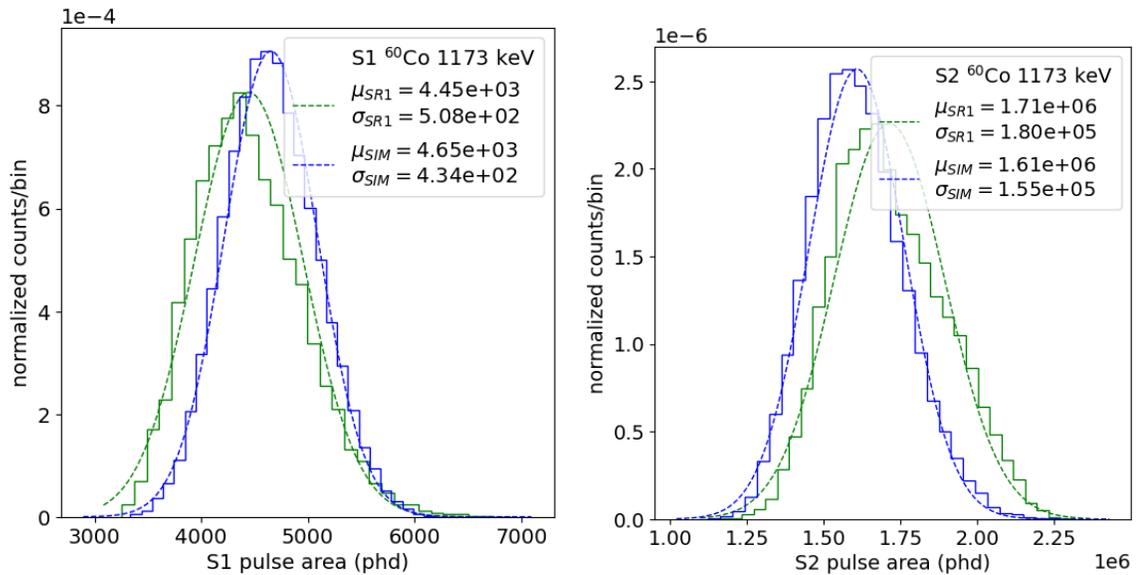


Figure 6.32: The Gaussian fits of the 1173 keV ^{60}Co events, from figure 6.31. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue). The stringent volume cut has been made.

The 1332 keV ^{60}Co peak follows the 1173 keV gamma. The 1332 keV gamma is released as ^{60}Ni relaxes from the 2+ angular momentum state to the 0+ ground state. These

gammas occur at the same time within the detector, but often one gamma exits the main detector volume. This leaves three distinct energy peaks at 1173 keV, 1332 keV, and a combined peak at 2505 keV. For this analysis we choose to only consider the two lower energy, single gamma peaks. The spatial distribution of the 1332 keV events is consistent with its lower energy counterpart and so a similar stringent volume cut is made to reduce the impact of poorly reconstructed events. The 1332 keV peak values in simulation changes from $\mu_{S1} = 5.73 \times 10^3$ phd and $\mu_{S2} = 1.65 \times 10^6$ to $\mu_{S1} = 5.25 \times 10^3$ phd and $\mu_{S2} = 1.82 \times 10^6$. After the stringent cut is applied, the SR1 peak values change from $\mu_{S1} = 5.22 \times 10^3$ phd and $\mu_{S2} = 1.81 \times 10^6$ to $\mu_{S1} = 4.72 \times 10^3$ phd and $\mu_{S2} = 1.91 \times 10^6$.

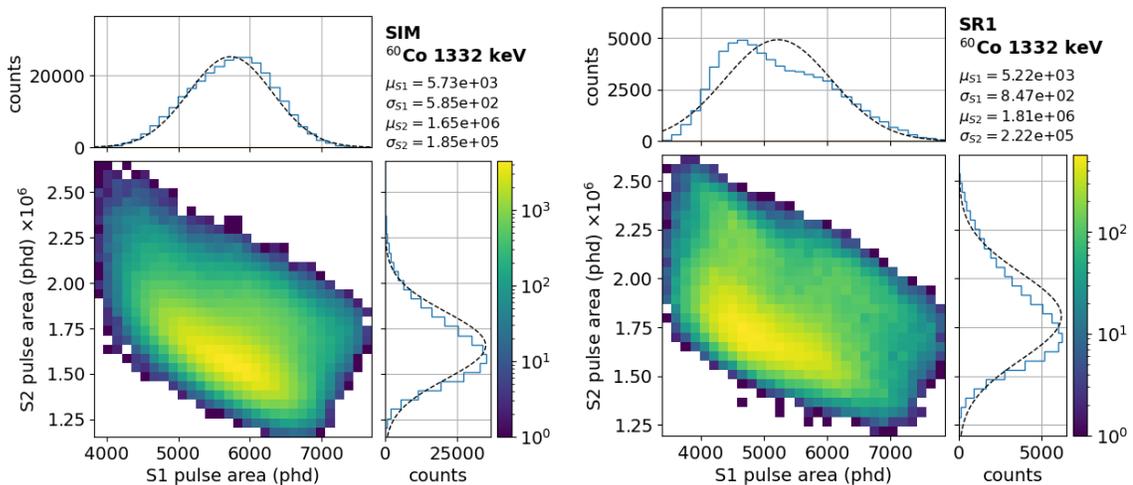


Figure 6.33: The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1332 keV ^{60}Co events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot. Only the standard fiducial cut has been made.

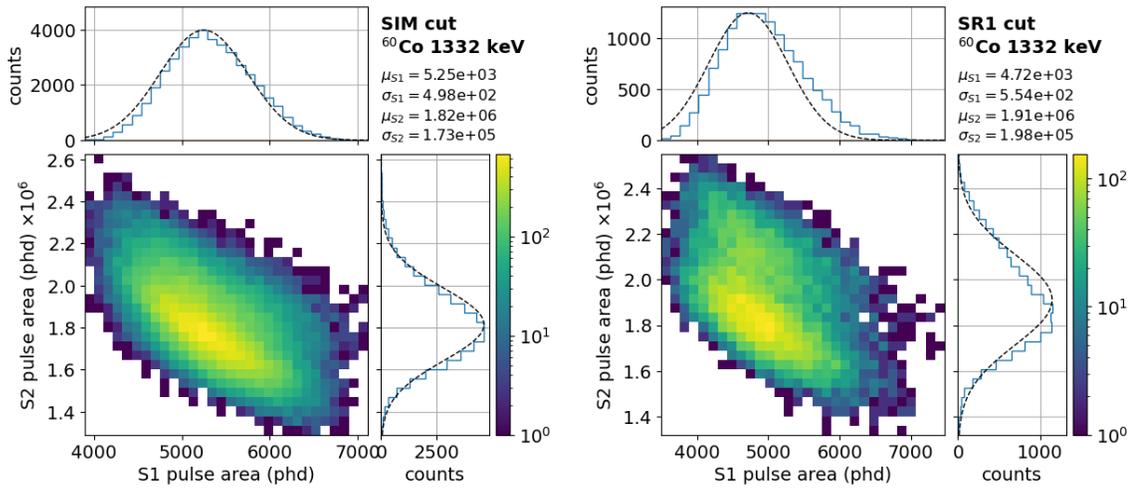


Figure 6.34: The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1332 keV ^{60}Co events. The stringent volume cut has been made.

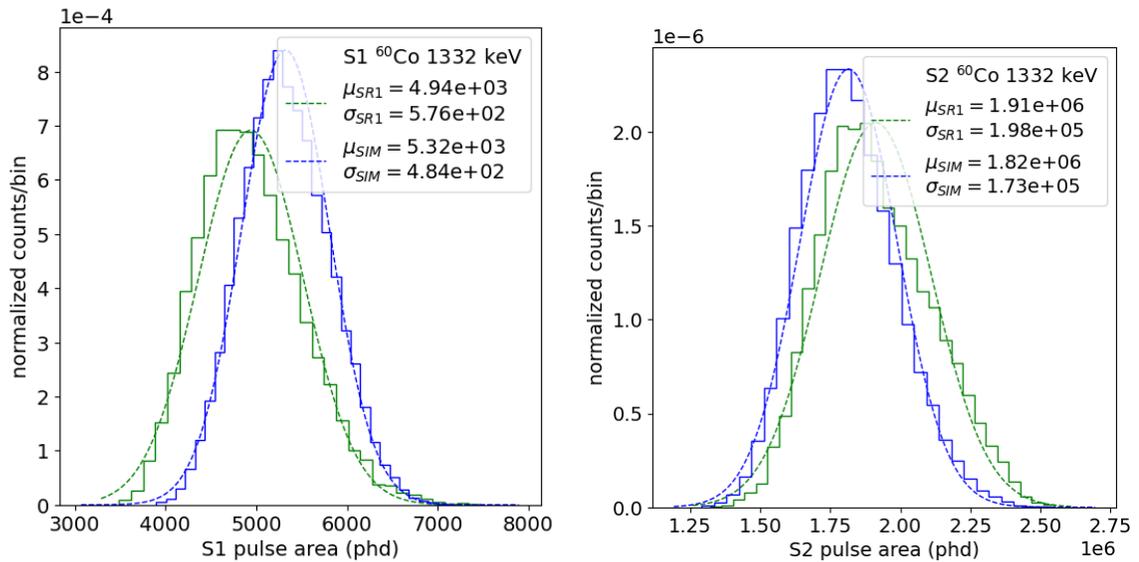


Figure 6.35: The Gaussian fits of the 1332 keV ^{60}Co events, from figure 6.34. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue). The stringent volume cut has been made.

The 1460 keV ^{40}K peak is caused by the beta decay of ^{40}K to the 2+ angular momentum state of ^{40}Ar which then relaxes to the 0+ ground state, releasing an 1460 keV gamma. ^{40}K is present in materials throughout the the detector and has a half-life of 1.25 billion years.

The spatial distribution of ^{40}K is similar to that of ^{60}Co . The 1460 keV peak values in simulation changes from $\mu_{S1} = 6.17 \times 10^3$ phd and $\mu_{S2} = 1.86 \times 10^6$ to $\mu_{S1} = 5.78 \times 10^3$ phd and $\mu_{S2} = 1.97 \times 10^6$. After the stringent cut is applied, the SR1 peak values change from $\mu_{S1} = 5.49 \times 10^3$ phd and $\mu_{S2} = 1.96 \times 10^6$ to $\mu_{S1} = 5.16 \times 10^3$ phd and $\mu_{S2} = 2.04 \times 10^6$.

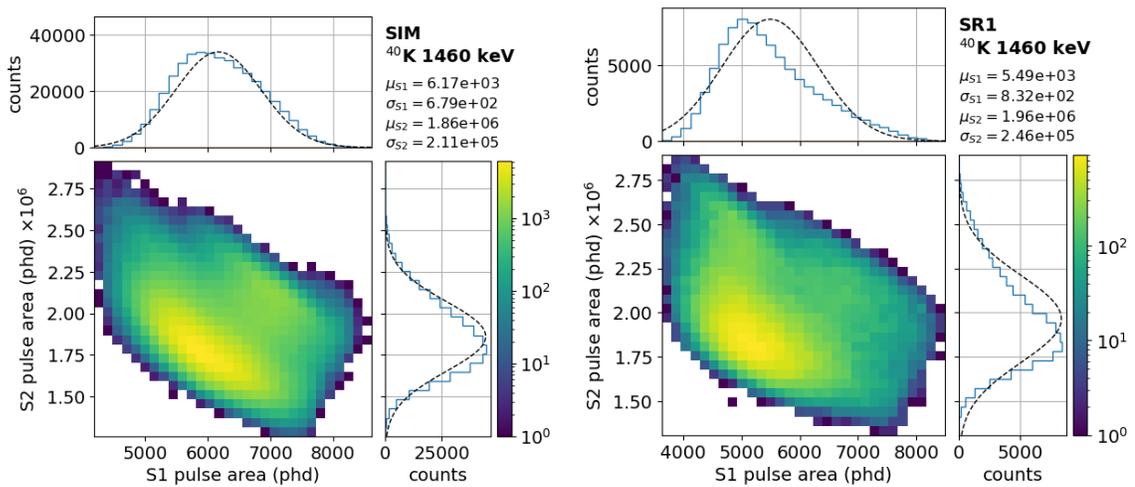


Figure 6.36: The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1460 keV ^{40}K events. The $S1$ and $S2$ values have been projected onto their respective axes and fit to a Gaussian. The fit parameters are shown on the top right corner of each plot. Only the standard fiducial cut has been made.

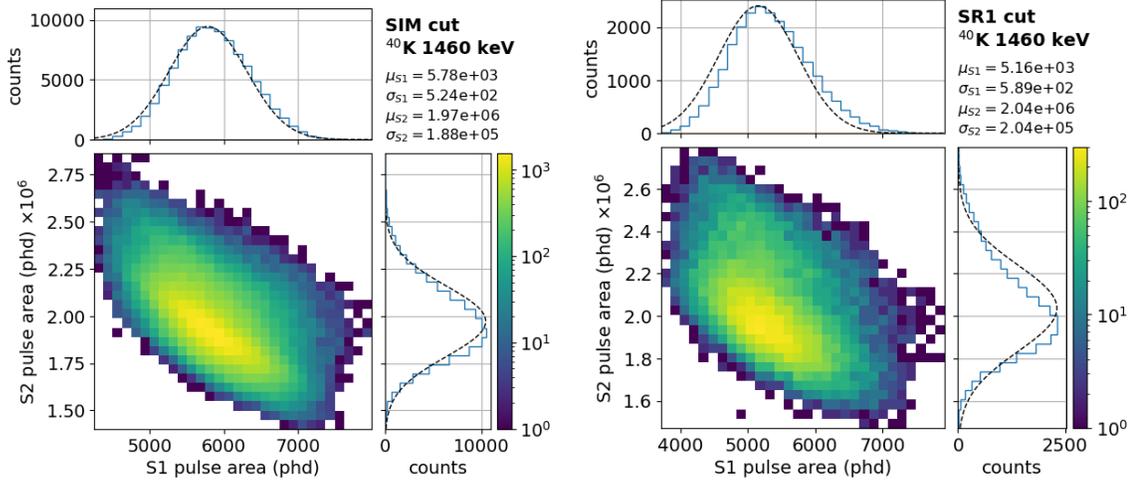


Figure 6.37: The $S1$ and $S2$ pulse areas from simulations (left) and SR1 (right) of the 1460 keV ^{40}K events. The stringent volume cut has been made.

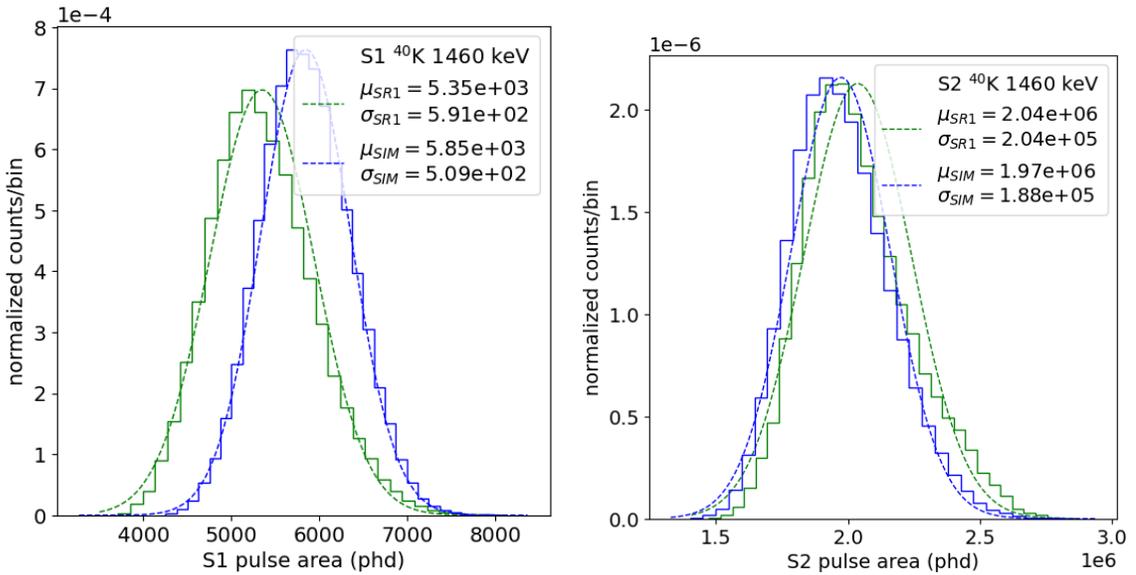


Figure 6.38: The Gaussian fits of the 1460 keV ^{40}K events, from figure 6.37. Fits show the $S1$ (left) and $S2$ (right) pulse area peaks from the SR1 data (green) and simulations (blue). The stringent volume cut has been made.

The same treatment is applied to the simulated excited state decay of ^{85}Kr . We consider the $S1_\gamma$ and $S2_{total}$ peaks identified within the standard fiducial volume, as shown in figure 6.39.

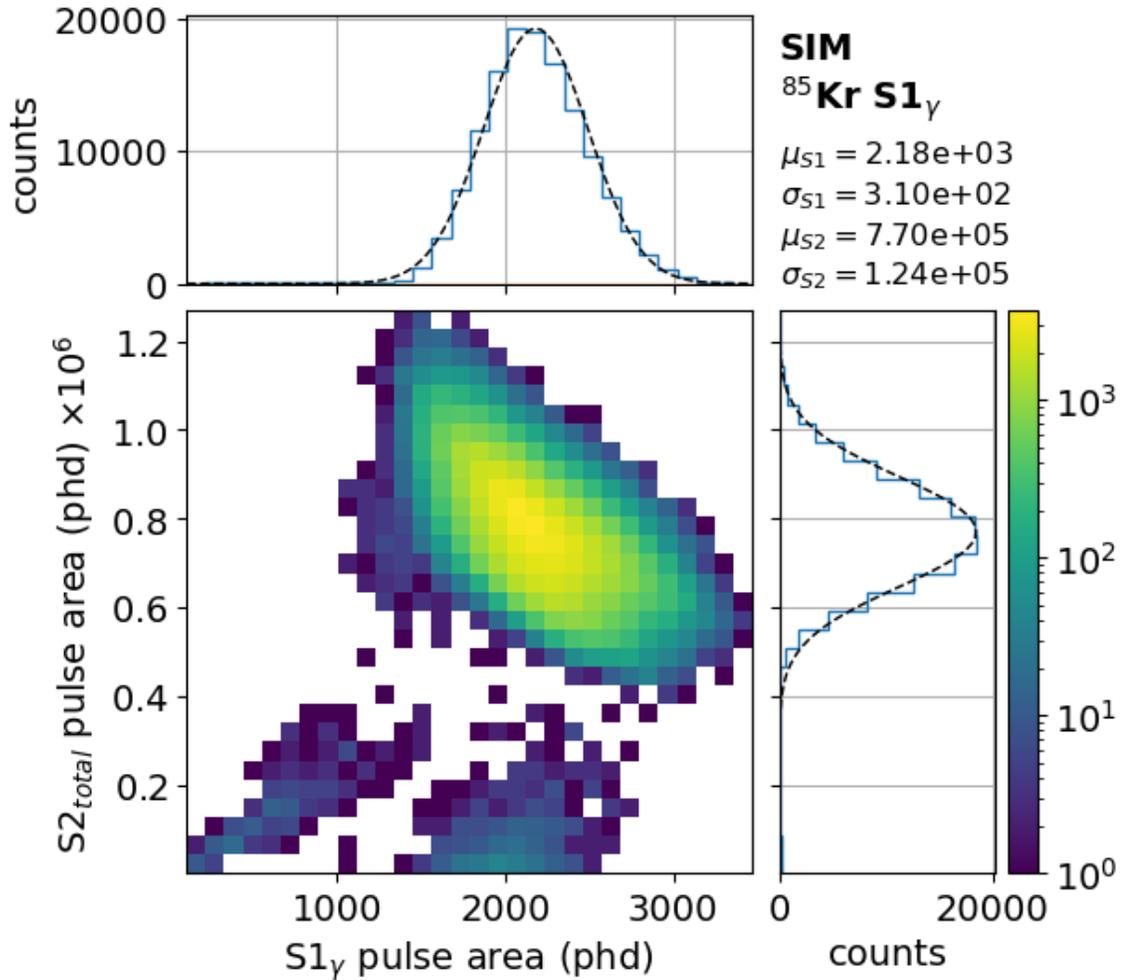


Figure 6.39: The $^{85}\text{Kr } S1_\gamma$ (514 keV) and $S2_{total}$ (514 to 687 keV) pulse areas from simulations of ^{85}Kr events. The standard fiducial volume cut is taken.

In figure 6.40, the $S1_{SIM}$ peak values from the sources defined in table 6.1 are plotted against the $S1_{SR1}$ peak values found in SR1. A linear fit of this data is used to calculate an equivalent $S1_{SR1}$ peak value from the $S1_{SIM}$ peak value.

$$S1_{SR1} = 0.95 \cdot S1_{SIM} + 25.82 \quad (6.12)$$

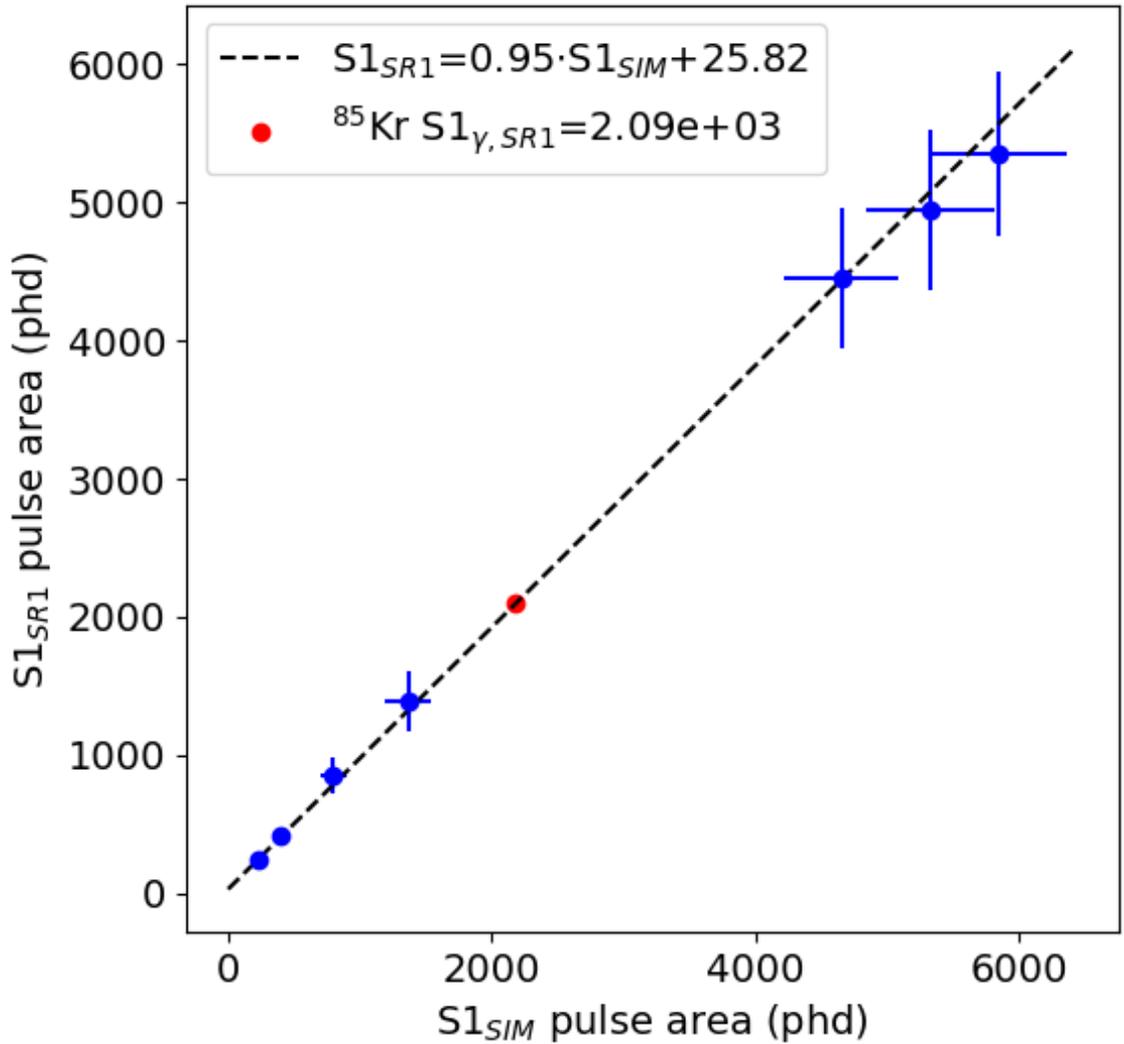


Figure 6.40: The ^{85}Kr $S1_{\gamma}$ peak value is shown in red. The black dotted line is a linear fit of the identified $S1$ peaks from table 6.1. The errors are the corresponding σ values found from the Gaussian fit of each peak.

For a simulated peak value of $S1_{\gamma,SIM} = 2181$ phd, the equivalent SR1 peak is $S1_{\gamma,SR1} = 2093$ phd. This amounts to an absolute difference of $|(1 - S1_{\gamma,SR1}/S1_{\gamma,SIM})| = 4\%$ disagreement between data and simulation. The percent difference is used to calculate a narrower and wider set of cuts for both $S1_{\beta}$ and $S1_{\gamma}$. We apply the same modification factor to the $S1_{\beta}$ case despite the 25.82 phd offset because the 39.5 keV is $S1$ peak

value differs by $|(1 - 233 \text{ phd}/229 \text{ phd})| = 1.7\%$ and the 67.3 keV $S1$ peak differs by $|(1 - 413 \text{ phd}/392 \text{ phd})| = 5.4\%$.

The original and modified cuts for $S1_\beta$ are shown in figure 6.41. The fractional error for cut2 is calculated using the number of surviving events from the two modified and the original cuts as described by equation (6.7).

$$\frac{\sigma_2}{N_2} = \frac{(N_{2,wide} - N_2) - (N_2 - N_{2,narrow})}{2 \cdot N_2} = \frac{117060 - 115855}{2 \cdot 116477} = 0.52\% \quad (6.13)$$

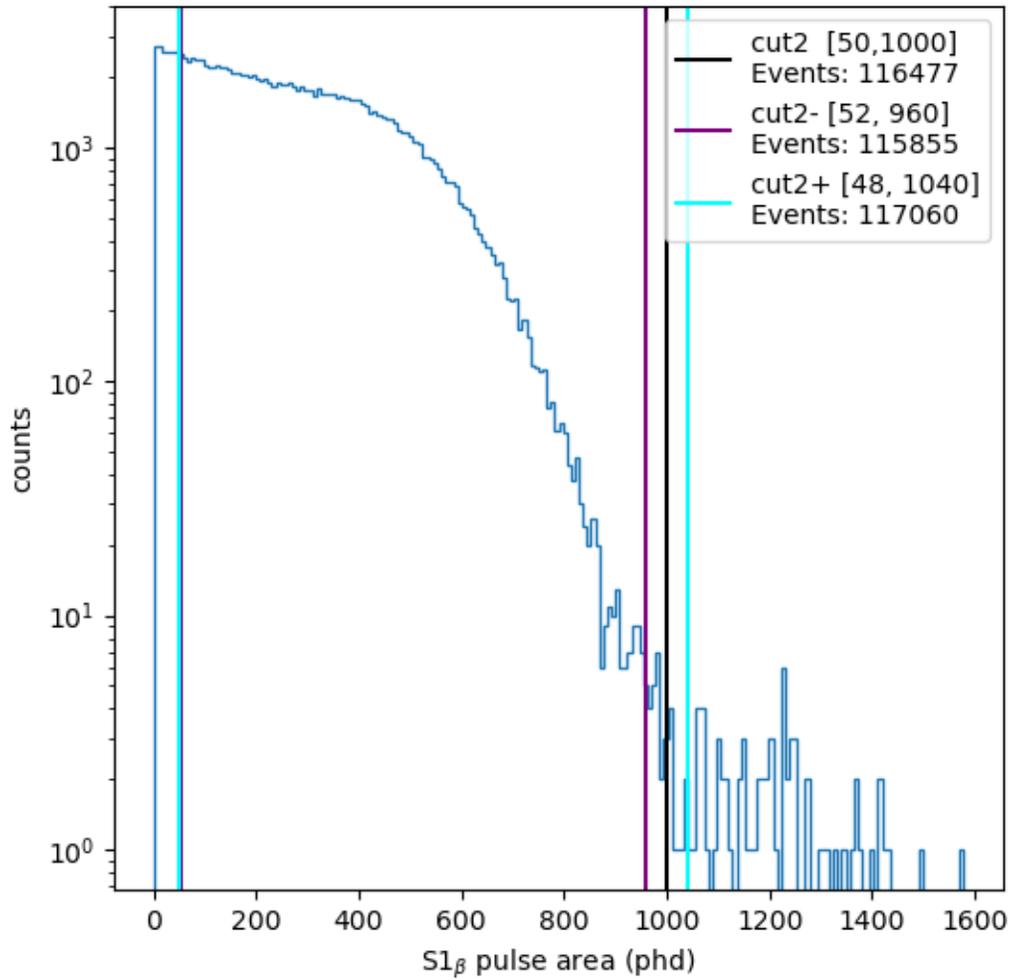


Figure 6.41: The original (black) and adjusted (cyan and purple) $S1_\beta$ cuts. The “cut-” and “cut+” denote the narrower and wider cuts respectively. The number of events satisfying each cut are given below the specified ranges.

The original and modified cuts for $S1_\gamma$ are shown in figure 6.42. The fractional error for cut3 is calculated using the same method.

$$\frac{\sigma_3}{N_3} = \frac{116298 - 115904}{2 \cdot 116215} = 0.17\% \quad (6.14)$$

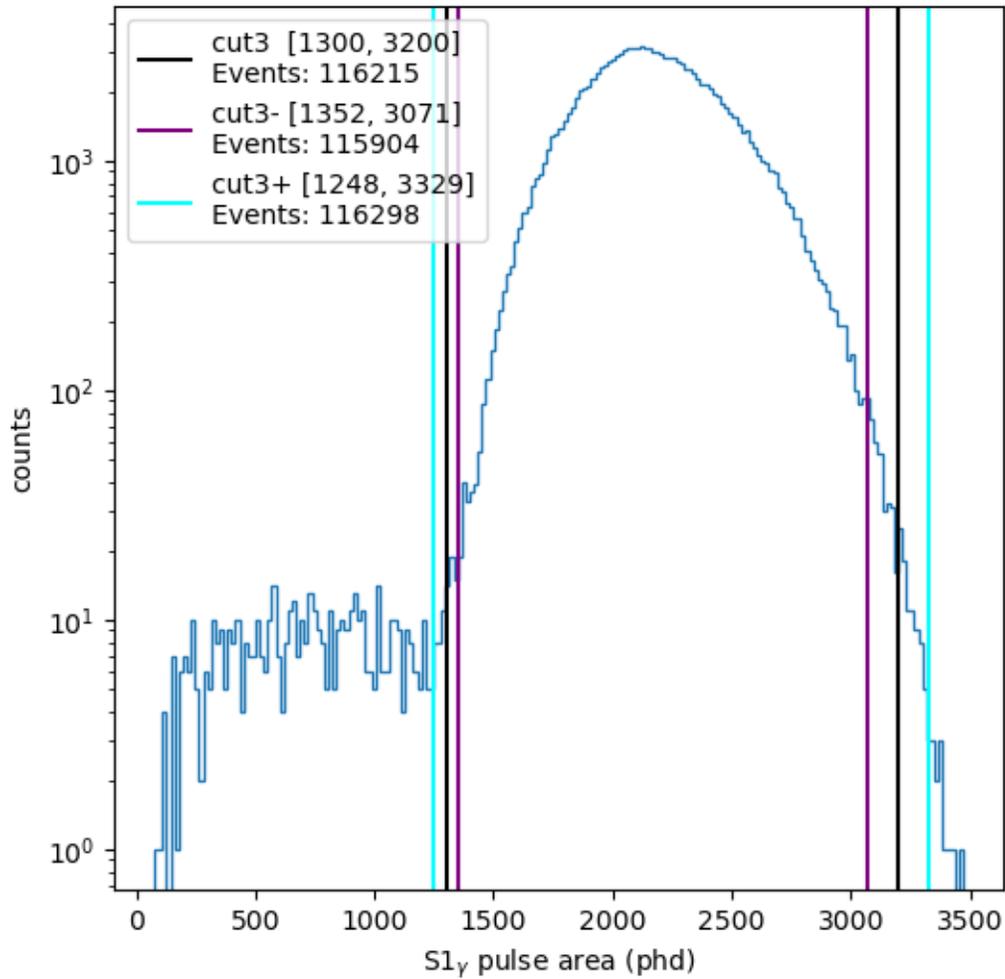


Figure 6.42: The original (black) and adjusted (cyan and purple) $S1_\gamma$ cuts.

In figure 6.43, the $S2_{SIM}$ peak values for the well defined peaks from table 6.1 are plotted against the $S2_{SR1}$ peak values for in the SR1 and DD data. A linear fit of this data is used to calculate an equivalent $S2_{SR1}$ peak value from the $S2_{SIM}$ peak value.

$$S2_{SR1} = 1.04 \cdot S2_{SIM} + 6483.93 \quad (6.15)$$

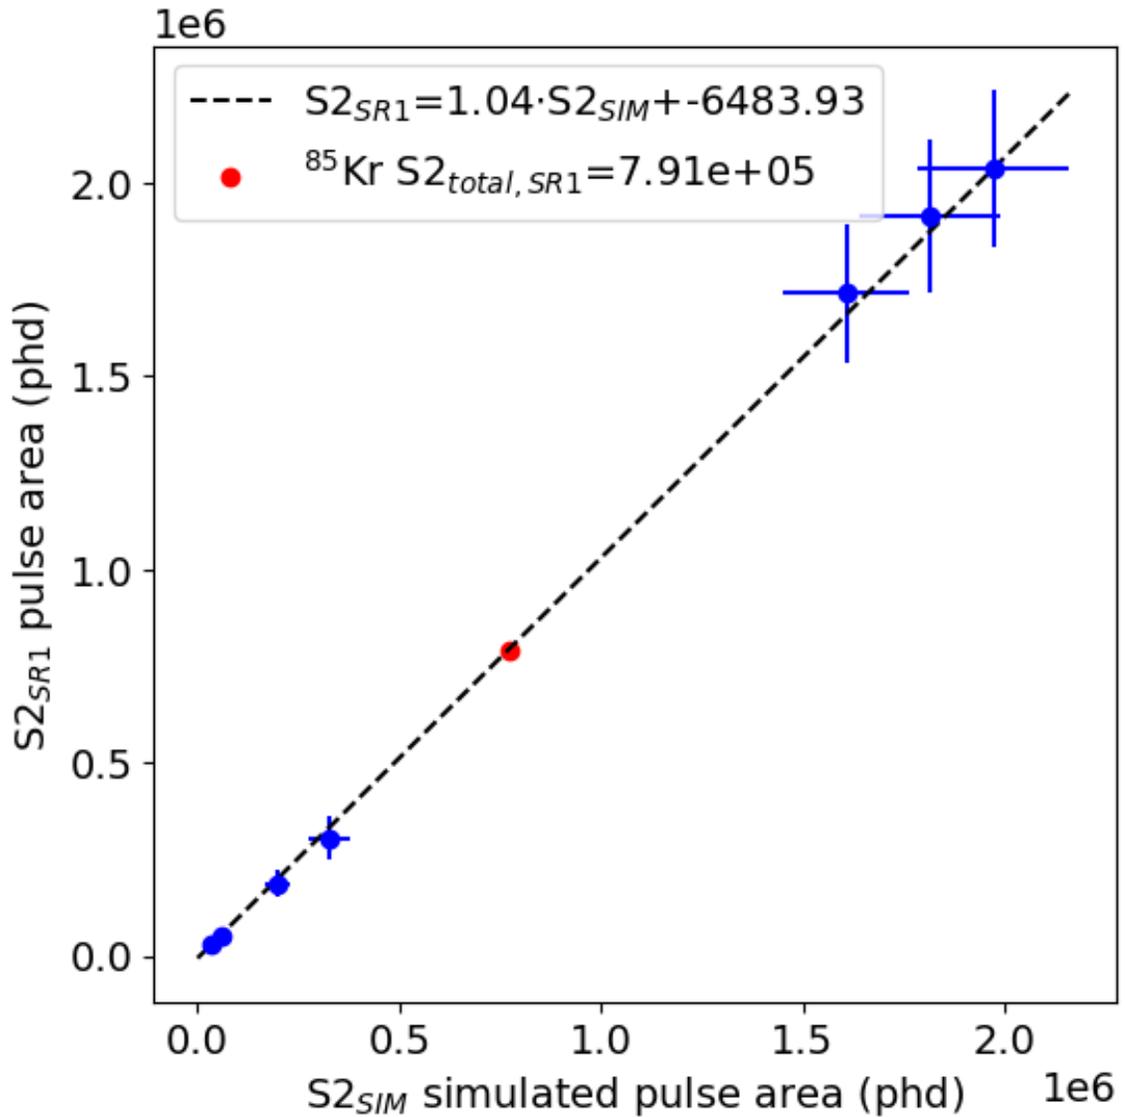


Figure 6.43: The ^{85}Kr , $S2_{total}$ peak value is shown in red. The black dotted line is a linear fit using the identified $S2$ peaks of each source in blue from table 6.1. The errors are the corresponding σ values found from the Gaussian fit of each peak.

For a simulated peak value of $S2_{total,SIM} = 7.70 \times 10^5$ phd, the equivalent SR1 peak is $S2_{total,SR1} = 7.91 \times 10^5$ phd. This amounts to an absolute difference of $|(1 - S2_{total,SR1}/S2_{total,SIM})| = 2.7\%$ disagreement between the peak values. The percent difference is used to calculate a narrower and wider set of cuts for $S2_{total}$. The fractional error for cut4 is calculated using

the same method.

$$\frac{\sigma_4}{N_4} = \frac{115814 - 115806}{2 \cdot 115814} \approx 0\% \quad (6.16)$$

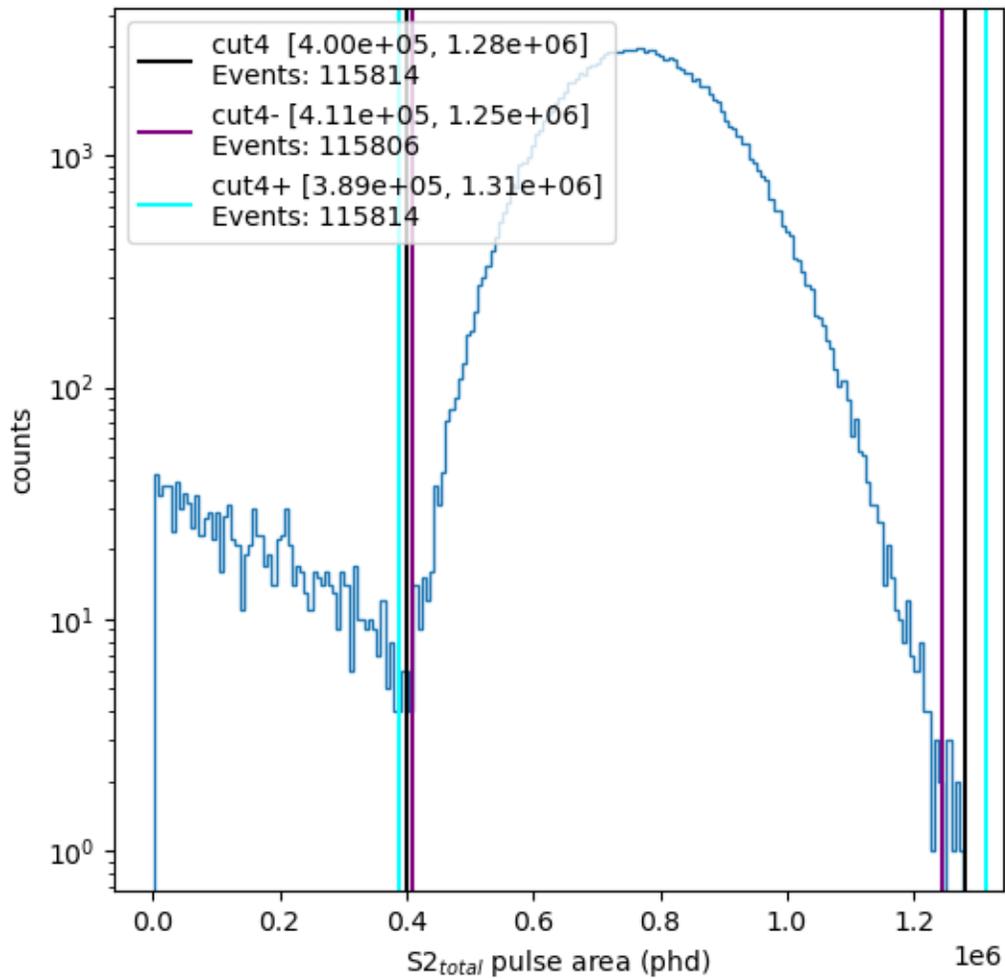


Figure 6.44: The original (black) and adjusted (cyan and purple) $S2_{total}$ cuts.

This validation analysis shows that the simulated $S1$ and $S2$ peak values match well with observed detector data. The analysis also shows that the errors introduced by cuts 2, 3, and 4 are all less than 1%.

6.4.5 Delay Time Window

The delay time window is defined to be between $0.15 \mu\text{s}$ and $100.15 \mu\text{s}$. The minimum time excludes events that are prone to pulse-merging where, LZap fails to recognize distinct $S1_\beta$ and $S1_\gamma$ pulses. The maximum time is selected to both maximize efficiency and also includes an interval in which random coincidence events could occur. Our final signal search will be carried out in this time delay space, where an excess of events with delay times around $1.015 \mu\text{s}$ are likely due to the excited state ^{85}Kr decays. Events with delay times significantly longer than our signal's $1.015 \mu\text{s}$ half-life will almost certainly be from random coincidence backgrounds. The delay time cut is referred to as “cut6”. Of the 114,964 that survive cut5, 109,019 events survive cut6.

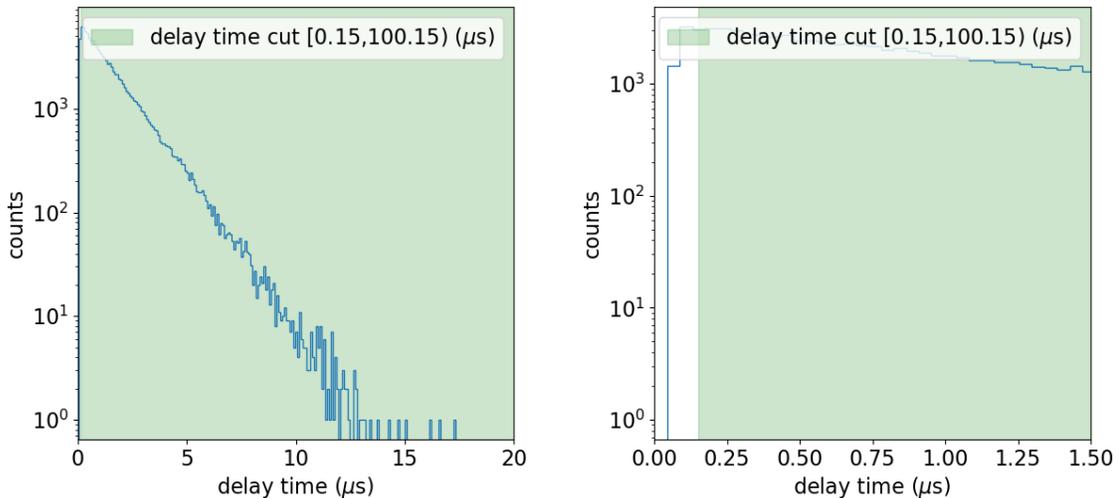


Figure 6.45: A histogram of events that survive up to cut5. All of the simulated events have a delay time less than $20 \mu\text{s}$ as shown of the left. The plot on the right is zoomed in and shows the first $1.5 \mu\text{s}$ to emphasize the $0.15 \mu\text{s}$ cut.

6.4.6 Localization

If an event has more than one $S2$ signal due to gamma scatters, the scatters should be clustered closely to the initial interaction. We define a localization variable to characterize how far the gamma has scattered. The localization (L) is calculated by taking the sum of the difference in position between and $S2_i$ pulse (\vec{r}_i for $i > 1$) and the $S2_{max}$ (\vec{r}_1) pulse. Each term is weighted by a the corresponding pulse are $S2_i/S2_{total}$. This is given in the equation below

$$L = \frac{\sum_{n=1}^5 S2_{area,n} \cdot |\vec{r}_1 - \vec{r}_n|}{\sum_{n=1}^m S2_{area,n}}. \quad (6.17)$$

L is zero for an event with only a single $S2$. L is larger for events with significant $S2$ pulse areas located further away the $S2_{max}$ interaction site. This cut helps to remove events due to random coincidences.

Following the initial fill and DD neutron calibrations, this analysis experiences an uptick in events caused by random coincidences due to xenon activation. Activated sources like ^{125}I (59.5 day half-life and 39.5 keV or 67.3 keV gamma plus x-ray), ^{131m}Xe (11.9 day half-life and 164 keV gamma) or ^{129m}Xe (8.9 day half-life and 196.6 keV plus 39.6 keV gammas) have energies that could satisfy the uncorrected $S1_\beta$ pulse area cuts and ^{127}Xe (36.5 day half-life and energies ranging from 204 keV to 408 keV due to various gamma and x-ray combinations) could be mistaken as either an $S1_\beta$ or $S1_\gamma$. Figure 6.46 shows the time structure of events over the course of the SR1 exposure in days.

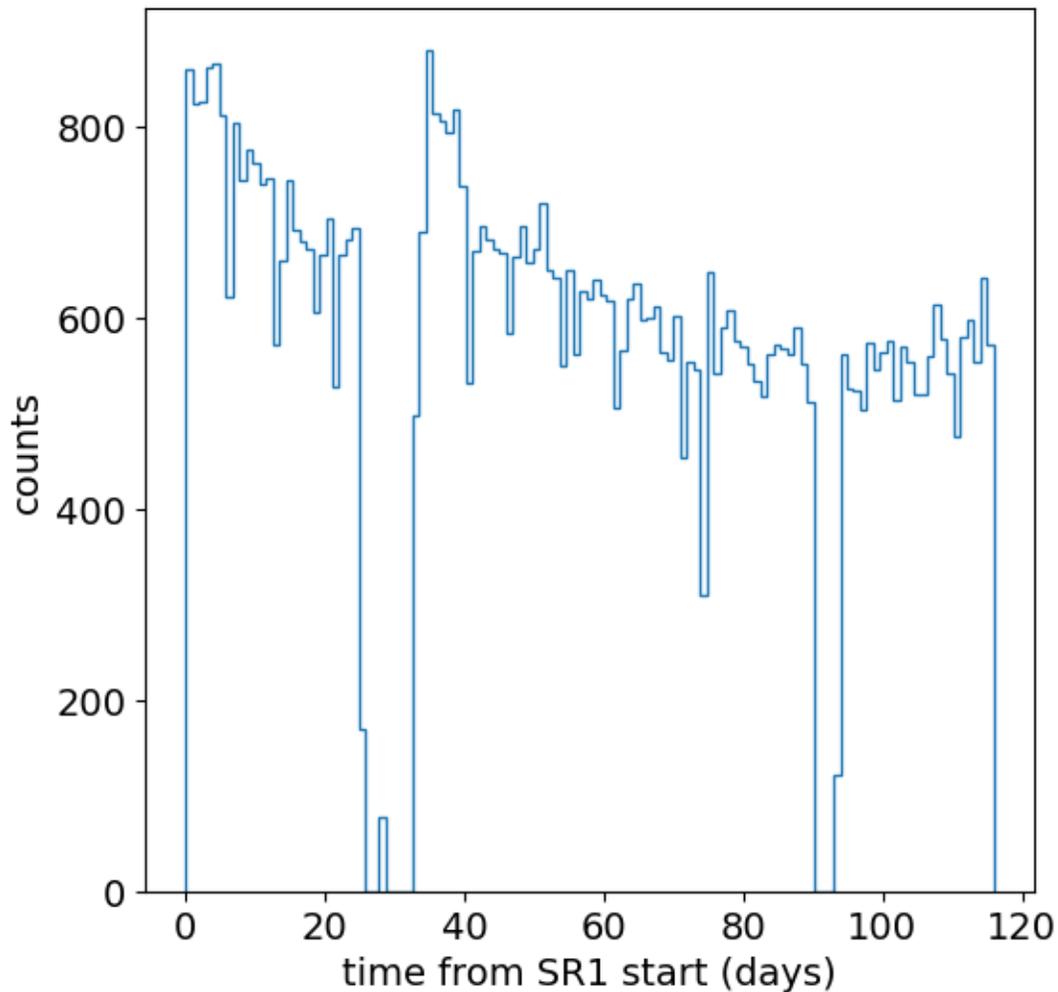
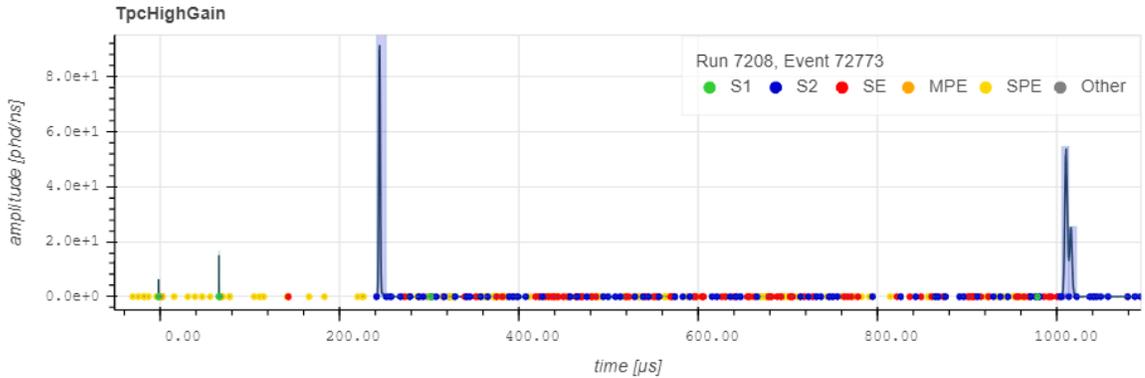


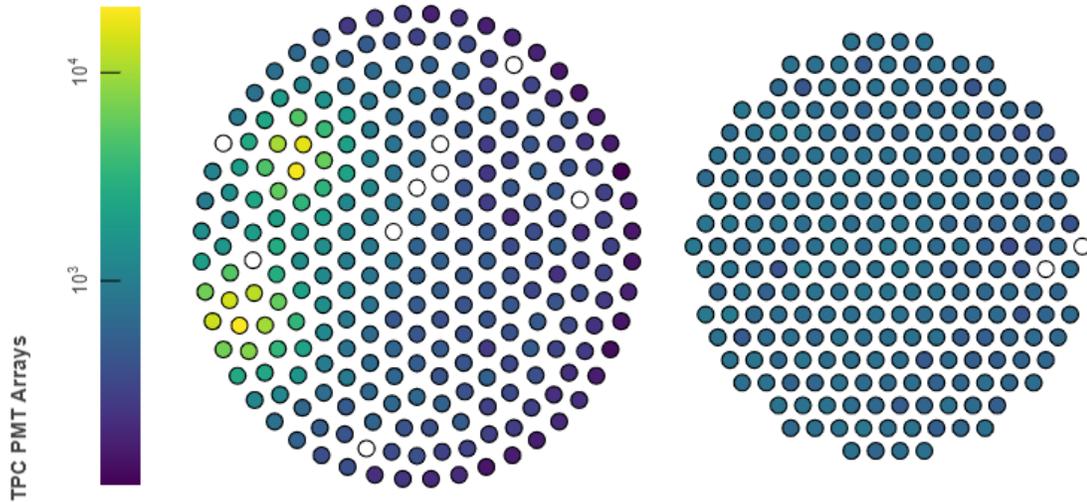
Figure 6.46: A histogram of events from SR1 data over the course of the 120 SR1 calendar days. The bin size is one day. The 68,916 events shown have been identified by the selection criteria, but fail cut7. The events excluded by the localization cut have a clear time structure that appears to correlate with random coincidences caused by cosmogenic and neutron activation following the initial fill (near day zero) and DD neutron calibrations (near day 30).

An event from day seven from the start of SR1 with a localization value of 76.1 cm is shown in figures 6.47a and 6.47b. This event satisfies all cuts prior to the localization cut but is removed due to the clear separation of the event locations. This is likely a random

coincidence event caused by increased activity from xenon activation.



(a) The $S1_{\beta}$ and $S1_{\gamma}$ pulses.



(b) The PMT hit map.

Figure 6.47: The pulse waveform and PMT hitmap of an event from SR1 day 7. Event 72773 from run 7208 occurred seven days into SR1 and satisfies cut0 through cut6 but fails localization. Note the significant separation between the two clusters of $S2$ light characterized by a localization value of 76.1 cm. The event has a delay time of 67.4 μs .

The ultimate cutoff is determined by simulations. We require that the localization be less than or equal to 9 cm, as shown in figure 6.48. Simulations also show that there are never more than five $S2$ pulses of significance. In the interest of increasing analysis speed, only the five largest $S2$ pulses are considered. Of the 109,019 that survive cut6, 108,995 events survive cut7.

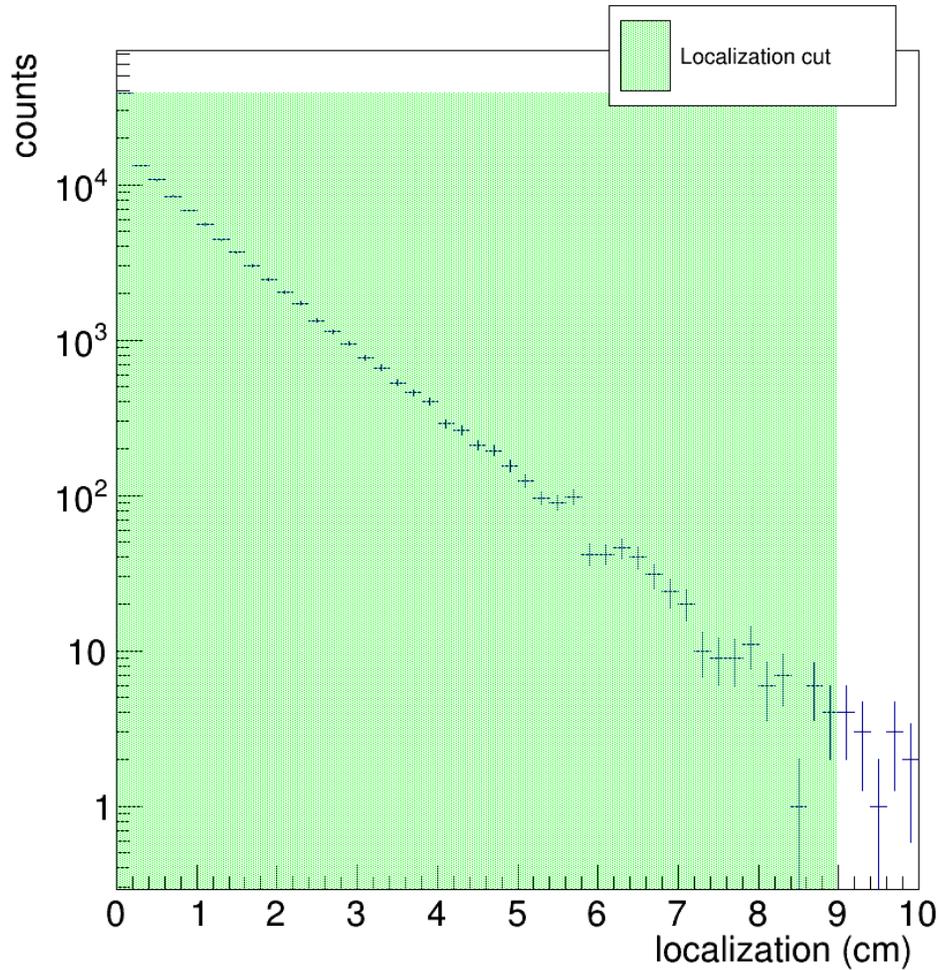


Figure 6.48: A histogram of event localization that survive up to cut6. The localization cut is shown in green.

This marks the final cut informed by the fast simulation. The cumulative efficiency of each cut and its error in number of events are given in table 6.2. From the cumulative cuts discussed this far we have determined the first term from equation (6.6)

$$\frac{N_{7,fast}}{N_{i,fast}} = \frac{108995}{147175} = 74.1\% \quad (6.18)$$

The final analysis used to identify signal events in the the surviving delay time histogram relies on a two-parameter likelihood fit of a decay exponential with a constant offset

$$f(t) = A \cdot e^{\frac{-t \cdot \ln(2)}{1.015}} + B \quad (6.19)$$

between $t = 0.15$ and $100.15 \mu\text{s}$. The fit floats the signal coefficient A and the constant background term B , both in units of events per bin.

To calculate the total number of signal events, the signal term in equation (6.20) is integrated over the entire event window and converted to the appropriate units through a scale term, c_{scale} , equal to number of bins per μs .

$$N_{signal} = c_{scale} \int_{0.15}^{100.15} A \cdot e^{\frac{-t \cdot \ln(2)}{1.015}} dt = 1.322 \cdot c_{scale} A \quad (6.20)$$

To calculate the total number of background events, the scaled background term is integrated over the entire event window.

$$N_{background} = c_{scale} \int_{0.15}^{100.15} B dt = 100 \cdot c_{scale} B \quad (6.21)$$

This technique is applied to events from the fast simulation that survive cut7, as shown in figure 6.49. For a $0.1 \mu\text{s}$ bin size $c_{scale} = 10$ bins per μs . The fit returns the scaled signal parameter of $A = 82478 \pm 250$ events per μs and the scaled background parameter of $B = 0 \pm 0.006$ events per μs . After integrating over the signal window, the calculated number of signal events is $N_{signal,fast} = 109016 \pm 330$ and background events is $N_{background,fast} = 0 \pm 0.6$. This is in agreement with the known number of surviving signal events $N_{7,fast} = 108995$ and zero background events.

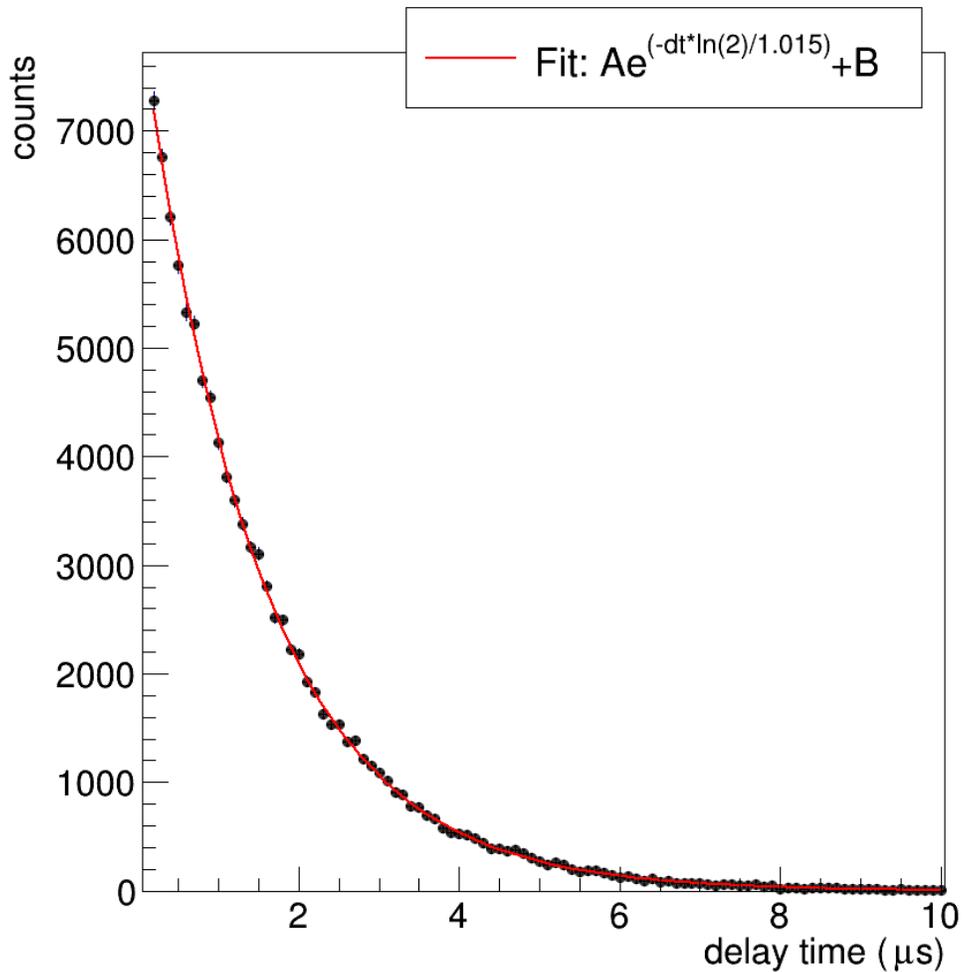


Figure 6.49: A delay time histogram of 108995 events that survive the fast simulation cut7. The bin size is $0.1 \mu\text{s}$ and the scale factor is $c_{scale} = 10$ bins per μs . The red line is fit to all events between 0.15 and $100.15 \mu\text{s}$, but the plot is zoomed in to the first $10 \mu\text{s}$. Integration of the scaled signal component accurately accounts for all of the signal events and a zero background.

6.4.7 ΔTBA

The TBA value is used as a proxy for the z position of interaction sites using only $S1$ signals. The ΔTBA value is a measure of the gamma and beta separation in z which should

be fairly small. The absolute difference in the TBA values of $S1_\beta$ and $S1_\gamma$ is

$$|\Delta TBA| = |TBA(S1_\gamma) - TBA(S1_\beta)| \quad (6.22)$$

Since TBA values require PMT hit information to calculate, the full simulation is used to inform the ΔTBA cut, or “cut8”.

Before continuing with the TBA analysis, we must apply the corresponding cuts from the fast simulation to the full simulation, up to cut7. Most cuts remain the same with the exception of 3 and 5 as shown in figures 6.51 and 6.53. This is due to the shift in the peak location of $S1_\gamma$, which in the full simulation has a Gaussian mean of 2671 phd, shown in figure 6.50, compared to the fast simulation value of 2181 phd shown in figure 6.13. Although the peak location of the $S2_{total}$ is centered at 7.23×10^5 phd, compared to the fast simulation value of 7.70×10^5 phd, we have chosen to leave the cut as is because the difference is negligible.

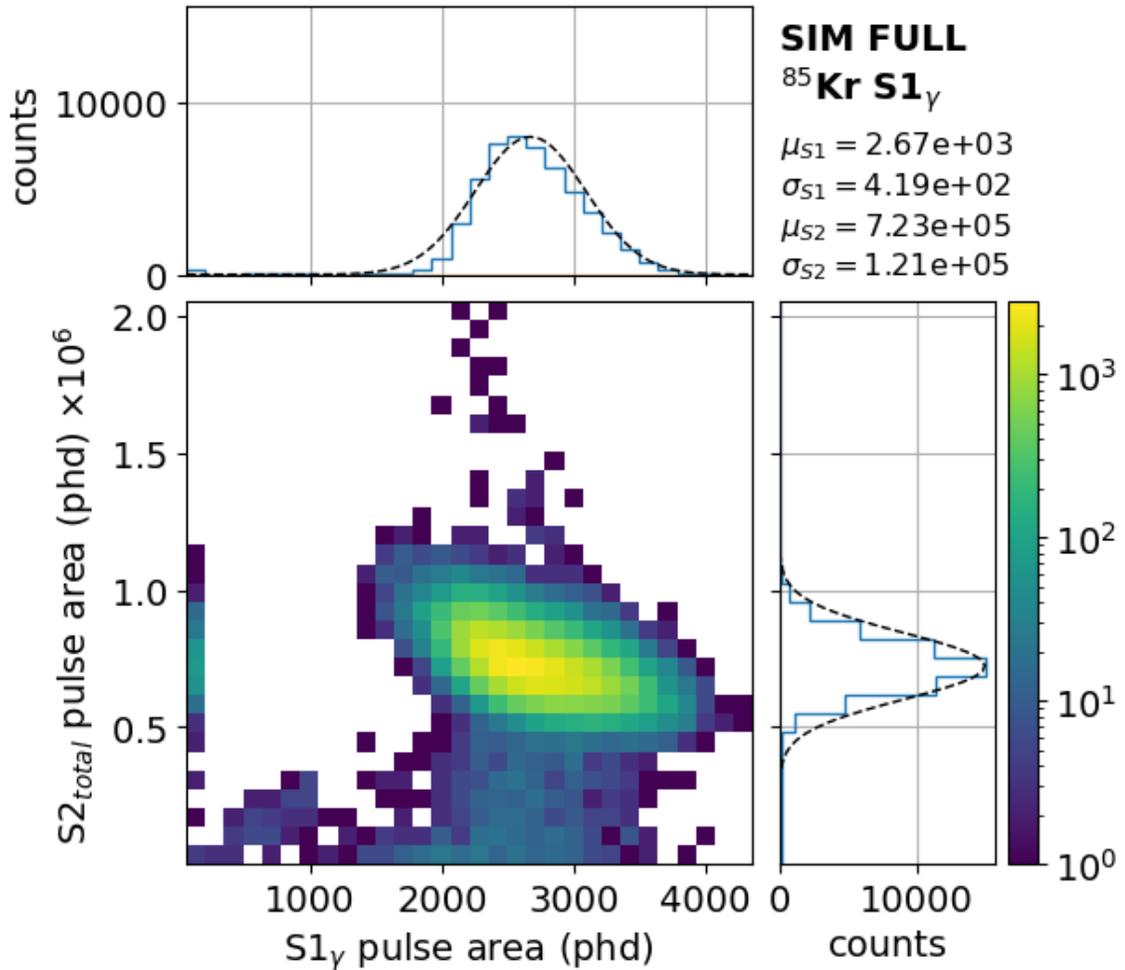


Figure 6.50: A histogram of $S1_\gamma$ pulse areas from the full simulation versus the $S2_{total}$ pulse area with Gaussian fits applied.

After validating the pulse areas between the fast simulation and the SR1 and DD data, it is clear that the full simulation is not as well tuned to detector response as the fast simulation. Figures 6.51, 6.52, 6.53, 6.54, and 6.55 show events from full simulation with their corresponding cuts overlaid.

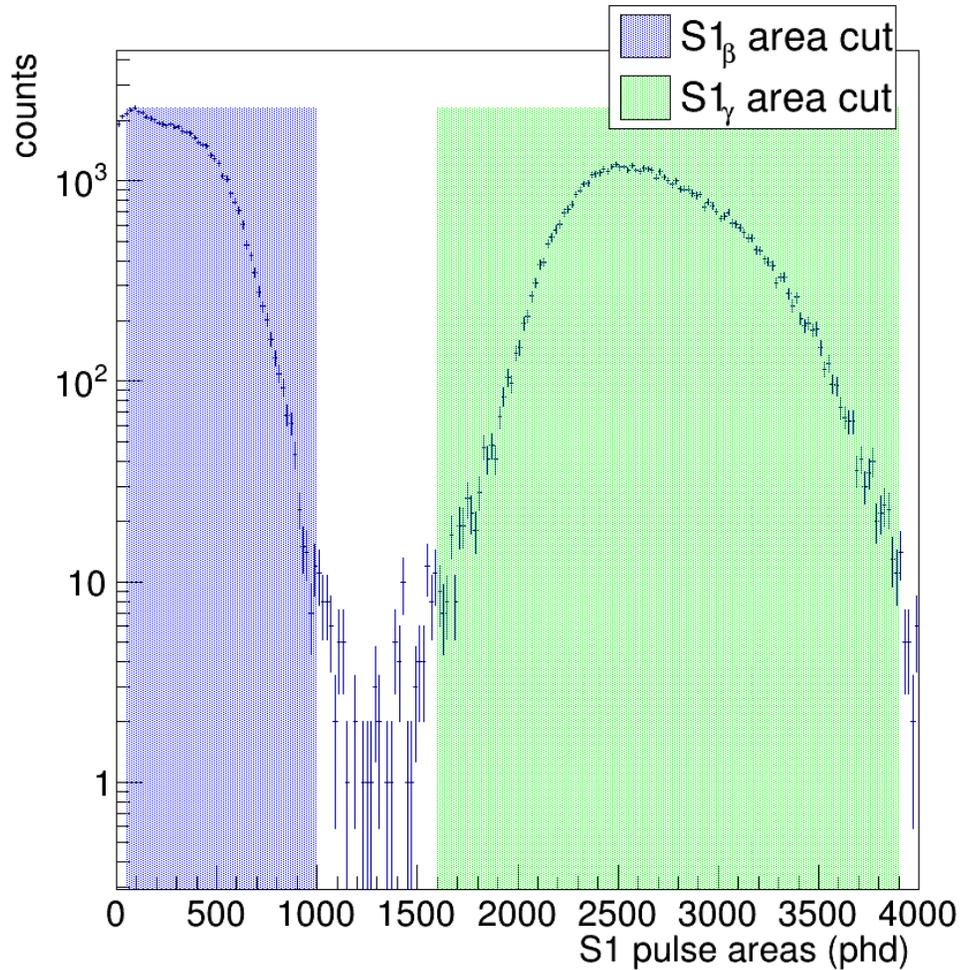


Figure 6.51: A histogram of $S1$ pulse areas from the full simulation. The blue shaded area accepts any $S1_\beta$ pulse with an area between 50 and 1000 phd. The green shaded area accepts any $S1_\gamma$ pulse with an area between 1600 and 3900 phd. The full simulation $S1_\gamma$ peak is centered at a pulse area 18.3% larger than the peak location identified and validated in the fast simulation. Only the selection criteria, or cut0, has been applied to emphasize events excluded by the cut overlay.

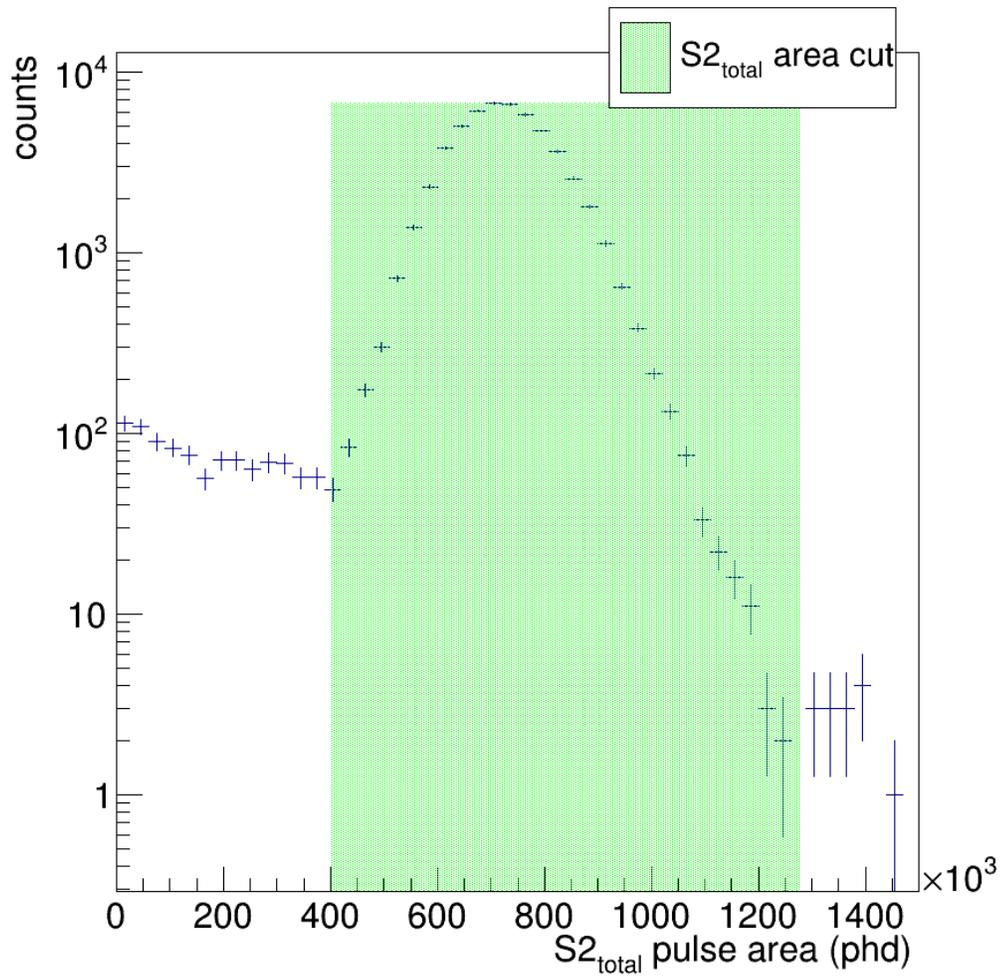


Figure 6.52: A histogram of $S2_{total}$ pulse areas from the full simulation. The green shaded area accepts all events with an $S2_{total}$ area between 0.4×10^6 and 1.28×10^6 phd. The full simulation $S2_{total}$ peak is centered at a pulse area 5.5% smaller than the peak location identified and validated in the fast simulation. Only the selection criteria, or cut0, has been applied to emphasize events excluded by the cut overlay.

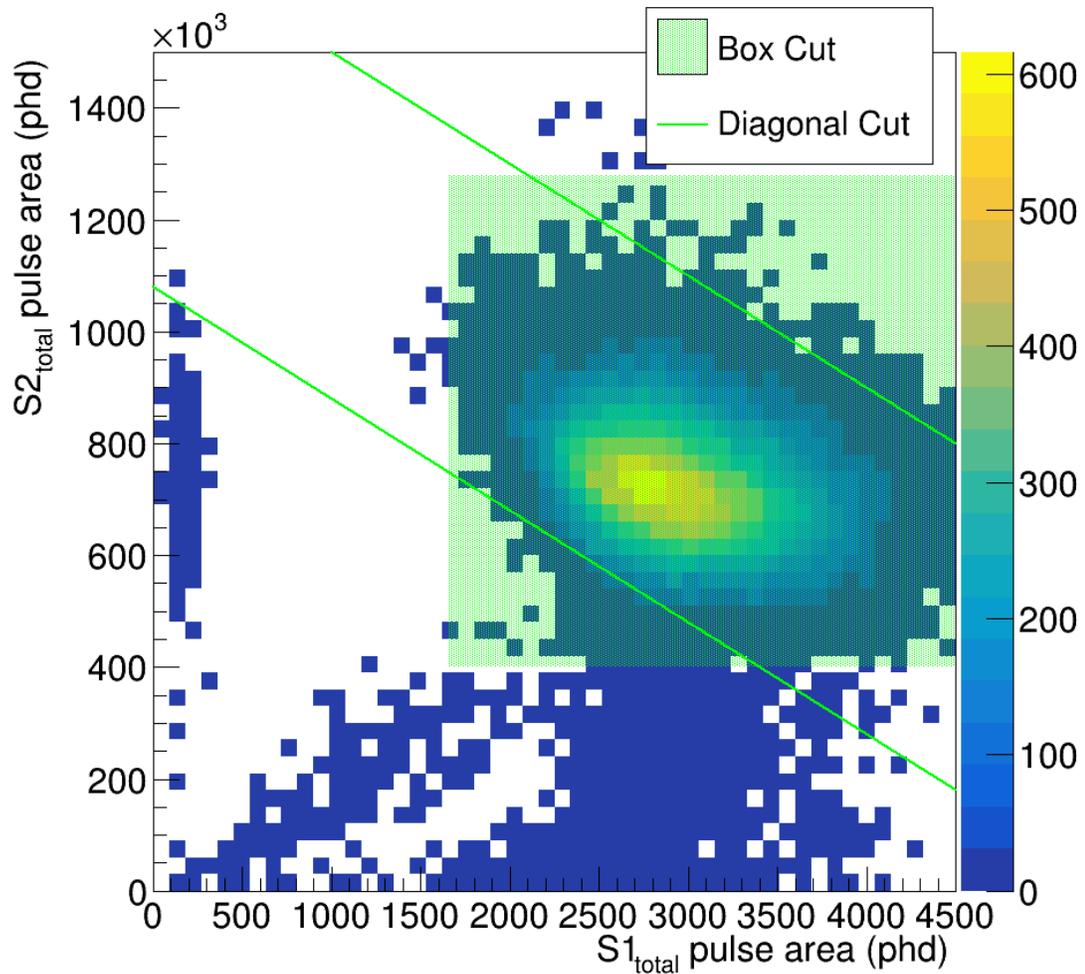


Figure 6.53: The $S1_{total}$ pulse area versus the $S2_{total}$ pulse area from the full simulation. The green shaded area is not explicitly a cut, but represents the combined $S1_{\beta}$, $S1_{\gamma}$, and $S2_{total}$ pulse area cuts. The $S1_{total}$ “cut” extends from 1650 to 4900 phd, which is beyond the x-axis limit shown here. The same $S2_{total}$ area cut from 0.4×10^6 and 1.28×10^6 phd is shown along the y-axis. The two diagonal lines represent cut4. Only the selection criteria, or cut0, has been applied to emphasize events excluded by the cut overlay.

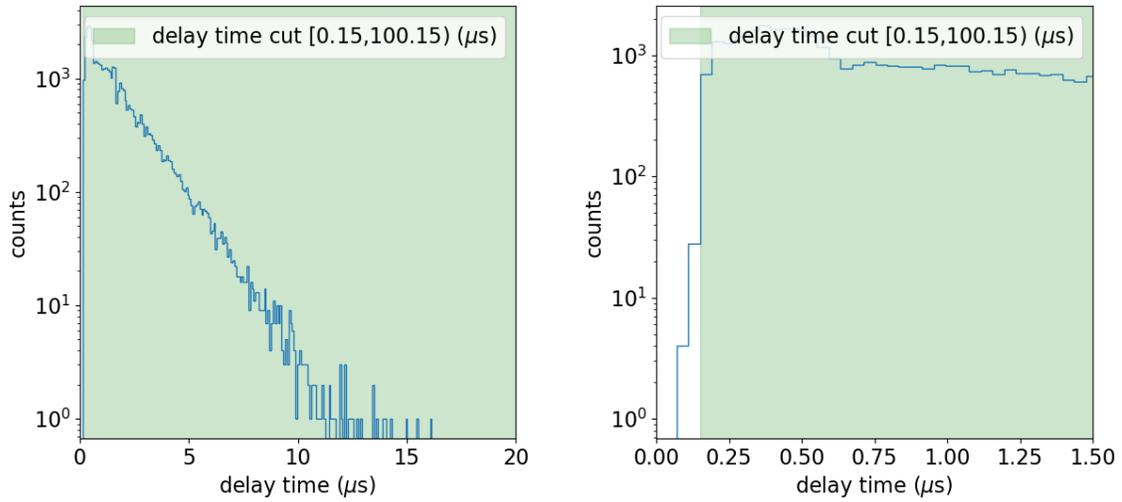


Figure 6.54: A histogram of full simulation events and their delay time. All of the events have a delay time less than $20 \mu\text{s}$, as shown on the left. The plot on the right is zoomed in and shows the first $1.5 \mu\text{s}$. Only the selection criteria, or `cut0`, has been applied to emphasize events excluded by the cut overlay.

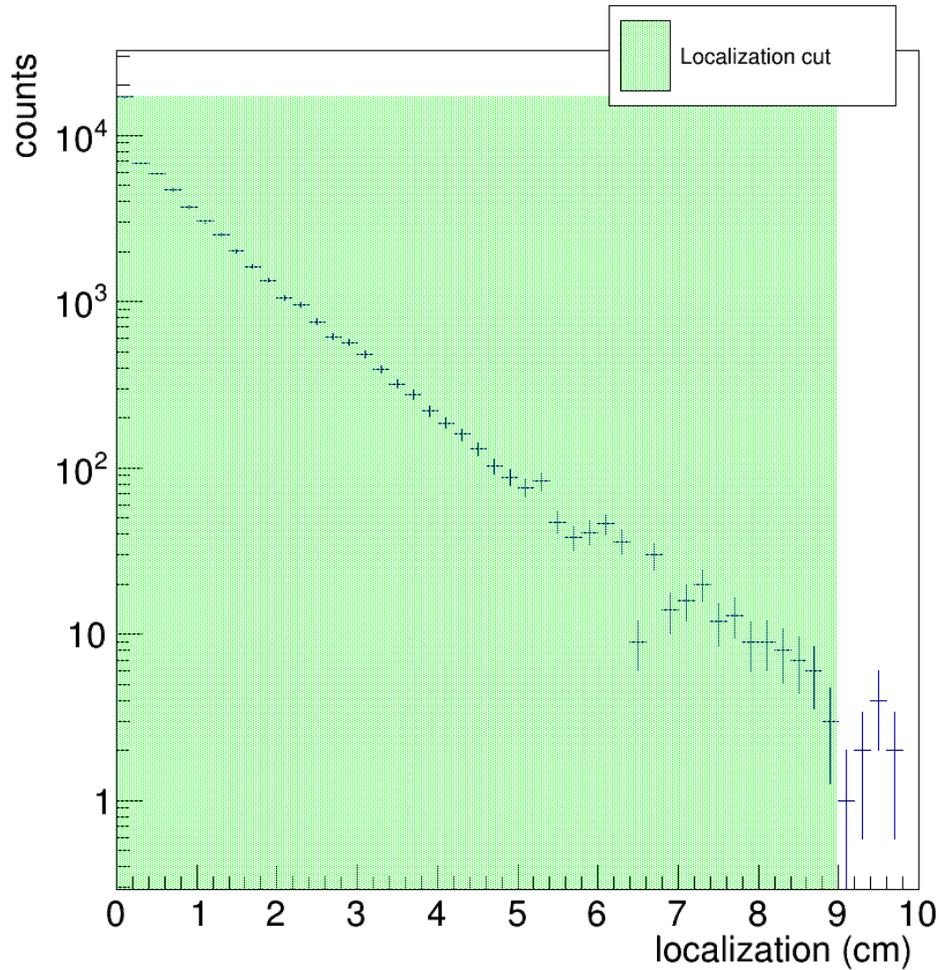


Figure 6.55: A histogram of full simulation event localization values. The localization cut is shown in green. Only the selection criteria, or cut0, has been been applied to emphasize events excluded by the cut overlay.

The full simulation initially contained 68,636 excited state ^{85}Kr events generated within a fiducial truth volume. After applying equivalent cuts up to cut7, the number of surviving events is found to be $N_{7,full} = 46,929$. This is the starting point for calculating the cumulative efficiencies of cut8 and cut9, as described in equation (6.6).

The ΔTBA cut value is determined from figure 6.56 as $|\Delta TBA| \leq 0.25$. Of the 46,929

events that survive cut7, 46735 survive cut8.

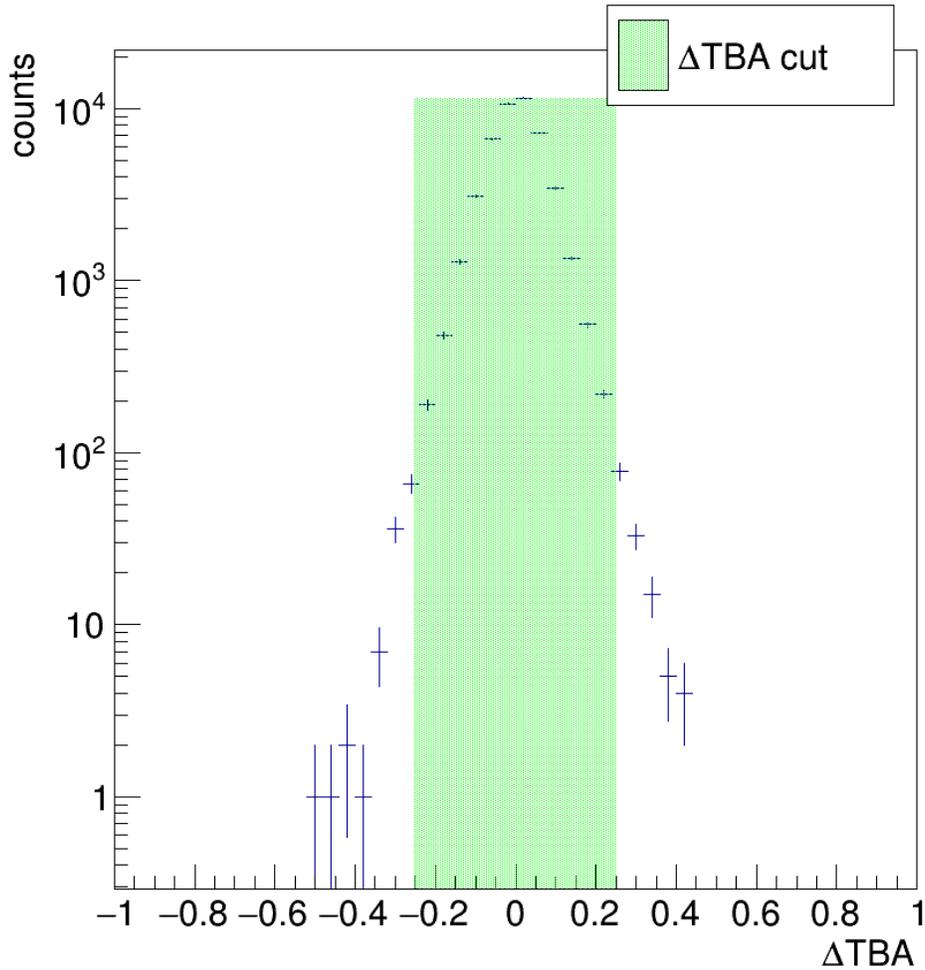


Figure 6.56: A histogram of ΔTBA values from full simulation events that survive cut7.

6.4.7.1 ΔTBA Validation

To validate the full simulation, the drift time and TBA distribution is compared between ^{85}Kr full simulation and ^{125}Xe (276 keV electron capture) detector data, as shown in figure 6.57. The purpose of this validation is to ensure that $S1$ TBA values can accurately represent

the reconstructed z positions in both simulation and SR1 data. The 514 keV $S1_\gamma$ TBA value and reconstructed position is used from ^{85}Kr and the analysis is carried out using uncorrected RQs. Once validated the ΔTBA cut can confidently be applied.

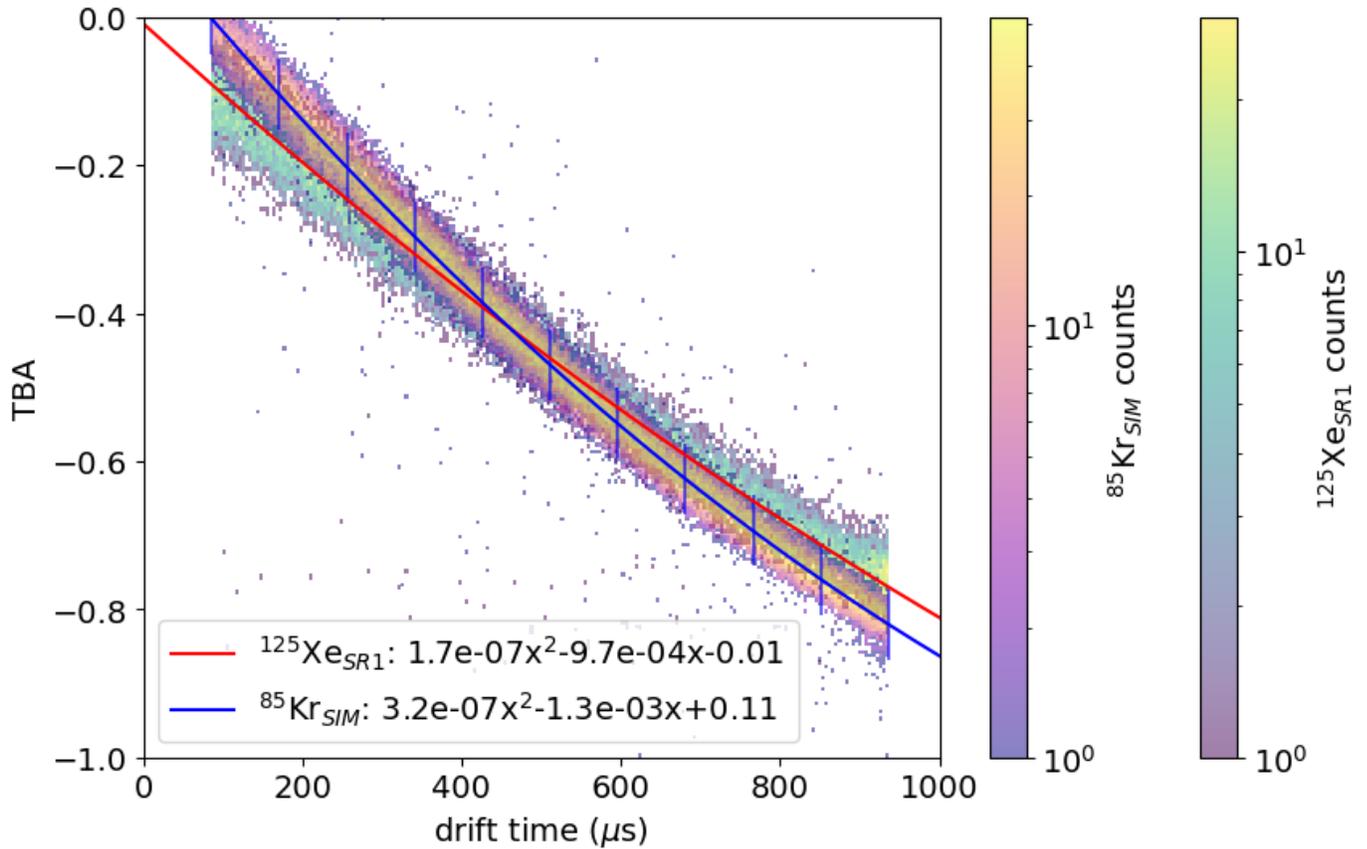


Figure 6.57: A 2-D histogram comparing full simulation ^{85}Kr reconstructed drift times and TBA values overlaid on SR1 ^{125}Xe reconstructed drift times and TBA values. The ^{85}Kr fit is shown in blue with vertical lines to identify binning used to calculate the difference in TBA values at a fixed drift time. The red line is the SR1 ^{125}Xe fit.

A second order polynomial is fit to both distributions, and the difference in TBA between the two at each drift time is a measure of how well the full simulation agrees with SR1 data. To quantify the agreement, the ^{85}Kr data is broken into 10 evenly spaced bins in drift time, and the average drift time \bar{t}_i for each bin, i , is calculated. This average drift time is then used to calculate a corresponding TBA value in both SR1 and full simulation,

as described by their polynomial fits ($TBA_{SIM,SR1}(\bar{t}_i)$). The error in the simulated TBA is the weighted root mean square of the difference between the calculated TBA values given by

$$\delta_{TBA} = \sqrt{\frac{\sum_{i=1}^{10} n_i (TBA_{SR1}(\bar{t}_i) - TBA_{SIM}(\bar{t}_i))^2}{N_{7,full}}} \quad (6.23)$$

where n_i is the number of ^{85}Kr events in bin i . The weighted RMS value is $\delta_{TBA} = 0.038$. The uncertainty in the efficiency is calculated by adjusting the ΔTBA cut to have a narrower and wider set of cuts, as shown in figure 6.58. We take the average difference in surviving counts from our selected cut value of 0.25 to find the fractional error. The fractional error for cut8 is calculated using the same method.

$$\frac{\sigma_8}{N_8} = \frac{46846 - 46415}{2 \cdot 46735} = 0.46\% \quad (6.24)$$

This validation could be improved by including ^{125}Xe multi-scatter events which would include some spatially dispersed $S1$ TBA signals from gamma scatters.

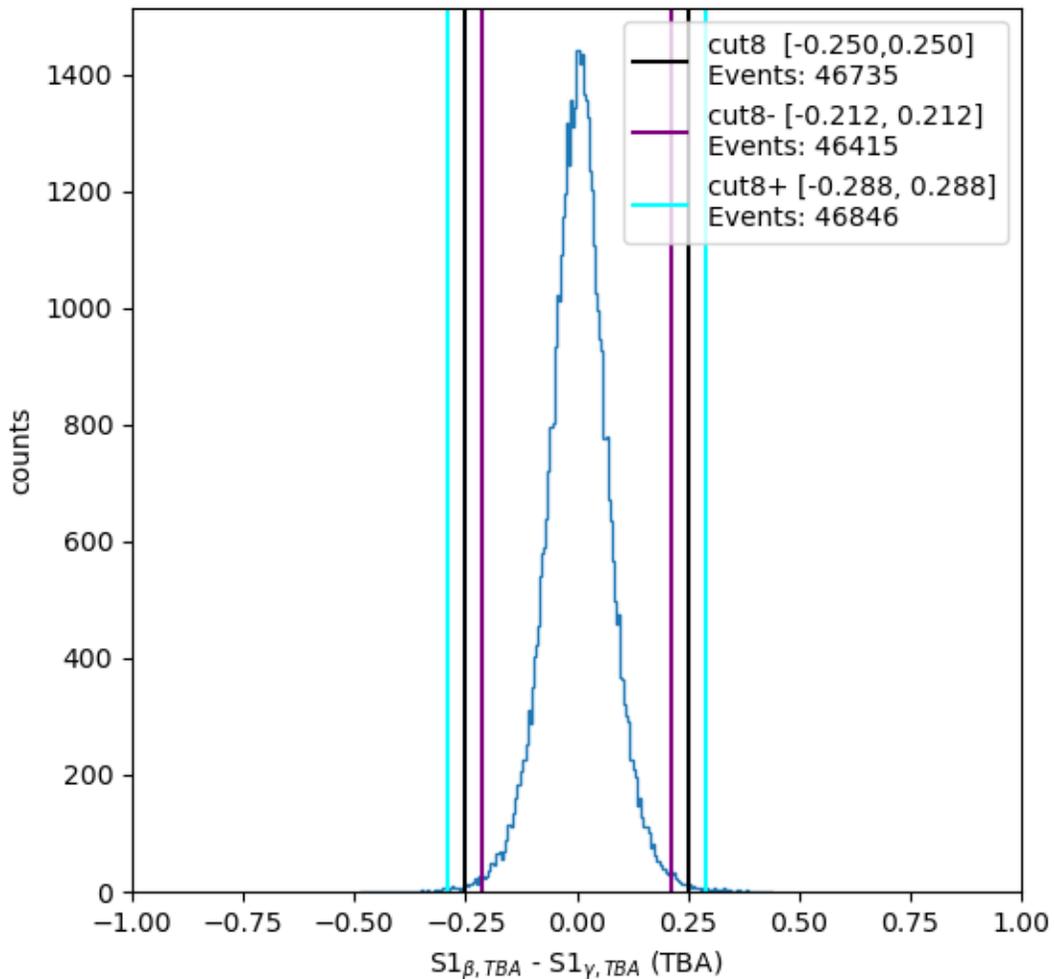


Figure 6.58: The original (black) and adjusted (cyan and purple) ΔTBA cuts.

6.4.8 Skin Veto

The final cut uses the skin veto, as described in section 3.1.2, to reject random coincidence events from the skin region. The efficiency of “cut9” must be evaluated as there is a chance that the ^{85}Kr gamma scatters in the skin and triggers a veto. The full simulation is used, as the fast simulation does not generate a skin detector RQs.

Applying the skin veto cut results in no additional events being excluded. This is because events that deposit enough energy outside of the TPC volume to trigger a skin veto lack sufficient $S1$ and $S2$ signals to satisfy all four of the pulse area cuts. Of the 68,636 full simulation excited state ^{85}Kr events generated within the fiducial truth volume, 55,369 met the minimum selection criteria and were reconstructed for further analysis. Of these events only 140 triggered a skin veto. Figures 6.59, 6.60, and 6.61, show the $S1$ and $S2$ distribution of these events failing to meet the required energy cuts. All of these events are clustered near the detector edge, shown in figure 6.62.

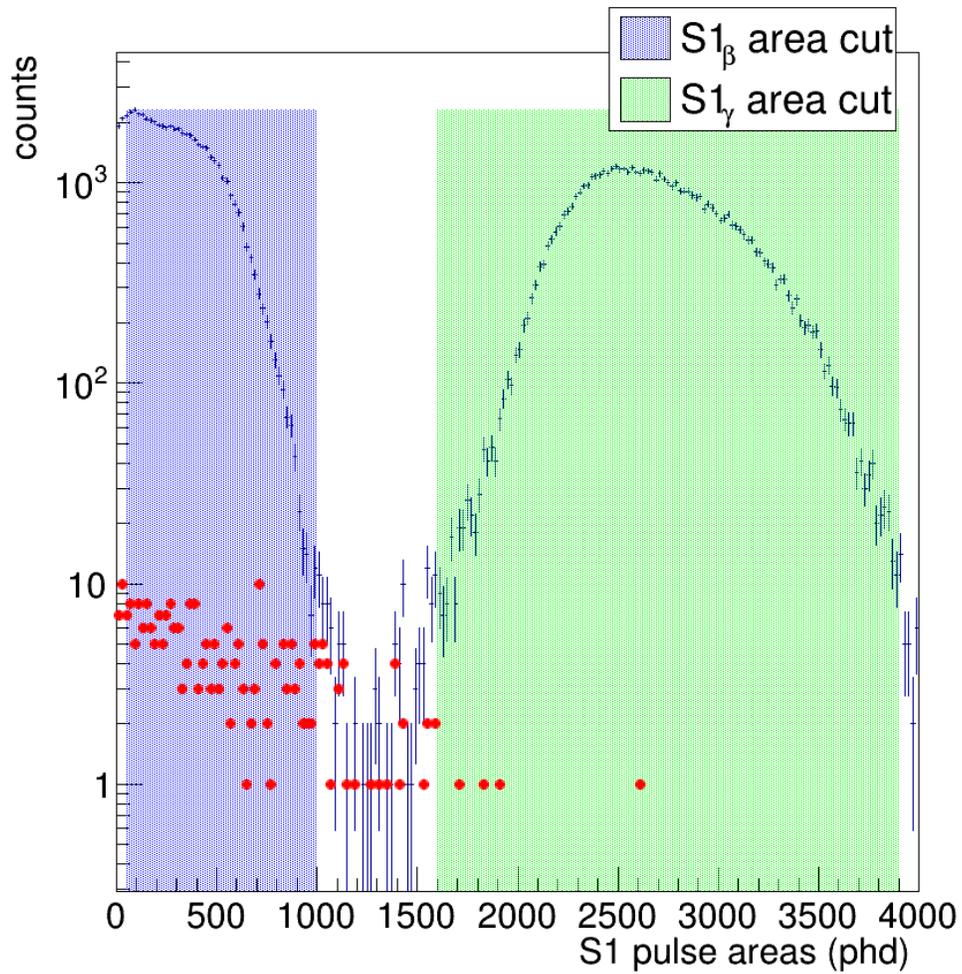


Figure 6.59: A histogram of full simulation $S1$ pulse areas with the $S1_{\beta}$ cut shown in blue and the $S1_{\gamma}$ cut shown in green. These events have been identified using the selection criteria. The red circles highlight events that fail cut9 (skin veto cut).

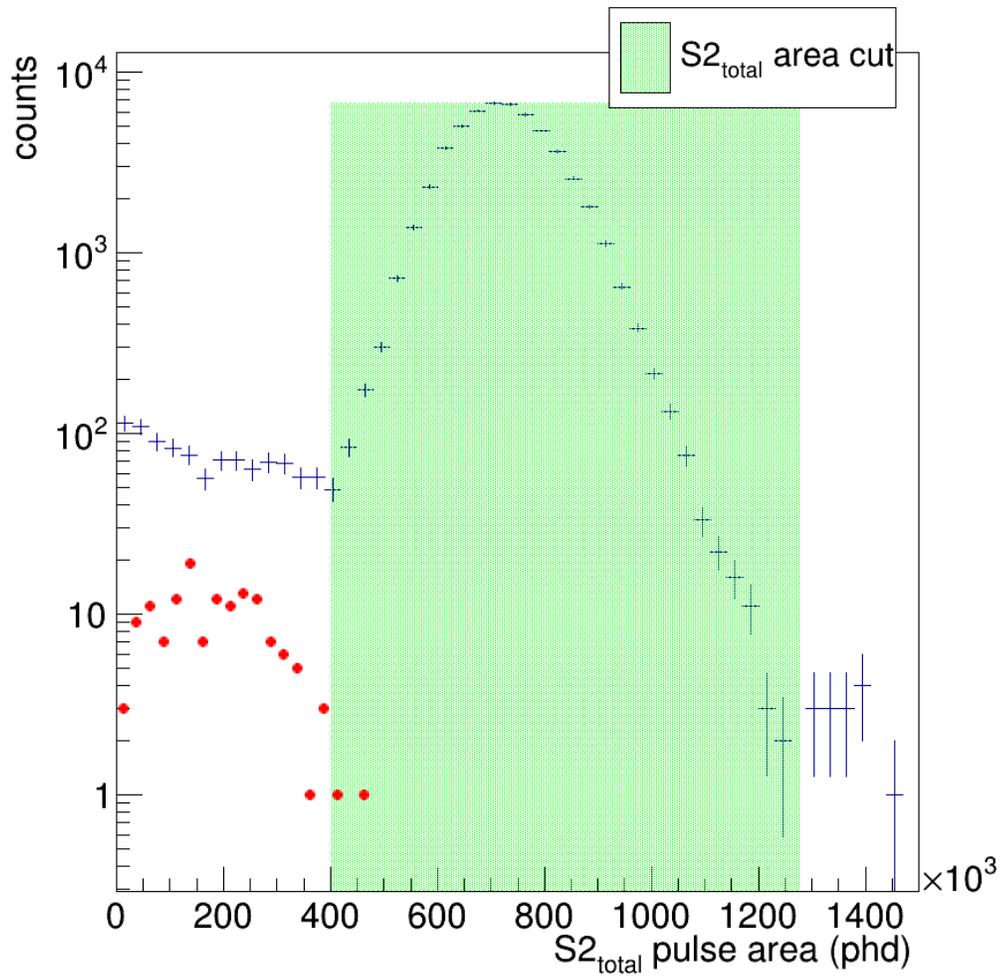


Figure 6.60: A histogram of full simulation $S2_{total}$ pulse areas. These events have been identified using the selection criteria. The red circles highlight events that fail cut9 (skin veto cut).

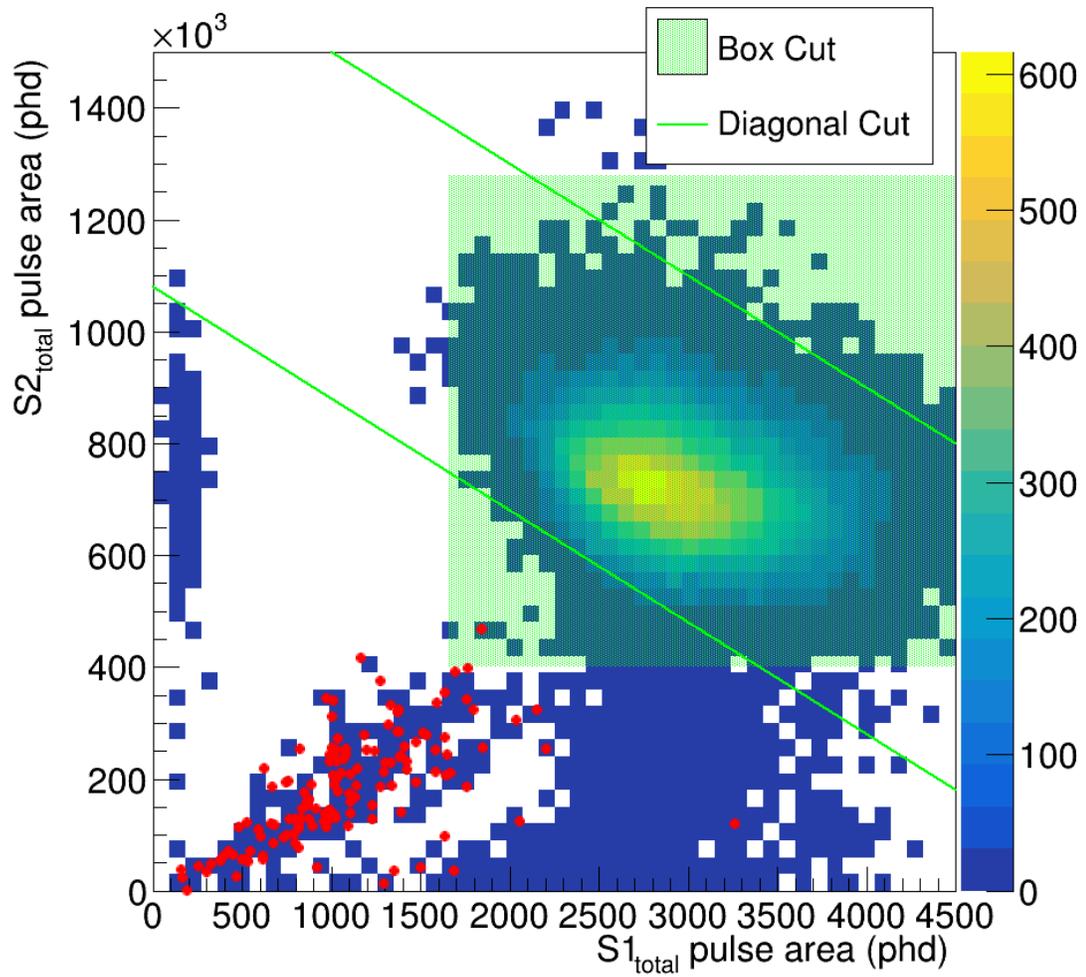


Figure 6.61: A 2-d histogram of full simulation $S1_{total}$ versus $S2_{total}$ pulse areas. These events have been identified using the selection criteria. The red circles highlight events that fail cut9 (skin veto cut).

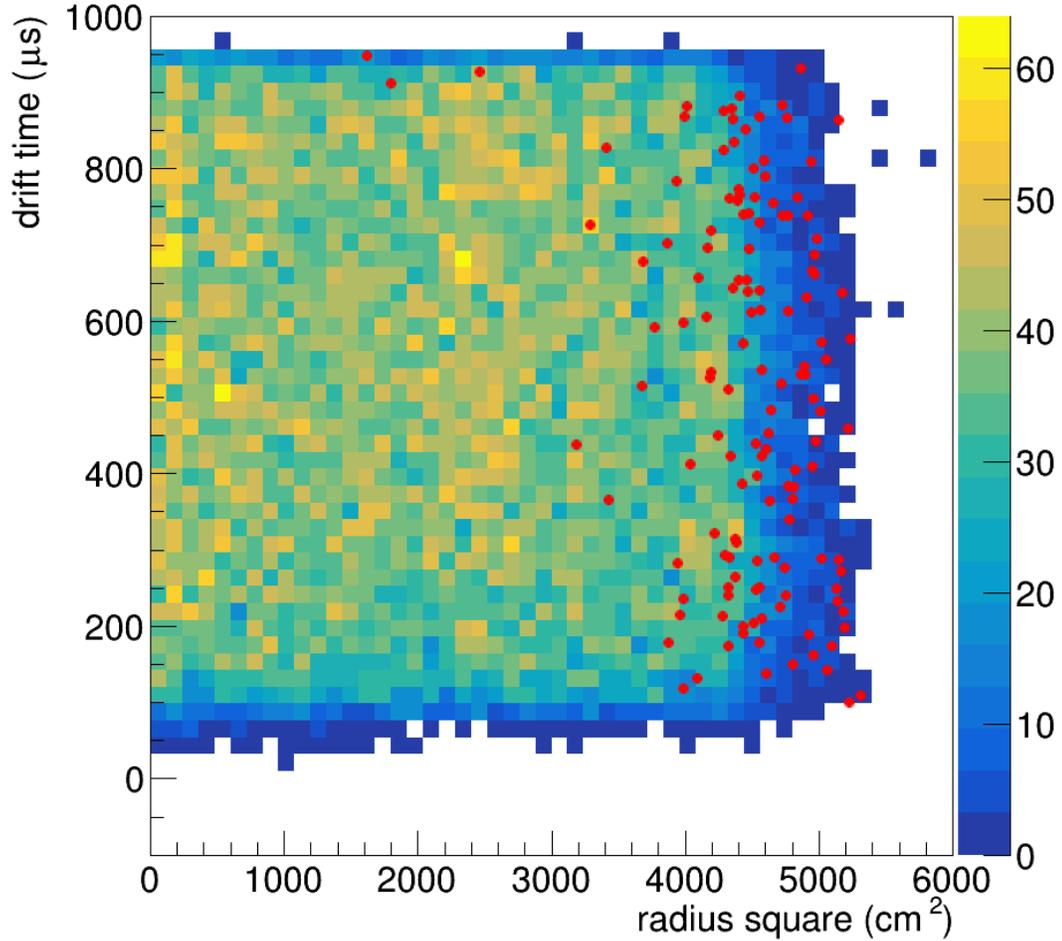


Figure 6.62: A 2-d histogram of full simulation event locations in squared radius (cm^2) versus drift time (μs) pulse areas. These events have been identified using the selection criteria. The red circles highlight events that fail cut9 (skin veto cut).

This marks the final cut informed by the full simulation. The cumulative efficiency of cut8 and cut9 as well as their error in number of events are given in table 6.3. The cumulative efficiency of the full simulation cuts is

$$\frac{N_{9,full}}{N_{7,full}} = \frac{46735}{46929} = 99.6\%. \quad (6.25)$$

To confirm that these events can be readily identified, we use a likelihood fit of surviving events' delay time to equation (6.19). The signal and background terms are again integrated and the values compared to the known number of signal events. This is applied to the results of the fast simulation following cut7, as shown in figure 6.63. For a $0.1 \mu\text{s}$ bin size $c_{scale} = 10$ bins per μs . The full simulation fit finds a scaled signal parameter of $A = 35365 \pm 164$ events per bin and a scaled background parameter of $B = 0 \pm 0.006$ events per bin. After computing the integral and applying c_{scale} , the calculated number of signal events is $N_{signal,full} = 46744 \pm 216$ and background events is $N_{background,fast} = 0 \pm 0.6$. This is in agreement with the known number of surviving signal events $N_{9,full} = 46735$ and zero background events.

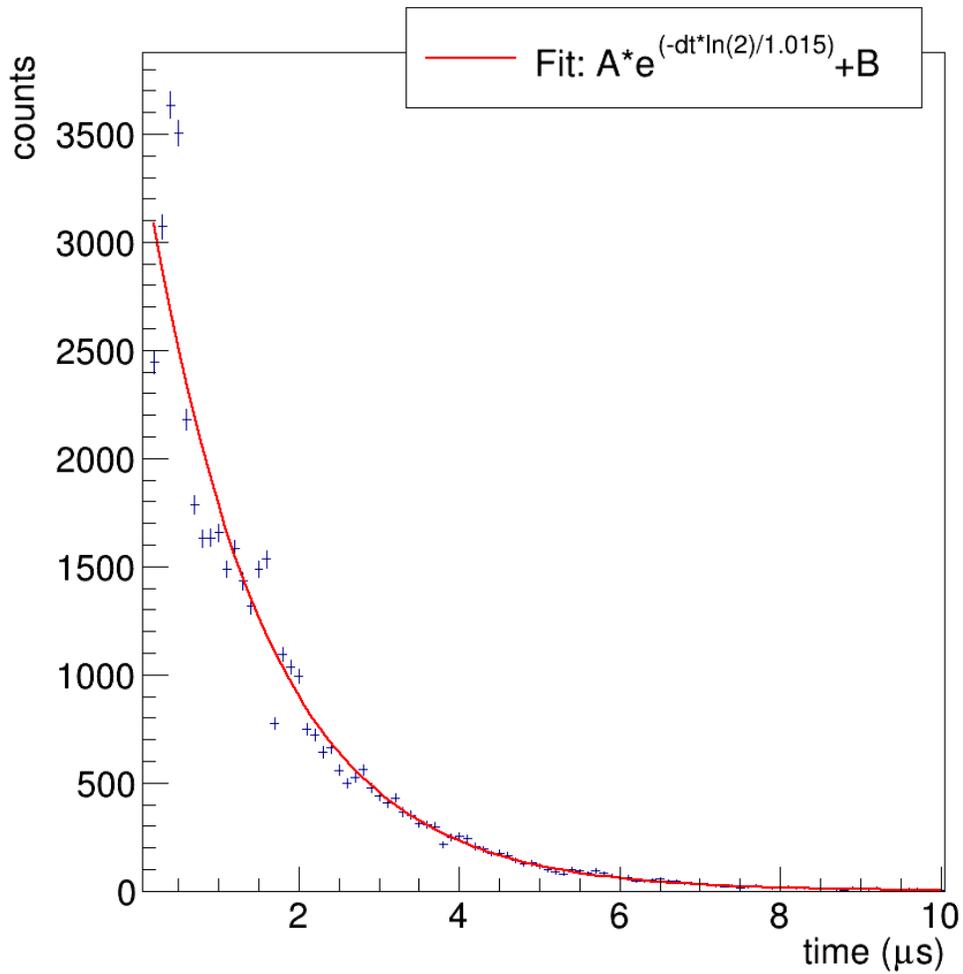


Figure 6.63: A delay time histogram of the 46735 full simulation events that survive cut9. The bin size is $0.1 \mu\text{s}$ and the scale factor is $c_{scale} = 10$ bins per μs . The red line is fit to all events between 0.15 and $100.15 \mu\text{s}$, but the plot is zoomed in to the first $10 \mu\text{s}$. Integration of the scaled signal component accurately accounts for all of the signal events and zero background events.

6.4.9 Signal Efficiency

Combining the cumulative efficiencies from each simulation technique gives a total signal efficiency of

$$\varepsilon_f = (74.1\%)(99.6\%) = 73.8\% \quad (6.26)$$

with the corresponding error

$$\sigma_f = \sqrt{(0.08\%)^2 + (0.52\%)^2 + (0.17\%)^2 + (0.003\%)^2 + (0.46\%)^2} = 0.72\% \approx 1\% \quad (6.27)$$

The signal identification technique which integrates the fit parameters applied to equation (6.19) has been used to calculate the expected number of events to within 0.02% of the known number of signal events for both simulations.

Cut Name	# of Events (fast)	Cumulative Efficiency
Events Generated	49984647	N/A
Excited Decays	214860	N/A
Initial Events	147175±118	N/A
0) Selection Criteria	137793	93.6%
1) Fiducial	131845	89.6%
2) Energy - $S1_{\beta}$ pulse area	116477±603	79.1%
3) Energy - $S1_{\gamma}$ pulse area	116215±197	79.0%
4) Energy - $S2_{total}$ pulse area	115814±4	78.7%
5) Energy - $S1_{total}$ vs $S2_{total}$	114964	78.1%
6) Delay Time Window	109019	74.1%
7) Localization	108995	74.1%
8) ΔTBA	N/A	N/A
9) Skin Veto	N/A	N/A

Table 6.2: Efficiencies from the fast simulation cuts. The fast simulation included ground state beta decays which were filtered using a particle ID truth RQ prior to identifying the initial number of events contained within the fiducial truth volume.

Cut Name	# of Events (full)	Cumulative Efficiency
Generated Excited Decays	99910	N/A
Initial Events	68636	N/A
0) Selection Criteria	55369	N/A
1) Fiducial	53062	N/A
2) Energy - $S1_{\beta}$ pulse area	48188	N/A
3) Energy - $S1_{\gamma}$ pulse area	47859	N/A
4) Energy - $S2_{total}$ pulse area	47223	N/A
5) Energy - $S1_{total}$ vs $S2_{total}$	46953	N/A
6) Delay Time Window	46939	N/A
7) Localization	46929	N/A
8) ΔTBA	46735 ± 216	99.6%
9) Skin Veto	46735	99.6%

Table 6.3: Efficiencies from the full simulation cuts. The full simulation only simulated excited state decays and was checked against a particle ID RQ. The cumulative efficiency up to cut7 is 68.4%. This reasonably close to the cumulative efficiency found from the fast simulation cuts despite a the 5.5% disagreement in the $S2_{total}$ pulse area peak and an 18.3% disagreement in the $S1_{\gamma}$ pulse area peak values.

6.5 SR1 data

Applying cuts one through nine to the 110 million SR1 events leaves 15 surviving events, shown in table 6.4 and figure 6.64. Fitting the remaining events' delay times to equation 6.19 with a bin size of $0.1 \mu\text{s}$ returns a scaled signal parameter of $A = 6.12 \pm 2.24$ and a background parameter of $B = 0.069 \pm 0.028$. After integrating over the entire delay time window, the calculated number of signal signal events is $N_{\text{signal},\text{SR1}} = 8.1 \pm 3.0$. The number of background events is $N_{\text{background},\text{SR1}} = 6.9 \pm 2.8$. Finally we apply the efficiency correction to determine the total number of excited state ^{85}Kr decays over the course of SR1

$$N_{\text{final}} = \frac{8.1 \pm 3.0}{0.738 \pm 0.01} = 11.0 \pm 4.0 \text{ events.} \quad (6.28)$$

The equivalent concentration of $^{\text{nat}}\text{Kr}$ to $^{\text{nat}}\text{Xe}$ is found by applying equation (6.3) and the appropriate branching fraction. The resulting concentration is

$$\Phi_{\text{Kr}} = \frac{N_{\text{final}}}{100.3 \text{ days}} \frac{10.856 \text{ years}}{\ln(2)} \frac{M_{85\text{Kr}}}{N_A} \frac{m_{\text{natKr}}}{m_{85\text{Kr}}} \frac{1}{0.434\% \cdot m_{\text{Xe}}} \quad (6.29)$$

$$\Phi_{\text{SR1}} = 183 \pm 67 \text{ ppq} \quad (6.30)$$

and is in agreement with current in situ XPM measurements, which record a $^{\text{nat}}\text{Kr}$ concentration of $144 \pm 22 \text{ ppq}$.

Cut Name	# of Events (SR1)
Initial Events	110393744
0) Selection Criteria	147848
1) Fiducial	40603
2) Energy - $S1_{\beta}$ pulse area	11437
3) Energy - $S1_{\gamma}$ pulse area	3404
4) Energy - $S2_{total}$ pulse area	2488
5) Energy - $S1_{total}$ vs $S2_{total}$	1704
6) Delay Time Window	93
7) Localization	55
8) ΔTBA	35
9) Skin Veto	15

Table 6.4: Data cuts

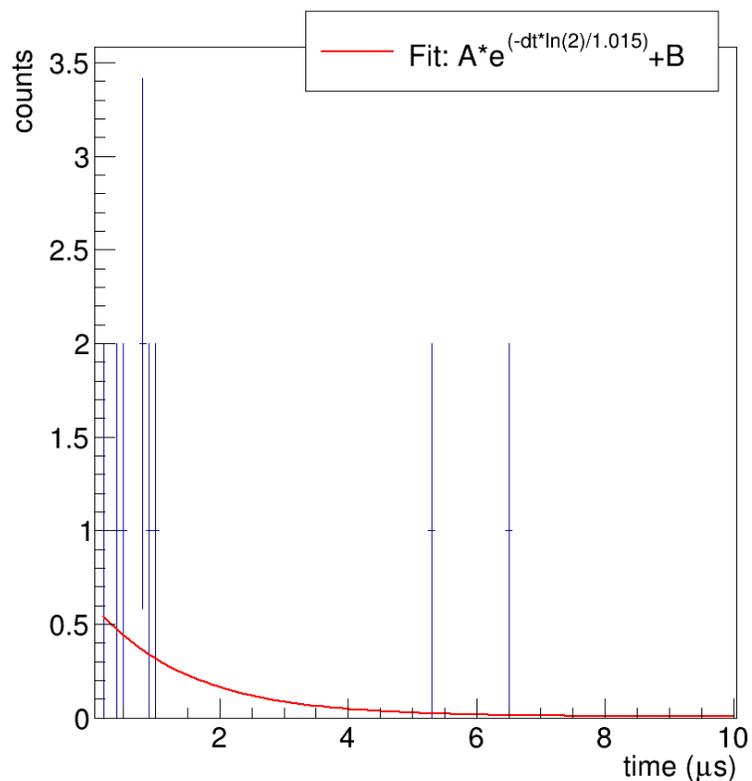


Figure 6.64: A histogram of the delay time for the remainin 15 SR1 events following cut9. The bin size is $0.1 \mu\text{s}$ and the scale factor is $c_{scale} = 10$ bins per μs . The red line is fit to all events between 0.15 and $100.15 \mu\text{s}$, but the plot is zoomed in to the first $10 \mu\text{s}$. The scaled and integrated fit parameters identify 8.1 ± 3.0 signal events and 6.9 ± 2.8 background events over the course of SR1 before taking efficiency into account.

6.6 Side-band Analysis

To further validate the analysis method, a selection of data known to lack any ^{85}Kr events is considered. In section 6.4.8, the veto cut was shown to have 100% relative efficiency because the pulse area cuts remove events that would have otherwise triggered a skin veto. This means that all 20 events removed from SR1 data by cut9 should be exclusively background events. These 20 events satisfy every cut except for the skin veto and make an

excellent choice for a side-band analysis. The vetoed side-band events are overlaid in red on the final 15 events shown below.

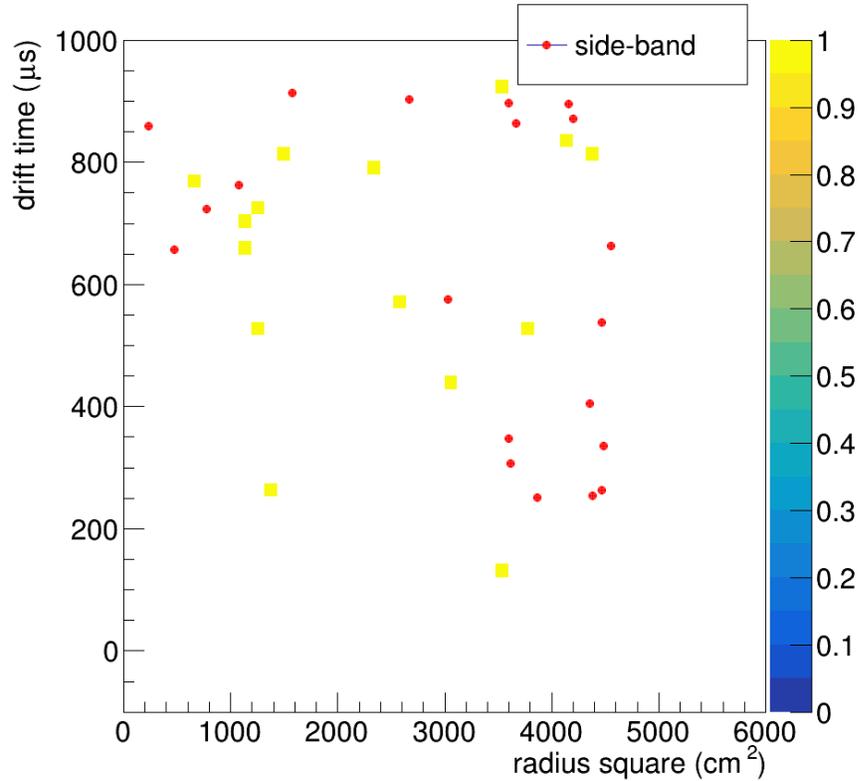


Figure 6.65: A histogram of the radial positions and drift times of the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.

Figures 6.65 and 6.66 show both the side-band events and final SR1 tend slightly towards the bottom of the detector. This is not reflected in simulations that have uniform distributions in drift time. Of the final SR1 events, those with delay times greater than 5 μs tend to be more heavily clustered near the bottom of the detector when compared to events with smaller delay times. These events are likely background events, characterized in our final likelihood fit by the constant term.

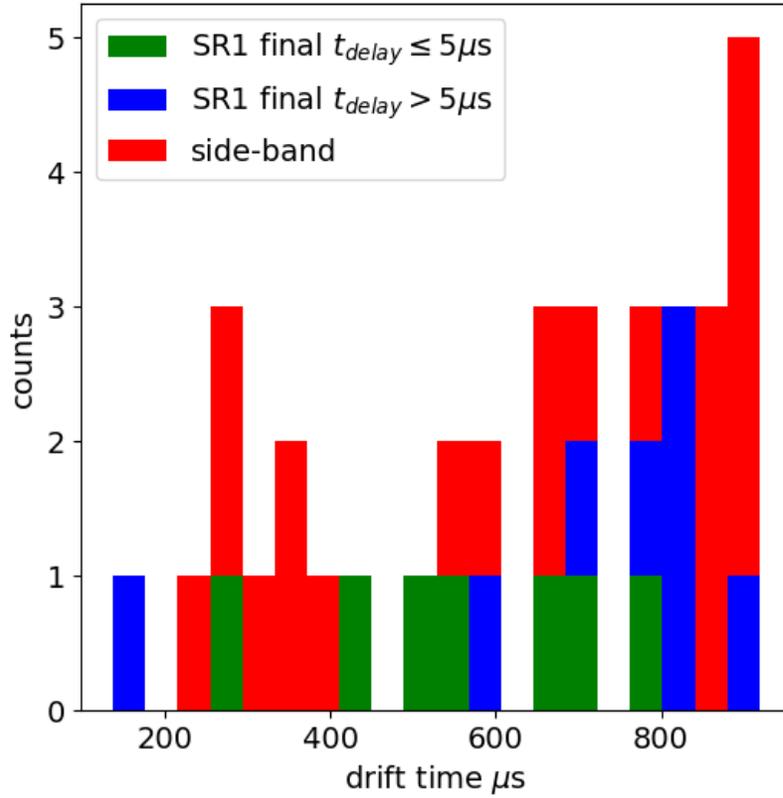


Figure 6.66: A stacked histogram of the drift times of the remaining 15 SR1 events following all nine cuts and side-band events. The final 15 SR1 events are shown in green for delay times $\leq 5 \mu\text{s}$ and blue for all other surviving delay times. The red corresponds to the 20 side-band events that failed only the skin veto cut. Events in both the side-band and final SR1 data appear to be cluster near the bottom of the detector.

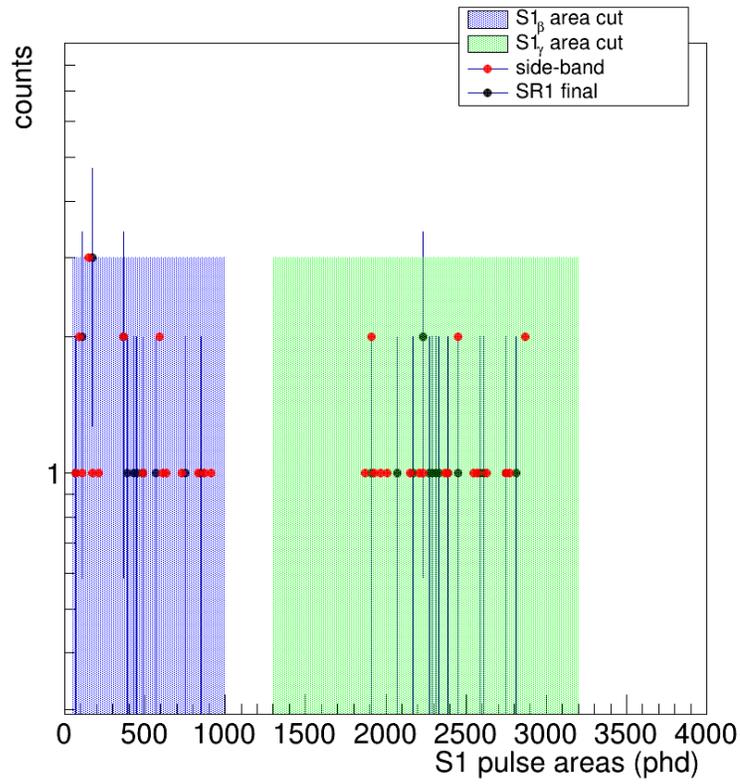


Figure 6.67: A histogram of $S1$ pulse areas for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.

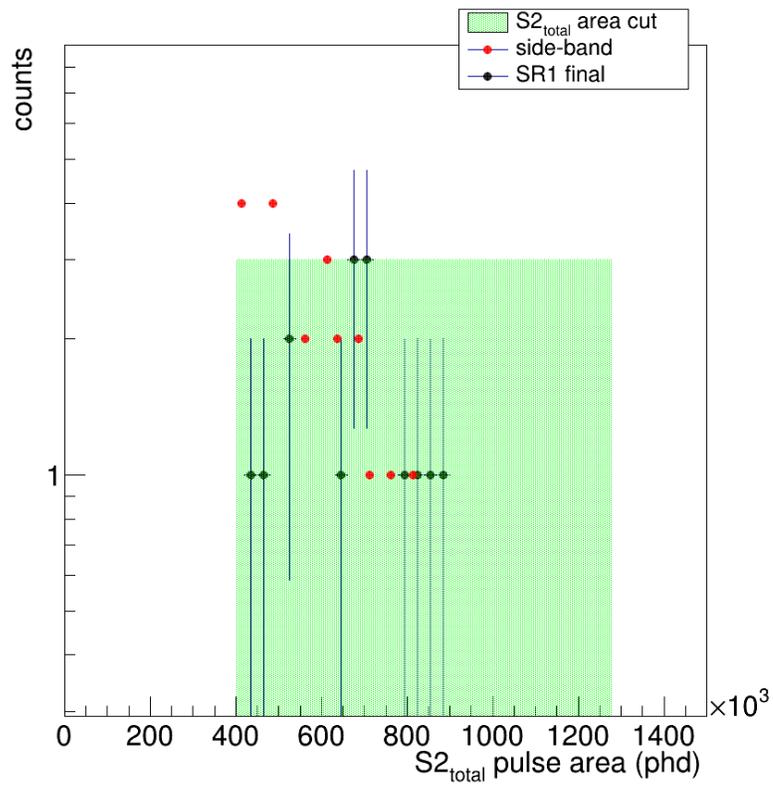


Figure 6.68: A histogram of $S_{2_{total}}$ pulse areas for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.

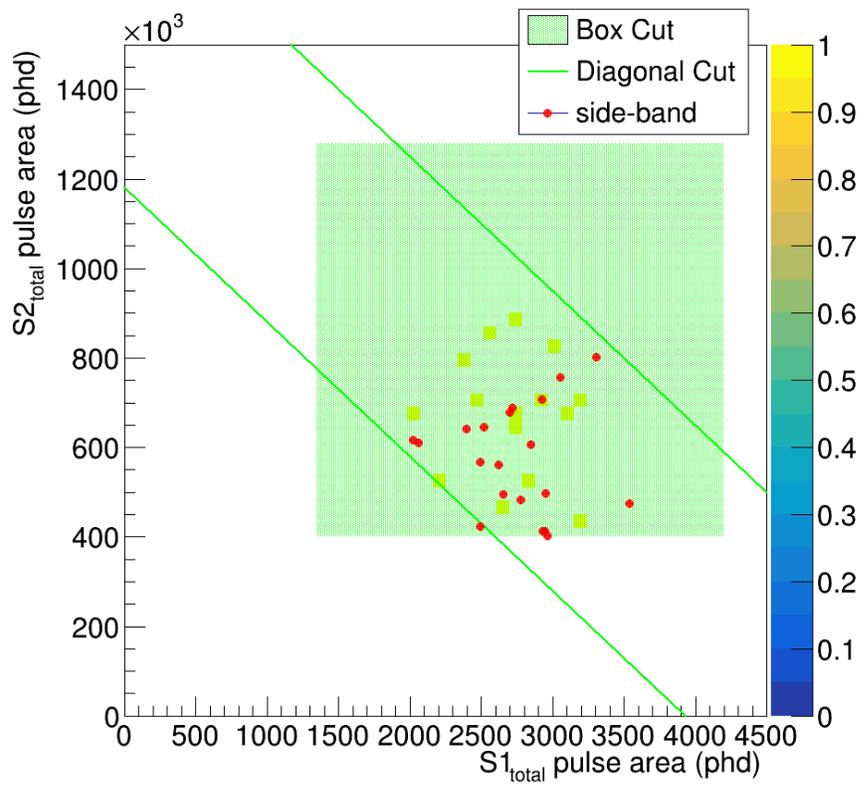


Figure 6.69: A histogram of $S1_{total}$ versus $S2_{total}$ pulse areas for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.

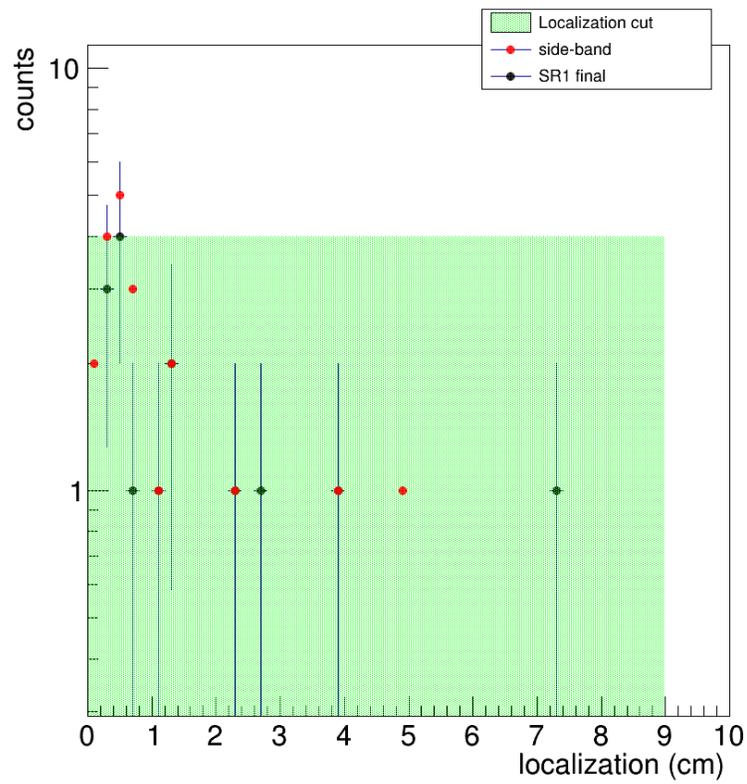


Figure 6.70: A histogram of localization for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.

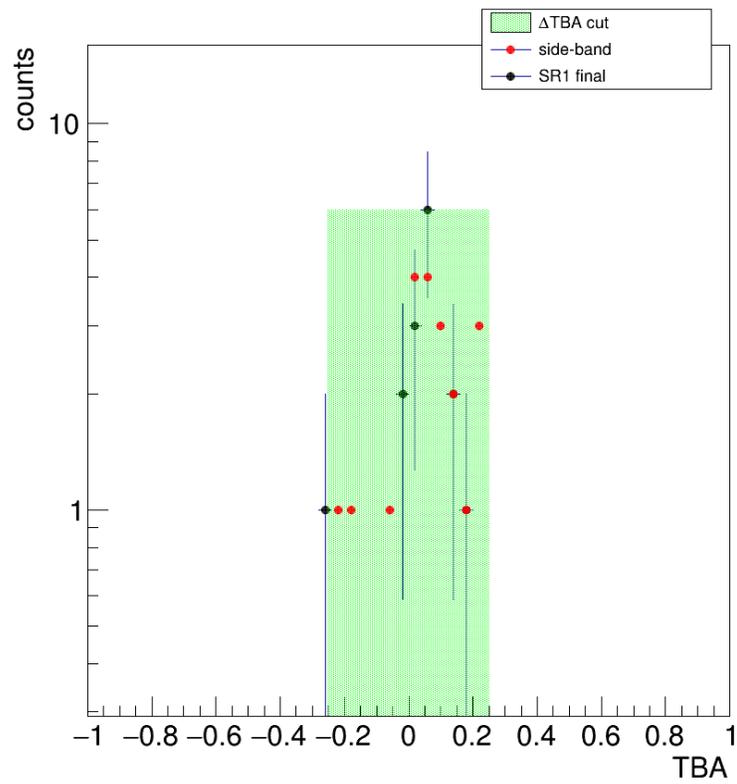


Figure 6.71: A histogram of TBA and ΔTBA cut for the remaining 15 SR1 events following all nine cuts. The red points correspond to the 20 side-band events that failed only the skin veto cut.

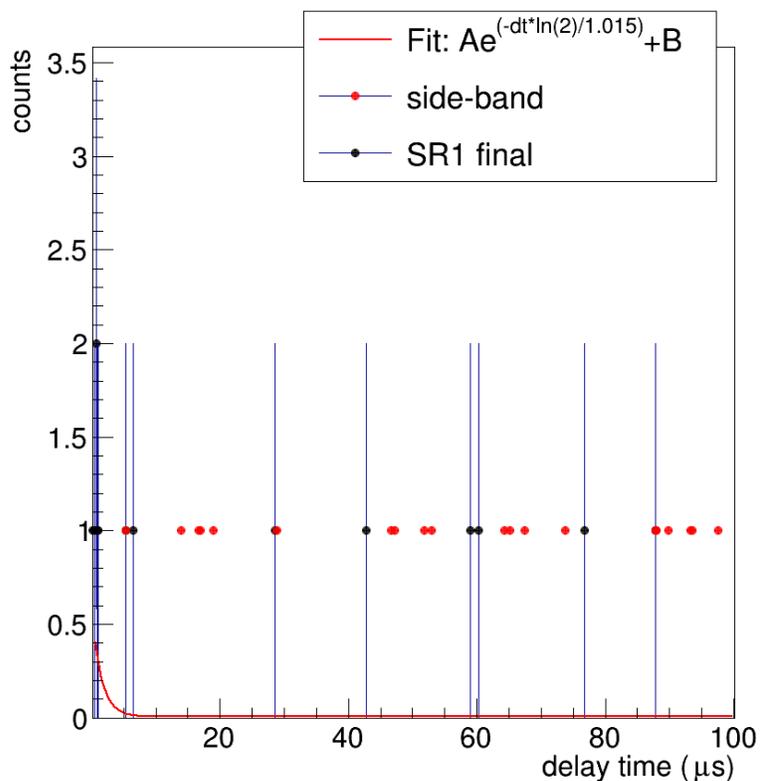


Figure 6.72: A histogram of delay time for the remaining 15 SR1 events following all nine cuts. The red fit is to the final 15 SR1 events. The red circles correspond to the 20 events in the side-band events and failed only the skin veto cut. The side-band events appear to be consistent with a flat background.

Applying the same likelihood fit of equation (6.19) with a bin size of $0.1 \mu\text{s}$ to the 20 side-band events, we find a scaled signal parameter $A_{side} = 0 \pm 0.4$ events per μs and a background parameter $B_{side} = 0.2 \pm 0.045$ events per μs . Taking efficiency and integrating over the background window, the side-band signal event count is 0.0 ± 0.8 events. The integrated background event count is 20 ± 4.5 events. There is no signal observed in the side-band analysis and all events are identified as backgrounds. Both the events from the side-band and the final 15 events have no obvious time structure throughout SR1, as shown in figure 6.73.

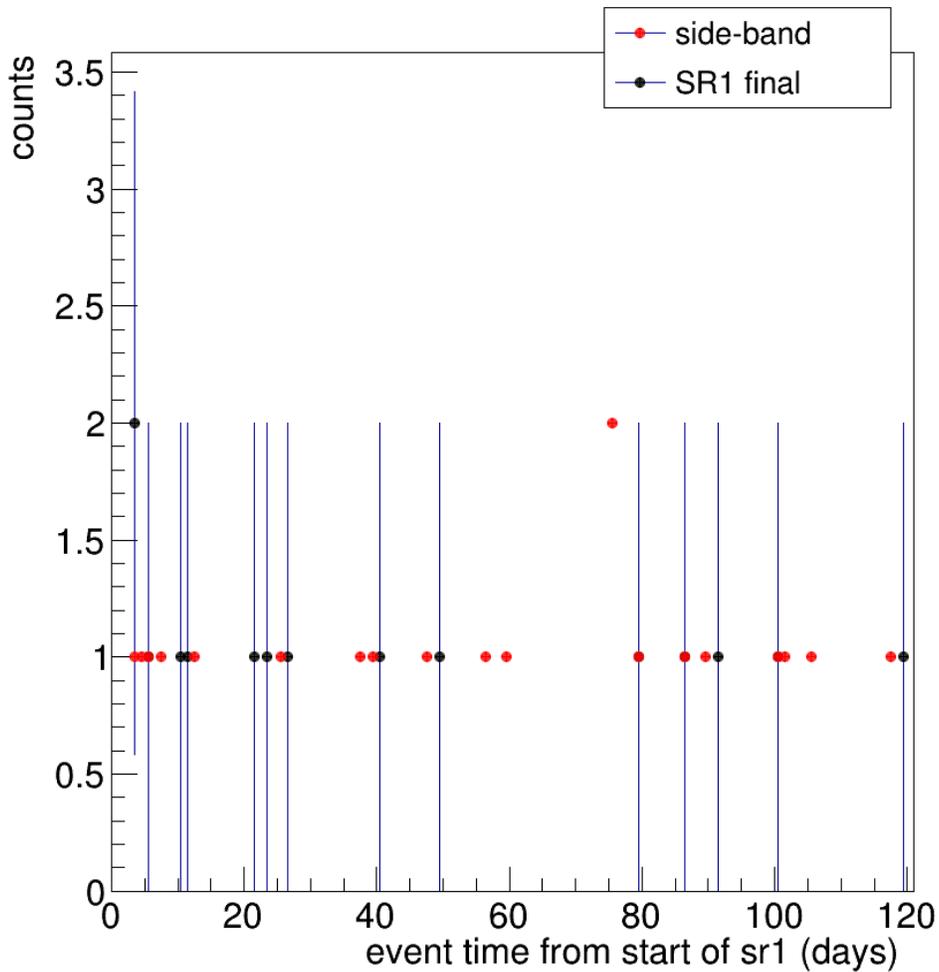


Figure 6.73: A histogram of events over the course of the 120 SR1 calendar days. The bin size is one day. Neither the 15 surviving events, nor the 20 side-band events appear to have any structure throughout SR1.

6.7 Toy Monte Carlo Model

The likelihood fit is further validated with a toy Monte Carlo study that generates mock delay times to imitate a similar signal search. We start by determining how many signal delay times, M_{sig} , to generate. M_{sig} is found by drawing from a Poisson distribution centered

at $N_{signal,SR1} = 8.1$ events, as shown in figure 6.74.

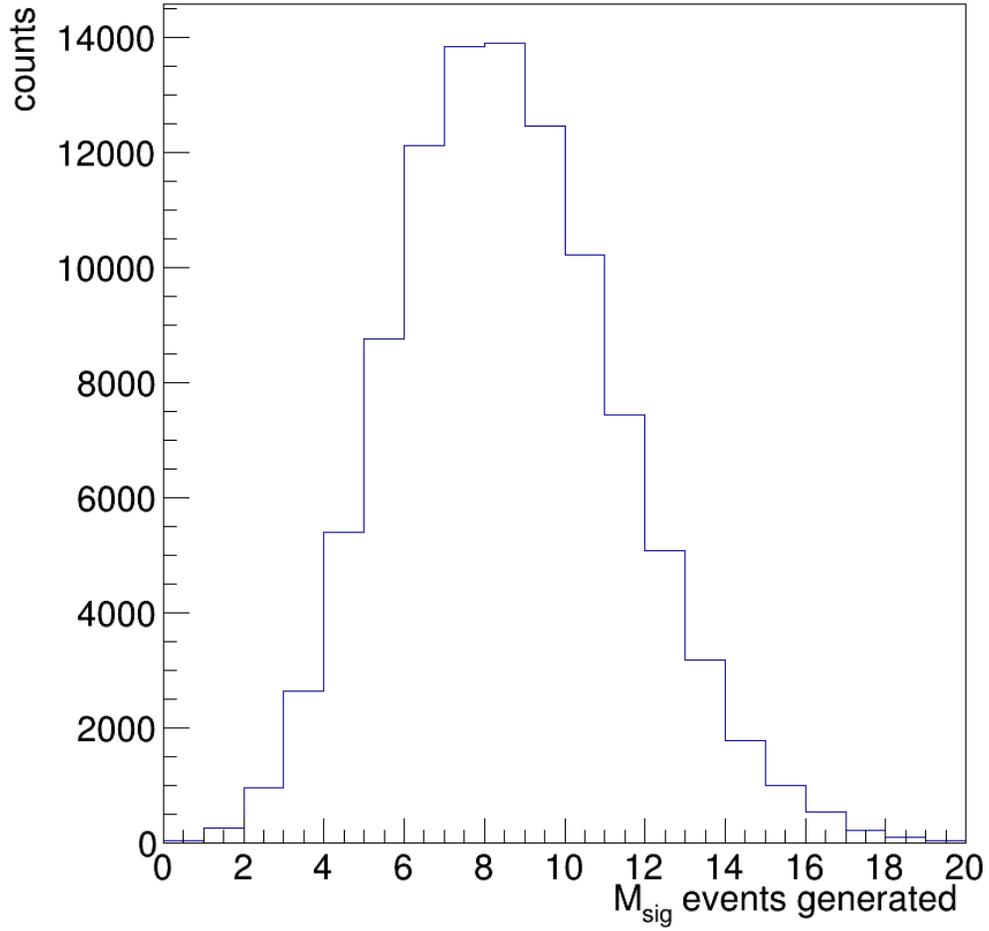


Figure 6.74: A histogram 100,000 M_{sig} values taken from a Poisson distribution centered at 8.1.

M_{sig} mock signal delay time values between 0.15 and 100.15 μs are then drawn from the exponential distribution

$$f(t) = e^{-t \ln(2)/1.015} \quad (6.31)$$

Next the number of mock background delay times, M_{bkg} , to generate is selected. M_{bkg}

is drawn from a Poisson distribution centered at $N_{background,SR1} = 6.9$ events, as shown in figure 6.75.

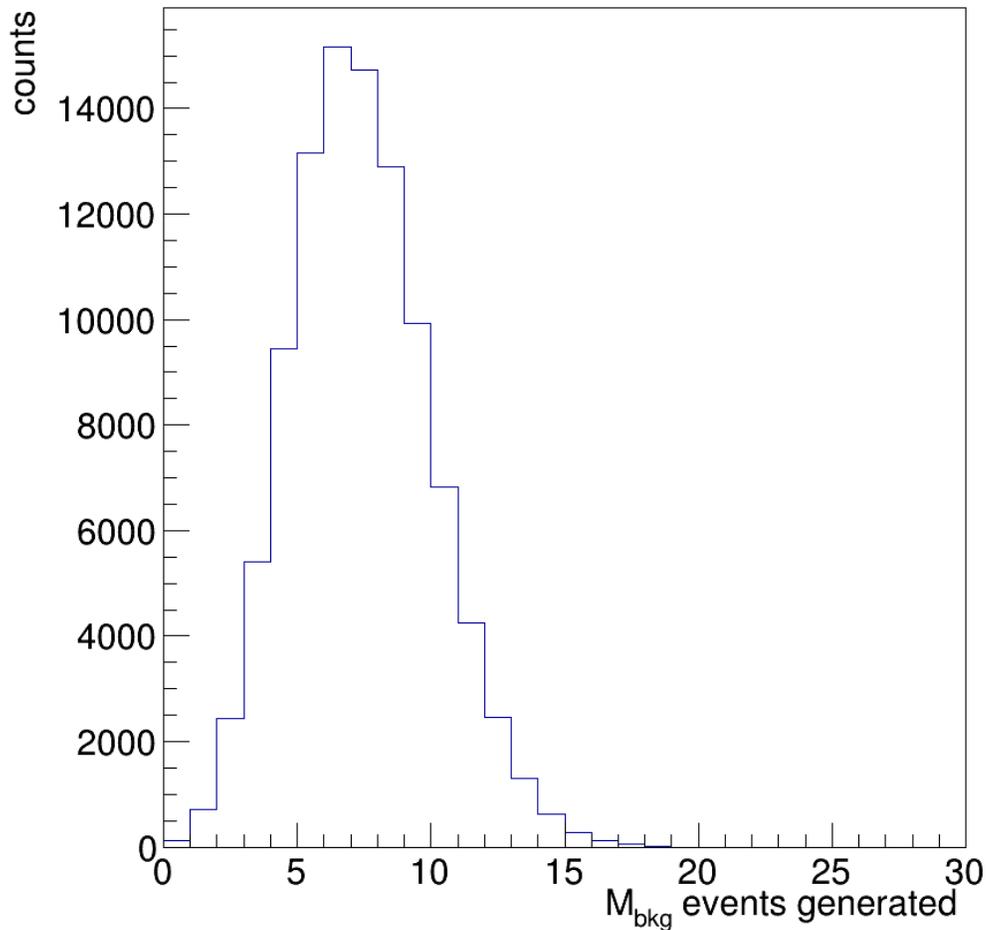


Figure 6.75: A histogram 100,000 M_{bkg} values taken from a Poisson distribution centered at 6.9.

Then M_{bkg} mock background delay time values between 0.15 and 100.15 are then drawn from a flat distribution. The signal and background mock delay times are combined and the values are fit using the same decaying exponential offset function with a $0.1 \mu s$ bin size. One toy histogram are shown in figure 6.76. The cumulative values of all 100,000 toy fits

are shown in figure 6.77.

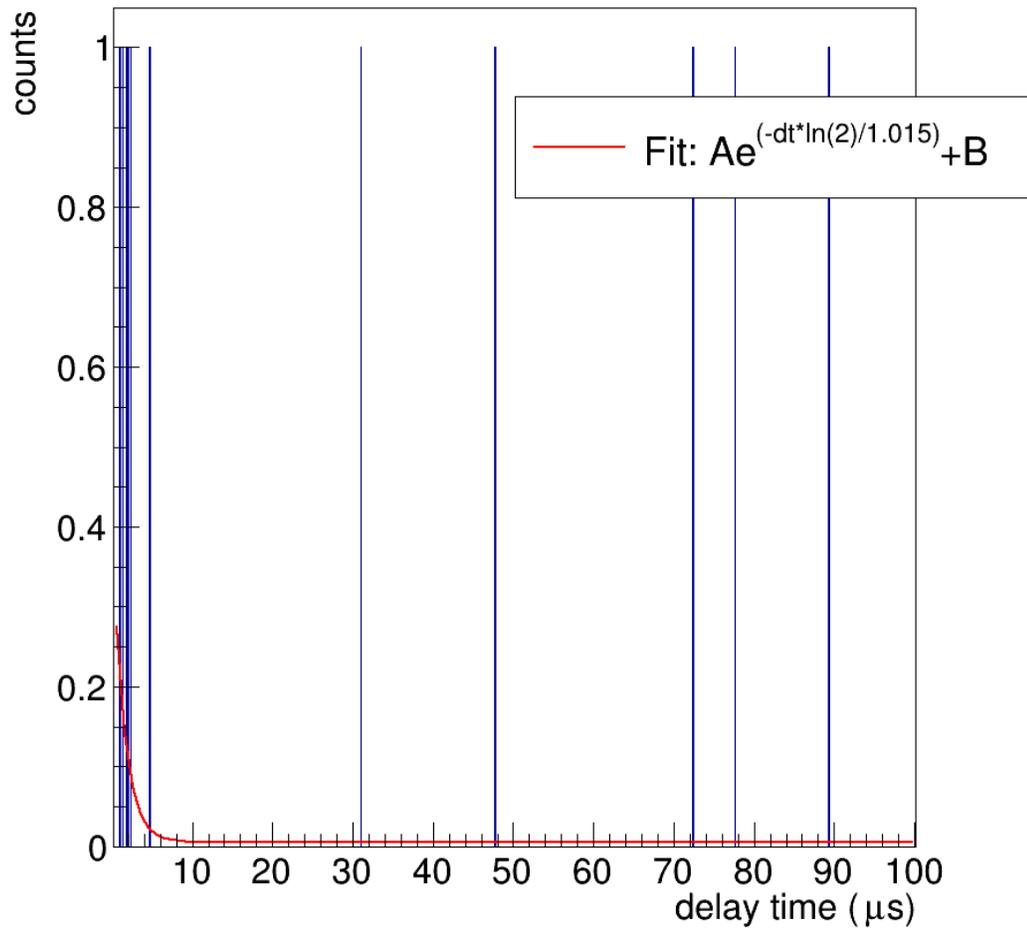


Figure 6.76: A histogram of toy delay time values from one iteration containing 5 signal and 6 background delay times. The number of events were generated using the M_{sig} and M_{bkg} Poisson distributions. The delay time values were generated using $e^{-t \ln(2)/1.015}$ and a flat distribution.

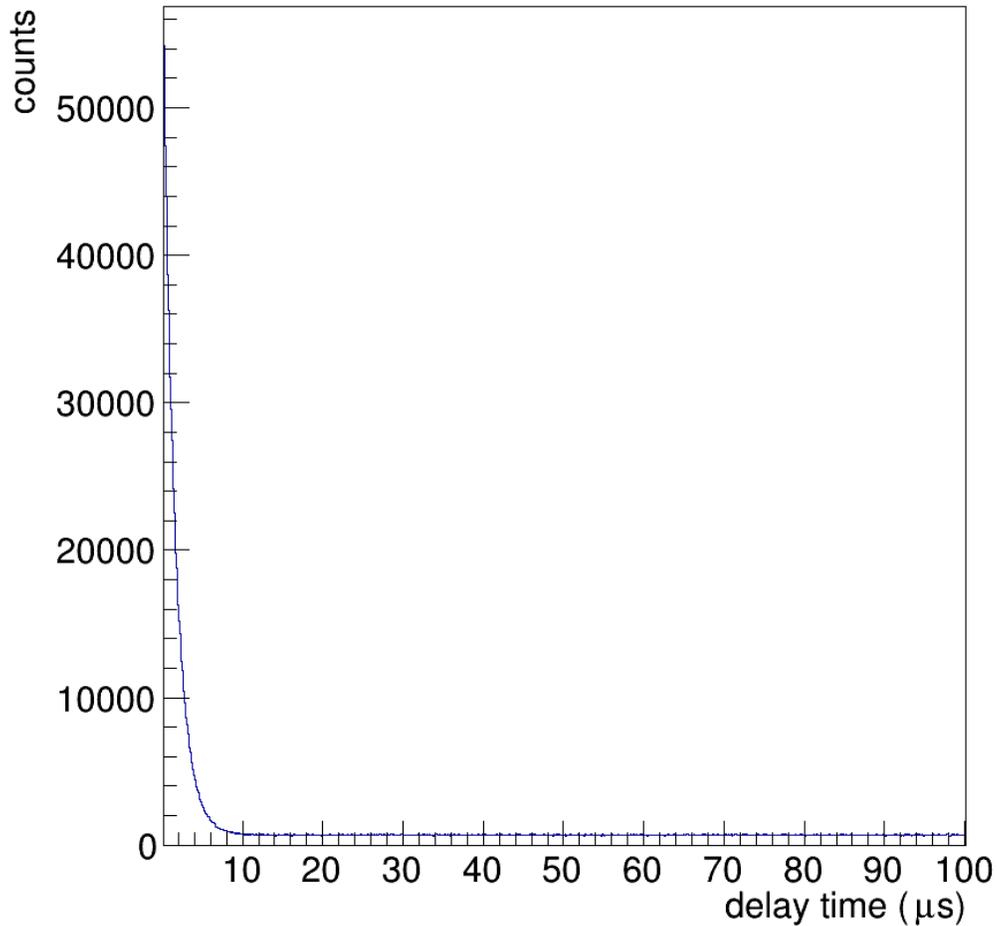


Figure 6.77: A histogram of all toy delay time values generated from 100,000 toy fits. The number of events were generated using the M_{sig} and M_{bkg} Poisson distributions. The delay time values were generated using $e^{-t \ln(2)/1.015}$ and a flat distribution.

The signal and background fit parameters (A_{toy} and B_{toy}) and their errors are determined using the same bin size used in the SR1 fit. The reported values have been scaled using $c_{scale} = 10 \mu$ per bin. 100,000 toy data sets are generated and fit. The scaled SR1 fit results find $A_{SR1} = 6.12 \pm 2.24$ events per μs and $B_{SR1} = 0.069 \pm 0.028$ events per μs that are consistent with the distribution of toy parameters and errors found. Equations (6.32) and

(6.33) are the reported mean values of the fit parameters and the mean values of the fit errors found from the data displayed in figures 6.78 and 6.79.

$$A_{toy} = 6.13 \pm 2.18 \quad (6.32)$$

Where the standard deviations of the reported A_{toy} values and its errors are 2.21 and 0.39 respectively.

$$B_{toy} = 0.069 \pm 0.027 \quad (6.33)$$

Where the standard deviations of the reported B_{toy} values and its errors are 0.027 and 0.006 respectively.

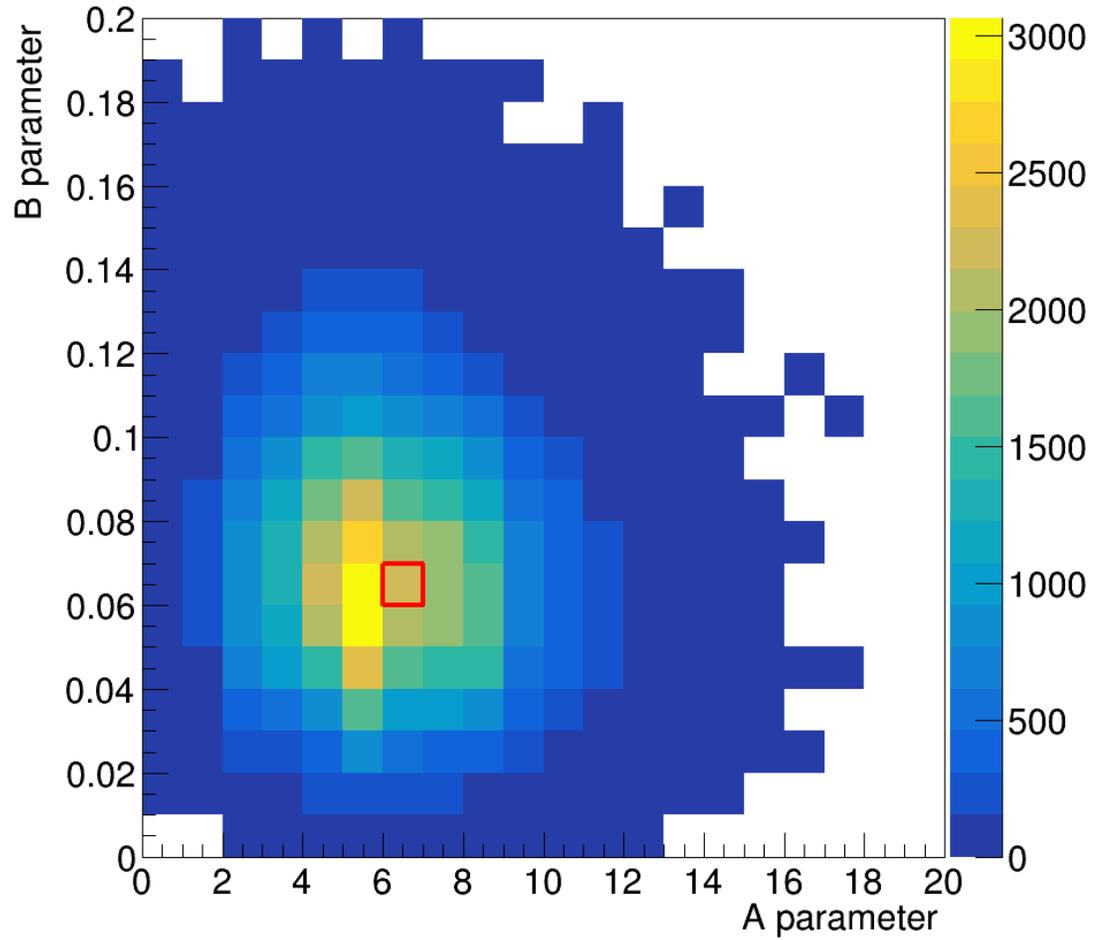


Figure 6.78: A 2-d histogram of A and B fit parameters from 100,000 toy fits. The red box highlights the bin that contains the fit values found in the SR1 data search. The mean and standard deviations of the fit parameters are $A_{toy} = 6.13$, $\sigma_{A_{toy}} = 2.21$ and $B_{toy} = 0.069$, $\sigma_{B_{toy}} = 0.027$.

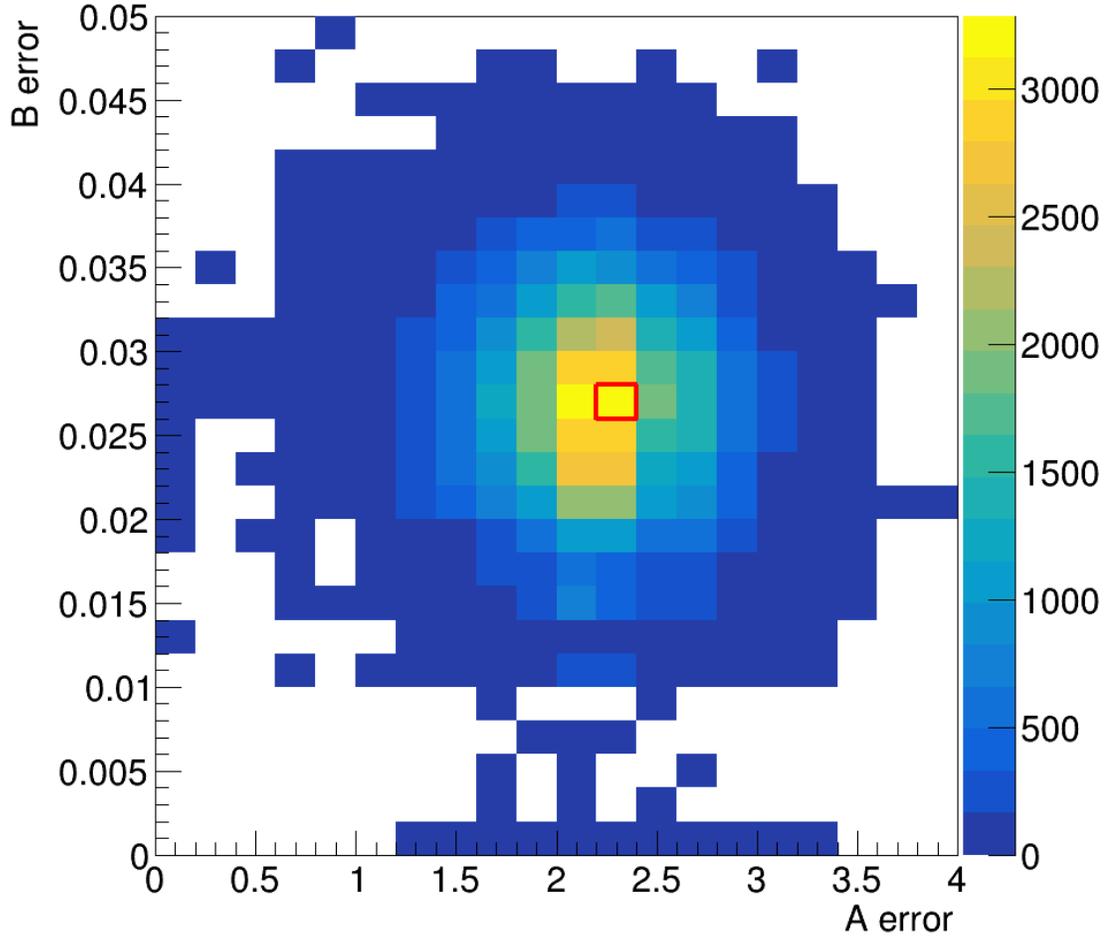


Figure 6.79: A 2-d histogram of $A_{MC,err}$ and $B_{MC,err}$ fit parameter errors from 100,000 toy fits. The red box highlights the bin that contains the fit error values found in the SR1 data search. The mean and standard deviations of the fit parameter errors are $A_{MC,err} = 2.18$, $\sigma_{A_{MC,err}} = 0.39$ and $B_{MC,err} = 0.027$, $\sigma_{B_{MC,err}} = 0.006$.

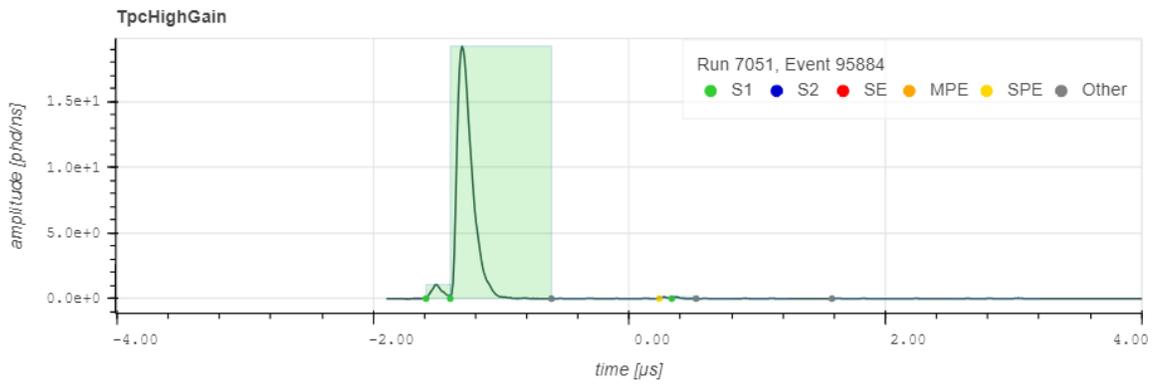
6.8 Final Event Hand Scanning

Of the 15 final events, seven have a delay time less than two μs . All 15 events have waveforms and PMT hit patterns consistent with the simulated events shown previously, but

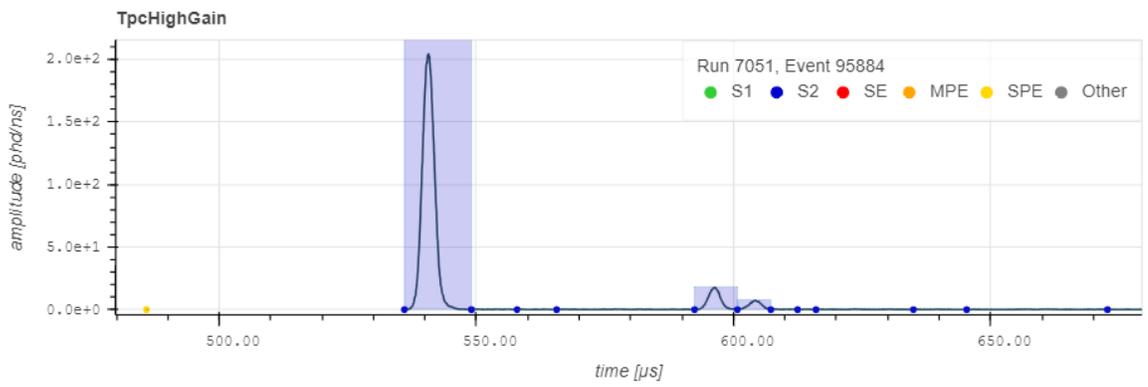
with varying delay times. Table 6.5 contains the final values of interest for the surviving events. Some figures from the event viewer show three events with the shortest delay time followed by the three events with the longest delay time below.

Table 6.5: Table of surviving SR1 events following all cuts.

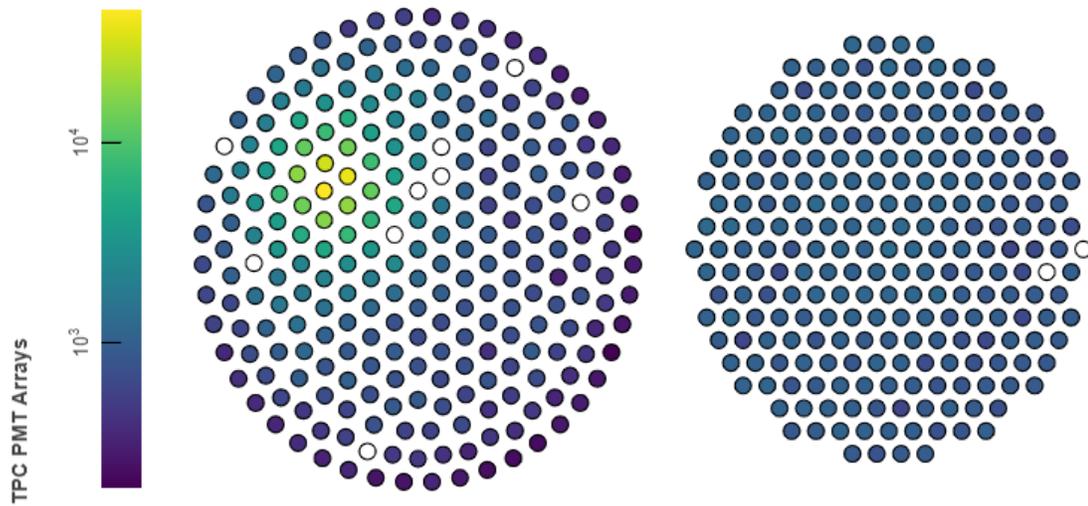
Run#	Event#	t_{delay} (μ s)	$S1_{\beta}$ (phd)	$S1_{\gamma}$ (phd)	$S2_{total}$ (phd)	ΔT_{BA}	t_{drift} (μ s)	x (cm)	y (cm)	local (cm)
7051	95884	0.19	104.7	2617.5	6.84×10^5	0.145	537.4	20.2	29.9	1.3
6986	87111	0.38	420.9	2814.3	7.10×10^5	0.055	717.4	33.7	9.3	2.7
7664	9834	0.50	163.2	2320.0	6.91×10^5	-0.038	429.1	-20.4	-51.2	1.1
8073	119504	0.77	105.3	2309.9	7.80×10^5	0.017	782.8	18.5	-44.4	0.5
6771	5135	0.78	174.0	2740.4	7.05×10^5	0.049	655.8	-29.0	-17.4	0.5
7746	90088	0.87	378.6	2396.3	6.56×10^5	-0.016	518.9	44.9	41.6	0.2
7879	70076	1.02	365.5	2238.1	8.51×10^5	0.158	265.5	-27.3	-26.3	0.3
6844	24790	5.3	483.4	2272.0	8.77×10^5	0.027	696.5	-5.3	-33.8	1.3
7337	93550	6.5	753.9	2442.5	4.22×10^5	0.061	883.5	-17.7	61.4	0.2
6856	11897	28.5	848.4	2223.1	6.69×10^5	0.060	577.1	29.4	41.7	7.3
7230	94959	42.8	76.0	1914.5	6.73×10^5	0.045	137.0	39.8	-44.3	2.3
7021	57314	58.9	452.3	3178.1	4.16×10^5	0.031	918.2	-29.2	51.8	0.7
6773	108440	60.3	388.2	2592.3	8.37×10^5	-0.247	774.2	25.7	1.7	4.0
6785	43340	76.8	570.6	2288.2	5.22×10^5	0.059	806.8	-35.6	-13.5	0.5
7790	13919	87.8	174.6	2063.8	5.29×10^5	0.187	817.3	-50.4	-42.6	0.5



(a) The $S1_{\beta}$ and $S1_{\gamma}$ pulses.

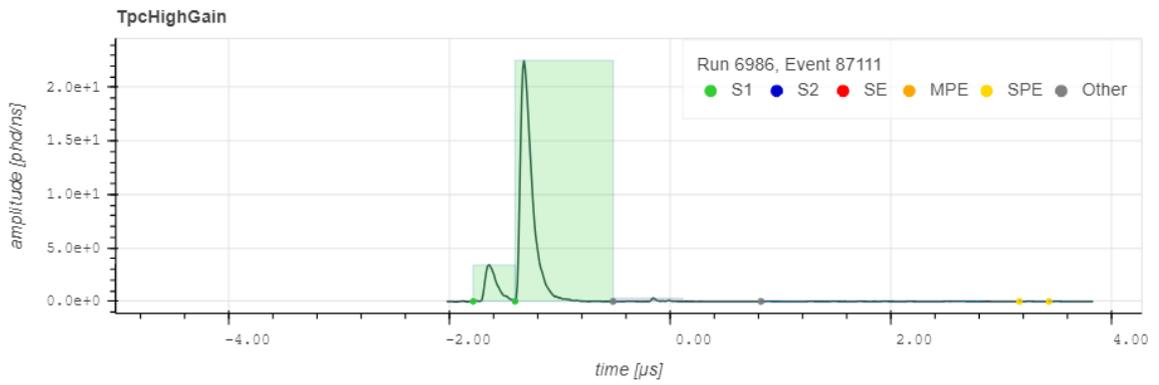


(b) The $S2$ pulses observed.

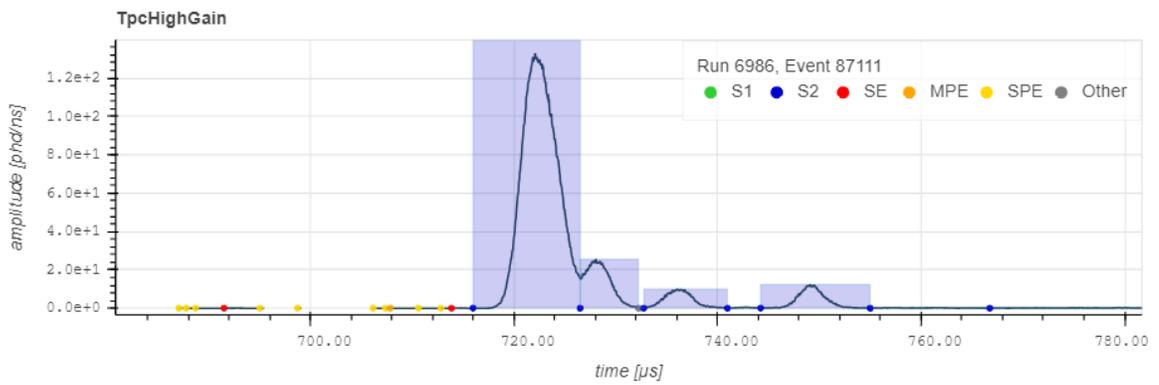


(c) The PMT hit map.

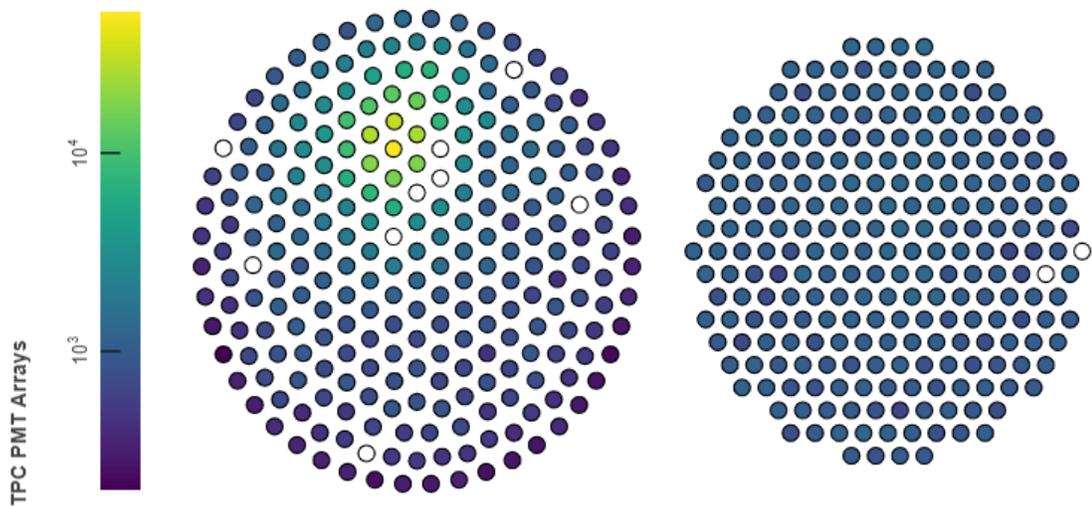
Figure 6.80: Final event 7051 from run 95884. The delay time is $0.19 \mu\text{s}$.



(a) The $S1_{\beta}$ and $S1_{\gamma}$ pulses.

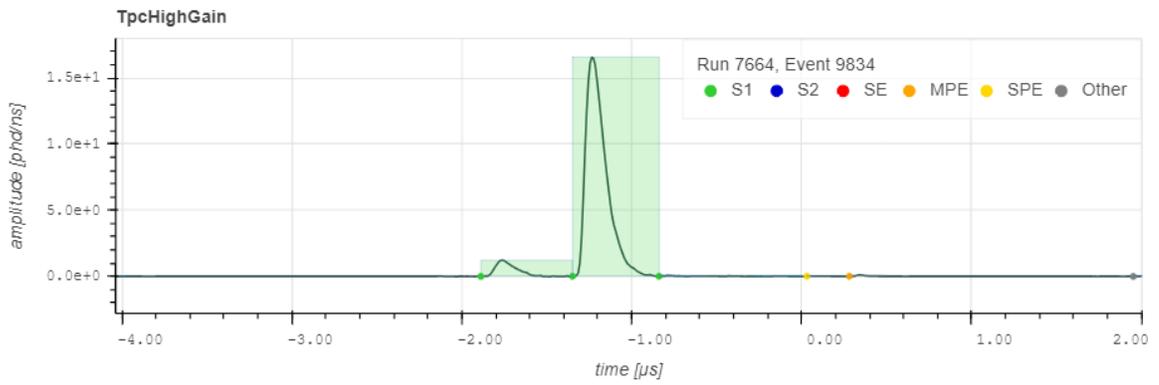


(b) The $S2$ pulses observed.

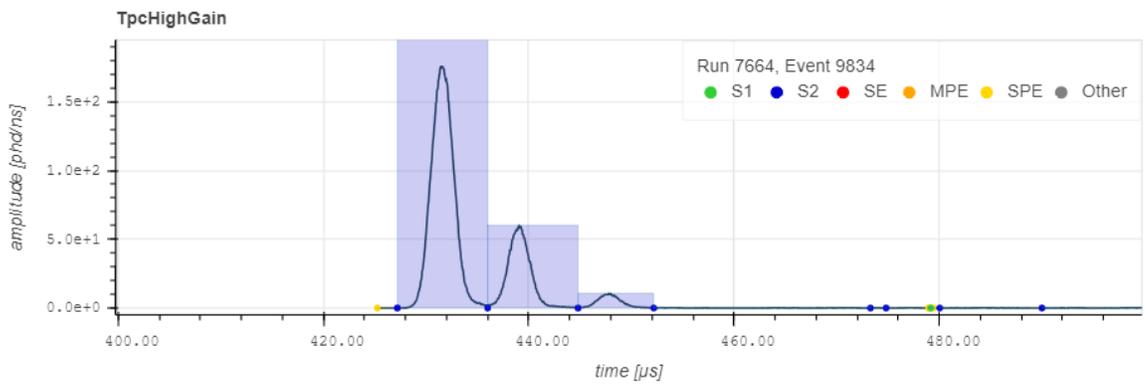


(c) The PMT hit map.

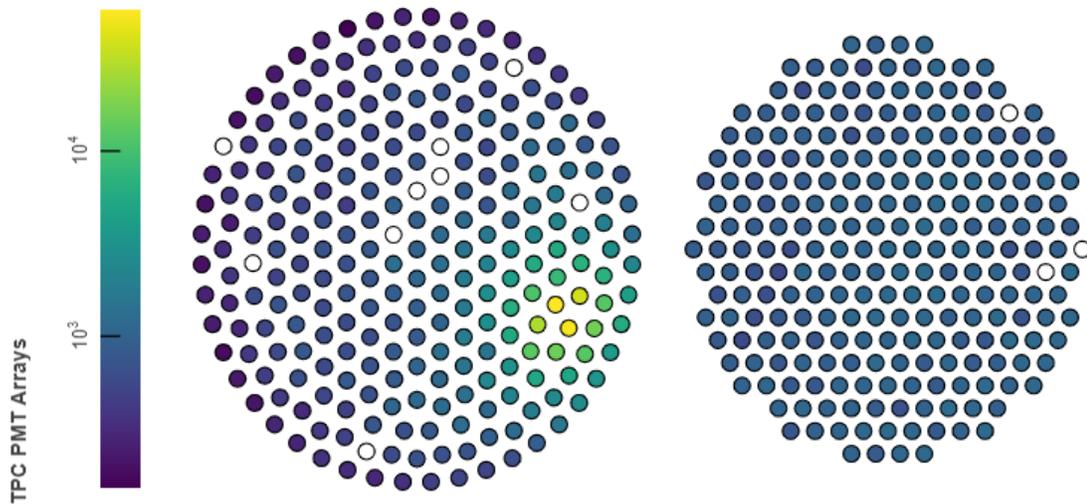
Figure 6.81: Final event 6986 from run 87111. The delay time is $0.19 \mu\text{s}$. The delay time is $0.38 \mu\text{s}$.



(a) The $S1_\beta$ and $S1_\gamma$ pulses.

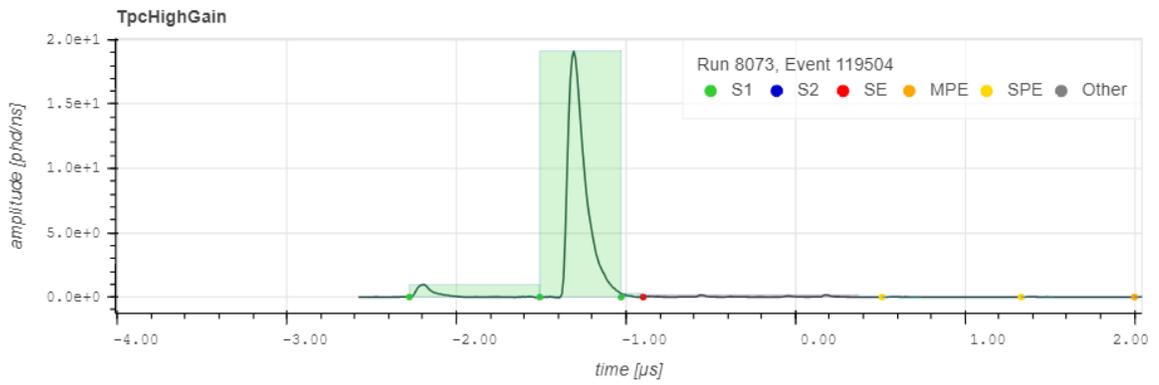


(b) The $S2$ pulses observed.

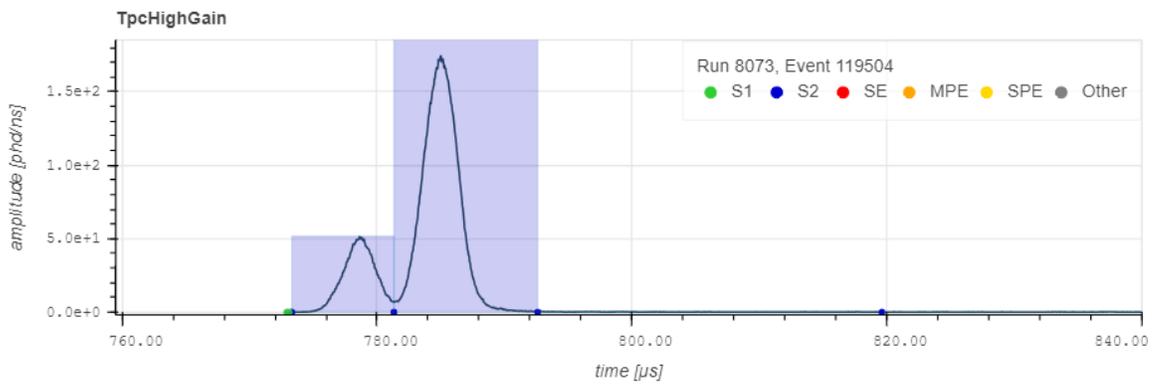


(c) The PMT hit map.

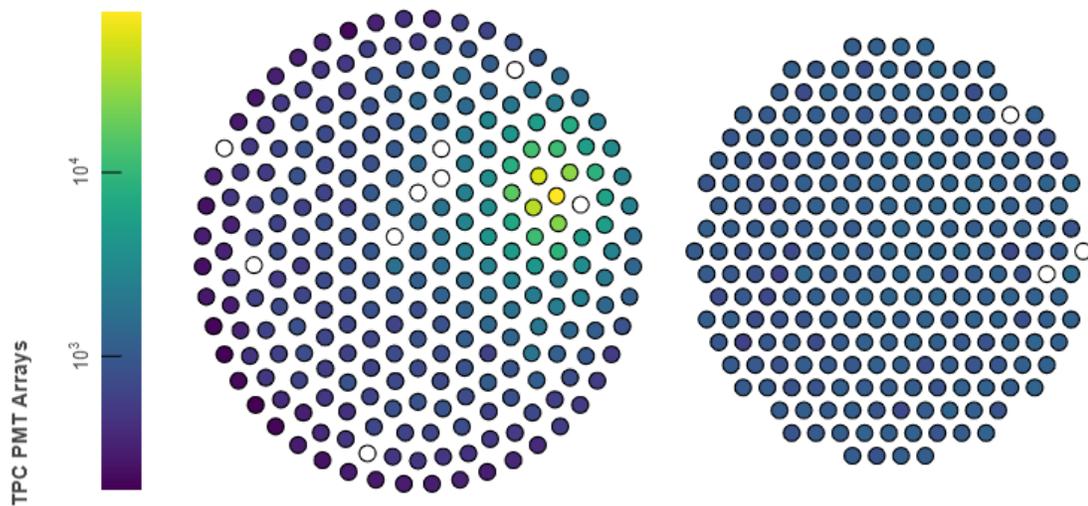
Figure 6.82: Final event 9834 from run 7664. The delay time is $0.50 \mu\text{s}$.



(a) The $S1_\beta$ and $S1_\gamma$ pulses.

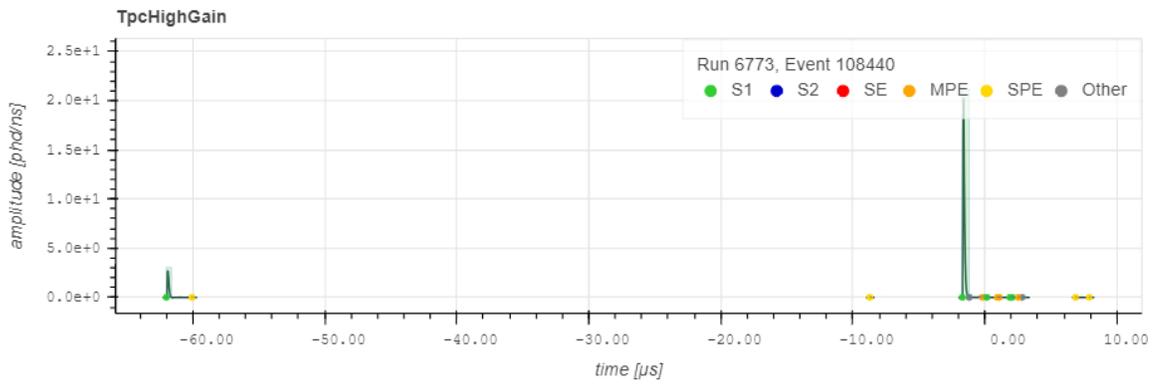


(b) The $S2$ pulses observed.

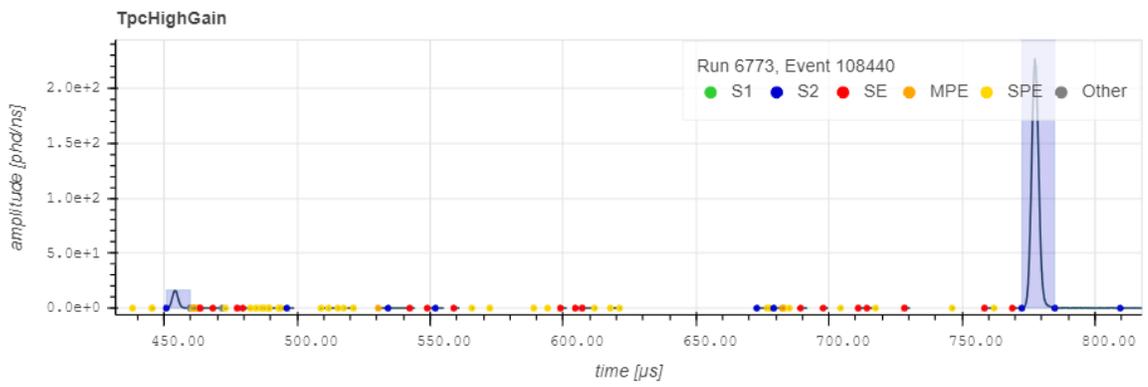


(c) The PMT hit map.

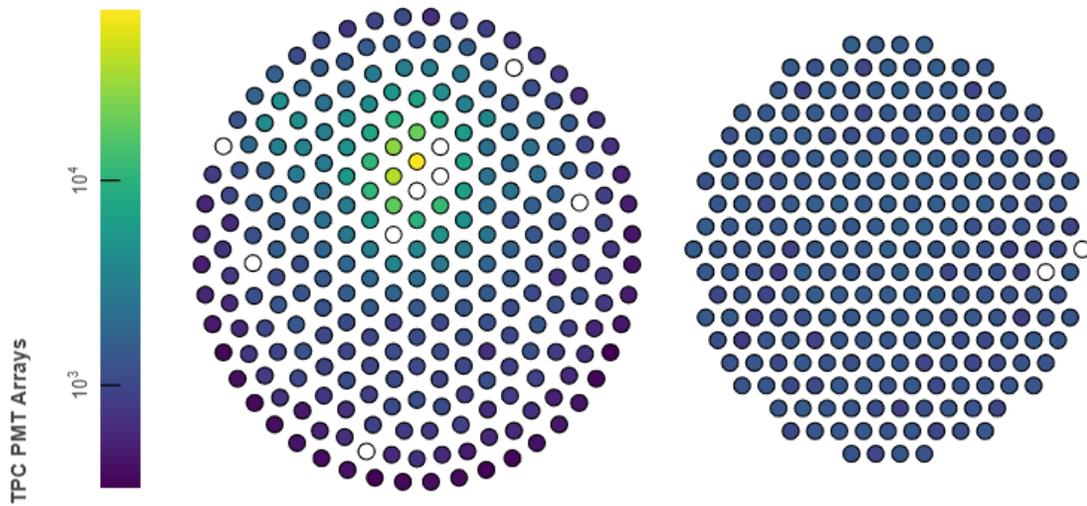
Figure 6.83: Final event 119504 from run 8073. The delay time is $0.77 \mu\text{s}$.



(a) The $S1_\beta$ and $S1_\gamma$ pulses.

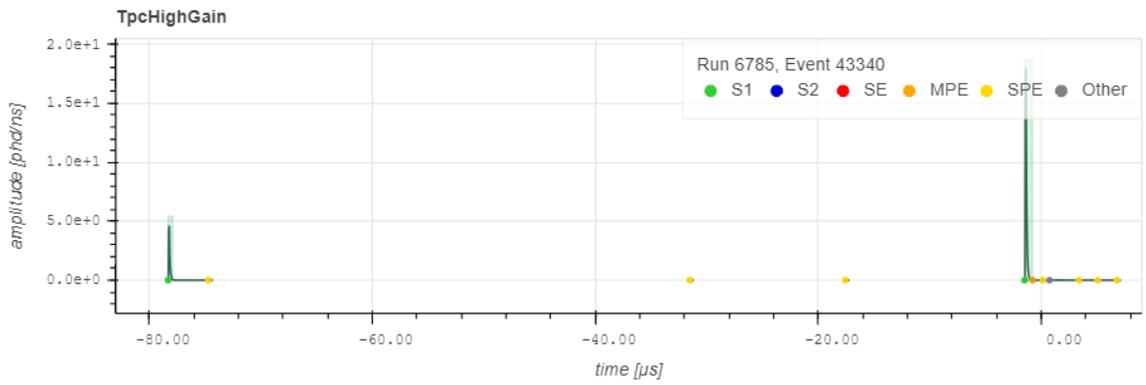


(b) The $S2$ pulses observed.

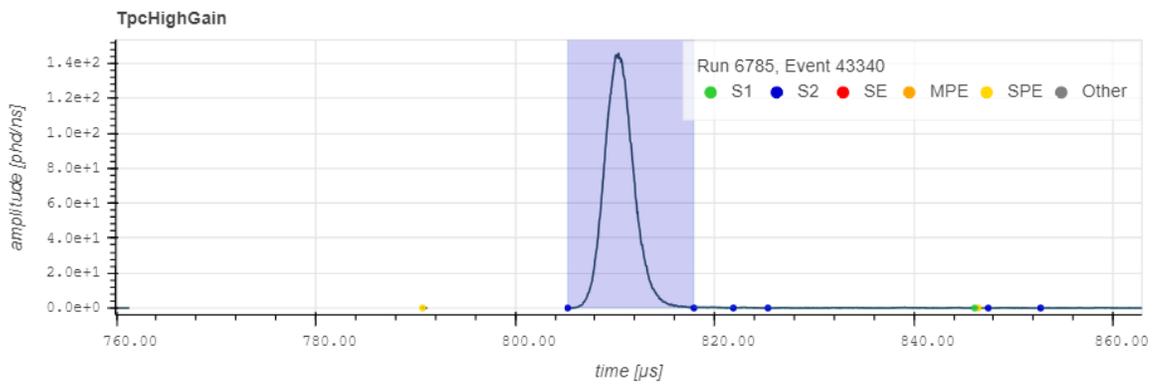


(c) The PMT hit map.

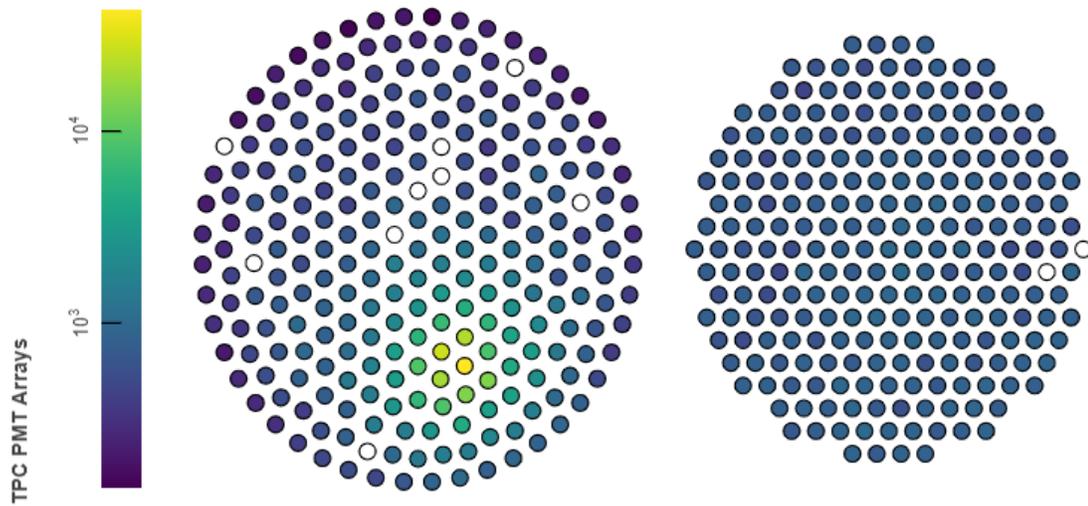
Figure 6.84: Final event 108440 from run 6773. The delay time is $60.3 \mu\text{s}$.



(a) The $S1_{\beta}$ and $S1_{\gamma}$ pulses.

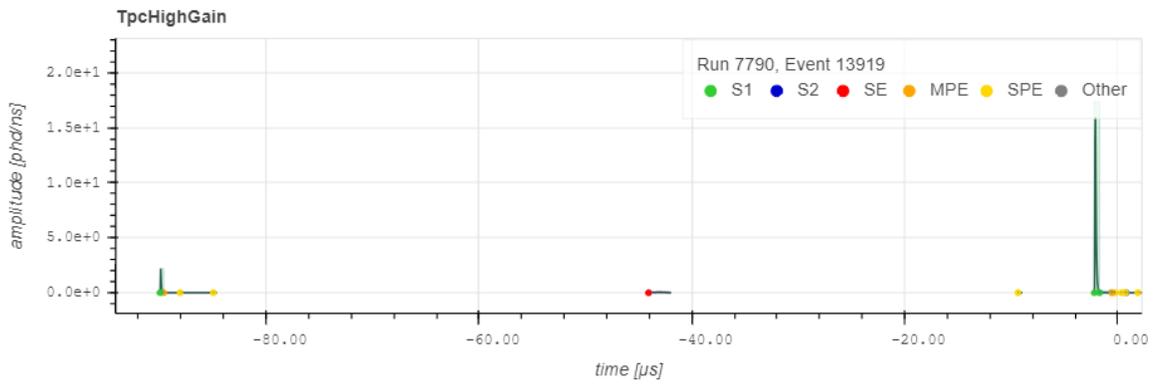


(b) The $S2$ pulses observed.

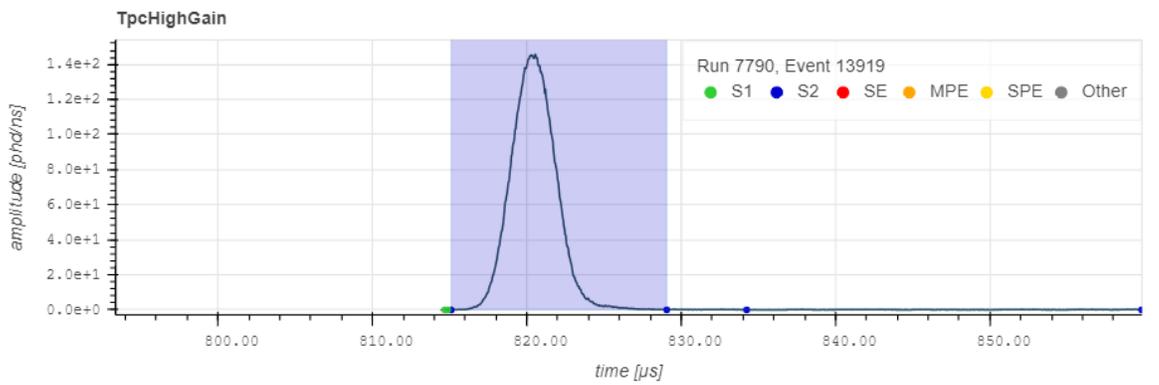


(c) The PMT hit map.

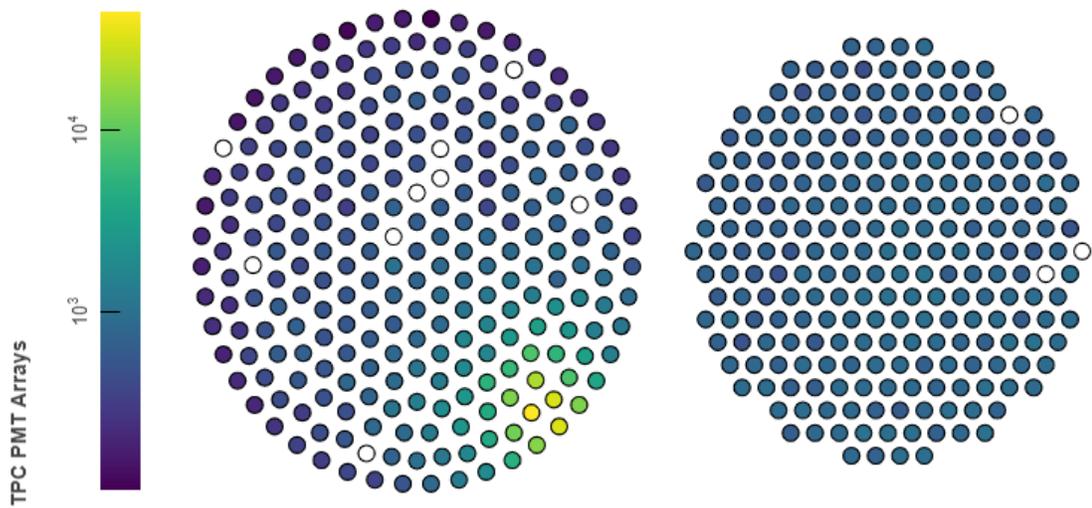
Figure 6.85: Final event 43340 from run 6785. The delay time is $76.8 \mu\text{s}$.



(a) The $S1_\beta$ and $S1_\gamma$ pulses.



(b) The $S2$ pulses observed.



(c) The PMT hit map.

Figure 6.86: Final event 13919 from run 7790. The delay time is $87.8 \mu\text{s}$.

6.9 Conclusion

The ^{85}Kr excited state decay search in LZ was carried out over 100 live days in SR1. The cut efficiencies and uncertainties were informed by two simulation techniques, both validated by comparisons to mono-energetic peaks identified in data. The analysis technique used for ^{85}Kr signal identification was able to identify all simulated events with null backgrounds, SR1 backgrounds with a null signal via side-band analysis, and a toy Monte Carlo model designed to mimic real signal and background delay times consistent with those observed in SR1. The final analysis of the SR1 data finds an excited state ^{85}Kr activity equivalent to a ^{nat}Kr to ^{nat}Xe concentration of 183 ± 67 ppq. This result is consistent with in situ measurements of xenon via the SSS at 144 ± 22 ppq.

Future versions of this analysis will benefit from updates to the full simulation chain as it is tuned to more accurately reflect the detector response. The two-simulation analysis could be combined into a single code-base reducing the burden of maintaining three separate codes (one for each simulation and one for the SR1 analysis). Future analysis will also have to re-validate pulse area cuts, in particular $S2$ areas, in the event the LZ extraction field is changed during future runs. This would motivate adding more mono-energetic sources to the simulation and SR1 validation analysis between 276 keV and 1173 keV where there is currently a gap. This will also allow proper characterization and validation $S2_{bottom}$ signals which do not cause PMT saturation at energies around 500 keV and are currently not explored due to poor peak tuning in the full simulation. LZap updates will also be beneficial, first through an external pulse corrections module and second by supporting pulse fitting identification. Currently pulses are corrected within LZap only if they are identified as a single scatter or multi-scatter event. If an external pulse corrections module is created or explicit pulse identification assignment is supported, detector corrections could be applied to the $S1_{\beta}$, $S1_{\gamma}$, and $S2$ pulses resulting in better energy resolution despite the event not

falling into either event classification. The LZap pulse identifier categorizes pulses as $S1$ s or $S2$ s based on pulse size and rise time. The pulse identifier does not fit to an $S1$ or $S2$ like shape, which if implemented could reduce pulse merging in both $S1$ and $S2$ signals that occur close in time and have slight overlap. If supported, pulse fitting would allow reductions in the minimum delay time window cut of $0.15 \mu\text{s}$. A deeper analysis of low pulse area $S1$ signals could also push the minimum $S1_{beta}$ pulse area cut lower, increasing efficiency.

Chapter 7: Summary

The evidence for dark matter, discussed in chapter 1, has resulted in a multi-pronged search effort to identify its makeup. A variety of candidates were discussed and in particular the very well motivated WIMP candidate. Chapter 2 explores the kinematics of WIMPs and establishes the potential for WIMP discovery via direct detection. An experimental overview outlines these efforts, and the unique detectors that have emerged within the field. The dual-phase xenon TPC and its mechanics are discussed. Chapter 3 introduces the LZ detector, composed of a TPC, OD and skin veto. The data handling and simulations are reviewed followed by a discussion of detector calibrations and backgrounds. Finally an SR1 WIMP search result that sets new limits on the spin-independent WIMP-nucleon cross-section for WIMP masses above 9 GeV is presented.

To characterize the ^{85}Kr backgrounds LZ relies on measurements performed by an XPM, as discussed in chapter 4. A detailed review of the XPM operating principle, hardware, procedures, and measurement analysis is given to serve as a reference for future system designs and deployment. The construction of the three XPMs were reviewed. Each XPM successfully demonstrated measurement capabilities to krypton concentrations on the order of 10 ppq. The theoretical model describing XPM response is expanded to include a variable temperature cold-trap and verified with data. The new model suggests significant gains to krypton sensitivity can be made through continued reduction of cold-trap temperatures and RGA pumping speeds.

Due to the non-interacting nature of krypton impurities a krypton removal campaign

was completed prior to filling the LZ detector. Chapter 5 described the krypton removal system and components that were used to purify the 10.4 tonnes of xenon to LZ specifications. The processed xenon was sampled prior to, throughout, and following krypton removal by the sampler XPM which measured impurities spanning nine orders of magnitude. This required a robust and complex automated system that operated in cadence with the krypton removal system. The sampler handled a wide variety of custom procedures that could be queued by non-specialized operators while minimizing any potential xenon loss. The sampler completed over 1000 measurements which proved integral in troubleshooting critical infrastructure within the krypton removal system. The sampler also implemented a run viewer with data history to keep track of each measurement result. In addition to identifying a critical xenon purification technique via freezer distillation, the sampler measured a mass averaged krypton concentration of 123 ± 10 ppq upon shipment from SLAC. This concentration met the design specifications of LZ and is in good agreement with the 144 ± 22 ppq krypton concentration measured via the SSS at SURF during SR1.

The krypton measurements presented in chapter 5 are independently confirmed through the direct observation of ^{85}Kr excited state decays discussed in chapter 6. This measurement technique relied on the characterization of a delayed coincidence β - γ signal distributed throughout the LZ fiducial volume. Simulations were used to inform detector response. The simulations were validated with comparisons to monoenergetic calibration sources identified in SR1. The efficiency of each cut and where appropriate its uncertainties are evaluated and applied to the SR1 data. The analysis is further validated by a side-band study and toy Monte Carlo model. The final 11.0 ± 4.0 events found correspond to a $^{nat}\text{Kr}/^{nat}\text{Xe}$ concentration of 183 ± 67 ppq which is in good agreement with the SSS result.

This result showing agreement between two independent measurement techniques is the culmination of nearly a decade of work I put into LZ. Some of the greatest challenges I

faced occurred during the krypton removal campaign. Many of the suggestions and explicit operations outlined in chapter 4 are included to offer support and a sanity check to future operators who may find themselves contending with a misbehaving XPM. Commissioning a high sensitivity measurement apparatus in a semi-outdoor environment introduced many complications from unexpected environmental factors. Choosing to sacrifice ultimate sensitivity in favor of measurement reliability required careful tuning of the system to satisfy all of the operational needs. During the mixing panel's commissioning the krypton removal system was also under construction. As subsystems were completed they relied heavily on measurements from the mixing panel to ensure proper operation. This was a logistics challenge of managing operations between two critical systems and avoiding unnecessary downtime. Once both the mixing panel and krypton removal system had finally achieved the final commissioning milestones COVID caused SLAC, and the world, to shut down. During this frustrating time, automation was the only task I could actively pursue under lockdown, but resulted in the robust control software presented in chapter 5. Despite the challenges discussed here, many more remain hidden in the information given in this thesis. Some I was able to address through stubbornness, but a majority were only resolved through collaboration with my fellow researchers.

As the LZ experiment continues to collect data an accurate account of the krypton within it remains crucial. The XPMs I constructed, commissioned, and operated were critical to the success of the krypton removal campaign and subsequent xenon measurements performed during SR1. The ability to directly monitor the real time krypton concentration of LZ's xenon is a powerful tool that cannot be understated. The XPM measurements were used to accurately predict background contributions from ^{85}Kr events and is validated by the result of the independent excited state decay search I completed for SR1. These krypton identification tools will continue to gather data and characterize ^{85}Kr backgrounds as LZ works toward a target exposure of 1000 live days. Improvements to either krypton

identification technique through sensitivity gains from cold-trap temperature reduction or increased statistics for the excited state decay search, will yield more accurate background models for LZ.

Appendix A: List of Publications

1. The LZ collaboration, 'A search for new physics in low-energy electron recoils from the first LZ exposure', arXiv:2307.15753.
2. The LZ collaboration, 'Background Determination for the LUX-ZEPLIN (LZ) Dark Matter Experiment', arXiv:2211.17120.
3. The LZ Collaboration, 'First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment', arXiv:2207.03764.
4. The LZ Collaboration 'A Next-Generation Liquid Xenon Observatory for Dark Matter and Neutrino Physics', arXiv:2203.02309.
5. The LZ Collaboration, 'Cosmogenic production of ^{37}Ar in the context of the LUX-ZEPLIN experiment', Phys.Rev.D 105 (2022) 8, 082004, arXiv:2201.02858.
6. The LZ Collaboration, 'Projected sensitivities of the LUX-ZEPLIN (LZ) experiment to new physics via low-energy electron recoils', Phys.Rev.D 104 (2021) 9, 092009, arXiv:2102.11740.
7. The LZ Collaboration, 'The LUX-ZEPLIN (LZ) radioactivity and cleanliness control programs', Eur.Phys.J.C 80 (2020) 11, arXiv:2006.02506.
8. The LZ Collaboration 'Projected sensitivity of the LUX-ZEPLIN (LZ) experiment to the two-neutrino and neutrinoless double beta decays of ^{134}Xe ', Phys.Rev.C 104

- (2021) 6, 066501, arXiv:2104.13374.
9. The LZ Collaboration, 'Projected sensitivity of the LUX-ZEPLIN experiment to the $0\nu\beta\beta$ decay of ^{136}Xe ', Phys.Rev.C 104 (2021) 1, 014602, arXiv:1912.04248.
 10. The LZ Collaboration, 'Simulations of Events for the LUX-ZEPLIN (LZ) Dark Matter Experiment', Astropart.Phys. 125 (2021) 102480, arXiv:2001.09363.
 11. The LZ Collaboration, 'The LUX-ZEPLIN (LZ) Experiment', Nucl.Instrum.Meth.A 953 (2020) 163047, arXiv:1910.09124.
 12. The LZ Collaboration, 'Measurement of the Gamma Ray Background in the Davis Cavern at the Sanford Underground Research Facility', Astropart.Phys. 116 (2020) 102391, arXiv:1904.02112.
 13. The LZ Collaboration, 'Projected WIMP sensitivity of the LUX-ZEPLIN (LZ) dark matter experiment', Phys.Rev.D 101 (2020) 5, 052002, arXiv:1802.06039.
 14. The LZ Collaboration, 'LUX-ZEPLIN (LZ) Technical Design Report', arXiv:1703.09144.

Bibliography

- [1] V. Rubin, J. Ford, W.K., and N. Thonnard, *apj* **238**, 471 (1980), URL <https://ui.adsabs.harvard.edu/abs/1980ApJ...238..471R>.
- [2] P. M. Salucci P., *Dark matter halos around galaxies* (1997), arXiv:astro-ph/9703027, URL <https://arxiv.org/abs/astro-ph/9703027>.
- [3] D. Tong, *Lectures on general relativity* (2019), URL <http://www.damtp.cam.ac.uk/user/tong/gr.html>.
- [4] S. Mollerach and E. Roulet, *Gravitational lensing theory* (World Scientific, 2002).
- [5] M. Limousin, J. Richard, and et al., *The Astrophysical Journal* **668**, 643 (2007), URL <https://dx.doi.org/10.1086/521293>.
- [6] R. Massey, T. Kitching, and J. Richard, *Reports on Progress in Physics* **73**, 086901 (2010), URL <https://doi.org/10.1088%2F0034-4885%2F73%2F8%2F086901>.
- [7] M. J. Jee and H. Ford, *Hubble finds ring of dark matter* (2007), URL <https://hubblesite.org/contents/news-releases/2007/news-2007-17.html#section-id-2>.
- [8] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, *The Astrophysical Journal* **648**, L109 (2006), URL <https://doi.org/10.1086%2F508162>.
- [9] zzzz, *Planck cmb* (2013), URL https://www.esa.int/ESA_Multimedia/Images/2013/03/Planck_CMB.
- [10] Planck Collaboration, N. Aghanim, Y. Akrami, M. Ashdown, J. Aumont, C. Bacigalupi, and et al., *Astronomy & Astrophysics* **641**, A6 (2020), URL <https://doi.org/10.1051%2F0004-6361%2F201833910>.
- [11] J. L. Feng, *Annual Review of Astronomy and Astrophysics* **48**, 495 (2010), URL <https://doi.org/10.1146%2Fannurev-astro-082708-101659>.
- [12] J. Balajthy, Ph.D. thesis, University of Maryland (2018), URL <https://drum.lib.umd.edu/handle/1903/20824>.

- [13] R. W. SCHNEE, in *Physics of the Large and the Small* (WORLD SCIENTIFIC, 2011), URL https://doi.org/10.1142%2F9789814327183_0014.
- [14] V. Chepel and H. Araújo, *Journal of Instrumentation* **8**, R04001 (2013), URL <https://dx.doi.org/10.1088/1748-0221/8/04/R04001>.
- [15] P. Cushman, C. Galbiati, D. N. McKinsey, H. Robertson, T. M. P. Tait, D. Bauer, A. Borgland, B. Cabrera, F. Calaprice, J. Cooley, et al., *Snowmass cFI summary: Wimp dark matter direct detection* (2013), URL <https://arxiv.org/abs/1310.8327>.
- [16] K. M. Stifter, Ph.D. thesis, Stanford U. (2021), URL <https://searchworks.stanford.edu/view/13974184>.
- [17] L.-Z. Collaboration, *Lux-zeplin (lz) technical design report* (2017), URL <https://arxiv.org/abs/1703.09144>.
- [18] E. Aprile, K. L. Giboni, P. Majewski, K. Ni, M. Yamashita, R. Hasty, A. Manzur, and D. N. McKinsey, *Phys. Rev. D* **72**, 072006 (2005), URL <https://link.aps.org/doi/10.1103/PhysRevD.72.072006>.
- [19] E. Aprile, L. Baudis, B. Choi, K. L. Giboni, K. Lim, A. Manalaysay, M. E. Monzani, G. Plante, R. Santorelli, and M. Yamashita, *Phys. Rev. C* **79**, 045807 (2009), URL <https://link.aps.org/doi/10.1103/PhysRevC.79.045807>.
- [20] V. Chepel, V. Solovov, F. Neves, A. Pereira, P. J. Mendes, C. P. Silva, A. Lindote, J. Pinto da Cunha, M. I. Lopes, and S. Kossionides, *Astropart. Phys.* **26**, 58 (2006), [physics/0512136](https://arxiv.org/abs/physics/0512136).
- [21] A. Manzur, A. Curioni, L. Kastens, D. N. McKinsey, K. Ni, and T. Wongjirad, *Physical Review C* **81** (2010), URL <https://doi.org/10.1103%2Fphysrevc.81.025808>.
- [22] L. Collaboration, D. S. Akerib, S. Alsum, H. M. Araújo, X. Bai, A. J. Bailey, J. Balajthy, P. Beltrame, E. P. Bernard, A. Bernstein, et al., *Low-energy (0.7-74 keV) nuclear recoil calibration of the lux dark matter experiment using d-d neutron scattering kinematics* (2016), [1608.05381](https://arxiv.org/abs/1608.05381).
- [23] R. Linehan, R. Mannino, A. Fan, C. Ignarra, S. Luitz, K. Skarpaas, T. Shutt, D. Akerib, S. Alsum, T. Anderson, et al., *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **1031**, 165955 (2022), URL <https://doi.org/10.1016%2Fj.nima.2021.165955>.
- [24] L.-Z. Collaboration, *First dark matter search results from the lux-zeplin (lz) experiment* (2022), URL <https://arxiv.org/abs/2207.03764>.

- [25] L. collaboration, *Background determination for the lux-zepplin (lz) dark matter experiment* (2022), URL <https://arxiv.org/abs/2211.17120>.
- [26] B. B., *Custom one gallon swagelok bottle analysis* (2019), URL <http://teacher.pas.rochester.edu:8080/wiki/bin/view/Lz/LzCtdDB0240>.
- [27] SRS, *Operating manual and programming reference: Models rga100, rga200, and rga300 residual gas analyzer* (2009), URL <https://www.thinksrs.com/downloads/pdfs/manuals/RGAm.pdf>.
- [28] D. Ames, *Krypton removal via gas chromatography for the lz experiment* (2022), URL https://indico.sanfordlab.org/event/29/papers/403/files/147-LRT_2022_Proceedings-2.pdf.
- [29] Z. V., N. K.N., V. M., and V. L., *Livechart of nuclides - advanced version* (2014), URL <https://www.iaea.org/resources/databases/livechart-of-nuclides-advanced-version>.
- [30] A. Musalhi and D. Kordoff, *A look at xe activations induced by dd calibrations* (2021).
- [31] *Air - composition and molecular weight* (2003), URL https://www.engineeringtoolbox.com/air-composition-d_212.html.
- [32] F. Zwicky, *Helvetica Physica Acta* **6**, 110 (1933), URL <https://ned.ipac.caltech.edu/level5/March17/Zwicky/frames.html>.
- [33] F. Zwicky, *apj* **86**, 217 (1937), URL <https://ui.adsabs.harvard.edu/abs/1937ApJ...86..217Z>.
- [34] S. Smith, *apj* **83**, 23 (1936), URL <https://ui.adsabs.harvard.edu/abs/1936ApJ...83...23S>.
- [35] F. D. Kahn and L. Woltjer, *apj* **130**, 705 (1959), URL <https://ui.adsabs.harvard.edu/abs/1959ApJ...130..705K>.
- [36] H. W. Babcock, *Lick Observatory Bulletin* **498**, 41 (1939), URL <https://ui.adsabs.harvard.edu/abs/1939LicOB...19...41B>.
- [37] J. H. Oort, *apj* **91**, 273 (1940), URL <https://ui.adsabs.harvard.edu/abs/1940ApJ...91..273O>.
- [38] J. P. Ostriker and P. J. E. Peebles, *apj* **186**, 467 (1973), URL <https://ui.adsabs.harvard.edu/abs/1973ApJ...186..467O>.
- [39] R. Herman, *Derivation of the geodesic equation and defining the christoffel symbols* (2008), URL <http://people.uncw.edu/hermanr/GR/geodesic.pdf>.

- [40] A. Einstein, *Science* **84**, 506 (1936), URL <https://ui.adsabs.harvard.edu/abs/1936Sci...84..506E>.
- [41] F. Zwicky, *Phys. Rev.* **51**, 290 (1937), URL <https://link.aps.org/doi/10.1103/PhysRev.51.290>.
- [42] P. Young, J. E. Gunn, J. Kristian, J. B. Oke, and J. A. Westphal, *apj* **241**, 507 (1980), URL <https://ui.adsabs.harvard.edu/abs/1980ApJ...241..507Y>.
- [43] L. J. Goicoechea, R. Gil-Merino, and A. Ullan, *Monthly Notices of the Royal Astronomical Society: Letters* **360**, L60 (2005), URL <https://doi.org/10.1111/2Fj.1745-3933.2005.00046.x>.
- [44] M. J. Jee, H. C. Ford, G. D. Illingworth, R. L. White, T. J. Broadhurst, D. A. Coe, G. R. Meurer, A. van der Wel, N. Benitez, J. P. Blakeslee, et al., *The Astrophysical Journal* **661**, 728 (2007), URL <https://doi.org/10.1086%2F517498>.
- [45] D. Tong, *Lectures on cosmology* (2019), URL <http://www.damtp.cam.ac.uk/user/tong/cosmo.html>.
- [46] S. Dodelson, *Modern Cosmology* (Academic Press, Amsterdam, 2003), ISBN 978-0-12-219141-1.
- [47] B. Ryden, *Introduction to cosmology* (Cambridge University Press, 1970), ISBN 978-1-107-15483-4, 978-1-316-88984-8, 978-1-316-65108-7.
- [48] C. J. Copi, D. N. Schramm, and M. S. Turner, *Science* **267**, 192 (1995), URL <https://doi.org/10.1126%2Fscience.7809624>.
- [49] W. Hu, *Cmb introduction* (2001), URL <http://background.uchicago.edu/~whu/intermediate/intermediate.html>.
- [50] C. Alcock, R. A. Allsman, D. R. Alves, T. S. Axelrod, A. C. Becker, D. P. Bennett, K. H. Cook, N. Dalal, A. J. Drake, K. C. Freeman, et al., *apj* **542**, 281 (2000), *astro-ph/0001272*, URL <https://ui.adsabs.harvard.edu/abs/2000ApJ...542..281A>.
- [51] T. D. Brandt, *The Astrophysical Journal* **824**, L31 (2016), URL <https://doi.org/10.3847%2F2041-8205%2F824%2F2%2F131>.
- [52] S. D. M. White, C. S. Frenk, and M. Davis, *apjl* **274**, L1 (1983), URL <https://ui.adsabs.harvard.edu/abs/1983ApJ...274L...1W>.
- [53] L. D. Duffy and K. van Bibber, *New Journal of Physics* **11**, 105008 (2009), URL <https://doi.org/10.1088%2F1367-2630%2F11%2F10%2F105008>.

- [54] N. Du, N. Force, R. Khatiwada, E. Lentz, R. Ottens, L. J. Rosenberg, G. Rybka, G. Carosi, N. Woollett, D. Bowring, et al. (ADMX Collaboration), *Phys. Rev. Lett.* **120**, 151301 (2018), URL <https://link.aps.org/doi/10.1103/PhysRevLett.120.151301>.
- [55] T. Braine, R. Cervantes, N. Crisosto, N. Du, S. Kimes, L. J. Rosenberg, G. Rybka, J. Yang, D. Bowring, A. S. Chou, et al. (ADMX Collaboration), *Phys. Rev. Lett.* **124**, 101303 (2020), URL <https://link.aps.org/doi/10.1103/PhysRevLett.124.101303>.
- [56] C. Bartram, T. Braine, E. Burns, R. Cervantes, N. Crisosto, N. Du, H. Korandla, G. Leum, P. Mohapatra, T. Nitta, et al. (ADMX Collaboration), *Phys. Rev. Lett.* **127**, 261803 (2021), URL <https://link.aps.org/doi/10.1103/PhysRevLett.127.261803>.
- [57] J. Feng, *Standard wimps*, URL <https://indico.cern.ch/event/949654/contributions/4025911/attachments/2288409/3896337/feng1.pdf>.
- [58] A. Boveia and C. Doglioni, *Annual Review of Nuclear and Particle Science* **68**, 429 (2018), URL <https://doi.org/10.1146%2Fannurev-nucl-101917-021008>.
- [59] ATLAS Collaboration, arXiv e-prints arXiv:2104.13240 (2021), 2104.13240, URL <https://ui.adsabs.harvard.edu/abs/2021arXiv210413240A>.
- [60] CMS Collaboration, *European Physical Journal C* **81**, 13 (2021), 2008.04735, URL <https://ui.adsabs.harvard.edu/abs/2021EPJC...81...13S>.
- [61] R. K. Leane, *Indirect detection of dark matter in the galaxy* (2020), URL <https://arxiv.org/abs/2006.00513>.
- [62] A. Olivares-Del Campo, C. Boehm, S. Palomares-Ruiz, and S. Pascoli, *Phys. Rev. D* **97**, 075039 (2018), URL <https://link.aps.org/doi/10.1103/PhysRevD.97.075039>.
- [63] C. Armand, E. Charles, M. di Mauro, C. Giuri, J. P. Harding, D. Kerszberg, T. Miener, E. Moulin, L. Oakes, V. Poireau, et al., *Combined dark matter searches towards dwarf spheroidal galaxies with fermi-lat, hawc, h.e.s.s., magic, and veritas* (2021), URL <https://arxiv.org/abs/2108.13646>.
- [64] A. Collaboration, *Phys. Rev. Lett.* **113**, 121101 (2014), URL <https://link.aps.org/doi/10.1103/PhysRevLett.113.121101>.
- [65] O. Adriani, G. C. Barbarino, G. A. Bazilevskaya, R. Bellotti, and et al, *Nature* **458**, 607 (2009), URL <https://doi.org/10.1038%2Fnature07942>.
- [66] DAMPE Collaboration, *Nature* **552**, 63 (2017), URL <https://doi.org/10.1038%2Fnature24475>.

- [67] Y.-X. Zhang, X.-L. Zhou, Y.-H. Zhao, and X.-B. Wu, *The Astronomical Journal* **145**, 42 (2013), URL <https://doi.org/10.1088%2F0004-6256%2F145%2F2%2F42>.
- [68] J. Bovy and H.-W. Rix, *The Astrophysical Journal* **779**, 115 (2013), URL <https://doi.org/10.1088%2F0004-637x%2F779%2F2%2F115>.
- [69] J. I. Read, *Journal of Physics G: Nuclear and Particle Physics* **41**, 063101 (2014), URL <https://doi.org/10.1088%2F0954-3899%2F41%2F6%2F063101>.
- [70] J. Lewin and P. Smith, *Astroparticle Physics* **6**, 87 (1996), ISSN 0927-6505, URL <https://www.sciencedirect.com/science/article/pii/S0927650596000473>.
- [71] M. Schumann, *Journal of Physics G: Nuclear and Particle Physics* **46**, 103003 (2019), URL <https://doi.org/10.1088%2F1361-6471%2Fab2ea5>.
- [72] C. Collaboration (CDEX Collaboration), *Phys. Rev. Lett.* **120**, 241301 (2018), URL <https://link.aps.org/doi/10.1103/PhysRevLett.120.241301>.
- [73] C. E. Aalseth, P. S. Barbeau, N. S. Bowden, B. Cabrera-Palmer, J. Colaresi, J. I. Collar, S. Dazeley, P. de Lurgio, J. E. Fast, N. Fields, et al. (CoGeNT Collaboration), *Phys. Rev. Lett.* **106**, 131301 (2011), URL <https://link.aps.org/doi/10.1103/PhysRevLett.106.131301>.
- [74] A. Aguilar-Arevalo, D. Amidei, X. Bertou, M. Butner, G. Canelo, A. Castañeda Vázquez, B. A. Cervantes Vergara, A. E. Chavarria, C. R. Chavez, J. R. T. de Mello Neto, et al. (DAMIC Collaboration), *Phys. Rev. D* **94**, 082006 (2016), URL <https://link.aps.org/doi/10.1103/PhysRevD.94.082006>.
- [75] M. Crisler, R. Essig, J. Estrada, G. Fernandez, J. Tiffenberg, M. S. Haro, T. Volansky, and T.-T. Yu (SENSEI Collaboration), *Phys. Rev. Lett.* **121**, 061803 (2018), URL <https://link.aps.org/doi/10.1103/PhysRevLett.121.061803>.
- [76] R. Bernabei, P. Belli, A. Bussolotti, F. Cappella, V. Caracciolo, R. Cerulli, C. Dai, A. d'Angelo, A. D. Marco, et al., *Nuclear Physics and Atomic Energy* **19**, 307 (2018), URL <https://doi.org/10.15407%2Fjnpae2018.04.307>.
- [77] H. Lee, B. J. Park, J. J. Choi, O. Gileva, C. Ha, A. Iltis, E. J. Jeon, D. Y. Kim, K. W. Kim, S. H. Kim, et al., *Frontiers in Physics* **11** (2023), URL <https://doi.org/10.3389%2Ffphy.2023.1142765>.
- [78] T. Shutt, B. Ellman, P. D. Barnes, A. Cummings, A. Da Silva, J. Emes, Y. Giraud-Héraud, E. E. Haller, A. E. Lange, R. R. Ross, et al., *Phys. Rev. Lett.* **69**, 3425 (1992), URL <https://link.aps.org/doi/10.1103/PhysRevLett.69.3425>.
- [79] S. collaboration, *Physical Review Letters* **120** (2018), URL <https://doi.org/10.1103%2Fphysrevlett.120.061802>.

- [80] P. Collaboration (PICO Collaboration), *Phys. Rev. Lett.* **118**, 251301 (2017), URL <https://link.aps.org/doi/10.1103/PhysRevLett.118.251301>.
- [81] J. Angle, E. Aprile, F. Arneodo, L. Baudis, A. Bernstein, A. Bolozdynya, P. Brusov, L. C. C. Coelho, C. E. Dahl, L. DeViveiros, et al., *Physical Review Letters* **100** (2008), URL <https://doi.org/10.1103/PhysRevLett.100.021303>.
- [82] D. Collaboration, *Physical Review D* **98** (2018), URL <https://doi.org/10.1103/PhysRevD.98.102006>.
- [83] D. Collaboration, *The European Physical Journal Plus* **133** (2018), URL <https://doi.org/10.1140/epjp/i2018-11973-4>.
- [84] X. Collaboration, *Phys. Rev. C* **106**, 024328 (2022), URL <https://link.aps.org/doi/10.1103/PhysRevC.106.024328>.
- [85] E. Collaboration, *Phys. Rev. C* **89**, 015502 (2014), URL <https://link.aps.org/doi/10.1103/PhysRevC.89.015502>.
- [86] V. Solovov, V. Chepel, M. Lopes, A. Hitachi, R. Ferreira Marques, and A. Policarpo, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **516**, 462 (2004), ISSN 0168-9002, URL <https://www.sciencedirect.com/science/article/pii/S0168900203024331>.
- [87] L. Collaboration, *Physical Review D* **97** (2018), URL <https://doi.org/10.1103/PhysRevD.97.112002>.
- [88] T. Shutt, C. Dahl, J. Kwong, A. Bolozdynya, and P. Brusov, *Nuclear Physics B - Proceedings Supplements* **173**, 160 (2007), ISSN 0920-5632, proceedings of the 7th UCLA Symposium on Sources and Detection of Dark Matter and Dark Energy in the Universe, URL <https://www.sciencedirect.com/science/article/pii/S092056320700669X>.
- [89] C. E. Dahl, Ph.D. thesis, Princeton U. (2009), URL <https://inspirehep.net/literature/1374815>.
- [90] P. Sorensen and C. E. Dahl, *Phys. Rev. D* **83**, 063501 (2011), URL <https://link.aps.org/doi/10.1103/PhysRevD.83.063501>.
- [91] J. Lindhard, V. Nielsen, M. Scharff, and P. Thomsen, *Mat. Fys. Medd. Dan. Vid. Selsk* **33**, 1 (1963).
- [92] E. Aprile, C. E. Dahl, L. de Viveiros, R. J. Gaitskell, K. L. Giboni, J. Kwong, P. Majewski, K. Ni, T. Shutt, and M. Yamashita, *Phys. Rev. Lett.* **97**, 081302 (2006), URL <https://link.aps.org/doi/10.1103/PhysRevLett.97.081302>.

- [93] T. Takahashi, S. Konno, T. Hamada, M. Miyajima, S. Kubota, A. Nakamoto, A. Hitachi, E. Shibamura, and T. Doke, *Phys. Rev. A* **12**, 1771 (1975), URL <https://link.aps.org/doi/10.1103/PhysRevA.12.1771>.
- [94] L. Collaboration, *Lz software docs: Sr1 core cuts* (2023), URL <https://luxzeplin.gitlab.io/docs/software/docs/analysis/analysiswithrqs/alpaca/tools/SR1CoreCuts.html#od-veto>.
- [95] A. Monte, *Tuned data quality cuts for sr1* (2022).
- [96] A. Fan, *Sr1 ws updates v* (2022).
- [97] A. Monte, *Skin studies for sr1* (2022).
- [98] D.-M. Mei, C. Zhang, K. Thomas, and F. Gray, *Astroparticle Physics* **34**, 33 (2010), ISSN 0927-6505, URL <https://www.sciencedirect.com/science/article/pii/S0927650510000782>.
- [99] V. Tomasello, M. Robinson, and V. Kudryavtsev, *Astroparticle Physics* **34**, 70 (2010), ISSN 0927-6505, URL <https://www.sciencedirect.com/science/article/pii/S0927650510001015>.
- [100] L. Collaboration, *Physical Review D* **101** (2020), URL <https://doi.org/10.1103/PhysRevD.101.052002>.
- [101] L.-Z. Collaboration, *Physical Review D* **105** (2022), URL <https://doi.org/10.1103/PhysRevD.105.082004>.
- [102] H. H. Loosli, B. E. Lehmann, and W. M. Smethie, *Noble Gas Radioisotopes: 37Ar; 85Kr; 39Ar; 81Kr* (Springer US, Boston, MA, 2000), pp. 379–396, ISBN 978-1-4615-4557-6, URL https://doi.org/10.1007/978-1-4615-4557-6_12.
- [103] D. Leonard, A. Dobi, L. K. C. Hall, T. Langford, S. Slutsky, and Y.-R. Yen, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **621**, 678 (2010), URL <https://doi.org/10.1016/j.nima.2010.04.152>.
- [104] A. Dobi, C. Davis, C. Hall, T. Langford, S. Slutsky, and Y.-R. Yen, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **665**, 1 (2011), URL <https://doi.org/10.1016/j.nima.2011.11.043>.
- [105] A. Dobi, C. Hall, S. Slutsky, Y.-R. Yen, B. Aharmin, M. Auger, P. Barbeau, C. Benitez-Medina, M. Breidenbach, B. Cleveland, et al., *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **675**, 40 (2012), ISSN 0168-9002, URL <https://www.sciencedirect.com/science/article/pii/S016890021200126X>.

- [106] *Air - maximum moisture carrying capacity* (2008), URL https://www.engineeringtoolbox.com/maximum-moisture-content-air-d_1403.html.
- [107] *Swagelok welding system m200 power supply* (2022), URL <https://www.swagelok.com/downloads/webcatalogs/en/ms-02-342.pdf>.
- [108] *Swagelok welding system series orbital welding system* (2022), URL <https://www.swagelok.com/downloads/webcatalogs/en/ms-02-129.pdf>.
- [109] S. J., *Mobile sampling system qa doc* (2020), URL <http://teacher.pas.rochester.edu:8080/wiki/bin/view/Lz/LzCtdDB0149>.
- [110] S. J., *Stationary sampling system qa doc* (2020), URL <http://teacher.pas.rochester.edu:8080/wiki/bin/view/Lz/LzCtdDB0150>.
- [111] Omega, *Platinum series temperature, process and strain controllers* (2023), URL https://assets.omega.com/spec/CNPT_SERIES.pdf.
- [112] G. A. and S. J., *Nature* **204**, 1084–1085 (1964), URL <https://doi.org/10.1038/2041084a0>.
- [113] D. Akerib, H. Araújo, X. Bai, A. Bailey, J. Balajthy, P. Beltrame, E. Bernard, A. Bernstein, T. Biesiadzinski, E. Boulton, et al. (LUX), *Astroparticle Physics* **97**, 80 (2018), URL <https://doi.org/10.1016%2Fj.astropartphys.2017.10.014>.
- [114] C. Lee, Ph.D. thesis, Case Western Reserve University (2015), URL http://rave.ohiolink.edu/etdc/view?acc_num=case1427482791.
- [115] SRS, *Bga244 binary gas analyzer user manual* (2018), URL <https://www.thinksrs.com/downloads/pdfs/manuals/BGA244m.pdf>.
- [116] E. Aprile, , J. Aalbers, F. Agostini, M. Alfonsi, F. D. Amaro, M. Anthony, F. Arneodo, P. Barrow, L. Baudis, et al., *The European Physical Journal C* **77** (2017), URL <https://doi.org/10.1140%2Fepjc%2Fs10052-017-4757-1>.
- [117] E. Aprile, K. Abe, F. Agostini, S. A. Maouloud, L. Althueser, B. Andrieu, E. Angelino, J. Angevaare, V. Antochi, D. A. Martin, et al., *Physical Review Letters* **129** (2022), URL <https://doi.org/10.1103%2Fphysrevlett.129.161805>.
- [118] E. Aprile, K. Abe, F. Agostini, S. A. Maouloud, M. Alfonsi, L. Althueser, E. Angelino, J. R. Angevaare, V. C. Antochi, D. A. Martin, et al., *Progress of Theoretical and Experimental Physics* (2022), URL <https://doi.org/10.1093%2Fptep%2Fptac074>.
- [119] B. M.J., H. J.H., S. S.M., C. J., C. J.S., S. R., Z. D.S., and O. K., *Xcom: Photon cross section database (version 1.5)* (2010), URL <http://physics.nist.gov/xcom>.