

## ABSTRACT

Title of Dissertation:     STUDIES OF IONIZATION BACKGROUNDS  
                                  IN NOBLE LIQUID DETECTORS FOR DARK  
                                  MATTER SEARCHES

Eli Mizrachi (they/them)  
Doctor of Philosophy, 2024

Dissertation Directed by:  Dr. Jinkge Xu  
                                  Lawrence Livermore National Laboratory

Dark matter is believed to make up almost 85% of the total mass of the universe, yet its identity remains unclear. Weakly Interacting Massive Particles (WIMPs) have historically been a favored dark matter candidate, and dual-phase noble liquid time projection chambers (TPCs) have set the strongest interaction limits to date on WIMPs with a mass greater than several GeV. However, because no definitive interactions have been observed, the parameter space for conventional WIMPs is highly constrained. This has sparked greater interest in new sub-GeV dark matter models. At this mass scale, dark matter interactions with xenon or argon target media may still produce detectable signals at or near the single electron limit. However, these signals are currently obscured by delayed ionization backgrounds (“electron trains”) which persist for seconds after an ionization event occurs. Electron trains have been observed in many different experiments and exhibit similar characteristics, but their cause is only partially understood.

This work examines the nature of electron trains in various contexts, as well as possible

strategies to mitigate them. First, a characterization of electron trains in the LZ experiment is presented, including new evidence of a dependence on detector conditions. The characterization also informed the development of an electron-train veto for LZ's first WIMP search, which set world-leading limits on the spin-independent and spin-dependent WIMP-nucleon cross-sections for medium and high-mass WIMPs.

Next, to complement the analysis of LZ data, hardware upgrades were performed in XeNeu, a small xenon TPC at Lawrence Livermore National Lab. These included replacing plastics with low-outgassing metal and machinable ceramic components, as well as a replacement of XeNeu's photomultiplier tube array with silicon photomultipliers. The resulting reduction in the intensity of electron-trains and better position resolution from the respective upgrades will improve future studies of low energy interactions and phenomena. Concurrent with this work, XeNeu was used to perform a nuclear recoil calibration and a search for the Migdal effect, the latter of which can substantially enhance an experiment's low-mass dark matter sensitivity.

Finally, the development of Coherent Ionization Limits in Liquid Argon and Xenon (CHILLAX), is reported. CHILLAX is a new xenon-doped, dual-phase argon test stand that has the potential to have a higher sensitivity to low-mass dark matter interactions and lower backgrounds than current liquid xenon TPCs. The system is designed to handle high (percent level) xenon concentrations in liquid argon that can enable a range of ionization signal production and collection benefits. CHILLAX demonstrated the feasibility of such concepts by achieving a world record xenon doping concentration with stable operation.

STUDIES OF IONIZATION BACKGROUNDS IN NOBLE LIQUID DETECTORS FOR  
DARK MATTER SEARCHES

by

Eli Mizrachi (they/them)

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  
2024

Advisory Committee:

Professor Carter Hall, Chair/Advisor  
Research Professor Anwar Bhatti  
Assistant Professor Brian Clark  
Professor Sarah Penniston-Dorland  
Dr. Jingke Xu, Co-Advisor

© Copyright by  
Eli Mizrachi  
2024

Source Code: [https://gitlab.com/emiz/thesis\\_public](https://gitlab.com/emiz/thesis_public)

## Acknowledgements

The phrase “standing on the shoulder of giants” always comes to mind when I think of the people who have made my work and life possible. I’ve tried to be fairly comprehensive in who I mention, but If I missed you, I apologize. I am incredibly fortunate in having been so well-supported; there are many more people I’ve interacted with through the years who have left their mark on me in some way.

I would like to thank Carter Hall, Anwar Bhatti, Brian Clark, Sarah Penniston-Dorland, and Jingke Xu for serving on my dissertation committee. Carter, thank you for your support as my on-campus advisor. When the time came for me to find a home in LZ, your assistance was instrumental. Jingke, thank you for your unending technical guidance and financial support through my years at LLNL. I’m grateful you granted me the opportunity to contribute to such a prolific group, which has been responsible for so many critical measurements and interesting projects within the noble liquid community.

Next, I would like to thank the rest of our group members at LLNL. Teal, thank you for your mentorship within the LZ S2-only group and our lab. Joining LZ at the height of the pandemic and MDC3 was an immense challenge, and I can’t imagine navigating my early analysis efforts without you. Ethan, it was a pleasure to finally have a chance to work with you after only seeing

you in passing at LBNL. I'll always treasure our mutual appreciation for neat hardware hacks and niche machining knowledge. To Jimmy (AKA "bigjim"), thank you for always stepping up to lend a hand, whether on XeNu or CHILLAX. I'm proud to leave the lab in your hands. Godspeed, and always remember that Gielinor Awaits. Same to your Alec. Speaking of which: Alec and Jianyang, I look forward to seeing the fruits of your labor within the noble liquid group. To Rachel, I consider myself fortunate to have been able to work alongside you at both SLAC and LLNL. Our lunchtime conversations have been nothing less than illuminating, and you always seem to be asking the questions I wish I had about detectors and lab management. Finally, I owe my sincere thanks to Danny Naim. I cherished our time together at Berkeley, and when I was scrambling to find a group in LZ to work with, you had my back in helping me join you at LLNL. I'm not sure how I can ever repay the favor, but I'm sure someday I'll be reading headlines about your success.

In no particular order, I would also like to thank some of the other folks in the LLNL RED Group. Thank you to Sergey for your efforts as an architect of the XeNu high voltage system. Thank you Tom for always being a friendly presence in our office before the pandemic made offices irrelevant. To Sean and Angelique, thank you for helping create all of the parts our group has needed over the years. It's been enlightening hearing from both of you about your experiences, and I have taken many of our conversations to heart about how to design components. To Felicia and Tomi, you both exude a level of brilliance and care that I can only hope to someday match. I look forward to remaining friends over the years.

Over the past 9 (!) years, I've been fortunate to be connected to and contribute to LZ in a variety of capacities. Starting at LBNL, I have Peter Sorensen, Kate Kamdin, and James Morad to thank for helping me get off the ground with high vacuum cryogenic systems. To Lucie Tvrznikova,

thank you for your help with COMSOL when I was struggling to build meshes and simulate electric fields. I would also like to thank Mani Tripathi for hosting me during the summer of 2015 at UC Davis, and Aaron Manlaysia for mentoring me during that time. Next, I want to give a shoutout to Michael Williams, Yue Wang, Jose Soria, Vetri Velan, and Ryan Gibbons for joining me at the Berkeley grad student beer hours. I'm looking forward to getting more chances to hang out with you all and everyone else in the LBNL LZ group.

At SLAC, I have Dan Akerib and Tom Shutt to thank for their skills as hosts and chefs, and for leading an incredible group of people that developed the LZ System Test and Krypton removal system. My time in the group instilled values and principles in me that I have fiercely clung onto. Of course, I would not have had that experience without contributions from the rest of the group's members. To TJ, I consider myself an expert in designing and configuring high vacuum cryogenic systems because I managed to absorb a mere 1% of your knowledge. Someday, I might also understand Crusader Kings to the same extent. To Kelly, thank you for getting LZ GETUP off the ground. It's one thing to be so instrumental to so many technical facets of LZ, but to be equally instrumental to the wellbeing of the collaboration is a tremendous accomplishment that places you above many of LZ's members. It's been empowering to see you become a rising-star physicist who is also a force for good within the field. Ryan, I will forever curse your name for sucking me into working on electron trains, but I still managed to enjoy our time together in the S2-only group. You also have my eternal gratitude for not only appreciating my puns, but contributing your own. Alden, I still struggle to grasp the magnitude of your contributions to LZ, and I can only hope to fill a sliver of your shoes on the offline team. Tomasz, LZ would be a frozen pile of rubble if it were not held together by your collaboration-spanning efforts and zip-ties. Christina, thank you

for being such a gracious host over the years, for donating your liquor collection to me, and for your boundless expertise in overseeing the development of the krypton removal system. Steffen, thank you for deftly managing LZ's software infrastructure. I also have you to thank for what is possibly the best professional headshot of my life. Shaun, your cheerful presence IR2 was always appreciated, and I always felt like I could approach you if I needed to bounce ideas off of someone. Eric, I commend you for seeing the krypton removal process to its completion. Drew, thank you for your friendship, muon veto, and hotspot veto. I always seem to notice you laugh harder than anyone else at my jokes, and I really appreciate it. Ann, thank you for being such a great postdoc to work with in the S2-only group. Without you, I would still probably be tearing my hair out over tracking livetime following S2s. I can say the same for Micah, who I could always ping when I had questions about python or problems at NERSC. Nicole, even though you were only with us for a brief rotation, it was always a joy to see you in the lab and hang out with you after work. I consider it a privilege that I was able to bring some of your ice cream dreams to fruition and share them with everyone. Finally, I consider myself incredibly blessed to be rejoining the group under Maria Elena Monzani, where I look forward to attaining a similar level of success and knowledge with software and computing. I'm optimistic that our shared enthusiasm for best practices and sense of humor will lead to a productive and memorable time for the benefit of LZ.

Within LZ, I have other members I would like to thank. Anh, I can't think of a better person to share a meal with. I treasure our friendship and the time we spent together in South Dakota, the Bay Area, and Edinburgh. You've made important contributions to LZ, as well as to me personally. Emily, Anna, Amar, and Tom: thank you for being such a great group to hang out with. I think a lot about how we instantly hit it off, and I'm looking forward to future opportunities to do so

again. To Emily, thank you for your efforts to pick up the e-train veto where Ryan left off. Doing so was a monumental task, and I'm honored to have been able to provide support where I could. I am unabashedly looking forward to the chill vibes you'll bring to the Bay Area. I would also like to give special recognition to the electronics team on LZ. Dev, thank you for teaching me that a good Manhattan should smell like "death and the French Alps". Elise, I'm glad we can share in our enthusiasm for mechanical keyboards. I hope I was able to repay a shred of your team's efforts to LZ by refactoring the run spreadsheet. To Georgia and Lindsey, it was a pleasure to meet you both and have the chance to mentor you in South Dakota. You both show so much promise as new graduate students, and I know your future advisors will be lucky to have you as students. Nicolas, thank you for going out of your way to check in on me from time to time. Your sense of fashion is iconic, as is the level of care you have for the students of LZ. Your sense of hospitality is also difficult to top, and I appreciate you, Cosmo, and your roommates for welcoming me into your home. Simran and Issy, thank you for the enlightening conversations about biscuits to go with coffee during my trip to London. Aiham, I'm sorry we only briefly interacted, but I'm sure we'll have more opportunities in the future. Josh, I'll never forget the time you told me "Don't quote me in your acknowledgements, I don't deserve it." Your humility is unmatched and I will always look up to you for that. Makayla, I felt like we immediately became friends at the collaboration meeting in South Dakota. One of these days, I'll stop by Santa Barbara so I won't have to rely on you for pictures of the sunset down there. Harvey, thanks for being so enthusiastic about VSCode and my efforts to make work in LZ less painful. I eagerly look forward to the new heights and depths you continue to (literally) reach. To the underground team, especially Alex, Julia, Doug, and Gavin: it was a blast (not the mine shaft kind) working with you and building the cabinets for the control

room. I'm glad I was able to leave such a lasting impact on the hardware side of LZ. Lastly, thanks to the S2-only group, especially Greg, Dan, Kelsey, and Sparshita for your efforts and feedback.

At UMD, I'd like to thank Anwar Bhatti for his efforts to get me started on analysis when I barely knew how to use a terminal. I'd also like to thank Jon Balajthy for hosting me when I visited UMD for the first time. Thanks also to John Armstrong for taking on the sampling system, and to Tim Edberg for hosting poker nights and always being a friendly presence in the group. To John Silk: thanks for being a great friend while I was at UMD, a great officemate, and letting me siphon coffee from the Moccamaster, as well as your deep knowledge of RGAs. I'm especially thankful to Chris Lobb, Fred Wellstood, and Sudeep "Serenity Now" Dutta for supervising me in the production of transmon qubits during the summer of 2018. To Dave Buehrle, Matt Severson, and the rest of the UMD Physics Education Research Group: teaching Physics 131/132 was an incredibly rewarding experience. The skills I developed have been invaluable in my growth as a mentor. I also want to acknowledge the friends I made at UMD who helped me navigate my first few years in the program. Ramsey and Milena, thank you for helping me feel at home in the astronomy department's parties and events. Even when things seemed like they were falling apart, you both always made me feel welcome and included. I'm grateful that we'll be able to keep in touch, and see where our careers lead us. Luc, I'm so glad we've been able to stay in touch! Discussing data engineering over brunch in Niles is a tradition I would love to keep going. Your perspectives on work outside of academia were some of the first I encountered, and they were pivotal in shaping my current trajectory. To Sophie, Jamie, Donny, Anastasia, Jessie, and Sergio: we all parted ways years ago, for our own reasons. It's a little absurd that UMD didn't work out for any of us, and I'm admittedly still sore about that. Nevertheless, I still find myself caught up in my memories

of sharing lunch and bonding in the PSC lobby with you all. Jamie, I cannot thank you enough for being by my side throughout the pandemic, and shepherding me through the first few years (!) of my transition. I'm sure someday we'll be able to celebrate the energy you selflessly gave to me. Sophie, it brings me so much joy and peace to know our friendship has endured in these intervening years. More than anything else, the discovery I'm most proud of is the one you had the biggest role in: discovering myself. You have my deepest gratitude for that, and so, so much more.

Coming out of the pandemic, I felt incredibly isolated in part because I had left my cohort behind. Fortunately, I was able to find a home again at Berkeley (Go Bears!) through Igenspectrum. For that, I'd like to give a huge thanks to Madge, Olive, and everyone else there for advancing the gay agenda and providing a much-needed safe space for a queer physics graduate student like myself. The rest of the graduate student community in the Berkeley physics department has also been wonderfully open and welcoming. Being able to crash the occasional beer hour was crucially uplifting when I was writing my dissertation, and I feel whole once again knowing that I've started to make new friends (hi Hana!).

To the rest of my friends: I'm so happy you're all a part of my life. Cheyenne, thank you for including me in my very first pride month, powwow, and other celebrations. It's always an adventure when I get to intersect with your circles of friends. Eric, I'm so glad you've been around these past few years. I envy your drive to work towards a better world, and I'm grateful that you and Becky have been so close by for lunch, coffee, and support. Radha, your indomitable spirit is second-to-none. I'm always taking notes on how you're breaking cycles and blazing your own path forward. The support and perspectives you shared with me throughout the pandemic were a beacon of light in an otherwise difficult and isolating time. Lauren: staying with you, Dylan,

and Lorraine was the perfect way to kick off life after my PhD. I'm not bothered being a continent away, because I know we'll have the rest of our lives to make up for the distance. I could say the same for you, Marion. I'm so happy that in the span of a week in Amsterdam, we closed a gap that was years apart. Nick, if I manage to get a hold of a paper copy of this dissertation I'll make sure to send it over to you, as long as you'll pay for the shipping. To John and the rest of the game night crew: every opportunity we had to get together was a breath of fresh air. I'm glad I could always count on playing some great games in exchange for bringing some great snacks. Reina, Nour, and Rimma: sharing a meal with you on my trips to San Diego has always been something I've looked forward to when I needed to get away from work at home. Because of you, I cherish every visit to Convoy Street more than the last. To the rest of my friends from San Diego: Zev, Luz, Bryce, Quinn, Daniel, Andy, Katrina, and so many more! Each one of you has had an impact on me throughout my PhD. You've kept me grounded in a way that physics hasn't, and I'm happy for every chance I get to reconnect.

There are too many Mizrachis to list all of you here, so for those of you not mentioned, please know I appreciate how you contribute to our vast family. Howard, Pam, Naomi, and Jeff: I'm grateful that we've been more in touch these past few years. I felt your encouragement every time we've been around each other. I'll be sure to look for more opportunities to take a detour to the east coast, or just come and visit outright. Emily and Michael, thank you for being so accommodating on my trip after my defense. Kelev is such a good boy and I'm happy I got to help take care of him. Noah and Kat, Gabe and Sarah, Dan, Alissa, and Brooke: I'm glad I got to see each of you and have you as my guides for New York. I can't wait for the next time we get together for Shabbat. Tia Laura, gracias por cada cena de Shabbat, hasta ahora y siempre. Yosef, Yael, and Grace: thank

you so much for always having my back, from growing up together to making it through graduate school. Tammy, thank you for being one of my strongest allies and always offering me beauty tips; I'm glad I finally have the free time to try some of them. Tio Ralphie and Tio Victor: thank you for all of your intriguing questions about dark matter, and giving me the chance to explain my work. Tio Jamie, Diane, and Tia Shirley (in memoriam): every visit to Panama, to see you, Eli, Leo, Danny, Jessica, and the rest of the cousins has been wonderful and memorable because of your hospitality. Whether for a Seder, Rosh Hashanah, a wedding, or another occasion, a trip there has always been a welcome respite.

Dan and Shelley, thank you so much for always welcoming me into your home on my visits and stays near Stanford. Dan, I love being able to talk about my work with you and hear your stories of growing up with Oma and Grampa. Both you and Shelley have helped me redefine my idea of what my relationships towards my family and career will look like. Ava and Adele, I'm always excited to see either of you when I get the chance. You both show so much promise in your endeavors, and I can't wait to see where they lead you. Dave: thank you for hosting me in SLO, showing me your latest plane builds, explaining HAM radio to me, and taking an interest in my work. Sue, Robert, Jimmy and Sally (in memoriam): I'm also looking forward to seeing you all in the future, whether I'm in Tucson, Washington, or Switzerland!

Finally, I'd like to thank my family. Mom, thank you for making sure I was fed, clothed, and sheltered. Dad, thanks for all of your enthusiastic support over the years. You always gave me the encouragement and financial support I needed to push myself to greater heights. Lori, thank you for all of your thoughtful care, and willingness to talk. Every trip home is so much brighter because of you, and it's been a gift to spend time with the rest of your family as well. Ben, I'm so

glad we got to grow up together for a few years. You and Meghan are always a blast to be around, and I'm looking forward to shredding some slopes with you both in the future. Nathan, Caitlin, Lilah, Clyde and Jasper: my love for you is uncountably infinite. You've made Portland my second home, where there's no greater joy than planning our next meal over the one we're having. I have so much to look forward to in all of our lives, and the role I'll be able to play in them as a sister and an aunt.

# Table of contents

<b>Acknowledgements</b>	<b>ii</b>
<b>1. Introduction to Dark Matter</b>	<b>1</b>
1.1. Evidence for Dark Matter . . . . .	2
1.1.1. Mass Estimates of Galactic Clusters . . . . .	2
1.1.2. Galactic Rotation Curves . . . . .	3
1.1.3. Cosmic Microwave Background (CMB) . . . . .	5
1.1.4. Big Bang Nucleosynthesis (BBN) . . . . .	6
1.1.5. The Bullet Cluster Merger . . . . .	7
1.1.6. Structure Formation . . . . .	10
1.2. Dark Matter Candidates . . . . .	11
1.2.1. Standard Model Neutrinos . . . . .	11
1.2.2. Axions . . . . .	12
1.2.3. Weakly Interacting Massive Particles (WIMPs) . . . . .	13
1.2.4. Sub-GeV Dark Matter . . . . .	17
1.3. Detection Methods . . . . .	20
1.3.1. Production . . . . .	21
1.3.2. Indirect Detection . . . . .	21
1.3.3. Direct Detection . . . . .	22
<b>2. Dual-Phase Noble Element Time Projection Chambers</b>	<b>33</b>
2.1. Theory of Operation . . . . .	33
2.2. Signal Production . . . . .	36
2.2.1. Excitation . . . . .	36
2.2.2. Recombination . . . . .	37
2.2.3. Scintillation and Electroluminescence . . . . .	39
2.2.4. Recoil Discrimination . . . . .	40
2.2.5. Recoil Energy Determination . . . . .	41
2.3. Lowering Detection Thresholds in Xenon . . . . .	45
2.3.1. CENNS . . . . .	48
2.3.2. S2 Only Dark Matter Searches . . . . .	50
2.4. Low Energy Electron Backgrounds . . . . .	52
2.4.1. Nomenclature . . . . .	54
2.4.2. Overview of Photoionization, Grid Emission, and Electron Trains . . . . .	56

<b>3. Characteristics of Delayed Electron Emission in the LZ Experiment</b>	<b>67</b>
3.1. The LZ Experiment . . . . .	67
3.1.1. First Dark Matter Search Results . . . . .	68
3.2. Measuring Pulse Rates with Deadtime . . . . .	75
3.2.1. Live, Dead, and Wall Time . . . . .	77
3.2.2. Factors Affecting Deadtime . . . . .	79
3.2.3. Pulse Rates vs. Time . . . . .	81
3.2.4. Repeated Trials . . . . .	84
3.3. LZ Data Extraction and Processing . . . . .	85
3.3.1. Anatomy of an Electron Train . . . . .	87
3.3.2. ALPACA Data Extraction Algorithm . . . . .	89
3.3.3. Livetime and RQ Processing . . . . .	94
3.4. Electron Train Behavior in LZ . . . . .	98
3.4.1. Position and Time Dependence of SEs . . . . .	102
3.4.2. Progenitor Size, Drift Time, and Electron Lifetime Dependence . . . . .	108
3.4.3. Electron Loss Normalization . . . . .	112
3.5. Evidence of a Drift Field Dependence . . . . .	116
3.5.1. Electron Train Origins . . . . .	119
3.5.2. Dependence on TPC Region . . . . .	121
3.5.3. Previous Extraction Region and Drift Field Studies . . . . .	123
3.5.4. Rates in TPC Regions of LZ . . . . .	126
3.5.5. Fits to Rates . . . . .	129
3.5.6. Varying the Extraction Field . . . . .	135
3.5.7. Varying the Drift Field . . . . .	140
3.5.8. Conclusions . . . . .	142
<b>4. Experimental Studies with a Compact Xenon TPC</b>	<b>145</b>
4.1. The XeNeu Detector . . . . .	146
4.1.1. Central Detector Components . . . . .	146
4.1.2. Detector Top Flange . . . . .	154
4.1.3. Detector Cart . . . . .	161
4.1.4. Circulation Cart . . . . .	167
4.1.5. Data Acquisition and Slow Control . . . . .	170
4.2. Physics Measurements with XeNeu . . . . .	171
4.2.1. Measurement of High Energy Nuclear Recoils . . . . .	171
4.2.2. Search for the Migdal Effect . . . . .	173
4.3. Plastic Reduction Efforts in XeNeu . . . . .	178
4.3.1. Overview of Outgassing . . . . .	179
4.3.2. Design of a Plastic-Free Bottom PMT Holder . . . . .	182
4.3.3. Single Electron Trigger Efficiency . . . . .	187
4.3.4. Electron Background Rates . . . . .	191
4.4. S2 Detection Improvements with Silicon Photomultipliers . . . . .	202
4.4.1. Introduction to SiPMs . . . . .	203
4.4.2. Replacement SiPM Array Design . . . . .	206

4.4.3. Electronics Re-cabling . . . . .	207
4.4.4. SPE Gain Calibration . . . . .	218
4.4.5. Position Calibration . . . . .	222
<b>5. Towards Lower Detection Thresholds with Xenon-Doped Argon TPCs</b>	<b>232</b>
5.1. Benefits and Challenges of Xenon Doped Argon TPCs . . . . .	233
5.2. Operating Conditions for a Xenon-Doped Liquid Argon Mixture . . . . .	237
5.2.1. Eutectic Point Estimate . . . . .	238
5.2.2. Predicted Xenon Gas Concentrations . . . . .	243
5.3. CHILLAX Operational Concept and Thermosyphons . . . . .	249
5.3.1. Theory of Operation . . . . .	250
5.3.2. Thermosyphon Design . . . . .	252
5.4. Evaluation of Hamamatsu VUV4 SiPM Liquid Argon Scintillation Photon Detection Efficiency . . . . .	259
5.5. The CHILLAX Gas Sampling System . . . . .	264
5.5.1. Introduction to Sampling Systems . . . . .	265
5.5.2. Design and Data Acquisition . . . . .	271
5.5.3. Commissioning and Calibration . . . . .	277
5.6. Demonstration of High Concentration Xenon Doping in CHILLAX . . . . .	295
<b>Appendices</b>	<b>305</b>
<b>A. Information on Stripping Kapton Cables</b>	<b>305</b>
A.1. Cable Assembly Procedures . . . . .	310
A.1.1. Sub-C Connector . . . . .	310
A.1.2. MMCX Connector . . . . .	311
A.1.3. D-sub Connectors . . . . .	314
A.1.4. Ring Terminal (Bottom PMT) . . . . .	317
A.1.5. Sub-C Connector (FEP Cable) . . . . .	318
<b>B. Supplementary Information for XeNeu</b>	<b>320</b>
B.1. Table of Slow Control Sensors and Instruments . . . . .	320
B.2. Lessons Learned . . . . .	321
B.2.1. The XeNeu Vacuum Box . . . . .	321
B.2.2. CAD Software and Chain-of-Custody . . . . .	323
<b>C. THERANOS Procedures</b>	<b>324</b>
C.1. RGA Stand P&ID . . . . .	324
C.2. XeNeu P&ID . . . . .	324
C.3. RGA Calibration Procedure . . . . .	324
C.3.1. Hardware Connections . . . . .	324
C.3.2. Pressure Guidelines . . . . .	325
C.4. Initial Vacuum . . . . .	327
C.4.1. XeNeu Pumpout . . . . .	327

C.4.2. RGA Stand Pumpout . . . . .	327
C.5. Gas Mixing . . . . .	328
C.5.1. XeNeu Gas Panel Setup . . . . .	328
C.5.2. Xenon Addition to MV1 . . . . .	329
C.5.3. First MV1 Expansion to MV2 . . . . .	329
C.5.4. Second MV1 Expansion to MV2 and Argon Addition . . . . .	330
C.6. Sampling . . . . .	331
C.6.1. MV2 Pressure Reduction to 10 torr . . . . .	331
C.6.2. Priming RGA and LV1: MV2 Pressure Reduction to 1.5 torr . . . . .	332
C.6.3. Finishing a Scan . . . . .	332
C.6.4. Finishing for the Day . . . . .	333
C.7. Calculations . . . . .	333
C.7.1. Dilution Factor . . . . .	333
C.7.2. Expected Concentration - Double Expansion . . . . .	333
C.7.3. Expected Concentration - Drain Refill . . . . .	334
C.7.4. Calibration Factor . . . . .	334
C.8. Recording Data . . . . .	335
C.8.1. Recorded Parameters . . . . .	335
C.8.2. Rough Calculations . . . . .	335
C.9. RGA Software Configuration and Setup . . . . .	335
C.9.1. Scan Speed . . . . .	335
C.9.2. Electron Multiplier (CEM or CDEM) Operation . . . . .	336
C.9.3. Pressure vs. Time Scans . . . . .	337
C.9.4. Analog Scans . . . . .	337
C.10. Using the Keithley 6510 for Pressure Readout . . . . .	338
C.10.1. Pressure Sensors . . . . .	338
C.10.2. Key Front Panel Features . . . . .	339
C.10.3. Real-time Readout . . . . .	340
C.10.4. Channel Scan . . . . .	340
C.10.5. Scan Options . . . . .	341
C.10.6. Saving Data from a Channel Scan . . . . .	341
C.11. Parsing Scan Data with <code>rgaplot</code> . . . . .	341
C.12. Summary of Recommendations for Future Development of THERANOS . . . . .	341
<b>D. Historical Review of Electron Trains</b>	<b>344</b>
D.1. 2008 - ZEPLIN-II . . . . .	344
D.2. 2009 - Burenkov <i>et al.</i> . . . . .	348
D.3. 2011 - ZEPLIN-III, XENON10 . . . . .	349
D.3.1. XENON10 . . . . .	352
D.4. 2012 - Akimov <i>et al.</i> . . . . .	352
D.5. 2014 - XENON100 . . . . .	353
D.5.1. Photoionization Rates . . . . .	354
D.5.2. Rate Dependence on Purity and S2 Size, Uncorrelated Rates . . . . .	356

D.6.	2016 - Akimov <i>et al.</i> (RED-1)	358
D.6.1.	“S3” Pulses	358
D.6.2.	Trapped SEs After Muons (Power Law)	359
D.7.	2017 - Sorensen	361
D.7.1.	Trapped Electron Model	362
D.8.	2018 - Sorensen and Kamdin	363
D.8.1.	Trapped Electrons?	365
D.9.	2020 - Akimov <i>et al.</i>	367
D.10.	2020 - LUX	368
D.10.1.	Photoionization - Neutral Impurities	369
D.10.2.	Electron Bursts	370
D.10.3.	Delayed Backgrounds	372
D.10.4.	Photon Trains	375
D.10.5.	Grids - Multiple Electrons	375
D.11.	2021 - PIXeY	376
D.11.1.	Pre-S1	377
D.11.2.	S1 Photoionization	378
D.11.3.	Post S2 and Post Drift	378
D.12.	2021 - ASTERiX	380
D.12.1.	Extraction Field	382
D.12.2.	Drift Field	383
D.12.3.	Drift Time	385
D.12.4.	Infrared Photoionization of Impurities	385
D.12.5.	Conclusions - Liquid Surface Electrons?	385
D.12.6.	Vibration Tests of ASTERiX	387
D.13.	2021 - XENON1T	387
D.13.1.	Position Dependence	388
D.13.2.	Progenitor Size	389
D.13.3.	Delay Time Since Progenitor	391
D.13.4.	Extraction Field	391
D.13.5.	Drift Time	392
D.13.6.	Purity	392
D.13.7.	Conclusion - Liquid Impurities	393
<b>E.</b>	<b>Repositories and Resources</b>	<b>394</b>
E.1.	Repositories	394
E.1.1.	This repository	395
E.1.2.	LZ Gitlab	395
E.1.3.	LLNL Gitlab	395
E.2.	Software	396
E.2.1.	CAD	396
E.2.2.	Zotero	396
E.2.3.	VSCode	396
E.2.4.	Inkscape	399

E.2.5. Notion . . . . .	399
E.2.6. Draw.io . . . . .	399
E.2.7. Python . . . . .	399

<b>References</b>	<b>402</b>
-------------------	------------

## List of Tables

1.1.	Nomenclature for Equation 1.8 . . . . .	24
2.1.	Light emission properties of liquid xenon and argon. Overall scintillation response times are dependent on the proportion of states created with excitation and recombination, which are themselves dependent on the concentration of impurities, initial interactions, and applied electric field. . . . .	39
2.2.	Table of electron train power law exponents from previous studies. . . . .	62
3.1.	Table of child pulse rate normalization factors. . . . .	102
3.2.	Definitions of position-correlated and uncorrelated child pulses. . . . .	103
3.3.	Coordinates used to determine bounds for averaging fields within the TPC. Regions were chosen by eye to be representative of values in the bulk of each region. . . . .	118
3.4.	Grid voltages and corresponding nominal field values from simulations. For some groups the cathode voltage difference was only <i>nominally</i> 28 kV, leading to slightly varied drift fields of 197, 194, and 191 kV/cm when the gate voltage was adjusted. For the sake of this analysis, these three drift fields are grouped together as nominally 194 kV/cm. . . . .	118
3.5.	Table of runs, grouped by voltage settings. Tuples indicate consecutive sets of runs e.g (9199, 9201) corresponds to 9199, 9200, 9201. LZ Users: see Fields Investigation Data spreadsheet for more information and <code>sv_SEDecayTime_lr_EdV_7_8_9_DdV_16_28_35.list</code> in ALPACA module [89] <code>inputs</code> . . . . .	119
3.6.	Dependencies of fast and slow regions in a summed waveform of alpha events [99]. As an example, the amplitude of the fast component increased at a higher purity and decreased with an increasing applied electric field. . . . .	124
3.7.	Table of electron train power law exponents from previous studies. . . . .	125
3.8.	Table of values shown in Figure 3.48. A 15% systematic uncertainty has been added to the errors for data taken at the 3900 and 3410 kV/cm (7 and 8 kV) fields to compensate for uncertainties introduced in the choice of fit regions and cuts. Reported field values for this work were lowered slightly from those calculated in Table 3.4 to account for a lower potential on the gate grid as a result of deflection between the gate and anode grids. . . . .	144

4.1.	DataFrame.head() of merged trigger and pulse dataframe. The join is performed on event_id and event_timestamp which uniquely identifies events. peak_time is the peak time of a trigger signal relative to a main event trigger, while npe, start_time, and end_time are quantities belonging to a pulse in the same event. The first three rows correspond to pairings of three different pulses with the same trigger pulse. Subsequent steps detailed in the text explain how distinct pulses with triggers are selected. . . . .	190
4.2.	Key operating parameters for data taken during Runs 11 and 12. Run 11 had a slightly smaller gas gap, leading to a higher field. . . . .	192
4.3.	Fit values and uncertainty for each value shown in Figure 4.34. Uncertainties are statistical only. . . . .	199
4.4.	Fit values from Equation 4.2 applied to each rate curve shown in Figure 4.35. Fits were performed in the range of 40 ms-1 s. Uncertainties are statistical only. Because the fit does not account for the presence of a possible decaying exponential, the magnitudes of $\beta$ may be elevated by ~10%. There is an additional ~10% systematic error on $\beta$ in the Run 11 rates that stems from the choice of the left edge of the fit. . . . .	200
5.1.	Gas flow regimes as defined by their Knudsen number. . . . .	265
5.2.	Selected sampling system components . . . . .	273
5.3.	Tables of pumps studied for gas recirculation at low pressures. Helium leak rates were independently checked and generally found to be in agreement with quoted leak rates. Some pumps were modified from their stock configuration in an attempt to lower leak rates. Of the three modified pumps, only the ACP-15's modification produced a substantial improvement by eliminating detectable leaks through the ballast port. . . . .	289
A.1.	Strippers tested on AWG 30 Kapton coaxial cable. Note that stripper 6 has a continuously adjustable blade which accommodates wires up to 0.125" in diameter. Stripper 4 has the same operating mechanism as stripper 5, but uses custom components for gripping and cutting Kapton insulation. . . . .	306
B.1.	Table of XeNeu Sensors and Instruments . . . . .	320
C.2.	RGA scan speed table from SRS [186]. Note how the scan speed and noise floor are inversely related. . . . .	336
C.4.	Possible testing matrix. Rows are xenon concentrations in ppm. . . . .	342
D.1.	Dependencies of fast and slow ROIs in a summed waveform of alpha events observed by [99]. As an example, the amplitude of the fast component increased at a higher purity and decreased with an increasing applied electric field. . . . .	365

## List of Figures

1.1.	Rogstad and Shoshak compared observed visible densities of galaxies with their rotational velocities. A few years later, Albert Bosma demonstrated the flat rotation curves of 25 different galaxies which added to the consensus surrounding dark matter at galactic scales. . . . .	4
1.2.	2018 Angular power spectrum constructed from Planck CMB data [14]. The first three peaks provide information on the critical density, baryon density, and dark matter density respectively. The blue line is a fit from dark energy-dark matter models, and the red points are observed data. The dotted gray line marks a change in the scale of the x-axis. The CMB is extremely uniform, with temperature fluctuations $\frac{\Delta T}{T} \sim 10^{-5}$ only visible at small scales. . . . .	6
1.3.	Primordial abundances of light elements predicted by BBN, from [15]. Yellow boxes correspond to observed abundances. The narrow vertical band corresponds to constraints from CMB measurements, while the wider band is a limit set by observations of deuterium and $^4\text{He}$ abundances. Relative abundances of $^3\text{He}$ have yet to be reliably measured, and the disagreement between the measured abundance of $^7\text{Li}$ and that predicted by CMB measurements and standard BBN is known as the cosmological lithium problem. . . . .	8
1.4.	A comparison of optical and x-ray images confirmed that galaxies had passed through each other relatively unimpeded, while the ICM had lagged behind due to higher interaction rates. Additional microlensing calculations resulted in contours (green lines) which increased towards the centers of the galactic distributions [18]. . . . .	9
1.5.	A simplified snapshot of constraints on the axion-photon coupling strength [36]. Other couplings are possible (e.g. axion-electron) and plots for those couplings can be found at the same source. . . . .	13
1.6.	A numerical solution of Equation 1.5 demonstrating the freeze-out process where $Y = n_\chi/s$ and $x = m_\chi/T$ . Solid curves approach a relic density as the temperature of the universe $T$ decreases. An increasing dark matter annihilation rate $\Gamma$ (indicative of a stronger SM coupling) lowers this relic density. The dashed curve tracks the abundance of dark matter particles that remain in thermal equilibrium. From [40]. . . . .	16
1.7.	A numerical solution of the Boltzmann equation demonstrating the freeze-in process where $Y = n_\chi/s$ and $x = m_\sigma/T$ . Solid curves approach a relic density as the temperature of the universe $T$ decreases. An increasing dark matter production rate $\Gamma$ increases the relic density. The dashed curve tracks the abundance of SM particles that remain in thermal equilibrium. From [40]. . . . .	19

1.8.	Detection methods for particle dark matter rely on three distinct classes of particle interaction; each are denoted with a corresponding arrow which signifies the flow of time for that process [46]. . . . .	20
1.9.	Observations of specific processes at the LHC can help set limits on various dark matter models, and can be compared to results from direct detection experiments. Furthermore, the model-independent nature of collisions grants the ability to sweep larger mass ranges of parameter space in comparison to more targeted searches [48] .	22
1.10.	Fermi-LAT limits on the dark matter annihilation cross-section from a gamma ray survey of 15 dwarf spheroidal galaxies. Limits are set by searching for gamma rays from the decay of products of specific DM annihilation channels. Observed spectra are compared to simulated spectra, which are swept over a range of DM masses that roughly correspond to WIMP models [52]. The thermal relic cross-section is shown as a dashed gray line for the sake of comparison. . . . .	23
1.11.	Prominent dark matter direct detection experiments, grouped by the detection technology used and measured signals. Single phase detectors and dual-phase time projection chambers (TPCs) utilize noble elements such as xenon or argon. Measuring exclusively charge signals, DAMIC uses silicon wafers while NEWS-G uses a mixture of neon and methane gas. Gas TPCs are constructed to search for a directional WIMP “wind”. Scintillating bolometers, cryogenic detectors, and scintillating crystals all utilize crystals of various compositions to probe different WIMP mass ranges. Each experiment under these categories has unique hardware suitable for reading out the desired quantities. PICO and DRIFT are notable for using fluoridated compounds to set the best spin-dependent WIMP interaction limits. Figure from [54]. . . . .	24
1.12.	The above plot shows simulated scattering rates for a 100 GeV WIMP with a spin-independent cross section of $10^{-45}$ cm <sup>2</sup> and the kinematic properties of a typical thermal relic. At lower nuclear recoil energies, xenon is an excellent choice of target material due to its high mass. However at higher nuclear recoil energies, scattering becomes incoherent as a WIMP interaction would only probe a portion of the xenon nucleus. Thus, lower recoil thresholds enable detectors to benefit from coherent scattering amplification of the nuclear interaction cross-section. Circles on each curve indicate typical detection thresholds. Taken from [58]. . . . .	28
1.13.	Differential ionization rates versus electron recoil energy for different target materials [59]. Notable here is that for events below a few hundred eV, the irreducible solar neutrino background is not expected to meaningfully obscure signals. . . . .	30
1.14.	Recent limits set by WIMP direct detection experiments [60]. The neutrino fog is where an irreducible background caused by coherent scattering of solar, atmospheric, and supernova neutrinos (Section 2.3.1) is expected to become dominant. Searches in this region require a $10^n$ increase in exposure for in order to obtain a corresponding factor of 10 increase in cross-section sensitivity. . . . .	31
1.15.	A chart highlighting the increasing levels of sensitivity reached by contemporary WIMP search experiments. Different colors indicate different technologies used. Note that some experiments pictured are more sensitive to WIMPs in a significantly lower mass range. See [61] for the figure shown here and [62] for a previous version. . . . .	32

2.1.	Mockup of a TPC and event with S1 and S2 pulses. A particle interaction produces light ( $\gamma$ ) which is detected as the S1 pulse at time $t = 0$ in an event window. Some electrons are also ionized in the initial interaction; an external electric field $\vec{E}$ prevents some of them from recombining by drifting them upwards. After some drift time $\Delta t$ , the electron cloud reaches the liquid surface where the electric field extracts them from the liquid. Upon extraction, the electrons generate light via electroluminescence in the gas, which is detected as the S2 pulse in an event window. Original image courtesy of C. Faham and D. Malling. . . . .	34
2.2.	Comparison of a simple two-grid TPC (Figure 2.2a) and a three-grid TPC (Figure 2.2b). The three-grid TPC has independently tunable drift and extraction fields. These are more suitable for TPCs with larger drift lengths where a high extraction field may be otherwise difficult to maintain. . . . .	35
2.3.	The liquid-gas interface in a dual-phase TPC forms a potential barrier that prevents electrons from being extracted. The extraction field modifies the shape of the potential barrier and increases the energy of drifting electrons so they are able to overcome it. This results in a higher electron extraction efficiency. . . . .	45
2.4.	Overview of liquid xenon nuclear recoil charge and light yield measurements from [74] at nominally similar electric field strengths. Charge yield results from [73] are shown in the bottom figure as purple triangles. . . . .	47
2.5.	Predicted CENNS rates for a detector with a 10 m standoff from two sources: a 3 GW <sub>th</sub> nuclear reactor, and the ISIS neutron spallation source. In general, scattering rates for a neutron spallation source are lower due to the smaller neutrino flux, but are detectable at higher threshold energies because they are produced at higher energies than in a nuclear reactor [58]. . . . .	49
2.6.	Event rate as a function of electrons produced by dark matter-electron scatters in xenon [78]. Colored spectra correspond to contributions from individual electron shells and the gray band represents the overall spectrum when varying the effect of higher-order processes. Left axes on the plot correspond to the maximum allowable cross-sections for benchmark models, while right axes are cross-sections constrained by the correct relic abundance. . . . .	51
2.7.	Low energy S2 spectrum in the 2019 XENON1T light dark matter search [79]. The analysis threshold was cut off at the vertical dotted line where events below $\sim 4$ -5 electrons in size were not considered due to the presence of “unmodeled backgrounds”. In a 2021 reanalysis (Section D.13), an empirical model of electron trains was constructed which enabled limits to be set on light dark matter candidates based on the observed single-to-few electron background rate. Additional backgrounds in this plot include the expected CENNS rate (shown in red) from $^8\text{B}$ solar neutrinos and a flat ER background from $^{214}\text{Pb}$ decays (blue). The expected signal model from 4 GeV and 20 GeV spin-independent DM models are also shown in orange and purple respectively. . . . .	53

2.8.	Hardware-related electron background signals in dual-phase xenon TPCs can be grouped by the timescales in which they occur relative to an event. Each relates to unique mechanisms and sources, some of which are speculative. Dependencies are factors which have been investigated in studies of the respective backgrounds. Solid lines indicate a strong connection has been observed. Dotted lines indicate a weak dependence, no dependence, or a suspected dependence that requires further study. . . . .	55
2.9.	Distinct timescales in a single scatter event where low energy electron backgrounds are most prominent. Prior to an S1, low energy backgrounds are typically dominated by spontaneous pulses (dashed outline) which are not associated with a event in the TPC. These are often due to grid emission. After the S1 and S2, prompt photoionization backgrounds (solid outline and fill) are most common up to the maximum drift time in a TPC $t_{drift}$ . As both the S1 and S2 can act as progenitors, they each have their own timeframes for prompt backgrounds. Note also that prompt backgrounds can be intense enough that single electron (SE) pulses often combine and pile-up to form smaller S2 pulses; this is shown after the progenitor S2. After prompt backgrounds from the S2 pulse have died off, electron trains in the form of delayed SEs and small S2s (solid outline and no fill) are still visible. It has been shown that delayed small S2s are not well-explained by SE pile-up [84]. . . . .	57
3.1.	The TPC (1) sits at the center of the LZ experiment. Segmented acrylic tanks (2), shown in green, form the OD which encircles the titanium outer vessel. 8" PMTs (3) are mounted on a frame to detect light from the OD and water shielding. The frame is lined with Tyvek to improve light collection. Water shielding (4) is contained by the stainless steel water tank. High voltage for the TPC grids is provided to the TPC via a feedthrough (5) which connects to the cathode. Neutron calibrations are conducted using two conduits (only one shown, 6) which provide a clear path to the TPC. Figure taken from [90]. . . . .	69
3.2.	Calibration events plotted in $\log_{10}(S2c) - S1c$ space. $CH_3T$ events are represented by dark blue points, and the median of their simulated distribution is represented by the solid blue line. The dotted blue lines represent 10% and 90% quantiles. DD events are represented by orange points and the same properties of their simulated distributions are shown with red lines. Contours of constant NR and electron-equivalent recoil energies are shown in gray. From [92]. . . . .	70
3.3.	Map of the TPC in reconstructed depth and radial coordinates. The dotted black line represents the maximum extent of the active volume, which is slightly distorted because of electric field nonuniformities near the walls. The solid black line represents the fiducial cut, which is used to reject events which are too close to the walls. Events marked with gray points fail the fiducial cut, while events with black points pass it. Red crosses and blue circles represent events which were rejected by skin or OD veto detectors, respectively. From [92]. . . . .	72
3.4.	Data from Figure 3.3 plotted in $\log_{10}(S2c) - S1c$ space. Events in the WIMP ROI are fairly sparse. From [92]. . . . .	73
3.5.	SR1 ER energy spectrum passing all cuts, superimposed with background estimates. From [92]. . . . .	74

3.6. Limits on the spin-independent cross-section for WIMPs with masses between $10^1$ and $10^4$ GeV. Results from DEAP-3600, LUX, XENON1T, and PandaX-4T are also shown. From [92]. . . . .	74
3.7. An LZ event after an S2 pulse; the S2 pulse is not shown. Single electron (SE) pulses are marked with red dots. The taller pulses at 4200 and 6600 $\mu$ s are multiphoton-electron (MPE) pulses, which are ignored here. The time axis in the event window is relative to the <i>trigger</i> , which was fired at random; $t=0$ on this axis is when the trigger was fired. In this trigger configuration, 1 ms of data was captured before the trigger and 11 ms was captured after. An additional axis is shown at the bottom denoting time since the end of the S2 pulse. . . . .	82
3.8. Counts of single electron pulses between 15 and 25 ms after an S2 pulse. Bins are 1 ms wide. the last two SE pulses from the event shown in Figure 3.7 are captured in this histogram. . . . .	83
3.9. Livetime between 15 and 25 ms after an S2 pulse. Bins are 1 ms wide. . . . .	83
3.10. Pulse count and livetime histograms for a progenitor S2 different from the one shown in Figure 3.7. The same period of time since the progenitor S2 is shown as in Figure 3.8 and Figure 3.9. . . . .	85
3.11. Pulse count and livetime histograms for pulses following thousands of different progenitors. Counts of single electrons pulses appear to decrease exponentially more than 100 ms after a progenitor (Figure 3.11a). This decrease is mostly an artificial result of livetime <i>also</i> decreasing exponentially (Figure 3.11a), and should serve as a motivating example for livetime corrections. . . . .	86
3.12. Overview of data extraction and processing structure . . . . .	87
3.13. Mockup of an electron train captured with a random trigger. The sequence of events 1 and 2 illustrate how a trigger holdoff prevents events from overlapping as triggers are blocked until the holdoff expires. Livetime is only within the bounds of event windows and is further limited to the time between vetoing pulses (pink triangles). This is because the trigger efficiency for any pulse outside of an event is manifestly zero. If a pulse-based trigger is used, then livetime includes time inside of events between vetoing pulses <i>and</i> outside of events when a trigger <i>could</i> have occurred. This is shown with a dotted line extending the rightmost livetime arrow. . . . .	88
3.14. Flow chart illustrating the event loop of the BigDEB algorithm. . . . .	90
3.15. Flow chart illustrating the pulse loop of the BigDEB algorithm. . . . .	91
3.16. General overview of steps in livetime and RQ post processing. . . . .	95
3.17. The first two events of an electron train in LZ, showing the presence of SEs for multiple milliseconds after an S2 pulse. . . . .	99
3.18. PMT hitmaps from the first and second events shown in Figure 3.17. Hitmaps on the left correspond to the top PMT array, while hitmaps on the right correspond to the bottom PMT array. . . . .	100
3.19. The rate of SEs after S2s, per square centimeter of liquid surface in the TPC. The x-axis is the radial distance between between the child SE pulse and the most recent progenitor S2 pulse. The stack axis denotes the time since the end of the progenitor S2. . . . .	103

3.20. Diagram showing geometric constraints needed to correct for the lack of pulses observed relative to a progenitor when pulses are observed at a radius $R_{SE}$ which extends beyond the walls of the TPC. A child pulse may be observed anywhere inside the TPC along the circle defined by $R_{SE}$ . Counts of pulses along this arc are then used to correct for the lack of pulses observed along the arc outside of the TPC.	105
3.21. Rates of position-correlated and uncorrelated SE pulses versus time since the end of their progenitor S2. The rates prior to 1 ms are a known artifact of pulse pile-up from photoionization (see also Figure 3.22). Position-uncorrelated rates drop off quickly after the full drift time has elapsed, in comparison to the rates from position-correlated pulses.	107
3.22. Rates of position-correlated SE and S2 pulses versus the size of their respective progenitors. For the smallest progenitors, SE rates appear to decay monotonically. As the progenitor size increases, pulses lose their SE classification, causing a distortion in the rate prior to 1 ms. This feature is not present when selecting for S2 pulses because pulse classification is conserved for S2s undergoing pile-up.	108
3.23. Rates of SEs versus the size of their progenitor, $e_R$ , in the number of electrons extracted. A time window of 3-300 ms was chosen to avoid integrating any residual photoionization. The position-uncorrelated rates do not exhibit a dependence on the size of the progenitor; this behavior was also observed in XENON1T [80]. Most progenitors are between $10^3$ and $10^4$ $e_R$ in size; SR1 data was used in this plot to maximize the available statistics.	109
3.24. Rates of SEs versus the drift time of their progenitor. The position-correlated rate appears to linearly depend on the drift time of the progenitor even after normalizing by $e_I$ , the size of the progenitor at the interaction site. The position-uncorrelated rates do not exhibit a dependence on the drift time of the progenitor. This behavior was also observed in XENON1T [80], and highlighted as an indication that uncorrelated delayed electrons may be correlated with previous progenitors. SR1 data was used in this plot to maximize the available statistics.	110
3.25. Drift time dependence of SE rates versus time since the progenitor S2. The position-uncorrelated rates do not exhibit a clear drift time dependence at any timescale. Position-correlated photoionization rates have a counterintuitive dependence on the drift time. At shorter drift times, S2s should lose fewer electrons and photoionize the TPC more. However, lower photoionization rates are observed in this case, and this phenomenon was also observed in LUX [81]. This behavior remains unexplained, and should be the subject of future work.	112

3.26. Rates of SEs versus the electron lifetime recorded at the time of the progenitor pulse. The rate at higher electron lifetime values appears to taper off, in agreement with behavior observed by XENON1T [80] and their explanation that the impurity or impurities responsible for electron trains does not perfectly track the concentration of impurities which affect the electron lifetime. A sharp drop in the position-correlated rate is apparent after the electron lifetime exceeds 5 ms. This may be an artifact of interpolating electron lifetime data, as explained in the text. The position-uncorrelated rate also exhibits a slight dependence on the electron lifetime, which would be consistent with the explanation from XENON1T [80] that these backgrounds are mostly from prior electron trains. . . . .	113
3.27. Electron lifetime dependence of SE rates versus time since the progenitor S2. A weak dependence on the electron lifetime is visible for both the position-correlated and uncorrelated rates at all timescales after the full drift time. The position-uncorrelated rate in the 2-4 ms electron lifetime bin is noticeably uneven due to wrongly associated pairs of SEs and progenitor S2s. This occurs for all other position-uncorrelated rates in both this figure and Figure 3.25, but the effect is enhanced with lower statistics. . . . .	114
3.28. The same rates from Figure 3.25 and Figure 3.27. instead normalized by $e_L$ , the number of SEs lost by the progenitor as it drifts through the liquid. Both the drift time dependence and electron lifetime dependence that were visible in Figure 3.25 and Figure 3.27 are negated, indicating that impurities which capture electrons in the liquid are primarily responsible for electron trains. . . . .	115
3.29. Simulated field strengths in the extraction region of the LZ TPC. Key features such as the anode and cathode are indicated in the plot. The cathode is not pictured; it is located at $Z = 0$ mm and is the lower boundary for the drift liquid. The TPC walls are located at $R \approx 740$ mm; only the first 50 mm are shown here to highlight the scale of the grid spacing in R. . . . .	117
3.30. Electron trains from progenitors in the gas, drift, and below cathode regions of XENON1T [80]. A 200 ms holdoff was imposed after interactions to avoid overlapping electron trains between progenitors. The SE rate during this 200 ms holdoff period is shown with the solid purple line. It served to verify that the holdoff period was adequate in allowing SE rates to return to a baseline level before SEs from progenitor pulses were tracked. It also showed that rates following progenitors in the gas and reverse field regions returned fairly quickly to the pre-holdoff rates, indicating that electron trains are linked to drift region interactions. . . . .	122
3.31. Power law in SE rate observed by Akimov et al. after muon S2s. . . . .	123
3.32. Low-pass filtered summed waveform of electron trains from alpha events with a “fast” and “slow” exponential component. . . . .	124
3.33. SE rates following single scatter events in the gas above and below the anode, extraction liquid above the gate grid, and drift liquid between the gate and cathode of LZ. Delayed electron rates in the gas and extraction liquid are much less intense than those in the drift liquid. This plot combines the rates from both position-correlated and uncorrelated pulses, which is indicated by No Cut (NC) to the relative pulse positions in the legend. No area normalization is used in this plot in order to enable direct comparisons to Figure 3.30. . . . .	127

3.34. Rate of SEs shown in Figure 3.19, projected in time since the progenitor S2. Normalizing rates by the area subtended by each interval in $\Delta r$ reveals that the flux of position-uncorrelated SEs is fairly uniform for $\Delta r > 20$ cm. . . . .	128
3.35. SE rates from Figure 3.33, split into the position-correlated and uncorrelated components. The position-correlated rates in the gas are largely from drift liquid events which leaked into the gas event selection. The relatively large fluctuations in these rates over time compared to the rates in other regions also indicate mis-associated S2 normalizations. Position-correlated rates in the extraction liquid appear to have a power law component, but with a much lower amplitude than those in the drift liquid. The exponents in these two regions are visually similar, but fits (Section 3.5.5) indicate that the power law in the extraction liquid may be slightly steeper. . . . .	130
3.36. Power law fit to the position-correlated SE rate following events originating in the drift liquid shown in Figure 3.35. . . . .	131
3.37. Power law fit to the position-correlated SE rate following events in the extraction liquid shown in Figure 3.35. A constant term in the fit approximates the presence of position-uncorrelated backgrounds at long timescales, but does not account for residual photoionization backgrounds at short timescales. As a consequence, there is a large systematic uncertainty in $\beta$ associated with the left edge of the fit. . . . .	132
3.38. Subtracting the position-uncorrelated SE rate in Figure 3.35 from the position-correlated rate in Figure 3.35 removes many of the photoionization backgrounds in the position-correlated SE rate between 100us and 1ms. It also removes the background component after 100ms. . . . .	133
3.39. Fitted value of $\beta$ in Figure 3.38 as a function of fit limits. Subtracting the position-uncorrelated rate leaves behind some discontinuities in the rate. First, rates before 1ms are still slightly elevated. Second, rates after $\sim 20$ ms are highly dependent on the intensity of the position-uncorrelated flux, which increases with the field. . . . .	133
3.40. Power law fits to position-correlated rates in Figure 3.25 show minimal variation in the exponent and a clear trend in the amplitude that increases with drift time. . . . .	134
3.41. Rate of position-correlated and uncorrelated single electrons from extraction liquid progenitors, at three different extraction voltages. Rates are normalized by $e_S$ , the size of the progenitor S2 at the liquid surface. . . . .	136
3.42. S2-normalized rate flux of position-correlated single electrons following extraction liquid progenitors, at three different extraction voltages. The background flux from position-uncorrelated electrons has been subtracted, revealing a slightly clearer power law in each dataset. . . . .	136
3.43. S2-normalized rate flux of position-correlated SEs following extraction liquid progenitors, without background subtraction. A power law is fit with the addition of a constant term as an alternative to account for the underlying position-uncorrelated background. The power law exponents are only marginally affected. The left edge of the fit is set to 2ms because of residual photoionization backgrounds at 1ms which are not well-modeled by the constant terms. . . . .	138

3.44.	S2-normalized rate flux of position-correlated single electrons following drift liquid progenitors, at three different extraction voltages. The position-uncorrelated background has been subtracted. The electron lifetime in the 8 and 9kV datasets was identical, and lower than the electron lifetime in the 7kV dataset. This resulted in a steeper slope from photoionization backgrounds between ~1-5ms. Because the position-correlated rates from the drift liquid are so intense, subtracting the position-uncorrelated background has no effect on the features in the rates at short (few ms) and long (few hundred ms) timescales. . . . .	139
3.45.	Extraction field dependence of the delayed electron fraction from drift region electron trains in XENON1T [80]. The delayed electron fraction was calculated by summing the number of electrons observed between 2-200ms after a progenitor S2 in the drift region, then normalizing by the size of the extracted S2. . . . .	140
3.46.	S2-normalized rate flux of position-correlated and uncorrelated single electrons from drift liquid progenitors, at three different drift voltages. . . . .	141
3.47.	S2-normalized rate flux of position-correlated single electrons following drift liquid progenitors, at three different drift voltages. The background flux from position-correlated extraction liquid electrons at the same extraction field has been subtracted. The power laws only appear to be marginally affected. . . . .	142
3.48.	Drift field variation of the electron train power law exponent in past studies, compared to results from this work. The 9kV dataset is excluded due to a combination of low statistics and qualitative similarity to gas events. Data from this figure is reproduced in Table 3.8. . . . .	143
4.1.	XeNeu (left) and its accompanying circulation cart (right). The detector is housed in a cylindrical can which hangs underneath the main outer vacuum insulation system. The circulation cart holds the gas circulation pump, getter, and xenon gas bottles. The gas bottles and getter are on the backside of the cart and not visible in this picture. Note that this picture is not up to date with the P&IDs shown in Figure 4.13 and Figure 4.16. . . . .	147
4.2.	XeNeu detector components. At the bottom of the detector assembly is a large cylindrical PTFE piece which houses the bottom PMT (Figure 4.5a). The field cage (Figure 4.4a) is mounted to the top of the bottom PMT holder. The weir reservoir is machined from a single block of PEEK and sits in a dedicated cutout on the side of the bottom PMT holder. The weir height is independently adjustable with clamps on two of the vertical support rods. Directly above the field cage is the top PMT assembly (Figure 4.6), which clamps to all three support rods. High voltage feedthroughs are welded to the detector top flange, and descend into slots machined into the top of the bottom PMT holder. . . . .	148

4.3.	2D schematic (drawn to scale) of XeNeu showing dimensions of the active volume and labels of key components. Several components are configurable and the subject of this work. The PMT holder has a PTFE variant (Figure 4.5a, pictured here) as well as one which is constructed from low-outgassing materials (Figure 4.23). The top photosensor array originally used PMTs (Figure 4.6), which were replaced with a SiPM array (Figure 4.38, pictured here). The field cage (Figure 4.4) can also be assembled in two configurations, either with a PTFE reflector (pictured here) to improve S1 light collection, or without one to reduce outgassing rates. . . . .	150
4.4.	Field cage components for the XeNeu TPC. The field cage is held together by three nylon screws cut to the appropriate length such that they fully engage with tapped holes in the base ring, which is made of PEEK. . . . .	151
4.5.	Two views showing key features of the bottom PMT holder assembly. . . . .	153
4.6.	The XeNeu top PMT assembly. . . . .	155
4.7.	Ports on the top flange of the detector. Clockwise from the top of the detector flange: PMT high voltage feedthrough (see also Figure 4.40), PMT signal feedthrough, liquid outlet port, gate HV feedthrough, and liquid inlet port. Continuing clockwise: pumpout port, RTD feedthrough, cathode HV feedthrough tube, and another RTD feedthrough. The gate and cathode feedthroughs, as well as the pumpout and central ports are accessible from the vacuum box mounting flange. This flange and the central port on the detector top flange are transparent for visibility of other components. . . . .	156
4.8.	Cutaway of tee showing capillary and thermosyphon paths. The capillary is nested inside the tee and protrudes into the detector space through the 1/4" tube which is welded to the detector top flange. The capillary is used to draw liquid from the weir into the heat exchanger. The second branch of the tee (bottom) is used to volume share the detector gas with the condenser, forming a crude thermosyphon. . . . .	157
4.9.	XeNeu detector with HV feedthroughs made visible. The HV feedthrough tubes are mated to the ceramic feedthroughs via a copper reducer. End caps are attached to the gate (left) and cathode (right) feedthroughs which contain sockets for running wires to the gate and cathode. The brass end cap belongs to the cathode feedthrough, and the copper end cap belongs to the gate feedthrough. The caps are different because of an incident where a solder joint on the cathode feedthrough failed, which leaked silicone oil over the TPC components. Ethan Bernard (LLNL) performed the repair and created the replacement brass cap, which is fixed in place with a set screw. . . . .	159
4.10.	Schematic for high voltage delivery and monitoring circuit. $V_{IN}$ is the connection point for the high voltage power supply, Precursors to a full electrical breakdown are monitored with an oscilloscope at the output of a high-pass filter, $V_{MON}$ . $V_{MON}$ and the high voltage output $V_{OUT}$ are fed out of the box through coaxial connections. $R_{CL}$ is a current-limiting resistor, which protects components in the event of an electric short. . . . .	160

4.11. Schematic showing three different designs for a high voltage resistor chain on a field cage. At the left is a naive design that is unsafe because a short from the gate or cathode grid would have an unrestricted current. In the middle is a safer design that installs current-limiting resistors before the gate and cathode. A tradeoff of this design is a reduced voltage drop between the gate and cathode, requiring higher input voltages to match the unsafe design. At the right is the design of the resistor chain in XeNeu. Current-limiting resistors are in series with gate and cathode grids, however the grids are disconnected from other components. Consequently, the gate and cathode voltages are equal to the voltage of their respective electrodes. All resistors are 2 GΩ and the resistors between rings plug directly into holes machined into the outer diameter of the rings. Resistors adjacent to the gate and cathode plug into ports on the caps of the respective feedthroughs. . . . .	162
4.12. Rendering of XeNeu with portions of the vacuum box and the outer can made transparent. Electrical connections are fed out through the flange on the far right side, which has 5 SHV, 5 BNC, and 5 DB15 connectors. Directly above this flange, the condenser and heat exchanger (Figure 4.15) are mounted in the volume contained by an 11 in ISO-K nipple. . . . .	163
4.13. P&ID for components on the XeNeu detector cart. Arrows on the heat exchanger lines denote the normal circulation path. The blue dotted line demarcates the vacuum vessel. The dotted line from the condenser to the detector represents a connection which joins the detector and condenser gas spaces (Figure 4.8), acting as a simple thermosyphon. PT: pressure transducer, PI: pressure (dial) indicator. . . . .	164
4.14. Rear view of XeNeu cart. The condenser and heat exchanger (Figure 4.15) are housed in an 11" ISO-K nipple, offset from the axis of the detector. The flange on this nipple has connections to the cryocooler as well as gas inlet and outlet connections to the heat exchanger. In the foreground, oil filled plastic boxes are mounted on top of the detector and connect to the high voltage feedthroughs. . . . .	166
4.15. Hardware used for the thermal management in XeNu is housed in an 11" ISO-K flange offset from the axis of the detector. Pre-cooled gas exits the heat exchanger and enters the condenser. The gas is then condensed, where it enters the detector space as a liquid. Liquid exits the weir reservoir via a capillary line, which enters the heat exchanger. The condenser also has a third port which couples to the detector gas space. The gas exchange line is capable of cooling the detector volume if circulation is halted. . . . .	168
4.16. P&ID for the circulation cart. The SRV is not mounted on the cart, but included here for clarity. V32 and V31 are intentionally dead ends which lead to a volume that previously contained an Ar37 calibration source. . . . .	169
4.17. Drawing showing top view of shielding and background detectors arranged around XeNeu for measuring high energy nuclear recoils. A mockup of the XeNeu detector cart is visible at the right, surrounded by backing detectors (blue rectangles). Backing detectors at high scatter angles were positioned closer to XeNeu to make up for the decreased elastic scattering cross section of 14.1 MeV neutrons at those angles. . . . .	172

4.18. Results from the XeNeu recoil energy measurement compared to previous values. Values at the two highest recoil energies were averaged over their electric fields due to low statistics. From [106]. . . . .	174
4.19. Illustration of the Migdal effect. In a nuclear recoil with a dark matter particle, the nucleus is displaced from its electron shells. As a result, some of these electrons may escape, producing a coincident electron recoil. From [118]. . . . .	175
4.20. Scattering cross sections as a function of scattering angle, and expected Migdal ER spectrum. From [107]. . . . .	176
4.21. S2 spectrum from interaction at a scattering angle of 16°. Three background components are shown in the spectrum: single scatters, passive scatters, and multiple scatters. These are overlaid with predicted counts of Migdal interactions as well as the measured counts from the best-fit model. Fit residuals are shown in the bottom plot for the best-fit model (black points) and the predicted model (dotted line). The shaded grey band represents the systematic uncertainty in the NR background models.	177
4.22. Diagram showing key processes for outgassing. The bulk material is represented by the grid of larger green circles. These processes are not mutually exclusive. For example a species which is permeating the material may adsorb and desorb repeatedly as it diffuses into the vacuum space. Figure adapted from [119]. . . . .	180
4.23. Exploded view of the final design of the new bottom PMT holder. From top to bottom: the Shapal retaining ring, the grid contact ring, two Shapal high voltage pockets, and the aluminum base. White components are made from Shapal. The grid contact ring is colored green here to improve its visibility, but is made from aluminum like the rest of the PMT holder. Screws used were made of nylon due to their proximity to high voltage components. . . . .	183
4.24. View of PMT holder assembly from above, with the field cage removed except for the PEEK base ring. Dark tracks are visible in the cathode (left) and gate (right) pockets. These tracks are formed from rubbing of metal surfaces on Shapal components.	185
4.25. Screenshot of slow control readout during condensation of xenon with redesigned bottom PMT holder. When liquid xenon made contact with the aluminum bottom PMT holder, over the course of two minutes the detector temperatures rose almost 10 K. . . . .	186
4.26. Trigger schematic for studies of SE rates. Signals needed to be amplified twice in order to exceed the minimum discriminator threshold. Output from the 3-fold coincidence module was recorded at the digitizer for trigger efficiency studies. The 1-fold coincidence module provided the main event trigger. A pulse from the 3-fold coincidence would also fire both the 1-fold coincidence and a gate generator. A long cable was used for the gate generator signal to stagger the arrival times of the 3-fold coincidence and vetoing pulses at the 1-fold coincidence so the first coincidence pulse in an event would not veto itself. . . . .	188
4.27. Plot of AFT 75-25 vs pulse area showing counts of pulses. Single electron pulses are roughly within the closed red dashed curve. The band of pulses below the dashed orange line corresponds to S1-like pulses, which have a very low AFT 75-25. . . . .	189

4.28. Example of a trigger efficiency histogram. The first peak at a pulse area of roughly 60 photoelectrons (npe) corresponds to SEs. The trigger efficiency approaches 100% for pulses larger than four electrons in size. . . . .	191
4.29. Rates of child pulses following progenitors in Runs 11 and 12. Rates are normalized following the convention established in Equation 3.8. Child pulses are additionally binned by their size along the y-axis. Pulse sizes are normalized by the size of the single electron pulse in each run. . . . .	193
4.30. Total integrated rates from Figure 4.29 over pulses up to 10 SE in size. For this plot, rates were integrated from 100 $\mu$ s to 1 s. At all circulation rates, approximately 60% of the pulse rate in Run 11 is due to SEs, whereas SEs make up almost 80% of the child pulse rate in Run 12. With low-outgassing components in Run 12, the circulation rate had a much smaller impact on reducing background rates. This suggests that plastic components in the TPC are a dominant source of impurities responsible for ionization backgrounds. . . . .	195
4.31. Projection of Figure 4.29 showing integrated rates of child pulses up to 10 SE in size at different circulation rates. For this plot, rates were integrated from 100 $\mu$ s to 1 s. Rates for pulses of all sizes are affected by the circulation rate in Run 11. This is in contrast with rates for Run 12 in Figure 4.31, where only SE rates appear to have a slight circulation rate dependence. . . . .	195
4.32. Projection of Figure 4.29 examining rates of single electron pulses over time at different circulation rates. Rates prior to 100 $\mu$ s are excluded due to the high prevalence of merged pulses. . . . .	197
4.33. Example of a double exponential fit between 400 $\mu$ s and 30 ms in the single electron rate time dependence. . . . .	197
4.34. Fit values for each fit parameter and rate curve shown in Figure 4.32. Error bars represent statistical uncertainties. Values are plotted along the x-axis according to their run number and circulation speed; they are also printed in Table 4.3. . . . .	198
4.35. Single electron rates from Run 11 and Run 12. The rates in both runs appear to converge on a power law with a constant component (Equation 4.2) after 30-40 ms. Fit parameters for these power laws are presented in Table 4.4. . . . .	201
4.36. Typical dimensions of a SiPM and equivalent circuit for a single SiPM. . . . .	204
4.37. Rendering of the new XeNeu SiPM array. Each S13371 unit has four cells which can be read out individually, but in this design the cells in each unit were connected in parallel. A PEEK spacer was also designed to sit on the perimeter of the board and reduce wear between the edge of the board and the metal components in the tray and bracket assembly Figure 4.38. . . . .	206

4.38. Exploded view of the SiPM bracket and tray assembly. The tray is in the foreground, and has a cutout in the center that accepts the SiPM array and PEEK array spacer. Two L-shaped brackets on opposing ends of the tray are screwed into place, capturing the SiPM array. Two additional holes at the rear corners of the tray are designed to be concentric with tapped holes in the bracket so the tray can be securely mated to the bracket. The bracket is in the background and shares a similar design as the previous iteration. Two new holes in the same plane as the holes for the rear mounting clamp are included as optional electrical ground points at the corners of the bracket. Cutouts for the forward mounting posts on the bracket were intentionally asymmetric so the bracket alone could be installed by sliding it laterally onto the mounting posts. Through holes were used in all cases to ensure the absence of any trapped volumes. . . . .	208
4.39. Pieces of metal were observed on the sealing face of an o-ring (orange) belonging to a leaking feedthrough. The metal was believed to originate from threads on the feedthrough, which was likely overtightened. At the bottom left, a feedthrough as it was normally installed can be seen. . . . .	210
4.40. Right: Rendering of the 7-pin Mil-spec feedthrough. Left: Pinout of the 7-pin feedthrough, as used on XeNeu. Numbers and colors corresponded to labels used on the XeNeu vacuum box. . . . .	211
4.41. Power delivery cable assembly. At the left, crimp on terminals are held together with a 2-56 screw and two nuts. These terminals offered a simplified way to temporarily and securely attach multiple ground connections throughout the detector. The main ground connection (visible at the top) initially used a stainless steel braid; this was later replaced with a copper braid. . . . .	212
4.42. Example of a custom PMT pin connection on a coaxial cable. The pin is soldered to a ring terminal, which is crimped to the inner dielectric of the cable. The inner conductor is inside the crimp sleeve, which is then flooded with solder. The solder-covered tip of the inner conductor can be seen in this picture emerging from the body of the crimp sleeve, to the left of the base of the PMT pin. Not shown here are the copper wick to extend the outer shield, and heatshrink around the body of the ring terminal and cable in order to provide additional strain relief. . . . .	213
4.43. 23-pin Sub-C connector cutaway and pin mapping. This connector was used to feed signals from the SiPM array and bottom PMT through the detector top flange (Figure 4.7). . . . .	215
4.44. Comparison of top mounting bracket designs used to suspend components in the XeNeu detector space. The new bracket (Figure 4.44b) has a profile which occupies almost 50% less space than the old bracket (Figure 4.44a). . . . .	217
4.45. Trigger configuration for SPE calibrations. When the bottom PMT (CH10) crosses a threshold set at the discriminator, the discriminator fires a pulse which is picked up by the coincidence module. If the coincidence module is not being vetoed, it fires a pulse (TRIG) to the trigger channel on the DAQ and to the gate generator. The gate generator vetoes the coincidence module for a set period of time after receiving a pulse from the coincidence module. . . . .	219

4.46. Example waveform of an SPE captured during preliminary SiPM tests at the time indicated with the orange “T”. The SiPM was biased to 55 V and the signal was filtered with a 5 MHz low-pass filter. A voltage and time scale of 2.00 mV and 2 $\mu$ s per division were used; these are printed at the top and bottom of the figure, respectively. . . . .	220
4.47. Example of an ROI integral spectrum from a single SiPM at a single voltage. The sum of four gaussians was fit to the spectrum. This fitting procedure was repeated for each SiPM channel at each bias voltage in order to determine the minimum voltage needed to resolve SPEs. . . . .	221
4.48. Gain curves for each SiPM channel. “ALL” represents the gain curve obtained after fitting to the sum of the integral spectra for all channels. Note that CH 8 is missing: this channel was broken on the DAQ so CH 9 was used instead. . . . .	223
4.49. Coordinate system used for positioning sources around XeNeu. $R$ , $Z$ , and $d_{laser}$ were measured with a standard measuring tape or calipers, where possible. . . . .	223
4.50. Trigger for single electrons using the bottom PMT. The bottom PMT signal was doubled using two FIFO modules to make use of the full range of the discriminator. The main event trigger was configured such that Discriminator 1 would need to be coincident with a delayed copy of itself. This would ideally happen for a pulse with a width greater than the delay time. . . . .	225
4.51. Example waveform showing the sum of the bottom PMT and SiPM outputs, with two SE pulses. . . . .	226
4.52. Comparison of reconstructed pulse areas for electroluminescence signals in the SiPM array vs. the bottom PMT. SE pulses are prominently featured in this plot between 20 to 80 npe on the PMT axis. . . . .	227
4.53. Reconstructed pulse positions using the CoM approach. The square distribution is unrealistic and a consequence of pulses at the edge of the TPC being misreconstructed towards the center. . . . .	228
4.54. Comparison of positions generated by the CoM algorithm and true positions from the simulation. The CoM algorithm is accurate for pulses with a true position within 1 cm of the TPC center, but misreconstructs positions thereafter. . . . .	228
4.55. Reconstructed pulse positions using the LRF. . . . .	230
4.56. Comparison of positions generated by CoM and LRF. . . . .	231
5.1. Emission spectra from xenon-doped argon gas mixtures generated via electron beam excitation. Light production from xenon excimers centered around 174 nm virtually saturates at concentrations as low as 10 ppm. From [141]. . . . .	235
5.2. Emission spectra from xenon-doped argon gas mixtures generated via heavy ion beam excitation. Subplots are arranged in columns of the same xenon concentrations and rows of the same total mixture pressures. Rows <i>a</i> through <i>e</i> correspond to total mixture pressures of 400, 780, 1000, 1200, and 1400 mbar respectively. Efthimiopoulos <i>et al.</i> note that emission at concentrations of 100 ppm is predominantly from atomic transitions. From [145]. . . . .	236
5.3. Examples of solid xenon formation resulting from excessive distillation of xenon in a xenon-doped liquid argon mixture. . . . .	237

5.4.	Plot of temperature composition curves for xenon and argon (Equation 5.5) at atmospheric pressure. This plot shows the eutectic point of the binary mixture, which is the intersection point between the two curves. . . . .	241
5.5.	Empirical temperature vs. composition plot reproduced from [157]. Labels on the plot designate regions where solid, liquid, and vapor phases are in equilibrium with each other. . . . .	243
5.6.	Theoretical temperature vs. composition plot of a xenon-argon mixture from [160]. Phase boundaries (solid lines) indicate predicted solid (S), liquid (L), and vapor (V) phases of xenon (2) and argon (1) in equilibrium (E) in the mixture. Data from Heastie and Lefebvre (Figure 5.5) are shown with red squares while data from Yunker and Halsey (Figure 5.8) are shown as purple circles. . . . .	244
5.7.	Argon and xenon vapor pressures in the temperature range of 85 to 110 K. Data sources: NIST Argon [164], Gilgen <i>et al.</i> [161], Ferreira and Lobo (F&L) [156], CRC [165], NIST Xenon [166], Leming and Pollack [163]. . . . .	245
5.8.	Deviation of argon partial pressure $P_{Ar}$ from pure argon pressure $P^0$ as a function of the xenon liquid mole fraction and temperature. The solid black line indicates and “ideal” solution described by Raoult’s Law. Reproduced from [159]. . . . .	245
5.9.	Comparison between predicted xenon vapor mole fractions using Raoult’s Law and empirical data from Yunker and Halsey [159]. Compared to Figure 5.9a, Figure 5.9b relies on an additional assumption that the partial pressure of xenon in the mixture should approach the vapor pressure of solid xenon when the mixture saturates (Equation 5.13). Points above the saturation mole fraction (Equation 5.14) are excluded. . . . .	249
5.10.	Schematic showing key components of CHILLAX. The left volume is the condenser where incoming gas is condensed by TSL. The right volume is the main volume which houses the capacitance meter for measuring the liquid composition. The condenser and main volume are connected by a lower liquid transfer line and upper gas interconnect. Gas enters via an inlet tube heated by h2, and exits at the top of the main volume. Stainless steel mesh in the condenser (represented by vertical black lines) traps gas bubbles from the inlet tube to motivate condensation in the liquid. Copper mesh is present at the top of the detector volume to mitigate the heat load of cables present in the main volume. From [135]. . . . .	251
5.11.	Diagram illustrating the loop (left) and counter-flow (right) thermosyphon concepts. Gravity ( $g$ ) assists in moving liquid downward, but alternative designs exist which use sintered material, fine mesh, or small channels to move liquid against gravity via capillary action [169]. . . . .	253
5.12.	Pool boiling curve. Different boiling regimes occur as the temperature of a wall $T_w$ deviates from the saturation temperature $T_{sat}$ of a liquid. These boiling regimes convey different quantities of heat per unit of contact area. The Leidenfrost Point (C) is noteworthy for being a local minimum in the pool boiling curve, where an excessive temperature difference above the critical heat flux point (B) results in suboptimal heat transfer. From [171]. . . . .	254

5.13. Pool boiling curves for Argon at different pressures showing the Peak Nucleate Boiling Flux (PNBF) at each pressure. At some pressures the heat flux was measured with a monotonically decreasing temperature difference instead of an increasing temperature difference to illustrate the effects of hysteresis. From [172]. . . . .	255
5.14. Rendering of the thermosyphon condenser head and funnel. A cross section of the funnel is shown to expose the interior of the condensing volume. In normal operation, gravity pulls liquid down the tapered end of the funnel, delivering LN to the inlet of the TSL. Gas returns from the TSL via one of two side ports on the funnel. A second gas return port accommodates gas returning from the TSU. The base of the condenser head is affixed to the bottom plate of the CH104 cryocooler using screws.	257
5.15. Cross-section of the CHILLAX condenser. Gas from the circulation system enters the condensation volume at the bottom right and bubbles out through a brass plate. Once condensed, liquid exits to the CHILLAX main volume at the bottom. The TSL was welded to the top CF flange to isolate it from the condensing volume; the boundary of the TSL is indicated with a dashed line. Liquid from the thermosyphon condenser enters the TSL via an inlet (2) where it makes contact with the vented aluminum helical inserts. These inserts were designed to convey liquid along the walls of the TSL to provide more uniform cooling and maximize the contact area of the liquid with the TSL walls. Vents in the inserts provide a path for gas in the evaporator to exit via the outlet (1) without entraining incoming liquid. The upper gas interconnect (3) is used to equalize gas pressures between the condenser and main volume of CHILLAX (Figure 5.10). . . . .	258
5.16. Rendering of a Hamamatsu S13370-6075CN. . . . .	259
5.17. Render and photos of SiPM cell. All three units are immersed in liquid argon during operation. . . . .	260
5.18. Plots showing averaged waveforms and pulse area spectra for windowed and windowless S13371 modules. . . . .	261
5.19. Upper bounds on SiPM dark count rates as a function of overvoltage. The breakdown voltage for both SiPM models was approximately 42.5 V. Figure modified from [131].	262
5.20. Plots showing afterpulsing and crosstalk estimates. . . . .	263
5.21. SiPM PDE as a function of overvoltage. The bands illustrate systematic uncertainties from varying cuts used to produce the PDE estimate. From [131]. . . . .	264
5.22. Study of fractionation at different pressures with mixtures of different compositions. Dotted lines are curves from a naive model whereas solid lines represent a more complete model. From [178]. . . . .	267
5.23. Pumped-bypass inlet configuration. The capillary reduces the gas pressure entirely within the viscous flow regime, and a molecular leak reduces the pressure entirely within the molecular flow regime to avoid gas fractionation. Excess gas is then routed away from the molecular leak to prevent a build-up of gas from the capillary inlet. From [180]. . . . .	268

5.24. Schematic of a Hiden HPR-70 batch sampling system. Supporting documentation could not be found on the exact operational sequences, but a schematic such as this is still useful for the purpose of reverse-engineering a batch-inlet sampling system. It is conceivable to assume that a sample is admitted into the volume between V1 and V2, then expanded to the 1 L volume downstream of V2. Following that, gas flows directly through LV1 to the quadrupole mass spectrometer (QMS). It is likely while sampling from the 1 L volume, V2, V3, and V4 are closed to prevent the sample from depleting too quickly. This would mean that a batch sampling system <i>can</i> operate without a pumped bypass (Figure 5.23), which was an important distinction that was not fully recognized when designing THERANOS (Figure 5.25). From [185]. . . . .	269
5.25. P&ID for the most recent version of THERANOS. See Table 5.2 for a list of key components. LV1 was left fully open during samples to avoid a redundant pressure reduction following AP1. V14 was installed in an attempt to regulate pressures ahead of AP1 but was never tested. . . . .	272
5.26. Front and rear views of THERANOS. . . . .	274
5.27. SRS RGA Software showing a P vs T scan of argon and xenon masses. A key feature which was missing was the ability to plot pressures on separate y-axes, which led to suppressed features on both curves. See also Figure 5.29. . . . .	275
5.28. Frontend of <code>rgaplot</code> showing part of a P vs T scan (Figure 5.29). P vs T scans include options for selecting masses to plot, which influence additional plots (not shown) like ratios and histograms of time series data. . . . .	276
5.29. Scan shown in Figure 5.28. The level of detail is much greater than what can be seen in Figure 5.27. . . . .	277
5.30. View of internal PDE components. The board at the top left was unused but left as a future option for reading out IG1. From top right to bottom right, the breakout boards connected to: Omega DPi 32 modules, PT1, CM1, and the DAQ6510. . . . .	278
5.31. THERANOS PDE wiring schematic. Capacity for reading out and powering the IGM402 was planned for, but never installed. . . . .	279
5.32. The first iteration of THERANOS had a very simple design where gas was fed directly into LV1. Regions upstream of LV1 could be pumped via V5 which acted as a bypass, however TP1 needed to be shut off and fully spun down to evacuate any high pressure samples. . . . .	281
5.33. MAGMA was installed in the second version of THERANOS. By opening V8, MAGMA could be used to agitate gas at LV1 instead of circulate gas past it. . . . .	281
5.34. Example scan with MAGMA. The CDEM was enabled for 129 (xenon) and disabled for argon (40). The bottom plot shows the ratio of the xenon and argon signals. The first mixing period is indicated with an arrow. It should also be noted that the argon pressure was nominally constant, compared to the xenon pressures which more than doubled in the same timeframe. . . . .	283

5.35. Scan taken with CM1 at 1 torr and IG1 at 5E-5 torr. LV1 was adjusted continuously to keep the value of IG1 as stable as possible. During the period between the dotted lines, LV1 was not adjusted and the pressure at IG1 was allowed to decrease naturally. Spikes in the values are from a high scan speed (4) at the RGA that introduced additional noise. These spikes were greatly diminished after switching to scan speed 2 in subsequent acquisitions. Because the scan speed was changed before the final calibration curve was generated, results shown here were not used to infer xenon gas concentrations. Note the use of a logarithmic scale in the ratio to capture the behavior at the end where the xenon pressure had bottomed out and the argon pressure decreased exponentially. . . . .	285
5.36. Sample taken from the outlet of early CHILLAX doping tests at 0.3% xenon liquid mole fraction and initial MV2 pressure of 1.5 torr. Nominal pressures from IG1 are indicated on the top plot. These pressures were set using LV1. . . . .	286
5.37. Example of scan instabilities occasionally seen when operating with gas agitation from the XeNeu circulation pump. The initial rise may have been from a xenon-rich dead volume trapped behind AP1. The gradual decline is not consistent with fractionation however, which doesn't appear to occur until almost 2500 s into the scan. This degree of variation within a scan was consistent with values found in other sources [184] [190]. However, inconsistencies of approximately 20-30% between scans persisted. . . . .	288
5.38. Cutaway of V11 and the HVCR-4 tee showing the location of the 1/4" copper inlet tube (highlighted in orange) added to reduce the size of the dead volume upstream of AP1. V14 is visible at the top. V11 (middle) was connected to LV1 (bottom) and V8 (lower right) via the HVCR-4 tee which housed the inlet tube. The inlet tube was suspended from the VCR gasket at the top of the HVCR-4 tee by flaring the tip of the tube and sliding the other end of the tube through the gasket. Note that VCR gaskets are not shown in this rendering. . . . .	291
5.39. Scan taken from early CHILLAX doping tests at 0.73 torr. Unlike in Figure 5.35, the measured ratio begins exponentially increasing before the xenon pressure reaches its baseline. . . . .	293
5.40. Example of a scan used for calibration purposes. The region inside the dashed box was used for calibration purposes due to it's stability. Subsequent calibration scans relied on similar procedure to select data where stable time periods of 5-10 minutes at the beginning of a scan were used in the final calibration curve (Figure 5.41). . . .	294

5.41. RGA calibration curve. To generate the calibration curve, mean values of 129/40 ratios during selected time intervals in calibration scans were plotted against known mixture concentrations. A linear fit of the form $y = mx + b$ is performed with the value of $b$ fixed at 0. Values with inlet pressures below 1.5 torr were excluded from this fit to maintain a more conservative estimate. Error bars on points represent the standard deviation of the selected scan period and serve as a rough guide of the variation present in scans; they were not used to generate the fit parameters. The ratio of the 129/40 RGA response can be multiplied by the calibration factor $m$ to obtain an estimate of the xenon concentration in a gas mixture. The error on $m$ is statistical only, and it dominates over systematic errors present from excluding nearby points with inlet pressures less than or equal to 1.5 torr. . . . .	296
5.42. Capacitive measurements of the CHILLAX liquid bath over time. Doping steps are outlined with the hatched red regions. Each point on the histogram represents a data point recorded by the capacitive meter in a specific time bin. . . . .	298
5.43. Photos of the CHILLAX main volume taken by a camera mounted on the underside of the top flange. Photos from the first, second, and third tests are in the top, bottom and middle rows, respectively. In the second test a prominent ice ring is visible as a bright arc spanning the top right corner to the bottom left corner of the photos. A fainter ice ring is present in the third test, while no ice ring was present in the first test. These photos provided substantial evidence that ice formation was strongly linked to the presence of a thermal gradient between the liquid bath and walls of the main volume. The thin dark arc present in all of the photos is the edge of the bubble router at the bottom of the main volume. . . . .	299
5.44. Capacitive censor data (top) and camera data (bottom) showing the xenon concentration in the liquid bath and ice ring thickness, respectively. The blue squares correspond to data from the first test, while the red circles correspond to data from the first test. Data from the black test is displayed as black triangles. The growth of the ice ring neatly corresponds with a depletion of the xenon concentration in both the second and third tests. In the first test, no xenon depletion or ice ring formation was observed beyond the range of the uncertainties of the camera analysis and capacitive sensor. . . . .	300
5.45. Xenon concentration over time after a test that formed an ice ring around the liquid bath. The top flange was cooled to condense gas on the walls of the main volume, which dissolved the ice back into the liquid bath. Chunks of ice fell back into the liquid at times, resulting in abrupt rises in the xenon concentration measured by the capacitive sensor. . . . .	301

5.46. Gas samples acquired with THERANOS at a liquid doping concentration of 2.35%. Error bars on points represent the standard deviation of the ratios over time. These can be understood as an indication of LV1 being changed during scans, or slow drifts of measurements over time during scans. Ice was intentionally formed on June 7th, which led to a sharp rise in the measured xenon gas concentration. The ice appeared to be fully dissolved by June 14th, but the vapor mole fraction remained elevated when compared to values measured prior to ice formation. This may have been a consequence of systematic errors, but could be a result of an enhanced vapor mole fraction from a nearly saturated liquid film at the walls of the main volume. . . . .	303
A.1. In this diagram, a set of cutting edges (typically two halves of wire stripping die) requires a significant shear impulse relative to the wire to pierce the kapton film. The cutting edge pries apart the insulation (colored orange) with some impulse $J_{C \rightarrow I}$ . Because the insulation does not slide along the wire, this impulse is directly transmitted to the wire as $J_{I \rightarrow W}$ . Thus, despite not making direct contact with the conductor, the cutter transmits a high tension impulse via the insulation, which can break the conductor. . . . .	306
A.3. Various stripping techniques were attempted with different strippers. Techniques indicate layers of insulation and/or shield removed, the stripper ID used, setting, and relative difficulty. Fully stripping the inner dielectric was the most difficult step, which could be accomplished with some practice using strippers 6 and 2. . . . .	309
A.4. After a sleeve is crimped to the inner conductor, the cable is captured in the body of the connector. A piece of 28 AWG wire is fed into the appropriate receptacle on the connector, then cut to length to match the location of the outer conductor from the cable. The 28 AWG wire is then stripped and soldered into place. It is vital that when soldered, the cable and lead wires have matching lengths and are not bent out of place in order to avoid unnecessary strain on either. . . . .	312
A.5. A coaxial cable with a layer of heatshrink partially covering the outer conductor and knurled portion of the MMCX plug. The plug is also connected to a jack in order to check for continuity. Not shown here is a crimp sleeve to the left of the heatshrink, which is pushed over the heatshrink tube, pushing the heatshrink to the right in the process. The layer of heatshrink prevents the edge of the crimp sleeve from biting into cable and severing the insulation layers. It also helps immobilize and protect the outer conductor wires which have a tendency to sever after being crimped. . . .	315
A.6. Prepared cable prior to insertion in female D-sub connector. This photo shows an early iteration where the outer conductor was peeled back, instead of being soldered to a copper braid. . . . .	316
C.1. RGA Stand Xenon Inlet connected to V9 on XeNeu . . . . .	325
C.2. Front panel of the Keithley DAQ 6510. Image courtesy of Keithley. . . . .	340
D.1. Early studies of electron backgrounds focused mainly on observations of electrons within a drift length or less following progenitors. “Spontaneous” pulses were still observed prior to events however this spontaneity was disproved in 2011 by ZEPLIN-III.	345

D.2. Later studies looked at backgrounds well after a full drift time and became increasingly detailed. A majority of studies have found a power law dependence in the rate, at timescales up to 1000 ms after a progenitor. . . . .	346
D.3. Example of a single electron pulse observed between an S1 and S2 in ZEPLIN-II. . .	347
D.4. SE counts following RFR S1 pulses in ZEPLIN-III. Counts followed a decreasing trend up to the full drift time where a spike in the counts was observed. The cause of this spike was inferred to be from photoionization of the cathode due to its consistent location at the maximum drift time in the TPC. The decreasing trend was ascribed to electrons being captured by electronegative impurities. From a previously calculated value for the photon mean free path for photoionization they were able to estimate the original number of photoelectrons produced following an interaction. This led them to obtain an estimate of the electron lifetime from the slope of the trend. . . . .	350
D.5. SE count rates observed by Akimov <i>et al.</i> dropped when tilting their TPC in either direction, possibly indicating that electrons trapped at the liquid surface were being removed before being able to emit later. . . . .	354
D.6. S2 and S1 photoionization rates in XENON100. . . . .	355
D.7. XENON100 was able to confirm the energy dependence of photoionization rates seen in prior experiments like ZEPLIN-II, and demonstrate that the rates were dependent on the liquid purity. . . . .	357
D.8. “S3” pulses (green) observed by Akimov <i>et al.</i> appeared no more than 200 $\mu$ s after a muon S2 (pink). In the intervening time, light (red) appeared to trace a path between the S2 and S3 across the TPC. S3 pulses preferentially appeared at a specific location on the edge, consistent with a known tilt of the TPC. These S3 pulses were concluded to be trapped electrons on the liquid surface. . . . .	359
D.9. Akimov <i>et al.</i> fit a power law of the form $at^\beta + c$ to the SE rate for over 10 ms following muon interactions in their TPC. . . . .	360
D.10. Events from a $^{210}\text{Po}$ source deposited on the cathode of a dual-phase xenon TPC. Electron trains were studied by looking at a summed waveform of signals from alpha recoils up to 1 ms after the S2 was detected. . . . .	364
D.11. Evidence from LUX showing that electronegative impurities were unlikely to be responsible for photoionization electrons produced in the TPC. . . . .	371
D.12. Time-correlated single electron backgrounds in LUX were observed at least a second after an event had occurred. The time dependence of the rate was found to fit a power law, and many of the background pulses appeared to be co-located with their progenitor. . . . .	373
D.13. Purity dependence of electron trains measured in LUX. . . . .	374
D.14. Transient periods of electron emission in LUX were visible even after aggressive vetoing on high energy events. Electrons from these emission cycles were localized to specific regions in the detector. . . . .	376

D.15.	Integrated SE rates vs. the drift field seem to follow the opposite trend of the same rates plotted vs. time since the progenitor S2 pulse. The origin of this discrepancy is unclear. It is possible that rates were accidentally integrated over a fixed window which included livetime after the maximum drift time. This could cause the integrated rate at higher drift fields to appear lower. . . . .	379
D.16.	S2-normalized rate flux of single-to-few electron pulses detected within a roughly 10 mm radius around the XY location of progenitor S2 pulses. . . . .	382
D.17.	Electron background fractions normalized by progenitor S2s corrected for drift losses. According to Kopec <i>et al.</i> the single electron fraction rises with drift field, possibly due to an increased charge yield or increased release probability from impurities in the liquid. They also claim the lack of response from multi-electron backgrounds is consistent with impurities in the liquid affecting only single electrons. . . . .	384
D.18.	Electron trains from progenitors in the gas, drift, and reverse field (below cathode) regions of XENON1T. A 200 ms holdoff was imposed after interactions to avoid overlapping electron trains between progenitors. The SE rate during this 200 ms holdoff period is shown with the solid purple line. It serves to verify that the holdoff period was adequate in allowing SE rates to return to a baseline level before SEs from progenitor pulses were tracked. It also shows that rates following progenitors in the gas and reverse field regions return fairly quickly to the pre-holdoff rates, indicating that electron trains are a phenomenon tied to drift region interactions. . . . .	389
D.19.	Larger progenitor S2s in XENON1T had more delayed single, double and multi-electron background pulses. Position correlated pulses were defined as those within $\Delta r < 15$ cm of their progenitor and position uncorrelated pulses were those outside of $\Delta r > 20$ cm. The apparent dependence on progenitor size for position-uncorrelated backgrounds was attributed to leakage from position-correlated backgrounds. . . . .	390
D.20.	Delayed electron fraction from 2-200 ms following a progenitor S2 in XENON1T. The raw S2 (uncorrected for drift losses) size was used for normalization. XENON1T determined that the weak correlation with extraction field would make trapped electrons at the liquid surface a subdominant component of electron trains at best. . . . .	391

# Chapter 1.: Introduction to Dark Matter

## Tip

The PDF copy of this dissertation is meant to be used for archival and note-taking purposes. The web version has plots which can be easily saved, and some of them are interactive. The repository also contains notebooks used for analyses. Links to these are detailed in [Appendix E](#).

Three-quarters of the total mass in the universe is believed to be “dark”. It does not interact with light, yet its presence has been widely confirmed thanks to astrophysical observations. Because many of its properties have remained opaque, the question of dark matter has been elevated from an observational anomaly, to one of the most prominent unresolved mysteries in physics. With many models ruled out, the detection of dark matter is certain uncover new physics beyond the Standard Model, making it an enticing area of study.

This chapter will begin by reviewing some of the more significant pieces of evidence for dark matter, with the second section broadly considering more prominent theoretical candidates. The third section looks into common strategies for detection, with brief synopses of detection techniques.

## 1.1. Evidence for Dark Matter

For early astronomers it was perfectly reasonable to expect that some objects in the night sky could not be easily observed. Prior to their discovery, planets from our own solar system were “dark”, in that their gravitational effects could still be observed elsewhere. Arguably the most noteworthy example of this came when the existence of Neptune was predicted by Urbain Le Verrier and John Couch Adams in 1846 [1]. At the time, numerous observations of the anomalous motion of Uranus had been made, but Adams and Le Verrier worked independently of each other with the hypothesis that the orbit of Uranus was being influenced by an unseen body. There was some controversy in the assigning of credit for the prediction, but ultimately Le Verrier had the most accurate calculations and confident outlook in his work.

Alongside and following the discoveries of planets, questions continued to arise in regards to the existence of “dark matter”—entities believed to be responsible for various gravitational effects but had not been directly observed. Typical culprits included dust clouds, nebulae, faint stars, and other objects which were not particularly luminous. To make determinations about these dark entities, astronomers in the early 20th century turned to estimation techniques which would offer some insight into how much mass was present in a given region of space—luminous or otherwise. For further reading on the history of dark matter detection, see [1].

### 1.1.1. Mass Estimates of Galactic Clusters

In 1933 Fritz Zwicky studied the Coma Cluster, and used the virial theorem to predict the velocity dispersion of galaxies in the cluster. By working with the number of observed galaxies in the cluster, an estimate for the average mass of a typical galaxy, and an estimate of the physical

size of the cluster, Zwicky predicted a velocity dispersion of 80 km/s. This was in stark contrast with the observed value of 1000 km/s, which implied to him that most of the cluster was populated with “dark matter” [2].

### 1.1.2. Galactic Rotation Curves

The advent of 21 cm astronomy enabled measurements of the rotational velocities of galaxies to their outer reaches. While these measurements had been performed in prior years and decades, Vera Rubin and Kent Ford published a significantly improved set of data in 1970 which measured the rotation curve of the Andromeda galaxy [3]. At roughly the same time, Ken Freeman asserted that previously measured rotation curves did not reflect their galaxy’s optical properties [4]. Luminous matter in galaxies followed distributions which would decay farther from the center, but largely static rotational velocities implied that significant mass was present at the outer reaches of galaxies.


This is because Newtonian dynamics predicts a rotational velocity  $v(r)$  of

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

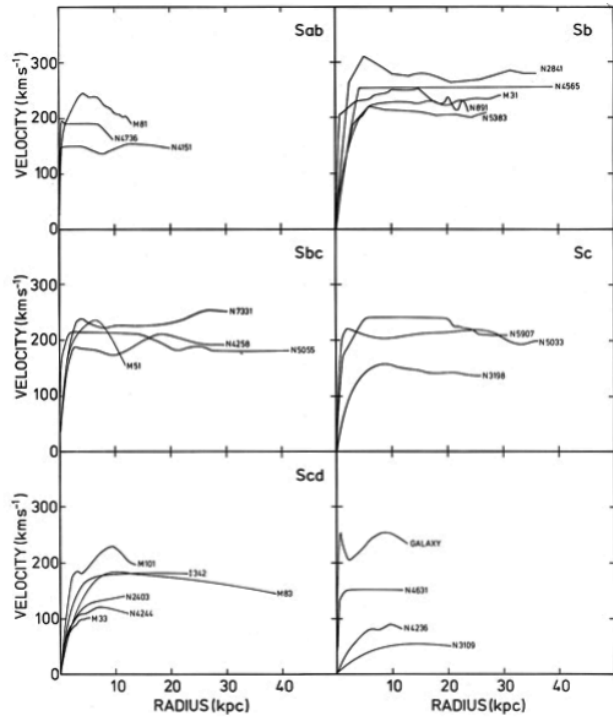
where

$$M(r) = 4\pi \int \rho(r)r^2 dr$$

$G$  is Newton’s gravitational constant, and  $\rho(r)$  is the mass density at radius  $r$ . It had been expected that  $\rho(r)$  was proportional to  $1/\sqrt{r}$  but because  $v(r)$  is approximately constant, we are led to conclude that  $M(r) \propto r$  and  $\rho \propto 1/r^2$  [5].



(a) Rogstad and Shoshak 1972



(b) Bosma 1978

Figure 1.1.: Rogstad and Shoshak compared observed visible densities of galaxies with their rotational velocities. A few years later, Albert Bosma demonstrated the flat rotation curves of 25 different galaxies which added to the consensus surrounding dark matter at galactic scales.

Data showcasing inexplicably flat rotation curves continued to emerge throughout the 70's. In 1972 Rogstad and Shoshak directly compared observed hydrogen surface densities of five galaxies to their rotation curves [6]. Between '72-'73, Roberts, Whitehurst, and Rots published data on three more galaxies—including an improvement on the Andromeda measurement made by Rubin and Ford [7] [8]. The field was beginning to accept the idea that there were vast quantities of “missing mass”, with Zwicky and Smith’s past observations of clusters further strengthening that position. By the end of the decade the missing mass problem had largely solidified, in part because of a radio-based analysis of 25 galaxies by Albert Bosma [9] see Figure 1.1, and optical observations

of 10 galaxies by Rubin, Ford, and Thonnard [10].

### 1.1.3. Cosmic Microwave Background (CMB)

Pensiaz and Wilson performed the first measurement of the CMB in the 1960's, and found a nearly constant signal throughout the entire universe. The smoothness of this signal motivated arguments for the presence of dark matter in the early universe, and with the launch of COBE in 1989, a more detailed glimpse of the CMB became possible. The CMB is the earliest picture of the universe, when matter had cooled down enough to become transparent to photons. Photons escaping from the cooled matter will be slightly redshifted if they originate from areas of high matter density, because of the slightly deeper gravitational wells they need to escape. Thus, fluctuations in the CMB are expected to reflect matter densities at the time of formation, and the density fluctuations are a result of opposing interactions between gravitational attraction and radiation pressure that eventually froze into place [11].

A map of the density fluctuations in the CMB can be decomposed into an angular power spectrum [5]. Distinct oscillation modes in the spectrum ultimately speak to the total energy density of the universe, the fraction of baryonic matter, and the fraction of dark matter. Because dark matter interacts only gravitationally, its existence will result in the strengthening of modes that are distinct from those caused by interactions between normal matter and radiation pressure [11].

Limitations in the precision of COBE resulted in an inability to detect fluctuations at the scales necessary to observe the imprints left behind by dark matter [12]. But the results from WMAP in the mid 2000's [13], and later from Planck in 2013 (Figure 1.2), were more than enough

to conclusively match predictions that had been made about the makeup of the critical density of the universe. Dark matter was determined to dominate approximately 27.5% of the critical density, whereas baryonic matter only contributed to about 5%.




Figure 1.2.: 2018 Angular power spectrum constructed from Planck CMB data [14]. The first three peaks provide information on the critical density, baryon density, and dark matter density respectively. The blue line is a fit from dark energy-dark matter models, and the red points are observed data. The dotted gray line marks a change in the scale of the x-axis. The CMB is extremely uniform, with temperature fluctuations  $\frac{\Delta T}{T} \sim 10^{-5}$  only visible at small scales.

#### 1.1.4. Big Bang Nucleosynthesis (BBN)

Throughout the 1950's it was held that stellar furnaces were the genesis of elements up to iron [1], however stellar nucleosynthesis models could not accurately predict observed quantities of helium or deuterium. Deuterium in particular is *consumed* by stars [1], which made previously relegated theories of BBN more appealing.

Standard BBN holds that prior to 1 second before the Big Bang, populations of neutrons and protons emerged that were in thermal equilibrium with each other. However as the temperature in the early universe dropped below 1 MeV, or roughly the mass difference between the neutron and proton, these populations froze out at a ratio of  $n/p \approx 1/6$  [15]. The  $n/p$  ratio dropped further to 1/7 as free neutrons underwent beta decay, during which time temperatures dropped to about 0.1 MeV, when deuterium could avoid near-immediate photodissociation.

From there reactions to form helium-4, helium-3, and tritium began to proceed, with helium-4 comprising the majority of mass content (aside from hydrogen) because it is the most stable light element. As temperatures dropped further, residual deuterium was frozen out of these processes and remained with some relic abundance.

Thus, observations of the abundances of these light elements can be used to infer the baryon density of the early universe. By 1973, a limit had been calculated that placed  $\Omega_b \lesssim 0.1\Omega_{crit}$  [16], and by 1990 the limit had dropped further thanks to improved measurements of the primordial deuterium abundance [1].

Ultimately, while measurements of the baryon density via the CMB have provided a stronger constraint on BBN than other astrophysical observations (see Figure 1.3), agreement<sup>1</sup> between the two is considered a strong cross-check in favor of the existence of non-baryonic dark matter.

### 1.1.5. The Bullet Cluster Merger

In 2006, another piece of evidence helped to demonstrate both the existence of dark matter, and offer some insight into its nature. During a merger of galactic clusters, galaxies are virtually

---

<sup>1</sup>There is a notable *disagreement* between observed abundances of  ${}^7\text{Li}$  and abundances which should be present based on the accepted value of the baryon density. This disagreement is known as the cosmological lithium problem. Solutions include: revised measurements, a better understanding of mass-7-producing nuclear processes, and the possibility of physics beyond the standard model [17].



non-interacting because of their exceptionally low density in a cluster [18]. However, the hot intracluster medium (ICM) of each will be highly interacting, slowing down and behaving like a fluid which will proceed to cool via x-ray emission. Clowe *et al.* recognized that because of these effects, a cluster merger should be a scenario where both stellar matter and the ICM are spatially separated. With the two baryonic contributions to the gravitational potential effectively decoupled, a weak microlensing survey could determine the contribution from dark matter.




Figure 1.4.: A comparison of optical and x-ray images confirmed that galaxies had passed through each other relatively unimpeded, while the ICM had lagged behind due to higher interaction rates. Additional microlensing calculations resulted in contours (green lines) which increased towards the centers of the galactic distributions [18].

As shown in Figure 1.4a, optical imaging of the Bullet Cluster confirmed that galaxies from each cluster had passed by each other unscathed. X-ray imaging (Figure 1.4b) from the Chandra Observatory also confirmed that the ICM from each cluster had collided closer to the spatial center of the system. In a cluster, the ICM is known to make up approximately five to seven times as much of the baryonic content in a cluster as galaxies. However, a microlensing survey showed that the gravitational potential was strongly concentrated at the centers of each distribution of galaxies. This was the critical evidence showcasing the effects of dark matter, which was deemed responsible

for shifting the gravitational center towards the galaxies and away from the ICM. This example and others [19] have demonstrated that dark matter composes the majority of the mass in galactic clusters, and they also provide weight to the claim that dark matter is weakly interacting and non-baryonic.

### 1.1.6. Structure Formation

To qualify as candidates, dark matter models must also be able to match expectations for the formation of large scale structures in the universe. Structure formation is useful to consider because it is a question which is largely agnostic of short-range particle interactions. Instead, the velocity distributions (i.e. the temperature) of a given candidate will dictate the way in which cosmological structures are formed. This gives rise to the designation of “hot” and “cold” dark matter, with hot particles behaving relativistically, and cold ones behaving non-relativistically.

While both classes can form similar structures at very large scales, density fluctuations at small scales would depend on the extent of a candidate’s thermal motion. A universe with hot dark matter will have a “top down” evolution where larger structures form first, and galaxies peel off later. On the other hand, cold dark matter (CDM) results in smaller structures forming first, and then merging later. Observations at the sub-cluster level have shown that there are distinct structures [20], largely excluding hot dark matter candidates [21]. On the other hand, structure formation simulations of CDM have been able to reproduce these observations [22].

## 1.2. Dark Matter Candidates

While dark matter has been widely observed at astrophysical scales, attempts to identify the exact identity and nature of it have proven fruitless so far. Being constrained by the aforementioned evidence, leading dark matter models have taken to describing particles beyond the Standard Model (SM) of physics. These particles are then subject to additional constraints like self-annihilation cross-sections, stability requirements, and interaction channels with SM particles [23]. Based on this, the following section will consider some of the more popular candidates.

### 1.2.1. Standard Model Neutrinos

Interest from a cosmological perspective on neutrinos dates back to the mid-20th century, when Zeldovich and Gershtien used the first CMB measurements to estimate the population of neutrinos that would persist as a thermal relic. To be consistent with measurements of the Hubble constant (i.e. the expansion rate of the universe), they estimated a 400 eV upper limit on the masses of the neutrinos [24]. Within about a decade, further estimates had been derived that excluded neutrino masses in the range of  $\mathcal{O}(10 \text{ eV})$  to  $\mathcal{O}(1 \text{ GeV})$  [25]. The elevation of neutrinos as a dark matter candidate came soon after, with the recognition that a stable and weakly interacting lepton could explain the missing mass problem in galaxies and clusters [26].

Following that, their downfall as a dark matter candidate came with the understanding that they would be too hot to reproduce the observed sequence of structure formation [21]. As of 2019, a terrestrial measurement by the KATRIN collaboration has placed an upper limit on neutrino masses of  $< 1.1 \text{ eV}$  [27], further confining a relic population to a relativistic regime in the universe.

### 1.2.2. Axions

Because a neutron contains both a positive charge (one up quark) and negative charges (two down quarks), we should expect it to have a permanent electric dipole moment (EDM). This permanent EDM will violate T-symmetry, so it should also violate CP-symmetry in order to preserve CPT. Thus, a neutron EDM (and hence CP-violation) is to be expected with contributions from both the weak and strong interactions. Estimations of these contributions can be computed [28], with the weak force comprising approximately  $10^{-5}$  of the best experimental upper limit of  $|d_n| 3.0 \times 10^{-26} e \text{ cm}$  [29]. In stark contrast to the weak force, the QCD Lagrangian contains a CP-violating term which, without any fine-tuning, should result in an EDM approximately  $10^{10}$  times larger than current experimental limits.

The fact that strong CP-violation appears to be drastically weaker than expected is known as the “strong-CP problem”. In 1977, Pecci and Quinn introduced a new broken symmetry which would suppress the CP-violating term [30]. Wilczek [31] and Weinberg [32] noted that this would give rise to the axion, a particle which had—at the time—an MeV scale mass. Constraints from experiments [33] and astrophysical observations [34] have radically lowered the mass and interaction rates of the axion (Figure 1.5), to the point that they can be considered candidates for dark matter. At a mass range of  $m_a \sim 10^{-6}$  to  $10^{-3}$  eV, primordial production of axions would result in them being ruled out as hot dark matter, but alternative production mechanisms have been proposed which would result in cold populations [35].




Figure 1.5.: A simplified snapshot of constraints on the axion-photon coupling strength [36]. Other couplings are possible (e.g. axion-electron) and plots for those couplings can be found at the same source.

### 1.2.3. Weakly Interacting Massive Particles (WIMPs)

WIMPs emerged as a strong contender to solve the dark matter problem because their existence was predicated on the much broader theory of supersymmetry (SUSY). The theory was originally developed to unify the fundamental nature of fermions and bosons behind a new broken symmetry. A basic proposition of the theory is that all fundamental particles have “superpartners”; each fermion is paired with a “superboson”, and each boson with a “superfermion”. This resulted in a new palette of stable, neutral, massive, and weakly interacting particles that led to the consideration of their cosmological implications throughout the 1980’s [37] [38]. Of these par-

ticles, the neutralino emerged as one of the most promising candidates [39] because it is typically the lightest particle in most models, and therefore expected to be the most stable. It is a linear combination of the photino, Z-ino, and Higgsino, the superpartners of the photon, Z-boson, and Higgs, respectively.

A general case that favors WIMPs (taken from [39]) can be built by first considering a calculation of the relic abundance of particle dark matter, where the particle is denoted as  $\chi$ . First, we assume that in the early universe a population of  $\chi$ 's is in thermal equilibrium with its environment. That is, at a temperature  $T \gg m_\chi$  (where  $m_\chi$  is the mass of  $\chi$ ),  $\chi$ 's are freely annihilated and pair-produced to give a number density of  $n_\chi^{eq} \propto T^3$ . At temperatures where  $T \ll m_\chi$ ,  $n_\chi^{eq}$  picks up a Boltzmann suppression factor and becomes:

$$n_\chi^{eq} \simeq g_\chi (m_\chi T / 2\pi)^{3/2} e^{-m_\chi/T} \quad (1.2)$$

where  $g_\chi$  is the number of internal degrees of freedom of the particle.

The  $\chi$  annihilation rate is given by  $\Gamma = \langle \sigma_A v \rangle n_\chi$  where  $\sigma_A$  is the annihilation cross-section,  $v$  is the relative velocity of a  $\chi\bar{\chi}$  pair, and the angled brackets denote a thermal average. When  $T < m_\chi$ ,  $n_\chi$  will be suppressed, and the expansion rate of the universe  $H = \dot{a}/a$  will exceed  $\Gamma$ . Here,  $a$  is a length scale of the universe, and the dot denotes a time derivative. The combination of a suppressed population, and expanding universe leads to a relic density of particles which freeze-out of thermal equilibrium with the rest of the universe.

This scenario can be quantitatively modeled by considering the time evolution of  $n_\chi$ . In the simplest case with no annihilation or creation,  $n_\chi \propto a^{-3}$ . That is, the number density will decrease as the universe expands. Taking a time derivative of this, we have

$$\dot{n}_\chi \propto -3a^{-4}\dot{a} \quad (1.3)$$

Then, multiplying the right-hand side of Equation 1.3 by  $n_\chi/a^{-3}$  gives:

$$\dot{n}_\chi \propto -3Hn_\chi \quad (1.4)$$

We can incorporate two additional terms to account for the mutual annihilation and creation of particles. To emulate a freeze-out process, we should expect that the rate of annihilated particles per unit volume will exceed that of the created particles. Then, from the static scenario in Equation 1.4 we have:

$$\dot{n}_\chi = -3Hn_\chi - \langle\sigma_A v\rangle[(n_\chi)^2 - (n_\chi^{eq})^2] \quad (1.5)$$

To interpret this, consider that the current number density should dictate the annihilation rate. As annihilation occurs, SM particles rapidly thermalize and fall below the threshold for pair production. There will be some exceptions, and these will form a backstop which gives rise to the third term in Equation 1.5. Finally, as  $n_\chi$  approaches  $n_\chi^{eq}$ ,  $\dot{n}_\chi$  is again solely dictated by the expansion of the universe. There is no closed form solution for Equation 1.5, but a rough approximation can be formulated where  $\langle\sigma_A v\rangle$  is assumed to be energy independent. For a plot of numerical solutions, see Figure 1.6.

Freeze-out likely occurred at some time in the early universe, when the critical density was radiation dominated. This grants us an expansion rate of  $H(T) = 1.66T^2\sqrt{g_*}/m_{pl}$  where  $m_{pl} \sim 10^{19}$  GeV is the Planck mass and  $g_*$  is the effective number of relativistic degrees of freedom. Freeze-out




Figure 1.6.: A numerical solution of Equation 1.5 demonstrating the freeze-out process where  $Y = n_\chi/s$  and  $x = m_\chi/T$ . Solid curves approach a relic density as the temperature of the universe  $T$  decreases. An increasing dark matter annihilation rate  $\Gamma$  (indicative of a stronger SM coupling) lowers this relic density. The dashed curve tracks the abundance of dark matter particles that remain in thermal equilibrium. From [40].

occurs when  $\Gamma(T_f) = H(T_f)$ , i.e. when:

$$n_\chi = \frac{1.66 T_f^2 \sqrt{g_*}}{m_{pl} \langle \sigma_A v \rangle} \quad (1.6)$$

At the weak scale,  $T_f \simeq m_\chi/20$  with a small dependence on  $m_\chi$  and  $\sigma_A$ . Furthermore, it is expected that the entropy density per co-moving volume remains constant. The entropy density  $s$  is approximately  $0.4g_*T^3$  so that  $n_\chi/s$  is also constant. Dividing both sides of Equation 1.6 by  $s$  and substituting in  $m_\chi/20$  for  $T_f$ , we obtain:

$$\frac{n_\chi}{s} \simeq \frac{100}{m_\chi m_{pl} \sqrt{g_*} \langle \sigma_A v \rangle} \simeq \frac{10^{-8}}{(m_\chi) (\langle \sigma_A v \rangle)} \quad (1.7)$$

where  $m_\chi$  is in units of GeV and  $\langle \sigma_A v \rangle$  is on the order of  $10^{-27} \text{ cm}^3 \text{ s}^{-1}$ . The current entropy density is roughly  $4000^{-3}$ , and the critical density is  $\rho_c \simeq 10^{-5} h^2 \text{ GeV cm}^{-3}$  with the Hubble constant  $h \sim 100 \text{ km s}^{-1} \text{ Mpc}^{-1}$ . In terms of the critical density, the present mass density of dark matter is then expected to be:

$$\Omega_\chi h^2 = \frac{m_\chi n_\chi}{\rho_c} \simeq \frac{3 \times 10^{-27} \text{ cm}^3 \text{ s}^{-1}}{\langle \sigma_A v \rangle}$$

For particles interacting on the weak scale, and at non-relativistic speeds, a volumetric annihilation rate of  $\langle \sigma_A v \rangle \sim 10^{-25}$  gives an estimate of  $\Omega_\chi h^2 \approx 0.03 \text{ GeV s}$ . This is strikingly close to the observed value of  $\Omega_{CDM} h^2 = 0.120$  [14] which strengthens the case that dark matter particles interact at the weak scale<sup>2</sup>.

#### 1.2.4. Sub-GeV Dark Matter

As early as 1977, Lee and Weinberg concluded that a lower bound of  $m_\chi > 2 \text{ GeV}$  was necessary to avoid an overabundance of WIMPs [41]. In simpler scenarios, such as the case proposed above, thermal relics with a mass below 10 GeV have largely been ruled out by CMB observations [42], leading to a communal focus on WIMPs that span a mass range from 10 GeV to a few TeV [43].

More recently however, detection limits for classical WIMPs have strengthened. This has placed significant strain on a hypothetical freeze-out process because a coupling to SM particles

---

<sup>2</sup>This apparent coincidence is also known as the ‘‘WIMP miracle’’.

which is too weak will result in a dark matter population that does not thermalize with SM particles [40]. Consequently, models of Feebly Interacting Massive Particles (FIMPs) that rely on a *freeze-in* process have seen renewed interest [44], especially those with a mass in the sub-GeV range.

As summarized in [40], freeze-in assumes a negligible initial number density of dark matter particles. These particles are then produced by decays or annihilations of standard model particles until the number density of parent particles freezes-out, resulting in a non-thermal population of dark matter particles that is said to have frozen-in as the universe cools (Figure 1.7).

This is in direct contrast with freeze-out, where the relic density is inversely related to the annihilation rate (and therefore the dark matter coupling strength to SM particles). A more concrete example of this can be seen in the predicted relic density of a simple case where a SM particle decays into dark matter particles [40]:

$$\Omega_\chi h^2 = 4.48 \times 10^8 \frac{g_\sigma}{g_{*s} \sqrt{g_*}} \frac{m_\chi}{\text{GeV}} \frac{M_P \Gamma_{\sigma \rightarrow \chi\chi}}{m_\sigma^2}$$

where  $\Gamma_{\sigma \rightarrow \chi\chi}$  is the production rate from a decaying SM particle  $\sigma$ ,  $m_\sigma$  and  $m_\chi$  are the respective masses of the SM and dark matter particles, and  $M_P$  is the Planck mass. The factors  $g_\sigma$ ,  $g_{*s}$  and  $g_*$  are the degrees of freedom intrinsic to the particle  $\sigma$ , and those related to the entropy and energy densities, respectively.

Like WIMPs, many FIMP models exist and they commonly describe an entire “dark sector” of new particles where new mediators act as “portals” to candidates at the MeV scale. Of special note here is the sterile neutrino, which was proposed as dark matter candidate in 1994 by Dodelson and Widrow [45]. This massive, right-handed neutrino could couple strictly to SM neutrinos via an additional degree of mixing and would complement its nearly-massless left-handed SM cousins.




Figure 1.7.: A numerical solution of the Boltzmann equation demonstrating the freeze-in process where  $Y = n_\chi/s$  and  $x = m_\sigma/T$ . Solid curves approach a relic density as the temperature of the universe  $T$  decreases. An increasing dark matter production rate  $\Gamma$  increases the relic density. The dashed curve tracks the abundance of SM particles that remain in thermal equilibrium. From [40].

While the original production mechanism has been ruled out, sterile neutrinos produced via a freeze-in mechanism have not been excluded [40].

### 1.3. Detection Methods

A broad spectrum of dark matter candidates exists far in excess of what has been covered above (see [23] [5] [40]). A complementary set of experiments and proposals likewise exists [43] [44] in order to better establish the existence of said candidates. This section will briefly cover two of the principal particle dark matter detection methods: production, and indirect detection. Then, it will focus on the direct detection of WIMPs and lighter dark matter. For a convenient diagram of all three techniques, see Figure 1.8.

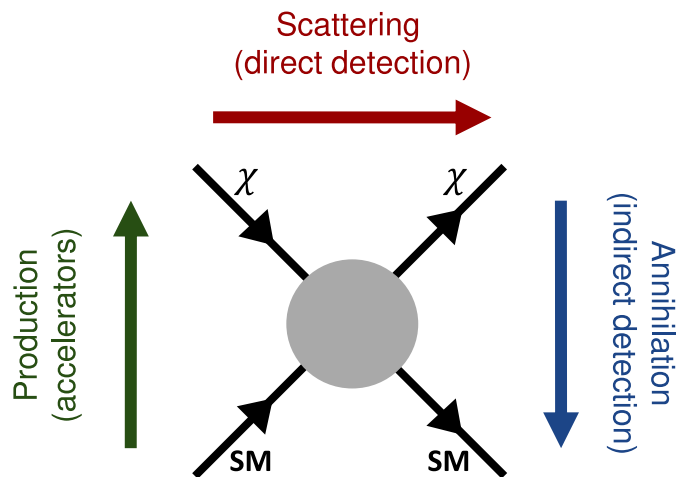


Figure 1.8.: Detection methods for particle dark matter rely on three distinct classes of particle interaction; each are denoted with a corresponding arrow which signifies the flow of time for that process [46].

### 1.3.1. Production

Similar to searches for the Higgs boson, or for events with neutrinos, WIMPs or FIMP mediators produced in a collision would not be directly detected. Instead, events that appear to violate momentum conservation can be used to identify new particles. As collision energies at accelerators increase, dark matter candidates can be further constrained. Furthermore, searches at colliders tend to be mass-independent for a given model because of the agnostic nature of collisions. Comparisons to models can be made by first searching for specific reactions, and then translating results into limits on parameters typically discussed in direct detection experiments. An example of this is seen in Figure 1.9. For further information on this topic, a review on dark matter searches at the LHC is suggested [47].

### 1.3.2. Indirect Detection

Indirect detection is a broad term, especially when the history of dark matter is considered. However as constraints have tightened, more targeted efforts involve cosmological gamma ray or neutrino surveys that set limits on dark matter annihilation rates as well as production rates from SM decays. Fermi-LAT is a prominent example of this type of search; surveys conducted have looked at the gamma ray flux from the center of the Milky Way [49], as well as a number of dwarf spheroidal galaxies [50]. While simulations predict the highest concentration of dark matter to be in the center of galaxies, the signal-to-noise ratio is poor because of interference from sources in the rest of the galactic plane. Dwarf spheroidal galaxies are a suitable counterpart to this because they are the most dark-matter dominated galaxies observed, with very high mass to luminosity ratios [51]. Recent limits on the dark matter annihilation cross-section from observations of dwarf




Figure 1.9.: Observations of specific processes at the LHC can help set limits on various dark matter models, and can be compared to results from direct detection experiments. Furthermore, the model-independent nature of collisions grants the ability to sweep larger mass ranges of parameter space in comparison to more targeted searches [48].

spheroidal galaxies are shown in Figure 1.10.

### 1.3.3. Direct Detection

Single particles of any kind can be directly detected via scintillation light, ionized charge, or excess heat (e.g. phonons) if they scatter off of and deposit energy in a target. Thus, the choice of a target defines a direct detection experiment, making the sensitivity of an experiment both detector and model specific. An overview of experiments which utilize various combinations of the three




Figure 1.10.: Fermi-LAT limits on the dark matter annihilation cross-section from a gamma ray survey of 15 dwarf spheroidal galaxies. Limits are set by searching for gamma rays from the decay of products of specific DM annihilation channels. Observed spectra are compared to simulated spectra, which are swept over a range of DM masses that roughly correspond to WIMP models [52]. The thermal relic cross-section is shown as a dashed gray line for the sake of comparison.

detection channels is shown in Figure 1.11, however a complete description of these technologies is outside of the scope of this discussion. For a review of these topics, see [53].

### 1.3.3.1. WIMP Interaction Rates in a Detector

In brief, a successful WIMP detection experiment must be able to measure a discernible number of events that correspond to a particle (or particles) of unknown origin depositing energy in a detector. That is, the experimental observable in a direct detection experiment is the scattering rate, which is a function of the target's recoil energy. For a WIMP, the differential scattering rate with respect to the recoil energy of a target nucleus in a detector is given by [55] as

$$\frac{dR}{dE_R} = \frac{\rho_\chi \sigma_0 F^2(E_R)}{2m_\chi \mu_\chi^2} \int_{v_{min}}^{v_{max}} \frac{f(\mathbf{v})}{v} d^3v \quad (1.8)$$

With the following nomenclature in Table 1.1.




Figure 1.11.: Prominent dark matter direct detection experiments, grouped by the detection technology used and measured signals. Single phase detectors and dual-phase time projection chambers (TPCs) utilize noble elements such as xenon or argon. Measuring exclusively charge signals, DAMIC uses silicon wafers while NEWS-G uses a mixture of neon and methane gas. Gas TPCs are constructed to search for a directional WIMP “wind”. Scintillating bolometers, cryogenic detectors, and scintillating crystals all utilize crystals of various compositions to probe different WIMP mass ranges. Each experiment under these categories has unique hardware suitable for reading out the desired quantities. PICO and DRIFT are notable for using fluoridated compounds to set the best spin-dependent WIMP interaction limits. Figure from [54].

Table 1.1.: Nomenclature for Equation 1.8

Variable	Significance
$R$	Scattering rate
$E_R$	Nuclear recoil energy
$\rho_\chi$	Local WIMP density

Variable	Significance
$\sigma_0$	WIMP interaction cross-section
$F(E_R)$	Nuclear form factor
$m_\chi$	WIMP mass
$\mu_\chi$	WIMP-nucleus reduced mass
$f(\mathbf{v})$	Local WIMP velocity distribution
$v$	WIMP velocity

Parameters pertaining to the WIMPs are model-dependent, and can also be constrained by astrophysical observations. Following those, the limits on the integral over the velocity distribution are established. First, because the thermal relic population has completely frozen-out,  $v_{max}$  must be less than the galactic escape velocity. Second,  $v_{min}$  must be great enough such that WIMPs will impart a recoil energy to a target nucleus of mass  $m_N$  that meets or exceeds a detector's threshold  $E_{th}$ . The constraint on  $v_{min}$  can be determined classically, since the relative velocity between WIMPs and a terrestrial target must be non-relativistic. Provided that the scattering is elastic, the recoil energy of a nucleus in the lab frame is given by [56] as

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

where  $E_i = 1/2 m_\chi v^2$  and

$$r = \frac{4m_\chi m_N}{(m_\chi + m_N)^2}$$

is a kinematic factor which is always less than 1, except when  $m_\chi = m_N$  in which case  $r = 1$ . In the case of a perfect backscatter where  $\theta = \pi$ ,  $E_R = E_i r \geq E_{th}$ , so  $v_{min} = \sqrt{2E_{th}/rm_\chi}$

As in the case above, the remaining parameters in Equation 1.8 can be “tuned” via the selection of a target nucleus, which effectively determines the technology used in the construction of the experiment and the sensitivity to a given WIMP mass. Prior to addressing the qualities of the cross-section, it should be noted that the scattering process has been assumed to be coherent. A coherent nuclear scattering process is one in which the incoming particle scatters uniformly off of all nucleons, and is expected when the deBroglie wavelength  $\lambda_b = h/p$  of the incoming particle exceeds the radius of a target nucleus. Consequently, coherent scattering is also regarded to be an elastic process. At higher energies incoherent scattering is expected because  $\lambda_b$  is lower and individual nucleons can be probed by the incoming particle. This is typically an inelastic process because the interacting nucleons are usually left in an excited state. In both scattering regimes, momentum is transferred from the incoming particle to the nucleus, necessitating the use of a nuclear form factor to encode the spatially varying density of the nucleus and its response to an input momentum. The authors of [57] note that (for neutrinos) the scattering cross-section for each process differs, with a quadratic dependence on atomic mass for coherent scattering and a linear dependence for incoherent scattering. The transition between these two regimes is smoothly driven by complementary nuclear form factors; for coherent scattering a form factor  $F^2$  is used, and  $1 - F^2$  is used for incoherent scattering. The form factor is normalized such that  $F^2(E_R = 0) = 1$ , with higher momentum transfers resulting in a lower value of  $F$  and increasingly incoherent scattering. Because WIMPs are especially non-relativistic, coherent nuclear scattering can be expected, and the incoherent form factor is disregarded for this calculation.

The cross-section is given in [55] with two components, to account for spin-independent and spin-dependent interactions with the target nucleus. The spin-independent cross-section for a WIMP interacting with a nucleus composed of  $Z$  protons and  $(A-Z)$  neutrons is

$$\sigma^{SI} = \frac{\mu_\chi^2}{\pi} |ZG_{SI}^p + (A-Z)G_{SI}^n|^2$$

where  $G_{SI}^p$  and  $G_{SI}^n$  denote the scalar-mediated WIMP-nucleon coupling strengths to protons and neutrons respectively. It is not uncommon to assume these couplings are roughly equal in strength in order to further simplify the spin-independent cross-section so that

$$\sigma^{SI} \simeq \frac{\mu_\chi^2 A^2}{\pi} |G_{SI}|^2$$

and the dependence of the target mass on the cross-section is readily apparent. High mass targets are not always desirable however, as the contribution from the reduced mass is maximized when the WIMP and target masses are equal. The nuclear recoil energy is likewise maximized in the case of matching scattering masses, which is critical when detectors have limited thresholds which dictate when an event is registered. For an example of the expected interaction rates, see Figure 1.12.

Finally, the spin-dependent cross-section at zero momentum transfer is

$$\sigma^{SD} = \frac{4\mu_\chi^2}{\pi} \frac{J+1}{J} |\langle S_p \rangle G_{SD}^p + \langle S_n \rangle G_{SD}^n|^2$$

where  $J$  is the nuclear spin,  $\langle S \rangle$  is the expectation value of the spin for either protons or neutrons in the nucleus, and  $G_{SD}$  is the axial coupling between a WIMP and a nucleon. The lack




Figure 1.12.: The above plot shows simulated scattering rates for a 100 GeV WIMP with a spin-independent cross section of  $10^{-45}$  cm<sup>2</sup> and the kinematic properties of a typical thermal relic. At lower nuclear recoil energies, xenon is an excellent choice of target material due to its high mass. However at higher nuclear recoil energies, scattering becomes incoherent as a WIMP interaction would only probe a portion of the xenon nucleus. Thus, lower recoil thresholds enable detectors to benefit from coherent scattering amplification of the nuclear interaction cross-section. Circles on each curve indicate typical detection thresholds. Taken from [58].

of atomic mass dependence in the spin-independent cross-section is noteworthy in that it serves to further differentiate experiments which probe a spin-dependent WIMP coupling. Such experiments will necessarily take nuclear spins into higher consideration than nuclear mass.

### 1.3.3.2. Dark Matter Electron Scattering

As shown in Figure 1.12, interaction rates for a 100 GeV WIMP are significantly improved provided that low energy (few keV) nuclear recoils are detectable. For sub-GeV dark matter candidates however, nuclear recoils are below detection thresholds with existing technology (see Section 2.3). As given in [59], a light dark matter particle can transfer on average:

$$E_{NR} = \frac{q^2}{2m_N} \approx 1 \text{ eV} \times \left(\frac{m_{DM}}{100 \text{ MeV}}\right)^2 \left(\frac{10 \text{ GeV}}{m_N}\right)$$

in an elastic nuclear recoil, assuming a momentum transfer  $q \approx m_{DM}v$  with the dark matter velocity  $v \approx 10^{-3}c$ . For dark matter particles with mass  $\mathcal{O}(100 \text{ MeV})$  and detection target with  $m_N \approx \mathcal{O}(100 \text{ GeV})$ , a nuclear recoil of  $\approx 0.1 \text{ eV}$  would be impossible to detect even with a best-case threshold of  $\mathcal{O}(100 \text{ eV})$ .

With nuclear recoils so strongly disfavored, other interaction channels should be considered. Given that the kinetic energy of a dark matter particle is:

$$E_{tot} = \frac{m_{DM}v^2}{2} \approx 50 \text{ eV} \times \frac{m_{DM}}{100 \text{ MeV}}$$

a low-mass candidate should be capable of ionizing, exciting, or dissociating a target. The case of ionization in xenon is of particular interest for this work; estimated dark matter ionization rates from [59] are shown in Figure 1.13.

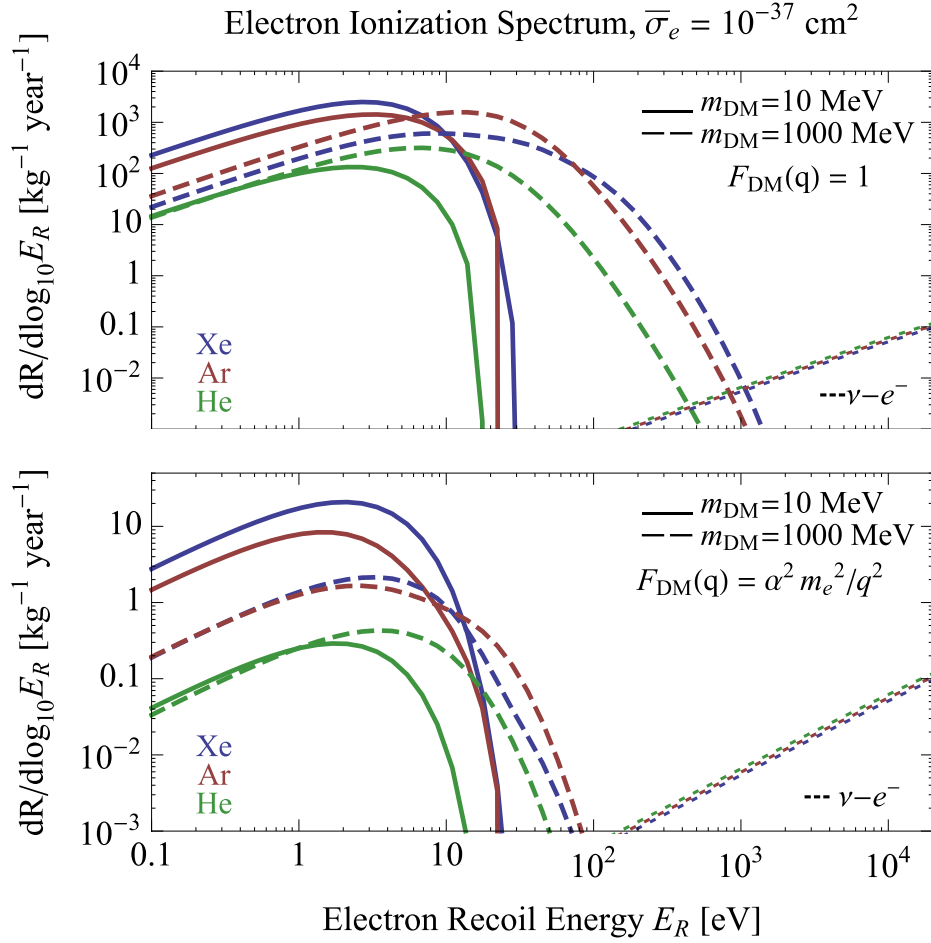


Figure 1.13.: Differential ionization rates versus electron recoil energy for different target materials [59]. Notable here is that for events below a few hundred eV, the irreducible solar neutrino background is not expected to meaningfully obscure signals.

### 1.3.3.3. Recent WIMP Detection Limits

The distinct characteristics of different detectors invariably make direct detection experiments more suitable for targeted searches that appeal to a relatively narrow range of dark matter models. The benefit of a more limited physics scope is an increased level of sensitivity; direct detection experiments are responsible for probing specific regions of dark matter masses to a greater extent than other methods (see Figure 1.14, Figure 1.5).




Figure 1.14.: Recent limits set by WIMP direct detection experiments [60]. The neutrino fog is where an irreducible background caused by coherent scattering of solar, atmospheric, and supernova neutrinos (Section 2.3.1) is expected to become dominant. Searches in this region require a  $10^n$  increase in exposure for in order to obtain a corresponding factor of 10 increase in cross-section sensitivity.

Most relevant to this work are dual-phase time projection chambers (TPCs). TPCs are currently responsible for setting the strongest limits on the spin-independent WIMP-nucleon cross-

section for WIMP masses in the range of 5 GeV to a few hundred TeV. TPCs commonly use xenon or argon as their detection medium, with each having their own unique benefits and drawbacks. At higher WIMP masses, xenon TPCs are notable for consistent improvement in setting limits (see Figure 1.15).




Figure 1.15.: A chart highlighting the increasing levels of sensitivity reached by contemporary WIMP search experiments. Different colors indicate different technologies used. Note that some experiments pictured are more sensitive to WIMPs in a significantly lower mass range. See [61] for the figure shown here and [62] for a previous version.

## Chapter 2.: Dual-Phase Noble Element Time Projection

### Chambers

#### 2.1. Theory of Operation

Figure 2.1 shows a mockup of an ideal single scatter in a dual-phase TPC. A dual-phase TPC is primarily liquid, with a gaseous layer on top. The liquid volume is the target medium, where incoming particles scatter and ionize a small region in the liquid. In the scattering process numerous excited dimers (excimers) are formed [53], which de-excite and release vacuum ultraviolet (VUV) photons. This prompt “S1” light pulse is read out by arrays of photomultiplier tubes (PMTs) or silicon photomultipliers (SiPMs) at the top and bottom of the TPC.

Some ionized electrons do not recombine in the liquid; they are instead drifted upward with an electric field. This field is created using wire grids which are maintained at a high voltage. After being drifted by the field, the electrons are eventually extracted into the gas where they collide with atoms and excite them. This process forms more excimers, which de-excite and generate a new “S2” pulse [53].

In a simple TPC, only two grids are used: an anode in the gas layer and a cathode in the liquid above the bottom PMTs. For more control over the fields, an additional “gate” grid is used.

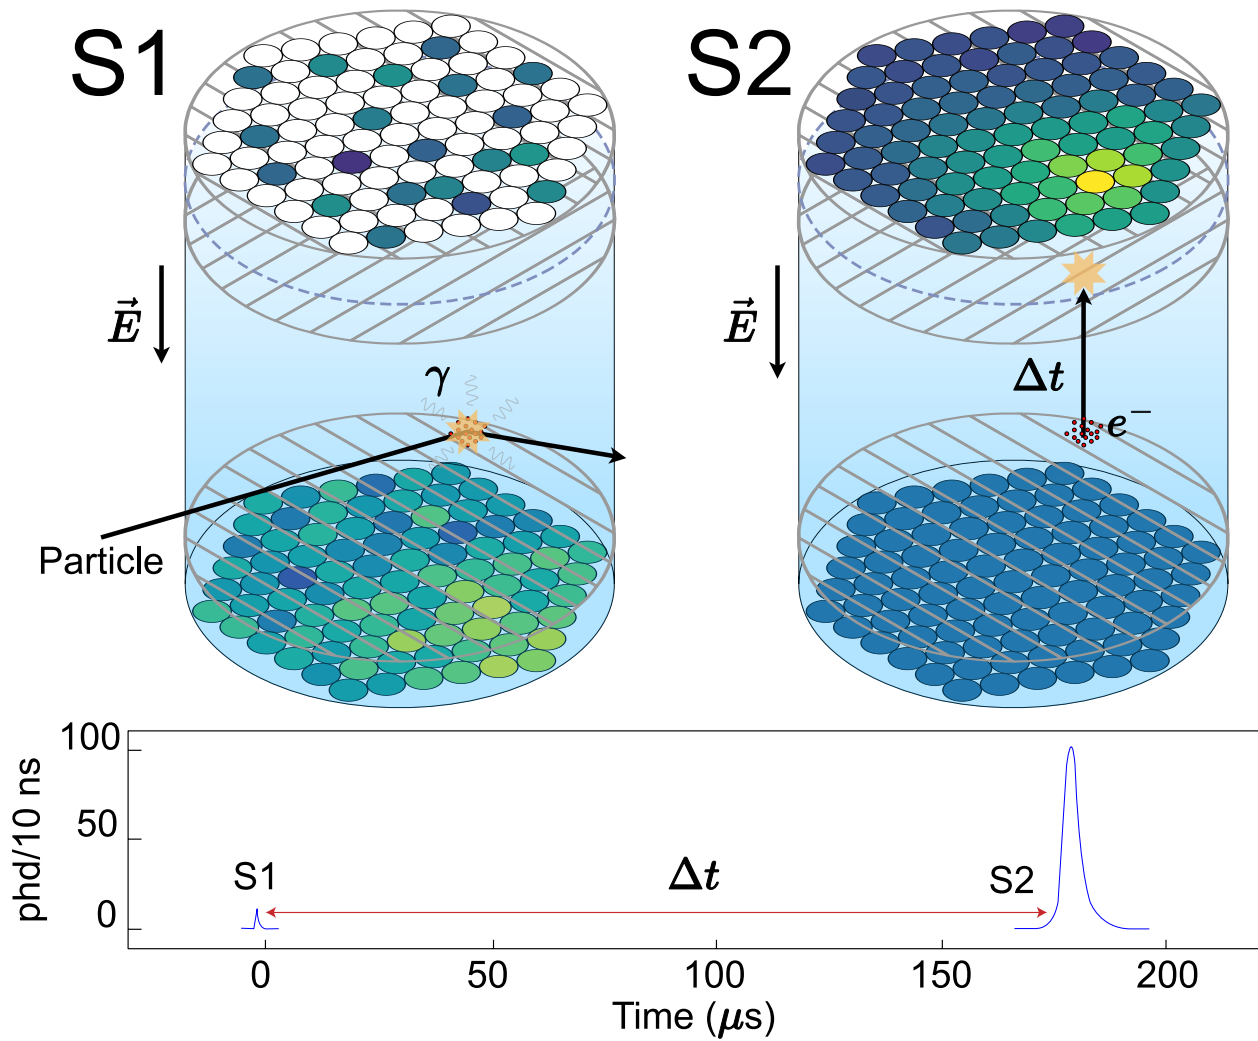
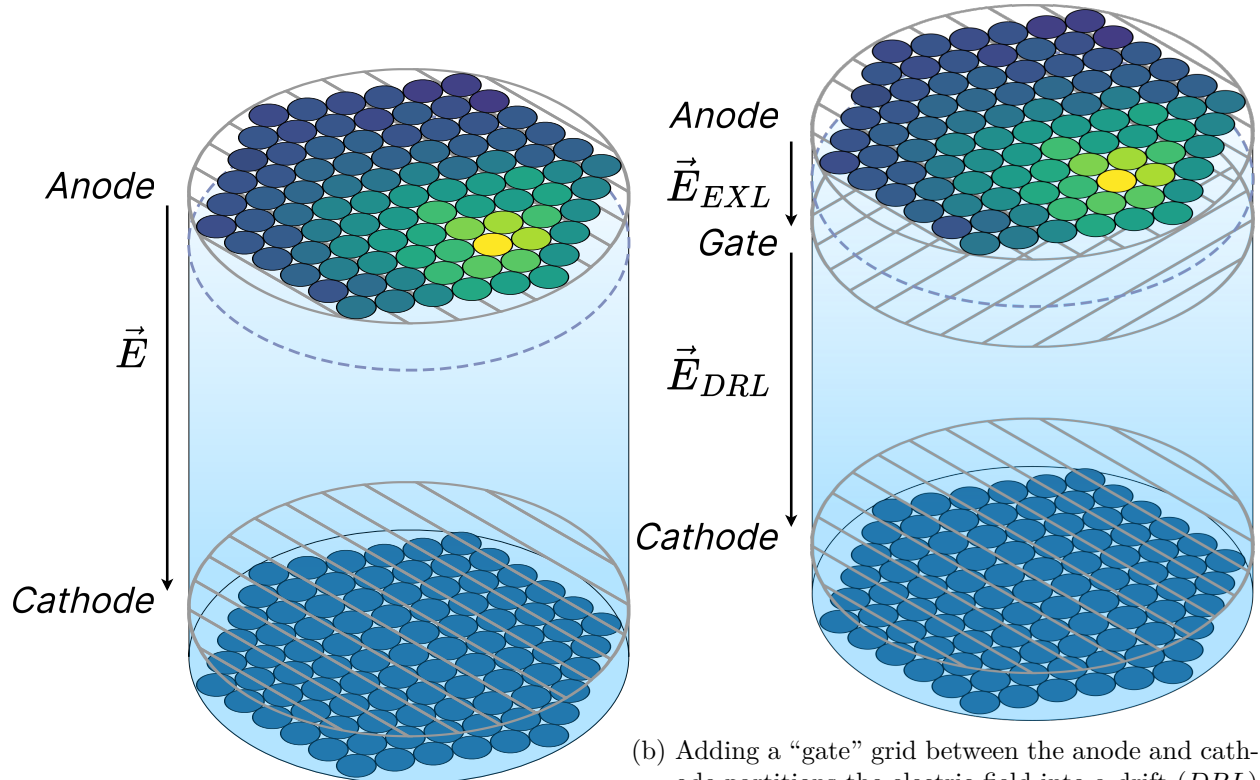


Figure 2.1.: Mockup of a TPC and event with S1 and S2 pulses. A particle interaction produces light ( $\gamma$ ) which is detected as the S1 pulse at time  $t = 0$  in an event window. Some electrons are also ionized in the initial interaction; an external electric field  $\vec{E}$  prevents some of them from recombining by drifting them upwards. After some drift time  $\Delta t$ , the electron cloud reaches the liquid surface where the electric field extracts them from the liquid. Upon extraction, the electrons generate light via electroluminescence in the gas, which is detected as the S2 pulse in an event window. Original image courtesy of C. Faham and D. Malling.

The gate is situated below the liquid surface and partitions the TPC into a “drift” (below the gate) and “extraction” (above the gate) region. This is illustrated in Figure 2.2. The addition of the gate allows the drift and extraction fields to be independently tuned by changing the anode, gate, and cathode voltages. This is especially useful in a large TPC where high extraction fields are more difficult to maintain longer drift lengths.



(a) A simple two-grid TPC has a single electric field which drifts electrons through the liquid, and extracts them once they reach the liquid surface.

(b) Adding a “gate” grid between the anode and cathode partitions the electric field into a drift (*DRL*) and extraction (*EXL*) field. Because the gate is much closer to the anode, it is less challenging to maintain a strong extraction field. In this configuration, an extraction field can typically be 10-20x stronger than the drift field.

Figure 2.2.: Comparison of a simple two-grid TPC (Figure 2.2a) and a three-grid TPC (Figure 2.2b). The three-grid TPC has independently tunable drift and extraction fields. These are more suitable for TPCs with larger drift lengths where a high extraction field may be otherwise difficult to maintain.

While photosensor arrays provide lateral position information on an event, a known electron


drift velocity is used in tandem with the time delay between S1 and S2 to vertically locate an event. This 3-D event position reconstruction is a useful feature for rejecting external backgrounds because events located at the edges of the detector can be rejected, while events inside the quieter “fiducial volume” are preserved. For events inside the fiducial volume, derived quantities from S1 and S2 pulses can help make further cuts and identify events. For example, in xenon the pulse area ratio of S2/S1 is commonly used to discriminate between electron and nuclear recoils Section 2.2.4. In argon, pulse shape discrimination (PSD) is used for recoil discrimination; it is also be especially valuable in an “S2 only” analysis (see Section 2.3.2) where S1 pulses are too weak to be detected. Pulse widths, rise times, and area fractions contained within those parameters are examples of useful quantities for a PSD analysis.

## 2.2. Signal Production


Light is the basis for any signal in a dual-phase TPC. In both argon and xenon-based TPCs, light is produced by the de-excitation of excimers. These excimers require an initially excited atom  $R^*$  to form, which itself is created in two ways: direct excitation, and recombination of ionized dimers with free electrons. A detailed review of these processes is given in [58] and partially reproduced here.

### 2.2.1. Excitation


The excitation process starts with an excited atom, produced via interaction with a passing particle, like an electron or recoiling nucleus. The excited atom then forms a dimer with a neighboring atom, with the dimer being in an electronically and vibrationally excited state:



The excimer then undergoes a non-radiative relaxation process where it is left in a state of only electronic excitation:




And finally it dissociates, releasing a VUV photon in the process.



### 2.2.2. Recombination

The recombination process to produce light involves the capture of electrons, making it more dependent on an applied electric field. It begins with the formation of a dimer from a pair of atoms, one of which is ionized:




Note that this is more likely to occur than immediate recombination of ions with charges because the spacing of atoms in the liquid is on the order of 0.5 nm, whereas the thermalization distance of an electron is a few microns. Thus from a statistical standpoint, the majority of neighbors for a xenon atom are other xenon atoms, not free electrons<sup>1</sup>.


---

<sup>1</sup>There is additional nuance here because a tight track structure left behind by a particle results in electrons being unable to escape the cloud of positive ions. Whereas in a loose track structure, it is easier for electrons to


The ionized dimer recombines with a free electron, promptly dissociating into a doubly-excited atom and ground-state atom:




The exact nature of the double-excited state is unclear, and does not appear to be defined [63]. Doubly-excited states typically refer to two electrons each being excited [64], but also do not preclude a single electron being excited two levels. In the latter case, the de-excitation spectrum is expected to be more complex. Following this, the doubly-excited atom releases energy in the form of heat ( $Q$ ):



From here, the recombination process bears resemblance to excitation, with the excited atom forming an excited dimer and releasing some heat in the process:



Finally, the excited dimer emits a VUV photon:




---

escape—even via diffusion without an external electric field. Still, a tight track structure does not imply perfect recombination because of bi-excitonic quenching. This occurs when two excited atoms collide to produce one ion and one atom, both in their respective ground states, resulting in fewer excimers being formed. Overall, the subject of recombination is a very complex topic with plenty of unanswered questions. See Section 4.2.2 in [58] for a more detailed review of this subject.

### 2.2.3. Scintillation and Electroluminescence

The excitation and recombination processes for both xenon and argon terminate with the emission of VUV photons from an excimer in either a singlet ( $^1\Sigma_u^+$ ) or triplet ( $^3\Sigma_u^+$ ) state. The respective lifetimes of these states for each element (Table 2.1) are one of a number of properties which distinguish xenon and argon as scintillators.

Table 2.1.: Light emission properties of liquid xenon and argon. Overall scintillation response times are dependent on the proportion of states created with excitation and recombination, which are themselves dependent on the concentration of impurities, initial interactions, and applied electric field.

Property	Xe	Ar
Triplet Lifetime (ns)	22	1600
Singlet Lifetime (ns)	4.3	7
Wavelength (nm)	178.1	129.6

With the application of an external electric field, recombination is suppressed as electrons can be drifted away from ions and extracted into a gaseous layer where they undergo electroluminescence by colliding with and exciting atoms. At lower gas pressures, wavelengths are shorter because excited atoms can fully decay to their ground state. As the pressure increases, the collision frequency between atoms increases, resulting in more excimer formation until around 1 bar where only decays from the singlet and triplet states can occur. In this regime, emitted photon wavelengths are shifted down only slightly to approximately 171 nm and 128 nm in xenon and argon due to an increase of the excimer energy levels in gas.

Until around 10 bar, the number of photons generated for each extracted electron traversing

a centimeter in the gas phase is linear with the electric field  $E$  (V/cm) and pressure  $P$  (bar):

$$\frac{dN_{ph}}{dx} = \alpha E - \beta P - \gamma$$

This lowers the overall scintillation light yield but in exchange, electroluminescence generates a few hundred photons per electron. This leads to significant benefits like improved fiducialization in all three dimensions, better determination of recoil energies, and a lower interaction detection threshold.

#### 2.2.4. Recoil Discrimination

The exact proportion of excitation and recombination processes depends on factors like the initial energy of the interacting particle, the identity of the particle, and whether an interaction initially occurred with an electron or nucleus. The last case is especially important: often times dark matter models predict interactions with *either* a nuclei or electrons, so being able to distinguish between the recoil type is a powerful method for background rejection.

In both argon and xenon, nuclear and electron recoils result in different proportions of excitation and recombination. Because the excitation and recombination processes produce different fractions of singlet and triplet excimers, ERs and NRs can have S1 signals with distinguishable timescale. Additionally, the extent of recombination that occurs will define an S2 pulse because electrons which fail to, or are prevented from recombining will eventually form an S2 pulse. This leads to two methods used to perform ER-NR discrimination: S1 PSD, and analysis of the ratio of S2 and S1 pulse areas in an event.

Since the triplet in argon is very long-lived relative to the singlet, ER-NR discrimination is

achieved with PSD of the S1 signal by analyzing how quickly it decays. This technique is very effective at electron recoil energies above  $\sim 10$  keV, affording a nominal ER leakage fraction (the fraction of ER events misclassified as NRs) below  $10^{-10}$  [65]. Below these energies however, it becomes necessary to incorporate measurements of S2 signals because S1 PSD quickly becomes background-dominated. This has proven to be challenging to accomplish in argon. Low energy nuclear recoil measurements have produced inconsistent results [66], and discrimination power using S1 and S2 has not been demonstrated below 20 keV [67]. Reasons for this are unclear [66] but one possibility could come from the need for wavelength shifter coatings in argon which cause a strong position dependence in S2 size, believed to be caused by unevenness in the coatings [68].

In xenon, the similar lifetimes of the singlet and triplet excimers make S1 PSD more difficult [69]. Instead, the ratio of S2 and S1 pulse areas are used to distinguish between ERs and NRs. This works reasonably well at lower energies, with ER leakage at 50% NR acceptance usually on the order of  $10^{-3}$  or better. In xenon, nuclear recoil energies at and below 1 keV have been measured (see Figure 2.4) while the lowest energy measurement of nuclear recoils to date in argon is 7 keV [70].

### 2.2.5. Recoil Energy Determination

Besides recoil classifications, many backgrounds also identifiable by their distinct recoil energy spectra. In a dual-phase TPC, recoil energy  $E_0$  is proportional to  $W_s^{min}$ , the minimum energy required to produce a scintillation photon with no external electric field. The two are related by  $N_{ph}$ , the number of scintillation photons produced in an interaction:

$$W_s = \frac{E_0}{N_{ph}} \quad (2.1)$$

Where under zero field it is assumed that  $N_{ph} = N_{ex} + N_i$ . This implies that all photons generated are produced by excitation and recombination, which further assumes that all ionized electrons recombine successfully.

In an external electric field some electrons are ionized, so the recombination fraction  $r$  is defined here as the fraction of electrons which successfully recombine. Following this is the assumption that the S1 signal is proportional to the number of scintillation photons generated:

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

And that the S2 signal is proportional to the number of electrons which do not recombine:

$$S2 \propto N_e = (1 - r)N_i$$

Based on these assumptions, it can be concluded that  $N_i + N_{ex} = N_{ph} + N_e$ , i.e. that the number of excitations and recombinations is equal to the number of photons and ionized electrons produced, regardless of recombination fraction.

Substituting back in to Equation 2.1, an expression is obtained relating electron recoil energy to the number of quanta produced:

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

For nuclear recoils, this is reduced by the Lindhard factor  $f_n$  to account for additional energy lost in the form of heat.

Critically, as these are numbers of quanta produced and not observed, the correction factors  $g_1$  and  $g_2$  are applied to relate back to  $S1$  and  $S2$ :

$$E_0 = \left( \frac{S1}{g_1} + \frac{S2}{g_2} \right) W_s^{min}$$

Together,  $g1$  and  $g2$  incorporate all of a detector's inefficiencies. Some of these are common to both signals:

- Quantum efficiency of photosensors: when a photon strikes a photosensor, the probability that a photoelectron will be generated is around 30%.
- Photosensor coverage: No photosensor has a 100% “active” surface and even then, mounting hardware and configurations will inevitably result in areas where light cannot be collected.
- Surface reflectivity: If light is not collected, then reflecting it is the next best option, in hopes that it might be collected. This has negative consequences for position reconstruction but is usually necessary to improve light collection efficiency. PTFE is a common choice of material due to its high reflectivity of VUV light from xenon, dielectric strength (making it suitable for use around high voltage components), and high inherent radiopurity. PTFE is less reflective of scintillation light in argon so it is typically coated with a wavelength shifter such as tetraphenyl butadiene (TPB) to improve reflectance. Aluminum coated with  $MgF_2$  is another option though it sees less use because it is electrically conductive and thus more challenging to implement around high voltage components.

$g_1$  can be thought of as the probability that a scintillation photon will actually be observed; it is essentially the ultimate detection efficiency of scintillation photon in the detector. It is usually valued around 10% or less for dual-phase xenon TPCs and there are a few factors specific to it:

- Internal reflection on liquid surface: Scintillation light is produced in the liquid, below a liquid-gas interface. This can result in photons being redirected or lost at the interface due to total internal reflection, and it is the main cause of an asymmetry in light collection between photosensor arrays above and below the liquid.
- Grid transparency: Electrical grids used to create drift and extraction fields can block the passage of light. Despite them being designed to be maximally transparent, photons which reflect multiple times off of other surfaces can end up being absorbed by grid wires.

Meanwhile,  $g_2$  translates the number of electrons drifted away from an interaction site into the number of observed S2 photons. Besides general light collection inefficiencies, there are additional losses which affect free electrons in the detector:

- Impurities: As electrons are drifted through the liquid bulk, electronegative impurities can capture electrons before they are extracted. The surviving fraction of electrons at the liquid surface  $F$  is a function of drift time  $t$ . Losses are exponential with drift time and are characterized by the length of time  $\tau_e$  that an electron can survive in the liquid before being absorbed by an electronegative impurity:

$$F(t) = e^{-t/\tau_e}$$

- Electron Extraction Efficiency (EEE): Electrons arriving at the liquid surface are faced with

a potential barrier (Figure 2.3) at the interface between phases. Overcoming this barrier requires a second “extraction” electric field which is stronger than the drift field that spans most of the TPC. The percentage of electrons extracted depends on the strength of the extraction field—this is known as the electron extraction efficiency (Figure 2.3b).

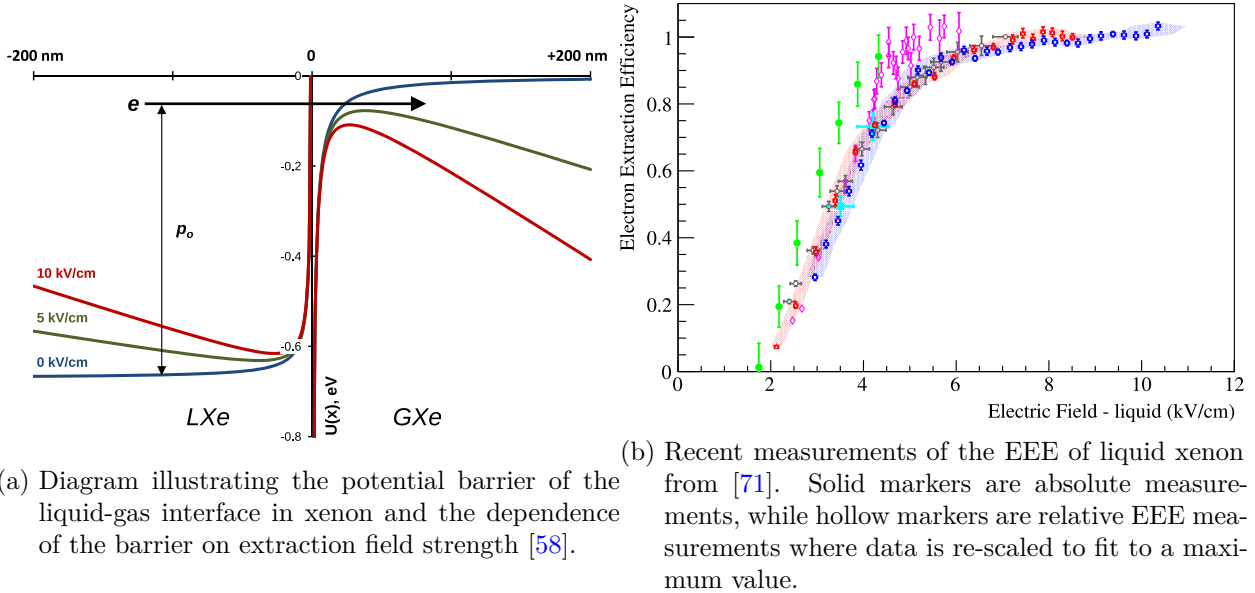


Figure 2.3.: The liquid-gas interface in a dual-phase TPC forms a potential barrier that prevents electrons from being extracted. The extraction field modifies the shape of the potential barrier and increases the energy of drifting electrons so they are able to overcome it. This results in a higher electron extraction efficiency.

### 2.3. Lowering Detection Thresholds in Xenon

Low-mass dark matter searches have historically been a secondary concern of liquid xenon TPCs due to kinematic constraints on nuclear recoil energies that are below current detection thresholds of a few keV. However as these detectors continue to produce null results, motivation to probe into lower mass ranges has increased [43]. At the same time, there is significant interest in

detecting coherent neutrino-nucleus scattering (CENNS) processes, which have only recently been observed for the first time [72]. In light of both of these issues, efforts to reach and characterize lower detection thresholds in dual-phase xenon TPCs are well-motivated.

Despite this, reaching lower thresholds is challenging because event detection in a dual-phase TPC hinges on being able to observe pairs of S1s and S2s. In a standard WIMP search, a TPC might operate with a 3-fold coincidence scheme where at least 3 photons need to be observed to constitute an S1. Given a reasonable  $g_1$  value of 10%, this puts the detection threshold around 30 photons produced, which corresponds to a nuclear recoil of a few keV (see Figure 2.4). Pushing for a lower threshold then requires either a higher value of  $g_1$  or relaxation of S1 coincidence requirements. While maximizing  $g_1$  is ideal, detecting S1 light is a fundamentally more inefficient process than detecting ionized electrons, which are readily detectable once extracted. Being the limiting factor in the relationship of event detection, removing the requirement of an S1 altogether has gained popularity as a way to push detection limits in TPCs lower, down to the fundamental level of a single electron .

In 2019, calibrations at this threshold in xenon were performed with the XeNeu TPC (Section 4.1) using a pulsed, monoenergetic neutron source at Triangle Universities Nuclear Laboratory (TUNL) [73]. The lack of S1 signal was circumvented by detecting scattered neutrons in external backing detectors, and then matching signals from those detectors with single-to-few electron pulses observed in the TPC some microseconds prior. The resulting measurement, shown in Figure 2.4, demonstrated that nuclear recoils as low as 296 eV in xenon produce charge signals on the order of 1.1-1.4 single electrons.




Figure 2.4.: Overview of liquid xenon nuclear recoil charge and light yield measurements from [74] at nominally similar electric field strengths. Charge yield results from [73] are shown in the bottom figure as purple triangles.

### 2.3.1. CENNS

The near-fundamental nuclear recoil threshold of liquid xenon measured in [73] raises the possibility of neutrino detection in dual-phase xenon TPCs via the CENNS process. CENNS is advantageous for detecting neutrinos because it receives an  $N^2$  enhancement to the cross-section where  $N$  is the number of neutrons in the target nucleus. This due to the weak charge term  $Q_W$  in the differential cross-section with respect to the nuclear recoil energy  $E_{NR}$ :

$$\frac{d\sigma}{dE_{NR}} = \frac{G_F^2 M}{4\pi} \left(1 - \frac{ME_{NR}}{2E_\nu^2}\right) Q_W^2 [F_W(q^2)]^2 \quad (2.2)$$

Where  $Q_W = Z(1 - 4\sin^2\theta_W) - N$ . With a value of approximately 0.23, the weak mixing angle  $\sin^2\theta_W$  virtually negates any contribution to the differential cross-section from protons in a nucleus, denoted by  $Z$ . Additional terms in Equation 2.2 are:  $G_F$ , the weak-scale coupling constant;  $M$ , the mass of the target nucleus;  $E_\nu$ , the incident neutrino energy;  $F_W$ , the weak nuclear form factor which is a function of momentum transfer  $q$ . In comparison to standard neutrino detection schemes which rely on elastic neutrino-electron scattering, the resultant rate of CENNS is orders of magnitude higher [58].

While this is enticing, one caveat is that for sufficiently low energy neutrinos, the maximum nuclear recoil energy for CENNS is approximately inversely proportional to the target's atomic mass:

$$E_{NR} \leq \frac{E_\nu^2}{M + E_\nu}$$

Thus while xenon is expected to have a high CENNS rate (see Figure 2.5), the maximum

nuclear recoil energy is almost three times lower than that of argon. Figure 2.5 reflects this, where lower detection thresholds are required in xenon to see comparable CENNS rates in argon.




Figure 2.5.: Predicted CENNS rates for a detector with a 10 m standoff from two sources: a 3  $\text{GW}_{\text{th}}$  nuclear reactor, and the ISIS neutron spallation source. In general, scattering rates for a neutron spallation source are lower due to the smaller neutrino flux, but are detectable at higher threshold energies because they are produced at higher energies than in a nuclear reactor [58].

If these concerns can be addressed, CENNS offers a new avenue for studies of neutrinos with dual-phase TPCs. As noted in Section 1.2.4, sterile neutrinos are a light dark matter candidate, whose existence could be inferred via measurements of neutrino oscillations. Being a flavor-blind process, all three flavors of neutrinos would contribute to the scattering rate. If oscillations also occur into a sterile channel, a precision measurement of neutrino scattering rates would show a deficit in the rate, indicating the existence of a sterile flavor [74]. Neutrinos are also the dominant channel

by which core-collapse supernovae radiate their energy [58], making their detection an important component of multimessenger astronomy. For a multi-ton xenon TPC like LZ, a few hundred neutrino interactions are expected [75], enabling LZ to contribute to the Supernova Neutrino Early Warning System which informs other observatories of impending optical and gravitational signals [76]. Another neutrino detection application is the solar neutrino flux. The energy spectrum of these neutrinos could provide information relevant to determining proportions of beryllium and boron in the sun that would have implications for BBN, stellar evolution, and other open questions in astrophysics [77]. Finally, the neutrino flux from a nuclear reactor could be monitored for treaty verification purposes [74]. In the case of  $\sim 3$  MeV antineutrinos typically produced in nuclear reactors, scatters in xenon at the endpoint energy are expected to generate charge signals in xenon near the single-ionization level [73].

### 2.3.2. S2 Only Dark Matter Searches

Small ionization signals in TPCs, similar in nature to those expected from CENNS<sup>2</sup>, are also relevant to dark matter searches where interactions in a TPC would not produce a detectable S1 signal. For example, a 5 GeV WIMP would on average produce nuclear recoils of around 250 eV in xenon. Even lighter dark matter particles, below 1 GeV in mass, could generate low energy electron recoils that would manifest as irreducible single and few-electron signals, shown in Figure 2.6.

The rates in Figure 2.6 are substantial. For an experiment on the scale of LUX-ZEPLIN (LZ), a few to single-electron rate on the order of 0.5 Hz could be consistent with a dark matter signal [78]. Predictions such as these have therefore motivated S2 only searches with dual-phase xenon TPCs.

---

<sup>2</sup>So similar, in fact, that CENNS signals are expected to form an irreducible background at higher sensitivities—see Figure 1.14.




Figure 2.6.: Event rate as a function of electrons produced by dark matter-electron scatters in xenon [78]. Colored spectra correspond to contributions from individual electron shells and the gray band represents the overall spectrum when varying the effect of higher-order processes. Left axes on the plot correspond to the maximum allowable cross-sections for benchmark models, while right axes are cross-sections constrained by the correct relic abundance.

The most recent example of this was conducted by the XENON1T experiment which observed an event rate  $< 1 \text{ event tonne}^{-1} \text{ day keV}_{\text{ee}}$  for events consisting of at least 12 extracted electrons [79]. No claims were made for events below this level due to the presence of unmodeled backgrounds which resulted in a rising event rate. A later study addressed these few-electron backgrounds with an empirical model of the few electron backgrounds (which are discussed throughout this work) that extended the results for the first search down to the single electron level. After accounting for this background, an event rate  $< 30 \text{ events electron}^{-1} \text{ kg}^{-1} \text{ day}^{-1}$  was obtained for events in the region of interest between 1-5 extracted electrons [80].

## 2.4. Low Energy Electron Backgrounds

Dual-phase liquid xenon TPCs are known to observe low-energy background signals in the form of spurious single-to-few electron pulses which are not a direct component of an interaction occurring the active volume of a TPC. These backgrounds are detrimental as they can pile-up and create accidental coincidences with S1s, mimicking a higher energy signal. Energy resolution of real signals is also diminished due to the increased uncertainty in extracted charge that forms an S2. Furthermore, they require operating at higher trigger thresholds as their relatively high rates would make data volumes otherwise difficult to cope with [81]. Needless to say, in a search for low energy interactions these backgrounds become signal-like when operating in “S2-only” mode where the requirement for an S1 signal in an event is dropped.

Unlike at higher energy levels (think: detectable S1s), there are few background interactions in a TPC below the few-electron level. Two exceptions to this are coherent elastic neutrino nucleus scattering (CENNS) [74] [58] and a beta decay background from  $^{214}\text{Pb}$  [79]. While both of these

are relevant to S2-only dark matter (DM) searches, they are outside the scope of this discussion. Instead, the low energy S2 background spectrum is dominated by pulses which are instrumental in origin (Figure 2.7).




Figure 2.7.: Low energy S2 spectrum in the 2019 XENON1T light dark matter search [79]. The analysis threshold was cut off at the vertical dotted line where events below  $\sim 4$ -5 electrons in size were not considered due to the presence of “unmodeled backgrounds”. In a 2021 reanalysis (Section D.13), an empirical model of electron trains was constructed which enabled limits to be set on light dark matter candidates based on the observed single-to-few electron background rate. Additional backgrounds in this plot include the expected CENNS rate (shown in red) from  $^8\text{B}$  solar neutrinos and a flat ER background from  $^{214}\text{Pb}$  decays (blue). The expected signal model from 4 GeV and 20 GeV spin-independent DM models are also shown in orange and purple respectively.

These instrumental electron backgrounds can be divided into three major groups: so-called “electron trains”<sup>3</sup>, grid emission, and photoionization. Electron trains are composed of delayed single-to-few electron pulses that are correlated in position and time with prior pulses. Their exact origins are unclear, but examined further in Chapter 3. They are a dominant component of

<sup>3</sup>First coined by Sorensen in 2017 [82]

instrumental electron backgrounds and are noteworthy for their impact on livetime. In the first LZ WIMP search, vetoing these backgrounds reduced livetime by roughly 30% [83]. This was by far the largest impact to livetime in the experiment, heightening the need to develop a better understanding of these backgrounds.

The second group, grid emission, originates on the gate and cathode high voltage grids. Electron background rates from grid emission are exacerbated by microscopic wire defects and dust. A detailed review of grid emission is beyond the scope of this work, but readers will note the mention of spontaneous grid electron emission in past papers (e.g. Section D.10.5, Section D.11).

Finally, VUV light from xenon scintillation and electroluminescence is known to photoionize grid wires and impurities throughout the liquid (Section D.3, Section D.5, Section D.10). Photoelectrons from this process are produced quickly and all of them arrive within one drift length of the TPC. This is a key factor which distinguishes them from electron trains.

All of these backgrounds have a shared history and were not well distinguished from each other for many years. One goal of this review is to clarify past studies that conflate them—an understandable consequence of their overlapping dependencies and timescales. See Figure 2.8 for an overview of these.

### 2.4.1. Nomenclature

One confusing aspect of this subject is the variety of terminology used, which has evolved over time. The following terms tend to be used interchangeably though they can have different meanings depending on the context:

**Electron backgrounds** Broadly speaking, this can mean any undesirable S2 signal in a TPC com-

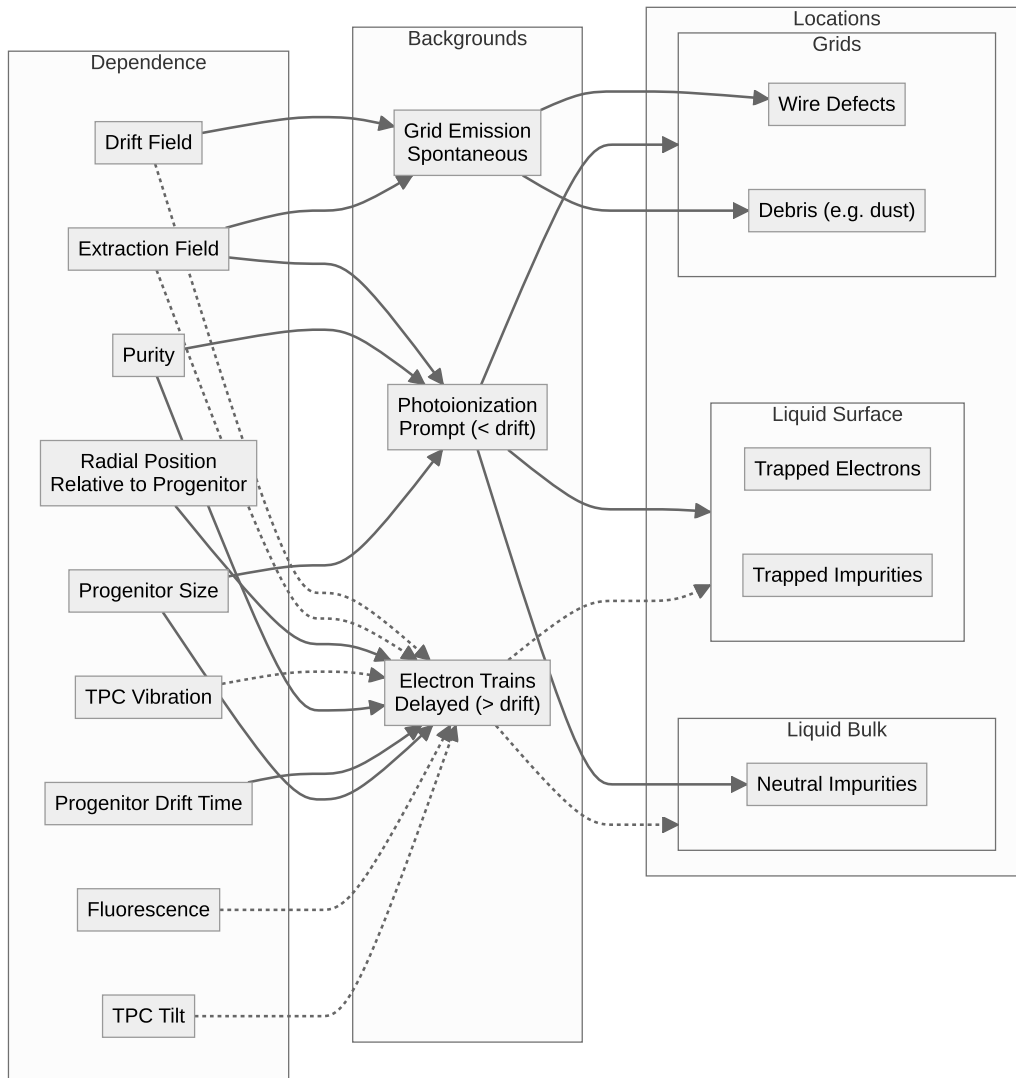


Figure 2.8.: Hardware-related electron background signals in dual-phase xenon TPCs can be grouped by the timescales in which they occur relative to an event. Each relates to unique mechanisms and sources, some of which are speculative. Dependencies are factors which have been investigated in studies of the respective backgrounds. Solid lines indicate a strong connection has been observed. Dotted lines indicate a weak dependence, no dependence, or a suspected dependence that requires further study.

posed of a dozen to a few electrons or less. In this review, only electron backgrounds which are instrumental in origin and smaller than a few electrons in size are being taken into consideration.

**Low energy electron backgrounds** More intentional than “electron backgrounds”, but less specific than other terms outlined here. Low energy typically implies “low enough energy that an S1 signal can’t be detected”.

**Single-to-few electron backgrounds** This adds a little more specificity as some electron background sources have higher rates at the single-to-few electron level. Again the hardware origin is implied, with perhaps some connotations of pile-up which is more common for signals of this size.

**S2-only backgrounds** Backgrounds which are specific to operations of a TPC in S2-only mode, which is defined by relaxing the requirement for an S1 signal in an event to achieve a lower event detection threshold. Electron backgrounds are the dominant background for S2-only event searches, but out of an abundance of caution they should not be regarded as synonymous,

**Ionization backgrounds** The term ionization is used as substitute for “S2” or “electron”. It is *probably* safe to use it interchangeably with “electron backgrounds”.

#### 2.4.2. Overview of Photoionization, Grid Emission, and Electron Trains

In this section, a basic overview of the three background types is presented, with a focus on their relation to electron trains. Grid emission is independent of other events in the TPC, so it is

considered “spontaneous”. Photoionization backgrounds are dominant within one drift length of a “progenitor”<sup>4</sup> pulse, making them “prompt” relative to an event. Finally, electron trains occur more than one drift length after a progenitor pulse; these are referred to as being “delayed”. Figure 2.9 is an illustration of these concepts in relation to a single scatter event.




Figure 2.9.: Distinct timescales in a single scatter event where low energy electron backgrounds are most prominent. Prior to an S1, low energy backgrounds are typically dominated by spontaneous pulses (dashed outline) which are not associated with a event in the TPC. These are often due to grid emission. After the S1 and S2, prompt photoionization backgrounds (solid outline and fill) are most common up to the maximum drift time in a TPC  $t_{drift}$ . As both the S1 and S2 can act as progenitors, they each have their own timeframes for prompt backgrounds. Note also that prompt backgrounds can be intense enough that single electron (SE) pulses often combine and pile-up to form smaller S2 pulses; this is shown after the progenitor S2. After prompt backgrounds from the S2 pulse have died off, electron trains in the form of delayed SEs and small S2s (solid outline and no fill) are still visible. It has been shown that delayed small S2s are not well-explained by SE pile-up [84].

Below are more details on each of the key terms in Figure 2.9:

**Spontaneous** High voltage grids in a TPC are known to produce single and multi-electron signals independent of any event in a TPC [83]. For the purpose of this review the term “spontaneous” will be adopted alongside “grid emission” to describe these, in line with existing literature on the subject.

---

<sup>4</sup>First coined by Linehan [83].

**Prompt** Another moniker for “prompt backgrounds” is simply “photoionization backgrounds”. Photoionization itself is virtually instantaneous, and the ensuing photoelectrons are seen at high rates for up to the full drift time in a TPC after a progenitor pulse.

**Delayed** For reasons which are still unclear, single electrons (SEs) and small S2s seemingly follow progenitor pulses in an “electron train” for as long as a second—much longer than the full drift time of any TPC<sup>5</sup>. Depending on a few factors like the spacing of events in time and the size of the progenitor, these backgrounds can persist into the next event and appear to be spontaneous.

**Progenitor** Progenitor pulses are any kind of pulse which is large enough to induce substantial prompt and/or delayed background pulses. Progenitor pulses may be referred to as “primary” or “main” pulses in the literature, but “progenitor” is favored here because it implies a correlation with child pulses.

**Children** “Child” pulses are those which are correlated with progenitor pulses, i.e. they make up a train of pulses following the progenitor. “Child” does not distinguish timescales, pulse type, or size, but is useful for categorizing any pulse that has an associated progenitor. Child pulses in electron trains are either SEs or small S2s up to a few electrons in size. Single photoelectrons (SPEs) can comprise a “photon train” which is distinct from but often times concurrent with an electron train<sup>6</sup>.

---

<sup>5</sup>For reference, the drift time in LZ is on the order of a millisecond, and for smaller R&D TPCs it can be tens of microseconds.

<sup>6</sup>Photon trains are mentioned here for completeness. They are arguably more intertwined with electron trains than grid emission, however they are not as well studied. While not the first to identify them, the LUX electron background paper was the first to characterize them (Section [D.10.4](#)) in a manner similar to electron trains.

### 2.4.2.1. Photoionization

Photoionization backgrounds were first documented in 2008 by ZEPLIN-II (Section D.1) where they provided an example in the form of a stray single electron between an S1 and S2 pulse (Figure D.3). The same study showed evidence of a rate dependence on the size of the progenitor pulse and proposed photoionization as mechanism responsible for producing these “unexpected” pulses.

This hypothesis was later confirmed in ZEPLIN-III (Figure D.4) and XENON100 (Figure D.6). In both studies from these experiments it was clear that photoionization of the liquid and cathode were responsible for a spatially homogeneous background that lasted until the maximum drift time of the respective TPCs. XENON100 in particular showed that the rates of these backgrounds depended exponentially on the time since the progenitor pulse. They also showed that photoionization rates had a linear dependence on the  $O_2$  equivalent concentration<sup>7</sup> of impurities in the liquid.

Following these studies, there appeared to be a consensus for a number of years that negative ions were the source of photoionization backgrounds. However, in 2020 LUX challenged this assumption (Section D.10.1). In the study it was argued that if negative ions were the main targets of photoionization, then:

1. The top of the detector should produce more photoelectrons because the drift field would concentrate negative ions there
2. Photoionization backgrounds should increase when with the event rate in the TPC, since more interactions would create more ionization throughout the liquid.

---

<sup>7</sup>The  $O_2$  concentration is not directly measured but instead calculated from a measurement of the electron lifetime.

In the LUX study, it was shown that the photoionization background was spatially homogeneous throughout the detector, when detector geometry and drift losses were accounted for. This was in agreement with results from ZEPLIN-III and XENON100. In addition, the ZEPLIN-III study explicitly assumed that photoelectron production was homogeneous throughout the detector.

The LUX study also performed tests of the event rate hypothesis. During detector calibrations there was no observable difference in photoionization background rates. This was true for  $^{83\text{m}}\text{Kr}$ , which is a spatially homogeneous calibration source in the TPC, and for neutron calibrations which were spatially localized in the detector (Figure D.11b). In the face of this evidence, the LUX study ruled out negative ions as possible photoionization centers, arguing that they were too short-lived to be candidates. The study instead concluded that some other neutral impurity was responsible for photoionization backgrounds in the liquid<sup>8</sup>.

This conclusion was affirmed in 2021, when Kopec *et al.* published a study of electron backgrounds that included tests where the TPC was flooded with 0.8 eV infrared light (Section D.12.4). This was roughly a factor of 2 larger than the electron affinity of  $\text{O}_2$  [85], but no significant increase in the electron background rate was observed.

#### 2.4.2.2. Grid Emission

Spontaneous emission of single and multi-electron pulses from high voltage grid wires is a well-documented effect [86]. While the exact mechanisms of electron emission are unclear<sup>9</sup>, two important factors are grid wire defects, and debris on grid wires [83]. Both behave similarly in that

---

<sup>8</sup>A subtle point: electronegative impurities are not necessarily negative ions; they may be electrically neutral.

<sup>9</sup>Fowler-Nordheim emission of electrons is a commonly cited mechanism, but Linehan [83] notes some inconsistencies like the need for surface fields on the order of MV/cm or monotonic behavior in emission current predicted by the Fowler-Nordheim equation. While enhancement is expected around wire defects, Linehan points out that an unrealistic aspect ratio of 20:1 would be required to achieve the fields necessary for Fowler-Nordheim emission.

they can cause drastic enhancement of fields next to grid wires. In the case of the latter, extensive cleaning and dust control protocols are necessary to reduce the chances of debris being deposited on grid wires during construction. Grid wire defects can be moderated with extremely robust quality assurance practices during wire and grid manufacturing. An additional step of acid passivation (nitric [86], citric [87]) induces the build-up of an oxide layer on wires which can counteract the presence of microscopic defects.

As a known production mechanism of isolated single and multiple electrons, grid emission was an early candidate to explain delayed electron backgrounds. Burenkov *et al.* first suggested the possibility in 2009 (Section D.2) after claiming that “spontaneous” pulses had been observed in a ZEPLIN-III prototype with no apparent connection to previous events.

ZEPLIN-III (Section D.3) observed these spontaneous SE pulses preceding S1 pulses in events, and performed a dedicated study with a  $^{137}\text{Cs}$  source that demonstrated the rate of these pulses increased with the event rate in the TPC. This was enough evidence that the so-called “spontaneous” pulses were correlated prior events, and not likely to be caused by grid emission.

In the context of electron backgrounds, the topic of grid emission was broached again in LUX (Section D.10.5). LUX confirmed results from previous works, showing evidence of transient yet localized (Figure D.14) multi-electron emission for seconds to minute-long periods of time.

### 2.4.2.3. Electron Trains

Burenkov *et al.* (Section D.2) had proposed that electrons trapped at the liquid surface may have also been a cause of the so-called “spontaneous” backgrounds observed in their ZEPLIN-III prototype. ZEPLIN-III determined these backgrounds were actually correlated with previous

events, but they did not rule out grid emission as a cause (Section D.3). Electrons trapped at the liquid surface later became the dominant hypothesis to explain delayed electrons following events. However in 2016, Akimov *et al.* cast doubt on this explanation. They observed correlated electron pulses following energetic events for over 200ms and fit a power law (Figure D.9) to the correlated electron background rate. The extremely long duration of delayed emission seemed incompatible with an electron lifetime on the scale of microseconds in TPCs, prompting them to suggest that electrons could be trapped on impurities in the liquid bulk and released later.

This power law-like dependence in the rate has proven to be a defining characteristic of electron trains. It has been observed in four different studies, summarized in Table 3.7.

Table 2.2.: Table of electron train power law exponents from previous studies.

Study	Exponent
Akimov <i>et al.</i> Figure D.9	$1.4 \pm 0.2$
LUX Figure D.12a	[1.0, 1.1]
Kopec <i>et al.</i> Figure D.16	$1.2 \pm 0.04$
XENON1T Figure D.18	1.1

As of the time of this writing neither the liquid bulk nor the liquid surface hypotheses have been conclusively ruled out. However, the most recent studies from LUX (Section D.10) and XENON1T (Section D.13) have leaned towards the liquid bulk impurity hypothesis. This is due to a strong dependence observed on five factors, outlined below.

**Purity** Rates from electron trains are known to decrease with an increasing electron lifetime. This was qualitatively observed by Sorensen and Kamdin in 2018 Section D.8, and measured more

quantitatively in LUX (Figure D.13). The dependence on electron lifetime was less clear in XENON1T (Section D.13) than in LUX, but still readily apparent.

**Drift Time** The LUX study showed that even after correcting for drift losses in progenitor S2s, delayed rates were larger at all observed timescales from progenitors at higher drift depths (Figure D.12a). This observation was confirmed again in XENON1T.

**Progenitor Size** As with photoionization rates, larger progenitors also produced electron trains with higher rates at all timescales. This was documented in XENON1T (Figure D.19).

**Radial Position Relative to Progenitor** LUX showed that XY positions of child pulses in electron trains exhibited a strongly correlation with the XY position of their progenitor (Figure D.12b). This dependence was also observed in XENON1T.

**TPC Region** XENON1T provided a positive confirmation that only progenitors borne of a TPC's drift region produced electron trains (Figure D.18). Events originating in the gas and reverse field region (RFR) below the cathode photoionized the TPC as expected. but did not produce electron trains.

These factors could indicate the presence of electrons trapped at the liquid surface, with the exception of the electron lifetime and progenitor drift time dependence. Removal of impurities should extend the electron lifetime at the surface, increasing the rates of electron trains. Additionally, since progenitor S2s at higher drift times would arrive with fewer electrons at the liquid surface, fewer electrons should be left behind.

On the other hand, an increase in rates at increased drift times is intuitive if electrons are captured by impurities in the liquid bulk and then released later via an unknown mechanism. The

low diffusion velocities of impurities ensures that electron trains are fairly localized in XY relative to the progenitor, and removal of impurities lessens the electron train intensity. The progenitor size dependence is also compatible with this explanation since larger S2s should lose more electrons as they are drifted through the TPC.

Beyond this, several other factors have been investigated which have an inconclusive or weak connection to electron trains.

**Extraction Field** Delayed electron backgrounds should be affected by a stronger extraction field if trapped electrons are present at the liquid surface. At Lawrence Berkeley National Lab (LBNL), Sorensen and Kamdin (Section D.8) observed a decrease in the amplitude of the delayed background rate with an increasing extraction field. The LBNL study claimed this was plausible if fewer electrons were left behind at the liquid surface when an S2 was initially extracted. They also observed a second delayed component which behaved in the opposite manner, where the background fraction of the progenitor increased with the field.

The LBNL study did not associate the second delayed background component with electrons trapped at this liquid surface. However, Kopec *et al.* (Section D.12) later observed a rate increase that was linear with the extraction field. They claimed that this was sensible as the chance of extracting trapped electrons from the surface should increase, and pointed to the second background component from the LBNL study as being in agreement with their findings. It is possible both interpretations are correct, though this apparent tension is difficult to resolve as both studies used different analysis methods: the LBNL study analyzed a summed waveform whereas Kopec *et al.* looked at discrete pulses. The TPC used at LBNL also consisted solely of an extraction region subject to fields 10-20x higher than the TPC used

by Kopec *et al.*, making comparisons even more opaque.

Finally it should not go without mention that in XENON1T (Figure D.20) a weak dependence on the extraction field was observed. At the very least it can be concluded here that further study of the rate behavior under a varying extraction field is merited.

**Drift Field** Drift field has been tested independently of the extraction field on two occasions, once by Kopec *et al.* (Section D.12) with ASTERiX and once by Bodnia *et al.* (Section D.11) using PIXeY. The results from both studies are difficult to compare. The PIXeY study only observed delayed backgrounds in the 30-40  $\mu\text{s}$  prior to an S1, so they could not be associated with a progenitor pulse. There also appears to be an inconsistency in how rates were calculated (Figure D.15). In the ASTERiX study the analysis was more contemporary, but the TPC was set back by a 3  $\mu\text{s}$  electron lifetime in a drift region of roughly 10  $\mu\text{s}$ .

The results from both studies appear to be in tension, with the PIXeY study observing a slight decrease in the un-normalized rate with drift field. ASTERiX observed an increase in the delayed SE fraction with drift field, but saw very little increase in the multi-electron pulse rates. This was interpreted as a sign that impurities could be releasing more single electrons under a more intense drift field or that increased charge yields could raise the delayed single electron background rates. While both of these explanations from ASTERiX are plausible, they would be working against the decreased drift times at higher drift fields which are strongly correlated with delayed background rates.

**TPC Tilt** A tilt test was performed by Akimov *et al.* in 2012 which showed that “spontaneous” SE rates decreased with a tilted TPC (Figure D.5). In 2016, Akimov *et al.* observed an

“S3” signal which they showed was likely due to unextracted electrons at the liquid surface (Section [D.6.1](#)). This S3 signal appeared within 200  $\mu\text{s}$  of muon S2s and was preceded by a period of electron emission consistent with an electron cloud migrating across the surface of the TPC. Further evidence of this was the preferential appearance of S3 signals at a specific position on the edge of the TPC, indicative of a substantially tilted liquid surface.

**Fluorescence** UV Fluorescence from PTFE is a known effect following illumination with VUV light [88]. It has been postulated in studies of electron trains (Section [D.8](#), Section [D.10](#)) that this could contribute to the formation of electron trains, but the phenomenon does not appear to have been investigated in TPCs.

**TPC Vibration** Kopec *et al.* hypothesized that lingering ripples in the liquid surface could liberate trapped electrons there. Depoian Section [D.12.6](#) conducted a primitive test of this hypothesis and saw no effect from attempts to vibrate a TPC. They concluded that further study of this was needed.

None of the factors listed above rule out the possibility of electrons being trapped at the liquid surface in dual phase xenon TPCs. Indeed the results from Akimov *et al.* (Section [D.6.1](#)) strongly demonstrate that electrons can be trapped at the liquid surface. Likewise in LUX, it was very convincingly argued that “electron-bursts” observed in the TPC were due to incomplete extraction of S2 pulses (Section [D.10.2](#)). However, based on the state of existing literature, it seems fair to conclude that surface electrons are not a dominant component of electron trains.

# Chapter 3.: Characteristics of Delayed Electron Emission in the LZ Experiment

## **i** GitLab Repository

Readers with LZ GitLab access should refer to the ALPACA module [89] used to generate plots in all sections for the chapter, except Section 3.1. The module contains a link to a dedicated GitLab Pages website where all other sections shown here are hosted as articles. These articles link to additional slides, and more importantly, to Jupyter notebooks used to generate plots of electron train rates. The website also contains a brief tutorial with instructions for generating plots.

## 3.1. The LZ Experiment

The LZ experiment is a collaborative effort between 37 institutions and over 250 scientists, engineers, and technicians. It is situated 4,850 ft underground in the Davis Cavern at the Sanford Underground Research Facility (SURF) in Lead, South Dakota [90]. At the heart of LZ is the world's largest dual-phase xenon TPC to date, with a total of 10 tons of xenon in circulation and a 7 ton active volume. The TPC is cylindrical (1.5 m x 1.5 m) and constructed mainly from interlocking PTFE pieces to improve the light collection efficiency. Light is read out by 494, 3”

PMTs, which are mounted in arrays at the top and bottom of the TPC. The TPC is housed in a double-walled cryostat made from specially-sourced radiopure titanium.

LZ is composed of multiple nested detectors which provide additional background rejection capabilities. The space between the TPC and inner vessel constitutes a liquid xenon “skin” detector which is fitted with 93 1” and 38 2” PMTs. Surrounding the outer titanium vessel are acrylic tanks which compose the outer detector (OD). The OD tanks contain 17 tons of gadolinium-loaded (0.1% by mass) liquid scintillator for neutron detection. All of these detection systems are submerged in 238 t of ultra-pure water, housed in a stainless steel tank fitted with 120 8” PMTs which detect OD and Cherenkov signals. A depiction of these systems is shown in Figure 3.1.

A more detailed overview of LZ’s subsystems can be found in [90]. For a comprehensive review, see the LZ Technical Design Report [91].

### 3.1.1. First Dark Matter Search Results

In its first science run (SR1), LZ collected 90 live days of data which were used to set world-leading limits on the spin-independent (SI) and spin-dependent (SD) WIMP-nucleon cross-section for WIMP masses between 9 and  $10^4$  GeV [92]. This required a systematic and thorough campaign to calibrate, simulate, and characterize the detector’s backgrounds as well as its response to NR and ER events. (Figure 3.6)

First, radioactive sources such as  $^{83\text{m}}\text{Kr}$  and  $^{131\text{m}}\text{Xe}$  were dispersed throughout the detector to gain an understanding of the spatial variations of the light collection efficiency and the electric field. Information from these events was used to produce corrections to normalize S1 and S2 signals to the geometric center and the center of the liquid surface of the TPC, respectively. Tritium was




Figure 3.1.: The TPC (1) sits at the center of the LZ experiment. Segmented acrylic tanks (2), shown in green, form the OD which encircles the titanium outer vessel. 8" PMTs (3) are mounted on a frame to detect light from the OD and water shielding. The frame is lined with Tyvek to improve light collection. Water shielding (4) is contained by the stainless steel water tank. High voltage for the TPC grids is provided to the TPC via a feedthrough (5) which connects to the cathode. Neutron calibrations are conducted using two conduits (only one shown, 6) which provide a clear path to the TPC. Figure taken from [90].

also dispersed throughout the TPC by injecting tritiated methane ( $\text{CH}_3\text{T}$ ) to observe the response to beta decay (ER) events. Information from tritium events was used to tune simulations of the detector response. These simulations were then used to predict the NR response, which was checked against localized calibrations of NR events. NR events were produced using neutron sources such as AmLi and a deuterium-deuterium (DD) neutron generator which could be deployed to specific ports surrounding the detector.

Figure 3.2 shows DD and  $\text{CH}_3\text{T}$  events plotted by their corrected S1 and S2 signals, S1c and S2c. The plot shows how NR and ER events can be distinguished from each other. This is critical since the signal region is in the NR band, while the most common background was expected to be from low-energy ERs.




Figure 3.2.: Calibration events plotted in  $\log_{10}(\text{S2c}) - \text{S1c}$  space.  $\text{CH}_3\text{T}$  events are represented by dark blue points, and the median of their simulated distribution is represented by the solid blue line. The dotted blue lines represent 10% and 90% quantiles. DD events are represented by orange points and the same properties of their simulated distributions are shown with red lines. Contours of constant NR and electron-equivalent recoil energies are shown in gray. From [92].

Next, numerous cuts were developed to exclude events which had:

- Occurred during periods of time with elevated noise from electronics or TPC activity
- Multiple S2 pulses, indicating a particle had scattered more than once in the TPC
- Deposited too much or too little energy in the TPC
- Contained accidental pairs of otherwise unrelated S1 and S2 pulses
- Coincident signals in any of the vetoing detectors
- Occurred too close to the outer boundaries of the TPC

These cuts were tuned to balance their effectiveness in removing backgrounds (false positives) against the possibility of rejecting a signal (false negatives). In particular, this work informed the development of the electron/photon train veto which determined time holdoffs for delayed electron and photon emission in SR1 [83]. At the conclusion of SR1, 24 of the 89 live days of data were excluded by the electron train veto to mitigate the risk of accidental coincidences from either pile-up photons or electrons following large S2s. After applying all cuts, a map of the remaining events in the active volume in the TPC was obtained (Figure 3.3).

The remaining events were plotted in a manner identical to Figure 3.2 to evaluate whether any signals were present in the region of interest (ROI) for a WIMP interaction. Various contours are shown in the plot for  $1\sigma$  and  $2\sigma$  bounds of the background model in gray,  $^{37}\text{Ar}$  events in orange, a 30 GeV WIMP in purple, and  $^8\text{B}$  solar neutrinos in green. The red lines correspond to the median as well as 90 and 10% quantiles defining the distribution of NR events.

Following this, an ER energy spectrum of the same events was produced (Figure 3.5) to compare to expected counts from nine different background populations. Most background spectra are too small to be visible, with the exception of those from  $^{37}\text{Ar}$ ,  $^{127}\text{Xe}$ , and a combined flat




Figure 3.3.: Map of the TPC in reconstructed depth and radial coordinates. The dotted black line represents the maximum extent of the active volume, which is slightly distorted because of electric field nonuniformities near the walls. The solid black line represents the fiducial cut, which is used to reject events which are too close to the walls. Events marked with gray points fail the fiducial cut, while events with black points pass it. Red crosses and blue circles represent events which were rejected by skin or OD veto detectors, respectively. From [92].




Figure 3.4.: Data from Figure 3.3 plotted in  $\log_{10}(S2c) - S1c$  space. Events in the WIMP ROI are fairly sparse. From [92].

background from low-energy  $^{85}\text{Kr}$ ,  $^{214}\text{Pb}$ , and  $^{212}\text{Pb}$  beta decays. Both  $^{127}\text{Xe}$  and  $^{37}\text{Ar}$  are cosmogenically produced. The two lead isotopes are products of decays from  $^{222}\text{Rn}$  and  $^{220}\text{Rn}$  respectively, which are naturally occurring isotopes.

The background spectrum was found to be consistent with zero WIMP events for WIMP masses above 9 GeV. Limits on the SI cross section as a function of WIMP mass are shown in Figure 3.6. The dotted black line shows the median sensitivity projection while the green and yellow bands respectively represent  $1\sigma$  and  $\sigma$  sensitivity thresholds around the median. The gray dot-dash line represents the 90% confidence limit before a power constraint is applied to account for an underfluctuation of backgrounds at masses between 13-36 GeV. The solid black line is produced after applying this constraint.

The sharp loss in sensitivity to interactions from low-mass WIMPs and the high impact to




Figure 3.5.: SR1 ER energy spectrum passing all cuts, superimposed with background estimates. From [92].




Figure 3.6.: Limits on the spin-independent cross-section for WIMPs with masses between  $10^1$  and  $10^4$  GeV. Results from DEAP-3600, LUX, XENON1T, and PandaX-4T are also shown. From [92].

livetime from the electron train veto highlights the ubiquity of ionization backgrounds and the importance of this work. After SR1, the high livetime losses from the first iteration of veto tuning led to efforts to re-tune the veto to significantly increase livetime. Among the options considered for re-tuning was an empirical model of electron trains based on numerous characteristics such as the progenitor area, drift time, and electron lifetime. This work delivers a comprehensive characterization of electron trains based on these parameters to inform such efforts. It also uncovered evidence of a drift field dependence that may have been previously overlooked in past experiments. Finally, significant efforts were dedicated to thoroughly documenting and developing `etrain`, an internal LZ python package for handling livetime calculations and generating plots of electron trains in noble liquid TPCs.

### 3.2. Measuring Pulse Rates with Deadtime

Naively speaking, a pulse rate is simple to measure. All that's needed is a stopwatch and a counter. The pulse rate  $r_0$  is the number of pulses  $n_0$  observed in a time interval  $t_0$ , divided by the length of that time interval.

$$r_0 = \frac{n_0}{t_0} \tag{3.1}$$

Where  $t_0 = t_{0f} - t_{0i}$ , which is the difference between two timestamps.

**i** Example

**Example 3.1.** A simple counting measurement where counting starts 2 ms after some arbitrarily decided start time, and ends 12 ms after the same arbitrarily decided start time. A UNIX timestamp, or some phenomenon like “after an actor steps onto the red carpet” are examples of a start time. Counts could be anything discrete like a number of pulses, or camera flashes.

$$\frac{10 \text{ counts}}{12 \text{ ms} - 2 \text{ ms}} = 1 \text{ kHz}$$

If counting is stopped and then started again, the pulse rate in a subsequent interval  $t_1$  is still given by Equation 3.1. Results from both measurements can be combined if it is physically sensible to do so:

$$r_{01} = \frac{n_0 + n_1}{t_{0f} - t_{0i} + t_{1f} - t_{1i}} \quad (3.2)$$

If the intervals are adjacent in time i.e.  $t_{0f} = t_{1i}$ , then  $r_{01}$  simplifies further:

$$r_{01} = \frac{n_0 + n_1}{t_{1f} - t_{0i}}$$

It is intuitive to see here that for adjacent time intervals, the result of two measurements is equal to the result of one measurement over a longer time interval. It should also be clear that rates over two different time intervals cannot be summed:

$$r_{01} \neq r_0 + r_1 \quad (3.3)$$

At best, if the intervals are the same duration ( $t_0 = t_1 = t$ ), then the rate across both intervals  $r_{01}$  is the arithmetic mean of  $r_0$  and  $r_1$ :

$$r_{01} = \frac{n_0 + n_1}{2t} \quad (3.4)$$

### 3.2.1. Live, Dead, and Wall Time

Returning again to Equation 3.1,  $t_0$  is said to be **livetime**, i.e. time when a count (e.g. a pulse) could have been observed. This is in contrast with **deadtime**, which is time when a pulse could *not* have been observed. In a simple one-off measurement like in Equation 3.1,  $t_0$  is continuous and there is no deadtime accrued. In this case, livetime is equal to the **wall** time, which is the total length of time elapsed on a hypothetical wall clock that continues to run, regardless of whether a counting measurement is taking place.

Stated mathematically, for a single counting measurement like in Equation 3.1:

$$t_{wall} = t_{dead} + t_{live} \quad (3.5)$$

Deadtime immediately becomes relevant if two measurements take place over non-adjacent time intervals. For example, if counting cannot take place while data from one measurement is being transferred, then the time between measurements can be taken as deadtime. Consider the case in Equation 3.2, except in this instance the start of the first counting interval is after the end

of the zeroth counting interval:  $t_{1i} > t_{0f}$  . The livetime is still:

$$t_{live} = t_{0f} - t_{0i} + t_{1f} - t_{1i}$$

while the deadtime is the time between measurements when counting cannot take place:

$$t_{dead} = t_{1i} - t_{0f}$$

and the wall time can be immediately recovered from Equation 3.5:

$$t_{wall} = t_{1f} - t_{0i}$$

### **i** Example

**Example 3.2.** An observer is assigned to determine the rate of camera flashes after a star steps onto the red carpet during an awards ceremony. After a star steps onto the red carpet, the observer starts their timer. While counting camera flashes, the observer blinks. While they blink, they cannot count any camera flashes; time spent blinking is deadtime.

So, in order to avoid misrepresenting the rate of camera flashes, the observer is trained to stop their timer when they blink, and start the timer again when their eyes are open.

Thus, the observer tracks livetime with their timer. The observer may also have a wrist watch which runs continuously, allowing them to track the wall time.

### 3.2.2. Factors Affecting Deadtime

It is important to note here that the given definitions of deadtime and livetime rely on knowing whether a pulse could have been observed. This is quite vague; in reality this condition is subject to many factors in an experiment, which in and of themselves can involve considerable ambiguity. Some of these factors are outlined below.

**Pile-Up** In a pulse-counting experiment, two or more pulses arriving close together in time can obscure each other. This phenomenon is known as pile-up and it results from inadequate timing resolution in either an experiment's physical characteristics or electronics chain. An experiment which measures the energy of pulses can mitigate this somewhat, as pulses which are piled up will appear to be one larger pulse with a higher energy. If the experiment can resolve the positions of pulses then pile-up can be mitigated further because pulses which pile up in time are not necessarily co-located. This raises the question of a position-dependent deadtime, where certain regions of a detector are vetoed. While this is possible in principle, it requires a significantly more complex deadtime calculations.

**Saturation** Saturation occurs when an experiment's ability to produce pulses, detect, or record pulses is overwhelmed from some kind of large signal. Excessive pile-up can cause saturation, but so can large one-off pulses, or prolonged periods of time where large numbers of pulses are detected. The former will typically saturate an experiment's detection hardware, while the latter may overload data acquisition systems. An experiment cannot accurately count or measure the energy of pulses while it is saturated.

**Trigger** As a way to reduce data volumes, data collection can be controlled by a trigger which

records events based on a set of predetermined conditions. Thus, the trigger conditions have direct ramifications for whether a pulse is observed and whether or not it should be counted.

The simplest condition for a trigger is that it fires randomly. This reduces the data volume and because the process is random, it avoids a sampling bias. With this trigger configuration, livetime is simply any time when data is being recorded.

If an experiment is searching for pulses with specific characteristics, a trigger can be configured to fire only when pulses that meet certain criteria are encountered. One example of this example involves the use of a discriminator which prevents the trigger from firing unless a pulse exceeds a preset amplitude threshold.

The consequence of even a simple discriminator is the introduction of a sampling bias and corresponding trigger efficiency, i.e. the probability that a pulse will cause the trigger to fire. If the trigger efficiency for a pulse is 100%, then livetime is any time within an event *and* any time when a trigger *could* fire but did not. This is because the absence of a trigger indicates that pulses of the desired kind were not present<sup>1</sup>.

If the trigger efficiency for a pulse is zero, then time within an event can be live if the pulse of interest is not correlated with the triggering pulse. Alternatively, if it is correlated with the triggering pulse, then such a configuration could still provide useful information for measuring the correlation between triggering pulses and other pulses of interest.

**Other Cuts** Cuts may be used to remove data which was subject to pile-up and/or saturation. But even when data is collected without issue, additional cuts are almost always applied to remove background signals or to simply select sub-populations of data. These actions can decrease

---

<sup>1</sup>As the saying goes: “No news is still news!”

livetime, counts, or both depending on the effect of the cut and how it is applied and how the rate is to be interpreted.

Elements from each one of these categories can overlap with each other, making a full determination of deadtime very complex. For example, even if the trigger efficiency for a pulse is 100%, deadtime can still increase if pile-up or saturation occur at any point in the detection hardware or data acquisition chain. Whether these factors *need* to be taken into account is a valid concern; often they can be neglected if their effects are small.

### 3.2.3. Pulse Rates vs. Time

When modeling physical processes, a pulse rate like the one given in Equation 3.2 is useful at-a-glance but it destroys information about how a pulse rate can vary with time. To preserve this information, two histograms are needed: one to record counts of pulses vs. time and one to record livetime vs. time.

Consider the following event from LZ shown in Figure 3.7.

This event takes place a few milliseconds after an S2 (a progenitor). Three single electron (SE) pulses are visible in the event. The counts histogram is fairly intuitive (Figure 3.8). A count is placed in a bin if a pulse is observed within the wall time interval that a bin corresponds to.

#### Caution

The bins on the time axis for these histograms correspond to wall time, which as defined in Equation 3.5 is *not* always equal to livetime. The wall time axis here also captures only part of the event window, and some time afterwards into a subsequent event window.




Figure 3.7.: An LZ event after an S2 pulse; the S2 pulse is not shown. Single electron (SE) pulses are marked with red dots. The taller pulses at 4200 and 6600  $\mu\text{s}$  are multiphotoelectron (MPE) pulses, which are ignored here. The time axis in the event window is relative to the *trigger*, which was fired at random;  $t=0$  on this axis is when the trigger was fired. In this trigger configuration, 1 ms of data was captured before the trigger and 11 ms was captured after. An additional axis is shown at the bottom denoting time since the end of the S2 pulse.

For the livetime histogram (Figure 3.9), a bin is filled with the value of the livetime elapsed *within* the wall time interval that a bin corresponds to.

In Figure 3.9 the livetime for wall time bins between 15 and 23 ms after the S2 should be fairly intuitive. These bins map almost directly to the event in Figure 3.7. For each ms of wall time elapsed between 15 and 23 ms, an SE pulse *could* have been observed at any point in time. After 23 ms, there is a dip in livetime. This is because the event in Figure 3.7 *ends*, and a new event does not start recording until approximately 4  $\mu\text{s}$  later. Thus, between 23 and 24 ms after the S2 pulse, there was only 0.96 ms of livetime, not 1 ms. This is analogous to the red carpet observer blinking, during which time they physically cannot observe any camera flashes.

The pulse rate vs. time is obtained when the two histograms are divided. After the two histograms are divided, care should be taken to avoid any operation which would inadvertently




Figure 3.8.: Counts of single electron pulses between 15 and 25 ms after an S2 pulse. Bins are 1 ms wide. the last two SE pulses from the event shown in Figure 3.7 are captured in this histogram.




Figure 3.9.: Livetime between 15 and 25 ms after an S2 pulse. Bins are 1 ms wide.

sum rates together (see Equation 3.3). This can include:

- Rebinning a rate histogram
- Creating an array of rate histograms in parallel that are then summed together
- Taking a projection of a higher-dimensional rate histogram over more than one bin along an axis orthogonal to the wall time axis.


The arithmetic mean (Equation 3.4) should also not be used to correct for rate summation operations *unless* livetime elapsed in every wall time bin is identical.

### 3.2.4. Repeated Trials

Deadtime in counting experiments is inevitable, and it can be difficult to reckon with. However, it is necessary to make an effort to account for deadtime in order to obtain results which represent the actual scope of the experiment. Clearly, if there is a large amount of deadtime between pulse rate measurements ( $t_{1i} \gg t_{0f}$ ) then there is ambiguity in whether the result from Equation 3.2 is representative of the pulse rate between  $t_{0i}$  and  $t_{1f}$ . This can be seen directly when creating histograms of pulse counts and livetime, where deadtime manifests as an un-physical lack of counts and livetime in some wall time bins.

One way around this problem is to repeat trials in such a way that deadtime does not occur in the same wall time bins for each trial. In the example covered in Section 3.2.3, there was  $\sim 4 \mu\text{s}$  of deadtime between 23 and 24 ms after an S2 pulse. Looking after a different S2 pulse (Figure 3.10), deadtime is present in a different time bin, in this case between 18 and 19 ms after the progenitor S2.

By summing together counts and livetime histograms across progenitors, an observer can




(a) Counts of single electron pulses following a second, different progenitor. (b) Livetime histogram for the associated progenitor.

Figure 3.10.: Pulse count and livetime histograms for a progenitor S2 different from the one shown in Figure 3.7. The same period of time since the progenitor S2 is shown as in Figure 3.8 and Figure 3.9.

paint a statistical portrait of the pulse rates which correctly accounts for livetime. Failing to account for livetime (i.e. only recording counts of pulses vs. time) will result in a distorted spectrum where pulse counts will exhibit an artificial exponential drop over time (Figure 3.11). This is because the arrival times of progenitors are exponentially distributed, meaning many pulses will be counted shortly after progenitors simply because another progenitor had not yet arrived. At longer timescales following a progenitor there is a greater chance that another progenitor will arrive, so the the amount of livetime spent counting pulses at long timescales will be much lower than it would be at shorter timescales.

### 3.3. LZ Data Extraction and Processing

The discussion until this point has been in somewhat general terms. Here, these terms will be applied to the case of electron trains in LZ. In this circumstance a trial can be thought of as one



(a) Counts of single electron pulses summed across many progenitors. (b) Livetime histogram summed across many progenitors.

Figure 3.11.: Pulse count and livetime histograms for pulses following thousands of different progenitors. Counts of single electron pulses appear to decrease exponentially more than 100 ms after a progenitor (Figure 3.11a). This decrease is mostly an artificial result of livetime *also* decreasing exponentially (Figure 3.11a), and should serve as a motivating example for livetime corrections.

electron train, and wall time for each trial is measured relative to the end of the progenitor pulse.

A key distinction to make here is that data extraction takes place within the LZ software environment using ALPACA, while post-processing and analysis are performed in a `conda` environment that is much easier to set up locally. This is especially advantageous for reproducibility and resiliency against infrastructure downtime.

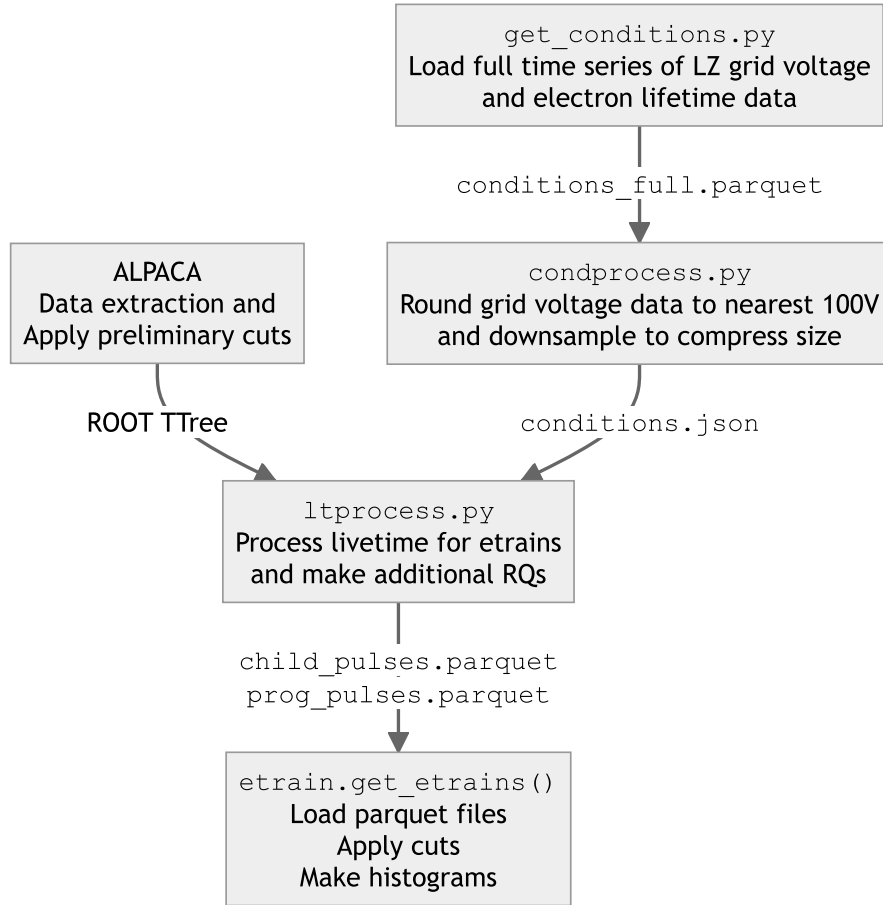


Figure 3.12.: Overview of data extraction and processing structure

### 3.3.1. Anatomy of an Electron Train

In randomly triggered data in LZ, events contain pulses up to 1.1 ms before and up to 10 ms after the trigger. Events are prevented from overlapping by a trigger holdoff that disables the trigger until 1.1 ms after an event ends. An electron train is composed of consecutive events (Figure 3.13):

- The first event contains the progenitor pulse and child pulses which follow
- Subsequent events contain more child pulses

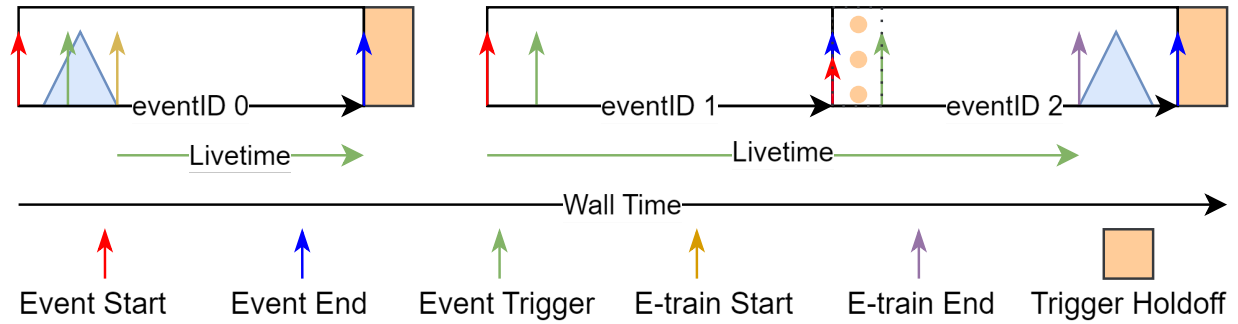


Figure 3.13.: Mockup of an electron train captured with a random trigger. The sequence of events 1 and 2 illustrate how a trigger holdoff prevents events from overlapping as triggers are blocked until the holdoff expires. Livetime is only within the bounds of event windows and is further limited to the time between vetoing pulses (pink triangles). This is because the trigger efficiency for any pulse outside of an event is manifestly zero. If a pulse-based trigger is used, then livetime includes time inside of events between vetoing pulses *and* outside of events when a trigger *could* have occurred. This is shown with a dotted line extending the rightmost livetime arrow.

- The last event contains child pulses and a vetoing pulse, which is any kind of pulse that disrupts the electron train. Typically, a vetoing pulse would be another S2 which may be a progenitor.

To calculate livetime relative to a progenitor pulse, the following timestamps are necessary:

- The time at the end of a progenitor pulse ( $t_{wall} \approx 0$ ) and end time of the event
- The start and end times of any intervening events
- The time at the start of an event and the time at the start of a vetoing pulse

Because event times are not necessarily in alignment with wall time bins, the livetime value for the first and last bins in the livetime histogram is expected to be less than the wall time (see again Figure 3.9).

### 3.3.2. ALPACA Data Extraction Algorithm

Data on electron trains is extracted from LZAP ROOT files using ALPACA. The Big Delayed Electron Backgrounds (BigDEB<sup>2</sup>) algorithm is shown in Figure 3.14 and Figure 3.15. The algorithm works by looping through pulses within events. Once a progenitor pulse is found, all S2-like pulses following the progenitor are counted as “child” pulses of that progenitor. If a pulse exceeds a set “veto” threshold, subsequent pulses are skipped until a new progenitor found. A variety of criteria determine whether a pulse is a progenitor. Typical progenitors must be large enough to induce sustained emission of child pulses. A pulse cannot be a progenitor if a vetoing pulse has occurred recently, to avoid associating child pulses with the wrong progenitor.

Details on each step in the event loop are given below:

**New Run** Before events are logged, check if the event is from a new run. This step exists to catch an edge case where pulse counting is enabled, but a new run has started and interrupted the sequence of events.

**Event Logging** An event is part of an e-train if counting is enabled, i.e. `counting_enable == true`. Event timestamps are saved here for livetime calculation later.

Details on each step in the pulse loop are given below:

**Veto Check** When looping over pulses, a vetoing pulse is any pulse which can disrupt an e-train either by photoionizing a substantial number of electrons or even starting a new one. Pulse area is used as a proxy for this as larger pulses will produce more secondary electrons.

---

<sup>2</sup>Pronounced “Big Debbie”

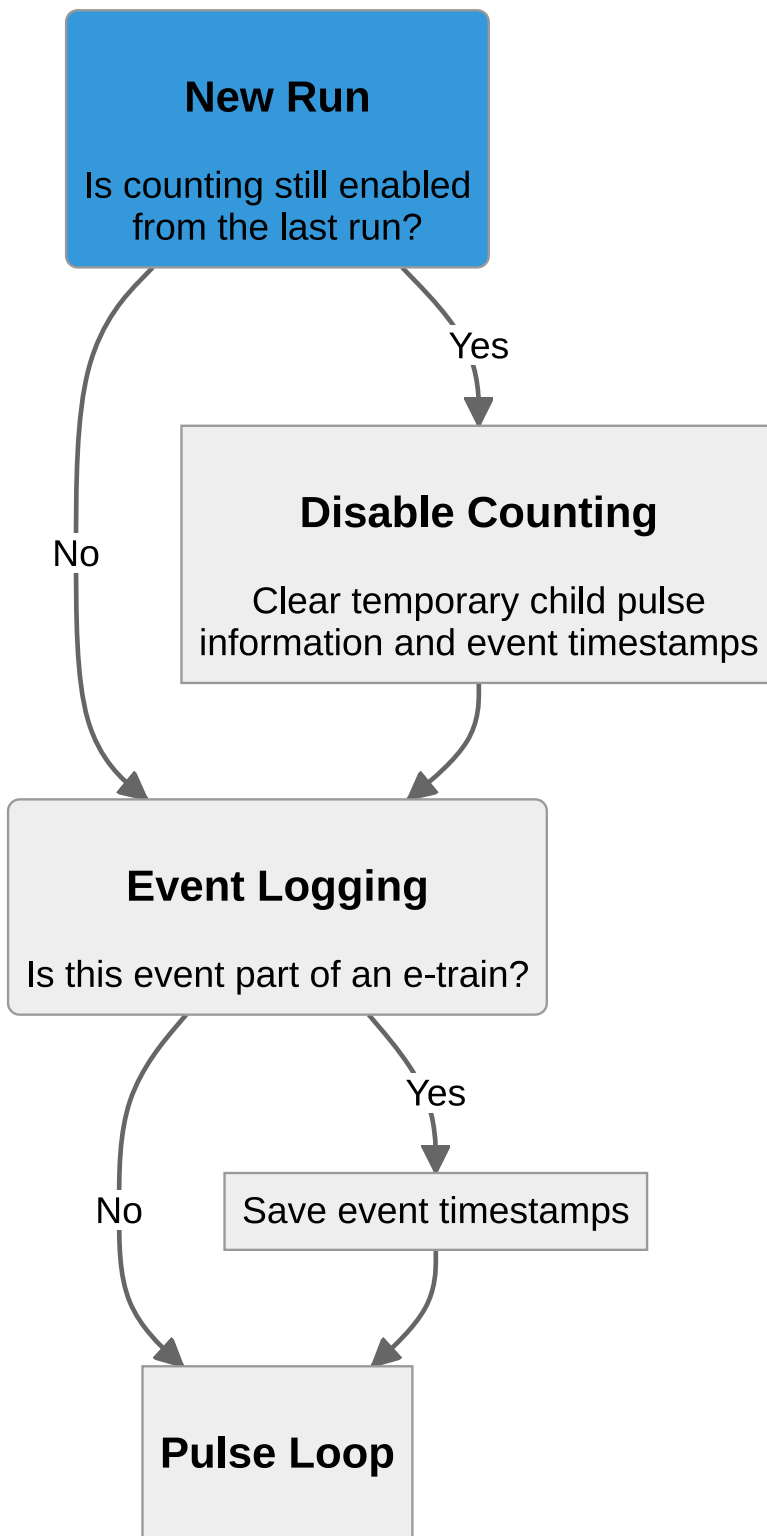


Figure 3.14.: Flow chart illustrating the event loop of the BigDEB algorithm.

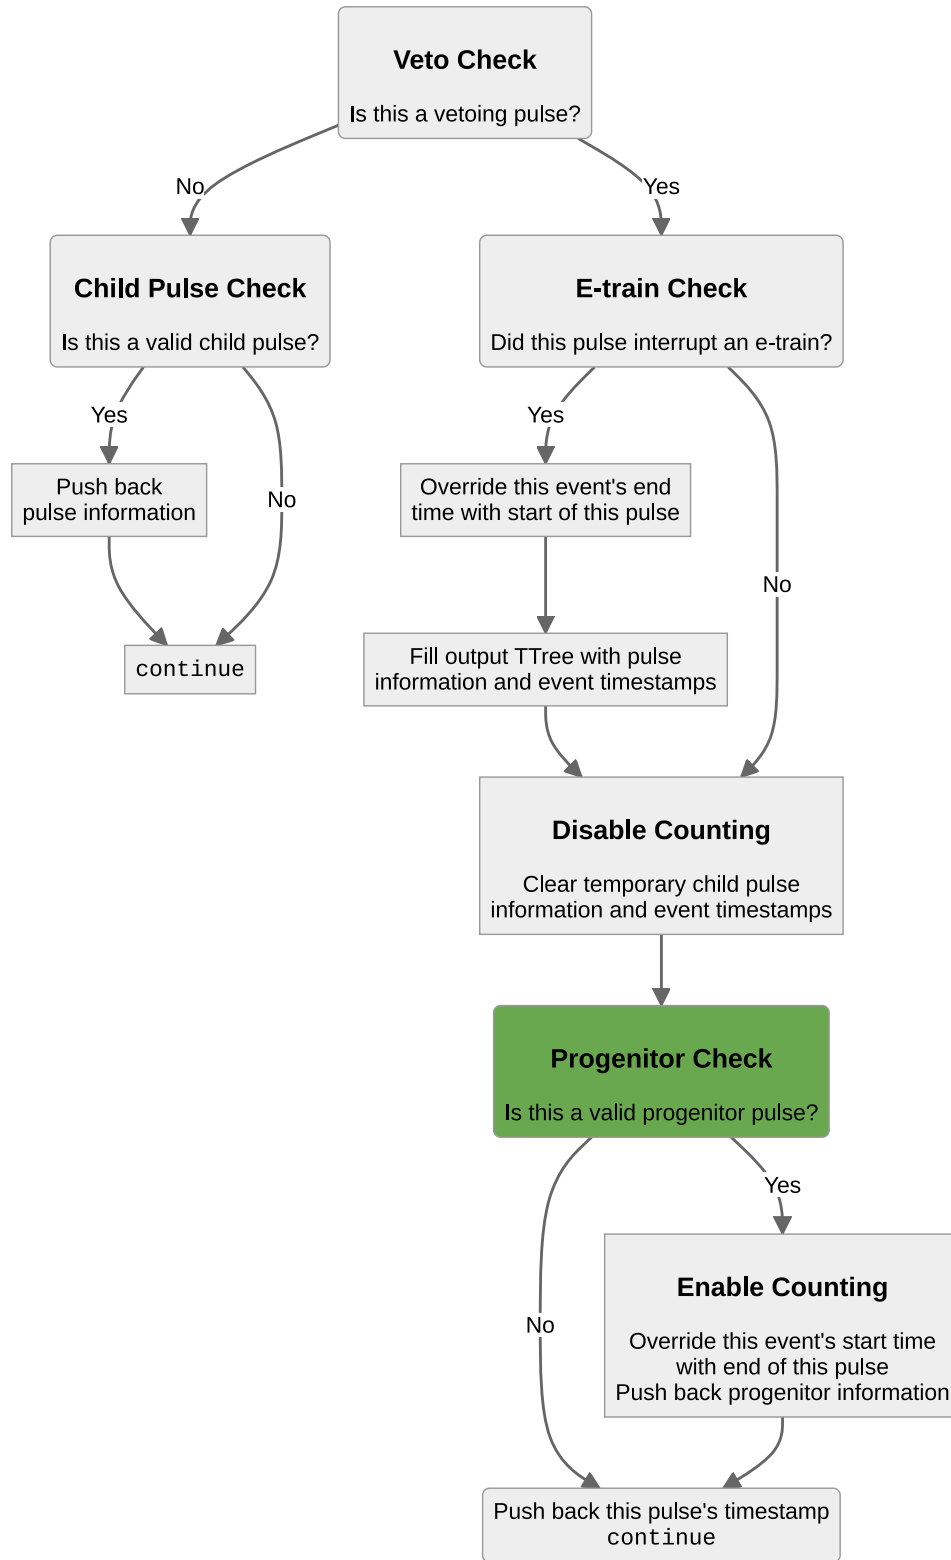


Figure 3.15.: Flow chart illustrating the pulse loop of the BigDEB algorithm.

**E-train Check** Data is collected if `counting_enable` and `IsIsolatedFromProgPulse(current_pulse)`

are both `true`. The first condition is a flag toggled by `EnableCounting` and `DisableCounting`.

The second condition checks whether enough time has elapsed ( $\mathcal{O}(ms)$ ) between the progenitor and the vetoing pulse. This ensures e-trains being collected have a long enough duration to be substantially representative and also excludes more complicated behavior like multiple scatters. In LZ, an additional cut is used here for excluding any e-train with grid “hotspots”—periods of sustained electron emission from the high voltage grids.

If data is collected, the event end timestamp is overridden by the start of the vetoing pulse to correctly track livetime.

**Child Pulse Check** Child pulses are saved if `counting_enable == true`.

**Disable Counting** In the event loop, electron train data is stored in temporary vectors before being offloaded. These vectors are always cleared when a vetoing pulse is encountered. The `counting_enable` flag is also set to `false` here.

**Progenitor Check** All progenitors are a subset of vetoing pulses, so when a vetoing pulse is encountered a check is also run to determine if it could be a progenitor pulse. An example of this check (in `c++`) is shown below:

```
return (nSingleScatters == 1
        && pulse.classification == "S2"
        && m_cutssv_SEDecayTime->EventIsNotSaturated(nSamplesInChannel) ①
        && m_cutssv_SEDecayTime->PassesRadialCut(pulse.x_cm, pulse.y_cm,
        ↪ R_max_cm) ②
```

```

    && (pulse.area_phd > 10000)

    && (pulse.AFT10_TS - vetoing_pulse_AFT99_TS.back() > he_holdoff_ns) ③

    && (eventendTS > pulse.AFT99_TS + 1e3)); ④

```

- ① Saturation can take many forms. In LZ, this particular check ensures that data buffers for holding PMT pulse digitizations are not full. Data buffer saturation is highly correlated with PMT output saturation, so an additional check for PMT output saturation is not used.
- ② Progenitor pulses at large radii are undesirable because they may be susceptible to unusual charge losses from electric fields or other effects. An additional background in LZ at high radii are so-called “glue-ring” events which are believed to originate from radioactive decays atop the gate grid “glue-ring”. When these events occur, detector components from the extraction region like the Kapton light baffles obscure incoming light and result in the S2 pulse being misreconstructed closer to the TPC [83].
- ③ It is important to include only progenitors which are sufficiently isolated from the most recent vetoing pulse. A static holdoff time (e.g. 200 ms) is sufficient in many cases but more advanced methods can be used which account for the relative areas of recent vetoing and progenitor pulses [80].
- ④ On rare occasions a progenitor pulse is cut off at end of event window. This can cause issues with area fraction time (AFT) parameters which may erroneously end after the end of an event window. This issue is mitigated by requiring that the progenitor pulse ends at least  $1\mu s$  before the end of an event window.

**Enable Counting** Valid progenitors set the `enable_counting` flag to `true` and their information is stored. The start and end timestamps of the event are recorded here; the end of the progenitor pulse is used for the start timestamp to correctly track livetime. Finally, a copy of the pulse timestamp is stored separately as it is still counted as a vetoing pulse which needs to be logged.

### 3.3.3. Livetime and RQ Processing

Event timestamps from electron trains are saved during the data extraction step for subsequent post-processing. Each e-train is saved with vectors of start and stop timestamps that span the duration of the e-train, and livetime is assumed to be any time between those time stamps when data is collected with a random trigger (see Section 3.2.2). In this work, this assumption is still used for data collected with a more complex trigger configurations, though in the future this could be improved upon by incorporating trigger efficiency measurements for corrections.

Processing files is mostly straightforward. As shown in Figure 3.16, datasets are split into chunks and then each chunk is processed in parallel. Within each chunk, livetime is processed for each electron train. Additional RQs are then derived and associated with each pulse. Some examples include relative pulse positions, and detector conditions like the grid voltages or the electron lifetime. Additional details on each step are given below.

**Chunk Files** ALPACA may output either a single large ROOT file or a collection of ROOT files. If a list of files is given, this step assumes the files are already evenly split and chunks correspond directly to files. In the case of a single ROOT file, the number of entries is read and a chunk consists of a start and stop index to be read. Chunks are generated to be nearly equal in size

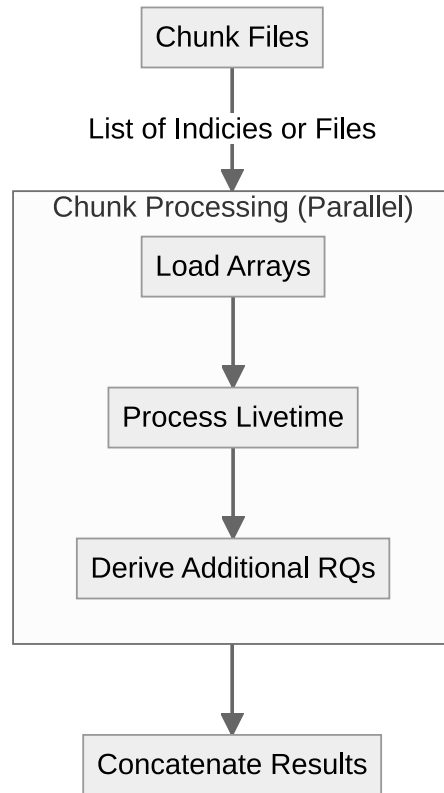


Figure 3.16.: General overview of steps in livetime and RQ post processing.

using `np.array_split`.

**Chunk Processing** Chunk processing is trivial to parallelize because chunks are independent from one another. Parallel processing is accomplished with `Dask`. `Dask` performs best when chunks fed in are some kind of pointer to data like a single file path or a start and stop index to read from a single file.

**Load Arrays** This is a simple step where arrays of event timestamps are extracted using `uproot.open`.

**Process Livetime** For each e-train, vectors of event timestamps are fed into a `numba-compiled` function that maps start and stop timestamps to livetime fill values and bin centers on a wall time axis. The wall time is measured relative to a progenitor pulse. When mapping a set of start and stop times, there are three cases to consider:

1. The start and stop times are in the same bin. The livetime in that bin is equal to the difference between the start and stop time.
2. The start and stop times are in adjacent bins: The livetime in the start bin is equal to the time difference between the start time and the right edge of the start bin. The same applies for the stop bin except the livetime is equal to the difference between the left edge of the stop bin and the stop time.
3. The start and stop times are in non-adjacent bins. The start and stop bins have the same livetime as they would in the second case, with intervening wall times bins containing livetime equal to their width.

It is worth emphasizing here that it is *necessary* to calculate livetime in the start and stop

bins instead of assuming that livetime is simply equal to the width of the wall time bin. This is especially true for bins on a logarithmically spaced wall time axis where the width of a bin in log-space can be *much* larger than the length of a single event. For example if five bins are specified on an axis ranging from -5 to 0 log-seconds, then each bin has a width of 1 log-seconds which can be as little as 90  $\mu$ s or as much as 90 ms.

**Derive Additional RQs** Additional RQs are processed for each pulse in an electron train as well.

`np.interp` is used to map progenitor pulse timestamps with conditions data like grid voltages and the electron lifetime in LZ<sup>3</sup>. Once grid voltage data is mapped to each pulse, a simple `groupby` operation on a `pandas` dataframe generates a curve of the single electron area dependence on the extraction voltage:

```
se_phd_vs_EdV = (  
    se_pulses_df                                     ①  
    .groupby("extraction_voltage")[["extraction_voltage", "pulse_area_phd"]]  
    ↪                                               ②  
    .mean()                                         ③  
    .to_dict(orient="list")                         ④  
)
```

① This is a dataframe where each row is an SE child pulse in a dataset which spans data across multiple extraction voltages.

---

<sup>3</sup>A left join between a conditions data table and pulse data table using progenitor pulse timestamps would also work here.

- ② The `groupby` method aggregates data based on the name of the column provided in the argument.
- ③ The aggregation operation `mean` in this case returns an object where SE pulse areas are averaged over each extraction voltage.
- ④ The resulting `GroupBy` object is exported to a python dictionary, comprising the SE pulse area vs. extraction voltage curve.


This curve is then used to convert S2 pulse areas in units of photons detected (phd) to a pulse area of an equivalent number of extracted single electrons. More concretely, if an S2 pulse consists of 200 photons detected, and a single electron generates 50 detectable photons, then it is assumed the S2 pulse consisted of  $\sim 4$  extracted single electrons. This additional conversion of pulse areas to a single electron-equivalent area is useful for making comparisons between datasets taken at different extraction fields.

### 3.4. Electron Train Behavior in LZ


This work will summarize the basic characteristics of electron trains in LZ. An example of a typical electron train in LZ is shown in Figure 3.17.

PMT hitmaps from both events reveal that the SEs in the following event are co-located with the lateral position of the progenitor S2.

For this analysis, progenitors are S2s at least 200 extracted electrons in size to induce significant delayed electron emission, and must be part of a single scatter event. A drift time cut of 50  $\mu\text{s}$  to 950  $\mu\text{s}$  and a radial cut of 55 cm are used to select progenitors with well-understood energy and position properties. The drift time cut follows from the fiducial cut implemented in LZ's




(a) A single scatter event in LZ. The progenitor S2 pulse at 6.2 ms is followed by a tail of misreconstructed pulses from photoionization up to approximately 7.5 ms. As the excess light subsides, SEs (SE, marked with red dots) are visible until the event ends.




(b) In the following event, more SE pulses are visible.

Figure 3.17.: The first two events of an electron train in LZ, showing the presence of SEs for multiple milliseconds after an S2 pulse.



(a) PMT hitmap from the event shown in Figure 3.17a. The excess of light at the 3 o'clock position in the top PMT array corresponds to the progenitor S2



(b) PMT hitmap from the event shown in Figure 3.17b. Light is still detected in the TPC in the form of stray single photoelectrons (SPEs) and SEs. The persistent excess of light at the 3 o'clock position is due to SEs which constitute an electron train.

Figure 3.18.: PMT hitmaps from the first and second events shown in Figure 3.17. Hitmaps on the left correspond to the top PMT array, while hitmaps on the right correspond to the bottom PMT array.

first science run, SR1 [92], while the radial cut removes “glue-ring” events which originate from radioactive decays on the gate grid at high radii but may have their positions mis-reconstructed. After each progenitor, all small electron pulses are identified in randomly triggered data acquisition windows, and the process stops when a vetoing pulse (any pulse larger than 100 extracted electrons) is detected. Progenitors followed by a vetoing pulse within 5 ms are discarded in order to ensure that a sufficient amount of time has elapsed to observe SEs. Progenitors also must be at least 200 ms after any vetoing pulse to mitigate the aftermath of earlier energy depositions. As an additional precaution, progenitors are required to end at least 1  $\mu$ s before the end of an event window to ensure that they are properly resolved. Finally, events with otherwise valid progenitors are rejected if any data buffer for storing PMT digitizations is saturated. This cut effectively excludes progenitor pulses larger than  $10^6$ phd and largely mitigates the effects of PMT saturation. Additional technical details on the data extraction algorithm and theory of pulse rate measurement are given in Section 3.3.2.

Plots shown in this work are created using data from either SR1, or runs 6726-6756 which were part of the pre-SR1 commissioning campaign. Grid voltages in both datasets were identical; the anode, gate, and cathode voltages were -4,-4, and -32 kV respectively. This corresponded to a nominal liquid extraction field of 3.9 kV/cm and drift liquid field of 193 V/cm [92]. A key difference between each set of data are the trigger settings. SR1 utilized an S2-based trigger which had a trigger efficiency of  $\sim 100\%$  for S2 pulses  $> 5$  SE in size, and approximately 10% for SE pulses [92]. Meanwhile, runs 6726-6756 utilized a purely random  $\sim 87$  Hz trigger with 11.1 ms long event windows, resulting in a pulse-agnostic  $\sim 95\%$  livetime. Despite this difference in trigger configurations, the similar SE trigger efficiencies resulted in a negligible impact to SE rate

measurements in SR1 data. Nevertheless, data from runs 6726-6756 is presented in this article unless the high statistics from SR1 were necessary.

Finally, data in this and other work is normalized by combinations of the following parameters shown in Table 3.1, depending on the circumstance.

Table 3.1.: Table of child pulse rate normalization factors.

Factor	Description
$e_R = S2_{phd}/n_{phd}/SE$	“Raw” (extracted) S2 area in units of SEs
$e_S = e_R/e_{eee}$	“Surface” S2 area, i.e. $e_R$ corrected for extraction efficiency
$e_I = e_S \exp(t_{drift}/\tau_{e^-})$	“Initial” S2 area, i.e. $e_S$ corrected for drift losses
$e_L = e_I - e_S$	Number of electrons lost while drifting
$\text{cm}^2$	Area of liquid surface subtended by radial selection of pulses

### 3.4.1. Position and Time Dependence of SEs

This position and time dependence is a hallmark of electron trains which are known to persist for at least a second after an S2 occurs. This can be seen in the SE rate relative to an S2, shown in Figure 3.19.

The stark disparity in rates within 10 cm of a progenitor lends itself to the definition of “position-correlated” and “position-uncorrelated” pulses (Table 3.2):




Figure 3.19.: The rate of SEs after S2s, per square centimeter of liquid surface in the TPC. The x-axis is the radial distance between between the child SE pulse and the most recent progenitor S2 pulse. The stack axis denotes the time since the end of the progenitor S2.

Table 3.2.: Definitions of position-correlated and uncorrelated child pulses.

Position	Distance to Progenitor [cm]
Correlated	$< 10$
Uncorrelated	$20 < \Delta r < 30$

These definitions were chosen by eye to capture the rates of each population without leakage between the two. Furthermore, the position-uncorrelated region was chosen to only minimally extend beyond the walls of the TPC in the worst-case where a progenitor occurs at a radius of 55 cm. To further correct for this scenario, an additional scaling factor is applied to counts of SE pulses at  $\Delta r$  which are commensurate with a circumference that extends beyond the TPC.

This scaling factor is inversely proportional to the fraction of the circumference of the circle defined by  $\Delta r$  that is contained inside the TPC. The intent of this correction rests on the assumption that a uniform number of counts per unit distance are expected along the circle defined by  $\Delta r$ . If a portion of this circle exists beyond the walls of the TPC, the additional weighting will appropriately scale up the number of detected counts at the value of  $\Delta r$ , effectively simulating a TPC with an infinitely large radius. The effects of this correction are minor for  $\Delta r$  up to approximately 30 cm, but result in a drastic flattening of the position-uncorrelated SE rate at higher  $\Delta r$ . This correction is especially elegant because it allows for progenitors to be treated on equal footing, regardless of their position in the TPC, eliminating the need for specialized corrections to livetime.

A diagram detailing this scaling factor is shown in Figure 3.20. For a progenitor at radius  $R_{prog}$  inside a TPC with radius  $R_{TPC}$ , and a SE pulse at  $R_{SE}$  relative to the progenitor, the following are always true:

A progenitor must be found within the bounds of the TPC:

$$0 \leq R_{prog} < R_{TPC}$$

And the farthest away a SE can be found is at the opposing end of the TPC:

$$0 \leq R_{SE} < R_{TPC} + R_{prog}$$

The angle  $\theta$  defined in Figure 3.20 subtends half of the arc inside of the TPC. Thus, the fraction of the circumference of the circle inside of the TPC is equal to:




Figure 3.20.: Diagram showing geometric constraints needed to correct for the lack of pulses observed relative to a progenitor when pulses are observed at a radius  $R_{SE}$  which extends beyond the walls of the TPC. A child pulse may be observed anywhere inside the TPC along the circle defined by  $R_{SE}$ . Counts of pulses along this arc are then used to correct for the lack of pulses observed along the arc outside of the TPC.

$$\frac{2\theta}{360}$$

While  $\theta$  is found using the law of cosines:

$$\theta = \arccos\left(\frac{R_{SE}^2 + R_{Prog}^2 - R_{TPC}^2}{2R_{SE}R_{Prog}}\right)$$


The position-correlated region captures a power law in the rate vs time, while the position-uncorrelated rate decays much faster (Figure 3.21a).

A power law of the following form can be fit to the position-correlated rates, as shown in Figure 3.21b:


$$\alpha t^\beta \tag{3.6}$$

A timescale of 3 ms to 200 ms was chosen for the fit range. 3 ms was chosen to avoid residual photoionization effects, while 200 ms was chosen as an upper bound because the power-law slope is shallower due to contamination from position-uncorrelated delayed electrons. The large difference in position-correlated and uncorrelated rates between 3-200 ms also obviates the need for background subtraction because the change in the position-correlated rate after subtraction is negligible in this region.

It is also worth emphasizing that the drop in the rates prior to 1 ms is a known artifact of pulse pile-up from photoionization. An example of this is visible in Figure 3.17a where many of the S2 pulses immediately after the progenitor S2 are piled-up SEs. This effect can be seen more concretely in Figure 3.22 when examining rates of S2 pulses after progenitors, instead of SEs. Unlike



(a) Rates of position-correlated and position-uncorrelated SE pulses at time  $\Delta t$  after a progenitor S2 pulse. The rates are converted to a flux using the definitions given in Table 3.2.



(b) Figure 3.21a with power law fit to rates of position-correlated SEs. The power law exponent is consistent with results from LUX [81] and XENON1T [80]. A timescale of 3 ms to 200 ms was chosen for the fit range. This range was chosen because the position-correlated rates are dominant in this region, to the extent that subtracting the position-uncorrelated rates has a negligible effect on the fit. Prior to 3 ms the slope is steeper from residual photoionization effects. After 200 ms, the power-law slope is shallower due to contamination from position-uncorrelated delayed electrons.

Figure 3.21.: Rates of position-correlated and uncorrelated SE pulses versus time since the end of their progenitor S2. The rates prior to 1 ms are a known artifact of pulse pile-up from photoionization (see also Figure 3.22). Position-uncorrelated rates drop off quickly after the full drift time has elapsed, in comparison to the rates from position-correlated pulses.

with SEs, S2 pile-up does not leak into a different classification in LZ. Rates may still be distorted because counts of S2s can be reduced, but the S2 classification is ultimately more resilient against the effects of pile-up.




Figure 3.22.: Rates of position-correlated SE and S2 pulses versus the size of their respective progenitors. For the smallest progenitors, SE rates appear to decay monotonically. As the progenitor size increases, pulses lose their SE classification, causing a distortion in the rate prior to 1 ms. This feature is not present when selecting for S2 pulses because pulse classification is conserved for S2s undergoing pile-up.

### 3.4.2. Progenitor Size, Drift Time, and Electron Lifetime Dependence

The rates of delayed position-correlated SEs depend on the size of the progenitor pulse, while position-uncorrelated SEs do not (Figure 3.23).

Progenitors drifting through more liquid have higher rates of delayed electrons following them. In Figure 3.23 rates are normalized by  $e_I$ , the size of the progenitor at the interaction site, in units of SEs.  $e_I$  is a function of the drift time  $t_{drift}$  and electron lifetime  $\tau_e$ :

$$e_I = e_S \exp(t_{drift}/\tau_e)$$

It is proportional to the size of the progenitor at the liquid surface  $e_S$ , which is the size of




Figure 3.23.: Rates of SEs versus the size of their progenitor,  $e_R$ , in the number of electrons extracted. A time window of 3-300 ms was chosen to avoid integrating any residual photoionization. The position-uncorrelated rates do not exhibit a dependence on the size of the progenitor; this behavior was also observed in XENON1T [80]. Most progenitors are between  $10^3$  and  $10^4$   $e_R$  in size; SR1 data was used in this plot to maximized the available statistics.

the progenitor corrected for the extraction efficiency  $e_{eee}$ :

$$e_S = e_R/e_{eee}$$




Figure 3.24.: Rates of SEs versus the drift time of their progenitor. The position-correlated rate appears to linearly depend on the drift time of the progenitor even after normalizing by  $e_I$ , the size of the progenitor at the interaction site. The position-uncorrelated rates do not exhibit a dependence on the drift time of the progenitor. This behavior was also observed in XENON1T [80], and highlighted as an indication that uncorrelated delayed electrons may be correlated with previous progenitors. SR1 data was used in this plot to maximize the available statistics.

It is worth highlighting here the lack of correlation exhibited by position-uncorrelated pulse rates in both Figure 3.23 and Figure 3.24. This would indicate that position-uncorrelated pulses originate from some source other than the most recent progenitor. This is compatible with the explanation put forward by the XENON1T collaboration that these may be delayed electrons from

previous electron trains [80].

Figure 3.25 shows the same data from Figure 3.24 except with rates plotted against the time since the progenitor S2. Counterintuitively, the position-correlated photoionization rates in Figure 3.25 following progenitors at lower drift times also appear to be lower. This is unexpected because with fewer electrons lost, S2s should be larger and the TPC should experience more photoionization. A possible explanation for this could be that a suppression of the observed SE rate occurs due to an increased pile-up rate (see again Figure 3.22). This seems unlikely though, as LUX [81] observed this even when small progenitors were explicitly studied to avoid the effects of pile-up masking SEs. Position-correlated photoionization rates are also lower at lower electron lifetimes (Figure 3.27), but this feature is less straightforward to understand because while more electrons are lost at lower electron lifetimes, more impurities are present which could be photoionized in the TPC, leading to pile-up suppression of the SE pulse classification. As this work focused on the behavior of delayed electrons, this question remains unaddressed, but may be the subject of future work.

During SR1, the electron lifetime reached 8 ms, offering a unique opportunity to measure the SE rate dependence at electron lifetimes that had not previously been measured. Figure 3.26 shows that the rate decreases with larger electron lifetimes, although the effect tapers off at higher electron lifetimes. This behavior is in agreement with that observed by XENON1T, where it was taken as an indication that the species responsible for electron trains is distinct from typical electronegative impurities like oxygen that affect the electron lifetime at all concentrations. The reason for the sharp drop in the position-correlated rate at electron lifetimes  $> 5$  ms is unknown. A possible reason could be the sampling of electron lifetime data in LZ multiple days apart. Electron lifetime data in the




Figure 3.25.: Drift time dependence of SE rates versus time since the progenitor S2. The position-uncorrelated rates do not exhibit a clear drift time dependence at any timescale. Position-correlated photoionization rates have a counterintuitive dependence on the drift time. At shorter drift times, S2s should lose fewer electrons and photoionize the TPC more. However, lower photoionization rates are observed in this case, and this phenomenon was also observed in LUX [81]. This behavior remains unexplained, and should be the subject of future work.

plot was generated by interpolating an empirical plot of electron lifetime versus run-time, and while it was typically stable over time, several temporary losses of of purity occurred. On these occasions the purity recovered quickly—much faster than the electron lifetime sampling rate—likely leading to an inaccurate interpolation of electron lifetime values. Meanwhile, the position-uncorrelated rate exhibits a slight dependence on the electron lifetime. This is to be expected if the position-uncorrelated SEs are from prior progenitors; the electron lifetime does not vary on the sub-second timescale between progenitors.

### 3.4.3. Electron Loss Normalization

The drift time (Figure 3.24) and electron lifetime (Figure 3.26) dependence persist in the SE rates despite normalization by  $e_I$ . These rates over time are explicitly shown in Figure 3.25 and Figure 3.27, which each compare the rates of position-correlated and position-uncorrelated SEs.




Figure 3.26.: Rates of SEs versus the electron lifetime recorded at the time of the progenitor pulse. The rate at higher electron lifetime values appears to taper off, in agreement with behavior observed by XENON1T [80] and their explanation that the impurity or impurities responsible for electron trains does not perfectly track the concentration of impurities which affect the electron lifetime. A sharp drop in the position-correlated rate is apparent after the electron lifetime exceeds 5 ms. This may be an artifact of interpolating electron lifetime data, as explained in the text. The position-uncorrelated rate also exhibits a slight dependence on the electron lifetime, which would be consistent with the explanation from XENON1T [80] that these backgrounds are mostly from prior electron trains.




Figure 3.27.: Electron lifetime dependence of SE rates versus time since the progenitor S2. A weak dependence on the electron lifetime is visible for both the position-correlated and uncorrelated rates at all timescales after the full drift time. The position-uncorrelated rate in the 2-4 ms electron lifetime bin is noticeably uneven due to wrongly associated pairs of SEs and progenitor S2s. This occurs for all other position-uncorrelated rates in both this figure and Figure 3.25, but the effect is enhanced with lower statistics.


A common hypothesis on the origin of electron trains is that they are generated from the release of electrons by electronegative impurities. After normalization by  $e_I$  however, the rates still vary linearly with drift time (Figure 3.24). Therefore, it is reasonable to consider an alternative normalization that better represents the mechanism responsible for the drift time dependence.

As it drifts through the liquid, a progenitor S2 loses electrons to electronegative impurities. The number of electrons lost,  $e_L$ , should be the difference between the number of electrons at the interaction site and at the surface:


$$e_L = e_I - e_S$$

Rates of SEs after normalizing by  $e_L$  instead of  $e_I$  are shown in Figure 3.28. This normalization negates both the drift time and electron lifetime dependence effectively unifying both. It further demonstrates the crucial role that electronegative impurities play in the formation of

electron trains.



- (a) Drift time dependence of SE rates versus time since the progenitor S2. The 50-230  $\mu\text{s}$  band remains elevated above the other bands, indicating that disproportionately fewer electrons may be lost in this region. It is unclear as to why this may be the case. Electric field non-uniformity is not plausible, as the drift field is well-established within a drift time of a few  $\mu\text{s}$ .



- (b) Electron lifetime dependence of SE rates versus time since the progenitor S2. The stratification of photoionization rates is expected because a larger number of electrons lost will depress rates more at low electron lifetimes, and high drift times.

Figure 3.28.: The same rates from Figure 3.25 and Figure 3.27. instead normalized by  $e_L$ , the number of SEs lost by the progenitor as it drifts through the liquid. Both the drift time dependence and electron lifetime dependence that were visible in Figure 3.25 and Figure 3.27 are negated, indicating that impurities which capture electrons in the liquid are primarily responsible for electron trains.

### 3.5. Evidence of a Drift Field Dependence

This work demonstrates hints of a new drift field dependence in the electron train power law. While a definitive conclusion is prevented by a combination of low statistics and competing backgrounds, the result is consistent with the hypothesis that electron trains originate in the liquid bulk of a TPC. This conclusion is reached via a novel analysis of electron trains from the extraction region liquid in LZ. Previous studies have examined electron trains with TPCs consisting solely of an extraction region. However, their drift and extraction fields were inextricably coupled, preventing them from isolating a drift field dependence. In TPCs with independent extraction and drift regions, the extraction liquid has been overlooked. This work demonstrates that at low drift fields the field dependence is minimal, and that electron trains from drift region events dominate the power law rate in a TPC. Thus any changes in the rate that may originate from the extraction region are washed out. A survey of past studies also shows the presence of a drift field dependence may have gone unrecognized.

The extraction liquid is an ideal location to study the drift field dependence of electron trains. It is isolated from the drift region and capable of sustaining much higher fields. At the same time, events from the drift region are readily accessible, offering an *in situ* comparison to electron trains at a much lower drift field and the same extraction field. A second benefit is that the high field and low depth make for an exceptionally short region where corrections for electrons lost to impurities are minimal.

For this study, five groups of data were used where grid voltages were varied between each group to examine the effects of electric fields on SE rates. Extraction voltage differences of 7, 8, and 9kV and drift voltage differences of roughly 16, 28, and 35 kV were used. The fields in each

region of interest in the TPC were simulated<sup>4</sup> for each voltage configuration. Figure 3.29 shows an example of one simulation in the extraction region of the TPC.




Figure 3.29.: Simulated field strengths in the extraction region of the LZ TPC. Key features such as the anode and cathode are indicated in the plot. The cathode is not pictured; it is located at  $Z = 0$  mm and is the lower boundary for the drift liquid. The TPC walls are located at  $R \approx 740$  mm; only the first 50 mm are shown here to highlight the scale of the grid spacing in  $R$ .

Grid voltages and the nominal fields in each region of interest are listed in Table 3.4. The values are elevated by  $\sim 0.5\text{-}2\%$  because a correction for a reduced gate voltage was not implemented when the simulations were re-run. The quoted field values were determined by averaging over the magnitude of the field in the respective regions of TPC. These regions are given in Table 3.3. Regions were chosen by eye to be representative of values in the bulk of each region.

<sup>4</sup>Details on simulations of LZs electric fields can be found in [93]. These simulations were repurposed for this work to determine fields at each of the grid voltage conditions.

Table 3.3.: Coordinates used to determine bounds for averaging fields within the TPC. Regions were chosen by eye to be representative of values in the bulk of each region.

Bound for Field			
Average	Radial [cm]	Extraction Liquid [cm]	Drift Liquid [cm]
Lower	5.0	145.8	25
Upper	55.0	146.0	125

It is worth clarifying that for some groups of voltage settings the cathode voltage difference was only *nominally* 28 kV, leading to slightly varied drift fields of 197, 194, and 191 kV/cm when the gate voltage was adjusted. For the sake of this analysis, these three drift fields are grouped together as being nominally 194 kV/cm.

Table 3.4.: Grid voltages and corresponding nominal field values from simulations. For some groups the cathode voltage difference was only *nominally* 28 kV, leading to slightly varied drift fields of 197, 194, and 191 kV/cm when the gate voltage was adjusted. For the sake of this analysis, these three drift fields are grouped together as nominally 194 kV/cm.

Anode [kV], Gate [kV], Cathode [kV]	Drift Liquid [kV/cm]	Drift Liquid [kV/cm]
(3.5, -3.5, -32.1)	197	3487
(4.0, -4.0, -39.1)	242	3986
(4.0, -4.0, -32.1)	194	3984
(4.0, -4.0, -20.1)	112	3981
(4.5, -4.5, -32.1)	191	4481

Finally, all data was acquired with a random trigger during pre and post-SR1 grid commissioning campaigns. The list of runs, grouped by voltage settings, is shown in Table 3.5.

Table 3.5.: Table of runs, grouped by voltage settings. Tuples indicate consecutive sets of runs e.g (9199, 9201) corresponds to 9199, 9200, 9201. LZ Users: see [Fields Investigation Data](#) spreadsheet for more information and `sv_SEDecayTime_lr_EdV_7_8_9_DdV_16_28_35.list` in ALPACA module [89] inputs.

---

Extraction  $\Delta V$ ,

Drift $\Delta V$ [kV]	Runs
(7.0, 28.6)	[(9174, 9175), (9178, 9180), (9183, 9185), (9188, 9191), (9194, 9196), (9199, 9201), 9204, (9206, 9209), (9211, 9213), (9216, 9218)]
(8.0, 16.1)	[(8959, 8961), (8964, 8966), 8972, (8974, 8976), (8979, 8981), (8984, 8986), (8989, 8992), (8996, 8997)]
(8.0, 28.1)	[(6726, 6734), (6736, 6748), (6750, 6762), (9370, 9371), (9375, 9378), (9383, 9385), (9388, 9390), (9393, 9395), (9574, 9576), (9579, 9581), (9584, 9586), (9589, 9591), 9595, 9597, (9600, 9602), (9605, 9606), 9608, (9611, 9613)]
(8.0, 35.1)	[(9496, 9498), (9501, 9503), (9510, 9511), (9516, 9518), (9521, 9523), (9530, 9532), (9535, 9537), (9541, 9543), (9546, 9547), 9549, (9553, 9555), 9558, 9570]
(9.0, 27.6)	[(6682, 6719), 6721]

---

### 3.5.1. Electron Train Origins

The position-correlated component of electron trains is known to depend on the drift time of the previous event (Figure 3.25) [80] [81]. This dependence persists even when rates are normalized by the drift-corrected S2 size  $e_I$ :

$$e_I = e_S \exp(t_{drift}/\tau_{e^-}) \quad (3.7)$$

Where  $e_S$  is the number of SEs at the liquid surface that comprise the S2,  $t_{drift}$  is the drift time of the progenitor S2, and  $\tau_{e^-}$  is the electron lifetime.  $e_S$  is related to the extracted S2 size  $e_R$  by the electron extraction efficiency  $e_{eee}$ :

$$e_S = e_R/e_{eee} \quad (3.8)$$

The rate increase with drift time has been taken as an indication that electron trains originate from some phenomenon in the liquid bulk of the TPC. A common hypothesis is that as electrons from an S2 are drifted upward, they are captured by impurities. After being captured, they are then released later via an unknown mechanism that exhibits a power law time dependence.

Thus at longer drift times, more electrons are captured and released later as electron trains. The low diffusion velocities of impurities ensures that electron trains are fairly localized in XY relative to the progenitor, and removal of impurities lessens the electron train intensity. Finally, more energetic events with larger S2 pulses would see more electrons lost in the liquid as they drift upwards. All of these qualities have been observed in LUX and XENON1T [81] [80].

An alternative hypothesis is that electrons could be trapped at the liquid surface and emitted later. This hypothesis has difficulty explaining the drift time [94] and purity dependence. Removal of impurities should extend the electron lifetime at the surface, increasing the rates of electron trains. Furthermore, since progenitor S2s at higher drift times would arrive with fewer electrons at the liquid surface, fewer electrons should be left behind.

Nevertheless, there is evidence that electrons can be trapped at the liquid surface in a TPC [95], resulting in delayed emission phenomena that are longer in duration than an S2, but much faster than electron trains. Despite the significantly different topologies of electron trains and “S3”

[96] or “electron burst” [81] pulses, the contribution of surface electrons to electron trains is still an open question.

### 3.5.2. Dependence on TPC Region

XENON1T showed that electron trains are virtually non-existent for events originating in the gas above the liquid surface and the liquid below the cathode<sup>5</sup>. This is compelling evidence that photoionization alone cannot induce electron trains. As shown in Figure 3.30, SE rates from photoionization drop sharply after the maximum drift time when the event occurs in the gas or below the cathode. The rates plateau near the background levels (denoted by a solid purple line) within 10 ms.

Figure 3.30 essentially confirms that electron trains emanate from liquid of the TPC. However, it does not include a distinction for liquid in the extraction region. Compared to the drift region, fields in the extraction region are 10-20x higher to facilitate the extraction of electrons into the gas to generate electroluminescence signals. It is difficult to maintain fields this high over long distances, so the extraction region is also exceptionally shallow in most TPCs, usually a few millimeters to centimeters deep depending on the scale of the TPC.

The high fields and short depth make for a region where S2 electrons produced in the extraction liquid spend an exceptionally short period of time being drifted out. Following from the known drift time dependence, this should result in less intense electron train rates than what would be expected from events in the drift region.

---

<sup>5</sup>This is commonly known as the “reverse field region” or RFR because the drift field below the cathode is reversed. Any S2 charge produced here will be drifted downward, away from the liquid surface. RFR S1 pulses are useful for studies of photoionization backgrounds [81] [97] [98].




Figure 3.30.: Electron trains from progenitors in the gas, drift, and below cathode regions of XENON1T [80]. A 200 ms holdoff was imposed after interactions to avoid overlapping electron trains between progenitors. The SE rate during this 200 ms holdoff period is shown with the solid purple line. It served to verify that the holdoff period was adequate in allowing SE rates to return to a baseline level before SEs from progenitor pulses were tracked. It also showed that rates following progenitors in the gas and reverse field regions returned fairly quickly to the pre-holdoff rates, indicating that electron trains are linked to drift region interactions.

### 3.5.3. Previous Extraction Region and Drift Field Studies

A study of the electron train time dependence in the extraction region has been conducted on two occasions. The first study, by Akimov *et al.* in 2016 [96] was the first to demonstrate a now widely-observed [81] [80] [94] power law-like rate (Figure D.9) following interactions in the RED-1 TPC.




Figure 3.31.: Power law in SE rate observed by Akimov et al. after muon S2s.

The RED-1 TPC was bounded by an anode and cathode 22 mm apart and a drift field of 3.75 kV/cm was used. The study did not vary any operating parameters, so a single power law fit was obtained. The exponent of 1.4 is noteworthy for being the steepest of any study which has observed a power law.

The second study of electron trains in the extraction liquid was performed at LBNL in 2018 [99]. The TPC consisted of only an anode and cathode 5 mm apart, with 4 mm of active liquid. An aqueous <sup>210</sup>Po source (2.6 MeV  $\alpha$ ) was deposited on the cathode to generate events at a consistent

depth and energy. The study notably observed a two-component exponential dependence in a summed waveform after the maximum drift time of hundreds of alpha events (Figure 3.32).




Figure 3.32.: Low-pass filtered summed waveform of electron trains from alpha events with a “fast” and “slow” exponential component.

Inspections of single waveforms showed SEs and photons appearing at any time within the full 1 ms length of the event window. The summed waveform was analyzed as a proxy of these spurious pulses. The electric field in the liquid was varied from 3-6 kV/cm and datasets were acquired at a “low” and “high” purity. The dependence of the time constants and their amplitudes is summarized in Table D.1.

Table 3.6.: Dependencies of fast and slow regions in a summed waveform of alpha events [99]. As an example, the amplitude of the fast component increased at a higher purity and decreased with an increasing applied electric field.

Component	Purity	Field
Fast Amplitude	Correlated	Anti-correlated
Slow Amplitude	Anti-correlated	Correlated
Fast Decay Time	Anti-correlated	Correlated
Slow Decay Time	Anti-correlated	Correlated

The summed waveform analysis and short observation timescale make it difficult to make direct comparisons to other studies. Furthermore, the coupled nature of the extraction and drift fields in these experiments limit their ability to definitively distinguish between extraction and drift field effects.

Two more studies have been conducted where the drift field was varied independently of the extraction field. The first, by Kopec *et al.* [94], found that the power law amplitude increased with the drift field between 100-1000 V/cm. The exponent was unchanged and reported to be 1.2 at a drift field of 500 V/cm. In the second study, Bodnia *et al.* [100] examined the rates of pre-S1 SE pulses and found that the rate decreased slightly with a drift field between 100-1000 V/cm.

Table 3.7 summarizes the results of studies which have fit a power law to electron train rates.

Table 3.7.: Table of electron train power law exponents from previous studies.

Study	Exponent	Drift Field (V/cm)
LUX [81]	$[-1.0, -1.1]$	180
XENON1T [80]	-1.1	125 [101]
Kopec <i>et al.</i> [94]	$-1.2 \pm 0.04$	500
Akimov <i>et al.</i> [96]	$-1.4 \pm 0.2$	3750

Highlighted in Table 3.7 are the respective drift fields of each study, demonstrating that the power law exponent may become more negative with an increasing drift field. Given this hint of a dependence, it seems reasonable to believe that events originating in the extraction region liquid of LZ should have a more negative power law exponent. This warrants a more detailed study of electron trains in the extraction region, which is the focus of this work.

### 3.5.4. Rates in TPC Regions of LZ

For this investigation, single-scatter progenitor S2s from the gas, extraction liquid (EXL), and drift liquid (DRL) were selected. Gas events were selected using two cuts developed to identify them in LZ's first WIMP search, Science Run 1 (SR1). The first cut was the `S2EarlyPeak` cut<sup>6</sup>, which was designed to select events where an S1 and S2 had merged. Because electroluminescence occurs almost immediately in the gas, S1 and S2 pairs in gas events were frequently merged. This led to a skewed full-width half-max in the primary S2, which formed the basis of the `S2EarlyPeak` cut. The second cut was the `GasS2TBACut`<sup>7</sup> which set a bound on the maximum top-bottom asymmetry (TBA) for an S2. S2s failing the `GasS2TBACut` were mostly located above the anode.

Following this, progenitor S2s between the gate and liquid surface were selected with a drift time between 0 and 2.5  $\mu\text{s}$ . While a 0  $\mu\text{s}$  drift time was technically unfeasible, the single scatter cut effectively required a minimum drift time of approximately 1  $\mu\text{s}$  for the S1 and S2 pulses to be properly resolved. Finally, a drift time cut of 50 to 950  $\mu\text{s}$  was used for selecting drift liquid events, similar to the drift time cut used to define LZ's fiducial volume in SR1 [92]. The overall rates in different TPC regions are shown in Figure 3.33.

The intention behind independently selecting for gas events was to enable a qualitative comparison to results from XENON1T (Figure 3.30). In agreement with their results, ionization backgrounds after gas events appear to be comprised mostly of photoionization. Delayed electron rates following gas-based progenitors are heavily suppressed. Delayed electron rates following extraction liquid events appear to be slightly higher, but are still much less intense than those from the drift region.

---

<sup>6</sup>Developed by R. Smith and S. Dey.

<sup>7</sup>Developed by R. Smith




Figure 3.33.: SE rates following single scatter events in the gas above and below the anode, extraction liquid above the gate grid, and drift liquid between the gate and cathode of LZ. Delayed electron rates in the gas and extraction liquid are much less intense than those in the drift liquid. This plot combines the rates from both position-correlated and uncorrelated pulses, which is indicated by No Cut (NC) to the relative pulse positions in the legend. No area normalization is used in this plot in order to enable direct comparisons to Figure 3.30.

The power law dependence in the drift region is obscured in Figure 3.33 because only position-correlated electron trains exhibit a power law dependence. Position-uncorrelated electrons are still captured in this rate, but because those are dominated by photoionization, the power law is warped downward. To recover the power law dependence, rates from each region are split into their respective position-correlated and uncorrelated components (Figure 3.35).

Position-correlated electrons are within 10 cm of a progenitor S2 and position-uncorrelated electrons are between 20 and 30 cm away from the location of the progenitor. This distinction was chosen by eye from the relative position distribution of SEs, which have a fairly minimal variation for  $\Delta r > 30$  cm (Figure 3.34). While position-uncorrelated electrons can be found more than 30 cm away from a progenitor, this limit was chosen to prevent regions outside of the TPC from contributing to the rates for progenitors which are closer to the TPC walls.




Figure 3.34.: Rate of SEs shown in Figure 3.19, projected in time since the progenitor S2. Normalizing rates by the area subtended by each interval in  $\Delta r$  reveals that the flux of position-uncorrelated SEs is fairly uniform for  $\Delta r > 20$  cm.

Rates are normalized by an additional factor equivalent to the area covered by the radial selection (Section 3.4.1). This “rate flux” enables a better comparison to be made between different sized zones of the detector. Without this area normalization, rates from the 10 to 50 cm region are substantially higher since that region comprises a larger space in the detector where electrons can be emitted from.

Figure 3.35 shows that gas events have the lowest rates of position-correlated electron trains. However, the relatively large fluctuations in these rates over time compared to the rates in other regions indicate mis-associated S2 normalizations. Handscanning also reveals that these position-correlated electron trains appear to be a consequence of leakage from events in the extraction liquid. Meanwhile, position correlated pulses from drift liquid events have a clear and prominent power-law like rate, as expected. Position-correlated rates in the extraction liquid appear to have a power law component, but with a much lower amplitude than those in the drift liquid. As such, while the power law exponents in both regions are visually similar, further consideration of the position-uncorrelated background in the extraction liquid events is merited.

### 3.5.5. Fits to Rates

Given the appearance of a power-law like structure in the extraction liquid SE rates, it is worth comparing with the rate from events in the drift region. For plots in this analysis, a power law of the following form is used (Equation 3.9):

$$\alpha t^\beta + c \tag{3.9}$$

where  $\alpha$  is the amplitude,  $t$  is the time since the progenitor S2,  $\beta$  is the rate exponent, and




Figure 3.35.: SE rates from Figure 3.33, split into the position-correlated and uncorrelated components. The position-correlated rates in the gas are largely from drift liquid events which leaked into the gas event selection. The relatively large fluctuations in these rates over time compared to the rates in other regions also indicate mis-associated S2 normalizations. Position-correlated rates in the extraction liquid appear to have a power law component, but with a much lower amplitude than those in the drift liquid. The exponents in these two regions are visually similar, but fits (Section 3.5.5) indicate that the power law in the extraction liquid may be slightly steeper.

$c$  is an optional constant. The constant  $c$  is added to approximate a constant background rate at long timescales, and is denoted on plots where it is used.

### 3.5.5.1. Drift and Extraction Liquid

In fitting to the position-correlated rate in the drift region (Figure 3.36), the exponent is very similar to that observed in previous studies (Table 3.7).

The rate from the extraction liquid (Figure 3.37) has a more clear position-uncorrelated background component. A constant term in the fit approximates the presence of position-uncorrelated backgrounds at long timescales, but does not account for residual photoionization backgrounds at short timescales. As a consequence, there is a large systematic uncertainty in  $\beta$  associated with the left edge of the fit. Thus while the power law appears to be steeper, its small amplitude relative to the amplitude of the position-uncorrelated background motivates subtracting the




Figure 3.36.: Power law fit to the position-correlated SE rate following events originating in the drift liquid shown in Figure 3.35.

position-uncorrelated background.

### 3.5.5.2. Background Subtraction

The position-uncorrelated area-normalized rate is fairly uniform in  $\Delta r$  (Figure 3.34). Given that the extraction region position-correlated and uncorrelated rates are similar in magnitude, it is worth trying to subtract the position-uncorrelated rate from the position-correlated rate in order to diminish the effects of photoionization backgrounds contaminating the power law fit in Figure 3.37.

After background subtraction (Figure 3.38), the power law appears to extend almost 100 $\mu$ s after the progenitor S2 pulse. However, an analysis of fit parameters shows that the residual photoionization background prior to 1ms still significantly influences the exponent of the power law (Figure 3.39).




Figure 3.37.: Power law fit to the position-correlated SE rate following events in the extraction liquid shown in Figure 3.35. A constant term in the fit approximates the presence of position-uncorrelated backgrounds at long timescales, but does not account for residual photoionization backgrounds at short timescales. As a consequence, there is a large systematic uncertainty in  $\beta$  associated with the left edge of the fit.




Figure 3.38.: Subtracting the position-uncorrelated SE rate in Figure 3.35 from the position-correlated rate in Figure 3.35 removes many of the photoionization backgrounds in the position-correlated SE rate between 100 $\mu$ s and 1ms. It also removes the background component after 100ms.




Figure 3.39.: Fitted value of  $\beta$  in Figure 3.38 as a function of fit limits. Subtracting the position-uncorrelated rate leaves behind some discontinuities in the rate. First, rates before 1ms are still slightly elevated. Second, rates after  $\sim 20$ ms are highly dependent on the intensity of the position-uncorrelated flux, which increases with the field.

The power law exponent of  $-1.33 \pm 0.09$  in Figure 3.38 is significantly steeper than the usual exponent of approximately -1.1 observed by LUX and XENON1T. It also bears some similarity to the exponent of -1.4 obtained by Akimov *et al.* at a field of 3.75 kV/cm with a drift depth of 2.2cm.

The change in the exponent is likely not due to the change in drift time. Fitting to the position-correlated rates shown in Figure 3.25 shows that the amplitude of the position correlated flux increases with drift time but the exponent remains virtually unchanged (Figure 3.40).




Figure 3.40.: Power law fits to position-correlated rates in Figure 3.25 show minimal variation in the exponent and a clear trend in the amplitude that increases with drift time.

Instead, the apparent change in the rate exponent may be an indication that the electron train production mechanism is affected in a unique way by the high fields. While a complete model of the power law is beyond the scope of this work, a qualitative comparison can be made to outgassing rates of materials in a vacuum. Outgassing rates over time are known to follow a power law [102], which is modeled by assuming that adsorbed gases are released exponentially over time.

Two more assumptions result in a sequence of exponentials spaced out in time, whose envelope forms a power law. The first assumption is that binding energies can vary between adsorption sites. The second assumption is that repumping may occur, i.e. that a desorbed gas may re-adsorb on the walls of a container. The slope of the power law is also material-dependent [103,104].

Assuming that electron trains are captured by impurities in the liquid bulk, the drastic increase in drift field could lead to electrons being released *sooner* back into the liquid. This explanation was previously proposed by Kopec *et al.* [94] however it was ruled out because they did not observe a change in the power law exponent when sweeping over a drift field of 0.1 to 1kV/cm. In light of this new information however, this could occur if multiple impurities are present, and if the increased drift field results in increased collisional energy between those impurities and neighboring atoms.

### 3.5.6. Varying the Extraction Field

The prior analysis of events in the extraction liquid can be extended to different extraction fields in order to further examine the dependence of the power law at higher and lower field values. To begin, the position-correlated and position-uncorrelated rates at extraction voltages of 7, 8, and 9kV and a fixed drift voltage are shown in Figure 3.41. Rates here are normalized by  $e_S$  to correct for differences in extraction efficiency and charge yield at different fields.

A few qualities are immediately apparent in Figure 3.41. First, the photoionization rates are almost independent of the field because of the S2 normalization. Second, the power law appears to steepen and extinguish substantially at the highest field setting. It is unclear if this is due to significant contamination from gas events or a real effect. Because of this, data is only fit at these




Figure 3.41.: Rate of position-correlated and uncorrelated single electrons from extraction liquid progenitors, at three different extraction voltages. Rates are normalized by  $e_S$ , the size of the progenitor S2 at the liquid surface.

field values for posterity, and otherwise excluded from this analysis.

The background subtraction procedure (Section 3.5.5.2) is performed on the position-correlated rates to better reveal the power law at each field (Figure 3.42). For the 9kV dataset, the power law appears to be completely obscured by position-uncorrelated backgrounds just 10ms after the S2.




Figure 3.42.: S2-normalized rate flux of position-correlated single electrons following extraction liquid progenitors, at three different extraction voltages. The background flux from position-uncorrelated electrons has been subtracted, revealing a slightly clearer power law in each dataset.

Despite the large statistical uncertainties, the values of the power law exponents for the 7

and 8kV datasets are remarkably robust. Choosing a drift time upper bound of 2.0  $\mu\text{s}$  instead of 2.5  $\mu\text{s}$  to define the extraction region has no effect on the reported slopes. In contrast to this, an upper bound of 3.0  $\mu\text{s}$  results in significant contamination from drift liquid electron trains and a shallower exponent.

It is also worth mentioning that the position-uncorrelated background is slightly elevated, resulting in unnaturally lower rates in *both* the first and last bins in the fit region after background subtraction. This minimally affects the power law slope because the effect is symmetric about the fit region. Nevertheless, a closer inspection of the data revealed the presence of a weak grid emissions contributing to the position-uncorrelated rate. In the 7kV dataset, choosing a different definition of  $40 < \Delta r < 50\text{cm}$  for position-uncorrelated pulses results in a position-uncorrelated background rate which is reduced by a few percent at timescales beyond a few ms. After subtracting the position uncorrelated rates, the value of  $\beta$  becomes  $-1.32 \pm 0.08$ , which is within the uncertainties presented in Figure 3.42.

Finally, rather than performing a background subtraction procedure, the model can be adjusted to include a constant term (Equation 3.9). Fitting this to the position-correlated rate (Figure 3.43) results in a larger systematic uncertainty due to the concavity of the rate curve at timescales close to 1ms. As a result, the left edge of the fit region is set to 2ms, to avoid some of the residual photoionization background.

### 3.5.6.1. Drift Liquid Events

Power laws are fit to the position-correlated rates from drift liquid events in each data set (Figure 3.44). As the extraction field is varied, the lifetime of the photoionization component




Figure 3.43.: S2-normalized rate flux of position-correlated SEs following extraction liquid progenitors, without background subtraction. A power law is fit with the addition of a constant term as an alternative to account for the underlying position-uncorrelated background. The power law exponents are only marginally affected. The left edge of the fit is set to 2ms because of residual photoionization backgrounds at 1ms which are not well-modeled by the constant terms.

increases between the 7 and 8kV datasets. This occurs because the 7kV data was taken after SR1 with a higher electron lifetime, while the 8kV data is predominantly from before SR1. This shift in photoionization rates results in a steeper slope from between roughly 1-5ms in the 8 and 9kV datasets in comparison to the 7kV dataset.

Unfortunately, even after subtracting the position-uncorrelated sideband, the high intensity of the DRL position-correlated rates leaves the slope unaffected at all timescales from 1-1000ms. This makes a direct comparison between fits to DRL and EXL electron trains tenuous. Ideally, the the DRL and EXL rates would be compared over the same time interval (i.e. 1-20ms), however the presence of intense photoionization backgrounds in the DRL rates obscures the slope of the delayed electron backgrounds from roughly 1-5ms. At best, the DRL fit can only be extrapolated to 1ms from its actual fit region of 3-300ms. This extrapolation approach was taken for comparing between results from different papers in Table 3.7, i.e. the nominal values of  $\beta$  were assumed to hold over the entire duration of the electron train. In hindsight, the photoionization profiles of

each experiment and the corresponding fit ranges were overlooked. Future work should consider the differences in fit region for each experiment, as well as alternative methods to improve the photoionization background subtraction in the DRL electron trains. Possible workarounds include selecting specific drift time bands or smaller progenitors, but these come at a cost of even lower statistics.

The detector conditions from the EXL and DRL datasets were otherwise identical at each extraction voltage. If the extrapolated DRL power law holds, it can be concluded here that the lower drift field ( $\sim 190$  V/cm) and much larger active volume are dictating the exponent of the power law.




Figure 3.44.: S2-normalized rate flux of position-correlated single electrons following drift liquid progenitors, at three different extraction voltages. The position-uncorrelated background has been subtracted. The electron lifetime in the 8 and 9kV datasets was identical, and lower than the electron lifetime in the 7kV dataset. This resulted in a steeper slope from photoionization backgrounds between  $\sim 1$ -5ms. Because the position-correlated rates from the drift liquid are so intense, subtracting the position-uncorrelated background has no effect on the features in the rates at short (few ms) and long (few hundred ms) timescales.

It is worth mentioning here that this weak extraction field dependence from drift region electron trains agrees with results from XENON1T (Figure 3.45). This demonstrates again that trapped electrons at the liquid surface are not likely to be a component of position-correlated

electron trains, as it has been previously hypothesized [94]. However, the higher extraction field might contribute an increased position-uncorrelated background flux (seen in Figure 3.41) that manifests as a larger delayed electron fraction for both position-uncorrelated and position-correlated SEs. This hypothesis is examined in the next section with an analysis of drift liquid events at different drift fields.




Figure 3.45.: Extraction field dependence of the delayed electron fraction from drift region electron trains in XENON1T [80]. The delayed electron fraction was calculated by summing the number of electrons observed between 2-200ms after a progenitor S2 in the drift region, then normalizing by the size of the extracted S2.

### 3.5.7. Varying the Drift Field

When the extraction field is varied, the position-uncorrelated background from extraction liquid events at long timescales increases slightly with the extraction field (Figure 3.41). This

position-uncorrelated background has been previously proposed as being due to electron trains from other progenitors [80]. While the increase in this component may be an effect of a higher extraction field acting on trapped electrons at the liquid surface, the same effect is seen when the extraction field is held constant and the drift field is varied (Figure 3.46). This would imply that the position-uncorrelated component is also drift field-dependent and independent of the extraction efficiency. This behavior does not appear to be a consequence of normalization. Determining the cause of this feature is beyond the scope of this work. It is speculated here though that it may be a result of weak grid emission where the increase seen in Figure 3.41 is from gate emitters, and the increase in Figure 3.46 is from cathode emitters.




Figure 3.46.: S2-normalized rate flux of position-correlated and uncorrelated single electrons from drift liquid progenitors, at three different drift voltages.

The appearance of a differently behaved power law for extraction liquid events has implications for modeling drift liquid electron trains. Rather than simply subtracting off the position-uncorrelated background from drift liquid events, it seems more sensible to regard position-correlated extraction liquid electron trains as a background for drift liquid electron trains. Subtracting off the respective position-correlated extraction liquid flux for each dataset results in the remaining drift liquid rates shown in Figure 3.47.




Figure 3.47.: S2-normalized rate flux of position-correlated single electrons following drift liquid progenitors, at three different drift voltages. The background flux from position-correlated extraction liquid electrons at the same extraction field has been subtracted. The power laws only appear to be marginally affected.

The effect of subtracting off extraction liquid electron trains appears to only marginally affect the power law of drift liquid electron trains. At best, it can be concluded that the power law is not sensitive to small changes at lower drift fields.

### 3.5.8. Conclusions

This work presents a novel analysis of extraction and drift liquid events from the same datasets. Along the drift path, high electric fields seem to result in a steeper power law exponent, but a decisive dependence is obscured by a lack of statistics at high fields. This comparison is also weakened further by relying on an extrapolated power law in the drift liquid events. Future work could improve on uncertainties present at high fields by collecting additional data, and/or by addressing trigger efficiency corrections in SR1 and SR2 data to enable their use in a future study. Additional work is also needed to account for the different photoionization profiles following events, which are affected by the electron lifetime. While the background subtraction process appears to be effective for extraction liquid events, photoionization backgrounds in drift liquid events persist even

after the subtraction is performed, obscuring comparisons to the drift liquid power law between 1-5ms.

An optimistic perspective is that this could be the first time this dependence has been demonstrated, and a review of past studies shows that it may have been overlooked (Figure 3.48, Table 3.8). This work should inform future modeling efforts which should consider the possibility of distinct power law components from the extraction and drift liquid regions respectively. Another implication of a possible field dependence is that the electron train power law results from some mechanism in the liquid bulk, and not on the liquid surface. Finally, it should be noted that position-uncorrelated electron trains also appear to be influenced by the drift field, increasing in intensity slightly with the strength of the field. However, it is unclear if this apparent rate increase is due to an unaccounted for factor in the analysis or datasets.




Figure 3.48.: Drift field variation of the electron train power law exponent in past studies, compared to results from this work. The 9kV dataset is excluded due to a combination of low statistics and qualitative similarity to gas events. Data from this figure is reproduced in Table 3.8.

Table 3.8.: Table of values shown in Figure 3.48. A 15% systematic uncertainty has been added to the errors for data taken at the 3900 and 3410 kV/cm (7 and 8 kV) fields to compensate for uncertainties introduced in the choice of fit regions and cuts. Reported field values for this work were lowered slightly from those calculated in Table 3.4 to account for a lower potential on the gate grid as a result of deflection between the gate and anode grids.

	Field [V/cm]	beta	yerr	Work	Year
0	100	-1.20	0.04	Kopec et al.	2021
2	110	-1.05	0.02	This Work	2023
1	125	-1.10	NaN	XENON1T	2021
3	180	-1.05	0.05	LUX	2020
4	193	-1.13	0.01	This Work	2023
5	200	-1.20	0.04	Kopec et al.	2021
6	241	-1.14	0.01	This Work	2023
7	500	-1.20	0.04	Kopec et al.	2021
8	1000	-1.20	0.04	Kopec et al.	2021
9	3410	-1.37	0.16	This Work	2023
10	3750	-1.40	0.20	Akimov et al.	2016
11	3900	-1.29	0.22	This Work	2023

## Chapter 4.: Experimental Studies with a Compact Xenon

### TPC

This chapter will cover the design, modifications, and studies conducted with XeNeu<sup>1</sup>, a dual-phase xenon TPC at Lawrence Livermore National Laboratory (LLNL). Compared to LZ, XeNeu offers a flexible platform for studies of liquid xenon particle detection properties and TPC performance.

The work contained in this chapter encompasses two major upgrades to XeNeu which will enable future studies of low energy phenomena. The first upgrade involved replacing the largest plastic components in the detector with low-outgassing materials (Section 4.3), leading to a preliminary study of ionization backgrounds (Section 4.3.4). The second upgrade saw the replacement of the 4 top PMTs (Figure 4.6b) with a more compact array of 9 silicon photomultipliers (SiPMs) (Figure 4.38) which improved XeNeu's position resolution. Minor contributions were made for two other studies, which are summarized in Section 4.2. The first study was a calibration of the highest energy nuclear recoils to date in liquid xenon [106], while the second was a search for electron recoils due to the Migdal effect accompanying nuclear recoils from neutrons [107].

---

<sup>1</sup>Short for Xenon/Neutron [105]). Definitely no relation to [the other Xenu](#).

## 4.1. The XeNeu Detector

This section will detail most of the subsystems and subassemblies in XeNeu, many of which have been modified since they were first built [105]. XeNeu was initially built for the purpose of measuring nuclear recoils from neutrons either at LLNL or specialized neutron sources across the country. To emphasize portability, XeNeu and additional circulation hardware are mounted on independent wheeled carts (Figure 4.1). Aside from electrical connections, the carts are connected by two flexible braided stainless steel hoses to convey xenon between the detector and circulation system. The TPC is compact, to minimize the possibility of multiple neutron scatters during calibrations. It was also designed to reach high voltages (Section 4.1.2.1) in order to maximize electron extraction efficiency and S2 gain.

### 4.1.1. Central Detector Components

The detector's central components are shown in Figure 4.2. There are three major subassemblies: the bottom PMT holder, field cage, and top PMT Assembly. All components are suspended from the detector top flange via three support rods which are connected to a mounting bracket (Figure 4.44) that bolts to the underside of the flange.

Figure 4.3 shows a cross-sectional schematic of XeNeu. The active region between the gate and cathode in the field cage (Figure 4.4a) contains ~100 g of liquid xenon. The liquid level in XeNeu is determined using an adjustable weir (Figure 4.5b). Light is detected with a PMT housed in the bottom PMT holder, and an array of photosensors at the top provides ~1 mm position resolution in the horizontal plane. Several components in XeNeu are configurable and the subject of this work. The PMT holder has a PTFE variant (Figure 4.5a) as well as one which is constructed from low-




Figure 4.1.: XeNeu (left) and its accompanying circulation cart (right). The detector is housed in a cylindrical can which hangs underneath the main outer vacuum insulation system. The circulation cart holds the gas circulation pump, getter, and xenon gas bottles. The gas bottles and getter are on the backside of the cart and not visible in this picture. Note that this picture is not up to date with the P&IDs shown in Figure 4.13 and Figure 4.16.




Figure 4.2.: XeNeu detector components. At the bottom of the detector assembly is a large cylindrical PTFE piece which houses the bottom PMT (Figure 4.5a). The field cage (Figure 4.4a) is mounted to the top of the bottom PMT holder. The weir reservoir is machined from a single block of PEEK and sits in a dedicated cutout on the side of the bottom PMT holder. The weir height is independently adjustable with clamps on two of the vertical support rods. Directly above the field cage is the top PMT assembly (Figure 4.6), which clamps to all three support rods. High voltage feedthroughs are welded to the detector top flange, and descend into slots machined into the top of the bottom PMT holder.

outgassing materials (Figure 4.23). The top photosensor array originally used PMTs (Figure 4.6), which were replaced with a SiPM array (Figure 4.38). The field cage (Figure 4.4) can also be assembled in two configurations, either with a PTFE reflector to improve S1 light collection, or without one to reduce outgassing rates.

#### 4.1.1.1. Field Cage

The field cage (Figure 4.4a) can be assembled in two configurations, either with or without a polished PTFE reflector (Figure 4.4b). The grids are hexagonal acid-etched stainless steel mesh, and are press-fit into assemblies consisting of two concentric rings. The gate grid ring assembly is made from stainless steel instead of copper to make it slightly more resistant to scratches that could compromise its integrity under high voltage. Specially made PEEK spacers with ridges separate the copper field shaping rings; the ridges are intended to hamper the migration of charge between the field shaping rings.

#### 4.1.1.2. Bottom PMT Holder

The bottom PMT holder (Figure 4.5a) consists of a monolithic 600 gram piece of PTFE. It acts as an electrical insulator between the high voltage electrodes, PMT, and field cage. It also displaces space in the vacuum chamber that would otherwise need to be filled with liquid xenon. The central cavity of the bottom PMT holder houses a 2 in Hamamatsu R8778 PMT. This PMT collects the majority of S1 light due to total internal reflection of light at the liquid surface. The PMT is affixed to a PTFE plate with a PEEK retaining ring and nylon #4-40 screws (Figure 4.5b), which is then inserted as one assembly into the bottom PMT holder. The PTFE plate is held in place with vented 1/4"-20 nylon screws. The bottom PMT shield grid rests on a lip carved into the





Figure 4.3.: 2D schematic (drawn to scale) of XeNeu showing dimensions of the active volume and labels of key components. Several components are configurable and the subject of this work. The PMT holder has a PTFE variant (Figure 4.5a, pictured here) as well as one which is constructed from low-outgassing materials (Figure 4.23). The top photosensor array originally used PMTs (Figure 4.6), which were replaced with a SiPM array (Figure 4.38, pictured here). The field cage (Figure 4.4) can also be assembled in two configurations, either with a PTFE reflector (pictured here) to improve S1 light collection, or without one to reduce outgassing rates.



- (a) The XeNeu TPC Field Cage. Ridged PEEK spacers separate the grids and field shaping rings. The gate grid is press-fit into the top ring, and the cathode is press-fit into the bottom ring.
- (b) The PTFE reflector is segmented so the field cage can be built around it.

Figure 4.4.: Field cage components for the XeNeu TPC. The field cage is held together by three nylon screws cut to the appropriate length such that they fully engage with tapped holes in the base ring, which is made of PEEK.


top of the central cavity of the PMT holder. The shield grid is grounded via an external wire. An additional PEEK spacer sits on top of the shield grid, slightly recessed from the top of the PMT holder, in order to locate the field cage.

The bottom PMT holder has additional cutouts for the high voltage electrodes, the weir reservoir, and the support rods. The weir reservoir was machined from a single piece of PEEK and contains its own set of clamps so it can be positioned independently along the support rods. This determines the liquid level of the TPC. The weir reservoir was designed to minimally interfere with incoming radiation from external sources. The wall thickness is only 2 mm (note this is not properly reflected in the CAD model), and its depth is such that the liquid level inside is beneath the active volume of the TPC. The liquid level in the weir is determined with three PT100 RTDs which are situated at different positions along the length of a plastic tube. The plastic tube is then inserted into the weir reservoir, and held in place with a small clamp on one of the support rods. As the xenon gas is slightly warmer than the liquid, the topmost RTD will read a few Kelvin warmer than the bottom two when the latter are fully submerged.

Replacing the PTFE PMT holder with low-outgassing materials is the subject of Section 4.3.

#### **4.1.1.3. Top PMT Assembly**

The top PMT assembly is shown in Figure 4.6. The assembly houses four, 1" Hamamatsu R8520 PMTs arranged in a square grid. For single electrons (SEs), the PMT array has a position resolution of  $\sim 1$  mm. The PMT assembly can be positioned along the support rods, allowing for the gas gap to be set independently of the liquid level. The PMT box, which houses the PMTs, can slide out laterally from the mounting bracket Figure 4.6b. However, cables which are soldered



- (a) Cutaway of the bottom PMT holder, showing the slots for the high voltage electrodes. At the upper far left a cylindrical PEEK nut is visible. Three of these nuts are used to level the TPC on the support rods. The hex nuts atop the PMT holder are used to fix the assembly in place on the support rods. Three additional screw holes (two of which are visible) atop the PMT holder sit close to the inner diameter. These holes are tapped so screws and washers can be used to hold the field cage in place.
- (b) Exploded view of components in the bottom PMT stack, as well as a cutaway view of the weir reservoir. From top to bottom: PEEK spacer ring, PMT shield grid, bottom PMT, PEEK retaining ring, PTFE mounting plate. Note that there is a circumferential ridge at the bottom of the PMT which is not pictured here, similar the one near the top of the PMT. The PEEK retaining ring clamps onto this ridge, which fully constrains the bottom PMT.


Figure 4.5.: Two views showing key features of the bottom PMT holder assembly.

to the PMT bases form a second constraint which prevent the PMT assembly from being easily detached from the detector. This concept of a modular bracket system is nevertheless extremely useful as it avoids having to remove the entire bottom PMT assembly in order to remove the top PMT assembly. Replacing this assembly with an array of SiPMs is the subject of Section 4.4.


#### 4.1.2. Detector Top Flange

A heavily customized 6" CF flange contains all of the ports required to interface with the detector. This flange is housed in a vacuum space, and is suspended from another flange which attaches to the outer vacuum box (Figure 4.12). All slow control and photosensor electrical connections are routed through four 1.33" CF ports Figure 4.7. 4-wire RTD measurements are made through two PEEK 9-pin Sub-C connectors from AccuGlass. High voltage for PMTs is delivered via 7-pin Mil-spec feedthrough (Figure 4.40) purchasable from MDC Vacuum or [MPF Products, Incorporated](#). Signals from photosensors are sent through an either 9-pin or 23-pin Sub-C (Figure 4.43b) connector on the forth 1.33" port. The creation of the 23-pin connector with 10 coaxial channels is detailed in Section 4.4.3.

The fifth CF port hosts a VCR-4 connection which acts as a return line from the condenser and delivers liquid to the bottom of the detector via a PTFE tube (not shown). Two more CF ports are present and accessible from outside of the vacuum box. The 1.33" CF port is the main pumpout port for the detector space, while the central 2.125" CF port is reserved for use with optical equipment. Lastly, a 1/4" tube is welded directly to the top flange which serves two purposes via a compression tee. One one branch of the tee is converted to a 1/16" capillary, which is connected to a PTFE tube (not shown) in the detector space that is submerged in the weir liquid. Liquid



(a) Fully collapsed view of the top PMT assembly. Not shown is a layer of copper foil which is affixed to the top of the PMT box by the four vertical screws and intended to reduce switching noise from slow control sensors in the TPC. An additional screw for attaching ground wires is visible at the top left. Note that all tapped holes have additional 1/64" through holes for venting otherwise blind volumes.



(b) Exploded view of the top PMT assembly. A bracket is mounted to the support posts, which mates with a plate that is fixed to the PMT box. The PMT box contains PTFE spacers to fix the PMTs in place and insulate them from any electrically conductive surface. The PMTs and PTFE spacers ultimately rest on top of a quartz window which acts as a shield in case of any discharge from the anode. The quartz window is supported by a lip on the bottom of the PMT box. Attached to the bottom of the PMT box is a PEEK spacer, followed by the anode itself. While the anode is usually grounded via a wire (not pictured), the spacer offers the option of biasing the anode without affecting the PMT box.

Figure 4.6.: The XeNeu top PMT assembly.




Figure 4.7.: Ports on the top flange of the detector. Clockwise from the top of the detector flange: PMT high voltage feedthrough (see also Figure 4.40), PMT signal feedthrough, liquid outlet port, gate HV feedthrough, and liquid inlet port. Continuing clockwise: pumpout port, RTD feedthrough, cathode HV feedthrough tube, and another RTD feedthrough. The gate and cathode feedthroughs, as well as the pumpout and central ports are accessible from the vacuum box mounting flange. This flange and the central port on the detector top flange are transparent for visibility of other components.

is drawn up into the capillary which exits the detector space into the two-phase heat exchanger. The second branch of the tee connects to the gas phase of the condenser, establishing a thermal connection that acts as a simple dual-phase counter-flow thermosyphon to the top of the detector space.




Figure 4.8.: Cutaway of tee showing capillary and thermosyphon paths. The capillary is nested inside the tee and protrudes into the detector space through the 1/4" tube which is welded to the detector top flange. The capillary is used to draw liquid from the weir into the heat exchanger. The second branch of the tee (bottom) is used to volume share the detector gas with the condenser, forming a crude thermosyphon.

#### 4.1.2.1. Grid High Voltage

The system which delivers high voltage to the grids was originally designed and built by Sergey Pereverzev (LLNL) and is capable of biasing the gate and anode to a potential difference

of more than 20 kV [71]. The inner conductors for the gate and cathode pass through 1/4" (outer diameter) stainless steel tubes which are welded in place on the detector top flange. The inner conductors are insulated with teflon, and the remaining volume inside the tubes is filled with silicone oil for further insulation. The inner conductors terminate approximately 12" below the top flange and are soldered to the ends of the ceramic feedthroughs. The ceramic feedthroughs are then soldered to the ends of the tubes (Figure 4.9), isolating the oil from the detector space.

Each feedthrough is biased with a Bertan Series 230 high voltage power supply. The power supplies are connected to boxes mounted on top of the vacuum vessel. These boxes contain circuitry for monitoring the high voltage connections to preempt an electrical breakdown and act as oil reservoirs for the feedthroughs. A diagram of the monitoring circuit is shown in Figure 4.10. Because an electrical breakdown is a transient process, the monitoring circuit consists of a high-pass filter so transient signals can be observed without directly coupling to the high voltage output. Prior to a large breakdown, smaller breakdowns may occur. These can be observed on an oscilloscope as spikes in the voltage at the output of the high pass filter, indicating that the voltage ramping process may be proceeding to quickly.

The field cage resistor chain (Figure 4.11) was designed to deliver the highest voltages possible while still maintaining current-limiting resistors upstream of each field-shaping ring. In a completely naive design, high voltage would be directly applied to the gate and cathode grids. This is dangerous because in the event of a short from either the gate or cathode to ground, current is unrestricted. A safer design places current-limiting resistors before the gate and after the cathode so that the flow of current is restricted in the event of a short. However, the current-limiting resistors will have a voltage drop during normal operation where a small DC current is expected. In order to




Figure 4.9.: XeNeu detector with HV feedthroughs made visible. The HV feedthrough tubes are mated to the ceramic feedthroughs via a copper reducer. End caps are attached to the gate (left) and cathode (right) feedthroughs which contain sockets for running wires to the gate and cathode. The brass end cap belongs to the cathode feedthrough, and the copper end cap belongs to the gate feedthrough. The caps are different because of an incident where a solder joint on the cathode feedthrough failed, which leaked silicone oil over the TPC components. Ethan Bernard (LLNL) performed the repair and created the replacement brass cap, which is fixed in place with a set screw.




Figure 4.10.: Schematic for high voltage delivery and monitoring circuit.  $V_{IN}$  is the connection point for the high voltage power supply, Precursors to a full electrical breakdown are monitored with an oscilloscope at the output of a high-pass filter,  $V_{MON}$ .  $V_{MON}$  and the high voltage output  $V_{OUT}$  are fed out of the box through coaxial connections.  $R_{CL}$  is a current-limiting resistor, which protects components in the event of an electric short.

achieve the same voltages at the gate and cathode in the unsafe design, the input voltages must be significantly higher. For example in the XeNeu field cage, the safer design will have a voltage difference between the gate and cathode which is roughly 70% of the of the unsafe design, assuming equal resistors and identical input voltages.

The XeNeu resistor chain uses a unique solution which recovers the full voltage drop of the unsafe design, and the current-limiting characteristics of the safer design. The grids are independently connected to the high voltage electrodes with a  $2\text{ G}\Omega$  current-limiting resistor. But because the grids are not connected to any other component, the current-limiting resistor has no voltage drop across it. In other words, with no current from the open circuit, the resistor connected to the grid is effectively invisible, and the grid is at the same potential as the high voltage electrode. In parallel to this, a second  $2\text{ G}\Omega$  resistor connects the electrode to the adjacent field-shaping ring. The resistor chain continues to the opposing electrode. In this way, the grids are kept at the same voltage as their adjacent electrodes, while still being downstream of a current-limiting resistor. The voltage drops between all components are identical to those in the unsafe design described above.

### **4.1.3. Detector Cart**

Figure 4.12 shows a rendering of the detector cart with portions of the vacuum vessel removed. The detector mates to the vacuum vessel via the vacuum box mounting flange, which is welded to the detector pumpout and optical ports (Figure 4.7).

Figure 4.13 shows the layout of different components on the detector cart. The detector is instrumented with PT100 RTDs for monitoring key temperatures, and resistive heaters for balancing heat loads. Pressures are measured using Setra 225 pressure transducers and analog dial indicators.

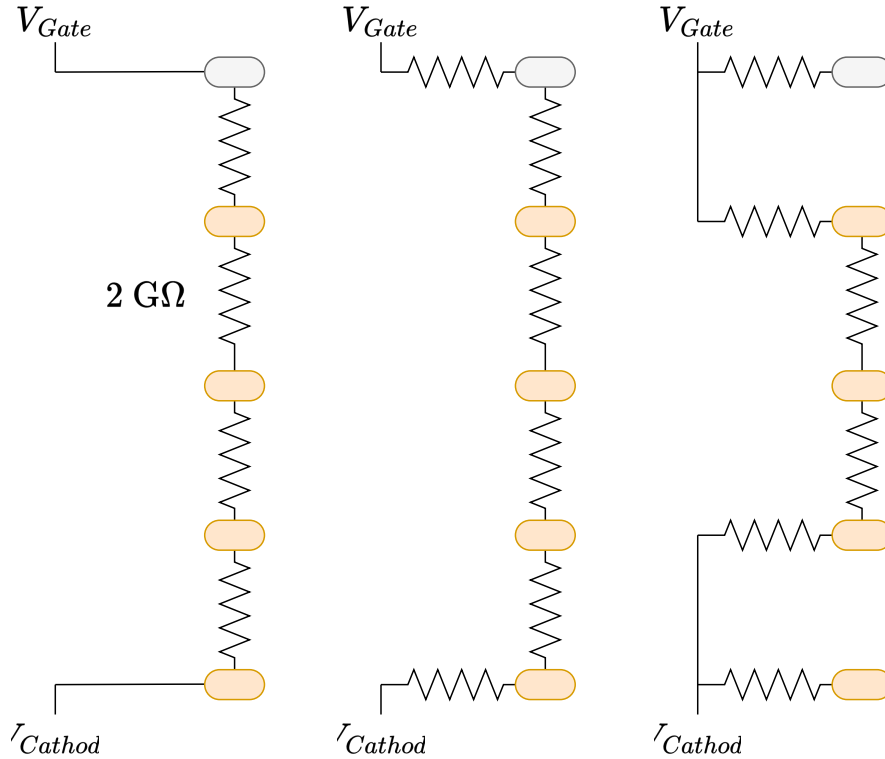


Figure 4.11.: Schematic showing three different designs for a high voltage resistor chain on a field cage. At the left is a naive design that is unsafe because a short from the gate or cathode grid would have an unrestricted current. In the middle is a safer design that installs current-limiting resistors before the gate and cathode. A tradeoff of this design is a reduced voltage drop between the gate and cathode, requiring higher input voltages to match the unsafe design. At the right is the design of the resistor chain in XeNeu. Current-limiting resistors are in series with gate and cathode grids, however the grids are disconnected from other components. Consequently, the gate and cathode voltages are equal to the voltage of their respective electrodes. All resistors are  $2\text{ G}\Omega$  and the resistors between rings plug directly into holes machined into the outer diameter of the rings. Resistors adjacent to the gate and cathode plug into ports on the caps of the respective feedthroughs.




Figure 4.12.: Rendering of XeNeu with portions of the vacuum box and the outer can made transparent. Electrical connections are fed out through the flange on the far right side, which has 5 SHV, 5 BNC, and 5 DB15 connectors. Directly above this flange, the condenser and heat exchanger (Figure 4.15) are mounted in the volume contained by an 11 in ISO-K nipple.

Additional information on each sensor shown is given in Table B.1.

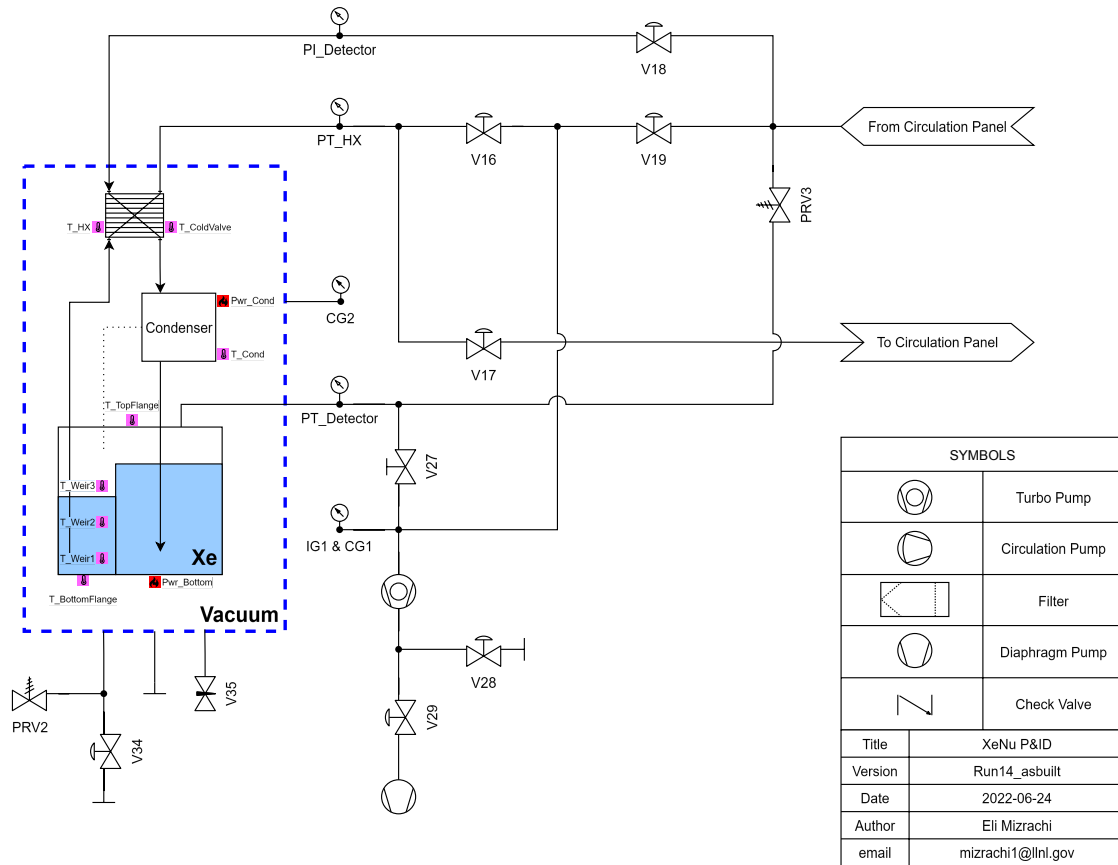


Figure 4.13.: P&ID for components on the XeNu detector cart. Arrows on the heat exchanger lines denote the normal circulation path. The blue dotted line demarcates the vacuum vessel. The dotted line from the condenser to the detector represents a connection which joins the detector and condenser gas spaces (Figure 4.8), acting as a simple thermosyphon. PT: pressure transducer, PI: pressure (dial) indicator.

To minimize heat loads, the inner detector can is surrounded by a vacuum vessel, indicated by a dotted blue line in Figure 4.13. This vessel is instrumented with a Hornet CVG101 convection gauge (CG2). During cooldowns, the vacuum vessel is continuously pumped on through a KF25 port by a Pfeiffer HiCube 80, which has a Hornet IGM402 ion gauge for reading out pressures below

the range of IG2. The vessel has two additional valves, V35 and V34, which are used for venting and filling the box with nitrogen respectively. In either case, increasing the pressure in the vessel is useful for increasing the heat load to the detector during the end of a warmup. A 5 psi pressure relief valve (PRV2) ensures that the vessel cannot be over-pressurized by nitrogen.

The vacuum vessel has feedthroughs which connect to the detector space, as well as the heat exchanger (Figure 4.14). The detector space is pumped on through V27 using a Pfeiffer TMP071 turbopump, backed by a Pfeiffer MVP-020 diaphragm pump. PRV3 is teed off of the detector pumpout port to allow a vent path for xenon that bypasses the condenser and heat exchanger. Normal circulation utilizes a gas feedthrough on the vacuum vessel which connects to V18 and V17. These are the gas inlet and outlet respectively. V16 and V19 are typically used for bypassing the heat exchanger and condenser when pumping down the detector space and plumbing, or recovering xenon at the end of a cooldown.

The design of the vacuum vessel is fairly unique in that it relies on a rectangular box to offset an 11" ISO-K flange which contains a two-phase heat exchanger and condenser. The intent behind this design was to leave vertical space above the detector free for optical studies of liquid xenon. Assembling the box was a challenging process which predated this work; additional information on some of the design's shortcomings are detailed in Section B.2.1.

Figure 4.15 details the components housed in the 11" ISO-K flange. The heat exchanger (HX) is a brazed plate heat exchanger from Kelvion<sup>2</sup>. Approximately 25W of cooling power are provided by an Advanced Research Systems ORCA MR-90 cryocooler coupled to the condenser. The condenser is a custom-made hollow block of copper consisting of two halves brazed together.

---

<sup>2</sup>This was noted with a model number of FR3x8-15 in [105] however the model numbers appear to have changed and the closest match seems to be from the GNS series of brazed plate heat exchangers.




Figure 4.14.: Rear view of XeNeu cart. The condenser and heat exchanger (Figure 4.15) are housed in an 11" ISO-K nipple, offset from the axis of the detector. The flange on this nipple has connections to the cryocooler as well as gas inlet and outlet connections to the heat exchanger. In the foreground, oil filled plastic boxes are mounted on top of the detector and connect to the high voltage feedthroughs.

The condenser has three ports: one port receives pre-cooled gas from the heat exchanger, one port delivers liquid to the TPC, and one port is connected to the gas space of the TPC (Figure 4.8). This gas connection acts as a passive thermosyphon which is capable of delivering cooling power to the detector independently of the circulation process.

#### 4.1.4. Circulation Cart

Circulation is important for removing impurities released by outgassing. Impurities are undesirable for multiple reasons. First, electronegative impurities reduce the lifetime of free electrons in the liquid, thereby attenuating S2 signals. These impurities also relate to the formation of electron trains, discussed extensively in Section 3.4 and Section 3.5. Finally, other impurities are known to ionize in the presence of VUV light from xenon (Section D.10.1), forming a prompt source of background electrons. Circulation is achieved by drawing liquid up from the weir reservoir and into the heat exchanger. Inside the heat exchanger, the liquid is evaporated by incoming warm gas, and pumped through the getter where it is purified. Gas returns into the heat exchanger, where it is cooled and condensed by outgoing liquid before entering the condenser.

Circulation rates of up to 1200-1500 sccm are maintained using an MKS G50A mass flow controller (MFC) and Senior Metal Bellows MB-601 pump, giving a roughly 2 hour turnover time for 1.5 kg of xenon. The pump is hermetically sealed<sup>3</sup> and has a minimum helium leak rate below 1E-10 mbar/L \*s. Trace impurities in the liquid are continuously removed by circulating through an Entegris/SAES MonoTorr PS3-MT3-R-1 getter. Gas is stored in bottles mounted on the cart. Space is left below the bottles so they can be used with open dewars to cryopump xenon during recovery.

---

<sup>3</sup>The manufacturer does not specify how the pump heads are sealed, but claims all wetted surfaces except internal valve seats are 300 Series stainless steel. Viton or Teflon valve seats are standard, with aluminum valve seats offered as an option.




Figure 4.15.: Hardware used for the thermal management in XeNu is housed in an 11" ISO-K flange offset from the axis of the detector. Pre-cooled gas exits the heat exchanger and enters the condenser. The gas is then condensed, where it enters the detector space as a liquid. Liquid exits the weir reservoir via a capillary line, which enters the heat exchanger. The condenser also has a third port which couples to the detector gas space. The gas exchange line is capable of cooling the detector volume if circulation is halted.

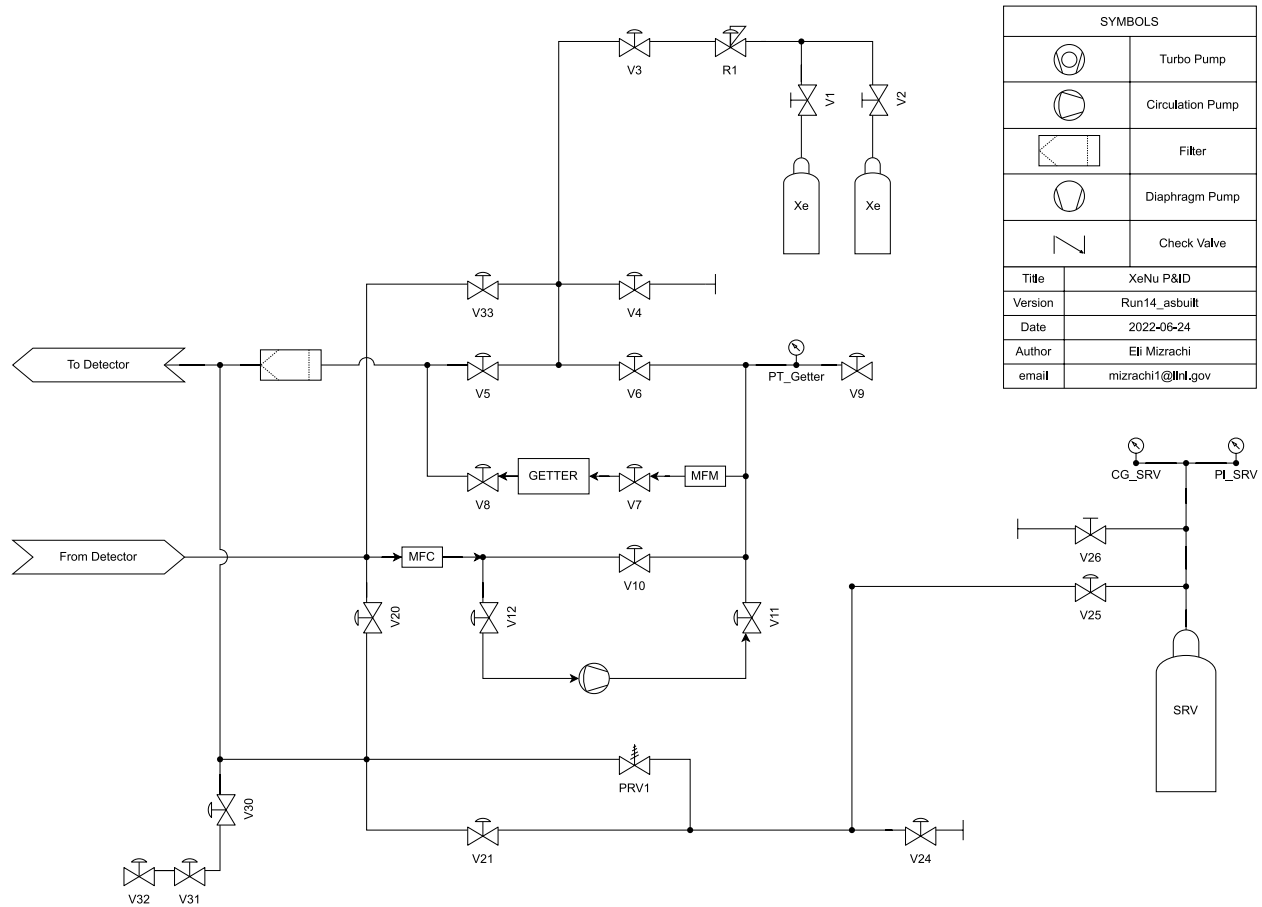


Figure 4.16.: P&ID for the circulation cart. The SRV is not mounted on the cart, but included here for clarity. V32 and V31 are intentionally dead ends which lead to a volume that previously contained an Ar37 calibration source.

During condensing, gas is fed from one of the two bottles through a regulator (R1). A major quality of life upgrade was the inclusion of V33 which routes gas through the MFC during condensation, allowing for a significantly less intensive condensation procedure. The previous procedure required constant manual tuning of V3 and V6 for multiple hours. The inclusion of V33 was prompted by the realization that xenon recovery from the detector could also be significantly automated with the MFC, which also saved multiple person-hours of attention.

In the event of a pressure excursion, a relief path exists from the detector, through PRV3, PRV1, and into the storage and recovery volume (SRV). V25 is locked open during detector operations, and kept closed otherwise to prevent atmospheric ingress into the SRV.

#### **4.1.5. Data Acquisition and Slow Control**

Data from photosensors is acquired at 250 MS/s using a Struck SIS3316 digitizer. The data acquisition system relies on the `daqman` library developed by Ben Loer [108] (PNNL), and modified by Jingke Xu (LLNL). `daqman` offers a wide variety of utilities for collecting, viewing, and processing waveform data from photosensors. During data acquisition, `daqman` interfaces with the digitizer to acquire waveforms which are written out for later analysis. Plotting features are also available for monitoring data quality. These waveforms can be viewed later, and optionally pre-processed to test configuration parameters for baseline detection, pulse finding, and other parameters. Finally, ROOT TTrees can be written out containing events with custom data classes. For unknown reasons, these data classes cannot be read out using `uproot`. As a result, a header-only file containing only the `daqman` data classes was developed [109] so the `daqman` TTrees could be read using ROOT6 to produce `uproot`-readable skimmed files (Section 4.3.4).

A custom Python slow control script written by James Morad (CashApp, formerly LLNL) and modified by Teal Pershing (LLNL) maintains the cryogenic stability of XeNeu. The script interfaces with a Keithley DAQ 6510 to obtain readings from instruments like power supplies and pressure transducers. Programmable instruments like the heater power supply and MFC are connected directly to the slow control computer with USB to RS232 and USB to RS485 adapters respectively. A PID loop in the program controls `Pwr_Cond` (Table B.1), which is used to reach temperature or pressure setpoints that are configured by the user. The system is exceptionally stable, generally drifting by only hundredths of a kelvin during normal operation.

## 4.2. Physics Measurements with XeNeu

XeNeu was used for two other experiments which were concurrent with this work. The first was a calibration of nuclear recoils in liquid xenon at energies up to 426 keV, the highest ever measured [106]. The second was a search for the Migdal effect, which is believed to be a probe for low-mass dark matter interactions in TPCs. The search produced a null result [107], informing future efforts to characterize this signal. This section will briefly summarize each experiment, and highlight contributions made to each in this work.

### 4.2.1. Measurement of High Energy Nuclear Recoils

Dark matter searches for nuclear recoils up to hundreds of keV have been motivated by Effective Field Theory (EFT) models [110,111]. These models have been difficult to test though, because of a lack of calibration data beyond roughly 100 keV [112]. Above this energy, the electric field dependence of the xenon NR response had also not been investigated, hampering detector

modeling efforts.

XeNeu was used to perform a calibration of the xenon NR response with 14.1 MeV neutrons generated by a deuterium-tritium (DT) neutron source. Because the scattering angle is a function of the recoil energy, seven liquid scintillator (LS) backing detectors were arranged around XeNeu to capture nuclear recoils at energies ranging from 39 to 429 keV. A diagram of the experimental setup is shown in Figure 4.17. Aside from general maintenance tasks, this work contributed designs for mounts to constrain the backing detectors, and detailed measurements of XeNeu's components from CAD models which were used in simulations.

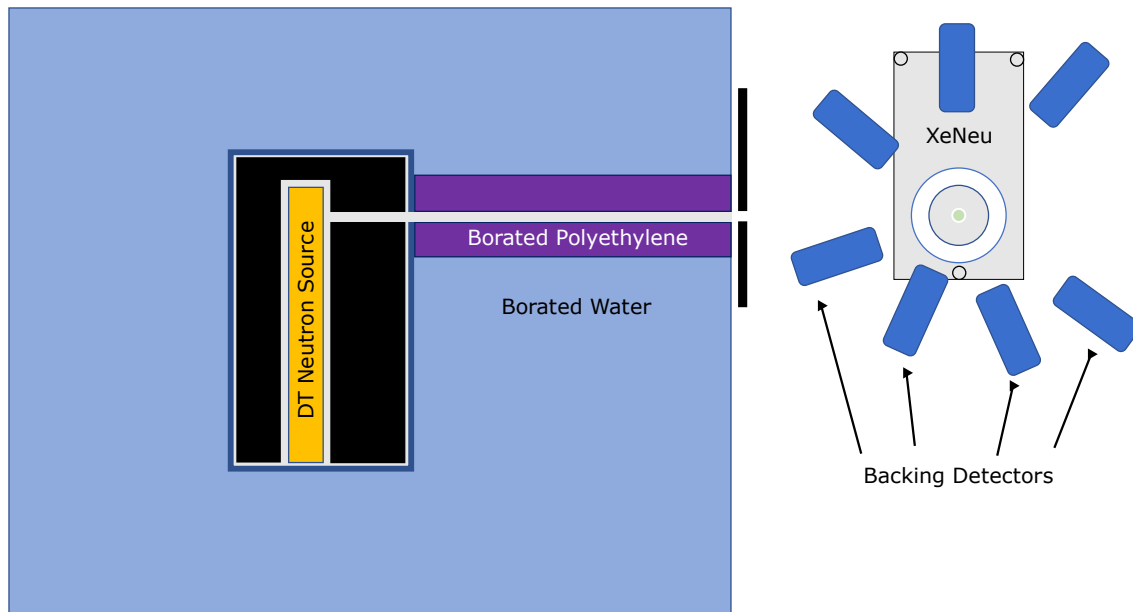


Figure 4.17.: Drawing showing top view of shielding and background detectors arranged around XeNeu for measuring high energy nuclear recoils. A mockup of the XeNeu detector cart is visible at the right, surrounded by backing detectors (blue rectangles). Backing detectors at high scatter angles were positioned closer to XeNeu to make up for the decreased elastic scattering cross section of 14.1 MeV neutrons at those angles.

Data was collected at three different drift fields:  $200 \pm 45$  V/cm,  $700 \pm 56$  V/cm, and 2000

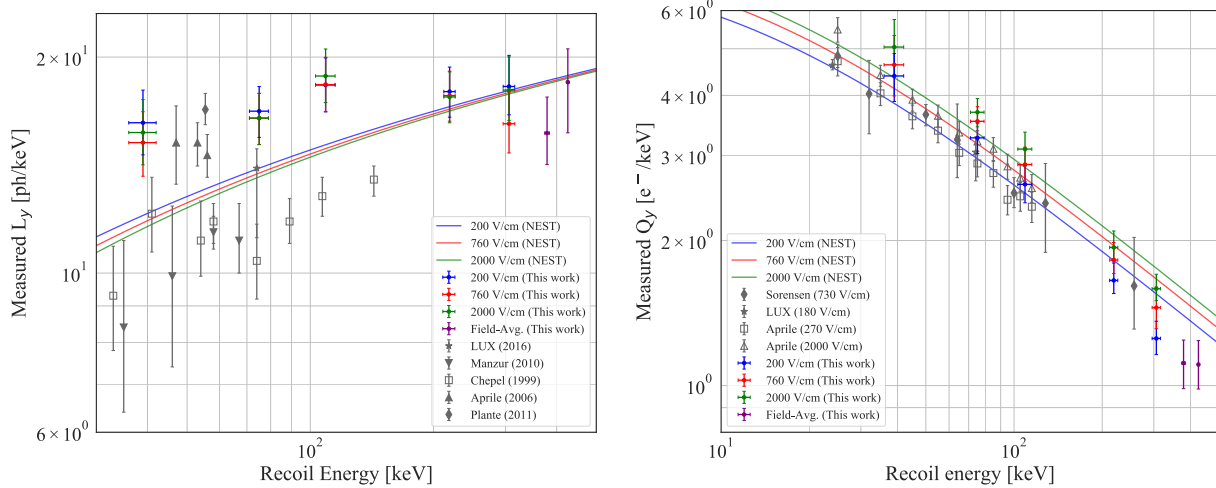
$\pm 81$  V/cm. The main event trigger relied on a 3-fold coincidence between the top PMTs, and any number of coincident LS signals. To avoid accidental triggers, the trigger threshold for the PMT signals was set above the level of typical SE pulses, and the LS thresholds were within 10% of the maximum energy that could be deposited by a neutron.

Neutron interactions were identified in the analysis by accepting only single scatter nuclear recoils within a drift time and x-y position consistent with the active volume. Interactions were rejected if more than one LS signal was present within a 30  $\mu$ s coincidence window, or if the LS pulse shape was inconsistent with that of a neutron interaction. Finally, interactions were required to have a valid time-of-flight (TOF) between XeNeu and the LS signals. At each field setting, drift time cuts were applied to avoid selecting events which were too close to the gate or cathode. No radial cut was used as the S1 and S2 sizes were not found to vary spatially. This uniformity was attributed to fiducialization from the built-in PTFE reflector (Figure 4.4b).

The S1 and S2 spectra from candidate events were modeled by simulating components from the experiment in BACCARAT, and then neutron interactions using Geant4. Simulation data was compared to measured data using a log-likelihood fit to estimate light yield ( $L_y$ ) and charge yield ( $Q_y$ ) values at each recoil energy. The resulting values are shown in Figure 4.18.

#### 4.2.2. Search for the Migdal Effect

The Migdal effect was first proposed in 1941 [113]. It predicts that an atom can be excited or ionized when its nucleus is displaced relative to its electron shells (Figure 4.19). This implies that some nuclear recoils could have a coincident electron recoil. For a very low energy nuclear recoil, this coincident electron recoil may raise the signal above the detector's low energy threshold. As



(a) Light yield values as a function of recoil energy. (b) Charge yield values as a function of recoil energy.

Figure 4.18.: Results from the XeNeu recoil energy measurement compared to previous values. Values at the two highest recoil energies were averaged over their electric fields due to low statistics. From [106].

a result, multiple experiments have proposed that the Migdal effect could boost their sensitivity to low-mass WIMP interactions [114,115]. However, because the Migdal effect has never been observed following a nuclear recoil [116], there is uncertainty surrounding these predictions. Thus, an observation of the Migdal effect would substantially improve the sensitivity of low-mass dark matter searches. This has motivated multiple proposals to detect it [116,117], as well as the search performed with XeNeu [107] that is outlined in this section.

Figure 4.20 shows information pertinent to the experimental setup used by XeNeu to search for the Migdal effect from a neutron scatter. The Migdal cross-section (Figure 4.20a) for neutrons peaks at a scattering angle of  $16^\circ$ . The ratio between the Migdal cross section and the elastic scattering cross section increases at higher scattering angles, but the absolute signal rate is lower at these energies and it is accompanied by higher backgrounds. Figure 4.20b shows the expected ER recoil spectrum for Migdal interactions in xenon, which is mostly independent of scattering





Figure 4.19.: Illustration of the Migdal effect. In a nuclear recoil with a dark matter particle, the nucleus is displaced from its electron shells. As a result, some of these electrons may escape, producing a coincident electron recoil. From [118].

angle and incident neutron energy. M and L-shell Migdal ionizations were considered the primary signal for this work as the optimal choice which balanced the expected detection thresholds and scattering probability.

A circular array of 14 LS backing detectors was placed behind XeNeu, and centered on the axis of a beam of 14.1 MeV DT neutrons. The array was designed to capture neutrons with scattering angles of  $15.0 \pm 1.5^\circ$ . Elastic neutron scatters were expected to deposit  $7.0 \pm 1.6$  keV in the TPC, producing on average 50 scintillation photons and 50 electrons. Of these, approximately 1.35% of them were expected to include Migdal ER signals. The corresponding M and L-shell ionizations were expected to produce 20 or 200 scintillation photons and 50 or 175 electrons, respectively. This would allow them to be distinguished from NRs in S1/S2 space.

Neutron single scatters were tagged in a manner similar to the high energy NR measurement. From simulations, pure single scatters, multiple scatters in the xenon, as well as “passive scatters” in detector components were the major background components. The signal model was generated by sampling simulated S1 and S2 signals from Figure 4.20b adding them to a corresponding simulated



(a) Migdal (black dash) and elastic (solid red) scattering cross sections for neutrons as a function of scattering angle. The grey band highlights the  $16^\circ$  scattering angle, including  $1\text{-}\sigma$  uncertainties for the range of scattering angles used in the experiment.

(b) Migdal ER spectrum showing the approximate recoil energies for M and L-shell ionizations in Xe. The solid black curve and dashed blue curve are generated from first-principle inputs and photoabsorption data respectively.

Figure 4.20.: Scattering cross sections as a function of scattering angle, and expected Migdal ER spectrum. From [107].

NR, and then scaling the result by the Migdal/elastic scatter ratio from Figure 4.20a.




Figure 4.21.: S2 spectrum from interaction at a scattering angle of 16°. Three background components are shown in the spectrum: single scatters, passive scatters, and multiple scatters. These are overlaid with predicted counts of Migdal interactions as well as the measured counts from the best-fit model. Fit residuals are shown in the bottom plot for the best-fit model (black points) and the predicted model (dotted line). The shaded grey band represents the systematic uncertainty in the NR background models.

A two-dimensional maximum likelihood fit in S1/S2 space was used to obtain the best-fit Migdal signal rate. The resulting S2 spectrum is shown in Figure 4.21. The best-fit signal rate was found to be  $16.3_{-16.3}^{+21.7}$  counts, almost an order of magnitude lower than the predicted  $148.2 \pm 16.3$  counts, and consistent with zero counts. Above an S2 size of 100 electrons, all three background distributions were below the predicted Migdal signal rate. This could occur when ER and NR tracks in the liquid xenon overlap, leading to enhanced electron-ion recombination. The resulting suppression of ionization signals and increase in scintillation signals would reduce

detection efficiencies from  $\sim 100\%$  (for ionization signals) to  $\sim 10\%$  (for scintillation signals), leading to a null result. Because this possibility could not be ruled out, it was not possible to conclude that Migdal ER rates have been overestimated.

This experiment has motivated future studies with XeNeu to conduct a more detailed search for the Migdal effect. Important aspects of a future search will be a decrease in ionization backgrounds, and increased position resolution to improve multiple scatter rejection. Both of these topics are covered extensively in the following sections. In addition to this, significant contributions were made to the general upkeep of XeNeu when the Migdal search took place. In particular, one of the high voltage feedthroughs (Figure 4.9) had suffered damage, coating many of the detector systems in silicone oil. The central detector components were fully removed from XeNeu and needed to be completely disassembled and cleaned while the high voltage feedthrough was repaired. This offered a unique opportunity to exhaustively measure and document all of the components which would otherwise be obscured.

### 4.3. Plastic Reduction Efforts in XeNeu

As covered in Section 2.4.2.3, electron train backgrounds are known to diminish in intensity with an increase in the electron lifetime. This is likely because impurities in the liquid bulk are responsible for capturing and releasing electrons along the drift path (Section 3.4.3). A major source of impurities in a vacuum system are those which are adsorbed or dissolved in hardware components. As a system is evacuated, these impurities desorb and diffuse into the vacuum space. This phenomenon is known as outgassing. In order to test the effects of outgassing on electron background rates, a campaign was conducted to measure electron background rates before and after

high-outgassing plastic components were replaced with ones made from low outgassing materials. Preliminary results indicated a reduction in electron trains, which should inform the development of future TPCs and benefit future studies of low energy interactions with XeNeu.

#### 4.3.1. Overview of Outgassing

Outgassing is well known to be the limiting factor in determining how high of a vacuum can be reached in a system [119]. Outgassing occurs because prior to evacuating, contaminants are attached to the surface of a material (adsorption) or dissolved in the bulk volume (absorption). When a system is brought to vacuum, adsorbed compounds may detach from surfaces, and absorbed compounds will diffuse through the bulk and into the vacuum space (Figure 4.22). Gases, especially lighter species like hydrogen and helium, can permeate a structure and diffuse into a vacuum space as well.

For a given hardware component, several varieties of contamination may be present [102]. Proper cleaning will remove oil and grime from manufacturing and commercial handling, after which an adsorbed layer of water and light hydrocarbons will remain on the surface of a component. This layer is recharged any time a component is exposed to the atmosphere. Metallic hardware components will have an oxide layer below this, and depending on the surface finish and porosity of the component, contaminants may diffuse deeper into a material. Existing oxide layers and surface finishes can be regenerated with electrochemical processes that form more compact oxide layers and smoother finishes to prevent the ingress of contaminants [120]. Heating components (baking out) to temperatures  $>200^{\circ}\text{C}$  significantly raises the outgassing rate of adsorbed contaminants, to the extent that after a bakeout, outgassing from metallic components is dominated by the diffusion

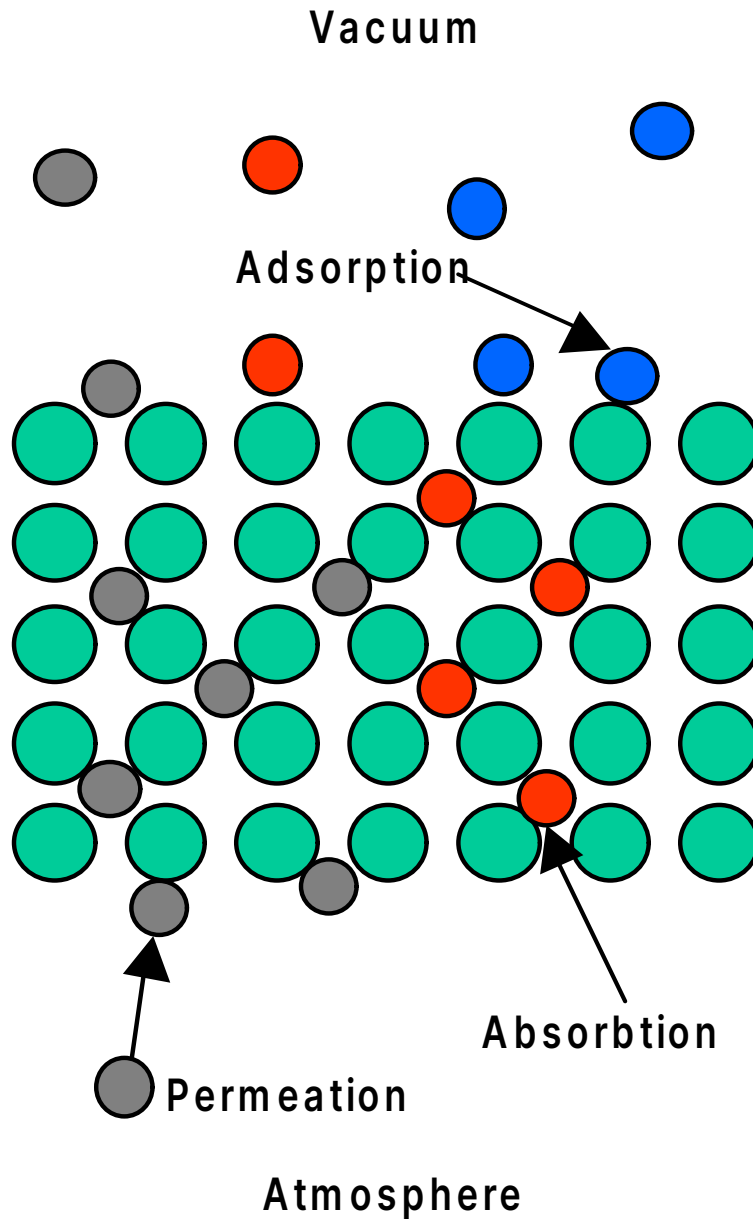


Figure 4.22.: Diagram showing key processes for outgassing. The bulk material is represented by the grid of larger green circles. These processes are not mutually exclusive. For example a species which is permeating the material may adsorb and desorb repeatedly as it diffuses into the vacuum space. Figure adapted from [119].

of hydrogen through the bulk of a metallic component.

Polymer components have a far more porous structure than metallic components. Whereas outgassing for metallic components largely depends on the conditions of the exposed surface, the overall mass and geometry of a polymer component is relevant to its outgassing profile. As noted in [102], oxygen atoms dissolved in stainless steel are expected to diffuse only a few microns per millennium, whereas water molecules in PEEK can diffuse the same distance in a day [121]. Consequently, dissolved concentrations of water in polymers can reach around a percent by weight at ambient levels of humidity. With much lower maximum bakeout temperatures than metals, outgassing contributions from polymers in a vacuum system will outweigh the effects of metals by orders of magnitude.

After the cleaning of surface contaminants, but absent any baking process, the dominant outgassing species in a system previously exposed to atmosphere is water [122]. After enough time under vacuum, outgassing of water becomes subdominant to other species like carbon dioxide, and finally hydrogen [102]. For metals, the outgassing rate is inversely proportional to the pumping time, whereas for polymers the outgassing rate is inversely proportional to the square root of the pumping time [123]. Despite this, PEEK, PTFE, and other plastics are ubiquitous in dual-phase TPCs. Both are relatively affordable and machinable, have good radiopurity, and are suitable for use as dielectrics around high voltage components. PTFE in particular is essential for attaining good light collection efficiency. However, the increase in delayed electron background rates from impurities released by plastics has negative consequences for dark matter searches near detection thresholds. The aim of this study is to establish the extent to which plastic components affect delayed electron background rates in dual-phase xenon TPCs.

### 4.3.2. Design of a Plastic-Free Bottom PMT Holder

The PTFE bottom PMT holder (Figure 4.5) in XeNeu is the single largest piece of plastic in XeNeu, with a mass of approximately 600 grams. Its primary function is to act as an insulator for the high voltage system, and it also displaces xenon from inactive regions of the inner volume. Because it has no role in improving light collection efficiency, it was an excellent candidate to replace with low-outgassing materials.

The final iteration of the plastic-free PMT holder is shown in Figure 4.23. The design consists of segmented components made from 6061 aluminum and Shapal, a ceramic which is machinable with standard tooling<sup>4</sup>. Regions near the field cage and high voltage feedthroughs needed a material with a high dielectric strength, so a machinable ceramic was an ideal choice for this purpose. For regions far from the high voltage feedthroughs, aluminum was chosen for its low cost and machinability. Macor was initially considered instead of Shapal due to its lower cost, however it was abandoned because it consists of 10% potassium oxide by weight. For a 600 gram piece of Macor, <sup>40</sup>K would make up roughly 8 milligrams. With a specific activity of 265.4 kBq/g, Macor replacement parts would have an activity of well over 1.5 kBq, making them unsuitable for use in XeNeu.

The PMT holder assembly consists of four key components: the aluminum base, Shapal pockets for the high voltage feedthroughs, the grid contact ring, and one Shapal retaining ring at the top of the assembly. The pockets drop into slots in the base, which have an angled bottom that slopes downward toward the center of the base. The tops of the pockets are coincident with the top of the aluminum base. The Shapal retaining ring is affixed to the top of the aluminum base with

---

<sup>4</sup>Shapal components were supplied by [Precision Ceramics](#), and authorized distributor of Shapal.




Figure 4.23.: Exploded view of the final design of the new bottom PMT holder. From top to bottom: the Shapal retaining ring, the grid contact ring, two Shapal high voltage pockets, and the aluminum base. White components are made from Shapal. The grid contact ring is colored green here to improve its visibility, but is made from aluminum like the rest of the PMT holder. Screws used were made of nylon due to their proximity to high voltage components.

screws. Once the retaining ring is in place, the Shapal pockets cannot be removed because of the sloped bottoms. The grid contact ring fits into the underside of the retaining ring; it is designed to provide a grounded surface for the PMT shield grid. This is accomplished by grounding the aluminum base, which in turn grounds the grid contact ring when the retaining ring is seated. Three additional screws protrude radially into the inner cavity of the PMT holder. These screws are useful for maintaining clearance between the bottom PMT, which is at high voltage, and the inside of the PMT holder, which is grounded.

Aside from constraints involving close part clearances, the only additional challenges with designing Shapal components included a minimum wall thickness of 2 mm, and restrictions on the surface-to-volume ratio of Shapal components due to a limited selection of stock part sizes. This was another reason for the segmented part design—the total length of the Shapal pockets could not exceed 3.5”.

The components arrived mostly as designed. Inside corners in the slots for the HV pockets were not accounted for during the design process, and had a tooling radius which had to be filed by hand. The slots for the pockets were also slightly too narrow due to a lack of consideration for proper tolerances; these were also filed by hand. The grid contact ring was intentionally oversized so that after soaking it in liquid nitrogen, it was easily seated in the Shapal retaining ring.

The Shapal components have an almost “chalky” surface finish, like that of a high grit abrasive. Being a ceramic, softer metals like aluminum, copper, and solder compounds easily embed themselves when rubbed against Shapal. This was a cause for concern, as insertion of the high voltage feedthroughs into the HV pockets left behind metal “tracks” (Figure 4.24) in the interiors of the pockets. There was some concern about these tracks and whether they could lead

to a short between the ends of the high voltage feedthroughs, or dislodge in the liquid xenon and create a short between other components in the field cage. To reduce the chance of these tracks forming, pieces of Kapton film were shimmed between the feedthroughs and walls of the pockets prior to insertion or removal of the feedthroughs.




Figure 4.24.: View of PMT holder assembly from above, with the field cage removed except for the PEEK base ring. Dark tracks are visible in the cathode (left) and gate (right) pockets. These tracks are formed from rubbing of metal surfaces on Shapal components.

It was also found that multiple 30 minute cycles in a 40°C ultrasonic bath would completely remove the embedded metal tracks from the Shapal components. [Detergent 8](#) was used for cleaning parts, however it was soon discovered to be unsuitable for use with aluminum due to its high pH. Aluminum parts cleaned with Detergent 8 developed a non-conductive layer with a tarnished

appearance. This was problematic for the connection needed between the grid contact ring and the aluminum base—both parts had to be lightly sanded to remove the oxide layer after it had been formed. For subsequent cleaning cycles, [Alconox](#) was used instead with no external heating applied<sup>5</sup>.

Detector operations with the new Shapal and aluminum components were almost completely identical to the previous PTFE PMT holder. The detector space reached a base pressure of 1.85E-5 torr after approximately 48 hours of pumping, a factor of two lower than the base pressure typically attained with the PTFE bottom PMT holder. The only anomaly during operations occurred during xenon condensing. Because aluminum is a better thermal conductor than PTFE and has a higher thermal mass, liquid xenon at the bottom of the vessel began to rapidly boil when it made contact with the PMT holder. This led to a 10 K temperature rise in the span of 1-2 minutes (Figure 4.25). Upon witnessing this, condensation was immediately stopped and recovery lines were opened to vent a few grams of xenon into the SRV.




Figure 4.25.: Screenshot of slow control readout during condensation of xenon with redesigned bottom PMT holder. When liquid xenon made contact with the aluminum bottom PMT holder, over the course of two minutes the detector temperatures rose almost 10 K.

<sup>5</sup>While a neutral pH cleaner like Luminox may be more appropriate for aluminum, Alconox was chosen as it is a more general purpose cleaner suitable for a wide variety of media found in XeNeu.

### 4.3.3. Single Electron Trigger Efficiency

In order to study rates of SEs after S2s, the trigger efficiency of S2-like pulses was characterized. A three-fold coincidence trigger between all four of the top PMTs was used. Pairs of the top PMTs were biased to either 680 V or 780 V, and the bottom PMT was biased to 1020 V to match their single photoelectron (SPE) sizes. All PMTs were powered with a CAEN 1470 high voltage power supply. A schematic of the trigger scheme is shown in Figure 4.26. PMT signals were amplified at the output of the detector before being fanned out and passed into the digitizer or another amplifier. The second amplifier was necessary in order to exceed the minimum thresholds of the discriminator modules. Each discriminator was set to fire on SPEs by comparing the discriminator and PMT outputs simultaneously on an oscilloscope. Outputs from the discriminator were fed into a 3-fold coincidence module, which was connected to the digitizer, a gate generator, and a 1-fold coincidence module.

Output from the 3-fold coincidence module was fed directly into the digitizer so all 3-fold coincidence signals could be captured without being vetoed. These signals were used for trigger efficiency calculations. To generate the main event trigger, signals from the 3-fold coincidence module were fed into a veto logic consisting of a 1-fold coincidence module and gate generator. At the 1-fold coincidence module input, 3-fold coincidence signals would arrive before the gate generator signal due to inherent ( $\sim 10$  ns) delays in the gate generator module. This meant that while the system was not being vetoed by the gate generator, an initial 3-fold coincidence would trigger the 1-fold coincidence, but subsequent 3-fold coincidence pulses within the veto period would not.

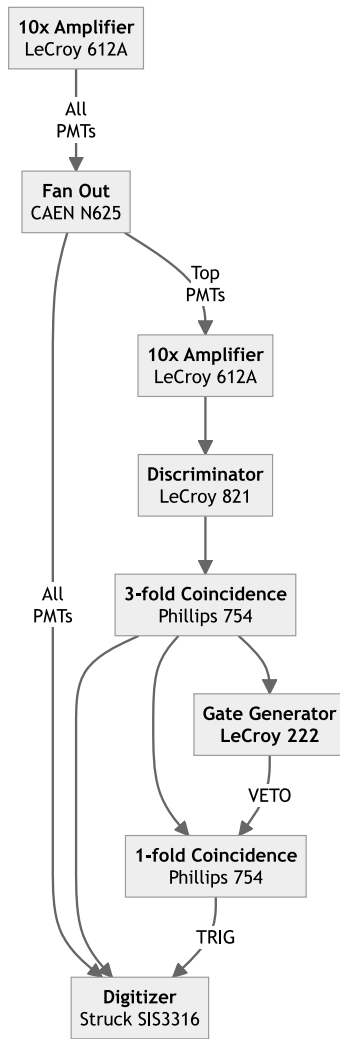


Figure 4.26.: Trigger schematic for studies of SE rates. Signals needed to be amplified twice in order to exceed the minimum discriminator threshold. Output from the 3-fold coincidence module was recorded at the digitizer for trigger efficiency studies. The 1-fold coincidence module provided the main event trigger. A pulse from the 3-fold coincidence would also fire both the 1-fold coincidence and a gate generator. A long cable was used for the gate generator signal to stagger the arrival times of the 3-fold coincidence and vetoing pulses at the 1-fold coincidence so the first coincidence pulse in an event would not veto itself.

The trigger efficiency analysis used area fraction times (AFTs) to reject S1-like pulses<sup>6</sup>. AFTs are timings where a specific area fraction of a pulse is first reached. For example, the AFT 50 of a pulse is the time taken for a pulse to reach 50% of its total area. For S1-like pulses, the AFT 75-25 is 50 ns, i.e. 75% of the total area is reached only 50 ns after 25% of the total area has been reached. SE pulses meanwhile have an AFT 75-25 greater than 250 ns. This can be seen in Figure 4.27, which shows the pulse area of pulses versus their AFT 75-25.




Figure 4.27.: Plot of AFT 75-25 vs pulse area showing counts of pulses. Single electron pulses are roughly within the closed red dashed curve. The band of pulses below the dashed orange line corresponds to S1-like pulses, which have a very low AFT 75-25.

For this analysis, S1-like pulses were rejected by accepting only pulses with an AFT 75-25  $> 250$  ns. Trigger signals were exempt from this cut because by design they have a very small pulse width. Second, pulses were required to be separated from each other by at least  $0.5 \mu\text{s}$  in order to ensure they were properly resolved. Because multiple trigger pulses could fire in quick

<sup>6</sup>Prompt fraction, which is the percentage of a pulse within a set time window (e.g. fraction of a pulse within 500 ns) can also be used for rejecting S1-like pulses. Lenardo [105] used  $F500 > 0.8$  to reject S1-like pulses.

succession for a single pulse of interest, they were also exempt from this cut. Finally, all pulses and trigger signals were required to be at least 5  $\mu$ s after the main event trigger. This was because if the trigger efficiency was greater than zero for a specific pulse, the pre-trigger window would have *fewer* pulses of interest since those same pulses would have caused the trigger to fire with some non-zero probability.

The trigger efficiency was evaluated by dividing the number of pulses which cause a trigger by the total number of pulses. Pulses which caused the trigger to fire were found first by digitizing all trigger signals independently of any veto. Then a `pandas` dataframe with rows corresponding to pairs of triggers and pulses on a per-event basis was produced (Table 4.1) with an inner join of pulse and trigger dataframes. `DataFrame.between()` was used to check if trigger times were between pulse start and end times. The inner join and multiple triggers per pulse produced some duplicate rows. These rows were removed by selecting rows corresponding to distinct pulses using event IDs and pulse start times as unique IDs for each pulse.

Table 4.1.: `DataFrame.head()` of merged trigger and pulse dataframe. The join is performed on `event_id` and `event_timestamp` which uniquely identifies events. `peak_time` is the peak time of a trigger signal relative to a main event trigger, while `npe`, `start_time`, and `end_time` are quantities belonging to a pulse in the same event. The first three rows correspond to pairings of three different pulses with the *same* trigger pulse. Subsequent steps detailed in the text explain how distinct pulses with triggers are selected.

	event_id	event_timestamp	peak_time	npe	start_time	end_time
0	5	1626812919	9.124	33.103590	6.736	10.736
1	5	1626812919	9.124	83.439986	16.524	21.828
2	5	1626812919	9.124	47.499900	23.792	28.088

An example trigger efficiency histogram is shown in Figure 4.28. The trigger efficiency for SE pulses is approximately 75%, which offers a compromise between digitizing all pulses and saturating the data acquisition system. Besides lowering discriminator thresholds, the trigger efficiency of SEs could be improved by increasing the extraction field strength in order to increase the number of photons produced by single electrons.




Figure 4.28.: Example of a trigger efficiency histogram. The first peak at a pulse area of roughly 60 photoelectrons (npe) corresponds to SEs. The trigger efficiency approaches 100% for pulses larger than four electrons in size.

#### 4.3.4. Electron Background Rates

To test the effect of plastics on electron background rates, preliminary data was taken in two runs: Run 11 and Run 12. In Run 11, the PTFE bottom PMT holder and PTFE field cage reflector Figure 4.4b were used. In Run 12, the field cage reflector was removed, and the plastic-free bottom PMT holder Figure 4.23 was used instead. The same trigger configuration (Figure 4.26) was used

in both runs. Table 4.2 shows configurations for each run:

Table 4.2.: Key operating parameters for data taken during Runs 11 and 12. Run 11 had a slightly smaller gas gap, leading to a higher field.

	Circulation	Gate	Cathode	Liquid Ext.		
	Rates	Voltage	Voltage	Field	Gas Field	Drift Field
Run	[sccm]	[-kV]	[-kV]	[kV/cm]	[kV/cm]	[V/cm]
11	300, 700, 1100	11	12	7.0	13.0	420
12	300, 1200	12	13	6.7	12.4	420

Both runs had an extraction efficiency approaching 100% [71], and a single electron size of approximately 55 photoelectrons (npe). Trigger efficiencies for SEs were typically centered around 75% in both runs. For the purpose of tracking livetime this was assumed to be 100%, but should be revisited in a future work with improved trigger settings or livetime corrections.

The analysis was conducted using the same method by porting over the ALPACA module developed in Section 3.3.2. The code was initially incompatible with ROOT6 because some classes in `daqman` relied on ROOT5. This incompatibility was resolved by creating a header-only file [109] with core `daqman` data classes, which was compatible with ROOT6. After removing all external ALPACA functionality, the core C++ code was compiled and loaded into `pyroot`. Fully multi-threaded data processing with `dask` in `pyroot` was possible because the global interpreter lock could be disabled on bindings generated by `cppyy`. The TTrees generated by the ALPACA module were then processed into analysis-ready parquet files for use with the `etrain` package that was custom-developed for use in the LZ analysis.

For this analysis, S2-like “child” pulses were examined following S2-like progenitors greater than 50,000 npe (~1000 electrons) in size. Progenitors were required to be fully contained in an event window, and occur more than 100 ms since the previous progenitor. This holdoff was chosen because increasing it to 200 ms, 500 ms, and 1000 ms had no effect with the 50,000 npe veto threshold. Pulses under this 50,000 npe threshold were considered child pulses. This large threshold was necessary due to the prevalence of extremely high energy events which would saturate the detector with S2-like pulses for tens of milliseconds or longer (Figure 4.29).




Figure 4.29.: Rates of child pulses following progenitors in Runs 11 and 12. Rates are normalized following the convention established in Equation 3.8. Child pulses are additionally binned by their size along the y-axis. Pulse sizes are normalized by the size of the single electron pulse in each run.

Many features in Figure 4.29 are worth remarking on. In both runs there are vertical bands from about 30  $\mu\text{s}$  to about 700  $\mu\text{s}$  where the pulse rates appear to be lower than in neighboring delay time bins. The large deficit of pulses at 30  $\mu\text{s}$  is a consequence of the length of the post-trigger window and 60  $\mu\text{s}$  holdoff period. The nearly even spacing of subsequent bands and their visibility at microsecond timescales is consistent with this assumption. Furthermore, inspection of the raw counts and livetime histograms shows either an under-counting of pulses in these time delay bins

or over-counting of livetime. The bands also appear to be visible across all child pulse area bins, so this issue is not due to the trigger efficiency which should be 100% for pulses  $> 2$  SE in size.

In the ensuing few milliseconds after a progenitor, child pulses up to 10s of electrons in size (~few thousand pe) appear to be highly prevalent for up to a few milliseconds in both runs. Many of the larger child pulses are associated with extremely energetic progenitor events, and occur in the midst of trains of smaller child pulses. During this time period the single electron size also shifts upwards. This feature is a sign of excessive PMT saturation resulting in temporary gain shifts.

Beyond 1 ms, additional vertical bands are present with rates higher than neighboring time delay bins. Handscanning some of these reveals “ghost” S2s which are S2 pulses that occur during a holdoff period and are not fully captured in an event window. The high veto threshold also enhances the possibility of these being from low energy events or isolated S1s which photoionize the TPC but do not produce an S2.

Figure 4.30 shows the pulse rate integrated over all child pulses up to 10 SE in size. Rates in Run 11 were 3-5 times higher than in Run 12, depending on the circulation rate. The circulation rate also had a more distinguishable impact in Run 11. This suggests that plastic components in the TPC are a dominant source of impurities responsible for ionization backgrounds. With low-outgassing components in Run 12, the circulation rate had a much smaller impact on reducing background rates.

The spectra in Figure 4.31 show more clearly that integrated rates for pulses of all sizes are affected by the circulation rate in Run 11. This is in contrast with rates for Run 12 in Figure 4.31, where only SE rates appear to have a slight circulation rate dependence when integrated from 100  $\mu$ s to 1 s.




Figure 4.30.: Total integrated rates from Figure 4.29 over pulses up to 10 SE in size. For this plot, rates were integrated from 100  $\mu$ s to 1 s. At all circulation rates, approximately 60% of the pulse rate in Run 11 is due to SEs, whereas SEs make up almost 80% of the child pulse rate in Run 12. With low-outgassing components in Run 12, the circulation rate had a much smaller impact on reducing background rates. This suggests that plastic components in the TPC are a dominant source of impurities responsible for ionization backgrounds.




Figure 4.31.: Projection of Figure 4.29 showing integrated rates of child pulses up to 10 SE in size at different circulation rates. For this plot, rates were integrated from 100  $\mu$ s to 1 s. Rates for pulses of all sizes are affected by the circulation rate in Run 11. This is in contrast with rates for Run 12 in Figure 4.31, where only SE rates appear to have a slight circulation rate dependence.

The presence of larger pulses along with intense trains of smaller pulses is consistent with observations following muons or other high energy events in LZ. The mechanism responsible for producing the larger child pulses at prior to approximately 1 ms is unclear. It is speculated here that imperfect extraction at the liquid surface may occur even at high extraction efficiencies due to field shielding effects from the vast quantities of charge liberated in a very energetic event. This could lead to electron-burst [81] or S3-like [96] phenomena where a few electrons at a time are continuously extracted over the course of milliseconds. This would repeatedly photoionize impurities in the TPC, explaining the strong circulation rate dependence for larger child pulses in Run 11.

While the total area and few  $\mu\text{s}$  widths of the smaller pulses do not match descriptions of electron-bursts and S3 pulses, this study lacked the capability to examine waveforms over long enough time periods to assess whether the periods of “sustained” emission were indeed electron bursts. Future work should be conducted to stitch together waveforms from multiple events and investigate the nature of these emission periods over longer timescales.

The time dependence of the single electron pulse rates is shown in Figure 4.32. The circulation rate dependence is visible in both runs, and more obvious for SE pulses in Run 12 after roughly 1 ms.

Rather than a power law, the time dependence between roughly 400  $\mu\text{s}$  and 30 ms in both runs is well-described by the sum of two decaying exponentials:

$$A_0 e^{-t/\tau_0} + A_1 e^{-t/\tau_1} \tag{4.1}$$

An example fit is shown in Figure 4.33 and fit parameters for each curve shown in Figure 4.32 are plotted in Figure 4.34.




Figure 4.32.: Projection of Figure 4.29 examining rates of single electron pulses over time at different circulation rates. Rates prior to 100  $\mu$ s are excluded due to the high prevalence of merged pulses.




Figure 4.33.: Example of a double exponential fit between 400  $\mu$ s and 30 ms in the single electron rate time dependence.

The lower rates in Run 12 are reflected in Figure 4.34 with lower amplitudes. In Run 11, both amplitudes and time constants decrease as the circulation rate increases. While  $\tau_1$  in Run 12 is more than 50% higher than  $\tau_1$  for Run 11, this may be an artifact of the rates in Run 11 and 12 approaching a similar baseline rate.




Figure 4.34.: Fit values for each fit parameter and rate curve shown in Figure 4.32. Error bars represent statistical uncertainties. Values are plotted along the x-axis according to their run number and circulation speed; they are also printed in Table 4.3.

Table 4.3.: Fit values and uncertainty for each value shown in Figure 4.34. Uncertainties are statistical only.

Run ID, Circulation				
Rate [sccm]	$A_0$	$\tau_0$	$A_1$	$\tau_1$
(11, 300)	$80.2 \pm 2.3$	$1.35 \pm 0.0775$	$0.000729 \pm$ $1.33e-05$	$0.00885 \pm$ $0.000348$
(11, 700)	$74.8 \pm 1.75$	$1.17 \pm 0.0541$	$0.000685 \pm$ $9.94e-06$	$0.00823 \pm$ $0.000251$
(11, 1100)	$70.5 \pm 1.65$	$1.05 \pm 0.0471$	$0.000648 \pm$ $9.19e-06$	$0.00785 \pm$ $0.00023$
(12, 300)	$14.3 \pm 0.337$	$0.291 \pm$ $0.00421$	$0.000449 \pm$ $5.62e-06$	$0.0149 \pm$ $0.000238$
(12, 1200)	$13 \pm 0.31$	$0.243 \pm 0.004$	$0.000446 \pm$ $5.65e-06$	$0.012 \pm 0.000185$

In LUX [81], electron bursts were found to have a time delay relative to an S2 up to 10s of milliseconds, with an exponentially decaying distribution that had a time constant of  $7.5 \pm 0.3$  ms. It is noteworthy that  $\tau_1$  in Run 11 is so similar to the result from LUX, which also heavily used PTFE and PEEK throughout its construction. This exponential also appears to decay faster in both runs at higher circulation speed. and it decays faster in Run 11 than in Run 12. Further investigation of this behavior is certainly warranted—at the very least, there are indications in the data that the composition of materials in a TPC have a strong effect on the fine structure of ionization backgrounds.

Because LUX did not present their results with a logarithmic time axis, it is unclear if a first exponential could have been present at shorter timescales. The only other work to observe an exponential decay in the context of electron/photon trains was the study conducted by Sorensen and Kamdin [99]. The time constants they measured were on the order of 10s of microseconds, and a few milliseconds. Their use of a  $^{210}\text{Po}$  source produced low energy  $^{206}\text{Pb}$  recoils, as well as 2.4 MeV alphas. The latter may have been similarly high enough in energy to create the conditions for exponential decays to form, but the nature of their summed waveform analysis again obscures possible comparisons.

Figure 4.35 compares the 300 sccm rate curves for both Run 11 and Run 12. Rates in Run 12 are slightly lower at long timescales, and the rates in Run 11 and 12 also appear to converge on a power law with a constant component after 30-40 ms:

$$\alpha t^\beta + c \tag{4.2}$$

Fitting Equation 4.2 to both spectra shown in Figure 4.35 produces the fit parameters in Table 4.4.

Table 4.4.: Fit values from Equation 4.2 applied to each rate curve shown in Figure 4.35. Fits were performed in the range of 40 ms-1 s. Uncertainties are statistical only. Because the fit does not account for the presence of a possible decaying exponential, the magnitudes of  $\beta$  may be elevated by  $\sim 10\%$ . There is an additional  $\sim 10\%$  systematic error on  $\beta$  in the Run 11 rates that stems from the choice of the left edge of the fit.

Run ID	$\alpha_0$	$\beta_0$	$c_0$
11	$0.000196 \pm 3.64\text{e-}05$	$-1.56 \pm 0.0648$	$0.00148 \pm 0.000153$
12	$0.000259 \pm 3\text{e-}05$	$-1.43 \pm 0.0399$	$0.000655 \pm 0.000106$




Figure 4.35.: Single electron rates from Run 11 and Run 12. The rates in both runs appear to converge on a power law with a constant component (Equation 4.2) after 30-40 ms. Fit parameters for these power laws are presented in Table 4.4.

This piece-wise fit could be justified due to the exceptional amounts of energy deposited in the detector at earlier timescales. Assuming that electron trains result from electrons trapped on liquid bulk impurities, then repeated photoionization from electron bursts may disrupt the trap-and-release process that takes place in a conventional electron train.

While there appears to be a disparity in  $\beta$  between runs, this is an artifact of systematic errors in choosing the fit bounds. Furthermore, this fit does not account for the presence of a decaying exponential, which would soften the power law slope. In total, the “true” value of  $\beta$  may be between -1.2 to -1.3. The data presented here was taken with a drift field of approximately 420 V/cm, so this result would be consistent with the value of  $\beta = -1.2$  observed by Kopec et al. [94].

Nevertheless, it is important to highlight that the discussion presented here is highly speculative. Future studies of electron backgrounds with XeNeu should take into account a number

of factors which were not considered here. First, a more detailed study of a proper trigger configuration which eliminates trigger bias or raises the trigger efficiency to be closer to 100% would resolve doubts of rates being improperly calculated at all timescales. Second, the phenomenon of high energy depositions could be partially addressed by acquiring data at an underground location. If an underground location is inaccessible, then data acquired at lower fields and photosensor gains would also reduce the possibility of saturation, at a cost of detection efficiency for smaller signals.

In conclusion, these preliminary results indicate that a factor of 3-5 reduction in electron backgrounds is possible by replacing plastics with low-outgassing components (Figure 4.30). The electron train rates may be consistent with a power law, however there are more complex features present which prevent this from being conclusively determined. Future studies of electron backgrounds should include a more careful consideration of trigger bias, increased exterior shielding, lower gain values, and make use of the increased position resolution from the silicon photomultiplier array detailed in Section 4.4.

## 4.4. S2 Detection Improvements with Silicon Photomultipliers

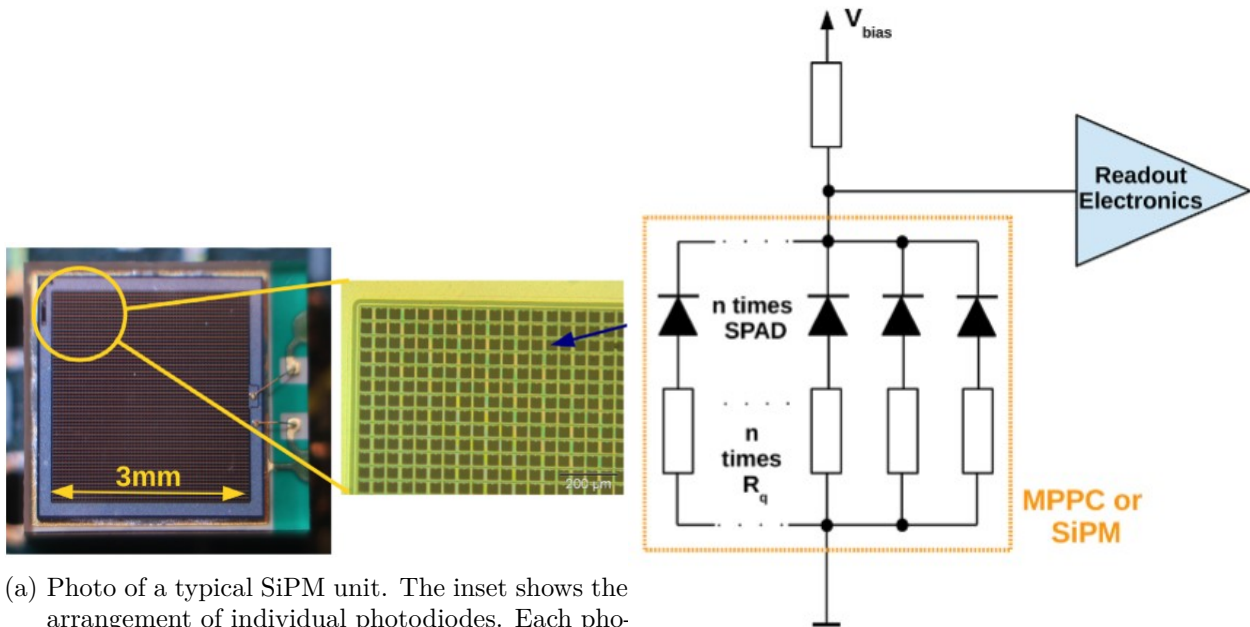
This section will detail the replacement of the top PMT array (Section 4.1.1.3) with an array of silicon photomultipliers (SiPMs) to improve XeNeu's position sensitivity. A basic characterization of the SiPM array was performed, including an evaluation of the position resolution of the new array. Among other benefits, the improved position resolution will allow for better identification of SEs in studies of low-energy ionization backgrounds, and improve rejection of multiple scatter backgrounds in searches for the Migdal effect.

#### 4.4.1. Introduction to SiPMs

SiPMs are a class of solid-state photosensors with properties that make them attractive replacements for PMTs. Compared to PMTs, SiPMs are compact and can be manufactured with almost any profile to maximize the photon detection area in a given space. This leads to the possibility of detectors with increased position resolution and better background rejection capabilities. SiPMs also have a lower inherent radioactivity than PMTs [124], and their bias voltages are on the order of dozens of volts instead of kilovolts. Finally, SiPMs are resistant to the effects of external magnetic fields, enabling applications in environments which would be prohibitive to PMTs.

A SiPM consists of numerous single-photon avalanche diodes (SPADs) arranged in parallel (Figure 4.36). Each SPAD is sensitive to a single photon, so a single SiPM is capable of resolving hundreds of photons simultaneously assuming they strike different SPADs. Figure 4.36a shows an example of a 3 x 3 mm SiPM, with an inset showing the arrangement of individual photodiodes.

A typical photodiode is operated with a reverse bias of a few volts. When a photon strikes the active region of the photodiode, an electron-hole pair is formed. The pair is split by the applied voltage, forming a current. At a bias voltage of a few dozen volts, the photodiodes operate as avalanche photodiodes (APD) where an electron moving through the diode can produce additional electron-hole pairs in the diode bulk. In this scenario, the initial electron creates an avalanche of electrons and a larger detected current. While holes are also created, the electron avalanche is not self-sustaining because the larger effective mass of holes prevents a corresponding avalanche of hole-electron pairs [125]. It is only when an even stronger bias is applied that a self-sustaining avalanche cascade can occur. In this regime, an avalanche of electrons creates a countervailing avalanche of holes, which create another avalanche of electrons, multiplying the current *ad infinitum*. This



(a) Photo of a typical SiPM unit. The inset shows the arrangement of individual photodiodes. Each photodiode is sensitive to single photons, so a SiPM is capable of resolving hundreds of photons simultaneously. From [125].

(b) Simplified schematic of a SiPM. During normal operation, a SPAD is biased with  $V_{bias}$  because no current is flowing through the quenching resistor  $R_q$ . Photons striking a SPAD induce a current, resulting in a voltage drop across  $R_q$  which subsequently lowers the bias voltage across the SPAD. This process prevents avalanches in a SPAD from becoming self-sustaining and allows them to reset to state where photons can be detected again. From [125].

Figure 4.36.: Typical dimensions of a SiPM and equivalent circuit for a single SiPM.

process has been documented to occur within a few hundred picoseconds [126]. Once a signal is produced, the bias voltage in the photodiode is lowered below the breakdown voltage by a quenching resistor (Figure 4.36b) to halt the avalanche cascade over the course of hundreds of nanoseconds. The onset of this avalanche cascade typically occurs above 50 V in SiPMs, and is called the “breakdown voltage”. This process is what distinguishes the operating regime of APDs from SPADs, where a gain of 100 or so can be expected from the former, and  $10^5$  to  $10^6$  can be expected from the latter [125].

Sources of noise in SiPMs include crosstalk, afterpulsing, and dark counts. Crosstalk occurs when charges in an avalanche generate secondary photons which can trigger an avalanche in a neighboring cell [127]. This can result in an erroneous doubling or tripling of SPE signals. Afterpulsing can occur when a charge carrier in a SiPM is trapped on a defect in the avalanche production region. When the charge carrier is released after some time, a distinguishable avalanche can be detected. Finally, dark counts are pulses formed by charge carriers which are spontaneously liberated by thermal fluctuations. Dark counts are the dominant source of noise in SiPMs, typically on the order of kHz to MHz per square millimeter of active area at room temperature [127]. These sources of noise become more intense at a higher overvoltage, i.e. voltage above the breakdown voltage.

Dark counts in particular also exhibit a strong dependence on the operating temperature of the SiPM. This dependence is exploited in noble liquid scintillators, where the cryogenic operating temperatures result in a dark count rate as low as 0.05 Hz/mm<sup>2</sup> (Figure 5.19). This is still a factor of more than 10 larger than dark count rates of PMTs in an experiment like LZ [128], but it is adequate in other scenarios where larger signals are expected. For example, efforts are underway to design large-scale SiPM arrays for experiments like nEXO [129] and DUNE [130]. Their viability

has already been demonstrated in small liquid xenon TPCs [124], and manufacturers have recently produced SiPMs which are capable of detecting argon VUV light with appreciable ( $\sim 15\%$ ) photon detection efficiencies (PDE) without the use of wavelength shifting compounds [131].

#### 4.4.2. Replacement SiPM Array Design

The new SiPM array in XeNeu had a square arrangement with a side length of approximately 1.8" (Figure 4.3). Nine Hamamatsu S13371-6050CQ-02 SiPMs were biased in groups of three, resulting in a total of 12 channels (9 signal, 3 power). All 12 channels were fed out from the SiPM array through MMCX coaxial jacks mounted on the printed circuit board (PCB). A rendering of this array is shown in Figure 4.37. Teal Pershing designed the PCB which was manufactured and supplied by OSHPark , while James Kingston completed the final assembly of the PCB and array.




Figure 4.37.: Rendering of the new XeNeu SiPM array. Each S13371 unit has four cells which can be read out individually, but in this design the cells in each unit were connected in parallel. A PEEK spacer was also designed to sit on the perimeter of the board and reduce wear between the edge of the board and the metal components in the tray and bracket assembly Figure 4.38.

The SiPM array was mounted in the detector space with a tray and bracket assembly, shown in Figure 4.38. All parts for the tray and bracket were made from 6061 aluminum, machined and supplied by Xometry. The array was held in the tray by two L-shaped brackets which fastened to the tray with spring washers and screws. The anode grid and PEEK anode grid spacer also fastened to the underside of the tray. This subassembly consisting of the SiPM array, anode grid, and tray could be easily assembled on a lab bench before being slid into place on the bracket. Because of the compact form factor of the SiPMs, the MMCX jacks were easily exposed so that signal and power connections could be made *in situ*. This was a substantial upgrade in usability. The previous design with PMTs (Figure 4.6) had cables soldered to bases on the PMTs, which were entirely enclosed in the top PMT assembly. The signal outputs for the PMT cables originated at a 9-pin Sub-C connector on the detector top flange, which was recessed in the underside of the flange to the point of being physically inaccessible. Consequently, removing the top PMTs was a tedious process which involved undoing a 1.33" CF flange to gain access to the 9-pin connector. The new mounting bracket was also slightly improved, with a profile which was easier to slide laterally onto the support rods and additional tapped holes tapped for electrical grounding purposes.

#### 4.4.3. Electronics Re-cabling

Previously, high voltage for all five PMTs was delivered via the 7-pin Mil-spec feedthrough (Figure 4.40), while all signals from the PMTs were transmitted through a 9-pin Sub-C connector (both are shown in Figure 4.7). For signals, the shield and signal pairs were separated and fed through independent channels on the 9-pin connector<sup>7</sup>, whereas high voltage connections shared a

---

<sup>7</sup>A single shield connection for one of the PMT signals was grounded to the feedthrough mating screw on the Sub-C connector as a pseudo-“10th” channel.




Figure 4.38.: Exploded view of the SiPM bracket and tray assembly. The tray is in the foreground, and has a cutout in the center that accepts the SiPM array and PEEK array spacer. Two L-shaped brackets on opposing ends of the tray are screwed into place, capturing the SiPM array. Two additional holes at the rear corners of the tray are designed to be concentric with tapped holes in the bracket so the tray can be securely mated to the bracket. The bracket is in the background and shares a similar design as the previous iteration. Two new holes in the same plane as the holes for the rear mounting clamp are included as optional electrical ground points at the corners of the bracket. Cutouts for the forward mounting posts on the bracket were intentionally asymmetric so the bracket alone could be installed by sliding it laterally onto the mounting posts. Through holes were used in all cases to ensure the absence of any trapped volumes.

single ground on each end of the 7-pin feedthrough. Despite neither feedthrough being designed for coaxial connections, noise from outside sources was fairly minimal thanks to the vacuum vessel acting as a Faraday cage. Inside of the TPC, a significant amount of fast switching noise was present from RTD measurements made by the slow control system. This was mitigated with an extensive patchwork of copper shielding and steel braids around the otherwise unshielded PMT power wires, while the signal cables were entirely coaxial. The new SiPM array required a considerable re-cabling effort to feed more electrical channels into the detector, which is covered in this section.

#### **4.4.3.1. Vacuum Box Feedthroughs**

With the new array, a total of 10 signal channels were needed: one for the bottom PMT, and nine for each of the SiPMs. A total of four power channels were needed, as the SiPMs were biased in groups of three. In order to accommodate the new channels on the vacuum box, 17 new BNC feedthroughs (Pasternack PE9420) were added to a previously blank flange. The excess channels were added for future-proofing and did not significantly increase the complexity of the design or leak rate. All feedthroughs were installed on the vacuum flange with a 5/8" socket and torqued to 60 in-lbs. Using a helium leak checker, three feedthroughs were found to leak substantially at their o-ring seals. They were torqued to 70 in-lbs and still found to leak. These feedthroughs were removed, and replaced with new ones which did not leak beyond  $10^{-8}$  mbar \* L/s after being torqued to 60 in-lbs. While the exact cause of the leaks was not determined, some metal debris (Figure 4.39) was observed on the sealing face of the o-rings after the leaking feedthroughs were removed. This debris may have originated from the feedthroughs being over-torqued, which partially stripped some of their threads.



Figure 4.39.: Pieces of metal were observed on the sealing face of an o-ring (orange) belonging to a leaking feedthrough. The metal was believed to originate from threads on the feedthrough, which was likely overtightened. At the bottom left, a feedthrough as it was normally installed can be seen.

#### 4.4.3.2. Cable Connections

In the detector space, the power delivery cables were connected to pins of the 7-pin Mil-spec feedthrough (Figure 4.40) using standard density female D-sub connectors (Amphenol L17RRD1F01100).




Figure 4.40.: Right: Rendering of the 7-pin Mil-spec feedthrough. Left: Pinout of the 7-pin feedthrough, as used on XeNeu. Numbers and colors corresponded to labels used on the XeNeu vacuum box.

Ground connections on the detector side of the power delivery cables were made with a 0.075” flux-free copper wick (Techspray 1832-5F) soldered to the outer shield (Figure 4.41) of the cables. In order to facilitate the assembly and disassembly of the ground connections between cables, ring terminals (McMaster 7113K263) were crimped onto the free ends of the wicks. These ring terminals were then attached to a main ground terminal with a 2-56 screw and two nuts. The main ground terminal was crimped to a piece of wick that was attached to the ground pin on the 7-pin Mil-spec connector. The modular nature of these ground connections was extremely useful, allowing for a significant consolidation of cable clutter throughout the detector space.

All cables for the SiPM board terminated with crimp-on RG-178 MMCX plugs (Amphenol




Figure 4.41.: Power delivery cable assembly. At the left, crimp on terminals are held together with a 2-56 screw and two nuts. These terminals offered a simplified way to temporarily and securely attach multiple ground connections throughout the detector. The main ground connection (visible at the top) initially used a stainless steel braid; this was later replaced with a copper braid.

262100; crimp tool: Amphenol CTL-13) or a custom set of pins for the bottom PMT (Figure 4.42).

The previous connections to the bottom PMT were notoriously fragile because the pins were soldered directly to the inner conductor of the cable with no strain relief. An improved design was developed which used a ring terminal crimped to the inner dielectric of the cable to protect the inner conductor. The inner conductor was then “potted” with solder, fully immobilizing it in the body of the ring terminal. The large surface of the ring terminal could be bent into a shape that would cradle a pin; the pin could then be soldered and de-soldered to the ring terminal as needed. The ring terminal also doubled as a convenient point for gripping with pliers or tweezers when attaching the pin to the bottom PMT base. This crimp and “pot” technique was also used for the female D-sub connectors on the 7-pin Mil-spec feedthrough.

All signal cables in the detector originated on a female PEEK 23-pin Sub-C connector (Figure 4.43a). The connector is composed of two halves, which capture crimp sleeves that accept pins





Figure 4.42.: Example of a custom PMT pin connection on a coaxial cable. The pin is soldered to a ring terminal, which is crimped to the inner dielectric of the cable. The inner conductor is inside the crimp sleeve, which is then flooded with solder. The solder-covered tip of the inner conductor can be seen in this picture emerging from the body of the crimp sleeve, to the left of the base of the PMT pin. Not shown here are the copper wick to extend the outer shield, and heatshrink around the body of the ring terminal and cable in order to provide additional strain relief.

from the male connector. Crimps were made with a DMC AFM8 crimp tool at setting 3. A scrap piece of PTFE was used to suspend and center the sleeves in the crimp tool, though a properly sized die would have been preferable. During assembly of the connector it was important to install cables starting from the center of the connector in order to avoid blocking access to any of the holes. Intuition would suggest that the connector sleeves could be crimped onto wire leads, which would dangle out of the top half of the connector where they could be soldered to the inner and outer conductors of a coaxial cable. However because of the fragility of the cables' inner conductors, the sleeves needed to be crimped directly to the inner dielectric of the cables in order to minimize strain on the inner conductors. The outer conductors were connected to the pins using 28 AWG stranded core Kapton wire as a lead for pins in the detector space, and 24 AWG stranded core wire in the vacuum space. Because the body of the Sub-C connector does not offer any shielding from interference, effort was made to avoid having an abundance of adjacent signal wires. The resulting wire mapping can be seen in Figure [4.43b](#).

#### **4.4.3.3. Stripping Kapton Cables**

New sets of coaxial cables were made in order to avoid the complications of applying ad-hoc pieces of shielding to cables in the TPC space. As before, the signal cables were broken into shield and signal pairs and fit on a single 23-pin Sub-C connector, while power channels re-used the 7-pin Mil-spec connector. One anticipated tradeoff of using standard FEP-insulated RG-178 coaxial cables was an increased level of outgassing in the TPC space. Consequently, a substantial amount of effort was invested in making cable sets from 30 AWG Kapton-insulated coaxial cable (AccuGlass 112130). The intent behind this effort was to have sets of each kind of cable in order



(a) Cross-section of the 23-pin Sub-C connector. The two halves of the connector are held together with two flathead screws (only one shown). An additional socket head screw is captured between the two halves; this screw fastens the connector to a feedthrough flange. Wires enter the connector at the top, and are crimped to the sleeves, which are captured between both halves of the connector.

(b) View of the as-built connector on the detector side of the detector top flange, with cables (not shown) coming out of the page. Coaxial cables were split such that signal (+) and shield (-) pairs were individually mapped to a pin on the 23-pin Sub-C connector. Channels 1-9 correspond to their respective SiPM PCB jacks (J1-J9), while channel 10 connects to the bottom PMT.


Figure 4.43.: 23-pin Sub-C connector cutaway and pin mapping. This connector was used to feed signals from the SiPM array and bottom PMT through the detector top flange (Figure 4.7).

to complement a future study of electron backgrounds with high and low-outgassing materials. 30 AWG cable was chosen because of its similar dimensions to FEP RG-178 cables. It also exhibited a good compromise between its minimum bend radius (1/4") and quantity of inner conductors (7 x 0.004 in). Due to time constraints, only the FEP set of cables was completed and used successfully in preliminary studies with the SiPM array. The Kapton set of cables was nearly completed, but ultimately suffered from connectivity issues due to recessed MMCX pins. Stripping and preparing the cable was an intensive process of trial and error that is detailed in Appendix A for posterity and is broadly relevant to the creation of the FEP set of cables.

#### 4.4.3.4. Cable Routing

The detector space is very compact. Despite significant space savings from the SiPMs, routing the new cables was not trivial. Kapton coaxial cables are especially stiff and impart destructive levels of strain on joints. FEP cables are much more flexible, but the sheer number of cables ultimately limits the bend radius of any cable bundle. Cables connecting to the bottom PMT also need to make a 90 degree bend beneath the bottom PMT holder without causing any pins to come loose. Three improvements significantly lessened the strain on all of the cables. First, the top mounting bracket was redesigned with a profile that had a nearly 50% reduction in area (Figure 4.44). This left more space for cables to pass through, and exposed more of the recessed 1.33" CF ports on the underside of the detector top flange.

For cables connecting to the bottom PMT, a 90 degree bend was imparted using a small segment of 3/16" copper tube. This tube was bent using a tube bender with a 7/16" bend radius—almost double the minimum bend radius of the Kapton cable. The copper “conduit” was slid onto



- (a) The original top mounting bracket, seen from below. Excess material on either side of the support rod holes (holes on the outermost radius) blocked access to 1.33" CF flanges inside of the detector space.
- (b) The new top mounting bracket removed large amounts of excess material, easing the cable routing process.

Figure 4.44.: Comparison of top mounting bracket designs used to suspend components in the XeNeu detector space. The new bracket (Figure 4.44b) has a profile which occupies almost 50% less space than the old bracket (Figure 4.44a).

the cable before the PMT pins were soldered to the end of the cable. The final major improvement consisted of tightly coiling annealed steel wire around the cables. The wire was used for affixing copper tags to the cables which were engraved with identifying labels. Then, unlike the cables, the steel wire would hold its shape after being bent. This meant that the cables could be bent into more favorable positions which removed almost all strain from key areas like the 23-pin Sub-C connector.

#### 4.4.4. SPE Gain Calibration

Calibration of the SiPM array began with a determination of the SPE response. As this was the first time this array had been operated, it was important to determine the minimum operating voltage of the SiPMs in order to resolve SPEs while minimizing sources of noise like afterpulsing and dark counts. For this calibration, background scintillation data was taken using a single-coincidence trigger on the bottom PMT (Figure 4.45).

This trigger configuration was chosen because the bottom PMT typically sees the largest fraction of S1 light in the TPC due to total internal reflection of light at the liquid surface. In each dataset, the bottom PMT was biased to 1060 V, while the SiPMs were swept from 50-52 V in quarter volt steps. The bottom PMT was powered with a CAEN 1470 power supply, while each of the three SiPM power banks were powered with a CAEN R8031. Signals from the SiPMs (Figure 4.46) were filtered with 5 MHz low-pass inline filters (MiniCircuits BLP 5+) before being fed into a 10x amplifier (PS 776), and then into the digitizer. Signals from the bottom PMT were also amplified, but not filtered.

Pulse finding required knowledge of the SPE response, so data on SPEs was collected by

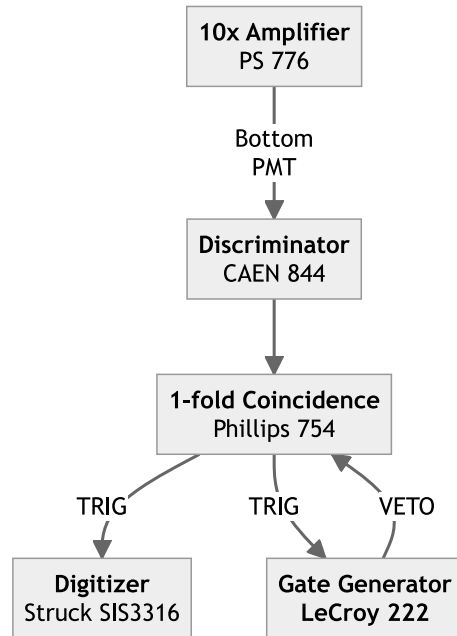


Figure 4.45.: Trigger configuration for SPE calibrations. When the bottom PMT (CH10) crosses a threshold set at the discriminator, the discriminator fires a pulse which is picked up by the coincidence module. If the coincidence module is not being vetoed, it fires a pulse (TRIG) to the trigger channel on the DAQ and to the gate generator. The gate generator vetoes the coincidence module for a set period of time after receiving a pulse from the coincidence module.




Figure 4.46.: Example waveform of an SPE captured during preliminary SiPM tests at the time indicated with the orange “T”. The SiPM was biased to 55 V and the signal was filtered with a 5 MHz low-pass filter. A voltage and time scale of 2.00 mV and 2  $\mu$ s per division were used; these are printed at the top and bottom of the figure, respectively.

manually specifying a Region of Interest (ROI) in an event window. Because `daqman` treated ROIs like pulses and generated pulse quantities for each ROI, they could be analyzed like pulses without assuming any prior knowledge of pulse characteristics.

An example of an ROI integral spectrum from one SiPM channel at a specific bias voltage is shown in Figure 4.47. The peak near 0 is from noise captured in the ROIs, while the next peak near 500 corresponds to SPEs in the ROIs. The next two peaks correspond to double and triple photoelectron pulses. The relative height of the peaks in the spectrum suggests that double and triple photoelectron rates were nearly equal to SPEs rates. This is somewhat counter-intuitive, but does not appear to be unique; see also [127]. This spectral shape could result from a complex interplay between the detection light spectrum, SiPM response, and low detection statistics. In a subsequent run with higher statistics, the noise pedestal dominates, followed by the SPE, double,




Figure 4.47.: Example of an ROI integral spectrum from a single SiPM at a single voltage. The sum of four gaussians was fit to the spectrum. This fitting procedure was repeated for each SiPM channel at each bias voltage in order to determine the minimum voltage needed to resolve SPEs.

and triple photoelectron peaks as expected.

The sum of four independent gaussians was fit to the spectrum to obtain an estimate of the SPE size. This could be refined in a future work with fewer free parameters, but was sufficient for this analysis. The fitting procedure was repeated for each SiPM channel at each bias voltage. In order to automate the fitting process, the first guess for each fit was determined using the `scipy.signal.find_peak` function on histograms with coarse binning to accentuate the approximate location of each peak. The SPE mean was estimated by taking the arithmetic mean of the single, double, and triple photoelectron peaks:

$$\mu_{overall} = (\mu_{SPE} + \frac{\mu_{DPE}}{2} + \frac{\mu_{TPE}}{3}) \div 3$$

The gain curves for each channel are shown in Figure 4.48. They are fairly well-behaved above 50.75 V, below which the resolution in all channels begins to falter. Based on these curves, it was determined that a bias voltage of 51 V was acceptable for use in all channels. The close grouping of SPE values across all channels is also evidence that individual channels were gain-matched to a satisfactory extent considering they were biased in groups of three.

#### 4.4.5. Position Calibration

The final tests accomplished in this work determined the position reconstruction capabilities of the SiPM array. In Run 16, data was taken with XeNeu using  $^{133}\text{Ba}$  and  $^{137}\text{Cs}$  sources, which were placed at varying positions around the detector. A coordinate system (Figure 4.49) was developed to take advantage of a self-leveling cross-line laser.

A pulse-width based trigger on the bottom PMT was used to capture S2-like pulses (Fig-

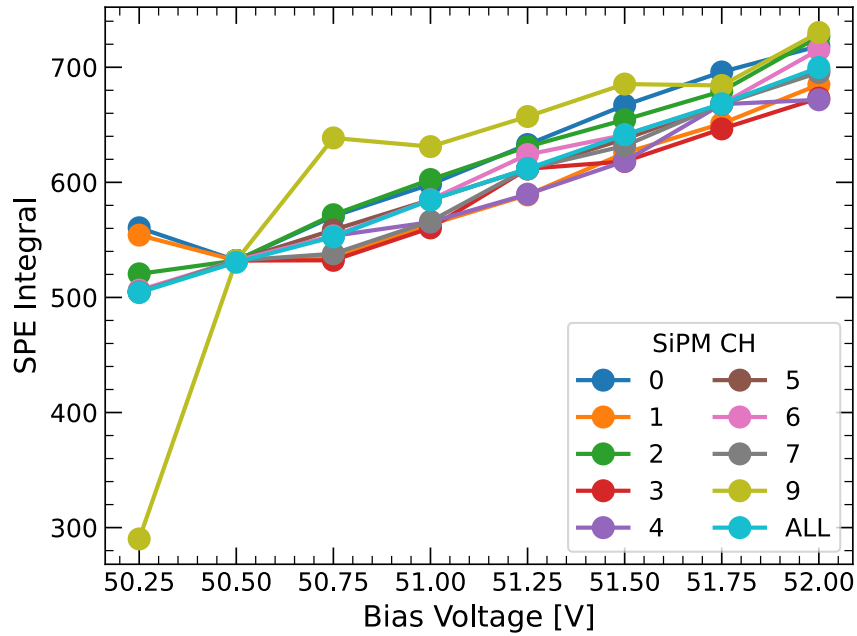


Figure 4.48.: Gain curves for each SiPM channel. “ALL” represents the gain curve obtained after fitting to the sum of the integral spectra for all channels. Note that CH 8 is missing: this channel was broken on the DAQ so CH 9 was used instead.




Figure 4.49.: Coordinate system used for positioning sources around XeNeu.  $R$ ,  $Z$ , and  $d_{laser}$  were measured with a standard measuring tape or calipers, where possible.

ure 4.50). Signals from the bottom PMT were amplified, before being passed into the first of two fan-in fan-out (FIFO) modules. This module was used to duplicate the bottom PMT signal, which was then passed into a second FIFO module to sum the two signals together and effectively double the amplitude of the PMT signal. This step was necessary to take advantage of the dynamic range of the first discriminator, which could only be adjusted in 1 mV increments.

A crude pulse-width trigger was established by requiring a discriminator to be coincident with its own “echo”. This was accomplished by feeding the output from Discriminator 1 into a 2-fold coincidence module and Discriminator 2. The output from Discriminator 2 would be delayed by the delay generator. If Discriminator 1 was still firing by the time the delayed output reached the 2-fold coincidence, then an event would be triggered. This scheme was devised so that a pulse which was at least as wide as the delay would fire the trigger, but other sequences of pulses would not. Discriminator 2 would also briefly veto the 2-fold coincidence module to mitigate accidental coincidences between Discriminator 1 and the output of the delay generator. To trigger on SEs, the delay was set such that the rising edge of Discriminator 1 and the falling edge of Discriminator 2 were approximately 300 ns apart. This configuration was able to achieve a reasonable 80% trigger efficiency on single electron pulses. One background captured by this trigger was accidental coincidences between pairs of single or double photoelectrons (DPE). These were hard to avoid since many SE pulses were a loose collection of these same SPE and DPE pulses, especially at lower extraction fields.

A preliminary analysis was carried out to demonstrate the SiPM array’s capabilities. A example waveform in Figure 4.51 shows the summed output from the bottom PMT and SiPM array, with two SE pulses.

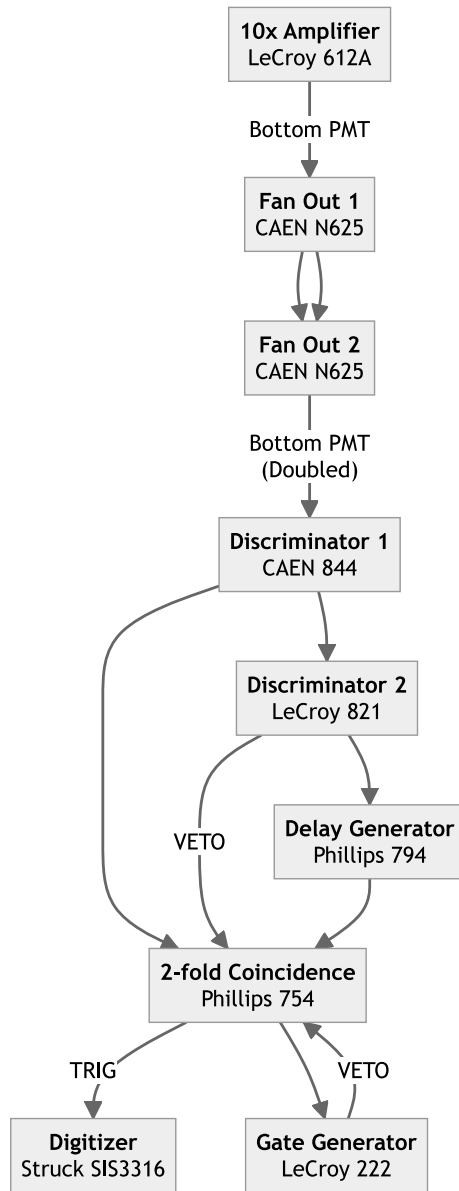


Figure 4.50.: Trigger for single electrons using the bottom PMT. The bottom PMT signal was doubled using two FIFO modules to make use of the full range of the discriminator. The main event trigger was configured such that Discriminator 1 would need to be coincident with a delayed copy of itself. This would ideally happen for a pulse with a width greater than the delay time.




Figure 4.51.: Example waveform showing the sum of the bottom PMT and SiPM outputs, with two SE pulses.

Figure 4.52 compares reconstructed areas of electroluminescence pulses for the bottom PMT, and the summed output of the SiPMs. The new SiPM array had slightly less photocathode coverage compared to the bottom PMT, but subtends a larger solid angle than the bottom PMT when collecting light from these signals. The plot confirms that the SiPM array can resolve SE pulses almost as well as the bottom PMT, with a mean SE area of approximately 40 npe compared to the 55 npe from the bottom PMT.

The effectiveness of two different position reconstruction algorithms was compared. Previous studies with the PMT array [71,73] used a center-of-mass (CoM) approach to calculate positions. This algorithm is relatively simple, relying on a weighted average of the light collected by each pixel [124]. One drawback of this algorithm is that it will reconstruct pulses at the edge of the TPC towards the center [71,124]. Pulses at the TPC edges can be excluded [71] with cuts, or corrections




Figure 4.52.: Comparison of reconstructed pulse areas for electroluminescence signals in the SiPM array vs. the bottom PMT. SE pulses are prominently featured in this plot between 20 to 80 npe on the PMT axis.

can be applied to map pulses to a circular coordinate system [124]. In either case however, the fiducial volume will be limited by the cuts used or nonlinearity of the correction.

Figure 4.53 shows SE pulse positions from the SiPM array using a CoM algorithm. The distribution of pulses is unrealistic: the TPC is circular, yet the extent of pulse positions appears to be a square. Furthermore, the maximum extent of pulse positions should be almost 3 cm in any given direction, not 1 cm; this shows how events at the edges of the TPC are mapped back to the center by the CoM algorithm. Figure 4.54 confirms this feature, where a simulation of pulses with the CoM algorithm has a maximum extent of 1 cm in either direction, compared to the true extent of 3 cm.

More complex position reconstruction algorithms use light response functions (LRFs) for each pixel in the detector. LRFs are distributions relating the distance of a light source to the fraction




Figure 4.53.: Reconstructed pulse positions using the CoM approach. The square distribution is unrealistic and a consequence of pulses at the edge of the TPC being misreconstructed towards the center.




Figure 4.54.: Comparison of positions generated by the CoM algorithm and true positions from the simulation. The CoM algorithm is accurate for pulses with a true position within 1 cm of the TPC center, but misreconstructs positions thereafter.

of light seen by a given pixel, i.e. the value of an LRF is always between 0 and 1. LRFs can be 1-dimensional, however a 2-D LRF will be more suitable for pixels near walls where the radial symmetry is broken. Multiple methods can be used to determine pulse positions with LRFs. For example in the case of LUX, the MERCURY algorithm maximized the likelihood that a set of PMTs observed a given number of photons from a pulse at a location in the TPC [132]. DarkSide-50 meanwhile, used LRFs to minimize a chi-squared value between the light detected in a PMT  $M_i$ , and the product of the LRF  $L_i$  and total amount of light detected  $M_{tot}$  [133]:

$$\chi^2 = \sum_i^{\text{PMTs}} \frac{1}{M_i^2} (M_i - L_i(x, y)M_{tot})^2$$

In smaller detectors, LRFs can be generated<sup>8</sup> using radioactive sources at known positions. However in larger detectors, central locations of the TPC might not be easily accessible in part due to self-shielding of the detection medium [132]. In this case, LRFs may be generated iteratively, using simulated data [133] or a corrected CoM algorithm as a seed [134]. In this study, initial LRFs were generated using simulated data.

Figure 4.55 shows the reconstructed pulse positions using algorithm similar to MERCURY. The “rings” with a radius of  $\sim 1$  cm are an artifact of a discontinuity in the LRFs at a radius of 1 cm. The plot shows pulses at the outer reaches of the TPC, including a nominally circular central distribution with a radius of approximately 2.8 cm. This corresponds almost exactly to the radius of the anode grid spacer on the underside of the SiPM array holder Figure 4.38.

The distribution of pulses in Figure 4.55 appears fairly uniform. The positions of pulses found using LRFs are compared to positions found using the CoM algorithm in Figure 4.56. The result

---

<sup>8</sup>Generating a set of LRFs is a complex process which is summarized here for the sake of brevity. The interested reader is encouraged to pursue the provided sources for more information.




Figure 4.55.: Reconstructed pulse positions using the LRF.

strongly resembles Figure 4.54, indicating that pulse positions found with LRFs closely match the true positions of the pulses.




Figure 4.56.: Comparison of positions generated by CoM and LRF.

## Chapter 5.: Towards Lower Detection Thresholds with Xenon-Doped Argon TPCs

Utilizing a parts-per-million (ppm) level of xenon dopant in the gas region of a dual-phase argon TPC presents the enticing prospect of enhanced electroluminescence, improving low-energy particle detection with argon. Compared to pure xenon, argon's lower mass makes it a more favorable scattering target for light particles such as neutrinos and sub-GeV dark matter candidates. Furthermore, the lower operating temperature of argon can result in reduced outgassing from components, in turn contributing to lower delayed electron background rates. However, xenon-doped argon's potential as a detection medium is hampered by several drawbacks. At low temperatures, xenon is known to remain in the liquid bulk due to its low vapor pressure, and it is liable to deposit on surfaces when exposed to cold components in the gas phase. This further complicates circulation, which must maintain a stable mixture and composition without disturbing equilibrium between the liquid and gas phases.

The possibility of realizing the benefits from both argon and xenon compelled the development of Coherent Ionization Limits in Liquid Argon and Xenon (CHILLAX), a new dual-phase xenon-doped liquid argon test stand. In 2023, CHILLAX demonstrated stable operation of a liquid argon mixture with a xenon liquid mole fraction of 2.35% and a xenon vapor mole fraction in the desired

range of 30-50 ppm [135]. This chapter will discuss the known properties of a xenon-argon liquid mixture (Section 5.2), the design of the cooling system (Section 5.3), and the campaign to develop a gas sampling system to measure the xenon vapor concentration in CHILLAX (Section 5.5). An evaluation of the liquid argon photon detection efficiency (PDE) in new Hamamatsu VUV4 SiPMs was also conducted during preliminary cryogenic tests with CHILLAX [131]. This evaluation was concurrent with this work, and is briefly summarized in Section 5.4.

### **i** Note on Concentrations

Doping concentrations in this work are given as mole fractions, i.e. ratios of xenon atoms per total number of atoms:

$$\chi_{Xe} = \frac{n_{Xe}}{n_{tot}} = \frac{n_{Xe}}{n_{Xe} + n_{Ar}} = \frac{1}{1 + \frac{n_{Ar}}{n_{Xe}}} \quad (5.1)$$

Molar masses can be used to convert ratios of masses to molar masses:

$$\frac{m_{Xe} \frac{1}{131.3 \text{ g/mol}}}{m_{Ar} \frac{1}{39.99 \text{ g/mol}}} = \frac{n_{Xe}}{n_{Ar}}$$

This is useful for comparing to values reported in other works, where concentrations may be denoted as a mass of xenon in the total mass of a mixture. Unfortunately, most works do not seem to specify the convention used.



## 5.1. Benefits and Challenges of Xenon Doped Argon TPCs

Compared to xenon, argon's lower mass makes it a more favorable scattering target for light particles such as neutrinos and sub-GeV dark matter candidates (Section 2.3). Furthermore, in a

dual-phase argon TPC the small ionization signals generated from these interactions are in principle easier to detect. One reason is that lower operating temperatures result in reduced outgassing from components, in turn contributing to lower delayed electron background rates. In addition, argon's smaller dipole moment leads to an enhanced electron extraction efficiency compared to xenon's at the same fields [136].

However, argon's potential as a detection medium is hampered by several drawbacks. First, unlike the 172 nm light from xenon dimers, argon emits 128 nm light which cannot be detected by PMTs with quartz windows. This results in a need for wavelength shifting coatings on detector surfaces. The coatings are difficult to distribute evenly, leading to a non-uniform position resolution which complicates the detection of small ionization signals [68]. Second, the microsecond-long decay time of the argon triplet state can also lead to dead time in larger detectors [137]. Small signals are also impacted by the triplet state lifetime; a few dozen photoelectrons from an SE distributed over 20  $\mu$ s will obscure low energy interactions [138].

Doping a dual-phase argon TPC with xenon can mitigate both problems because of xenon's lower excitation energy. In a manner similar to quenching argon scintillation light with impurities, xenon excimers are produced via an intermediate argon-xenon exciplex<sup>1</sup> [140]:




---

<sup>1</sup>An exciplex is an *excited complex* consisting of heterogeneous components, whereas an excimer is entirely homogeneous [139]. That is, an excimer is a special type of exciplex.

Due to the preferential production of xenon excimers, 174 nm scintillation light is predominant at concentrations as low as 10 ppm [141] (Figure 5.1), obviating the need for other wavelength shifting compounds. Furthermore, because xenon excimers de-excite orders of magnitude faster than the argon triplet (Table 2.1), a strong reduction in the event pile-up rate is expected. Finally in addition to these effects, Penning ionization has been shown to increase charge yields by roughly 10-15% [142] [143]:





Figure 5.1.: Emission spectra from xenon-doped argon gas mixtures generated via electron beam excitation. Light production from xenon excimers centered around 174 nm virtually saturates at concentrations as low as 10 ppm. From [141].

As in the liquid, light production in the gas phase from xenon and xenon-containing excimers is preferential because of xenon's lower excitation energy. The resulting spectral shift is peaked strongly at 147 nm (Figure 5.2), corresponding to transitions from the xenon  $^3P_1$  excited state

[144] [145]. This state has a lifetime of less than 4 ns [146], similar to that of xenon excimers. The argon-xenon exciplex also radiates at 150 nm [147], however this transition occurs on the order of microseconds in the liquid [148], making its role in gas phase light production questionable.


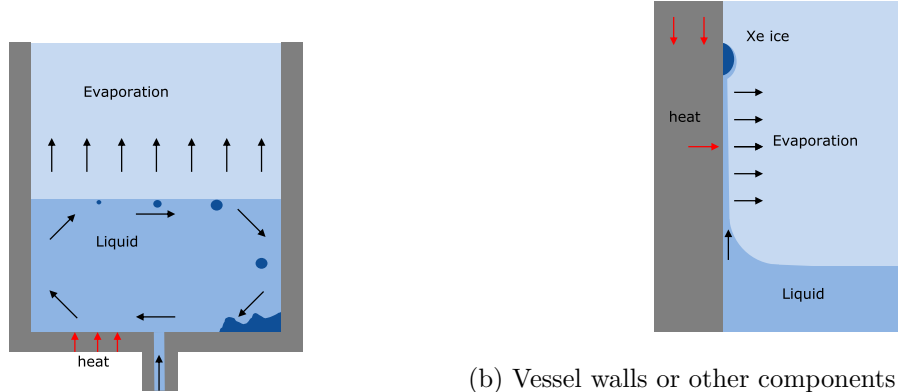


Figure 5.2.: Emission spectra from xenon-doped argon gas mixtures generated via heavy ion beam excitation. Subplots are arranged in columns of the same xenon concentrations and rows of the same total mixture pressures. Rows *a* through *e* correspond to total mixture pressures of 400, 780, 1000, 1200, and 1400 mbar respectively. Efthimiopoulos *et al.* note that emission at concentrations of 100 ppm is predominantly from atomic transitions. From [145].

While these benefits are promising, the addition of xenon presents new technical challenges. The Gibbs phase rule stipulates that the additional degree of freedom can allow for up to three phases to be in equilibrium with each other across a range of either temperatures or pressures at any given doping concentration. In a dual-phase xenon-doped argon TPC, the formation of solid xenon can result in the depletion of xenon from the active volume or even clog circulation systems [149].

Solid xenon can form when a xenon-argon liquid mixture evaporates because argon's high volatility distills xenon in the remaining liquid (Section 5.2.2). Evaporation in a TPC can happen anywhere along the liquid surface (Figure 5.3a), especially at the juncture between the liquid and warm components immersed in the liquid (Figure 5.3b). In the latter case, the thin film that forms

due to the surface tension of the liquid has a high heat input relative to its volume. When this liquid evaporates, xenon is deposited, the liquid film replenishes, and xenon accumulates as the process repeats. Xenon ice can also form when a gas mixture is cooled. This is most pertinent in circulation schemes where a cold tube may form a xenon ice blockage [149], however it could also lead to the formation of xenon frost over other cold components.



(a) A vessel which is evaporating xenon from the liquid mixture (dark blue) surface can produce solid xenon (darkest blue) when the xenon vapor pressure from the liquid mixture exceeds the vapor pressure of solid xenon in the same conditions.

(b) Vessel walls or other components immersed in the liquid mixture will be coated with a thin film because of the liquid's surface tension. If these components are warmer than the mixture, the film will evaporate, leaving behind solid xenon. This process can repeat, resulting in the accumulation of solid xenon and corresponding depletion of xenon from the liquid.

Figure 5.3.: Examples of solid xenon formation resulting from excessive distillation of xenon in a xenon-doped liquid argon mixture.

## 5.2. Operating Conditions for a Xenon-Doped Liquid Argon Mixture

At the start of the project, a feasibility study was conducted to produce rough estimates of possible operating temperatures and pressures that would:

1. Avoid unintended xenon crystallization and deposition
2. Result in 10s of ppm of xenon gas above the liquid surface

This study encompassed the topics of eutectic diagrams and predicted xenon gas concentrations above a xenon-doped liquid argon mixture. Ultimately, a xenon vapor mole fraction of 30-50 ppm was predicted with a liquid mole fraction of around 2% xenon. This was anticipated to be possible across a range of operating temperatures between 90-95 K with total pressures under 2 bar.

### 5.2.1. Eutectic Point Estimate

A binary eutectic mixture is composed of two substances in quantities such that the freezing point is lower than either component. This lowest freezing temperature of the mixture is known as the eutectic point. Having some understanding of where this point could be for a mixture of xenon and argon is helpful for establishing operating limits for a future xenon-doped liquid argon TPC.

The eutectic point can be found by generating a plot of mixture composition vs. temperature at some fixed pressure. On the plot are solid-liquid equilibrium curves which describe the freezing point of the mixture as a function of its composition. If the curves intersect, then the intersection will be at the eutectic point. Robinson [150] gives a brief overview of this process for cryogenic refrigerant mixtures. This section will attempt to follow his approach to create one of these diagrams for a mixture of argon and xenon.

Robinson ultimately cites Sandler's [151] method for finding the temperature vs. composition curves. In section 12.3 of Sandler, equation 12.3-2 is needed for both argon and xenon:

$$\ln(\chi_1\gamma_1) = -\frac{\Delta H_f}{R} \left[ \frac{T_m - T_f}{T_m T_f} \right] - \frac{\Delta C_p}{R} \left[ 1 - \frac{T_m}{T_f} + \ln\left(\frac{T_m}{T_f}\right) \right] \quad (5.2)$$

Nomenclature for this equation follows:

- $\chi_1$  is the mole fraction of substance 1 in a mixture
- $R$  is the universal gas constant
- $T_m$  is the melting point of the pure substance
- $T_f$  is the *new* freezing temperature of substance 1 (to be calculated)
- $\Delta H_f$  is the molar enthalpy of fusion of substance 1 at  $T_m$
- $\Delta C_p$  is the difference in heat capacity for the solid and liquid phases of substance 1
- $\gamma_1$  is the activity parameter which is equal to 1 in the case of an ideal mixture

There are a few assumptions that are made in this equation:

- $T_m$ ,  $\Delta H_f$ ,  $\Delta C_p$ , and  $R$  can be estimated or found elsewhere
- $\Delta C_p$  is independent of temperature, or assumed to be zero

What follows is an attempt to understand Robinson's approach and his path from Sandler.

Robinson's expression for the temperature-composition curve is seen below:

$$\ln \chi_1 = -\frac{\Delta H_f}{R} \left[ \frac{T_m - T_f}{T_m T_f} \right] - \frac{V_1 \phi_2^2}{RT_f} (\delta_1 - \delta_2)^2 \quad (5.3)$$

$$\phi_2 = \frac{(1 - \chi_1)V_1}{(1 - \chi_1)V_1 + \chi_1 V_2}$$

Where  $V_i$  is the molar volume, and  $\delta_i$  is a physical quantity known as the solubility parameter.

Sandler claims that a constant liquid molar volume can be used if thermal expansion is neglected (see Eq 12.1-10 in Sandler). Sandler also notes that the solubility parameter can be calculated as follows:

$$\delta = \left( \frac{\Delta U_v}{V} \right)^{1/2}$$

Where  $\Delta U_v$  is the is the internal energy of vaporization, and  $V$  is again the molar volume.

Data on solubility parameters seems limited [152] to temperatures far in excess of the predicted operating temperatures of CHILLAX. Likewise, the data required to calculate solubility parameters is quite sparse. Eq 12.1-9 from Sandler can be substituted into Equation 5.3 to make Robinson's approach more manageable:

$$RT_f \ln \gamma_1 = V_1 \phi_2^2 (\delta_1 - \delta_2)^2 \quad (5.4)$$

Doing so produces an equation identical to Sandler's 12.3-2 in the case where  $\Delta C_p$  is zero. Robinson rearranges Equation 5.3 into the following<sup>2</sup>:

$$T_f = \frac{\Delta H_f + V_1 \phi_2^2 (\delta_1 - \delta_2)^2}{\frac{\Delta H_f}{T_m} - R \ln \chi_1}$$

Substituting in Equation 5.4 and setting  $\gamma_1 = 1$  gives:

$$T_f = \frac{\Delta H_f T_m}{\Delta H_f - RT_m \ln \chi_1} \quad (5.5)$$

Robinson uses values of  $\Delta H_f$  and  $T_m$  measured at atmospheric pressure for eutectic point calculations of cryogenic refrigerant mixtures; this approach is also taken here. Equation 5.5 is plotted in Figure 5.4 for argon and xenon. A eutectic point is predicted at 76 K with a 15% xenon mole fraction.

---

<sup>2</sup>There appears to be a typo on  $T_m$  as  $T_f$  in Robinson's paper.




Figure 5.4.: Plot of temperature composition curves for xenon and argon (Equation 5.5) at atmospheric pressure. This plot shows the eutectic point of the binary mixture, which is the intersection point between the two curves.

#### **i** Note on Pressure Dependence

$T_m$  is largely invariant with pressure between 1 to 10 bar for both Argon and Xenon, so taking into account a possible pressure dependence might not be warranted. This is likely why Robinson used values of  $T_m$  and  $\Delta H_f$  at atmospheric pressure. Nevertheless, it is worth commenting on the relation between  $T_m$  and  $\Delta H_f$  for the sake of a reader interested in extending this work.

$\Delta H_f$  can be found from the Clausius-Clapeyron equation [153]:

$$\ln P = -\frac{\Delta H_f}{RT} + \text{const.}$$

which can be differentiated to obtain  $\Delta H_f$ :

$$\Delta H_f = -R \left( \frac{d \ln(P)}{d(1/T)} \right)$$

Both  $\Delta H_f$  and  $T_m$  are found with the fusion (solid-liquid) curve: the pressure and temperature curve that separates the liquid and solid phases of a substance on a pressure-temperature diagram. Whereas the Antoine equation [154] is used to describe the vapor pressure of a pure substance, the similarly empirical Simon-Glatzel equation [155] exists to describe the fusion curve:

$$P_m = a[(T_m/T_0)^c - 1] + P_0$$

Here,  $P_m$  and  $T_m$  are the melting point temperature and pressure.  $P_0$  and  $T_0$  are a reference point, usually the triple point, while  $a$  and  $c$  are empirically derived constants. Ferreira and Lobo [156] examined experimental data, and determined the Simon equation parameters for rare heavy gases, including argon and xenon.

This prediction is in loose agreement with experimental results from Heastie and Lefebvre [157], who found a eutectic point of approximately 82.5 K at a xenon mole fraction of 23% (Figure 5.5) and pressure of 556 mBar. In a much more detailed theoretical study, Campestrini *et al.* also calculated the phase-equilibria of a xenon-argon mixture over a temperature range of 74-94 K [158]. A eutectic point was predicted at a temperature near 80 K and xenon mole fraction of 16%, with the caveat that the behavior documented by Heastie and Lefebvre was not well-modeled. In a further disagreement, experimental data from Yunker and Halsey [159] shows solid xenon forming at liquid mole fractions between 3-5% with temperatures in the range of 84-87 K (Figure 5.6). This

discordant behavior between the modeled and experimental results has motivated contemporary phase equilibrium studies of argon-xenon mixtures [160].




Figure 5.5.: Empirical temperature vs. composition plot reproduced from [157]. Labels on the plot designate regions where solid, liquid, and vapor phases are in equilibrium with each other.

### 5.2.2. Predicted Xenon Gas Concentrations

Sufficient doping of the gas phase is important to ensure a high detection efficiency of ionization signals. Figure 5.2 shows that a minimum of 30-50 ppm in the gas phase is necessary to shift electroluminescence to 147 nm light produced by xenon. One factor that complicates this is the drastically lower vapor pressure of xenon compared to argon at the same temperature. Figure 5.7 shows vapor pressure curves for both argon and xenon from a number sources. Data from this figure consists of fits to empirical data or raw empirical data. An exception to this are the curves




Figure 5.6.: Theoretical temperature vs. composition plot of a xenon-argon mixture from [160]. Phase boundaries (solid lines) indicate predicted solid (S), liquid (L), and vapor (V) phases of xenon (2) and argon (1) in equilibrium (E) in the mixture. Data from Heastie and Lefebvre (Figure 5.5) are shown with red squares while data from Yunker and Halsey (Figure 5.8) are shown as purple circles.

from Ferreira and Lobo [156] who present a calculation derived from the Claperyron equation. This work uses the argon vapor pressures measured by Gilgen *et al.* [161] due to their agreement with other sources and on the recommendation of [162]. Xenon pressures from Leming and Pollack [163] are used because their measurement is the most conservative.

As a consequence of these disparate vapor pressures, stable operation of a TPC may be threatened because a liquid concentration near the saturation limit may be required to produce an adequate xenon vapor concentration. Yunker and Halsey’s [159] study reported the most conservative results for the saturation limit (Figure 5.8), so they are used here as a baseline assumption for subsequent calculations.

Figure 5.8 shows that at low xenon mole fractions in the liquid, the vapor pressure of argon above the mixture is reduced by a percentage equal to the xenon mole fraction:




Figure 5.7.: Argon and xenon vapor pressures in the temperature range of 85 to 110 K. Data sources: NIST Argon [164], Gilgen *et al.* [161], Ferreira and Lobo (F&L) [156], CRC [165], NIST Xenon [166], Leming and Pollack [163].

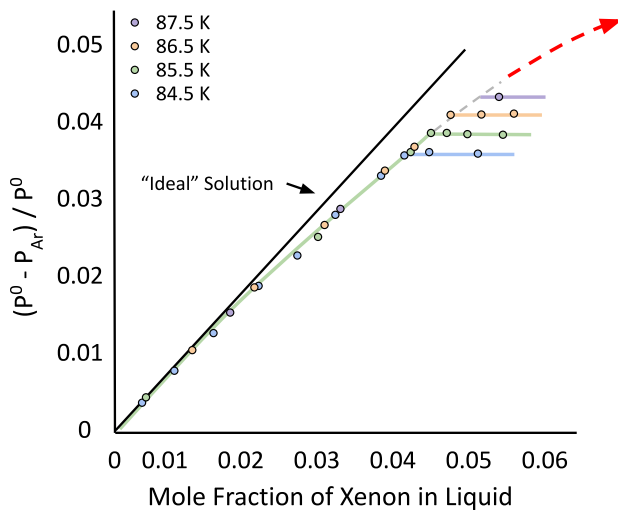


Figure 5.8.: Deviation of argon partial pressure  $P_{Ar}$  from pure argon pressure  $P^0$  as a function of the xenon liquid mole fraction and temperature. The solid black line indicates an “ideal” solution described by Raoult’s Law. Reproduced from [159].

$$\frac{P_{Ar}^0 - P_{Ar}}{P_{Ar}^0} = \chi_{Xe} \quad (5.6)$$

Where  $\chi_{Xe}$  is the xenon liquid mole fraction,  $P_{Ar}^0$  is the vapor pressure of pure argon, and  $P_{Ar}$  is the vapor pressure of argon over the liquid mixture at some temperature  $T$ . Note that because this is a two component mixture, the xenon and argon mole fractions are directly related:

$$\chi_{Xe} = 1 - \chi_{Ar} \quad (5.7)$$

Substituting Equation 5.7 into Equation 5.6 produces a restatement of Raoult's Law which predicts the vapor concentration of a mixture component given its liquid mole fraction and pure pressure:

$$P_{Ar}^0 \chi_{Ar} = P_{Ar} \quad (5.8)$$

At higher doping concentrations Figure 5.8 shows that the argon vapor pressure depression deviates from Raoult's Law, where it decreases *less* than Raoult's Law predicts. Yunker and Halsey determined the following modification to the argon chemical potential:

$$\mu_{Ar} - \mu_{Ar}^0 = RT \ln(1 - \chi_{Xe}) + \alpha \chi_{Xe}^2 \quad (5.9)$$

Where  $\alpha$  was experimentally found to be  $535 \pm 38$  cal./mole. This can be quickly converted to a modified version of Raoult's Law using Equation 7.8.6 from Devoe [167], which relates the deviation in chemical potential for an ideal gas to its pressure:

$$\mu_{Ar} - \mu_{Ar}^0 = RT \ln\left(\frac{P_{Ar}}{P_{Ar}^0}\right) \quad (5.10)$$

Substituting Equation 5.10 into Equation 5.9 gives:

$$\ln\left(\frac{P_{Ar}}{P_{Ar}^0}\right) = \ln(1 - \chi_{Xe}) + \frac{\beta\chi_{Xe}^2}{T} \quad (5.11)$$

where  $\beta = \alpha/R = 269.2 \text{ K}^3$ .

Eventually as enough xenon is added, the mixture saturates and nearly pure solid xenon forms. At this point, the argon vapor pressure stops decreasing, and by the same token the xenon vapor pressure stops increasing. This is because in the solubility limit, the vapor pressure of xenon above the liquid is essentially the sublimation pressure of solid xenon<sup>4</sup>:  $P_{Xe}^0$ . Raoult's Law (Equation 5.8) does not accurately predict this because it assumes the solution is ideal, i.e., that all components experience the same intermolecular forces. This is a good approximation for argon atoms, which usually only see other argon atoms. However, the opposite is true for xenon atoms, which have virtually no xenon nearest-neighbors. For extremely dilute components in a solution, Henry's Law describes the vapor pressures:

$$\chi_{Xe} = H^{cc}\phi_{Xe}$$

where  $\phi_{Xe}$  is the mole fraction in the vapor and  $H^{cc}$  is an empirical constant. Unfortunately, data on Henry's constant for this mixture is lacking, so Equation 5.11 is used instead to predict

---

<sup>3</sup>It is helpful to recall here that  $R$  is the product of Boltzmann's constant and Avogadro's constant, which has units of  $\frac{\text{Joule} \cdot \text{mole}}{\text{Kelvin}}$ .

<sup>4</sup>This neglects interactions between xenon and argon in the gas phase, which are negligible [168].

the xenon gas concentration. The xenon gas mole fraction is calculated by noting that it is equal to the pressure fraction in the vapor via the ideal gas law:

$$\phi_{Xe} = \frac{n_{Xe}}{n_{tot}} = \frac{P_{Xe}}{P_{tot}} \quad (5.12)$$

where  $P_{tot} = P_{Ar} + P_{Xe}$  is the total pressure. The result is shown in Figure 5.9a.

Obtaining a plot of the xenon liquid mole fraction as a function of the vapor mole fraction requires solving Equation 5.12 for the liquid mole fraction. If Equation 5.11 is used for both the xenon and argon pressures, this solution would need to be found numerically. Instead of doing this, a few approximations can be made. First, the xenon partial pressure is scaled by the saturation liquid mole fraction:

$$P_{Xe} = \frac{P_{Xe}^0 \chi_{Xe}}{\chi_{Xe}^{sat}} \quad (5.13)$$


This scaling reproduces the limiting behavior of the xenon partial pressure, which can be plugged back in to Equation 5.12. Yunker and Halsey measured  $\chi_{Xe}^{sat}$  as a function of temperature:

$$\ln \chi_{Xe}^{sat} = 1.6463 - 406K/T \quad (5.14)$$

Assuming Raoult's Law holds for  $P_{Ar}$  (Equation 5.8), Equation 5.12 can be solved for  $\chi_{Xe}$  as a function of  $\phi_{Xe}$  and temperature  $T$ :

$$\chi_{Xe} = \frac{\phi_{Xe}^s P_{Ar}^0}{P_{Ar}^0 \phi_{Xe}^s + P_{Xe}^0 (1 - \phi_{Xe}^s)}$$

where  $\phi_{Xe}^s = \phi_{Xe} \times \chi_{Xe}^{sat}$ . The results are plotted in Figure 5.9b. A xenon vapor mole fraction of 30-50 ppm is expected for a liquid mole fraction around 2%. The results in Figure 5.9b also agree to within 5% with those in Figure 5.9a for the region of interest between temperatures of 90-95 K and doping fractions between 2-3%. These results helped inform operating conditions for the first stability test of a xenon-doped liquid argon mixture [135].



(a) Predicted xenon vapor mole fraction as a function of the xenon liquid mole fraction and mixture temperature. (b) Predicted xenon liquid mole fraction as a function of an assumed vapor mole fraction and mixture temperature.

Figure 5.9.: Comparison between predicted xenon vapor mole fractions using Raoult's Law and empirical data from Yunker and Halsey [159]. Compared to Figure 5.9a, Figure 5.9b relies on an additional assumption that the partial pressure of xenon in the mixture should approach the vapor pressure of solid xenon when the mixture saturates (Equation 5.13). Points above the saturation mole fraction (Equation 5.14) are excluded.

### 5.3. CHILLAX Operational Concept and Thermosyphons

This section will present how CHILLAX was designed to avoid uncontrolled xenon accumulation. Additional details on the design of the cooling mechanism also presented. Readers interested in further information on CHILLAX may refer to Bernard *et al.* [135].

### 5.3.1. Theory of Operation

CHILLAX primarily consists of a main volume and a separate condenser to liquefy incoming gas. The condenser delivers liquid to the main volume via a lower tube, while an upper tube connects the gas spaces of both volumes to equalize their liquid levels. A block diagram of these components is shown in Figure 5.10. Prior to any doping, pure argon is condensed in the system and stable circulation is established. Xenon gas is then injected into the circulation path where it is thoroughly mixed in by the circulation pump (Senior Metal Bellows MB-601). This xenon-rich gas (~0.6% vapor mole fraction) is delivered to the condenser via an inlet tube heated to over 200 K (h1 in Figure 5.10) which ensures xenon can reach the condenser liquid without freezing prematurely. This also ensures that liquid in the condenser is continually mixed via convective and mechanical action from the heater and bubbling respectively. A fine steel mesh is present in the condenser volume to trap bubbles and improve their chance of condensing in the liquid.

Xenon accumulates in the liquid as a consequence of its solubility in argon and low vapor pressure (Section 5.2.2). This xenon-rich liquid flows towards the main volume because the liquid there is continually evaporated both by a heater (h3 in Figure 5.10) and pumping action from the circulation pump. Condensation in the condenser also serves to remove gas from the main volume via the upper gas interconnect. Once doping ceases, xenon concentrations are maintained in a steady state because gas evaporated from the main bath surface has the same concentration as gas returning to the condenser to be re-condensed.

Cooling power is provided by a lower (TSL) and upper (TSU) thermosyphon loop. Both thermosyphons originate on a condenser head connected a Sumitomo CH104 cryocooler, which provides 50 W of cooling power between 85-95 K. The design of the condenser and condenser head

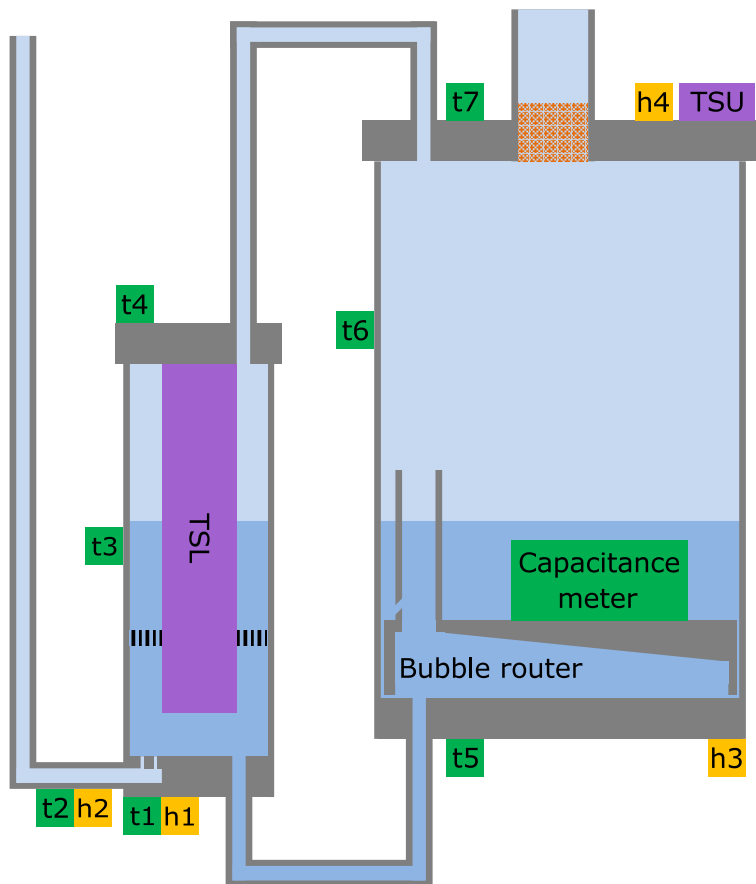


Figure 5.10.: Schematic showing key components of CHILLAX. The left volume is the condenser where incoming gas is condensed by TSL. The right volume is the main volume which houses the capacitance meter for measuring the liquid composition. The condenser and main volume are connected by a lower liquid transfer line and upper gas interconnect. Gas enters via an inlet tube heated by h2, and exits at the top of the main volume. Stainless steel mesh in the condenser (represented by vertical black lines) traps gas bubbles from the inlet tube to motivate condensation in the liquid. Copper mesh is present at the top of the detector volume to mitigate the heat load of cables present in the main volume. From [135].

are described in Section 5.3.2. The upper thermosyphon and heater (h4 in Figure 5.10) afford the possibility of regulating the main volume wall temperatures to avoid xenon accumulation depicted in Figure 5.3b. In the event that solid xenon does accumulate, the walls of the detector can be cooled to induce condensation of liquid argon and dissolve any xenon ice Figure 5.45.

Additional features in the detector space include a bubble router, capacitance meter, and camera. The bubble router is designed to funnel bubbles from heating (h3 in Figure 5.10) into the gas phase without disturbing the liquid surface. The capacitance meter measures the capacitance in the liquid, which changes with the composition of the liquid mixture. Finally, the camera is situated on one of the top flange ports, looking into the main volume. More details on the capacitance meter and camera are presented in [135]; both the camera and capacitance meter will be the subject of future work.

### 5.3.2. Thermosyphon Design

A thermosyphon consists of a closed system containing a working fluid which exchanges heat between two otherwise isolated thermal reservoirs. A dual-phase thermosyphon exploits a liquid-gas phase change in the working fluid to convey even more heat between the two reservoirs. The simplest dual-phase thermosyphon is a counter-flow design where the working fluid moves along the walls of the tube, while evaporated gas passes through a central column. This design is very economical and often used in heat pipes [169]. Typically liquid flows downwards with the assistance of gravity, but some counter-flow designs can operate with an inverted liquid flow where capillary action from sintered material, channels, or fine mesh on tube along the walls can carry liquid upwards [169]. The designs have an added benefit of preventing liquid entrainment in the gas

phase, which improves their heat transfer capability [170]. Alternatively, a loop design completely isolates the phases so that cold liquid flows down one leg of the loop and returns as warm gas in the other. Both of these concepts are illustrated in Figure 5.11.

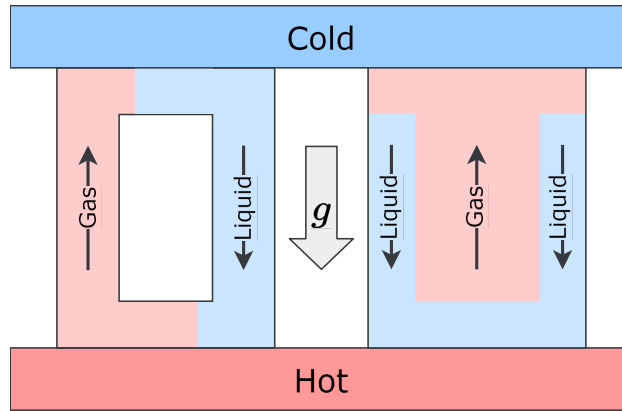


Figure 5.11.: Diagram illustrating the loop (left) and counter-flow (right) thermosyphon concepts. Gravity ( $g$ ) assists in moving liquid downward, but alternative designs exist which use sintered material, fine mesh, or small channels to move liquid against gravity via capillary action [169].

When designing a thermosyphon, certain characteristics influence the time constants of the system. The thermal mass of the components will determine whether the system is thermally stable, or more responsive. The contact areas between components and reservoirs will also play a role in determining this. The heat flux removed by a dual-phase thermosyphon is dependent on the boiling regime of the working fluid. Figure 5.12 shows this, where the heat flux is low at the onset of boiling but then begins to rise as new boiling regimes are entered. A local maximum in the heat flux occurs as the fluid undergoes nucleate boiling, before falling and reaching the Leidenfrost point where heat flux is *minimized*. This brief decrease occurs because optimal heat transfer relies on a balance between liquid being in contact with the wall and the liquid actually boiling.

This motivates adequate thermal controls for a thermosyphon through the use of heaters or a gas gap [91] to a modify the thermal conductivity at the point of contact. Experimental




Figure 5.12.: Pool boiling curve. Different boiling regimes occur as the temperature of a wall  $T_w$  deviates from the saturation temperature  $T_{sat}$  of a liquid. These boiling regimes convey different quantities of heat per unit of contact area. The Leidenfrost Point (C) is noteworthy for being a local minimum in the pool boiling curve, where an excessive temperature difference above the critical heat flux point (B) results in suboptimal heat transfer. From [171].

data on pool boiling curves for different working fluids is also critical for designing thermosyphon components which can convey the appropriate quantities of heat across a range of possible operating temperatures. Nitrogen or argon are common choices for working fluids at cryogenic temperatures due to their low cost. In CHILLAX, argon was chosen as the working fluid because it would ideally only require a small difference from the planned operating temperatures in the main volume.

Pool boiling curves for argon at various pressures are shown in Figure 5.13. Between 1 and 4 atmospheres the maximum heat flux is 20-30 W/cm<sup>2</sup> with a temperature difference of about 20 K. For designing the thermosyphon components, a conservative estimate of 2 W/cm<sup>2</sup> was chosen, assuming a temperature difference of 3-4 K. This was done to ensure that none of the components would be a limiting factor in delivering the 50 W of cooling power from the cryocooler.

The thermosyphon condenser head was machined from copper, and a steel funnel was brazed




Figure 5.13.: Pool boiling curves for Argon at different pressures showing the Peak Nucleate Boiling Flux (PNBF) at each pressure. At some pressures the heat flux was measured with a monotonically decreasing temperature difference instead of an increasing temperature difference to illustrate the effects of hysteresis. From [172].

on to it to create a hollow cavity for condensing the working fluid. Both parts are shown in Figure 5.14. The head has a diameter of 2.5” at its widest point, and the base of the head to the end of the funnel measures approximately 3”. The base of the condenser head attaches to the CH104 with eight screws; thermal grease was used to ensure a good thermal connection. Gas entered the condensing volume through one of the side ports on the funnel where it condensed on the exposed condenser head surface. The bottom of the condenser head was intentionally made concave to help guide condensing liquid to the outer edges of the funnel. The condensing surface area was maximized by milling holes into the bottom face of the condenser head in a hexagonal pattern. The walls of these cavities added 73 cm<sup>2</sup> of surface area to the exposed portion of the condenser head, more than tripling the 22 cm<sup>2</sup> from other exposed surfaces on the head.

The minimum surface area needed was based on a rough estimate following a derivation of the mean heat transfer coefficient  $\bar{h}$  for a vertical condensing surface found in Appendix A of Bradley’s thesis [170]. The heat flux  $q$  removed via condensation is proportional to  $\bar{h}$ :

$$q = \bar{h}\Delta T$$

Where  $\Delta T$  is the difference between the saturation and wall temperatures. Bradley estimated  $\bar{h} = 0.225 \text{ W}/(\text{K cm}^2)$  for liquid nitrogen, and it was assumed this value would be similar for argon. Using this value, it was estimated that the total exposed area of 95 cm<sup>2</sup> in the condenser could remove approximately 135 W of heat via condensation with a temperature difference of 6 K.

The condenser consists of two 2.75” CF half-nipples which are individually capped off and nested inside each other to form a condensing volume for xenon-doped argon gas mixtures (Figure 5.15). The lower half has gas inlet and liquid outlet ports for xenon-doped argon mixtures. The




Figure 5.14.: Rendering of the thermosyphon condenser head and funnel. A cross section of the funnel is shown to expose the interior of the condensing volume. In normal operation, gravity pulls liquid down the tapered end of the funnel, delivering LN to the inlet of the TSL. Gas returns from the TSL via one of two side ports on the funnel. A second gas return port accommodates gas returning from the TSU. The base of the condenser head is affixed to the bottom plate of the CH104 cryocooler using screws.

tube in the upper half contains the TSL, which was welded to the upper CF flange. The upper CF flange has inlet and outlet ports for delivering and removing argon from the TSL, as well as a third port which serves as the gas interconnect between the condenser and CHILLAX main volume.

The interior surface of the TSL had an area of approximately  $235 \text{ cm}^2$ . Vented helical aluminum inserts distribute liquid along the interior of the evaporator while maintaining some segregation between phases. The end cap alone had a surface area of roughly  $10 \text{ cm}^2$ , so only a fraction of the inner tube needed to be in contact with liquid to transmit the maximum possible cooling power into the condensation volume.




Figure 5.15.: Cross-section of the CHILLAX condenser. Gas from the circulation system enters the condensation volume at the bottom right and bubbles out through a brass plate. Once condensed, liquid exits to the CHILLAX main volume at the bottom. The TSL was welded to the top CF flange to isolate it from the condensing volume; the boundary of the TSL is indicated with a dashed line. Liquid from the thermosyphon condenser enters the TSL via an inlet (2) where it makes contact with the vented aluminum helical inserts. These inserts were designed to convey liquid along the walls of the TSL to provide more uniform cooling and maximize the contact area of the liquid with the TSL walls. Vents in the inserts provide a path for gas in the evaporator to exit via the outlet (1) without entraining incoming liquid. The upper gas interconnect (3) is used to equalize gas pressures between the condenser and main volume of CHILLAX (Figure 5.10).

## 5.4. Evaluation of Hamamatsu VUV4 SiPM Liquid Argon Scintillation

### Photon Detection Efficiency

Preliminary cryogenic tests with CHILLAX using pure argon led to the opportunity to evaluate the performance of new SiPMs [106] manufactured by Hamamatsu which were billed as having improved photon detection efficiencies (PDEs) in the regime of 128 nm argon scintillation light.

Two SiPM modules were evaluated: the S13371-6050CQ-02 (Figure 4.37), and S13370-6075CN (Figure 5.16). The S13371 is a module containing four active cells which can be read out individually, while each S13370 module contains a single cell. A major difference between these units is the inclusion of a quartz window on the S13371 which is opaque to light below 170 nm. Consequently, a second S13371 was included in the evaluation which had its window removed to improve its sensitivity to argon light. The windowed module was intentionally mounted opposite to the module with the window removed, to provide an *in situ* measurement of backgrounds with wavelengths greater than 170 nm.




Figure 5.16.: Rendering of a Hamamatsu S13370-6075CN.

A small testing cell (Figure 5.17) was constructed to isolate the SiPMs in the CHILLAX main volume. The cell was surrounded by a layer of Kapton film to absorb possible UV backgrounds

from fluorescence of other components in the main volume. It also improved the quality of optical models by reducing the possibility of reflections from argon scintillation light.

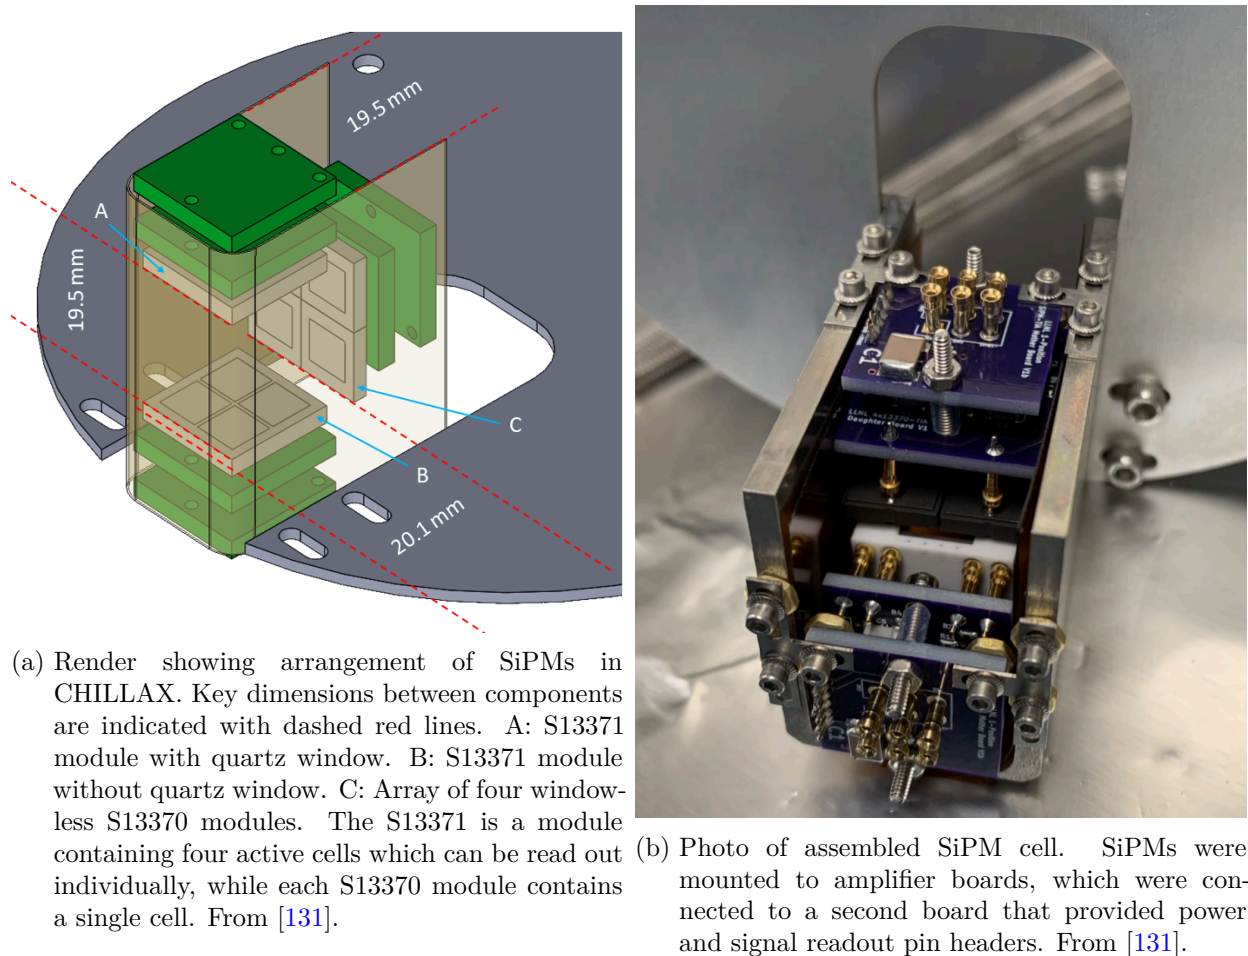



Figure 5.17.: Render and photos of SiPM cell. All three units are immersed in liquid argon during operation.

An  $^{241}\text{Am}$  source was used to generate events in the main volume. Average waveforms (Figure 5.18a) were analyzed to verify the presence of the argon scintillation light. A time constant of  $1447 \pm 20$  ns was measured for the triplet state, and the singlet-to-triplet intensity was found to be 0.29. Both of these values were in agreement with prior works, confirming that the detected light was due to argon scintillation. A pulse area spectrum (Figure 5.18b) also showed the attenuating

effects of the quartz window. In total, the windowed S13371 detected only 1.5% of the light seen by the windowless S13371, confirming that possible backgrounds were negligible.



- (a) A comparison of averaged waveforms shows a much more prominent prompt light component in the windowed module. This was taken to suggest the presence of longer wavelength light in the prompt component. The source of light at longer wavelengths is unclear. Possible causes include other wavelength shifting processes in the detector and fluorescence from detector components. From [131].
- (b) The windowed S13371 module significantly attenuated UV light, resulting in only pulses a few photoelectrons in size being visible. From [131].

Figure 5.18.: Plots showing averaged waveforms and pulse area spectra for windowed and windowless S13371 modules.

Gain curves, dark noise rates, afterpulsing, and cross-talk probabilities were also measured. Upper limits on dark noise (Figure 5.19) were measured by counting single-coincidence pulses in 12  $\mu\text{s}$  event windows. The possibility of collecting single photoelectrons (SPEs) from genuine argon scintillation could not be discounted, leading to higher upper limits in both windowless modules. Consequently, the rates from the windowed S13371 were expected to be more realistic.

SPE sizes in the SiPMs were initially evaluated by choosing a mean pulse area within a range that was slightly asymmetric about the statistical mean. This was done because afterpulsing was




Figure 5.19.: Upper bounds on SiPM dark count rates as a function of overvoltage. The breakdown voltage for both SiPM models was approximately 42.5 V. Figure modified from [131].

expected to contribute to a statistical mean SPE area which was slightly higher than the true mean SPE size. The effects of afterpulsing were estimated as the fractional difference between the manually chosen SPE means and the statistical SPE means (Figure 5.20a). The number of crosstalk pulses per SPE (Figure 5.20b) was estimated by scaling down the area of each SPE pulse by the ratio of the afterpulse-shifted and manually chosen mean SPE sizes. This result was then averaged across each SPE pulse. A subpopulation of SPEs more than 7  $\mu$ s after  $^{241}\text{Am}$  events was chosen for the crosstalk and afterpulsing analysis to reduce the probability of a coincidence between SiPM modules.

The PDE of a SiPM is function of its overvoltage  $V$  and the wavelength of the incident light  $\lambda$ :

$$PDE(\lambda, V) = \nu(\lambda)\eta(V)F \quad (5.15)$$

where  $\nu$  is the quantum efficiency of the silicon in each photodiode,  $\eta$  is the avalanche initiation

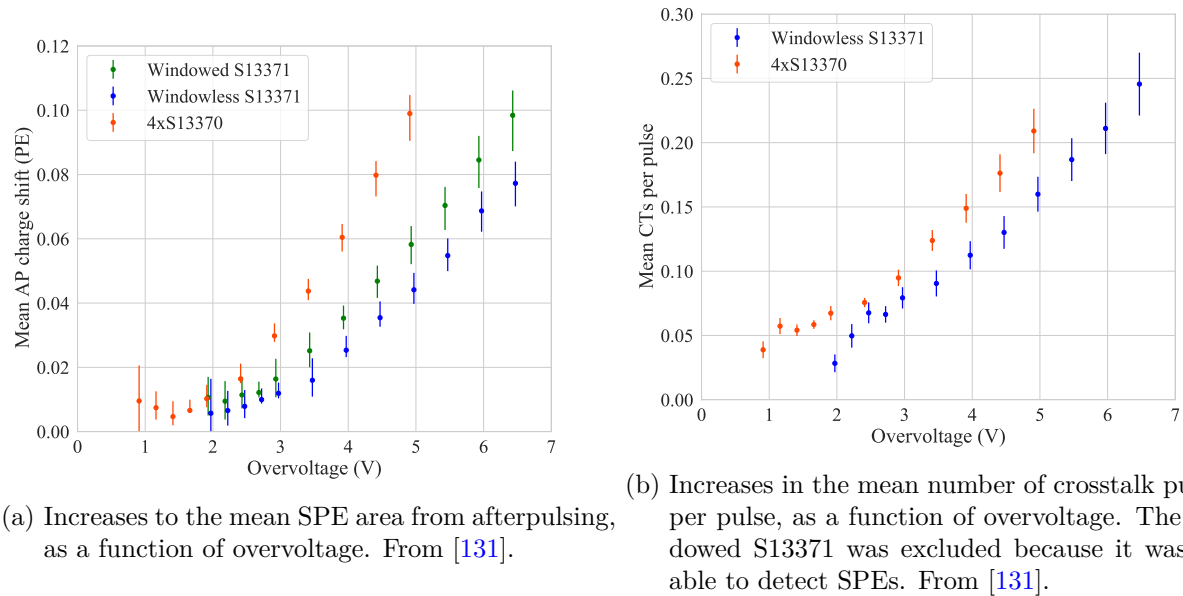


Figure 5.20.: Plots showing afterpulsing and crosstalk estimates.

probability, and  $F$  is the fill factor of the SiPM. The quantum efficiency is the probability that a photon incident on an active photodiode will create a pair of charge carriers. Not all pairs of charge carriers will lead to an avalanche, so the avalanche initiation probability accounts for this chance. Lastly, the fill factor accounts for dead space between active photodiodes (Figure 4.36a) [127].

To determine the PDE, optical simulations using Geant4 were first used to estimate the numbers of incident photons on SiPMs in each event. A binomial distribution with the PDE as a free parameter was then used to compute the number of photoelectrons produced in the SiPMs. Because the optical simulations could not account for crosstalk and afterpulsing, shifts in the SPE sizes from these sources of noise (Figure 5.20) were used to convert simulated pulse sizes to more realistic sizes. A simulated pulse area distribution was then compared to the measured pulse area distribution, and a chi-squared minimization was used to derive the PDE as a function of overvoltage (Figure 5.21).




Figure 5.21.: SiPM PDE as a function of overvoltage. The bands illustrate systematic uncertainties from varying cuts used to produce the PDE estimate. From [131].

At 3.97 V overvoltage, the S13371 measured PDE was  $17.2^{+1.6}_{-3.0}\%$ . At 3.91 V overvoltage, the S13370 measured PDE was  $14.7^{+1.1}_{-2.4}\%$ . Both values were consistent with those quoted by Hamamatsu at 4 V overvoltage operation [173], and amount to a roughly 25-40% increase in PDE in comparison to previous generation (VUV3) SiPMs from Hamamatsu. A more recent work [174] measuring the PDE across a range of wavelengths for S13370 SiPMs was also in agreement with these results.

## 5.5. The CHILLAX Gas Sampling System

The High Efficiency Residual gAs aNalysis nOble Sampler<sup>5</sup> (THERANOS) was a custom-built gas sampling system that used a Stanford Research Systems Residual Gas Analyzer (RGA) to determine the concentration of xenon in gas samples from CHILLAX. From its inception, THERANOS suffered from significant (20-50%) and often unpredictable measurement drift attributable to

<sup>5</sup>Named as an homage to the now defunct company which set high expectations for their sampling capabilities.

a number of factors like improper pressure reduction schemes and small sample volumes. Over the course of its commissioning, this issue was resolved to the extent that xenon concentrations of 10s of ppm could be reliably identified. This section will first review literature that proved to be crucial in understanding some of the pathologies THERANOS displayed. The most recent design iteration will be presented next, followed by a summary of the calibration campaign that was undertaken to characterize the system’s response to xenon-doped argon gas mixtures. This section concludes with data and methodologies from the CHILLAX stability study.

### 5.5.1. Introduction to Sampling Systems

Gas flows in a vacuum system can be characterized by the Knudsen number [175]:

$$K_n = \frac{\lambda}{d}$$

where  $\lambda$  is the mean free path of a gas constituent and  $d$  is the characteristic size of a container.

Three flow regimes are commonly defined [176,177], they are presented in Table 5.1 below:

Table 5.1.: Gas flow regimes as defined by their Knudsen number.

Flow Regime	Knudsen Number
Viscous	$K_n < 0.01$
Transition	$0.01 < K_n < 1$
Molecular	$K_n > 1$

For a gas undergoing viscous flow, the velocity of the gas is independent of its constituents. In molecular flow however, the velocities of constituents begin to dominate, and these are inversely

proportional to the square root of their mass  $m_i$ <sup>6</sup>. In transition flow, the dependence is more complex [176,178]. Both molecular and transition flow lead to a phenomenon known as fractionation, where the composition of sampled gas mixture can change because gas constituents are effectively pumped at different rates.

The SRS RGA200 operates at pressures of  $10^{-4}$  mbar or lower, offering 6 decades of dynamic range, or 9 decades with the continuous-dynode electron multiplier (CDEM) enabled. At this low operating pressure, gas flow is virtually guaranteed to be molecular, which has significant implications for how gas compositions are analyzed. For example, if a gas sample at high pressure (usually  $> 10$  mbar) is reduced to the RGA operating pressure in a single step via an orifice, then the gas mixture will fractionate [176,178,179]. This is because gas entering an inlet will be undergoing viscous flow, while gas flow exiting the inlet into the RGA space will be molecular.

Blessing *et al.* [176] provide an example of the resulting mixture composition where the mixture initially consists of 75% He and 25% N<sub>2</sub>. At the entrance to the inlet, the velocities are equal. At the exit however, the ratio of the gas velocities is:

$$\frac{v_{\text{N}_2}^{\text{out}}}{v_{\text{He}}^{\text{out}}} = \sqrt{\frac{4}{28}} = 0.38$$

Whereas previously the ratio of He to N<sub>2</sub> was 3, the increased He velocity relative to the N<sub>2</sub> enriches the outgoing mixture in N<sub>2</sub>. The corresponding He to N<sub>2</sub> ratio is then:

$$3 \times \sqrt{\frac{4}{28}} = 1.13$$

Bach *et al.* demonstrated this effect very convincingly using a crimped capillary [178]. In

---

<sup>6</sup>Also known as Graham's Law.

Figure 5.22, the measured ratios of different gas mixtures are plotted against the upstream pressure of a capillary. The effects of fractionation become evident beyond 10 torr, and mixtures do not appear to fully fractionate until the upstream pressure is greater than 1000 torr. Between this range of pressures, the mixture is undergoing transition flow, where the effect of fractionation is pressure dependent.




Figure 5.22.: Study of fractionation at different pressures with mixtures of different compositions. Dotted lines are curves from a naive model whereas solid lines represent a more complete model. From [178].

Because of this effect, it is widely recommended [175,176,178,180,181] to deliver gas to the RGA in two stages, where the flow regime changes between stages, not within them. The first stage reduces the gas pressure entirely in the viscous regime, commonly via a capillary, valve, or expansion into a large volume. The second stage uses some kind of molecular leak like an orifice or sintered plug [182] to deliver gas to the RGA entirely within the molecular flow regime. Variable leak valves may be used for the second stage, however an orifice will always be more stable and not subject to hysteresis [175]. Given the effects of fractionation can also depend on the geometry of a molecular leak [178,179], components with a fixed conductance are considered ideal [183].

Nevertheless, variable leak valves will always have a place in specialized systems where adjustments to the orifice need to be made [176,184,185].

In the case where a gas sample is continuously delivered (e.g. through a capillary), the pressure at the inlet will rise until it equilibrates with the pressure upstream because the gas flow rate up to the inlet will be much larger than the flow through it. To avoid this situation, a pumped-bypass configuration is used where excess gas at the inlet is exhausted using a dedicated vacuum pump [176,180,186]. A diagram of this configuration is shown in Figure 5.23.




Figure 5.23.: Pumped-bypass inlet configuration. The capillary reduces the gas pressure entirely within the viscous flow regime, and a molecular leak reduces the pressure entirely within the molecular flow regime to avoid gas fractionation. Excess gas is then routed away from the molecular leak to prevent a build-up of gas from the capillary inlet. From [180].

An added bonus of this two-stage pressure reduction system is that the volume of gas ahead of the inlet is pumped much faster, leading to a response time on the order of seconds [186] instead of minutes [179]. This is desirable because a faster response time decouples the response of the system from other time varying effects and true changes in a mixture's composition. The RGA manual [186] notes that a properly configured inlet should have gas from the first stage impinging on the second stage to further improve response times and ensure there are no dead volumes ahead of the leak.

The case where a finite sample of gas is expanded into a larger volume is often referred to as a batch inlet [175,176,180,183]. This method of sampling is appropriate for small gas samples that cannot be introduced continuously. A schematic of this type of system is shown in Figure 5.24.

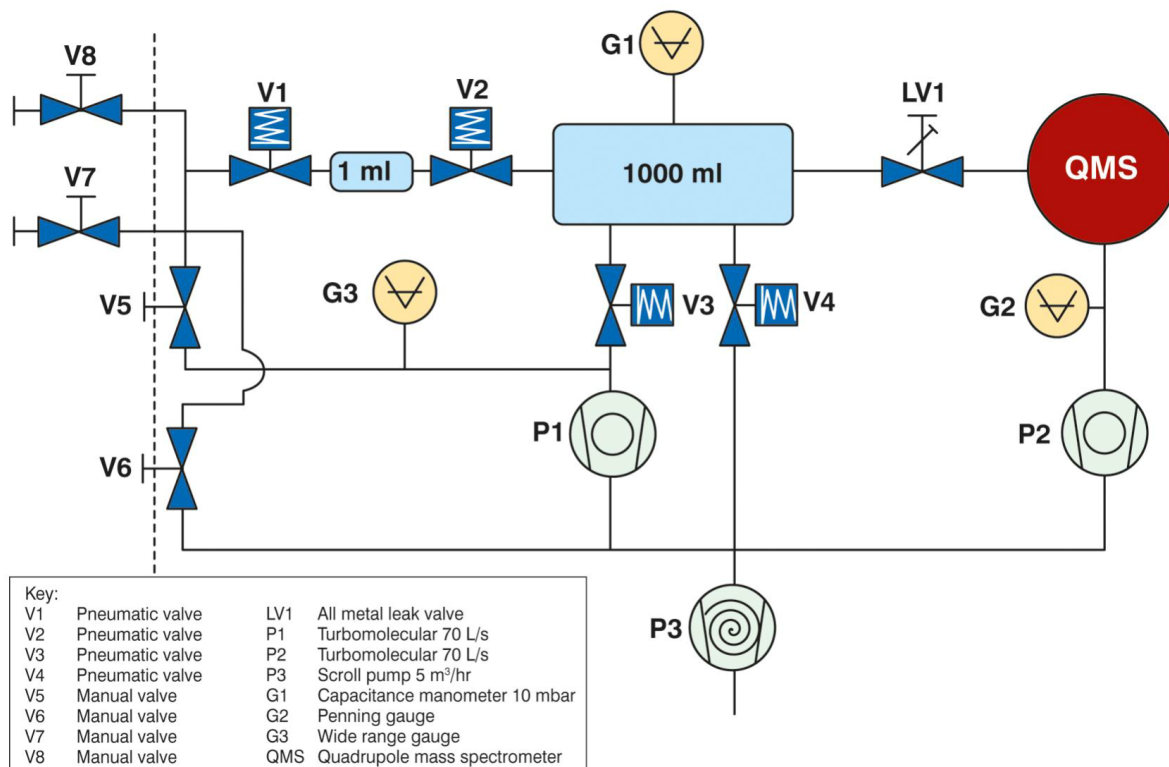


Figure 5.24.: Schematic of a Hidden HPR-70 batch sampling system. Supporting documentation could not be found on the exact operational sequences, but a schematic such as this is still useful for the purpose of reverse-engineering a batch-inlet sampling system. It is conceivable to assume that a sample is admitted into the volume between V1 and V2, then expanded to the 1 L volume downstream of V2. Following that, gas flows directly through LV1 to the quadrupole mass spectrometer (QMS). It is likely while sampling from the 1 L volume, V2, V3, and V4 are closed to prevent the sample from depleting too quickly. This would mean that a batch sampling system *can* operate without a pumped bypass (Figure 5.23), which was an important distinction that was not fully recognized when designing THERANOS (Figure 5.25). From [185].

A consequence of the finite gas sample in a batch inlet system is that the effects of fractionation are unavoidable because lighter components will be pumped out of the sample faster

than heavier components. Ellefson *et al.* [183] give equations modeling the partial pressures of constituents in a batch inlet system:

$$P_i = P_i^0 e^{-Lt/\sqrt{m_i}} \quad (5.16)$$

where  $L/\sqrt{m_i}$  is the ratio of the leak constant of the system and  $m_i$  is the mass of the constituent. If the total pressure of the system were being considered, this would correspond to the conductance-to-volume ratio  $C/V$ . Ellefson *et al.* recommend determining  $L$  via the introduction of He and extracting  $L$  from a plot of  $\ln P$  vs  $t$ . From Equation 5.16, the pressure ratio of two constituents will be:

$$\frac{P(m_1)}{P(m_2)} = e^{-Lt(\sqrt{m_2}-\sqrt{m_1})/\sqrt{m_1m_2}} \quad (5.17)$$

The authors further note that the effects of fractionation in a batch inlet system can be reduced if large sample volumes are used and sample times are kept short. While sample leak rates could be minimized, impacts on response times may be undesirable. Blessing *et al.* [176] also remark that samples can be integrated over time to calculate total quantities of constituents in a sample in order to work around the effects of fractionation.

Calibrating a system typically entails measuring the concentration of a gas mixture of known composition. MFCs [184,187] as well as gas expansion into static volumes [95,188] can be used to create gas mixtures for calibration purposes. Li *et al.* provide a useful summary of the theory behind calibrating a sampling system. In essence, ratios of mass peaks are useful because variations in certain operating parameters can effectively be divided out. During scans taken with THERA-

NOS, it was noticed that trends in  $^{129}\text{Xe}$  most closely matched those of  $^{40}\text{Ar}$ ; for this reason the “129/40” ratio is frequently quoted throughout this work. The issue of backgrounds and chemical reactivity influencing scans is also thankfully minimized because xenon and argon are noble gases. Furthermore, xenon is expected to be present in mixtures well above 1 ppm, making it easier to measure above background levels.

### 5.5.2. Design and Data Acquisition

THERANOS was designed to be a batch sampling system due to the limited quantities of gas that could be withdrawn from CHILLAX without perturbing it. For each sample, approximately 1 mBar of gas was stored in a 1 L sampling volume, and then pumped through an aperture or leak valve to an RGA for analysis. Calibrations could also be prepared by releasing small quantities of xenon into the larger storage volume, and then diluting the xenon by adding argon. Gas samples could also be mixed by hand or with a circulation pump.

The most recent P&ID for THERANOS (at the time of writing) is shown in Figure 5.25 and specific components for THERANOS are listed in Table 5.2. LV1 was initially used for pressure reduction upstream of the RGA. It was later supplanted by AP1, a specially produced VCR gasket with a 75  $\mu\text{m}$  aperture installed at the inlet of LV1. With AP1 installed, LV1 was left fully open during samples and it was confirmed to not interfere with gas pressures downstream. AP1 was installed because it was more consistent at reducing pressures and more exposed to agitation by the circulation pump CP1. V14 was installed to regulate pressures upstream of AP1 but was never tested; it is displayed in Figure 5.25 for posterity. More details on the performance of THERANOS are given in Section 5.5.3.

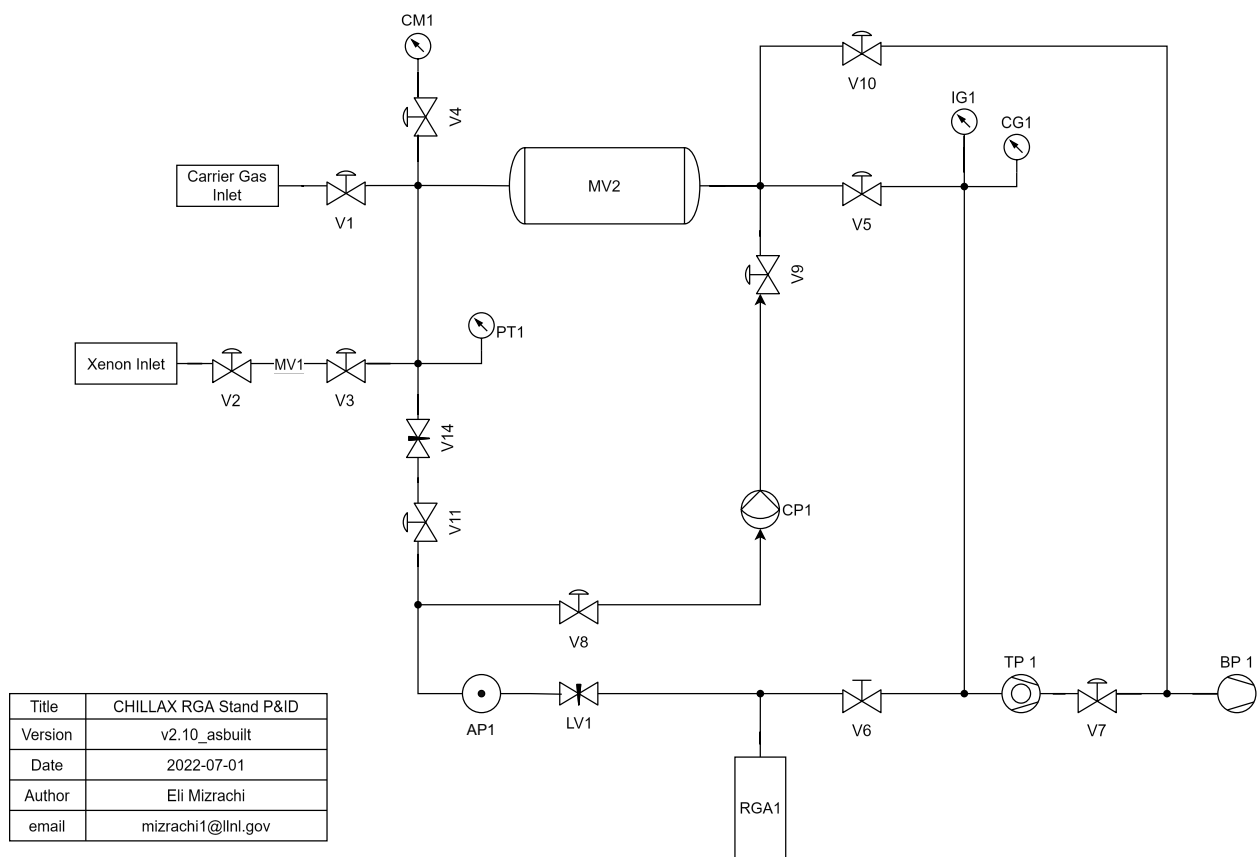


Figure 5.25.: P&ID for the most recent version of THERANOS. See Table 5.2 for a list of key components. LV1 was left fully open during samples to avoid a redundant pressure reduction following AP1. V14 was installed in an attempt to regulate pressures ahead of AP1 but was never tested.

Table 5.2.: Selected sampling system components

ID	Description	Vendor	SKU
AP1	75 $\mu\text{m}$ aperture	Lenox Laser	SS-4-VCR-2-75
CP1	Circulation Pump	Metal Bellows	MB-601
LV1	Variable Leak Valve	Pfeiffer-Balzers	UDV 140
PT1	3.4 bar Pressure Transducer	Setra	225
CM1	100 Torr Capacitance Manometer	MKS	626C
RGA1	Residual Gas Analyzer	Stanford Research Systems	RGA200
V14	Needle Valve	Swagelok	SS-SVR4-VH
IG1	Ion Gauge	Hornet	IGM402
CG1	Convection Gauge	Hornet	CVG101

The THERANOS cart (Figure 5.26) is constructed from two orthogonal subframes. This design enables the cart to be broken down and “shipped flat” in the event that the system will be moved to an exotic radiation source. The cart was found to be remarkably sturdy, and had enough empty space at the base to accommodate multiple pumps and other equipment. The entire plumbing manifold is mounted on three points of contact, to allow any point to be adjusted without excessive torque or load being applied on any of the others. Custom-machined Delrin blocks are used to offset valves on the gas panel so they sit in the same plane, easing the assembly process.



- (a) Front view of THERANOS. The top-most valve is V1, while V2 is at the top right. CM1 is downstream of V4 and located directly below V2. The CF plumbing manifold is supported by the green circular pipe clamp at the center of the frame, a U-bolt on the 4-way cross at the RGA, and a second U-bolt which acts as a pedestal (not visible) for the turbo.
- (b) Rear view of THERANOS. V7 can be seen at the bottom left, which was used to isolate TP1 and allow MV2 (the large cylinder) to be roughed through V10. The power and data enclosure (PDE) which contained a power supply for PT1 and CM1, as well as readout channels, can be seen in the center. See Figure 5.30 for more information on the PDE.

Figure 5.26.: Front and rear views of THERANOS.

Data was collected from the RGA via an RS-232 to USB adapter cable using the provided SRS RGA program (Figure 5.27). This program was used to control the RGA and acquire partial pressure spectra over a range of masses (analog scans) or time series data of specific mass peaks (P vs T scans). The SRS RGA program had a number of user interface/user experience (UI/UX) deficiencies that prompted the development of `rgaplot` [189] (Figure 5.28), a Python application designed for viewing and analyzing scans in real-time.




Figure 5.27.: SRS RGA Software showing a P vs T scan of argon and xenon masses. A key feature which was missing was the ability to plot pressures on separate y-axes, which led to suppressed features on both curves. See also Figure 5.29.

The backend of `rgaplot` uses `pandas` to read text files output by the RGA software, and a simple `streamlit` script works as a frontend for quickly making interactive plots with `plotly`. Documentation of classes and functions in `rgaplot` is managed with `mkdocs-material` and an API reference is generated automatically with the `mkdocstrings` extension for `mkdocs-material`.

The impact of a custom analysis frontend for RGA data was drastic. Because of the poor user interface in the official software, the initial workflow with the RGA involved downloading scans at the end of the day and analyzing them offline. This effectively added a 24-hour lead time

to analyses, which negatively impacted decision-making capabilities while calibrations were being performed. `rgaplot` completely eliminated any turnaround time between analyses, with a user-friendly interface that could be customized as needed. At the conclusion of work on this project, `rgaplot` could rapidly catalog analog and P vs T scans in a folder and assign each an integer ID. The user would then be presented with a table to select scans from. Given an integer ID, the scan type would be automatically detected and relevant plots would be drawn depending on the scan type. Filenames would also be parsed for information like initial sampling pressures and expected concentrations.

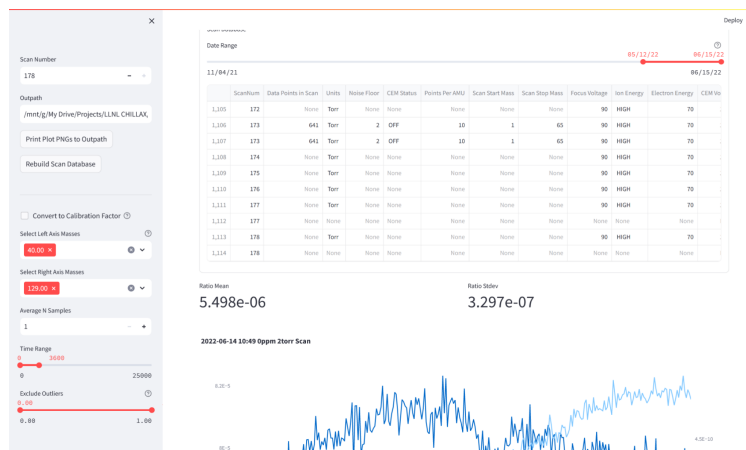


Figure 5.28.: Frontend of `rgaplot` showing part of a P vs T scan (Figure 5.29). P vs T scans include options for selecting masses to plot, which influence additional plots (not shown) like ratios and histograms of time series data.

An incomplete feature of `rgaplot` was the incorporation of readouts from sensors like pressure transducers to augment analyses and calibrations. These sensors were read out from the THERANOS PDE using a Keithley DAQ6510 which saved data to a USB drive. Two Omega DPi 32 modules were also connected to the PDE to provide simultaneous readouts of PT1 and CM1 on the THERANOS cart because the DAQ6510 was limited to only displaying one channel at any given time. A rendering is shown in Figure 5.30, and a wiring schematic of the PDE is shown in




Figure 5.29.: Scan shown in Figure 5.28. The level of detail is much greater than what can be seen in Figure 5.27.

Figure 5.31. The PDE was custom designed to provide a compact solution for integrating wiring between instruments on THERANOS. Rather than use a bulky benchtop power supply for PT1 and CG1, both could be powered from a 30 V DC power supply (see Figure 5.31 for part numbers). Breakout boards from Winford Engineering also offered a convenient way to route wires between different ports on the box.

### 5.5.3. Commissioning and Calibration

THERANOS was commissioned over a series of many months. Throughout that time, the design was iterated in order to mitigate issues with repeatability and stability of calibrations. While some problems remained at the time of writing, a reasonable calibration curve was eventually produced. This section will attempt to chronicle some of the problems experienced, however some open questions remain which will likely be the subject of future work.

THERANOS was calibrated by creating mixtures of known concentrations using MV1 and MV2. MV2 had a volume of 1 L, while MV1 had a volume of a few mL, formed by the interstice between V2 and V3. When gases in MV1 were expanded into MV2, their pressures were reduced




Figure 5.30.: View of internal PDE components. The board at the top left was unused but left as a future option for reading out IG1. From top right to bottom right, the breakout boards connected to: Omega DPi 32 modules, PT1, CM1, and the DAQ6510.

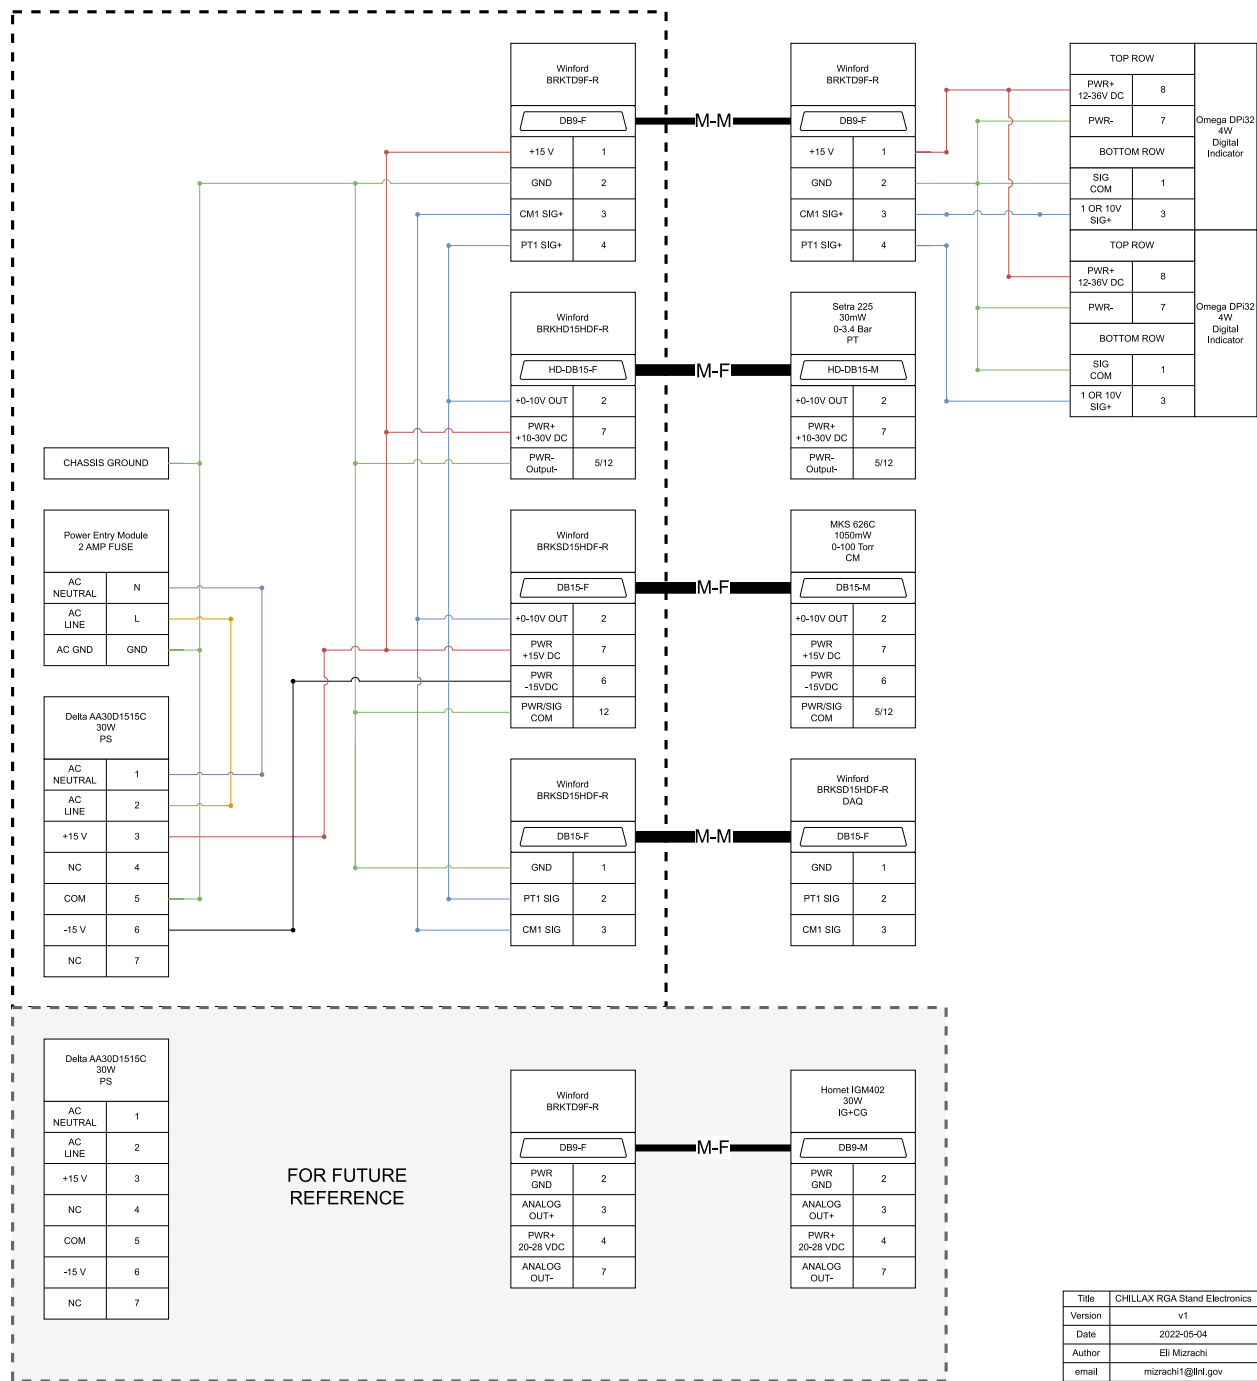


Figure 5.31.: THERANOS PDE wiring schematic. Capacity for reading out and powering the IGM402 was planned for, but never installed.

by a factor of approximately 600. Gas samples could be prepared in two ways. The first method involved expanding gas from MV1 twice by pumping out MV2 in between expansions. With this method, xenon gas injected into MV1 at 1 bar would have its pressure reduced by a factor of 360,000. Mixtures of varying concentrations at the ppm level could then be created by adding argon while monitoring CM1 or PT1. With the second method, xenon was only expanded once, then much more argon (typically ~1000 torr) was added to create a mixture usually with concentrations of ~1000 ppm or more. Following this it was also possible to reduce the pressure of the system by some desired factor, then refill with pure argon to effectively dilute the mixture concentration by that factor. For a more detailed explanation of these procedures, see Appendix C.

Version 1 of THERANOS (Figure 5.32) fed gas directly into LV1, with pressures in MV2 usually on the order of hundreds of torr. This was highly ineffective, and resulted in signals usually drifting upward on the order of 20% or more within a few minutes. Quantifying these issues was also difficult because data was recorded using the RGA's analog scan mode. Any time xenon partial pressures were measured, the CDEM had to be manually switched on, and a new scan had to be initiated. Scans were later switched to P vs T mode which could automatically switch on the CDEM without needing to initiate a new scan.

Initial issues with drifting scans were believed to be due to poor mixing of the prepared gas samples. To test this, Version 2 of THERANOS (Figure 5.33) saw the addition of the Manually Actuated Gas Mixing Apparatus (MAGMA). MAGMA consisted of a set of edge-welded bellows and 20 torr check valves to generate a one-way flow of gas when the user extended and contracted the bellows. V8 could also be opened to agitate gas instead of circulate it.

An example of a scan taken with MAGMA is shown in Figure 5.34. A 1540 ppm mixture

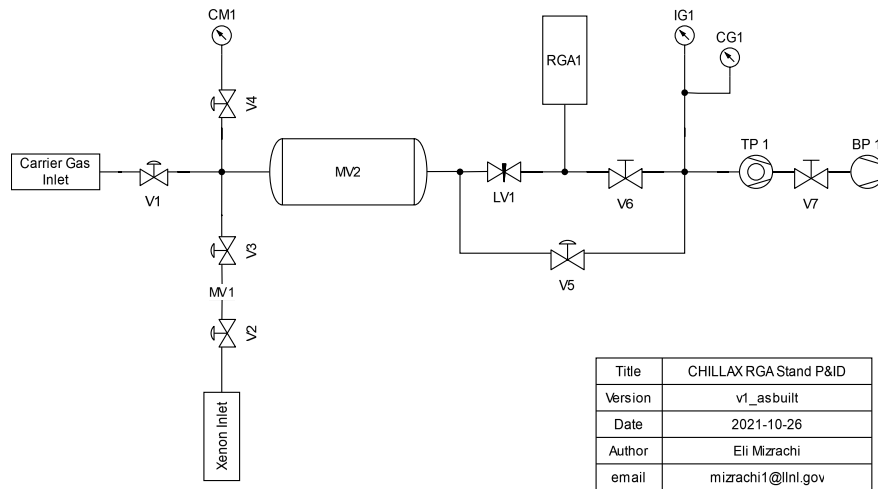


Figure 5.32.: The first iteration of THERANOS had a very simple design where gas was fed directly into LV1. Regions upstream of LV1 could be pumped via V5 which acted as a bypass, however TP1 needed to be shut off and fully spun down to evacuate any high pressure samples.

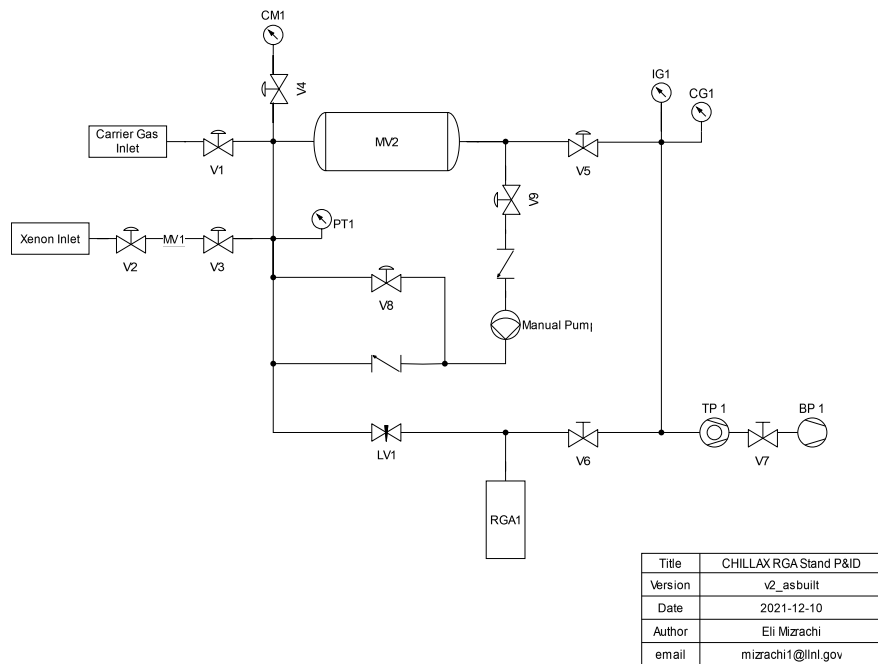


Figure 5.33.: MAGMA was installed in the second version of THERANOS. By opening V8, MAGMA could be used to agitate gas at LV1 instead of circulate gas past it.

was prepared with an initial pressure of 772 torr and LV1 was set to reduce the pressure at IG1 to 5E-5 torr. Initially, MAGMA was not used to mix any gas. The ratio appeared to increase rapidly within the first minute, before continuing to increase at a nearly constant rate. The mixture was agitated with MAGMA within the first few minutes, and a few times thereafter. These agitations are visible as small oscillations in the ratio, especially at 2000 s.

Agitation produced no clear effect, and counter-intuitively, pressures of both masses rose over time. This behavior is hard to explain. It is not consistent with the model of fractionation given by Ellefson *et al.* (Equation 5.17) but may have been due to more exotic effects from an excessive pressure reduction through LV1. The effects did not appear to be consistent with saturation at high pressures, which had been observed in previous scans. The possibility of nonlinear behavior near saturation however could not be ruled out. It should also be noted that the argon pressure was nominally constant, compared to the xenon pressures which more than doubled in the same timeframe.

The scan showed in Figure 5.34 was emblematic of other scans taken at this point in the campaign, which generally displayed large drifts over timescales greater than an hour. These issues during scans compounded with other concerns about repeatability *between* scans. Errors inherently present in the mixture preparation process were difficult to determine because the leak valve needed to be reset with each sample to pump out MV2 and the RGA space. Since both the upstream pressure and the position of LV1 would determine the pressure at IG1, it was difficult to reproduce conditions for each scan. Backgrounds from previous scans were a secondary concern because scans on the RGA showed that pumping for several minutes was enough to reduce xenon partial pressures by a factor of 100-1000 below typical signal values.




Figure 5.34.: Example scan with MAGMA. The CDEM was enabled for 129 (xenon) and disabled for argon (40). The bottom plot shows the ratio of the xenon and argon signals. The first mixing period is indicated with an arrow. It should also be noted that the argon pressure was nominally constant, compared to the xenon pressures which more than doubled in the same timeframe.

By far the largest improvement to the sampling procedure came from lowering inlet pressures (measured with CM1) to less than 2 torr. At these pressures it was found that mixing prepared samples with MAGMA—and later the XeNeu circulation pump—significantly lessened drifts that were seen previously. It was also found that at low inlet pressures, LV1 could be manipulated in different ways to further stabilize measurements.

At inlet pressures ranging from 0.75-1 torr, few adjustments were needed to produce stable results. Figure 5.35 shows an example where LV1 was slowly and continuously manipulated for the first 900 seconds to maintain a relatively constant pressure at IG1. The pressure at IG1 was then allowed to drop naturally for another 300 seconds, until LV1 was abruptly opened to return IG1 back to a nominal pressure of  $\sim 5 \times 10^{-5}$  torr. This produced a slight, but noticeable change in the measured ratio. Following this, the pressure was maintained again using LV1 until the entire sample was depleted. The measured ratio again did not appear to change until the xenon pressure had fully bottomed out. This procedure, while it appeared to be successful, was also incredibly demanding on the operator's attention. Alternative procedures were sought out to relieve operator fatigue, and investigation of this technique was left as a possible avenue for future work (Section 5.5.3.1).

At pressures between 1.5-2 torr, abrupt adjustments were necessary to stabilize the measured ratio (Figure 5.36).

Unfortunately at these pressures, consistency between samples remained frustratingly out of reach. This was in part because measured ratios of samples depended on system pressures as well as the setting of LV1. Figure 5.36 shows this effect quite clearly, where the ratio values differ almost 30% depending on the setting of LV1. Within a single scan this level of variation might be acceptable, however it was difficult to characterize this behavior between scans. To make matters




Figure 5.35.: Scan taken with CM1 at 1 torr and IG1 at  $5E-5$  torr. LV1 was adjusted continuously to keep the value of IG1 as stable as possible. During the period between the dotted lines, LV1 was not adjusted and the pressure at IG1 was allowed to decrease naturally. Spikes in the values are from a high scan speed (4) at the RGA that introduced additional noise. These spikes were greatly diminished after switching to scan speed 2 in subsequent acquisitions. Because the scan speed was changed before the final calibration curve was generated, results shown here were not used to infer xenon gas concentrations. Note the use of a logarithmic scale in the ratio to capture the behavior at the end where the xenon pressure had bottomed out and the argon pressure decreased exponentially.




Figure 5.36.: Sample taken from the outlet of early CHILLAX doping tests at 0.3% xenon liquid mole fraction and initial MV2 pressure of 1.5 torr. Nominal pressures from IG1 are indicated on the top plot. These pressures were set using LV1.

worse, it was later found that abrupt changes in LV1 were subject to hysteresis. When LV1 was opened, closed, and then opened again to reach the same pressures on IG1, gas ratios did not return to previously measured values. Because LV1 adjustments between scans were fairly arbitrary, the coupled nature of all of these problems was too complicated to navigate.

LV1 was made redundant by AP1 in later iterations of THERANOS in an attempt to overcome these challenges. It was thought that scans with AP1 would be easier to reproduce with a fixed conductance and a pressure at IG1 that would be determined solely by the pressure in MV2. Unfortunately, this did not reduce instabilities in scans. Consequently, an extensive effort was undertaken to reproduce the qualities of a pumped-bypass sampling inlet (Figure 5.23). Gas was continuously recirculated through V8, V9, MV2, and V11<sup>7</sup> (see Figure 5.25) during samples so gas would impinge on AP1 and prevent “dead” volumes from forming.

The XeNeu circulation pump (MB-601) reduced the magnitude of instabilities during scans to an order of roughly 10% (Figure 5.37). This degree of variation within a scan was consistent with values found in other sources [184] [190]. However, inconsistencies between scans persisted, with typically one scan per day disagreeing by 20-30%. At this point in time, reducing instabilities during scans seemed like the most feasible way to decouple uncertainties between and during scans. This led to the pursuit of a higher gas circulation rate to increase agitation at AP1. But even when the MB-601 was configured with both pump heads in series to maximize its compression ratio, it was not designed for circulating gas at pressures on the order of 1 torr or less. Using the MB-601 also meant sacrificing XeNeu’s operational capabilities, which added an incentive to switch to a more conventional pump for recirculating gas in THERANOS. As a result, vacuum pumps were studied as an alternative because of their high pumping speeds at low pressures. A table of pumps

---

<sup>7</sup>Note that V14 had not been installed yet.




Figure 5.37.: Example of scan instabilities occasionally seen when operating with gas agitation from the XeNeu circulation pump. The initial rise may have been from a xenon-rich dead volume trapped behind AP1. The gradual decline is not consistent with fractionation however, which doesn't appear to occur until almost 2500 s into the scan. This degree of variation within a scan was consistent with values found in other sources [184] [190]. However, inconsistencies of approximately 20-30% between scans persisted.

investigated is shown in Table 5.3.

Table 5.3.: Tables of pumps studied for gas recirculation at low pressures. Helium leak rates were independently checked and generally found to be in agreement with quoted leak rates. Some pumps were modified from their stock configuration in an attempt to lower leak rates. Of the three modified pumps, only the ACP-15's modification produced a substantial improvement by eliminating detectable leaks through the ballast port.

			Max He Leak	
Vendor	SKU	Type	Rate [mbar L/s]	Modifications
Pfeiffer	ACP-15	Roots	1E-8	Blanked ballast
Pfeiffer	MVP 035-2	Diaphragm	3E-5	Metal transfer tubes
Varian	949-9451 <sup>8</sup>	Diaphragm	5E-8	
Varian	949-9432 <sup>9</sup>	Diaphragm	2.5E-7	Metal transfer tubes
Edwards	NXDSi6	Scroll	2E-7	
Agilent	IDP-3	Scroll	2E-5	
Senior Metal	MB-601	Reciprocating	<1E-10	
Bellows		Bellows		

Of all the pumps studied, no pumps were found to be as leak tight as the MB-601. While any of these pumps likely would have been suitable for mixing gas for brief periods of time, their leak rates were too high to be useful for sustained low pressure operations during a sample. The ACP-15 was the closest pump to being effective, however it could not be fully evacuated after

<sup>8</sup>AKA Vaccubrand MD-4

<sup>9</sup>Essentially identical to the Pfeiffer MVP 015-2.

circulating gas because of an internal check valve on its exhaust port. This check valve leaked at a rate of approximately  $1\text{e-}3$  mbar L /s, raising concerns of systematic errors being introduced into the sample preparation process. In addition to this, the ACP-15 also suffered from a known issue of hydrocarbon contamination via labyrinth seals between its oil-filled gearbox and the process space.

To focus the flow of gas and reduce the dead volume ahead of AP1 an additional tube was added between V11 and LV1 (Figure 5.38.) An 1/4" copper tube was suspended from the VCR gasket at the outlet of V11 by flaring the tip of the copper tube and sliding it through the VCR gasket. This tube then extended downward through an HVCR-4 tee that connected V11 to V8 and LV1. The tube extended past the inlet of V8, leaving a few mm of space at the inlet of LV1.

The inlet tube did not appear to have a significant effect on samples taken using the XeNeu circulation pump. To try and take advantage of more directed flow at AP1, subsequent tests were conducted using the IDP-3 scroll pump connected to V8. Rather than recirculate gas, the IDP-3 was allowed to directly exhaust to atmosphere, essentially creating a pumped-bypass configuration (Figure 5.23). Results from this configuration were promising at eliminating spontaneous instabilities, however it was ultimately not feasible to implement because the high flow rates would deplete MV2 in a few minutes or less. During this timeframe it was also difficult to manually maintain stable pressures using V8 and V11.

#### **5.5.3.1. Recommendations For Future Work**

The body of work conducted suggests that in this system, other instabilities may form while fractionation from a small sampling volume is subdominant. These instabilities can be dispelled if the gas is agitated by the flow of more gas, or in the case of a leak valve, by varying the orifice itself.




Figure 5.38.: Cutaway of V11 and the HVCR-4 tee showing the location of the 1/4" copper inlet tube (highlighted in orange) added to reduce the size of the dead volume upstream of AP1. V14 is visible at the top. V11 (middle) was connected to LV1 (bottom) and V8 (lower right) via the HVCR-4 tee which housed the inlet tube. The inlet tube was suspended from the VCR gasket at the top of the HVCR-4 tee by flaring the tip of the tube and sliding the other end of the tube through the gasket. Note that VCR gaskets are not shown in this rendering.

A promising direction for future work would be collecting scans with conditions similar to those in Figure 5.35. Inlet pressures below 1 torr, and possibly as low as 0.5 or 0.25 torr should be examined. At these pressures, samples can be fully exhausted in a reasonable (~30-40 min) timeframe. This would be beneficial for studying the effects of fractionation and to assess whether the mixture is behaving according to existing models (Equation 5.17). Other concentrations should be examined as well; only a few scans were taken at inlet pressures below 1 torr, and all of them were taken at similar xenon concentrations of 10-20 ppm. A concise summary of recommendations is given in Section C.12.

It is unclear why fractionation did not appear to occur in Figure 5.35 even as the sample was rapidly being depleted. In the scan shown in Figure 5.39, LV1 was only opened initially with an inlet pressure of 0.73 torr. The measured ratio varied slightly more over time compared to the scan in Figure 5.35. It also appeared to fractionate according to Equation 5.17 after approximately 2000 s, when the argon partial pressure appeared to decrease at a slightly faster rate.

Differences between the scans in Figure 5.35 and Figure 5.39 could be attributed to manipulation of LV1 to produce a stable inlet pressure in Figure 5.35. In a paper by Little and Gordon [191], the authors developed a novel inlet valve controlled by stepper motors in order to maintain a stable inlet pressure as process pressures varied across roughly 10 orders of magnitude. This paper, along with the work by Ellefson *et al.* [183] could serve as a starting point for future modeling efforts. At the very least, they offer some more motivation for testing the behavior of THERANOS with a weakly varying conductance at sub-torr inlet pressures.

An example of a calibration scan is shown in Figure 5.40. Calibration scans were taken using AP1 for pressure reduction and the XeNeu circulation pump to provide some form of agitation.




Figure 5.39.: Scan taken from early CHILLAX doping tests at 0.73 torr. Unlike in Figure 5.35, the measured ratio begins exponentially increasing before the xenon pressure reaches its baseline.

The dashed box indicates the region of data used for calibrations. This region was chosen due to its ~10 minute-long period of stability. The very beginning of the scan was excluded because it was not apparent that the scan had stabilized yet. The mixture began to visibly fractionate at around 1500 s, so this period was also excluded.




Figure 5.40.: Example of a scan used for calibration purposes. The region inside the dashed box was used for calibration purposes due to its stability. Subsequent calibration scans relied on similar procedure to select data where stable time periods of 5-10 minutes at the beginning of a scan were used in the final calibration curve (Figure 5.41).

Subsequent calibration scans relied on similar procedure to select data where stable time periods of 5-10 minutes at the beginning of a scan were used in the final calibration curve (Figure 5.41). To generate the calibration curve, mean values of 129/40 ratios during selected time intervals in calibration scans were plotted against known mixture concentrations. Adding the approximate pressure at CM1 during a scan as a third (color) axis reveals a possible systematic effect stemming

from the approximate inlet pressure as measured by CM1. At inlet pressures of around 1.5 torr, scans occasionally appear to be more enriched in xenon. This may be because samples at these pressures began to experience fractionation faster, but some samples at these input pressures did not appear to follow this trend. The possibility of other systematic effects between scans cannot be ruled out here, and should be the subject of future work.

A linear fit of the form  $y = mx + b$  is plotted on Figure 5.41 with the value of  $b$  fixed at 0. In order to provide a more conservative estimate of the xenon vapor mole fraction, data points with inlet pressures less than or equal to 1.5 torr are excluded from the calibration curve. The ratio of the 129/40 RGA response can be multiplied by the calibration factor  $m$  to obtain an estimate of the xenon concentration in a gas mixture. The curve was adequate for predicting concentrations of samples above roughly 5 ppm. Below this region the curve was non-linear, possibly due to unaccounted-for effects from adsorptive backgrounds in scans.

## 5.6. Demonstration of High Concentration Xenon Doping in CHILLAX

This section will summarize a study conducted in 2023 [135] that demonstrated prompt doping of xenon in liquid argon and the stability of the mixture after the doping process was completed. The mixture was doped up to a mole fraction of 2.35%, and it was estimated based on measurements from THERANOS that ~2% doping was sufficient to maintain 30-50 ppm of xenon in the gas phase of the system. These results from THERANOS were nominally consistent with expected concentrations calculated in Section 5.2.2.

The study began by condensing 454 standard liters of pure argon in the main detector volume, forming a 6.6 cm deep liquid bath. After establishing circulation at 1.5 standard liters per minute




Figure 5.41.: RGA calibration curve. To generate the calibration curve, mean values of 129/40 ratios during selected time intervals in calibration scans were plotted against known mixture concentrations. A linear fit of the form  $y = mx + b$  is performed with the value of  $b$  fixed at 0. Values with inlet pressures below 1.5 torr were excluded from this fit to maintain a more conservative estimate. Error bars on points represent the standard deviation of the selected scan period and serve as a rough guide of the variation present in scans; they were not used to generate the fit parameters. The ratio of the 129/40 RGA response can be multiplied by the calibration factor  $m$  to obtain an estimate of the xenon concentration in a gas mixture. The error on  $m$  is statistical only, and it dominates over systematic errors present from excluding nearby points with inlet pressures less than or equal to 1.5 torr.

(SLPM), the detector pressure was stabilized at  $1.8 \pm 0.005$  bar with a liquid temperature of roughly 93 K by delivering  $\sim 0.5$  W to H3 (Figure 5.10).

Xenon was injected into the circulation path at a rate of 8.8 standard cubic centimeters per minute (SCCM), corresponding to an incoming gas concentration of approximately 0.6%. To prevent the formation of clogs, H2 was used to heat the condenser inlet tube to 200 K before and during xenon doping periods. Throughout the study, xenon was injected in four separate doping steps, raising the mole fraction by 0.26%, 0.86%, 0.59%, and 0.64%.

Measurements from the capacitive sensor are shown in Figure 5.42. During the first doping step, the temperature of the main volume walls and top flange was set slightly below the liquid bath temperature. This resulted in liquid condensing on the walls of the main volume. The xenon concentration in the main liquid bath also rose slowly during this period of time. This was thought to occur because gas from the condenser was being pulled into the main volume through the upper interconnect, preventing xenon-rich liquid in the condenser from moving into the main volume.

During the last three doping steps, the temperature of the main volume walls and top flange was set slightly above the condenser and liquid bath temperature. This promoted gas flow from the main volume into the condenser via the upper gas interconnect, thereby encouraging xenon-rich liquid to flow from the condenser into the main liquid bath. As a result, the xenon concentration in the liquid rose quickly during these doping steps.

Upon doping the liquid mixture to 2.35%, a sequence of tests was conducted where different thermal gradients were established between the top flange of the main volume and liquid bath. These tests demonstrated how ice formation in the main volume could be controlled and reversed.

In the first test, the temperature of the top flange was maintained 0.3-0.5 K higher than the




Figure 5.42.: Capacitive measurements of the CHILLAX liquid bath over time. Doping steps are outlined with the hatched red regions. Each point on the histogram represents a data point recorded by the capacitive meter in a specific time bin.

liquid bath temperature. In the second test, the TSU was disabled and the top flange temperature stabilized around 170 K or approximately 75 K higher than the liquid bath temperature. In the third test, the top flange temperature was maintained at 10 K above the liquid bath temperature to simulate poor temperature regulation capabilities.

Figure 5.43 shows pictures from the camera during each of the three tests. A large ice ring formed rapidly in the second test, while the ice ring was slower to form in the third test. No ice ring was observed during the first testing period.

Photos from the camera were analyzed and the size of the ice ring was measured in each photo. The size of the ice ring over time in each test is shown in Figure 5.44, alongside data from the capacitive sensor during the same time periods. The data from the photos further demonstrated that the test with the largest thermal gradient resulted in an ice ring which formed rapidly compared to the other two tests. It is also apparent that even with a minor 10 K temperature difference, ice




Figure 5.43.: Photos of the CHILLAX main volume taken by a camera mounted on the underside of the top flange. Photos from the first, second, and third tests are in the top, bottom and middle rows, respectively. In the second test a prominent ice ring is visible as a bright arc spanning the top right corner to the bottom left corner of the photos. A fainter ice ring is present in the third test, while no ice ring was present in the first test. These photos provided substantial evidence that ice formation was strongly linked to the presence of a thermal gradient between the liquid bath and walls of the main volume. The thin dark arc present in all of the photos is the edge of the bubble router at the bottom of the main volume.

was liable to form over time. The inverse behavior from the capacitor during the testing periods also suggests that the ice ring depleted xenon from the liquid bath, consistent with the behavior described in Figure 5.3b.




Figure 5.44.: Capacitive sensor data (top) and camera data (bottom) showing the xenon concentration in the liquid bath and ice ring thickness, respectively. The blue squares correspond to data from the first test, while the red circles correspond to data from the first test. Data from the black test is displayed as black triangles. The growth of the ice ring neatly corresponds with a depletion of the xenon concentration in both the second and third tests. In the first test, no xenon depletion or ice ring formation was observed beyond the range of the uncertainties of the camera analysis and capacitive sensor.

After the ice was formed in each test, the top flange was cooled to about 1 K below the temperature of the liquid bath. The resulting film of liquid that condensed on the walls of the main volume dissolved the xenon ice and restored the concentration of the liquid back to the original value that had been measured by the capacitive sensor. At times, chunks of xenon ice fell into the liquid bath, resulting in a rapid rise in the measured xenon concentration. The behavior is shown

in Figure 5.45.




Figure 5.45.: Xenon concentration over time after a test that formed an ice ring around the liquid bath. The top flange was cooled to condense gas on the walls of the main volume, which dissolved the ice back into the liquid bath. Chunks of ice fell back into the liquid at times, resulting in abrupt rises in the xenon concentration measured by the capacitive sensor.

Once a xenon liquid mole fraction of 2.35% was reached in the main volume, gas samples were evaluated with THERANOS in parallel with the ice formation tests. From June 1, 2022 to the beginning of June 7th, no ice was present in main volume. Starting on Jun 7th ice was intentionally formed by allowing a thermal gradient to form between the top flange and the walls of the main volume. The ice was removed again by June 14th. Gas samples were typically acquired at 2 torr and 10 torr inlet pressures. In some scans, measured ratios drifted 30-50% over time. It was estimated based on prior experiences during calibration that variations in procedures and changes in the state of LV1 during scans also contributed to shifts of 20-30% between scans.

Despite these uncertainties, an analysis of the data in aggregate (Figure 5.46) shows that

these effects are likely subdominant to true changes in the xenon vapor mole fraction between June 1st and June 14th. For example between June 1st and June 7th, samples were consistently within a range of 30-50 ppm. This provides an estimate of systematic errors induced from different procedures used to sample gas. On June 7th however, the measured concentration rose substantially, in a manner commensurate with ice formation in the main volume. Following this the measured gas concentration remained elevated, even up to June 14th when ice did not appear to be present in pictures taken with the camera. This could be a result of systematic errors dominating over a lack of data. However another possibility worth examining is the effect on the vapor mole fraction from evaporation at the liquid meniscus.

Ice forms at the meniscus because it is more difficult for the liquid there to evenly mix, and there is an influx of heat from the walls of the vessel. By the same principle, the xenon vapor mole fraction should be enhanced by this effect, leading to a higher-than-predicted vapor mole fraction if only the characteristics of the bulk liquid were being taken into consideration. Future studies could examine the possibility of this effect by comparing temperatures between the bottom and walls of the main volume. At a xenon liquid mole fraction of roughly 6%, the predicted vapor mole fraction is almost 100 ppm with a liquid temperature of 94 K (Figure 5.9b). This is near saturation conditions, which seem reasonable for a liquid film present at the meniscus.

Ultimately, the doping test with CHILLAX was successful in demonstrating that the xenon concentration in an argon mixture could be reliably controlled. Showcasing the conditions required to produce and eliminate instabilities will have future implications for experiments such as the Scintillating Bubble Chamber Collaboration [192], ProtoDUNE [193], and LEGEND [194] which are considering xenon-doped argon as a scintillator. The measurement of 30-50 ppm of xenon in



the gas phase also paints a promising picture for dual-phase experiments, motivating future work with CHILLAX to measure the electroluminescence light spectrum.

## Appendix A.: Information on Stripping Kapton Cables

Kapton cables are challenging to work with because of Kapton's high cut resistance, and the nearly absurd fragility of the stranded inner conductors<sup>1</sup>. Both the inner dielectric and outer insulation on Kapton coaxial cables consist of multiple layers of Kapton film, not unlike a roll of adhesive tape. These qualities generally make standard plier-style strippers (Figure A.2b) unsuitable for use with Kapton cables. Because a significant amount of force is required to sever the layers of insulation, these strippers have a tendency to snap cables entirely since the rapid and vigorous cutting action will pry apart the insulation, in turn applying catastrophic amounts of tension to the inner conductor. See Figure A.1 for a diagram describing these forces.

Prior to the assembly process, a total of six strippers were evaluated. Evaluation of these strippers included repeatability tests and checks of cuts under 10X magnification to ensure that there was no damage to the underlying conductor or layers of insulation composing the inner dielectric. These tests were crucial to perform as the eventual construction of a cable set would require over 100 cuts at varying depths across all of the cables in the set. The strippers tested are listed in Table A.1.

---

<sup>1</sup>The inner conductors are listed as silver coated copper wire, like with FEP cables. They likewise have similar diameters to those in FEP cables. However for unknown reasons, they have a tendency to fatigue and break *much* faster—typically within 1-2 cycles of being bent.




Figure A.1.: In this diagram, a set of cutting edges (typically two halves of wire stripping die) requires a significant shear impulse relative to the wire to pierce the kapton film. The cutting edge pries apart the insulation (colored orange) with some impulse  $J_{C \rightarrow I}$ . Because the insulation does not slide along the wire, this impulse is directly transmitted to the wire as  $J_{I \rightarrow W}$ . Thus, despite not making direct contact with the conductor, the cutter transmits a high tension impulse via the insulation, which can break the conductor.

Table A.1.: Strippers tested on AWG 30 Kapton coaxial cable. Note that stripper 6 has a continuously adjustable blade which accommodates wires up to 0.125" in diameter. Stripper 4 has the same operating mechanism as stripper 5, but uses custom components for gripping and cutting Kapton insulation.

ID	Description	AWG	Vendor	SKU	Figure
1	Red Standard Strippers	16-26	Generic	N/A	Figure A.2b
2	Yellow Standard Strippers	22-30	Generic	N/A	Figure A.2b
3	Precision Strippers	26-36	AccuGlass	100191	Figure A.2a

ID	Description	AWG	Vendor	SKU	Figure
4	Black Stripmaster	24-30	AccuGlass	100196	Figure <a href="#">A.2c</a>
5	Blue Stripmaster	10-22	Ideal	45-292	Figure <a href="#">A.2c</a>
6	Grey Coaxial Ringer	0-0.125"	Ideal	45-162	Figure <a href="#">A.2d</a>

The following types of cuts were needed on a given cable:

- Stripping the outer insulation to expose the outer shield
- Stripping the outer insulation and outer shield to expose the inner dielectric
- Fully stripping the inner dielectric to expose the inner conductor
- Partially stripping the inner dielectric to reduce its diameter
  - This was necessary in order to fit the cable inside the body of a crimp-on MMCX connector, and to fit the inner dielectric through the 23-pin connector body.

Sections of Kapton insulated 28 AWG stranded core wire (AccuGlass 100690) also needed to be fully stripped. These wires were used to bridge outer shield braids to pins in the 23-pin connector body.

Strippers [3](#), [4](#) and [5](#) were listed by AccuGlass specifically for stripping Kapton<sup>2</sup>. A summary of stripping techniques is shown in Figure [A.3](#). Stripper [6](#) performed exceptionally well largely

<sup>2</sup>Stripper 3 is claimed to be “Great for use with Kapton Insulated wire”, while stripper 4 is listed as being used for “the bulk of our production wire assembly needs”. AccuGlass further advertises that “WE HAVE NOT FOUND A BETTER WAY THAN THIS TO STRIP THE KAPTON® INSULATED WIRE”.



(a) So-called “precision stripper” with adjustable stripping length. These strippers work similar to a pair of tweezers, with a blade at the end that shears the insulation. Due to Kapton’s cut resistance, these strippers would routinely sever wires completely because of the force needed to initially pierce the Kapton insulation. Image from [zoro.com](http://zoro.com).

(b) Example of standard plier-style strippers. Image from [idealind.com](http://idealind.com).



(c) The Ideal Stripmaster clamps onto wires, cuts the insulation, and removes insulation when the handles are squeezed together. Custom variants (typically with a black-colored handle) can be purchased with clamping components and cutting dies specifically for materials like Kapton. Image from [idealind.com](http://idealind.com).



(d) A coaxial “ringer” style stripper. The ringer contains two continuously adjustable cutting blades. The cutting depth of each blade can be set using a depth gauge on a pair of calipers, improving repeatability. Cuts can be made parallel or perpendicular to the axis of the cable. Cuts parallel to the cable are useful for cutting long segments of insulation, or cutting segments without needing to remove the insulation from either end of a cable. Image from [idealind.com](http://idealind.com).

because of its gentle cutting action and its continuously adjustable blade, which could be set to a specific depth using a pair of calipers. Stripper 4 was the easiest to use, however it was only able to perform two tasks: partially stripping the inner dielectric, and fully stripping AWG 28 stranded wire. The plier style strippers (1, 2) were also useful for gripping and removing layers of insulation once incisions had been made with stripper 6. The high utility and low cost (~\$30) of stripper 6 made it worthwhile to purchase three copies to avoid repeatedly “dialing-in” a single unit to strip different layers of the cables.

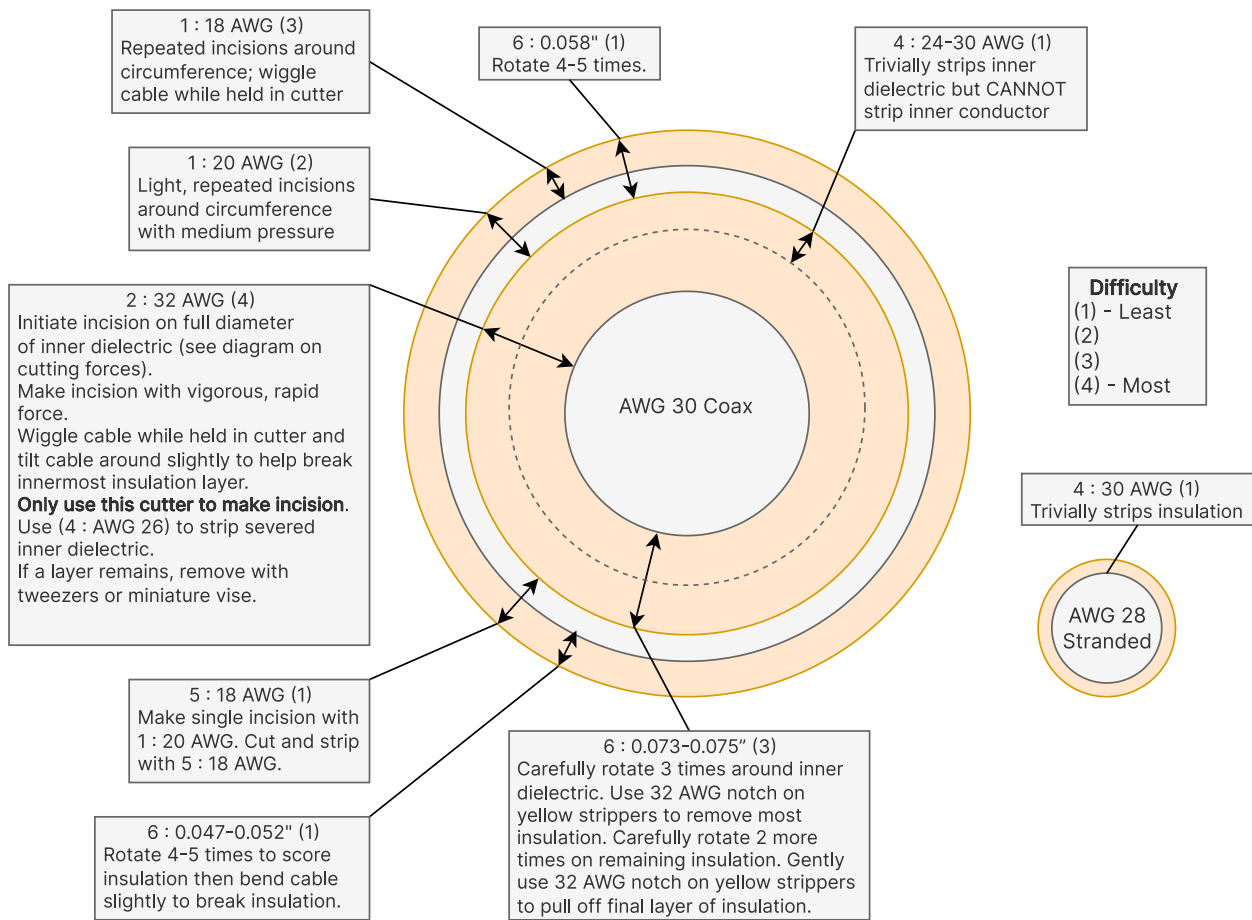


Figure A.3.: Various stripping techniques were attempted with different strippers. Techniques indicate layers of insulation and/or shield removed, the stripper ID used, setting, and relative difficulty. Fully stripping the inner dielectric was the most difficult step, which could be accomplished with some practice using strippers 6 and 2.

## A.1. Cable Assembly Procedures

This section details assembly procedures for the various cable connections on 30 AWG Kapton coaxial cables. Procedures were developed with the intent that a batch of cables could be processed at each step. This reduces careless errors since the same action is performed repeatedly, instead of a new action each time. Quoted lengths of insulation etc. to be removed from a cable are typically measured from the end of a cable unless otherwise noted. A depth gauge on a pair of calipers is helpful for this, if not essential.

### A.1.1. Sub-C Connector

#### *Step 1*

- Slide heatshrink over cable
- Make initial incision on outer insulation and braid with 20 AWG notch on red strippers 1.25” from cable end as a visual guide
- Remove 1.25” outer insulation and outer conductor with 18 AWG notch on blue stripmaster<sup>3</sup>

#### *Step 2*

- Remove an *additional* 0.375” outer insulation with 0.047” depth blade on ringer

#### *Step 3*

- Peel back outer conductor using needle nose tweezers

#### *Step 5*

---

<sup>3</sup>Note that this length is necessary to provide enough slack for the sleeve crimp tool.

- Fully strip 0.125” of inner dielectric from end:
  - Use 30 AWG notch on black stripmaster to remove most insulation.
  - Carefully rotate 2 times on final layer of insulation with grey ringer at 0.075”
  - Gently use 32 AWG notch on yellow strippers to pull off final layer of insulation.
- Partially strip inner dielectric 0.25” from end with 26 AWG notch on black stripmaster

*Step 6*

- Insert exposed inner conductor into connector
- Crimp pin onto inner dielectric, near base of sleeve, then crimp again onto inner conductor
- Use 30 AWG notch on black stripmaster to strip 0.125” of 28 AWG wire section
- Crimp pin onto insulation of 28 AWG wire, then conductor of 28 AWG wire
- Insert 28 AWG wire into corresponding hole for cable outer conductor connection, making sure sleeve is coincident with connector body
- Cut 28 AWG wire such that when stripped, an exposed portion of 28 AWG wire will overlap with the outer conductor stub of the cable
- Remove 28 AWG wire from connector, strip previously marked segment, and reinsert into connector
- Solder 28 AWG wire to outer conductor stub (Figure A.4)
- Cover inner and outer conductor wires with ~1” heatshrink

### **A.1.2. MMCX Connector**

*Step 1*




Figure A.4.: After a sleeve is crimped to the inner conductor, the cable is captured in the body of the connector. A piece of 28 AWG wire is fed into the appropriate receptacle on the connector, then cut to length to match the location of the outer conductor from the cable. The 28 AWG wire is then stripped and soldered into place. It is vital that when soldered, the cable and lead wires have matching lengths and are not bent out of place in order to avoid unnecessary strain on either.

- Slide on crimp sleeve, then heatshrink
- Expose 0.080" inner conductor using ringer at 0.083"
  - Rotate 3-4 ringer 3-4 times
  - Gently remove some of inner dielectric with 28 AWG notch on yellow strippers
  - Reinsert into ringer, then rotate *cable* 2 more times
  - Pull off remaining insulation with 28 AWG notch on yellow strippers
    - \* Fully closing strippers is unnecessary, only a "grip" is needed on the inner dielectric which can be achieved with moderate pressure. This avoids bending the inner conductor.

### *Step 2*

- Cut through outer insulation 0.355-0.360" from end using ringer at 0.053"

**! Warning**

This was found to be too short, leading to recessed MMCX pins which did not make a strong connection. It is suggested to strip at least 0.010-0.020 inches more than this so the pin can be inserted further into the body of the connector. Prior to crimping, comparison with a commercially produced connector should be performed.

- Do *not* remove outer insulation

*Step 3*

- Cut through outer insulation and outer conductor 0.160-0.165" from end
- Remove outer insulation and outer conductor on this section, exposing inner dielectric

*Step 4*

- Remove remaining layer of outer insulation
- Peel back outer conductor using tweezers
- Use 20 AWG notch on blue stripmaster to partially remove all (~0.345"-0.350") exposed inner dielectric

– If there is a strong adhesive bond between layers, this will fail and the stripmaster will only scrape some of the layers off. If this happens, a heatgun at 400 deg F aimed at the stripper and cable is essential to soften the adhesive. Remaining layers can be removed with the 22 AWG stranded notch on red strippers.

\* This seems to be an unknown factor between batches of Kapton cables.

- The 24 AWG notch on the black stripmaster can be used here, but it may damage the outer conductor slightly

*Step 5*

- Tin inner conductor with solder
- Solder inner conductor to pin

*Step 6*

- Insert cable into body of plug
- Push heatshrink such that heatshrink covers half of knurled portion of connector
- Slide crimp sleeve over heatshrink, pushing heatshrink forward slightly until crimp sleeve is butted up against body of plug
- Check pin depth by connecting plug to jack and use multimeter to verify continuity (Figure A.5)
- Crimp using 0.128" hex die

### **A.1.3. D-sub Connectors**

*Step 1*

- Remove 0.75" outer insulation and braid
- Remove *additional* 0.5" outer insulation to expose braid

*Step 2*

- Fully remove 1/16" inner dielectric from end using 0.075" setting on ringer




Figure A.5.: A coaxial cable with a layer of heatshrink partially covering the outer conductor and knurled portion of the MMCX plug. The plug is also connected to a jack in order to check for continuity. Not shown here is a crimp sleeve to the left of the heatshrink, which is pushed over the heatshrink tube, pushing the heatshrink to the right in the process. The layer of heatshrink prevents the edge of the crimp sleeve from biting into cable and severing the insulation layers. It also helps immobilize and protect the outer conductor wires which have a tendency to sever after being crimped.

*Step 3*

- Partially remove additional 1/16" inner dielectric (1/8" from end) using 24 awg notch on black stripmaster (Figure A.6)

*Step 4*

- Fully insert inner conductor into sleeve
- Gently crimp sleeve around inner dielectric

*Step 5*

- Add solder into gap between leaves of sleeve until inner conductor is encapsulated and continuity test is passed




Figure A.6.: Prepared cable prior to insertion in female D-sub connector. This photo shows an early iteration where the outer conductor was peeled back, instead of being soldered to a copper braid.

#### A.1.4. Ring Terminal (Bottom PMT)

##### *Step 1*

- Remove 0.75" outer insulation and braid
- Remove additional 0.5" outer insulation

##### *Step 2*

- Remove nylon sheath from ring tongue using heat gun and pliers

##### *Step 3*

- Expose 1/16" inner conductor

##### *Step 4*

- Insert inner conductor into crimp sleeve such that tip of inner conductor is roughly flush or slightly past end of sleeve
  - Inner dielectric should be inserted such that it enters approximately half of the crimp sleeve
- Gently crimp sleeve onto inner dielectric using pliers

##### *Step 5*

- Apply copious amounts of solder to fully immerse the inner conductor in solder
  - This is easiest when gravity assists to cause solder to flow into the sleeve

### A.1.5. Sub-C Connector (FEP Cable)

#### *Step 1*

- Slide heatshrink over cable
- Cut outer insulation at 1.5" using ringer at 0.055"

#### *Step 2*

- Cut outer insulation and braid at 1.25" using ringer at 0.065" (3 rotations)

#### *Step 3*

- Cut outer insulation, braid, and inner dielectric at 0.100" using ringer at 0.080" (3 rotations)

#### *Step 4*

- Prepare outer conductor leads by cutting 2.5" segments of 24 AWG stranded wire

#### *Step 5*

- Strip 0.100" insulation from one end of lead wire using 24 AWG notch on black stripmaster

#### *Step 6*

- Crimp sleeve to one end of ground wire

#### *Step 7*

- Remove all segments of insulation, braid, and inner dielectric from cable
  - Use parallel cutter to remove long sections of insulation and braid

*Step 8*

- Tin tip of inner conductor
- Insert into connector
- Crimp sleeve onto inner conductor
- Feed ground wire into adjacent hole
- Cut and strip end of ground wire to match location of braid on cable
- Tin braid
- Solder ground wire to braid
- Apply heatshrink

## Appendix B.: Supplementary Information for XeNeu

### B.1. Table of Slow Control Sensors and Instruments

The following table is presented for more information on the sensors and instruments shown in Figure 4.13 and Figure 4.16.

Table B.1.: Table of XeNeu Sensors and Instruments

Sensor	Description
T_BottomFlange	Useful during cooldowns and warmups for identifying the first and last quantities of liquid xenon to enter and exit the TPC.
T_Cond	Critical for determining the heat load to be delivered or removed from the detector. Its value is controlled by a PID loop on the Pwr_Cond heater. Typically set to 178.4 K, maintained with roughly 8-10 W of power from Pwr_Cond
Pwr_Cond	Offsets surplus cooling power from the cryocooler; effectively controls the detector temperature.
Pwr_Bottom	Provides additional heat during a warmup.
T_TopFlange	Intended for estimating the gas temperature in the detector space.

---

Sensor	Description
T_HX	Monitors the temperature of liquid exiting the weir and entering the HX.  Especially useful for determining when liquid has reached the heat exchanger, an important step during the cooldown process.
T_ColdValve	Monitors the temperature of incoming gas before the condenser.
PT_Detector	The nominal operating pressure is 1.2 Bar, which is monitored with  PT_Detector.
T_Weir	The three weir RTDs are used to coarsely detect the liquid level in the wier.

---

## B.2. Lessons Learned

### B.2.1. The XeNeu Vacuum Box

Details here on the box’s construction are provided secondhand following conversations with Sean Durham (LLNL), who was heavily involved with the machining and fabrication process of the box. While the information given here predates this work, it had yet to be documented elsewhere. In general, the majority of the challenges in

The vacuum box itself was made from 0.250 inch thick 316L stainless steel plates which were welded together. The thickness of the plate was driven by a desire to reduce weight, but it added severe constraints to the box’s construction. Detachable flanges for the box needed to be bolted to threaded studs (essentially set screws) on the walls of the box (Figure 4.1, Figure 4.14). Studs and nuts were chosen with the intent that the studs could be left in place to avoid unnecessary wear on the tapped holes. The holes for these studs required a custom tap be made to maximize the

usable threading depth in each of the holes. On multiple occasions, the tapping process punctured the walls of the box, requiring spot welds for repairs. 1/4"-28 threaded fasteners were chosen over 1/4"-20 fasteners to ensure as many threads as possible would be engaged, however the finer thread was more susceptible to wear. In the event that the tapped threads would be damaged, repairing with a threaded insert would not be possible because of the shallow depth of the threaded holes. Re-tapping to a larger size fastener would also be unlikely to work because of space constraints for the holes on the flanges.

The construction process was complicated further by the need for the sides of the box to seal against o-rings on the undersides of the detachable flanges. Typically, a sealing surface must be milled flat and brought to a specific finish, removing material from the surface in the process. But with a minimal wall thickness constraining the depth of the tapped holes, the requirements for a sealing surface added significant tension to the design. This tension came into play on numerous occasions. Despite efforts made to mitigate heat-induced distortion, initial welding of the box's panels resulted in warping which necessitated the removal of more material to achieve a sealing surface. On three separate occasions during vacuum tests, sealing surfaces on the box had to be refinished in a mill before internal supports were added to prevent the box from bowing inwards under pressure. The distorted walls of the box also added to the complexity of zeroing it on a mill because machining surfaces flat for an easier zeroing process would remove precious amounts of material. Finally, because the stock material had already been milled to size (0.250 inch thickness), excess mill scale could not be removed in some areas without significantly altering the material's dimensions. This mill scale again contributed to the difficulty of zeroing components on a mill.

The detachable flanges were cut to size from a piece of stock using a shear, which added

a significant degree of curvature to them. Normally, the curved sides could be machined flat, however the flanges required an o-ring groove which was constrained by the total thickness of the stock material. To maintain enough material for the o-ring groove, only the sealing side of the flange could be machined flat. The curvature on the opposing side and residual mill scale on the edges of the flanges were left in place, impeding the process of zeroing and squaring the parts on a mill. Reaching an acceptable vacuum with the flanges was a tedious process of trial and error. Ultimately, multiple half-thou finishing passes were required at certain spots in the grooves to correct for the curvature in the components that prohibited them from being squarely mounted on a mill.

### **B.2.2. CAD Software and Chain-of-Custody**

Replacing the PTFE bottom PMT holder was a fairly challenging process, taking roughly a year from the start of design iteration to installing the final components in XeNeu. Among other technical challenges, multiple months were spent painstakingly measuring dimensions of existing components, as CAD files for the detector had been lost, omitted, or not kept up to date with the as-built detector. Additional difficulties came in the form of sourcing properly licensed CAD software. During the design process, modeling was done with a student copy of SolidWorks. Designs were then communicated to Angelique Mangoba, a technician at LLNL, who possessed a properly licensed copy of SolidWorks and could recreate STEP files which were sent to her. These challenges spurred the migration of existing CAD files and future CAD work to OnShape, which has more permissive licensing for education use. OnShape has a much more polished workflow for collaborative work (including git-like version control) which reduces the chance of designs being lost.

## **Appendix C.: THERANOS Procedures**

### **C.1. RGA Stand P&ID**

See Figure 5.25 for the most recent THERANOS P&ID at the time of writing.

### **C.2. XeNeu P&ID**

See Figure 4.13 and Figure 4.16 for the detector and circulation panel P&IDs, respectively.

### **C.3. RGA Calibration Procedure**

#### **C.3.1. Hardware Connections**

The RGA Stand is connected to V9 on XeNeu for two reasons:

1. V9 on XeNeu is extremely close in height to V2 on the RGA Stand. This results in a very short transfer path.
2. V9 is next to the getter pressure transducer on XeNeu, which lets us monitor the input pressure of xenon to the sample volume MV1 in the RGA Stand.




Figure C.1.: RGA Stand Xenon Inlet connected to V9 on XeNeu

### C.3.2. Pressure Guidelines

**⚠ Warning**

The turbopump, ion gauge, and RGA should never be exposed to high pressures when they are operating! Make sure to turn them off before introducing large quantities of gas.

	Maximum Safe Sustained	
Item	Pressure During Operation	Comments
Turbopump	$10^{-3}$ torr	The turbopump can easily manage brief pressure spikes of up to $10^{-1}$ torr. Example: opening V5 with 1 torr of gas in MV2 is fine. If more gas is being pumped out, shut off power to the turbo and slowly feed gas in. A sustained pressure of $10^{-2}$ torr will safely spin it down.
RGA200	$10^{-4}$ torr	Always check: Is the filament off? Does it need to be on?
Ion Gauge (IG1)	$10^{-4}$ torr	Will automatically trip if pressure spikes too quickly, but this should not be relied on. If IG1 trips, it can be reset by turning it off through the menu, then turning on.

## **C.4. Initial Vacuum**

1. Vacuum in XeNeu gas delivery lines is established with an external pump on V4 of **XeNeu**.
2. Vacuum in the RGA Space is established with the turbo pump through the manual valve (V6).

1. The rest of the plumbing can be pumped through V5 and LV1

### **C.4.1. XeNeu Pumpout**

1. Start with all valves closed
2. Turn on the backing pump to establish a weak vacuum at V4 and mitigate air ingress
3. Open V6, V9, and V4 and pump out the gas delivery path up to the RGA stand with the backing pump, and then turn on the turbo.

1. V3 and R1 are left closed because their volumes only ever have xenon or are cryopumped out during xenon recovery

### **C.4.2. RGA Stand Pumpout**

For long term storage, close LV1 and V6 to isolate the RGA. If the RGA is being used multiple times per week, keep V6 open as it is better to leave the turbo pumping on it overnight. The same can be said for the RGA Electronics Control Unit (ECU) which requires 0.5-1 hr to fully stabilize when it is turned on.

1. If it is not on, turn on the RGA ECU so the electronics can stabilize. If you do not know how this differs from the RGA filament, DO NOT proceed, and instead refer to the RGA manual.

2. Keep the valve between the RGA Stand and XeNeu (V2) closed to isolate the two systems.
3. If the argon cylinder is not connected to V1:
  1. Purge the line from the argon cylinder with argon (~few psi) for a few minutes to clean and dry out the line.
  2. While the purge is ongoing, connect the line to the carrier gas inlet (V1), then close the regulator on the argon cylinder.
  3. Open V1 to expose the argon line to MV2
4. Open the valve in front of the capacitance manometer CM1 (V4) and V3, the valve which isolates MV1 from the rest of the sampling space.
5. Open V8 and V9 to expose the circulation pump to MV2.
6. Begin pumping out MV2 and connected spaces through V5 and/or V1 depending on how much excess argon is in MV2.
7. Once the pressure in MV2 is below  $10^{-4}$  torr, open LV1 and V6 slowly to let the RGA equalize in pressure with MV2.
8. The system should reach a pressure of  $10^{-6}$  torr or lower before proceeding.

## C.5. Gas Mixing

### C.5.1. XeNeu Gas Panel Setup

1. Close all valves on XeNeu gas panel
2. If a slow control instance on XeNeu has not been started:
  1. Start a slow control instance on `xenonsc` to read out pressures from  $P_{Getter}$ , the pressure

transducer upstream of V9 on XeNeu

3. Open the valve of a xenon bottle
4. Open the regulator (R1) and use the dial indicators to roughly set the gas pressure ahead of V3.

 Warning

If making multiple mixtures, start with the settings for the lowest concentration first so that the pressure for  $P_{Getter}$  can subsequently be increased.

1. Open V3 and V6 on XeNeu

### C.5.2. Xenon Addition to MV1

1. Isolate MV1 on the RGA Stand by closing V2 and V3
2. Open V9 on XeNeu and V2 on the RGA stand to fill MV1 with xenon.
3. Wait for the pressure to stabilize on  $P_{Getter}$ , then record the pressure.
4. Close V9 on XeNeu and V2 on the RGA stand once a sample has been stored in MV1

### C.5.3. First MV1 Expansion to MV2

The system should be at a pressure of  $\sim 10^{-6}$  torr or lower. 1-2 bar of Xe should be in isolated in MV1.

1. Close LV1 and V5 to isolate the turbo
2. V11, V8, and V9 and V4 should be open since they were being pumped on.
3. Carefully crack open V1 to add  $\sim 1$  torr of argon to “shield” MV2 from xenon adsorption.

4. Record the reading of the capacitance manometer  $V_{BCM1}$  (CM1).
5. Close V8 and V9 to isolate the circulation pump.
6. Open V3 to volume share MV1 and MV2
7. Record the first expansion xenon pressure,  $P_{Xe1}$  from CM1.

1. Make sure to subtract off the argon pressure,  $V_{BCM1}$ !

8. Close V3 to preserve this xenon-rich mixture in MV1
9. Fill MV2 with ~1000 torr of argon through V1 to some pressure  $P_{tot}^{i1}$
10. Calculate and record this xenon concentration:  $\phi_{Xe}^e = P_{Xe1}/P_{tot}^{i1}$
11. Open V8 and V9 to mix this with the circulation pump
12. Turn off the circulation pump and record the total pressure:  $P_{tot}^{f1}$
13. Close V7, then reduce this pressure by some factor  $x_r$  using V10
14. Close V10, open V7, then open V1 and refill back up to  $P_{tot}^{f1}$

1. The xenon concentration  $\phi_{Xe}^e$  will now have been reduced by  $x_r$

15. This gas can now be sampled starting from Section C.6.1.

#### C.5.4. Second MV1 Expansion to MV2 and Argon Addition

The system should be at a pressure of  $\sim 10^{-6}$  torr or lower. Approximately 1 torr of Xe should be isolated in MV1.

#### ! Important

February 27, 2022: PT1 and CM1 have not yet been cross calibrated. It is recommended to only rely on CM1 until PT1 has been calibrated to CM1.

1. Verify that V1, LV1, V5, V8, and V9 are closed.
2. Open V3 to volume share MV1 and MV2
3. **Very Carefully** open V1 to begin adding Ar to MV2
4. Close V1 when CM1 or PT1 read as desired. Record the argon pressure added,  $P_{Ar}$ .
5. Open V8 and V9 to expose the circulation pump to the sample
6. Turn on the circulation pump and left the system mix for at least 5 minutes.
7. This gas can now be sampled starting from Section [C.6.1](#).

## C.6. Sampling

### C.6.1. MV2 Pressure Reduction to 10 torr

System pressure will drop slightly after the circulation pump has been exposed to MV2. The circulation pump should be left on for the duration of the sample to keep the gas mixture agitated.

#### Warning

Continuously monitor the value of CM1 through this process. It has a max reading of 13.60 V, which is out of the normal 10 V range that corresponds to 100 Torr. Proceed slowly once CM1 reads under 13.60 V!

1. V1, LV1, and V5 should be closed.
2. Close V7 to isolate the turbo from the backing pump.
  1. You can leave the turbo on for a few minutes since it's not pumping any gas.
3. Open V10 to drain out excess gas from MV2

4. Close V10, then reopen V7.

### **C.6.2. Priming RGA and LV1: MV2 Pressure Reduction to 1.5 torr**

1. Turn off IG1. Make sure the RGA is also turned off.
2. Open LV1 to start clearing gas trapped between it and AP1.
3. Carefully open V5 to drain MV2 to 1.5 torr
  1. Because the input pressure is low, the turbo does not need to be turned off here.
4. Monitor CG1 and do not exceed  $10^{-3}$  torr for sustained periods while the turbopump is powered and running.
5. Close V5 when CM1 reads 1.5 torr
6. Turn on IG1 when CG1 is  $< 10^{-4}$  torr
7. Verify that IG1 is reading  $1 \times 10^{-5}$  torr. If not, adjust LV1 until it is.
8. Turn on RGA Filament and start scanning!

### **C.6.3. Finishing a Scan**

1. Turn off the RGA filament and IG1
2. With  $\sim 1$  torr of gas in the RGA space, V5 can be opened immediately
3. Pump out MV1, MV2, and the circulation pump
4. Close V8, V9, and V3 to isolate the circulation pump and MV1
5. Close LV1 and V6 to isolate MV2 from the RGA
6. Open V2 to load more xenon into MV1
7. If a sample was saved in MV1, it can be expanded again and diluted with argon to make a

new sample.

#### **C.6.4. Finishing for the Day**

1. Shut off XeNeu Slow Control
2. Close V3, V6 and V9 on XeNeu
3. Make sure V5 and LV1 are closed so the turbo can keep pumping on the RGA
  1. If the RGA is not going to be used for >week, it is better to close V6 and LV1 to isolate the turbo.
4. Power off turbo
5. Power off backing pump and close V7
6. Optionally power off the RGA ECU

### **C.7. Calculations**

#### **C.7.1. Dilution Factor**

- The reading from  $P_{Getter}$ , CM1 baseline, and CM1 reading after expansion are used to calculate the dilution factor:  $D_F = P_{Getter}/P_{Xe1}$
- The dilution factor lets us predict what the resulting pressures will be when gas in MV1 is expanded into MV2.

#### **C.7.2. Expected Concentration - Double Expansion**

- The xenon pressure after the second expansion can be calculated by taking the measured xenon pressure after the first expansion and dividing by the dilution factor:  $P_{Xe2} = P_{Xe1}/D_F$

- The expected xenon concentration after the second expansion can be calculated by dividing  $P_{Xe2}$  by the pressure of the argon added:  $\phi_{Xe}^e = P_{Xe2}/P_{Ar}$

### C.7.3. Expected Concentration - Drain Refill

- Instead of a double expansion, the xenon concentration can be lowered by sequential dilutions with argon.
- After the first expansion from MV1, the xenon pressure is  $P_{Xe1}$ . Adding more argon does not change this! Thus, argon can be added until a mixture at some intermediate concentration is made.
- Reducing the total pressure  $P_{tot}$  of this mixture by some factor  $x_r$  will reduce the pressure of  $P_{Xe1}$  by the same factor such that  $P_{Xe2} = P_{Xe1}/x_r$
- Thus, if the total pressure is restored to its former value ( $P_{tot}$ ) via the addition of more argon, the new concentration will be  $\phi_{Xe}^e = P_{Xe2}/P_{tot}$

### C.7.4. Calibration Factor

- A calibration factor is a value which when multiplied by the output of the RGA readout, yields the “true” reading. For example if the ratio of  $^{129}\text{Xe}$  and  $^{40}\text{Ar}$  are used as a proxy of the concentration of Xe in Ar, then a calibration factor  $k_{129/40}$  multiplied by the readout of the same ratio from the RGA,  $R_{129/40}$ , should yield the measured concentration of Xe in Argon,  $\phi_{Xe}^m$
- The calibration factor can be determined by comparing the expected  $\phi_{Xe}^e$  from a sample of known composition to the measured  $\phi_{Xe}^m$

## C.8. Recording Data

Talk to Eli about this, still a work in progress.

### C.8.1. Recorded Parameters

1. Add a new row to the spreadsheet for each acquisition (Ask Eli for access)
2. See notes in column headers (hover over them) for explanations

### C.8.2. Rough Calculations

1. Add a row to the spreadsheet for each acquisition (Ask Eli for access)
2. This sheet is simply used for “spot checks” during a scan. It should NOT be relied on for any analysis.

## C.9. RGA Software Configuration and Setup

### C.9.1. Scan Speed

**!** Important

Scans completed with different scan speeds will have different calibration factors.

Scan Speed is a setting that can be changed for any scan. It is the rate at which the RGA scans over masses, i.e. it has units of amu per second. A faster scan speed introduces more noise whereas a slower scan speed may not cover rapidly changing conditions at high pressures. For our purposes a scan speed of 2 is sufficient.

Table C.2.: RGA scan speed table from SRS [186]. Note how the scan speed and noise floor are inversely related.

Scan Speed	Noise Floor	Single Mass		
		Scan Rate [ms/amu]	Measurement Time [ms]	Baseline Noise (std dev) [Amps]
1	0	2000	2200	7E-15
2	1	1000	1100	1E-14
3	2	400	440	1.5E-14
4	3	200	220	2E-14
5	4	126	139	4E-14
6	5	45	50	1.2E-13
7	6	30	33	2.5E-13
8	7	15	16.5	5E-13

### C.9.2. Electron Multiplier (CEM or CDEM) Operation

#### ! Important

The CEM is a “wear-and-tear” part. Because of this you can expect the calibration factor to change over the lifetime of the CDEM.

The Electron Multiplier is a secondary amplifier designed to greatly increase the RGA’s sensitivity to masses below  $10^{-9}$  torr. It should **never** be used to scan over a mass or mass range where the pressure is expected to exceed this limit, as this will greatly increase wear on the part or burn it out immediately.

### C.9.3. Pressure vs. Time Scans

Pressure vs. Time scans allow you to scan specific masses of your choosing, and independently toggle the CEM for each mass. In a pressure vs. time scan, the RGA will scan over a 0.6 amu interval centered on the desired mass in steps of 0.1 amu and then report back the peak in that interval. This interpretation is based on the RGAs description of a “miniscan” in single mass measurements.

#### C.9.3.1. Saving Data

##### Warning

Changing scan parameters will erase displayed data. Make sure to save when a scan is complete!

Whether a scan is stopped or running:

1. File → Save ACSII Data
2. Specify filename: YYYYMMDD\_XX.Xppm\_X.Xtorr\_PvsT

Scans can be overwritten while running, and viewed in real-time with `rgaplot`.

### C.9.4. Analog Scans

Analog Scans let you scan over a “continuous” mass range. Given a range of two masses, you can specify the scan resolution, i.e. points per amu, and toggle the CEM on or off for the entire mass range to be scanned over. More points per amu will result in a slower scan but with a better ability to distinguish peaks. This should be contrasted with scan speed which can increase measurement accuracy, but cannot unilaterally distinguish peaks in a mass spectrum unless the scan resolution is fine enough.

### C.9.4.1. Saving Data

#### Warning

Changing scan parameters will erase displayed data. Make sure to save when a scan is complete!

Analog scans feature scan logging, i.e. successive analog scans will be grouped together after the scan has been started. When the scan is stopped, logging stops, and when another scan is started, a new scan log is created. The previous scan log will still be saved, but best practice is to export the scan log to a set of ASCII files as soon as it is finished.

After a scan has been stopped:

1. File → Open Scan Logs (open most recently saved file)

In the scan log viewer:

1. File → Save Scan Log
2. Specify directory and folder name: YYYYMMDD\_XX.Xppm\_X.Xtorr
3. Save ASCII data in chosen folder.

## C.10. Using the Keithley 6510 for Pressure Readout

The Keithley 6510 can record the outputs of CM1 and PT1. It is very convenient to use for real-time readout as well as logging pressure values during samples.

### C.10.1. Pressure Sensors

Pressure Sensor	Range	Scale	Torr / V	6510 Channel
PT1	0-3.4 Bar	0-10 V	255.021	109
CM1	0-100 Torr	0-10 V	10	110

### C.10.2. Key Front Panel Features

- Trigger Mode
  - Tap to access in top right
  - Continuous trigger mode displays values of individual channels in real-time
  - Scan trigger starts a scan with trigger settings programmed into scan (e.g. collect data once per second)
  
- Active (Closed) Channel
  - Select which relay to close to take a measurement from that channel
  - You can select which channels to close and what measurement to take (e.g. DC Voltage) from the specified channel.
  - Manually cycle through the active channels in continuous trigger mode to switch readout between PT1 and CM1.
  
- Watch Channel
  - Can be used to specify relative measurements between channels. Example: Set CH109 to active but display readout of CH110.
  - This is a confusing option but setting this to “all” seems result in “normal” behavior

where the active channel and watched channel are identical.

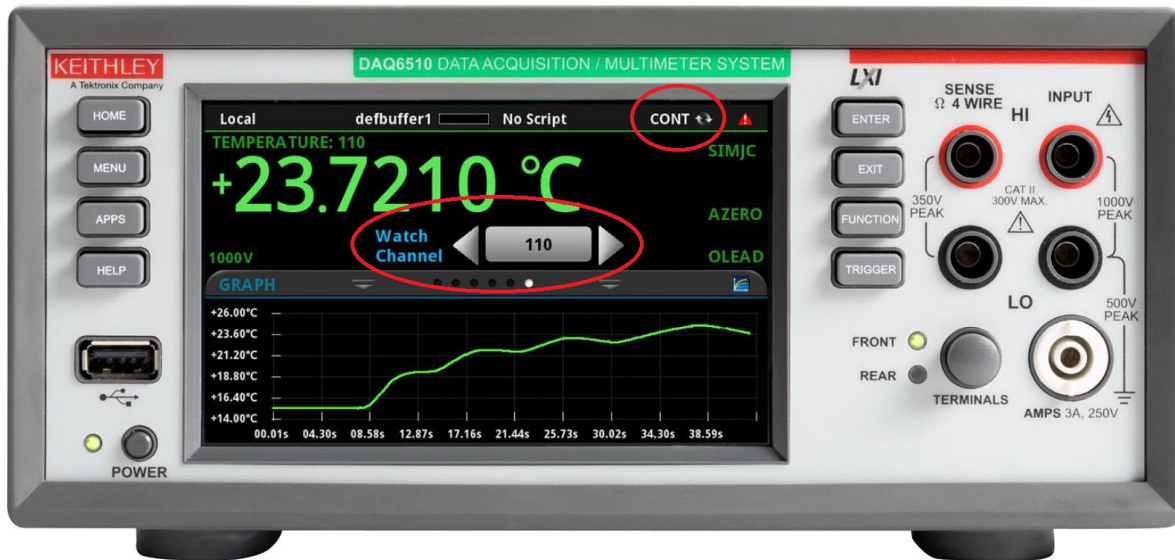


Figure C.2.: Front panel of the Keithley DAQ 6510. Image courtesy of Keithley.

### C.10.3. Real-time Readout

- Set the trigger mode to continuous for realtime readout of any one channel.
- This is the main option for reading out CM1 while draining MV2 or adding argon during calibrations, etc.

### C.10.4. Channel Scan

- Scan function under Menu for continuously logging data from specific channels
- Complex acquisition and triggering setups are possible with the interface on the 6510, but don't let this be discouraging—the interface is fairly user-friendly!
- A basic scan over PT1 and CM1 where data is read out once per second is useful during scans

to gain some insight into pressure dependence.

### C.10.5. Scan Options

- **Scan to Scan Time:** Time interval between data acquisitions. 1 s is usually sufficient.

### C.10.6. Saving Data from a Channel Scan

1. After a scan completes, the option to save the scan to a USB is offered in the scan menu.
2. If a USB drive is inserted, save the scan and select desired options for data formatting.

## C.11. Parsing Scan Data with `rgaplot`

Streamlit application written by Eli [189] for quickly parsing and viewing scan data. Requires LLNL Gitlab access, and VPN to view documentation.

## C.12. Summary of Recommendations for Future Development of THERANOS

The intent of this list is to concisely summarize future work with THERANOS described in Section 5.5.3.1.

- Inlet pressures should be 1 torr or lower. Lower inlet pressures will maximize the opening of LV1, hopefully reducing the chance of non-ideal behavior. Per Ellefson *et al.* [183], a shorter scan duration may also be useful here.

- Tests should be conducted where the leak valve is slowly and continuously opened to maintain a stable pressure at IG1. This will further shorten the duration of scans.
  - If successful, LV1 could be replaced with a stepper-controlled leak valve. Little and Gordon [191] provide more information on the control theory for such a device.
- Samples should be fully exhausted so that low pressure fractionation behavior Equation 5.17 can be assessed.
- Samples should be thoroughly mixed with the XeNeu circulation pump.
  - If possible, the plumbing or electronics should be configurable for reversible circulation to ensure there are no exposed dead volumes.
- A matrix should be used to test permutations of samples each day during a calibration run. This could potentially expose underlying issues with measuring higher concentrations first, or whether the first scan on a given day is more susceptible to adsorption.

Table C.4.: Possible testing matrix. Rows are xenon concentrations in ppm.

Day 1	Day 2	Day 3
100	25	50
50	100	25
25	50	100

- Backgrounds should be characterized with no gas injected into the system, and also with “pure” argon to establish baseline xenon levels.

- Further development should take place to record information from IG1, CM1, PT1, and the RGA simultaneously.
- A procedure should be developed to create calibrated mixtures of gas with CHILLAX and its MFCs, to reduce reliance on XeNeu and potentially improve reproducibility.
- A more compact plumbing circuit and adequate heating could be considered to reduce possible backgrounds.
  - The SRS Universal Gas Analyzer [181] notably uses a compact 6-way “cube” at the RGA and a short elbow directly to a turbopump to minimize the background surface area.

## Appendix D.: Historical Review of Electron Trains

The following review summarizes numerous works that have come to define electron trains as they are known today. At times, commentary is included in these summaries to draw comparisons or highlight discrepancies between works. Below (Figure D.1, Figure D.2) are condensed timelines highlighting notable findings or claims from each work. There is a sharp divide between early studies which looked mainly at photoionization at short timescales, and later studies which have studied delayed backgrounds more intently.

### D.1. 2008 - ZEPLIN-II

Electron emission from the gas-liquid interface of a noble liquid has been known as a possibility for decades. However it seems as though the first description of electron *background* signals was in 2008 with a characterization of the single electron response performed during the ZEPLIN-II experiment [195].

In this study, “unexpected” pulses were seen following S1 and S2 pulses (Figure D.3), and it was confirmed that these were in fact single electrons (SEs) by looking at a histogram of the pulse area. These SE pulses were taken advantage of to characterize the S2 gain of the detector, and additional analysis was performed in an attempt to characterize these signals.

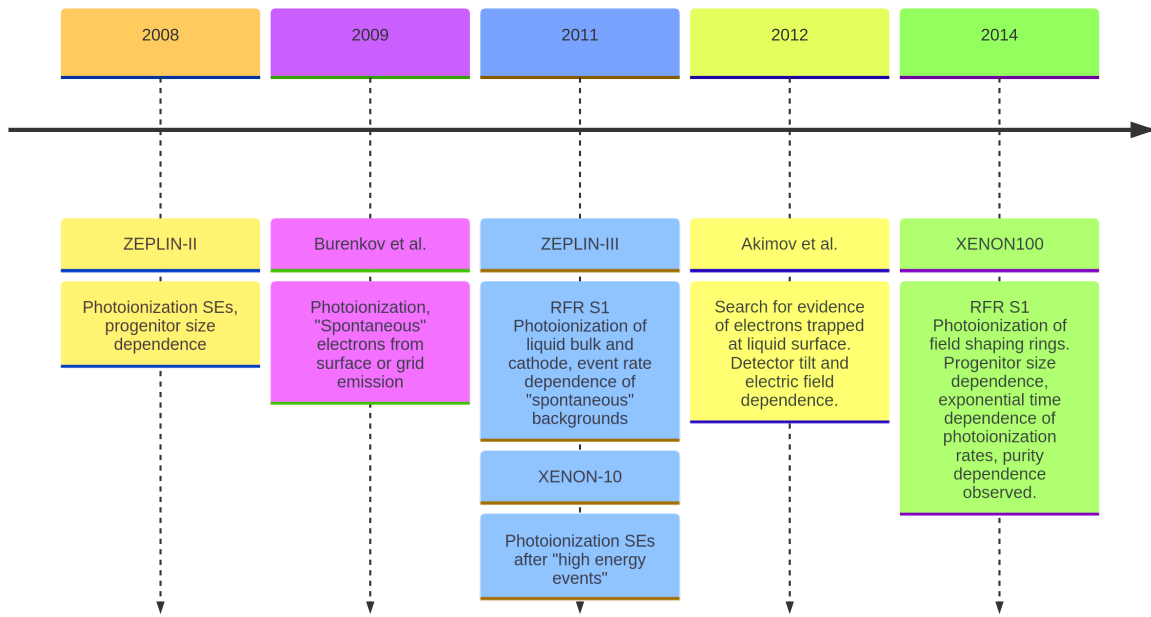


Figure D.1.: Early studies of electron backgrounds focused mainly on observations of electrons within a drift length or less following progenitors. “Spontaneous” pulses were still observed prior to events however this spontaneity was disproved in 2011 by ZEPLIN-III.

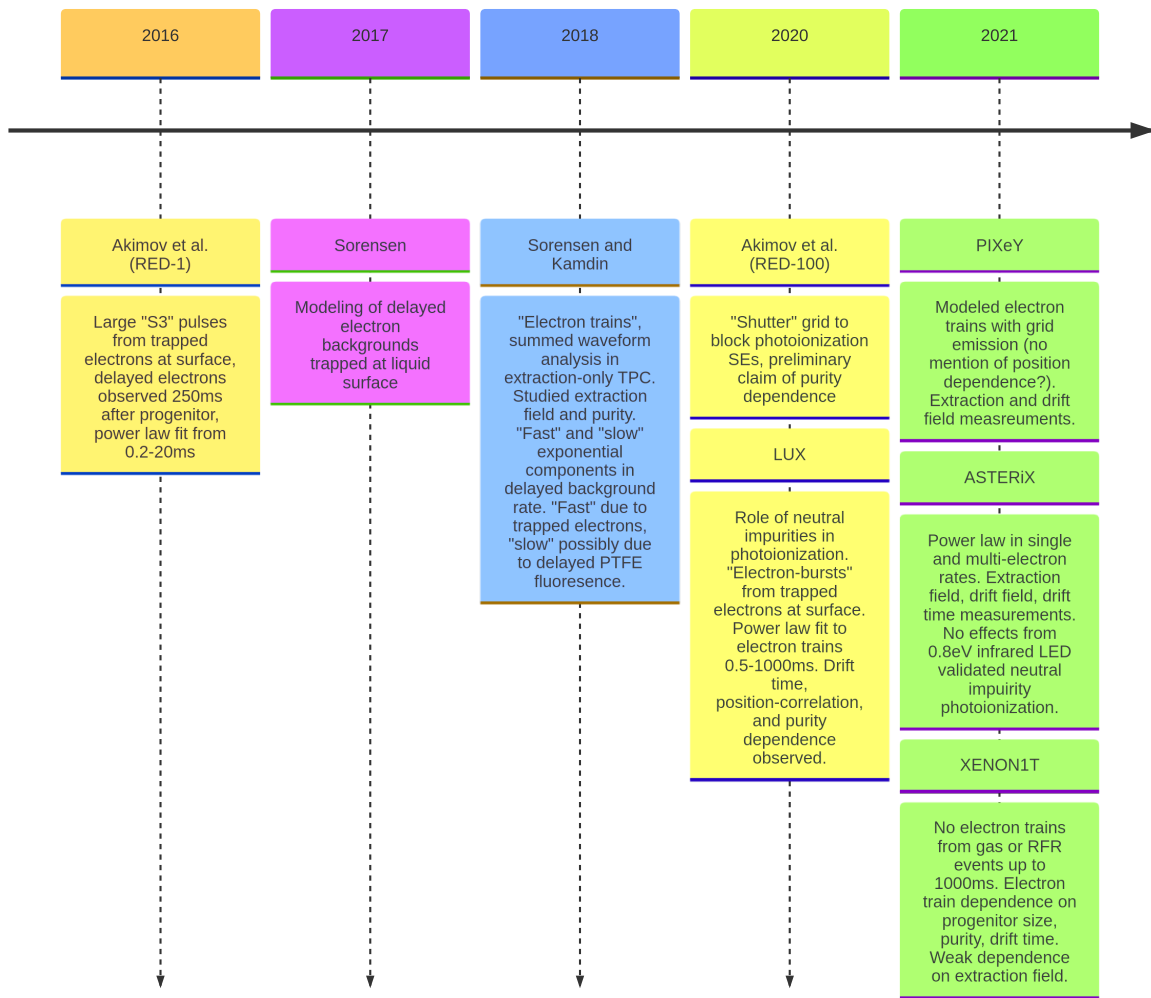


Figure D.2.: Later studies looked at backgrounds well after a full drift time and became increasingly detailed. A majority of studies have found a power law dependence in the rate, at timescales up to 1000 ms after a progenitor.




Figure D.3.: Example of a single electron pulse observed between an S1 and S2 in ZEPLIN-II.

The study concluded these SE pulses were anomalous in part because their radial position distribution throughout the detector seemed fairly even, in contrast to the radial position of other small S2 signals which were concentrated near the walls of the detector. The study also looked at SE pulses following S1 pulses and found that the delay time relative to the S1 progenitor was evenly distributed up to the full drift length of the TPC, about 70  $\mu\text{s}$ .

The study also examined the dependence of the SE rate on size of the progenitor S1. They acknowledged that this was motivated by observations of multiple single electron pulses following large S2s. They found a clear dependence on progenitor size, where more energetic events are followed by a higher number of SE pulses.

Based on their study they speculated some possible causes. One proposed cause was that electronegative contaminants could be photoionized since they had a very long drift time (0.7 cm/s at 1 kV/cm). However, they then claim this cannot explain the observed rate entirely because

it would not be consistent with concentration estimates from the electron lifetime. Another proposal was that VUV light from xenon could induce a specific excitation mode of another xenon atom (“n=1”) which could transfer energy to an impurity and ionize it in a Penning-type process. They concluded by suggesting that the spurious SE pulses could be produced at least in part by photoionization.

## **D.2. 2009 - Burenkov *et al.***

In a follow-up of the ZEPLIN-II study, Burenkov *et al.* observed SE pulses in an unnamed dual-phase xenon TPC [196]. In this TPC they saw SE pulses which they believed were originating from the photoionization of impurities. Specifically, they thought that residual calcium from a calcium getter used to purify xenon could be the source of the impurities.

The study also mentions a preliminary analysis for a prototype of ZEPLIN-III. Unlike in ZEPLIN-II (Section D.1), here they mention *two* distinct types of SE backgrounds. The first were believed be photoionization SEs. However the second type appeared to be spontaneous (~20 Hz) with no apparent dependence on previous events. Here they speculated on two possible causes for these so-called spontaneous backgrounds. The first cause of these spontaneous SEs could have been due to grid emission at the cathode, noting that the work function near the cathode would be reduced in liquid by 0.67 eV. They also mention that grid wire defects could increase the probability of emission. A second possible source is given as trapped electrons at the liquid surface which they believed would accumulate in large numbers due to imperfect extraction.

### D.3. 2011 - ZEPLIN-III, XENON10

ZEPLIN-III performed a dedicated study of “single electron emission” in 2011 [98]. They performed some short dedicated runs with 256  $\mu\text{s}$  event windows instead of their WIMP search event windows 36  $\mu\text{s}$  in length, and maximum drift time of the TPC was stated to be 14.5  $\mu\text{s}$ . A very high external trigger rate was used for the dedicated runs. In their study of SE backgrounds they also applied an upstream veto where events were not accepted if there were other pulses detected in the prior 20  $\mu\text{s}$ —intentionally longer than than the maximum drift time. They also applied a radial cut to exclude events closer to the walls of the TPC.

As in the previous study, no position dependence of the backgrounds was observed. The photoionization hypothesis was more rigorously tested by looking at S1 pulses in the reverse field region below the cathode. These events are useful for photoionization tests because S1s in the RFR do not have corresponding S2s since any liberated charge will be drifted downward and away from the extraction region of the TPC. Data from these progenitor pulses was very conclusive in demonstrating that photoionization was a significant cause of electron backgrounds. Following an RFR S1, SE pulses were observed up to the maximum drift time of the TPC (Figure D.4) and an additional spike in the rate was observed at the maximum drift time that was indicative of the cathode being photoionized. This was in addition to an observed increase in rate with larger S1 pulses that had previously been seen in ZEPLIN-II.

The topic of spontaneous SE pulses was also addressed in this paper<sup>1</sup>. To study them they looked for SEs *before* S1 pulses in WIMP search data for higher statistics. They found a rate of about 20 Hz which was 4 times higher than in their first dedicated study but they attributed

---

<sup>1</sup>The word spontaneous was used in quotes in this study (e.g. ‘spontaneous’), probably as a way to signify that pulses of this type may not have been entirely independent of previous events,




Figure D.4.: SE counts following RFR S1 pulses in ZEPLIN-III. Counts followed a decreasing trend up to the full drift time where a spike in the counts was observed. The cause of this spike was inferred to be from photoionization of the cathode due to its consistent location at the maximum drift time in the TPC. The decreasing trend was ascribed to electrons being captured by electronegative impurities. From a previously calculated value for the photon mean free path for photoionization they were able to estimate the original number of photoelectrons produced following an interaction. This led them to obtain an estimate of the electron lifetime from the slope of the trend.

this to having shorter veto periods. In their second WIMP search they saw the spontaneous SE rate decrease to  $\sim 1$  Hz after upgrading their PMTs to ones with much lower inherent radioactivity. Another difference was that the drift field was 13% lower and the field at the cathode wire was over 30% lower (38 kV/cm down from 60 kV/cm). Consequently, no conclusion was drawn in regards to the factor 20 rate decrease because of these potential confounding factors. In hindsight, it seems likely that the lower activity PMTs or some kind of grid background was suppressed.

Following their second WIMP search, a second short dedicated SE data set (i.e. long event windows and 20  $\mu$ s upstream veto) was taken with a  $^{137}\text{Cs}$  source. Here they found a large increase in the spontaneous rate to 120 Hz. Increasing the upstream veto in a reanalysis to 40  $\mu$ s and 60  $\mu$ s reduced this rate to 55 Hz and 45 Hz respectively. This was a major development in the subject of spontaneous backgrounds, as they were able to conclude that the so-called spontaneous rate depended significantly on the rate of energy deposition in the TPC and that the rate decreased with time. In the text this is referred to as a “delayed electron emission mechanism” and in hindsight this seems characteristic of electron trains as they are known today.

Again, two possible causes were given for these delayed electrons. First, they considered trapped electrons at the liquid surface. They believed this could occur in part because the electron extraction efficiency was 83% in initial their runs (first science run and dedicated SE study) and 66% in the latter two (second science run and  $^{137}\text{Cs}$ ). For electrons stuck at the surface they found it plausible that electrons could still be emitted later and that the emission rate decreased (as observed) because of the electron lifetime—45.4  $\mu$ s in  $^{137}\text{Cs}$  dataset.

Grid emission from the cathode was proposed as the second possible cause, as in the Burenkov study (Section D.2), because the field dependence of emission would be extremely steep according

to the Fowler-Nordheim model. They claimed that a reduction in the field at the cathode wire could be responsible for the factor of 20 decrease in rate observed after upgrades between first and second science runs.

### **D.3.1. XENON10**

In the same year as the ZEPLIN-III study, XENON10 published results of a low mass DM search [197]. They evaluated their S2 background spectrum and specifically noted the presence of SE backgrounds at least 20  $\mu\text{s}$  after “high energy events”. They cited ZEPLIN-II for the hypothesis of photoionization being a possible cause. In retrospect this is probably somewhat accurate given that XENON10 had a maximum drift time of about 75  $\mu\text{s}$ . They also gave trapped electrons at the liquid surface as another possible cause. There was little other discussion on these backgrounds as the focus of the paper was on dark matter search results.

### **D.4. 2012 - Akimov *et al.***

In 2012, another dedicated study of SE backgrounds was published by Akimov *et al.* [95]. In it, they proposed three processes were responsible for SE backgrounds: emission of trapped electrons at the liquid surface, photoionization of impurities, and grid photoionization. In regards to the first hypothesis, they modeled the emission of electrons past the potential barrier at the liquid gas interface. They noted that after thermalizing within  $\sim 1$  ps, electrons should be captured relatively quickly by impurities which persist at the surface.

For experimental data they selected events with no prior S2 more than multiple drift lengths away, and they also excluded any event with an S1 pulse or larger. Like in the ZEPLIN-III study,

the goal was to avoid looking at any photoionization backgrounds. Three tests were performed where they measured the dependence of SE rates on the applied electric field, with an external gamma source, and with a physical tilt applied to the detector.

From their measurement they believed that electrons they saw were from those trapped at the liquid surface. Two observations contributed to this. First, the rate increased with the electric field<sup>2</sup> and second, the rate lowered when a tilt was applied, where tilting in either direction resulted in a lower rate compared to the rate with no tilt (Figure D.5). The rate increase with electric field was in agreement with their model which predicted a rate  $f_{1e} \approx 1/t_e$  where  $t_e$  was the emission time of electrons trapped at the liquid surface, which decreased with the electric field.

In all this is a fairly remarkable study because it is the first test of a rate dependence on an electric field, and the only test of detector tilt angles. In 2016, the same group published a plot demonstrating a power law in the electron background rate after high energy events (Figure D.9).

## D.5. 2014 - XENON100

XENON100 released a major study of SE backgrounds in 2014 [97]. In their study they presented a measurement of the time delay of small S2 (1,2,3 SE) signals following S1 and S2 pulses in single scatters. The maximum drift time of the TPC was 180  $\mu\text{s}$  and the maximum time delay they looked at was 200  $\mu\text{s}$ , so the SE pulses they observed were heavily dominated by those produced by photoionization. They selected S2s from single scatter events “to avoid confusion in the time association” of SE pulses.

---

<sup>2</sup>This is given as the voltage difference between the anode and cathode in the paper.


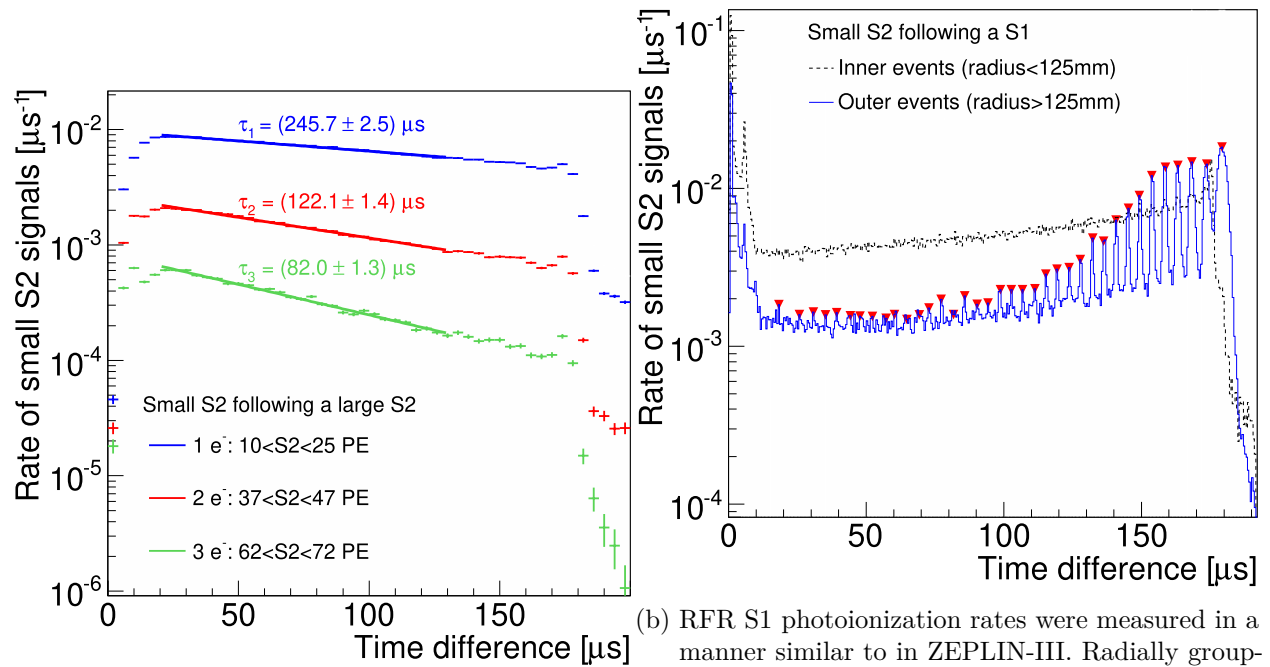


Figure D.5.: SE count rates observed by Akimov *et al.* dropped when tilting their TPC in either direction, possibly indicating that electrons trapped at the liquid surface were being removed before being able to emit later.

### D.5.1. Photoionization Rates

Fitting to the time delay of pulses after progenitor S2s, they found an exponential decrease in the rate with time for 1,2, and 3 SE pulses (Figure D.6a). They concluded that the rates from double and triple electron pulses were mostly due to pile-up in time because of the relation between their time constants and because of examinations of hit patterns of PMTs. Given that these were prompt backgrounds being studied, this is consistent with contemporary observations.

Looking only at RFR S1s, they divided small S2 populations in to two categories based on radial position in the TPC (Figure D.6b). Small S2s near the walls had periodic rate spikes up to the max drift time that were consistent with photoionization from field shaping rings while the rate for centrally located small S2s showed a gradual increase up to the maximum drift time. Interestingly, this seems to contradict the behavior of RFR S1 single electron rates obtained by



(a) Rates of small S2 pulses observed after large S2 pulses in XENON100 were observed to exponentially decay with time. Note that the timescale here is entirely within a drift length, so these are distinct from delayed electron trains as they are known today.

(b) RFR S1 photoionization rates were measured in a manner similar to in ZEPLIN-III. Radially grouping progenitor S1s showed that detector components as well as the liquid bulk were responsible for photoionization backgrounds. This can be seen in the rates from progenitors near the walls, which were peaked at locations consistent with the field shaping ring spacing. Interestingly the trend appears to increase with time instead of decrease like in ZEPLIN-III ( Figure D.4). It is likely this was due to factors such as a longer electron lifetime.

Figure D.6.: S2 and S1 photoionization rates in XENON100.

ZEPLIN-III (Figure D.4), and is not commented on in the study. In the case of both XENON100 and ZEPLIN-III, no mention is made of drift time corrections being applied to this rate. However in the XENON100 study they mention that the small S2 rates after S2s and RFR S1s should decrease and increase with time respectively due to the photoionization spatial distributions of the respective pulses. That is, S1s should produce more photoionization at the bottom of the TPC where they occur, and S2s should do the same closer to the top of the TPC.


### **D.5.2. Rate Dependence on Purity and S2 Size, Uncorrelated Rates**

XENON100 performed additional analyses of small S2s where no other pulses above 150 pe ( $\sim 7-8$  SE) were present within 400  $\mu\text{s}$  (the maximum waveform length) of a progenitor S2.

First, they examined the rate dependence for small S2s on S2 size and found that the number of small S2s per event was linearly dependent on the progenitor S2 size (Figure D.7a). This was in good agreement with the analysis of SE rates following RFR S1s performed by ZEPLIN-II. The XENON100 study also mentioned that the extrapolated rate to a nonexistent progenitor was nonzero, implying an uncorrelated component to the electron background signals. These uncorrelated backgrounds were examined in the event pre-trigger window and they found  $\sim 1\text{E-}3$  electrons per event. The number of pre-trigger electrons decreased exponentially with a time constant of 2 ms when the time since the previous event increased. Little else is claimed about these but they put forward that the decrease in rate with time since the previous event could be indicative of trapped electrons at the liquid surface.

In more alignment with previous hypotheses regarding impurities, the small S2 rate within 20 to 150  $\mu\text{s}$  after a progenitor S2 was found to linearly depend on the O<sub>2</sub> equivalent concentration

in the liquid (Figure D.7b). This was the first quantitative demonstration that photoionization backgrounds were dependent on the purity of the liquid, and a strong source of evidence that impurities were being photoionized.



- (a) Integrated rates of small S2s following progenitor S2s (designated as “Main S2” in the figure) were found to depend linearly on the size of the progenitor in XENON100.
- (b) Integrated photoionization rates were also shown to linearly depend on the O<sub>2</sub>-equivalent concentration of impurities in the liquid bulk. in XENON100

Figure D.7.: XENON100 was able to confirm the energy dependence of photoionization rates seen in prior experiments like ZEPLIN-II, and demonstrate that the rates were dependent on the liquid purity.

Ultimately, the study concluded by ruling out trapped electrons for *prompt* electron backgrounds due to the above evidence and the abrupt rate dropoff after the maximum drift time. A lot of attention was given to justifying the photoionization hypothesis, and a passing remark was made that the remaining background rate after the maximum drift time could be due a “delayed extraction phenomenon”.

## D.6. 2016 - Akimov *et al.* (RED-1)

In 2016, Akimov *et al.* produced another dedicated study of “electron emission” following S2 signals and muons with the RED-1 (originally ZEPLIN-III) TPC [96]. They found a rate which was strongly dependent on the time since previous event and claimed to find evidence of trapped electrons at liquid surface. The TPC itself had a max drift time of 9  $\mu\text{s}$  and a drift depth of 22 mm. The active volume was bounded only by an anode and cathode so the drift field of 3.75 kV/cm was also the extraction field. Waveforms *268 ms long* were recorded for 100 muons. By this point since photoionization had been well established as a cause for prompt backgrounds, they gave more attention to delayed and spontaneous single electron background signals.

### D.6.1. “S3” Pulses

In a significant departure from previous studies with TPCs, they observed an “S3” signal no more than 200  $\mu\text{s}$  after S2 signals from high energy muons Figure D.8. The S3 signals preferentially occurred at a specific edge of the detector and in the intervening time between an S2 and S3, there were many SE pulses emitted. Despite having poor position reconstruction capabilities they were able to demonstrate that the SE pulses “moved” from the S2 to the S3 position. They also showed that S3 pulses closer to initial interactions had a shorter time delay before emission. Furthermore, if S3 pulses were due to trapped electrons then it would make sense for their width (i.e. duration) to increase due to diffusion of the trapped electrons on the surface. They attempted to model this diffusion and showed that later S3s had a larger width, in rough agreement with their model. In all, they provided significant evidence that the S3 pulse was due to trapped electrons drifting tangentially to the surface due to both gravitational forces and tangential electric fields present in a

tilted detector. While S3 pulses *per se* have not been seen in other studies of electron backgrounds, the phenomenon described here seems reminiscent of so-called “electron burst” pulses detected by LUX (Section D.10.2).




Figure D.8.: “S3” pulses (green) observed by Akimov *et al.* appeared no more than 200  $\mu\text{s}$  after a muon S2 (pink). In the intervening time, light (red) appeared to trace a path between the S2 and S3 across the TPC. S3 pulses preferentially appeared at a specific location on the edge, consistent with a known tilt of the TPC. These S3 pulses were concluded to be trapped electrons on the liquid surface.

### D.6.2. Trapped SEs After Muons (Power Law)

After muon S3 pulses, the long event windows presented an opportunity to study delayed backgrounds over long timescales. In the study, these delayed backgrounds were referred to as being spontaneous in nature, in contrast with photoionization. First, they noted that the position

distribution of the SE signals after the S3 is was not spatially uniform. Second, they fit a power law (Figure D.9) to the rate distribution from 0.2 ms to 20 ms after the S3 signal and obtained a constant of  $-1.4 \pm 0.2$  averaged over all events. This is quite remarkable as it appears to be the first power law fit to an electron background rate at multi-millisecond long timescales.




Figure D.9.: Akimov *et al.* fit a power law of the form  $\alpha t^\beta + c$  to the SE rate for over 10 ms following muon interactions in their TPC.

Two possible causes of these delayed backgrounds were acknowledged: grid emission from cathode wires, and trapped electrons at the liquid surface. The discussion focuses on the latter hypothesis, highlighting the oddly long timescale of these delayed backgrounds. First, they seemed to defy the 10  $\mu\text{s}$  electron lifetime that was measured, though the study notes that the electron capture cross-section could be different for lower energy electrons moving along the liquid surface. Second, they expected trapped electrons to also be captured relatively quickly by the anode ring, which was partially immersed in the liquid at the edges of the TPC. Because of these two factors, they speculate that the delayed backgrounds could instead be due to impurities or negative ions which could be drifting up slowly that release trapped electrons later.

## D.7. 2017 - Sorensen

By this point in time the problem of delayed backgrounds had gone unresolved, and photoionization had been ruled out as a cause. Focus seemed to have turned to the idea of electrons being trapped at the liquid surface, and impurities in the liquid bulk were also in contention as an explanation. In 2017, Sorensen published an article motivating the hypothesis of trapped electrons at the liquid surface [82].

Here, the term “electron train” is coined, which is described as single electrons emitted over timescales longer than 100  $\mu\text{s}$ . In hindsight, this definition is a bit vague as depending on the detector 100  $\mu\text{s}$  is not necessarily a full drift length. It is later clarified in the article that electron trains are defined as electrons seen after photoionization.

### **i** On Timescales

It should be emphasized here that the original definition of “electron trains” seems like an understatement considering that Akimov *et al.* had observed time-correlated SE pulses for well over 100 ms. This remark is not made in criticism of the article, but rather to emphasize that conflating the timescales of these backgrounds can make it more difficult to compare results from different studies. It is understandable that prior studies had only observed backgrounds in time windows less than a millisecond long as knowledge of this topic was still evolving, but caution is advised in conducting future studies of delayed electron trains if observation windows are less than a few milliseconds in length.

### D.7.1. Trapped Electron Model

The article began with review of a previous model of electrons at a liquid-gas potential barrier, and then added corrections to the mean free path of electrons based on evaluations of a drift velocity model. An additional correction in the form of an assumption that multiple barrier crossing attempts could be made was incorporated as well. After adding these corrections, the electron extraction efficiency was modeled. The potential barrier value and the number of crossing attempts were two free parameters which were swept over to find the best fit to data from the only existing absolute measurement of the electron extraction efficiency. A best fit was found with 20 crossing attempts and a barrier height of 0.34 eV. The modeled barrier height value was approximately twice as small as the previously known value. As a sanity check, the emission time of a typical extracted electron was predicted to be less than 1 ns based on a mean free path of approximately 10 nm, a drift velocity of 3000 m/s at 5 kV/cm, and 10 collisions between crossing attempts.

For a thermalized electron which was trapped at the liquid surface, the predicted lifetime before emission was on the order of 10s of ms. In those 10 ms, it was noted that an electron may diffuse over a centimeter across the surface, based on a diffusion constant of  $80 \text{ cm}^2/\text{s}$ .

Caveats on this model were provided by mention of complications from other factors like an irregular electric potential and ripples in liquid surface. The topic of electrons at the liquid surface being captured on impurities was not covered. Some mitigation strategies were proposed including the use of infrared LEDs to stimulate to stimulate the emission of electrons from liquid surface.

## D.8. 2018 - Sorensen and Kamdin

In 2018, Sorensen and Kamdin published a dedicated study [99] of electron trains that looked for time-correlated SE signals following the decays of a  $^{210}\text{Po}$  source which had been deposited on the cathode. Similar to the 2016 study by Akimov *et al.* the TPC consisted of only a cathode and anode 5 mm apart with a liquid layer of 4 mm deep.


Analysis was performed on the 5.3 MeV alpha (Figure D.10a) and  $^{206}\text{Pb}$  recoils (Figure D.10b) which were easily identified based on their fixed drift times and energy spectra. A 1 ms long summed waveform (Figure D.10c) was used as the basis of the analysis. Two regions of interest (ROIs) were identified where the summed waveform exhibited a distinct exponential decay on two different timescales. The first ROI began 40  $\mu\text{s}$  after the S2 and was labeled the “fast” component. The second ROI began 500  $\mu\text{s}$  after the S2 and was labeled the “slow” component. The fast and slow components had decay constants of roughly 40  $\mu\text{s}$  and 2 ms, respectively.

Two datasets were studied, one at high purity and one at low purity. In each data set the electric field was varied from 3 to 6 kV/cm and the dependence of the exponential decay constants and amplitudes for each ROI were examined. The results are qualitatively summarized in Table D.1. The amplitudes were reported as a percentage of the original number of electrons in the event<sup>3,4</sup>.


---

<sup>3</sup>Because the parameters were extracted from a fit to a summed waveform, it’s somewhat unclear what was meant here by “original number of electrons in the event”. Based on the percentage values reported it is possible that amplitudes were normalized by the number of electrons in the summed event, possibly with some drift time correction applied.


<sup>4</sup>The relative amplitudes are reported as “delayed/prompt electrons [%]” in the paper which is somewhat confusing given other uses of the terms “delayed” and “prompt” in the literature. In the paper, “prompt” electrons are those from progenitor S2 pulses, whereas “delayed” electrons are any which occur after the S2.



(a) A  $^{206}\text{Pb}$  recoil with S1=130 detected photons and S2=67 electrons. A single electron was visible 800  $\mu\text{s}$  after the S2.



(b) A  $^{210}\text{Po}$  alpha recoil with large amounts of light and electrons following the S2.



(c) Low-pass filtered summed waveform of alpha events with a “fast” and “slow” exponential component. These fits were the object of study for this paper.

Figure D.10.: Events from a  $^{210}\text{Po}$  source deposited on the cathode of a dual-phase xenon TPC. Electron trains were studied by looking at a summed waveform of signals from alpha recoils up to 1 ms after the S2 was detected.

Table D.1.: Dependencies of fast and slow ROIs in a summed waveform of alpha events observed by [99]. As an example, the amplitude of the fast component increased at a higher purity and decreased with an increasing applied electric field.

Component of Summed Waveform	Purity	Field
Fast Amplitude	Correlated	Anticorrelated
Slow Amplitude	Anticorrelated	Correlated
Fast Decay Time	Anticorrelated	Correlated
Slow Decay Time	Anticorrelated	Correlated

### D.8.1. Trapped Electrons?

Discussion in the paper seemed to focus on the fast component of the summed waveform. A summary of discussion on each component is given below.

**Fast Amplitude:** They proposed that the fast component originated due to trapped electrons at the liquid surface because the amplitude decreased with electric field. Then they calculated the electron extraction efficiency and estimated that at lower fields, half of the trapped electron population was never emitted, which they ascribed to electrons being captured on impurities at the surface. This interpretation seems like it may align with the larger amplitude measured at a higher purity. This would be because fewer trapped electrons would be captured, so more would be seen at short timescales.

**Fast Decay Time:** The paper notes that the fast decay time could be interpreted as the time required for a thermalized electron to be extracted, and that a lower purity would cause this time to decrease artificially. Based on this interpretation, a higher purity should result in a larger

(i.e. slower) fast decay time as thermalized electrons would have more opportunities to be extracted. However the opposite was observed: the fast decay time was shortened slightly at higher purity, and this does not seem to have been remarked on in the paper.

**Slow Amplitude:** The paper notes that the slow component of the summed waveform is more difficult to interpret. Little explanation is offered on the slow amplitude increasing with electric field, only that it might be due higher S2 gain or the “increase in the electric field”. It is unclear what was meant by this. The slow amplitude does drop significantly with increased purity though, and this is acknowledged in the paper.

**Slow Decay Time:** It was found that the slow decay time increased with the electric field and decreased with purity. Hypotheses for these behaviors are not offered; instead the mechanism of fluorescence from PTFE or impurities in the liquid bulk is proposed. If this mechanism is true, they argued, then the slow decay time would depend on factors like the inherent purity and preparation of the PTFE.

The fluorescence hypothesis is also motivated by the presence of delayed single photons among the observed single electrons. However because the TPC lacks any position resolution, they caution that single photons could be from SEs at the edge of the detector.

A second hypothesis for electron trains in general is also proposed in that the electrons could be released by electronegative impurities after some period of time either by tunneling out or when collisions of negative ions occur with other atoms. Skepticism is cast on the collision mechanism due to a quoted timescale of roughly 10 seconds, though they note that the tunneling probability would be larger at higher electric fields.

Finally, the paper concludes by proposing future improvements including studies at longer

timescales, improved position resolution capable of resolving single electrons, studies with impurity dopants, and a removal of PTFE components to further modulate the xenon purity.

A comparison of results from this study to others is difficult. The summed waveform analysis is unique, and the short waveform duration may have obscured a power law that was visible to Akimov *et al.* in 2016. Furthermore, the electric fields are the highest used in any known study. This, and the “extraction-only” nature of the TPC are worth noting here as a configuration which could be useful in a future study to disentangle the effects of a typical drift region which has a factor of 10-20 lower fields.

## **D.9. 2020 - Akimov *et al.***

In an overview of a new TPC dubbed “RED-100”, Akimov *et al.* [198] debuted a patented “shutter” grid 3 mm below the gate, which is used to block electrons from reaching the liquid surface after a high energy event like a muon. Blocking is activated when a large S1 pulse from a muon is detected. A brief (1 ms) 300 V positive bias is applied to the shutter using a capacitor so it becomes less negative than the anode, making it opaque to electrons which attempt to cross it. As in a previous design, the anode ring is also thicker and partially immersed in the liquid in order to promote the capture of electrons trapped on the liquid surface.

They acknowledge that photoionization plays a role in the formation of SE backgrounds and they go on to mention that other SE backgrounds with no direct causal link to VUV photons have previously been described as “spontaneous”. They continue to use the “spontaneous” moniker (with quotes) to describe delayed SE backgrounds, and they agree with the idea of two possible causes: trapped electrons at the liquid surface, and trapped electrons released by impurities.

Preliminary observations of SE backgrounds are included in the overview. Some dependence of the SE rate on purity was observed but the authors claim that the rate didn't completely correlate with purity. They also note that SE count rates were 3 times higher when the shutter was deactivated. Interestingly they claim to have observed an exponential decay in the SE rate following energetic events like muons. This appears to disagree with the results from their past study where they observed a power law (Figure D.9), but this is not commented on in the paper.

Future efforts on a detailed study of SE backgrounds were planned, including a test at higher extraction fields, and runs at higher purity.

## **D.10. 2020 - LUX**

The most detailed study of instrumental electron backgrounds to date was published by the LUX experiment in 2020 [81].

Numerous potential causes of electron background signals are discussed first, starting with the release of electrons from impurities and trapped electrons at the liquid surface. Additional discussion follows on the possibility of positive ions drifting downwards (with higher mobility than negative ions) which can neutralize on the cathode and produce Auger electrons in the process. In another scenario they pose, positive ions may also accumulate on top of oxide or solid xenon monolayers on the grids, leading to a significant decrease in the work function via the Malter effect.

Recognition is also given to radioactive decays on grid surfaces which may produce S2-only signals in the presence of high fields. Finally, cathode grid emission due to surface defects on grid wires are noted as having caused high voltage instabilities or even prevented TPCs operating at their designed parameters.

The discussion of backgrounds starts with the designation of four categories: photoionization, electron bursts, delayed emission, and spontaneous emission. Delayed emission in this case refers to timescales from milliseconds of up to a second, while spontaneous emission is defined in the text as electrons which appear to be completely independent of prior interactions.

### **D.10.1. Photoionization - Neutral Impurities**

In this section, RFR S1 and low drift-time S2 signals from  $^{83\text{m}}\text{Kr}$  events were studied for their photoionization yields, i.e. the number of electrons observed were normalized by their respective progenitor pulse sizes. Photoionization rates followed an increasing trend with lower purity for both S1 and S2. Yields were higher for S2s and this was attributed to detector geometry from the point of view of an S2 vs an RFR S1. As in the XENON100 study (Section D.5), the rates following an RFR S1 increased, and this was especially noticeable after correcting rates of photoionization SEs for their respective drift times.

Most relevant to the subject of electron backgrounds in this section is a study of whether negative ions (i.e. electronegative impurities) are responsible for photoionization backgrounds in the liquid. This is an interesting challenge to an otherwise uncontested hypothesis. To shed some light on this issue, a simple assumption was made in the LUX study: more photoionization should occur with more negative ions present, and more negative ions should be produced during periods of high activity in the TPC, like when a calibration run occurred.

They showed that photoionization yield was not affected by the event rate in the detector during multiple different calibrations, even though significantly more negative ions should have been produced (Figure D.11a). They also noted that Z positions of photoionization electrons were

consistent with a homogeneous distribution of ionization centers, which was also consistent with results from optical simulations<sup>5</sup>. They remark that this was counter to the expected distribution of negative ions, which should have been concentrated closer to the top of the detector due the drift field. Finally, they compared photoionization yields in the Y direction to ionization yields during neutron calibrations (which were localized in y) and showed that photoionization yields did not increase in the region of the TPC where neutrons were interacting (Figure D.11b). Thus even though an overdensity of ions would have been expected in this region, no additional photoionization electrons were detected from RFR S1s.

They concluded that if neutral impurities were responsible for photoionization, they must have be present in a concentration that followed the concentration of electronegative impurities because of the photoionization yield dependence on electron lifetime. This would also be consistent with the XENON100 experiment which demonstrated a linear relationship between the O2 equivalent concentration in the liquid and photoionization SEs (Figure D.7b).


### D.10.2. Electron Bursts

A phenomenon observed in LUX was periods of electron emission which lasted much longer than an S2 (10  $\mu$ s to 1 ms) that seemed to be correlated with a preceding S2 and its size. These so-called “electron bursts” also exhibited a strong position correlation with the prior S2.


Extensive discussion and analysis of this phenomenon follows. It is convincingly argued that poor electron extraction efficiency led to trapped clusters of electrons which preserved their lateral position relative to progenitor and avoided capture from impurities via local deformation of the

---

<sup>5</sup>While not mentioned in the LUX study, this assumption of homogeneous photoionization is additionally made in the ZEPLIN-III study (Section D.3) where an estimate of the electron lifetime was obtained from SE rates following RFR S1 photoionization.



(a) S1 photoionization yields in LUX remained unchanged despite ongoing calibrations producing quantities of negative ions in the detector above background levels. This indicated that negative ions could not have been responsible for photoionization electron backgrounds as had been previously hypothesized.



(b) Neutron calibrations in LUX were localized in a specific portion of the TPC. This should have resulted in a skewed position distribution of photoionization electrons (shown as black circles) biased towards the location where an increased number of interactions were occurring (blue squares).

Figure D.11.: Evidence from LUX showing that electronegative impurities were unlikely to be responsible for photoionization electrons produced in the TPC.

liquid surface.


Of note here is how this background seems to bear qualitative similarity to the “S3” pulses observed by Akimov *et al.* in 2016 (Section [D.6.1](#)). While fewer experiments have seen these electron burst-like signals, it seems plausible to conclude that electrons *can* be trapped at the liquid surface in a TPC and produce background signals of some kind.

### D.10.3. Delayed Backgrounds


The section on delayed backgrounds is comprehensive. From 0 to 3 ms, SEs following  $^{83\text{m}}\text{Kr}$  S2s were analyzed to take advantage of the high statistics from calibrations. From 3 ms to 1000 ms—a much longer analysis window than in previous studies—WIMP search data was used due to the low background event rate. In both cases, a 100 ms upstream veto was implemented to minimize the effects of high energy events.

To correct for the dependence of the rate on progenitor size, rates were normalized by the raw progenitor S2 area. Beyond the prompt photoionization backgrounds, the rate was well-described by a power law between  $t^{-1.1}$  and  $t^{-1.0}$  (Figure [D.12a](#)). Rates from interactions in the bottom of the TPC were higher at all timescales, even when normalizing by drift-time corrected S2 areas. SE backgrounds also exhibited a strong position-correlation with their progenitors, with rates being significantly higher within approximately 10 cm of a previous S2 pulse (Figure [D.12b](#)). This, along with low photoionization yields, was used to rule out the possibility of photoionization cascades (SEs photoionizing other SEs) dominating the pulse rate.

The dependence of the SE rate was also studied vs purity and it was found that at all timescales up to 1000 ms after a progenitor, the rate decreased with purity (Figure [D.13a](#)). Coupled



(a) SE rates vs time since their respective progenitor S2 were found to follow a power law distribution. Rates were normalized by the raw S2 size (i.e. uncorrected for drift losses). The blue curve is for SEs with progenitors in the bottom 5 cm of the TPC and the red curve is for SEs with progenitors in the top 5 cm of the TPC. The dotted line designated the separation between the  $^{83\text{m}}\text{Kr}$  and WIMP search datasets.



(b) Rates of SEs were found to have a significant correlation with the position of their progenitor pulse. In this plot, only S2s from the central 5 cm of the TPC were selected as progenitors. The even distribution of prompt photoionization backgrounds can also be seen here, lasting less than a millisecond.

Figure D.12.: Time-correlated single electron backgrounds in LUX were observed at least a second after an event had occurred. The time dependence of the rate was found to fit a power law, and many of the background pulses appeared to be co-located with their progenitor.

with the rate dependence on drift time, this provided the basis for an argument that delayed electrons originated from impurities. The position correlation was also apparent over long timescales, qualitatively matching the slow diffusion of impurities on seconds-long timescales. It was further noted that this did not preclude the capture of electrons at the liquid surface, which was suggested as an explanation for the non-zero emission rates following interactions near the top of the liquid bulk.




Figure D.13.: Purity dependence of electron trains measured in LUX.

Among the delayed backgrounds, some SEs did not appear to be correlated in position with their progenitors. A component of these are presumed to be cascade SEs, and delayed SEs from previous interactions. In reference to the photoionization discussion, it was claimed that the low rate of the position-uncorrelated pulses was further evidence of negative ions having a short lifetime or otherwise unobtrusive presence in the TPC.

Finally, the section concludes with some speculation on the power law dependence. It was

claimed that the release of electrons from impurities could follow a power law, possibly due to dynamics similar to other processes which occur in the xenon like electron-ion recombination. The possibility of background optical or infrared photon emission was floated as cause, along with collisional ionization, and PTFE fluorescence.

#### **D.10.4. Photon Trains**

The subject of background photon emission was discussed in the next subsection. An increase in the photoelectron (PHE) rate after interactions was observed in LUX that was independent of purity or the XY position of previous interactions. Notably, these did not exhibit a double-PHE feature, which is a known response of PMTs to VUV light. This made these photons unlikely to have originated from xenon scintillation or electroluminescence. A power law rate which was slower than that of the SEs ( $t^{-0.6}$  to  $t^{-0.5}$ ) ruled them out as being produced by SE backgrounds, but made them candidates for producing SEs via photoionization. PTFE fluorescence was posed as an origin of these photons.

#### **D.10.5. Grids - Multiple Electrons**

The final section of the paper covered the topic of multi-electron backgrounds. In this case, grid emission is believed to be responsible for a significant fraction of multi-electron (ME) backgrounds. This was argued in part by noting that single electrons emitted from wires could undergo multiplication at high fields which were also present near wires. Pile-up of SE pulses was also excluded as a cause by implementing a time-based veto following periods of high activity, up to 5000 ms depending on the size of the preceding interaction. Whereas SE pulses appeared to be uniformly distributed after cuts were implemented, ME pulses appeared to have distinctive populations at

different locations in the detector Figure D.14. These populations were also found mostly during transient emission periods ranging 1-50 seconds in length.




Figure D.14.: Transient periods of electron emission in LUX were visible even after aggressive vetoing on high energy events. Electrons from these emission cycles were localized to specific regions in the detector.

## D.11. 2021 - PIXeY

A study with the PIXeY TPC was released in early 2021 by Bodnia *et al.* that examined SE backgrounds following decays of  $^{83\text{m}}\text{Kr}$  [100]. The paper largely focused on modeling Fowler-Nordheim emission of electrons from defects on grid wires by varying the extraction and drift fields. During the study, 20 Hz of  $^{83\text{m}}\text{Kr}$  decays were continually dispersed throughout the TPC. An electron lifetime of approximately 500  $\mu\text{s}$  was achieved, significantly larger than the maximum drift time of around 50  $\mu\text{s}$ .

SE pulses were studied at three timescales relative to a  $^{83\text{m}}\text{Kr}$  decay: before S1 pulses, before S2 pulses, after S2 pulses, and after the maximum drift time. 130  $\mu\text{s}$  event windows were used.

#### **D.11.1. Pre-S1**

Given a 20 Hz event rate, the pre-S1 rates are at best reflective of the last 20-30  $\mu\text{s}$  of delayed backgrounds which were at most 5 ms since the most recent event. Prior events were not tracked for this part of the analysis, so the rates reported are also not normalized to the size of the prior S2. Given the dispersed  $^{83\text{m}}\text{Kr}$  source however, it may be reasonable to assume that they are dominated by events from those decays.

The delayed SE rates seemed to decrease slightly with drift field and increased dramatically with extraction field, even after correcting for the electron extraction efficiency. The decrease with drift field is explained as being possible because the number of field lines originating on the gate and ending on the anode should decrease with an increased drift field, thus lowering the possibility of collection after emission.

The extraction field dependence of the rate is presumed to be due to Fowler-Nordheim emission. The Fowler-Nordheim equation predicts a specific emission current from an emitter based on numerous parameters like a scale factor to describe the emitter's geometry, the work function of the metal, and others. This emission current was predicted and compared to an estimate derived from rates of SEs being extracted from the liquid surface. A fit to the model was obtained, and Fowler-Nordheim emission was claimed to be a more favorable explanation for the observed backgrounds over the possibility of electrons tunneling away from neutral impurities in the liquid.

Given that the position dependence of delayed electron backgrounds seen in LUX had been

acknowledged in this study, it is puzzling that a positional analysis of backgrounds was not performed here. Such a study would have ideally been able to help distinguish between the effects of grid emitters and other delayed backgrounds, which were known to be independent of each other.


### **D.11.2. S1 Photoionization**

S1 photoionization rates are reported on next. They were found to increase with increasing extraction field, and decrease with a higher drift field. The increase with extraction field is attributed to an increase in the electron extraction efficiency, and most of the discussion in the section is dedicated to providing interpretations for the drift field dependence which is “not well understood”. One potential reason given for this is that a higher drift field will result in a lower drift time which would mean fewer photoionization targets would be created.

### **D.11.3. Post S2 and Post Drift**

Qualitatively, the S2 photoionization rates seem to behave similarly to those observed following S1 pulses. The rate dependence with respect to drift field decreases (Figure D.15a), and the rate clearly increases with the extraction field. However, a major discrepancy which is not explained in the paper is that the rate shown vs time since the S2 appears to follow the opposite trend when swept across different drift fields (Figure D.15b). By eye, it appears that rates at lower drift fields are consistently lower at every drift field tested. This does not appear to be a labeling error as the cathode photoionization spike behaves accordingly at lower drift fields, arriving later than at higher fields. One possibility could be that a fixed integration window was chosen to find the integrated rates based on the location of the cathode rate spike at the lowest drift field. This integration window would artificially drive down the observed rates at higher drift fields as photoionization

rates plummet after photoelectrons from the cathode are extracted.



(a) The rate of SE pulses integrated over a window from the end of the S2s up to the maximum drift time, which is quoted as approximately 35  $\mu\text{s}$ . The rate here is shown to decrease with increasing drift field  $E_d$  at a fixed extraction field voltage  $E_{ext}$ .

(b) When the rate dependence vs time is examined, it appears that rates *increase* with drift field. The color scale does not appear to be inverted here as rates at lower drift fields should have rate spikes from cathode photoionization occur at a longer time after an S2 pulse.

Figure D.15.: Integrated SE rates vs. the drift field seem to follow the opposite trend of the same rates plotted vs. time since the progenitor S2 pulse. The origin of this discrepancy is unclear. It is possible that rates were accidentally integrated over a fixed window which included livetime after the maximum drift time. This could cause the integrated rate at higher drift fields to appear lower.

Similar observations are made regarding the rate dependence of electrons after the rate spike from cathode photoionization. The rate falls with a higher drift field and rises with the extraction field. Grid emission and the release of electrons from impurities are given as possible reasons for this behavior.

Following a study of the cathode quantum efficiency, the total electron background rate is modeled as a function of the Fowler-Nordheim rate, electron extraction efficiency, progenitor size, mean free path of a photon in the liquid, and the densities of neutral and negative impurities. The impurities are granted an electric field-dependent photoionization cross-section which decreases

with the electric field, to account for the observed drift field dependence in the rate.

In all, it is fairly difficult to take the claims from the paper into account. Grid emission was known to be a localized phenomenon at the time of publishing, yet no attempt was made to distinguish from other uncorrelated backgrounds. The plots showing conflicting rate behaviors with drift time also hamper the conclusions made.

## D.12. 2021 - ASTERiX

Kopec *et al.* performed a dedicated study of delayed electron trains over millisecond timescales in 2021 with the ASTERiX TPC [94]. This study was the first published attempt to induce electron backgrounds using 0.8 eV infrared light. They also investigated the dependence on extraction field, drift field, and drift time.

The ASTERiX TPC was quoted as having a drift time of approximately 10  $\mu\text{s}$ , a 3  $\mu\text{s}$  electron lifetime<sup>6</sup>, and 5 ms event windows were used. A background trigger rate of 10 Hz was present after shielding the TPC with lead, and interactions to be studied were generated with a  $^{57}\text{Co}$  source. The final trigger rate for the studies performed was 20 Hz, effectively matching the duration of the event window. Runs were 15 minutes long, which produced about 4,000 event windows per run.

### **i** Event Rates in Small TPCs

It seems worth mentioning here that for a small research TPC, there is a question of optimizing event rates. Longer run times grant higher statistics, but can be disproportionately demanding on staff and students to operate. Therefore, there is a need to balance high enough statistics in a reasonable amount of time, which can constrain the duration of windows where

<sup>6</sup>This is quantified later in the paper as a 90% loss of electrons from interactions at the bottom of the TPC.

electron trains are observed. An additional factor to consider is that small TPCs will have higher background event rates than purpose-built TPCs with extensive background mitigation protocols.

Unlike in LUX where a 100 ms holdoff was required between progenitors, a dynamic holdoff based on the size of the progenitor S2 and the previous progenitor was implemented. The distribution of position-uncorrelated SEs was simulated, and subtracted from the general population of SEs. After this subtraction, a radial cut on single-to-few electrons was applied to select the remaining position-correlated background pulses.

Three populations of child pulses were identified: single electrons, double electrons, and small S2s 3-5 electrons in size. A power law was fit to the S2-normalized rate flux<sup>7</sup> for each population from 30  $\mu$ s to 1 ms (Figure D.16) after progenitor S2s. In a decent agreement with LUX, the SE power law fit went as  $t^{-1.20 \pm 0.04}$ . Rates for the double electron and small S2s were deemed to be too large to be consistent with SE pile-up. This was seen as an indication that the respective background populations were produced via differing mechanisms.

The quantity of interest for subsequent studies was the integrated fraction of each population over the duration of the power law fit, normalized by the size of the progenitor S2. Throughout the course of the paper, two hypotheses for delayed electrons are considered: electrons trapped at the liquid-gas interface, and electrons being released by impurities in the liquid.

---

<sup>7</sup>A flux (counts per unit area) was used here, where the area was determined by the size of the radial cut.




Figure D.16.: S2-normalized rate flux of single-to-few electron pulses detected within a roughly 10 mm radius around the XY location of progenitor S2 pulses.

### D.12.1. Extraction Field

They proposed that if electrons were trapped at the liquid surface, then they would expect the power law to turn more negative (implying faster emission) as the extraction field was increased. This was not observed; instead they found the amplitude increased, with no change in the power law. The increase in amplitude was reflected as a larger fraction of backgrounds observed in the same 30  $\mu$ s to 1 ms timeframe. For all background populations they normalized the delayed background fraction by the drift-time corrected S2. They found that the fraction increased linearly with the extraction field—and not linearly with the extraction efficiency.

This amplitude was compared to the amplitude of the slow component of the electron trains measured by Sorensen and Kamdin (Section D.8), which also increased with the electric field. In retrospect, this connection seems tenuous to make as Sorensen and Kamdin interpreted the *fast* component they saw as being related to electrons trapped at the liquid-gas interface. This

interpretation was predicated on the fact that the measured amplitude of the fast component decreased with increasing electric field, in contrast with the behavior demonstrated in this study.

### **D.12.2. Drift Field**

The effects of a higher drift field were investigated. In the study it was hypothesized that a higher drift field could steepen the power law if electrons in the liquid bulk were susceptible to being trapped and released later by impurities. This was argued because a higher drift field would result in a reduced chance for electrons to be caught by electronegative impurities and an increased chance for electrons to tunnel out once captured.

They observed an increase in the delayed background fraction with the drift field Figure D.17, and no change in the slope of the power law. The change in fraction was obvious for single electrons but not for multielectron pulses, prompting them to point out that the hypothesis of impurity capture should only be relevant to single electrons. As in the extraction field study, the delayed electron fraction was normalized by the drift time corrected S2.

This observation is particularly interesting because higher drift fields should globally lower drift times of progenitors. In other studies (Figure D.12a, Section D.13) it has been shown that events at lower drift times have significantly smaller delayed electron background fractions, even after correcting for drift losses in the progenitor S2. The study here noted that increased charge yield from interactions could play a role in the increased delayed backgrounds, but that does not seem to explain the apparent lack of dependence for multi-electron pulses.




Figure D.17.: Electron background fractions normalized by progenitor S2s corrected for drift losses. According to Kopec *et al.* the single electron fraction rises with drift field, possibly due to an increased charge yield or increased release probability from impurities in the liquid. They also claim the lack of response from multi-electron backgrounds is consistent with impurities in the liquid affecting only single electrons.

### **D.12.3. Drift Time**

They saw a clear linear increase in the delayed electron fraction with the drift time when normalizing by the size of the observed S2. The dependence is completely obscured when correcting the progenitor S2s for drift time losses of electrons. This is partially in conflict with the LUX study, which found that progenitors at larger drift times still produced higher background rates even after correcting for drift time losses. The study seems to make note of this, mentioning that the depth dependence of electron trains “cannot be disentangled”.

### **D.12.4. Infrared Photoionization of Impurities**

The final test involved flooding the TPC with 0.8 eV infrared light, which was explained as being a high enough energy to ionize any electronegative impurities. No increase in the S2 normalized rates are seen for any population, except for a slight increase in the position-uncorrelated SEs. The study’s conclusion from this test was that electronegative ions did not seem to be a component of photoionization backgrounds. While not mentioned in the paper, this does agree with the conclusion from LUX regarding the photoionization of electronegative impurities.

### **D.12.5. Conclusions - Liquid Surface Electrons?**

The study concludes that trapped electrons at the liquid surface are a dominant mechanism in the production of electron backgrounds at long time scales. First, they point out that impurity capture of electrons in the liquid cannot explain the rates of pulses larger than a single electron, as they had already ruled out pile-up of SEs as a cause. The nonlinear increase in background rates with extraction efficiency is also used as a justification that bulk impurities alone could not

be responsible—it is not made clear why this is the case, however. They acknowledge that the increasing single electron fraction with drift field could be a result of liquid impurities releasing electrons. However, they believe that the lack of change in the measured electron lifetime and power law indicate that capture rates are unchanged. They propose that the increased delayed electron fraction is better attributed to a shortening of drift times and an increase in charge yields. This assessment is difficult to follow, and later on in the discussion the drift field dependence is addressed again as having an “unclear” relation to liquid surface effects.

In further service to the liquid surface electron hypothesis, the study highlights the drop in delayed electron fraction at higher drift times, obtained after correcting the progenitor S2 sizes for losses to impurities. An argument regarding this is made earlier in the paper where a decrease in backgrounds with event depth is deemed plausible as it would result in fewer electrons eventually being trapped at the surface. The unchanging power law slope with extraction field is given as a final reason for why delayed backgrounds are likely to originate at the liquid surface<sup>8</sup>.

They move on to speculating about possible mechanisms at the liquid surface, and focus on ripples being caused by the progenitor S2 as it was extracted from the liquid. If electrons and negative ions were trapped at the liquid surface, they may be emitted due to the disturbance. They point out that the ripples dying off would lead to a collective emission rate in the form of an exponentially distributed set of exponentially decaying rates that would take on the form of a power law. Furthermore, as the relaxation of the ripples would not be affected by the extraction field, infrared light, or the initial position of the S2, the power law could remain relatively rigid as they claimed to observed.

---

<sup>8</sup>This is a perplexing statement to make as earlier in the paper they hypothesized that “higher extraction fields would cause the electrons to be emitted faster and make the power law **more negative**” (emphasis added).

The paper concludes by briefly mentioning that no grid hot spots (localized areas with sustained electron emission) were observed during the tests, and further points out that grid emissions should not exhibit any position correlation with progenitor signals in a TPC.

#### **D.12.6. Vibration Tests of ASTERiX**

In Depoian’s 2022 dissertation [199], qualitative tests of the ripple hypothesis raised by Kopec were performed. The tests consisted of attempting to vibrate the outside of ASTERiX with a massage gun at a frequency of  $\sim 30$  Hz, and an ultrasonic transducer at a frequency of 10 kHz. No major difference in the electron background rate was observed, though it was cautioned that this may have been because the oscillators were not strongly coupled to the TPC.

#### **D.13. 2021 - XENON1T**

The most recent major study of electron backgrounds to be published was a combined study by the XENON1T collaboration [80] that characterized single electron backgrounds and used that characterization to set new limits on light dark matter interactions in dual-phase TPCs. This review will cover the background characterization that was performed in first part of the paper.

In the electron background study, the electron lifetime varied over the course of year and was on the order of a few hundred microseconds, with a maximum of 1000  $\mu\text{s}$ . The maximum drift time was nominally 750  $\mu\text{s}$ . Two principal datasets were acquired, one using triggered data with  $\sim 2$  ms event windows that encompassed the first science run (SR1), and one using continuously acquired (triggerless) data. In the triggered dataset, the total livetime is roughly 1%. This is in part because only 1 ms pre-trigger windows are used for counting child pulses, and further exclusions are made

if S1 pulses or S2 pulses  $> \sim 5$  electrons are found in the pre-trigger window.

For the triggerless dataset, progenitor S2s were required to be at least 1 ms after any S1 to prevent contamination from S1 photoionization. Any S2 with a size greater than 5 electrons is considered to be a potential progenitor S2, so only one S2 larger than this threshold is allowed in the pile-up of single electrons which typically follows a high energy interaction.

The study entertains the two most popular causes of delayed electron backgrounds: electrons trapped at the liquid surface, or release of electrons from impurities in the liquid bulk. As a general test of both of these hypotheses, they first compare electron trains from events in the gas above the liquid surface, RFR S1s, and events in the drift region (Figure D.18). The drift region events display a power law-like rate up to 1s after S2s, whereas the rates after both the gas and RFR events exhibit a sharp drop off after the maximum drift time. This is an especially strong demonstration of the electron trains being independent of photoionization. In the study, this is taken as evidence that either hypothesis could be responsible.

Following this, the study focuses on progenitor S2s from events occurring in the drift region. A dynamic holdoff is devised to mitigate contamination between electron trains. If the electron train from a progenitor S2 comprised  $> 50\%$  of electrons in all electron trains within the previous 2 seconds, it was accepted. The number of electrons in a train was not counted directly for this holdoff, but rather modeled using the known power law dependence.

### **D.13.1. Position Dependence**

Building off the known position dependence from progenitors, the XENON1T study defined bounds for position-correlated and position-uncorrelated backgrounds, within 15 cm and outside




Figure D.18.: Electron trains from progenitors in the gas, drift, and reverse field (below cathode) regions of XENON1T. A 200 ms holdoff was imposed after interactions to avoid overlapping electron trains between progenitors. The SE rate during this 200 ms holdoff period is shown with the solid purple line. It serves to verify that the holdoff period was adequate in allowing SE rates to return to a baseline level before SEs from progenitor pulses were tracked. It also shows that rates following progenitors in the gas and reverse field regions return fairly quickly to the pre-holdoff rates, indicating that electron trains are a phenomenon tied to drift region interactions.

20 cm of the progenitor respectively. Rates from these selections were normalized relative to the XY area they encompass.

### D.13.2. Progenitor Size

Counts of single-to-few electron pulses were integrated from 2-200 ms following a progenitor and plotted against the size of the progenitors (Figure D.19). A dependence was observed vs. progenitor size for all position correlated pulses, and likewise for single and double position-uncorrelated pulses. The dependence of the position uncorrelated pulses was believed come from contamination from position correlated pulses. For pulse counts which exhibited a dependence with progenitor size, it was remarked that the relationship was slightly non-linear.

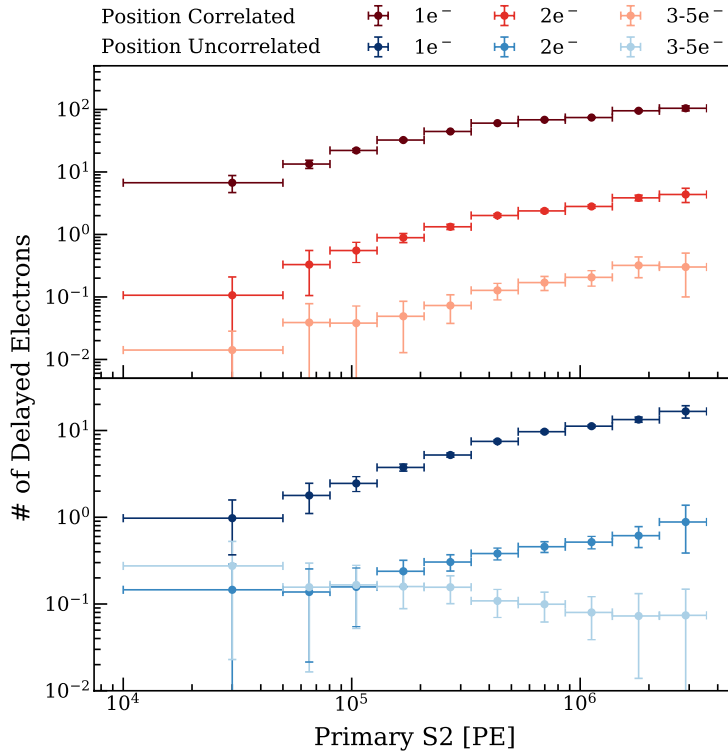


Figure D.19.: Larger progenitor S2s in XENON1T had more delayed single, double and multi-electron background pulses. Position correlated pulses were defined as those within  $\Delta r < 15$  cm of their progenitor and position uncorrelated pulses were those outside of  $\Delta r > 20$  cm. The apparent dependence on progenitor size for position-uncorrelated backgrounds was attributed to leakage from position-correlated backgrounds.

### D.13.3. Delay Time Since Progenitor

Area and S2 normalized rates of background pulses were plotted, and a power law was fit to all three position-correlated populations from 2-200 ms. The powers for single, double, and 3-5 electron pulses were -1.1, -1.3, and -1.6 respectively. The study went on to claim that the differing power laws were indicative of all three populations being a consequence of the same mechanism, and that the resulting multi-electron pulses could not be entirely explained by accidental pile-up of SEs. The same leakage of position-correlated pulses was observed in the time delay plots of position-uncorrelated pulses. No power law was fit to the rates of pulses in these populations.

### D.13.4. Extraction Field




Figure D.20.: Delayed electron fraction from 2-200 ms following a progenitor S2 in XENON1T. The raw S2 (uncorrected for drift losses) size was used for normalization. XENON1T determined that the weak correlation with extraction field would make trapped electrons at the liquid surface a subdominant component of electron trains at best.

The trapped electron hypothesis was evaluated by varying the extraction field and measuring

the relative fraction of delayed electrons (Figure D.20). In this circumstance, the extracted S2 size (uncorrected for drift losses) was used as a normalization factor. Position correlated backgrounds were shown to increase less than 0.02% over an extraction field of about 5-8 kV/cm, and position-uncorrelated backgrounds varied even less. This was taken to mean that trapped electrons at the liquid surface were a subdominant component of electron trains.

### D.13.5. Drift Time

A dependence on progenitor drift times was examined next. Position-correlated backgrounds were found to exhibit a strong dependence on the drift time of the progenitor S2, even after correcting the progenitor size for drift losses. This was in stark contrast with the 2021 ASTERiX study (Section D.12) which had to introduce very large corrections to the progenitor sizes due to poor xenon purity. Position-uncorrelated pulses of all sizes did not exhibit any dependence on drift time. This was understood in the paper as an indication that position-uncorrelated pulses were likely to be remnants of delayed electron backgrounds from previous interactions. It was also noted that the drift time dependence signified the role of impurities in producing delayed electron backgrounds.

### D.13.6. Purity

In most datasets, purity wavered around 600  $\mu\text{s}$ , however, it reached up 1 ms at later dates. The purity dependence of the S2-normalized fraction<sup>9</sup> of background pulses was less clear, but still plainly visible. Rates trended downwards for both position-uncorrelated and position-correlated populations over time as the electron lifetime increased. When plotted directly vs. electron lifetime,

---

<sup>9</sup>Using S2 sizes corrected for drift time losses

a dramatic drop in the position correlated rate was seen as the lifetime exceeds 600  $\mu\text{s}$ . The change in intensity for position-uncorrelated rates vs. electron lifetime was less dramatic, but still discernible.

#### **D.13.7. Conclusion - Liquid Impurities**

The study concluded by restating the lack of extraction field dependence and strong drift time dependence observed in the normalized electron fractions following their progenitors. These factors motivated a conclusion that some impurity in the liquid was responsible for capturing and releasing electrons at long timescales following progenitor pulses. They also interpreted the loosely correlated electron lifetime dependence as an indicator that the impurity was not electronegative, but still subject to removal as the TPC is operated. Finally, they remarked that progenitor pulses should photoionize most electronegative impurities, making it unlikely for them to be responsible for capturing and releasing the observed numbers of electrons hundreds of milliseconds later.

## Appendix E.: Repositories and Resources

Writing a dissertation is challenging. Perhaps a greater challenge though is reading it, and trying to reproduce the work. I’ve always found I was able to learn best by reverse-engineering a working solution, but a major pitfall of dissertations and papers I’ve read is that they don’t include any link back to the code used to generate plots. This motivated me to write my dissertation using Quarto [200]. Quarto is an open-source scientific and technical publishing system, not unlike LaTeX. But unlike LaTeX, Quarto can connect the outputs of my code directly to figures and tables in the text. It’s also possible to publish to multiple formats with Quarto—including to HTML—using the same source. This leads to some interesting benefits, like interactive plots (e.g Figure 5.41), and the ability to embed gifs like the one on the landing page.

### E.1. Repositories

The HTML output is hosted as a static website [here](#), and the repository is [here](#). A PDF copy is downloadable from the website. I find PDFs work best for archival purposes, but if you need to save a plot, for example, you can save it directly from the website. If you need the code I used to generate plots, my work can be found in three places:

### **E.1.1. This repository**

[The repository](#) for this dissertation does contain copies of the notebooks I used for my work. However these are disjointed copies which do not work on a standalone basis, i.e. they need to be run from a dedicated `conda` environment with the correct dependencies. Quarto books, at the time of writing, also do not link back to notebooks when their outputs are embedded in the source. Nevertheless, the notebooks can be found by browsing the repository with some savvy inspection of file paths.

### **E.1.2. LZ Gitlab**

LZ has a private gitlab instance. If you are looking for material pertaining to chapter 3 and have LZ gitlab access, I highly recommend looking in my [user folder](#), and my [ALPACA module repository](#). The latter has its own website like this dissertation, which includes a [tutorial notebook](#). Plots on the website also automatically link back to the notebook used to render them.

### **E.1.3. LLNL Gitlab**

LLNL has an on-premises [gitlab instance](#). Members of the LLNL RED group should know how to gain access. Once you have access, code and documentation for `rgaplot` can be found there, as well as notebooks used for producing plots shown in chapters 4 and 5.

## E.2. Software

### E.2.1. CAD

All renderings of components from XeNu and THERANOS were produced in OnShape; ask Eli for access to the files. See Section [B.2.2](#) for a brief discussion of why I chose to use OnShape over SolidWorks.

### E.2.2. Zotero

Zotero is a fantastic open source reference management software which integrates well with Quarto, VSCode, and Overleaf.

### E.2.3. VSCode

VSCode is an editor with a powerful suite of extensions that transform it into a jack-of-all trades IDE.

I currently use the following extensions:

```
albertopdrf.root-file-viewer  
bierner.markdown-mermaid  
bpruitt-goddard.mermaid-markdown-syntax-highlighting  
brunnerh.file-properties-viewer  
charliermarsh.ruff  
davidanson.vscode-markdownlint  
eamodio.gitlens
```

fcrespo82.markdown-table-formatter  
hediet.vscode-drawio  
james-yu.latex-workshop  
jeff-hykin.better-cpp-syntax  
jock.svg  
mads-hartmann.bash-ide-vscode  
mechatroner.rainbow-csv  
mhutchie.git-graph  
ms-azuretools.vscode-docker  
ms-dotnettools.vscode-dotnet-runtime  
ms-python.debugpy  
ms-python.python  
ms-python.vscode-pylance  
ms-toolsai.datawrangler  
ms-toolsai.jupyter  
ms-toolsai.jupyter-keymap  
ms-toolsai.jupyter-renderers  
ms-toolsai.vscode-jupyter-cell-tags  
ms-toolsai.vscode-jupyter-slideshow  
ms-vscode-remote.remote-containers  
ms-vscode-remote.remote-ssh  
ms-vscode-remote.remote-ssh-edit

```
ms-vscode-remote.remote-wsl
ms-vscode-remote.vscode-remote-extensionpack
ms-vscode.cmake-tools
ms-vscode.cpptools
ms-vscode.cpptools-extension-pack
ms-vscode.cpptools-themes
ms-vscode.live-server
ms-vscode.remote-explorer
ms-vscode.remote-server
ms-vsliveshare.vsliveshare
njpwerner.autodocstring
pomdtr.excalidraw-editor
quarto.quarto
redhat.vscode-yaml
s3gf4ult.monokai-vibrant
shd101wyy.markdown-preview-enhanced
streetsidesoftware.code-spell-checker
sumneko.lua
tamasfe.even-better-toml
timonwong.shellcheck
twxs.cmake
usernamehw.errorlens
```

```
vscode-icons-team.vscode-icons  
yzhang.markdown-all-in-one  
zoxon.monokai-light
```

#### E.2.4. Inkscape

Inkscape is an open source vector graphics editor which is excellent for editing and extracting SVG and PDF formatted plots from papers.

#### E.2.5. Notion

Notion is a closed source program for knowledge management, with a free premium tier for students. I found their databases to be extremely helpful for staying on top of projects.

#### E.2.6. Draw.io

Draw.io is an open source diagramming program which is much more powerful than PowerPoint and Google Slides. It is comparable to Microsoft Visio.

#### E.2.7. Python

Numerous Python packages were used to conduct the analyses presented in this dissertation. They are listed below. Citations are available for specific packages, denoted with annotations.

```
awkward = "2.1.0" ①
```

```
uproot = ">=5.0.11" ②
```

```
pyarrow = ">=13.0.0"
typer = ">=0.9.0"
numba = ">=0.57.1"
pandas = ">=2.1.0"
dask = ">=2023.9.0"
distributed = ">=2023.9.0"
hist = { version = ">=2.6.3", extras = ["plot"] }
fsspec = ">=2023.9.0"
ipywidgets = ">=8.1.0"
ipykernel = ">=6.25.1"
rich = ">=13.5.2"
lmfit = ">=1.2.2"
regex = ">=2023.8.8"
bokeh = ">=3.2.2"
git-cliff = ">=1.3.1"
git-filter-repo = ">=2.38.0"
ruff = ">=0.1.3"
nbformat = ">=5.9.2"
nbclient = ">=0.8.0"
tabulate = ">=0.9.0"
```

① [201]

② [202]

③ [203]

④ [204]

⑤ [205]

⑥ [206]

⑦ [207]

⑧ [208]

## References

- [1] G. Bertone and D. Hooper, *History of Dark Matter*, Rev. Mod. Phys. **90**, 045002 (2018).
- [2] F. Zwicky, *Die Rotverschiebung von extragalaktischen Nebeln*, Helvetica Physica Acta **6**, 110 (1933).
- [3] V. C. Rubin and Jr. Ford W. Kent, *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*, Ap. J. **159**, 379 (1970).
- [4] K. C. Freeman, *On the Disks of Spiral and S0 Galaxies*, Ap. J. **160**, 811 (1970).
- [5] G. Bertone, D. Hooper, and J. Silk, *Particle Dark Matter: Evidence, Candidates and Constraints*, Physics Reports **405**, 279 (2005).
- [6] D. H. Rogstad and G. S. Shostak, *Gross Properties of Five Scd Galaxies as Determined from 21-CENTIMETER Observations*, Ap. J. **176**, 315 (1972).
- [7] R. N. Whitehurst and M. S. Roberts, *High-Velocity Neutral Hydrogen in the Central Region of the Andromeda Galaxy*, Ap. J. **175**, 347 (1972).
- [8] M. S. Roberts and A. H. Rots, *Comparison of Rotation Curves of Different Galaxy Types*, A. Ap. **26**, 483 (1973).
- [9] A. Bosma, *The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types*, PhD thesis, -, 1978.

- [10] V. C. Rubin, Jr. Ford W. K., and N. Thonnard, *Extended rotation curves of high-luminosity spiral galaxies. IV. Systematic dynamical properties*, *Sa -&gt; Sc.*, Ap. J. **225**, L107 (1978).
- [11] H. K. Eriksen, *"An Introduction to the CMB Power Spectrum"*, (2011).
- [12] M. Tegmark, *The angular power spectrum of the 4 year COBE data*, Astrophys. J. **464**, L35 (1996).
- [13] T. Souradeep, R. Saha, and P. Jain, *Angular power spectrum of CMB anisotropy from WMAP*, New Astron. Rev. **50**, 854 (2006).
- [14] Planck Collaboration et al., *Planck 2018 Results. VI. Cosmological Parameters*, Astronomy and Astrophysics **641**, A6 (2020).
- [15] R. L. Workman and others, *Review of Particle Physics - Big Bang Nucleosynthesis*, PTEP (2022).
- [16] H. Reeves, J. Audouze, W. A. Fowler, and D. N. Schramm, *On the Origin of Light Elements*, Ap. J. **179**, 909 (1973).
- [17] B. D. Fields, *The Primordial Lithium Problem*, Annual Review of Nuclear and Particle Science **61**, 47 (2011).
- [18] D. Clowe, M. Bradac, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, *A direct empirical proof of the existence of dark matter*, Astrophys. J. **648**, L109 (2006).
- [19] S. M. Molnar, *Cluster Physics with Merging Galaxy Clusters*, Frontiers in Astronomy and Space Sciences **2**, (2016).
- [20] M. Davis, J. Huchra, D. W. Latham, and J. Tonry, *A survey of galaxy redshifts. II. The large scale space distribution.*, Ap. J. **253**, 423 (1982).

- [21] S. D. M. White, C. S. Frenk, and M. Davis, *Clustering in a neutrino-dominated universe*, Ap. J. **274**, L1 (1983).
- [22] J. Diemand and B. Moore, *The Structure and Evolution of Cold Dark Matter Halos*, Advanced Science Letters **4**, 297 (2011).
- [23] M. Tanabashi et al., *Review of Particle Physics*, Phys. Rev. D **98**, 030001 (2018).
- [24] S. S. Gershtein and Ya. B. Zeldovich, *Rest Mass of Muonic Neutrino and Cosmology*, JETP Lett. **4**, 120 (1966).
- [25] P. Hut, *Limits on Masses and Number of Neutral Weakly Interacting Particles*, Physics Letters B **69**, 85 (1977).
- [26] J. E. Gunn, B. W. Lee, I. Lerche, D. N. Schramm, and G. Steigman, *Some astrophysical consequences of the existence of a heavy stable neutral lepton.*, Ap. J. **223**, 1015 (1978).
- [27] M. Aker et al., *An improved upper limit on the neutrino mass from a direct kinematic method by KATRIN*, (2019).
- [28] S. Dar, *The Neutron EDM in the SM: A Review*, (2000).
- [29] J. M. Pendlebury et al., *Revised Experimental Upper Limit on the Electric Dipole Moment of the Neutron*, Phys. Rev. D **92**, 092003 (2015).
- [30] R. D. Peccei and H. R. Quinn, *CP Conservation in the Presence of Pseudoparticles*, Phys. Rev. Lett. **38**, 1440 (1977).
- [31] F. Wilczek, *Problem of Strong P and T Invariance in the Presence of Instantons*, Phys. Rev. Lett. **40**, 279 (1978).
- [32] S. Weinberg, *A New Light Boson?*, Phys. Rev. Lett. **40**, 223 (1978).

- [33] G. Mueller, P. Sikivie, D. B. Tanner, and K. van Bibber, *Detailed design of a resonantly-enhanced axion-photon regeneration experiment*, Phys. Rev. **D80**, 072004 (2009).
- [34] R. Mayle, J. R. Wilson, J. Ellis, K. Olive, D. N. Schramm, and G. Steigman, *Constraints on Axions from SN 1987A*, Physics Letters B **203**, 188 (1988).
- [35] K. Saikawa, T. Hiramatsu, M. Kawasaki, T. Noumi, R. Sato, T. Sekiguchi, and M. Yamaguchi, *Production and Evolution of Dark Matter Axions in the Early Universe*, in *Proceedings, 10th Patras Workshop on Axions, WIMPs and WISPs (AXION-WIMP 2014): Geneva, Switzerland, June 29-July 4, 2014* (2014), pp. 131–134.
- [36] C. O'HARE, *Cajohare/AxionLimits: AxionLimits* (Zenodo, 2020).
- [37] J. Ellis, J. S. Hagelin, D. V. Nanopoulos, K. Olive, and M. Srednicki, *Supersymmetric Relics from the Big Bang*, Nuclear Physics B **238**, 453 (1984).
- [38] G. Steigman and M. S. Turner, *Cosmological Constraints on the Properties of Weakly Interacting Massive Particles*, Nuclear Physics B **253**, 375 (1985).
- [39] G. Jungman, M. Kamionkowski, and K. Griest, *Supersymmetric Dark Matter*, Physics Reports **267**, 195 (1996).
- [40] N. Bernal, M. Heikinheimo, T. Tenkanen, K. Tuominen, and V. Vaskonen, *The Dawn of FIMP Dark Matter: A Review of Models and Constraints*, International Journal of Modern Physics A **32**, 1730023 (2017).
- [41] B. W. Lee and S. Weinberg, *Cosmological Lower Bound on Heavy-Neutrino Masses*, Phys. Rev. Lett. **39**, 165 (1977).
- [42] S. Galli, F. Iocco, G. Bertone, and A. Melchiorri, *Updated CMB Constraints on Dark Matter Annihilation Cross Sections*, Phys. Rev. D **84**, 027302 (2011).

- [43] M. Battaglieri et al., *US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report*, in *U.S. Cosmic Visions: New Ideas in Dark Matter College Park, MD, USA, March 23-25, 2017* (2017).
- [44] R. Essig et al., *Snowmass2021 Cosmic Frontier: The Landscape of Low-Threshold Dark Matter Direct Detection in the Next Decade*, in (arXiv, 2023).
- [45] S. Dodelson and L. M. Widrow, *Sterile Neutrinos as Dark Matter*, Phys. Rev. Lett. **72**, 17 (1994).
- [46] J. Balajthy, Purity Monitoring Techniques and Electronic Energy Deposition Properties in Liquid Xenon Time Projection Chambers, PhD thesis, University of Maryland at College Park, 2018.
- [47] A. Boveia and C. Doglioni, *Dark Matter Searches at Colliders*, Ann. Rev. Nucl. Part. Sci. **68**, 429 (2018).
- [48] ATLAS Collaboration, *Dark Matter Summary Plots for s-Channel, 2HDM+a and Dark Higgs Models*, (2022).
- [49] C. Karwin, S. Murgia, T. M. P. Tait, T. A. Porter, and P. Tanedo, *Dark Matter Interpretation of the Fermi-LAT Observation Toward the Galactic Center*, Phys. Rev. **D95**, 103005 (2017).
- [50] A. Albert et al., *Searching for Dark Matter Annihilation in Recently Discovered Milky Way Satellites with Fermi-LAT*, Astrophys. J. **834**, 110 (2017).
- [51] L. E. Strigari, S. M. Koushiappas, J. S. Bullock, M. Kaplinghat, J. D. Simon, M. Geha, and B. Willman, *The Most Dark Matter Dominated Galaxies: Predicted Gamma-ray Signals from the Faintest Milky Way Dwarfs*, Astrophys. J. **678**, 614 (2008).

- [52] M. Ackermann et al., *Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi Large Area Telescope Data*, Phys. Rev. Lett. **115**, 231301 (2015).
- [53] T. Marrodán Undagoitia and L. Rauch, *Dark matter direct-detection experiments*, J. Phys. **G43**, 013001 (2016).
- [54] L. Tvrznikova, *Sub-GeV Dark Matter Searches and Electric Field Studies for the LUX and LZ Experiments*, PhD thesis, Yale U., 2019.
- [55] P. Gondolo, J. Edsjö, P. Ullio, L. Bergström, M. Schelke, and E. A. Baltz, *DarkSUSY: Computing Supersymmetric Dark Matter Properties Numerically*, Journal of Cosmology and Astroparticle Physics **2004**, 008 (2004).
- [56] R. W. Schnee, *Introduction to dark matter experiments*, in *Physics of the large and the small, TASI 09, proceedings of the Theoretical Advanced Study Institute in Elementary Particle Physics, Boulder, Colorado, USA, 1-26 June 2009* (2011), pp. 775–829.
- [57] V. A. Bednyakov and D. V. Naumov, *Coherency and incoherency in neutrino-nucleus elastic and inelastic scattering*, Phys. Rev. **D98**, 053004 (2018).
- [58] V. Chepel and H. Araújo, *Liquid Noble Gas Detectors for Low Energy Particle Physics*, Journal of Instrumentation **8**, R04001 (2013).
- [59] R. Essig, J. Mardon, and T. Volansky, *Direct Detection of Sub-GeV Dark Matter*, Physical Review D **85**, 076007 (2012).
- [60] D. S. Akerib et al., *Snowmass2021 Cosmic Frontier Dark Matter Direct Detection to the Neutrino Fog*, (2022).
- [61] S. Shaw, S. J. Haselschwardt, and A. Fan, *Snowmass*, in (2021), p. 7.

- [62] R. Gaitskell, *Welcome to IDM and Brown University*, in *International Identification of Dark Matter Conference, Providence, USA, 2018* (n.d.).
- [63] T. Doke, *FUNDAMENTAL PROPERTIES OF LIQUID ARGON, KRYPTON AND XENON AS RADIATION DETECTOR MEDIA*, *Portugale Physica* **12**, 9 (1981).
- [64] F. Kossoski, M. Boggio-Pasqua, P.-F. Loos, and D. Jacquemin, *Reference Energies for Double Excitations: Improvement and Extension*, (2024).
- [65] P. Adhikari et al., *Pulse-Shape Discrimination Against Low-Energy Ar-39 Beta Decays in Liquid Argon with 4.5 Tonne-Years of DEAP-3600 Data*, *The European Physical Journal C* **81**, 823 (2021).
- [66] M. Szydagis et al., *A Review of Basic Energy Reconstruction Techniques in Liquid Xenon and Argon Detectors for Dark Matter and Neutrino Physics Using NEST*, *Instruments* **5**, 13 (2021).
- [67] T. Washimi, M. Kimura, M. Tanaka, and K. Yorita, *Scintillation and Ionization Ratio of Liquid Argon for Electronic and Nuclear Recoils at Drift-Fields up to 3 kV/Cm*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **910**, 22 (2018).
- [68] The DarkSide collaboration et al., *Calibration of the Liquid Argon Ionization Response to Low Energy Electronic and Nuclear Recoils with DarkSide-50*, *Physical Review D* **104**, 082005 (2021).
- [69] E. Hogenbirk, J. Aalbers, P. A. Breur, M. P. Decowski, K. van Teutem, and A. P. Colijn, *Precision Measurements of the Scintillation Pulse Shape for Low-Energy Recoils in Liquid Xenon*, *Journal of Instrumentation* **13**, P05016 (2018).

- [70] P. Agnes et al., *Measurement of the Liquid Argon Energy Response to Nuclear and Electronic Recoils*, Physical Review D **97**, 112005 (2018).
- [71] J. Xu, S. Pereverzev, B. Lenardo, J. Kingston, D. Naim, A. Bernstein, K. Kazkaz, and M. Tripathi, *Electron extraction efficiency study for dual-phase xenon dark matter experiments*, Phys. Rev. **D99**, 103024 (2019).
- [72] D. Akimov et al., *Observation of Coherent Elastic Neutrino-Nucleus Scattering*, Science **357**, 1123 (2017).
- [73] B. Lenardo et al., *Measurement of the ionization yield from nuclear recoils in liquid xenon between 0.3 - 6 keV with single-ionization-electron sensitivity*, (2019).
- [74] J. Xu, P. S. Barbeau, and Z. Hong, *Detection and Calibration of Low-Energy Nuclear Recoils for Dark Matter and Neutrino Scattering Experiments*, Annual Review of Nuclear and Particle Science **73**, 27 (2023).
- [75] D. A. Khaitan, *Supernova Neutrino Detection in LZ*, Journal of Instrumentation **13**, C02024 (2018).
- [76] S. Al Kharusi et al., *SNEWS 2.0: A Next-Generation Supernova Early Warning System for Multi-Messenger Astronomy*, New Journal of Physics **23**, 031201 (2021).
- [77] S. Randich and L. Magrini, *Light Elements in the Universe*, Frontiers in Astronomy and Space Sciences **8**, (2021).
- [78] R. Essig, T. Volansky, and T.-T. Yu, *New Constraints and Prospects for Sub-GeV Dark Matter Scattering Off Electrons in Xenon*, Physical Review D **96**, 043017 (2017).
- [79] E. Aprile et al., *Light Dark Matter Search with Ionization Signals in XENON1T*, Physical Review Letters **123**, 251801 (2019).

- [80] XENON Collaboration et al., *Emission of Single and Few Electrons in XENON1T and Limits on Light Dark Matter*, Physical Review D **106**, 022001 (2022).
- [81] D. S. Akerib et al., *Investigation of Background Electron Emission in the LUX Detector*, Physical Review D **102**, 092004 (2020).
- [82] P. Sorensen, *Electron Train Backgrounds in Liquid Xenon Dark Matter Search Detectors Are Indeed Due to Thermalization and Trapping*, (2017).
- [83] R. E. Linehan, T. Shutt, D. S. Akerib, P. Burchat, and A. Friedland, High Voltage Electrode Development and the LZ Experiment's WIMP Search, PhD thesis, 2022.
- [84] E. Aprile et al., *Emission of Single and Few Electrons in XENON1T and Limits on Light Dark Matter*, (2021).
- [85] *Oxygen Gas Phase Ion Energetics Data*, in (National Institute of Standards; Technology, Gaithersburg MD, 20899, n.d.).
- [86] A. Tomás, H. M. Araújo, A. J. Bailey, A. Bayer, E. Chen, B. L. Paredes, and T. J. Sumner, *Study and Mitigation of Spurious Electron Emission from Cathodic Wires in Noble Liquid Time Projection Chambers*, Astroparticle Physics **103**, 49 (2018).
- [87] K. M. Stifter, D. S. Akerib, T. Shutt, H. Tanaka, and N. Toro, The LZ Dark Matter Experiment: From Detector Development to Early Data, PhD thesis, 2021.
- [88] P.-S. Shaw, Z. Li, U. Arp, and K. R. Lykke, *Ultraviolet Characterization of Integrating Spheres*, Applied Optics **46**, 5119 (2007).
- [89] E. Mizrachi, *LUXZEPLIN / Analysis / ALPACA / Modules / Sv\_SEdecaytime* (2024).

- [90] The LZ Collaboration et al., *The LUX-ZEPLIN (LZ) Experiment*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **953**, 163047 (2020).
- [91] B. J. Mount et al., *LUX-ZEPLIN (LZ) Technical Design Report*, (2017).
- [92] J. Aalbers et al., *First Dark Matter Search Results from the LUX-ZEPLIN (LZ) Experiment*, Physical Review Letters **131**, 041002 (2023).
- [93] J. R. Watson, *High Voltage Considerations for Dark Matter Searches*, PhD thesis, 2023.
- [94] A. Kopec, A. L. Baxter, M. Clark, R. F. Lang, S. Li, J. Qin, and R. Singh, *Correlated Single- and Few-Electron Backgrounds Milliseconds After Interactions in Dual-Phase Liquid Xenon Time Projection Chambers*, Journal of Instrumentation **16**, P07014 (2021).
- [95] D. Yu. Akimov et al., *Measurement of Single-Electron Noise in a Liquid-Xenon Emission Detector*, Instruments and Experimental Techniques **55**, 423 (2012).
- [96] D. Yu. Akimov et al., *Observation of Delayed Electron Emission in a Two-Phase Liquid Xenon Detector*, Journal of Instrumentation **11**, C03007 (2016).
- [97] E. Aprile et al., *Observation and Applications of Single-Electron Charge Signals in the XENON100 Experiment*, Journal of Physics G: Nuclear and Particle Physics **41**, 035201 (2014).
- [98] E. Santos et al., *Single Electron Emission in Two-Phase Xenon with Application to the Detection of Coherent Neutrino-Nucleus Scattering*, Journal of High Energy Physics **2011**, 115 (2011).

- [99] P. Sorensen and K. Kamdin, *Two Distinct Components of the Delayed Single Electron Background Signals in Liquid Xenon Emission Detectors*, Journal of Instrumentation **13**, P02032 (2018).
- [100] E. Bodnia et al., *The Electric Field Dependence of Single Electron Emission in the PIXeY Two-Phase Xenon Detector*, Journal of Instrumentation **16**, P12015 (2021).
- [101] XENON Collaboration et al., *The XENON1T Dark Matter Experiment*, The European Physical Journal C **77**, 881 (2017).
- [102] P. Chiggiato, *Outgassing Properties of Vacuum Materials for Particle Accelerators*, (2020).
- [103] R. N. Peacock, *Practical Selection of Elastomer Materials for Vacuum Seals*, Journal of Vacuum Science and Technology **17**, 330 (1980).
- [104] S. G. Sammartano, I. Wevers, G. Bregliozzi, and P. Chiggiato, *Outgassing Rates of PEEK, Kapton® and Vespel® Foils*, (2020).
- [105] B. G. Lenardo, *Measurements and Modeling of Low Energy Nuclear Recoils in Liquid Xenon for Dark Matter and Neutrino Detection*, PhD thesis, 2018.
- [106] T. Pershing et al., *Calibrating the Scintillation and Ionization Responses of Xenon Recoils for High-Energy Dark Matter Searches*, Physical Review D **106**, 052013 (2022).
- [107] J. Xu et al., *Search for the Migdal Effect in Liquid Xenon with keV-Level Nuclear Recoils*, (2023).
- [108] B. Loer, *Bloer/Daqman: Citation Release* (n.d.).
- [109] E. Mizrachi, *Emiz/Daqman2py* (2022).
- [110] J. Fan, M. Reece, and L.-T. Wang, *Non-Relativistic Effective Theory of Dark Matter Direct Detection*, Journal of Cosmology and Astroparticle Physics **2010**, 042 (2010).

- [111] A. L. Fitzpatrick, W. Haxton, E. Katz, N. Lubbers, and Y. Xu, *The Effective Field Theory of Dark Matter Direct Detection*, Journal of Cosmology and Astroparticle Physics **2013**, 004 (2013).
- [112] B. Lenardo, K. Kazkaz, A. Manalaysay, J. Mock, M. Szydagis, and M. Tripathi, *A Global Analysis of Light and Charge Yields in Liquid Xenon*, IEEE Transactions on Nuclear Science **62**, 3387 (2015).
- [113] A. Migdal and others, *Ionization of Atoms Accompanying  $\beta$ - and  $\beta^+$ -Decay*, **4**, 449 (1941).
- [114] K. V. Berghaus, A. Esposito, R. Essig, and M. Sholapurkar, *The Migdal Effect in Semiconductors for Dark Matter with Masses Below 100 MeV*, Journal of High Energy Physics **2023**, 23 (2023).
- [115] LUX Collaboration et al., *Results of a Search for Sub-GeV Dark Matter Using 2013 LUX Data*, Physical Review Letters **122**, 131301 (2019).
- [116] H. M. Araújo et al., *The MIGDAL Experiment: Measuring a Rare Atomic Process to Aid the Search for Dark Matter*, Astroparticle Physics **151**, 102853 (2023).
- [117] K. D. Nakamura, K. Miuchi, S. Kazama, Y. Shoji, M. Ibe, and W. Nakano, *Detection Capability of the Migdal Effect for Argon and Xenon Nuclei with Position-Sensitive Gaseous Detectors*, Progress of Theoretical and Experimental Physics **2021**, 013C01 (2021).
- [118] M. J. Dolan, F. Kahlhoefer, and C. McCabe, *Directly Detecting Sub-GeV Dark Matter with Electrons from Nuclear Scattering*, Physical Review Letters **121**, 101801 (2018).
- [119] J. L. De Segovia, *Physics of Outgassing*, (1999).
- [120] K. Saito, Y. Sato, S. Inayoshi, and S. Tsukahara, *Measurement System for Low Outgassing Materials by Switching Between Two Pumping Paths*, Vacuum **47**, 749 (1996).

- [121] G. Mensitieri, A. Apicella, J. M. Kenny, and L. Nicolais, *Water Sorption Kinetics in Poly(aryl Ether Ether Ketone)*, Journal of Applied Polymer Science **37**, 381 (1989).
- [122] D. ., Jr. Edwards, *Upper Bound to the Pressure in an Elementary Vacuum System*, Journal of Vacuum Science and Technology **14**, 606 (1977).
- [123] P. Chiggiato, *CERN Accelerator School (CAS) on Vacuum for Particle Accelerators*, (2017).
- [124] L. Baudis, Y. Biondi, M. Galloway, F. Girard, S. Hochrein, S. Reichard, P. Sanchez-Lucas, K. Thieme, and J. Wulf, *The First Dual-Phase Xenon TPC Equipped with Silicon Photomultipliers and Characterisation with  $^{37}\text{Ar}$* , The European Physical Journal C **80**, 477 (2020).
- [125] S. Gundacker and A. Heering, *The Silicon Photomultiplier: Fundamentals and Applications of a Modern Solid-State Photon Detector*, Physics in Medicine and Biology **65**, 17TR01 (2020).
- [126] S. Gundacker, R. M. Turtos, N. Kratochwil, R. H. Pots, M. Paganoni, P. Lecoq, and E. Auffray, *Experimental Time Resolution Limits of Modern SiPMs and TOF-PET Detectors Exploring Different Scintillators and Cherenkov Emission*, Physics in Medicine & Biology **65**, 025001 (2020).
- [127] *Introduction to the Silicon Photomultiplier (SiPM)*, (2023).
- [128] B. L. Paredes, H. M. Araújo, F. Froborg, N. Marangou, I. Olcina, T. J. Sumner, R. Taylor, A. Tomás, and A. Vacheret, *Response of Photomultiplier Tubes to Xenon Scintillation Light*, Astroparticle Physics **102**, 56 (2018).
- [129] A. Jamil et al., *VUV-Sensitive Silicon Photomultipliers for Xenon Scintillation Light Detection in nEXO*, IEEE Transactions on Nuclear Science **65**, 2823 (2018).

- [130] A. Falcone et al., *Cryogenic SiPM Arrays for the DUNE Photon Detection System*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **985**, 164648 (2021).
- [131] T. Pershing et al., *Performance of Hamamatsu VUV<sub>4</sub> SiPMs for Detecting Liquid Argon Scintillation*, Journal of Instrumentation **17**, P04017 (2022).
- [132] D. S. Akerib et al., *Position Reconstruction in LUX*, Journal of Instrumentation **13**, P02001 (2018).
- [133] J. P. Brodsky, *Xy Position Reconstruction in DarkSide-50*, PhD thesis, 2015.
- [134] V. N. Solovov et al., *2011 IEEE Nuclear Science Symposium Conference Record*, in (2011), pp. 1226–1233.
- [135] E. P. Bernard et al., *Thermodynamic Stability of Xenon-Doped Liquid Argon Detectors*, Physical Review C **108**, 045503 (2023).
- [136] E. M. Gushchin, A. A. Kruglov, and I. M. Obodovskij, *"Hot" electron emission from liquid and solid argon and xenon*, Zh. Eksp. Teor. Fiz.; (USSR) **82:5**, (1982).
- [137] C. Vogl, M. Schwarz, X. Stribl, J. Griebing, P. Krause, and S. Schönert, *Scintillation and Optical Properties of Xenon-Doped Liquid Argon*, Journal of Instrumentation **17**, C01031 (2022).
- [138] P. Benetti et al., *Measurement of the Specific Activity of <sup>39</sup>Ar in Natural Argon*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **574**, 83 (2007).
- [139] M. A. Omary and H. H. Patterson, *Luminescence Theory*, in edited by J. C. Lindon (Academic Press, Oxford, 1999), pp. 1372–1391.

- [140] S. Kubota, M. Hishida, S. Himi, J. Suzuki, and J. Ruan, *The Suppression of the Slow Component in Xenon-Doped Liquid Argon Scintillation*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **327**, 71 (1993).
- [141] A. Neumeier, T. Dandl, T. Heindl, A. Himpsl, L. Oberauer, W. Potzel, S. Roth, S. Schönert, J. Wieser, and A. Ulrich, *Intense Vacuum-Ultraviolet and Infrared Scintillation of Liquid Ar-Xe Mixtures*, EPL (Europhysics Letters) **109**, 12001 (2015).
- [142] S. Kubota, A. Nakamoto, T. Takahashi, S. Konno, T. Hamada, M. Miyajima, A. Hitachi, E. Shibamura, and T. Doke, *Evidence of the Existence of Exciton States in Liquid Argon and Exciton-Enhanced Ionization from Xenon Doping*, Physical Review B **13**, 1649 (1976).
- [143] M. Suzuki, M. Hishida, J. Ruan(Gen), and S. Kubota, *Light Output and Collected Charge in Xenon-Doped Liquid Argon*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **327**, 67 (1993).
- [144] T. Takahashi, S. Himi, M. Suzuki, J. Ruan(Gen), and S. Kubota, *Emission Spectra from Ar-Xe, Ar-Kr, Ar-N<sub>2</sub>, Ar-CH<sub>4</sub>, Ar-CO<sub>2</sub> and Xe-N<sub>2</sub> Gas Scintillation Proportional Counters*, Nuclear Instruments and Methods in Physics Research **205**, 591 (1983).
- [145] T. Efthimiopoulos, D. Zouridis, and A. Ulrich, *Excimer Emission Spectra of Rare Gas Mixtures Using Either a Supersonic Expansion or a Heavy-Ion-Beam Excitation*, Journal of Physics D: Applied Physics **30**, 1746 (1997).
- [146] D. K. Anderson, *Lifetimes of the  $(5p)^5(6s)^1p_1$  and  $(5p)^3p_1$  States of Xenon*, Physical Review **137**, A21 (1965).

- [147] G. Nowak and J. Fricke, *The Heteronuclear Excimers ArKr\*, ArXe\* and KrXe\**, Journal of Physics B: Atomic and Molecular Physics **18**, 1355 (1985).
- [148] J. Soto-Oton and on behalf of the D. Collaboration, *Impact of Xenon Doping in the Scintillation Light in a Large Liquid-Argon TPC*, Journal of Physics: Conference Series **2374**, 012164 (2022).
- [149] C. G. Wahl, E. P. Bernard, W. H. Lippincott, J. A. Nikkel, Y. Shin, and D. N. McKinsey, *Pulse-Shape Discrimination and Energy Resolution of a Liquid-Argon Scintillator with Xenon Doping*, Journal of Instrumentation **9**, P06013 (2014).
- [150] L. B. Robinson, *Phase Equilibria in Cryogenic Mixtures: Part II*, in edited by R. G. Ross (Springer US, Boston, MA, 1995), pp. 559–567.
- [151] S. I. Sandler and S. I. Sandler, *Chemical, Biochemical, and Engineering Thermodynamics*, 4th ed (John Wiley, Hoboken, N.J, 2006).
- [152] Y. Marcus, *Solubility Parameters of Permanent Gases*, Journal of Chemistry **2016**, e4701919 (2016).
- [153] J. B. Ott and J. Boerio-Goates, *Chemical Thermodynamics: Advanced Applications* (Academic Press, London, UK ; San Diego, Calif, 2000).
- [154] M. C. Antoine, *Nouvelle Relation Entre Les Tensions Et Les Temperatures*, C. R. Seanc. Acad. Sci. Paris **107**, 681 (1888).
- [155] F. Simon and G. Glatzel, *Bemerkungen Zur Schmelzdruckkurve*, Zeitschrift Für Anorganische Und Allgemeine Chemie **178**, 309 (1929).
- [156] A. G. M. Ferreira and L. Q. Lobo, *The Sublimation of Argon, Krypton, and Xenon*, The Journal of Chemical Thermodynamics **40**, 1621 (2008).

- [157] R. Heastie and C. Lefebvre, *Phase Equilibria in Condensed Mixtures of Argon and Xenon*, Proceedings of the Physical Society **76**, 180 (1960).
- [158] M. Campestrini, P. Stringari, and P. Arpentinier, *Solid-liquid Equilibrium Prediction for Binary Mixtures of Ar, O<sub>2</sub>, N<sub>2</sub>, Kr, Xe, and CH<sub>4</sub> Using the LJ-SLV-EoS*, Fluid Phase Equilibria **379**, 139 (2014).
- [159] W. H. Yunker and G. D. Halsey, *THE SOLUBILITY, ACTIVITY COEFFICIENT AND HEAT OF SOLUTION OF SOLID XENON IN LIQUID ARGON*, The Journal of Physical Chemistry **64**, 484 (1960).
- [160] P. Stringari, *Phase Equilibrium Behavior of the Ar-Xe Mixture at Low T*, (2024).
- [161] R. Gilgen, R. Kleinrahm, and W. Wagner, *Measurement and Correlation of the (Pressure, Density, Temperature) Relation of Argon II. Saturated-Liquid and Saturated-Vapour Densities and Vapour Pressures Along the Entire Coexistence Curve*, The Journal of Chemical Thermodynamics **26**, 399 (1994).
- [162] Ch. Tegeler, R. Span, and W. Wagner, *A New Equation of State for Argon Covering the Fluid Region for Temperatures from the Melting Line to 700 K at Pressures up to 1000 MPa*, Journal of Physical and Chemical Reference Data **28**, 779 (1999).
- [163] C. W. Leming and G. L. Pollack, *Sublimation Pressures of Solid Ar, Kr, and Xe*, Physical Review B **2**, 3323 (1970).
- [164] L. I. Drii and V. A. Rabinovich, *Dependence of Vapor Pressure on Argon*, Zhurnal Fizicheskoi Khimii **40**, 709 (1966).

- [165] D. R. Lide, *CRC Handbook of chemistry and physics: 2005-2006 a ready-reference book of chemical and physical data*, 86th ed (CRC Press Taylor; Francis Group, Boca Raton New York Washington, 2005).
- [166] A. Michels and T. Wassenaar, *Vapour Pressure of Liquid Xenon*, *Physica* **16**, 253 (1950).
- [167] H. DeVoe, *Thermodynamics and Chemistry* (Prentice Hall, Upper Saddle River, NJ, 2001).
- [168] E. A. Guggenheim, *THE THEORY OF THE EQUILIBRIUM PROPERTIES OF SOME SIMPLE CLASSES OF MIXTURES SOLUTIONS AND ALLOYS* (The Clarendon Press, Oxford, 1952).
- [169] D. Jafari, A. Franco, S. Filippeschi, and P. Di Marco, *Two-Phase Closed Thermosyphons: A Review of Studies and Solar Applications*, *Renewable and Sustainable Energy Reviews* **53**, 575 (2016).
- [170] A. W. Bradley, *LUX Thermosyphon Cryogenics and Radon-Related Backgrounds for the First WIMP Result*, PhD thesis, 2014.
- [171] H. Hu, C. Xu, Y. Zhao, K. J. Ziegler, and J. N. Chung, *Boiling and Quenching Heat Transfer Advancement by Nanoscale Surface Modification*, *Scientific Reports* **7**, 6117 (2017).
- [172] P. G. Kosky and D. N. Lyon, *Pool Boiling Heat Transfer to Cryogenic Liquids; I. Nucleate Regime Data and a Test of Some Nucleate Boiling Correlations*, *AIChE Journal* **14**, 372 (1968).
- [173] *VUV-MPPC 4th Generation (VUV4) Product Flyer*, (2017).

- [174] R. Álvarez-Garrote et al., *Measurement of the Photon Detection Efficiency of Hamamatsu VUV<sub>4</sub> SiPMS at Cryogenic Temperature*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1064**, 169347 (2024).
- [175] D. M. Hoffman, B. Singh, and J. H. Thomas, editors, *Handbook of Vacuum Science and Technology* (Academic Press, San Diego, CA, 1998).
- [176] J. E. Blessing, R. E. Ellefson, B. A. Raby, G. A. Brucker, and R. K. Waits, *Recommended Practice for Process Sampling for Partial Pressure Analysis*, Journal of Vacuum Science and Technology A: Vacuum, Surfaces, and Films **25**, 167 (2007).
- [177] D. H. Adam et al., *Fundamentals of Vacuum Technology* (Leybold Vacuum, 2016).
- [178] H. T. Bach, B. A. Meyer, and D. G. Tuggle, *Role of Molecular Diffusion in the Theory of Gas Flow Through Crimped-Capillary Leaks*, Journal of Vacuum Science and Technology A: Vacuum, Surfaces, and Films **21**, 806 (2003).
- [179] H. C. Mattraw, R. E. Patterson, and C. F. Pachucki, *Mass Spectrometer Analyses Using a Viscous Leak*, Applied Spectroscopy **8**, 117 (1954).
- [180] R. Wright, *Mass Spectrometry Gas Analysis*, B9780124095472121985 (2016).
- [181] *Stanford Research Systems Universal Gas Analyzer*, (2018).
- [182] J. A. Basford, M. D. Boeckmann, R. E. Ellefson, A. R. Filippelli, D. H. Holkeboer, L. Lieszkovsky, and C. M. Stupak, *Recommended Practice for the Calibration of Mass Spectrometers*, Journal of Vacuum Science and Technology A: Vacuum, Surfaces, and Films **11**, A22 (1993).

- [183] R. E. Ellefson, D. Cain, and C. N. Lindsay, *Calibration of Mass Spectrometers for Quantitative Gas Mixture Analysis*, Journal of Vacuum Science and Technology A **5**, 134 (1987).
- [184] J. Li, M. Guan, Y. Gan, and P. Zhang, *A New Method to Measure the Concentration of Argon-xenon Gas Mixture*, Radiation Detection Technology and Methods **5**, 280 (2021).
- [185] *HPR-70 - Hiden Analytical*, (2020).
- [186] *Stanford Research Systems Residual Gas Analyzer*, (2009).
- [187] Y. Y. Gan, M. Y. Guan, Y. P. Zhang, P. Zhang, C. G. Yang, Q. Zhao, Y. T. Wei, and W. X. Xiong, *Using  $^{22}\text{Na}$  and  $^{83}\text{mKr}$  to Calibrate and Study the Properties of Scintillation in Xenon-Doped Liquid Argon*, Journal of Instrumentation **15**, P12007 (2020).
- [188] I. A. Platonov, O. V. Rodinkov, A. R. Gorbacheva, L. N. Moskvina, and I. N. Kolesnichenko, *Methods and Devices for the Preparation of Standard Gas Mixtures*, Journal of Analytical Chemistry **73**, 109 (2018).
- [189] E. Mizrachi, *Emiz/Rgaplot* (2023).
- [190] G. B. Bunyard, *Accurate Mass Spectrometry at Pressures of  $10^{-5}$  to  $10^{-9}$  Pa*, Vacuum **44**, 633 (1993).
- [191] J. C. Little and L. B. Gordon, *Automated, Differentially Pumped, Mass-Spectrometer Sampling System*, Review of Scientific Instruments **62**, 334 (1991).
- [192] E. Alfonso-Pita et al., *Snowmass 2021 Scintillating Bubble Chambers: Liquid-Noble Bubble Chambers for Dark Matter and CE\$ \$NS Detection*, in (arXiv, 2022).
- [193] N. Gallice and on behalf of the D. Collaboration, *Xenon Doping of Liquid Argon in Proto-DUNE Single Phase*, Journal of Instrumentation **17**, C01034 (2022).

- [194] N. McFadden, S. R. Elliott, M. Gold, D. E. Fields, K. Rielage, R. Massarczyk, and R. Gibbons, *Large-Scale, Precision Xenon Doping of Liquid Argon*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1011**, 165575 (2021).
- [195] B. Edwards et al., *Measurement of Single Electron Emission in Two-Phase Xenon*, Astroparticle Physics **30**, 54 (2008).
- [196] A. A. Burenkov, D. Yu. Akimov, Yu. L. Grishkin, A. G. Kovalenko, V. N. Lebedenko, V. N. Solovov, V. N. Stekhanov, F. Neves, and T. J. Sumner, *Detection of a Single Electron in Xenon-Based Electroluminescent Detectors*, Physics of Atomic Nuclei **72**, 653 (2009).
- [197] XENON10 Collaboration et al., *Search for Light Dark Matter in XENON10 Data*, Physical Review Letters **107**, 051301 (2011).
- [198] D. Y. Akimov et al., *First Ground-Level Laboratory Test of the Two-Phase Xenon Emission Detector RED-100*, Journal of Instrumentation **15**, P02020 (2020).
- [199] A. L. Depoian, *Few-Electron Signals and Their Implications in Liquid Xenon Time Projection Chambers*, PhD thesis, 2022.
- [200] J. J. Allaire, C. Teague, C. Scheidegger, Y. Xie, and C. Dervieux, *Quarto* (2024).
- [201] J. Pivarski, I. Osborne, I. Ifrim, H. Schreiner, A. Hollands, A. Biswas, P. Das, S. Roy Choudhury, N. Smith, and M. Goyal, *Awkward Array* (Zenodo, 2024).
- [202] J. Pivarski, H. Schreiner, A. Hollands, P. Das, K. Kothari, A. Roy, J. Ling, N. Smith, C. Burr, and G. Stark, *Uproot* (Zenodo, 2024).
- [203] S. Ramírez, *Typewriter* (2024).
- [204] S. K. Lam et al., *Numba/Numba: Numba 0.59.1* (Zenodo, 2024).

- [205] T. pandas development team, *Pandas-Dev/Pandas: Pandas* (Zenodo, 2024).
- [206] H. Schreiner, S. Liu, and A. Goel, *Hist* (Zenodo, 2023).
- [207] F. Perez and B. E. Granger, *IPython: A System for Interactive Scientific Computing*, *Computing in Science & Engineering* **9**, 21 (2007).
- [208] M. Newville et al., *Lmfit/Lmfit-Py: 1.3.1* (Zenodo, 2024).