

Shedding Light on The Neutrino's Nature: R&D for the nEXO Light Detection System

David Gallacher
Doctor of Philosophy



Department of Physics
McGill University
Montreal, Quebec, Canada

August 14, 2026

A thesis submitted to McGill University in partial
fulfillment of the requirements of the degree of
Doctor of Philosophy

©David Gallacher, 2026

Abstract

Searches for physics beyond the Standard Model increasingly rely on ultra-sensitive detectors operating at the intensity frontier, where the discovery potential is driven by the ability to observe extremely rare events in the presence of background signals. One of the most compelling targets of this approach is neutrinoless double beta decay, whose observation would establish the Majorana nature of neutrinos and demonstrate the violation of lepton number. The proposed nEXO experiment aims to reach world-leading sensitivity to this process, targeting half-lives of order 10^{28} years using a tonne-scale liquid xenon detector enriched in the isotope ^{136}Xe . Achieving such sensitivity requires exceptional control of detector backgrounds and a detailed understanding of detector response, particularly for scintillation light detection.

This thesis focuses on the characterization and modelling of light transport and detection in ultra-low background liquid xenon detectors employing silicon photomultipliers (SiPMs), using the R&D experiment known as, light only liquid xenon (LoLX) at McGill University. This work presents the first in-liquid measurements of correlated secondary light emission from SiPMs relevant to the nEXO design. To enable future measurements, the LoLX experiment was comprehensively upgraded, including its cryogenic systems, detector geometry, data acquisition, and analysis and simulation frameworks. These upgrades support precision measurements essential for validating optical models and informing the design of discovery-oriented experiments.

By combining improved experimental characterization with upgraded optical simulations, this work delivers a detailed description of the light detection system for nEXO. The results presented here provide critical input for performance projections and design optimization, and contribute to the broader effort of enabling next-generation rare-event searches with liquid xenon detectors.

Abrégé

Les recherches de physique au-delà du Modèle standard reposent de plus en plus sur des détecteurs ultrasensibles opérant à la frontière de l'intensité, où le potentiel de découverte dépend de la capacité à observer des événements extrêmement rares en présence de bruits de fond considérables. L'une des plus prometteuses de cette approche est la double désintégration bêta sans émission de neutrinos, dont l'observation établirait la nature de Majorana des neutrinos et démontrerait la violation du nombre leptonique. L'expérience proposée nEXO vise à atteindre une sensibilité de premier plan mondial à ce processus, ciblant des demi-vies de l'ordre de 10^{28} années à l'aide d'un détecteur au xénon liquide à l'échelle de la tonne, enrichi en isotope ^{136}Xe . L'atteinte d'une telle sensibilité exige un contrôle exceptionnel des bruits de fond du détecteur ainsi qu'une compréhension détaillée de la réponse du détecteur, en particulier pour la détection de la lumière de scintillation.

Cette thèse porte sur la caractérisation et la modélisation du transport et de la détection de la lumière dans les détecteurs au xénon liquide à très faible bruit de fond utilisant des photomultiplicateurs au silicium (SiPM), au moyen de l'expérience de recherche et développement Light only Liquid Xenon (LoLX) à l'Université McGill. Ce travail présente les premières mesures en milieu liquide de l'émission de lumière secondaire corrélée provenant des SiPM, pertinentes pour la conception de nEXO. Afin de permettre de futures mesures, l'expérience LoLX a fait l'objet d'une mise à niveau complète, incluant ses systèmes cryogéniques, la géométrie du détecteur, l'acquisition de données, ainsi que les cadres d'analyse et de simula-

tion. Ces améliorations soutiennent des mesures de précision essentielles à la validation des modèles optiques et à la conception d'expériences orientées vers la découverte.

En combinant une caractérisation expérimentale améliorée à des simulations optiques mises à niveau, ce travail fournit une description détaillée du système de détection de la lumière pour nEXO. Les résultats présentés ici constituent un apport essentiel aux projections de performance et à l'optimisation de la conception, et contribuent à l'effort plus large visant à rendre possibles les recherches d'événements rares de prochaine génération avec des détecteurs au xénon liquide.

Acknowledgements

This thesis would not have been possible without the support of many people, both professionally in my research work, and personally throughout my life. First and foremost, I would like to thank my supervisor Dr. Thomas Brunner for his support and confidence in me, and for his patience in listening to my never-ending stream of new ideas. Thomas helped me try to find the balance between chasing after new avenues, and being thorough and detailed in the work on-hand that needed to be done.

I would like to thank all the members of the Brunner Neutrino lab, past and present, for their camaraderie and support over the years. Special thanks go to Soud Al Kharusi for his enthusiastic welcome to the group and Montréal, and for bringing me into his circle of friends, which made my arrival at McGill during a global pandemic much more enjoyable. My time with the Brunner Neutrino group would have been much poorer without the laughs, both in joy and in disbelief, with my fellows. Especially office-mates over the years, Kevin Murray, Hussain Rasiwala, Lucas Darroch, Regan Ross, Samin Majidi, Riya Rai, post-docs Chris Chambers, Bernadette Ribeiro, Sarah Nowicki, and Frederic Girard, and who could ever forget the mysterious Tsvetelin Totev. To the many students I helped supervise, I hope I was able to teach you a thing or two, and if not, then I hope we at least had some fun along the way.

A big thank you to all our LoLX collaborators and PIs. Working in a small project comes with its challenges, but also with its pleasures, and working with all of you has been a pleasure, undoubtedly. I would like to give a special thanks to my long-time comrade in-arms Austin

de St. Croix, from our beginnings in undergraduate classes and research together at Carleton, to writing and publishing a paper together this year on SiPM performance, has been a constant source of laughs and encouragement. Special thanks to the nEXO light simulation team, especially Sierra Wilde and Molly Watts, for their excellent banter and for stomaching my long-winded rants on code maintainability. Thanks also to my nEXO mentors, Erica Caden and Richard Saldanha, for all your advice and support.

I would like to thank all the friends made, and colleagues I've connected with, through my volunteer and outreach coordinator work at McGill. Special thanks to Kim Metera for her unending support and encouragement, especially during our weekend-long Hackathons. I learned so much about teaching and science communication from my fellow outreach coordinators, Hannah Fronenberg, Alice Curtin, and from TSI outreach officer Carolina Cruz-Vinaccia. Being able to share my passion for physics, research work, and scientific computing with the next generation through our outreach events, especially the McGill physics hackathon which I had the pleasure to organize for two years, was incredibly rewarding. Big shoutout to all our participants and volunteers over the years, I hope you had as much fun as we did. During the writing of this thesis my friend and sometimes mentor Professor Dominic Ryan passed away. Dominic was always eager for a lively lunch-time conversation, especially on our two mutual favorite topics, films and poking fun at Thomas. Dominic was incredibly helpful in navigating all the institutional hurdles in securing a custom radioactive source for LoLX, and his help in preparing the samples and input in designing and selecting the source was invaluable. He will be greatly missed by his colleagues at McGill and by all the students who benefited from his wisdom, often delivered by his irreverent wit.

I have had the absolute pleasure of representing my fellow graduate students as an elected member of our graduate society, and then as President. Special thanks to my fellow negotiators and sometimes co-conspirators, Nick Viera, and Nayyer Raza, and to my protégé Catherine Boisvert. Through substantial efforts in surveying, meetings, and arguing, we made a

real impact on our graduate fellows by negotiating the largest raise in departmental history, much needed after the pandemic. Big thanks to all members past and present of the MGAPS council, for their hard work and support of all of us grads.

A big thank you to our physics intramural volleyball team (despite recent events), for your incredible company and fantastic friendship over the years. Especially to Aidan and Tony and Malte, who joined me on a surprise journey to the hospital one late November evening.

Finally, I would like to thank my family and friends for their support and encouragement over the long years of graduate education. Especially my sister Kelly, who, despite having absolutely no interest in physics, was a constant source of encouragement and an ear to listen. Our life was brought to a halt in June 2023, with the sudden passing of our mother, Tracy McRae. Mom was always so proud of us both, and while she never really understood why I was still in school after all these years (which Big Bang Theory character are you again?), she loved and encouraged me powerfully. Mom, I hope you would be proud of what I have accomplished so far and the person I have become. Thank you for never giving up on me, even when the guidance councilors and teachers already had, and my apologies to anyone who stood in your way. Thank you to my father, Jim, and brother John, who never wavered in their constancy.

Last but absolutely not least, I would like to thank my loving and supportive life partner Michelle. I do not know where I would be without you, and your never-ending support and belief in me is a constant source of inspiration. The impossibility of losing one's mother in the middle of a PhD, and continuing on, was only made possible through your care, love, support and encouragement. I am so excited about the next stage of our lives, and I cannot wait for more adventures in the future. Special thanks to our beloved Mr. Beans, for all his cuddles, and I'll even forgive his early-morning yowling.

Science may sometimes happen in a vacuum, but scientists are humans and we are nothing without the people around us and the contributions of those who came before us.

Statement of Contributions

This research work would not be possible without the efforts of many collaborators in both the nEXO and LoLX research groups. The design of the nEXO experiment was in an advanced state before I began this project, and the LoLX experiment was constructed and began commissioning with liquid xenon at McGill. Here I list explicitly the contributions of collaborators to the work presented here, and my own.

Chapter 1 - Introduction

My introduction to the history of neutrinos, and an overview of the Standard Model and neutrino physics relevant to neutrinoless double beta decay. I also include an overview of experimental interpretations for neutrinoless double beta decay measurements, and introduce key considerations in detection and background suppression in a general way with supporting citations for further reading.

Chapter 2 - LoLX Experiment Overview

My introduction to the LoLX experiment, which was developed as a collaboration through researchers at McGill University, Carleton University, TRIUMF national accelerator centre, and the Université de Sherbrooke. The upgraded analysis toolkit, MERCI, was created from the MIDAS `manaLyzer` framework developed by TRIUMF research scientists in the DAQ group, and adapted for our use. This codebase was created starting from the base implementation

done by TRIUMF researcher Andrea Capra, and adapted to LoLX by myself, with submodule contributions from past McGill group members, Bernadette Ribeiro, Soud Al Kharusi, and Liang Xie from TRIUMF. The development of the phase-2 analysis and online code was led and done by myself, and the upgraded simulation toolkit using Chroma was also led and carried out by myself. The hardware upgrades to the cryogenic system were led by myself, with initial cooling power estimates and advice provided by Thomas Brunner and former group member Lucas Darroch. I was the main designer for the upgraded cooling block, creating, reviewing and arranging for manufacturing of all parts. The installation of the upgraded components was carried out with assistance from former postdocs Bernadette Ribeiro and Sarah Nowicki.

The upgraded LoLX phase-2 detector was designed as a collaborative effort with contributions from myself, with design and fabrication carried out by TRIUMF engineering support staff. The installation of the upgraded detector would not have been possible without the extensive contributions from TRIUMF group members who travelled to McGill, including Stephanie Bron, Austin de St. Croix, Colin Hempel, and Carleton PhD student Bindiya Chana. The upgraded DAQ system was created by the MEG-II collaboration and on loan, with installation and setup carried out with assistance from Université de Sherbrooke group members Jessee LeFebvre and Marc-André Tetrault. DAQ integration and readout decoding was my work, alongside coordination and overseeing of all integration efforts.

Chapter 3 - Measuring SiPM ExCT with LoLX

In this chapter, I present a modified version of the work that was published in the European Physical Journal C in early 2025. In this work, we presented the measurement of SiPM secondary light detection, called SiPM external crosstalk, in a liquid xenon detector under nominal operational conditions. For this manuscript, I was the lead author and editor, and main contributor, with substantial support and writing provided by TRIUMF PhD student Austin

de St. Croix. The measurement method presented was conceived by myself after my preliminary work on measuring and constraining the performance-limiting fluorescence in the phase-1 detector, enabled by the new analysis framework outlined in the first chapter. After preliminary analysis showing the method’s viability, analysis proceeded collaboratively with significant contributions from former postdoc Bernadette Ribeiro, and PhD student Soud Al Kharusi. I led the data taking for the data used for this analysis, with assistance from Soud Al Kharusi, and LoLX collaborators for remote detector monitoring.

The simulation method would not have been possible without generous contributions from TRIUMF collaborators Austin de St. Croix and Kurtis Raymond, who carried out in-lab measurements of SiPM ExCT emission spectra and shared ANSYS simulations for use in our Geant4 model. The analytical model presented was developed iteratively in a back-and-forth way by myself, Austin de St Croix and Fabrice Retiere. Austin contributed substantially to modelling SiPM PDE in LoLX Geant4 as well, while code implementation, testing and execution was my work.

Chapter 4 - LoLX phase 2 commissioning

In this chapter, I outlined the commissioning runs taken to-date with the upgraded second phase of LoLX experimentation, which was introduced in chapter 2. As the LoLX run coordinator since 2021, the data was taken by myself, including DAQ configuration, run plan devising and oversight. Overnight detector monitoring was supported by LoLX collaboration members from multiple institutions. The upgraded DAQ was supported with efforts from Sherbrooke group members, with training and trigger implementation by INFN scientists Luca Galli and Marco Francesconi.

The analysis tool upgrades were my own work, and all analysis presented was done by myself solely. We received substantial assistance from McGill Professor Dominic Ryan when planning and developing an internal calibration source using hafnium. The hafnium source

characterization was done by summer intern student Ethan Coulthard, under the supervision of myself and Thomas Brunner.

Chapter 5 - Introduction to nEXO

In this chapter, I review the proposed nEXO experiment, which was developed over almost a decade by the nEXO collaboration. I reference heavily on the extensive publications describing nEXO R&D and system designs, and the 2021 work on sensitivity projections. I add no new work to this, only present the work done by collaborators with my interpretations, with special attention to the light detection system. The nEXO collaboration builds on the efforts from the previous EXO-200 experiment collaboration, with extensive sharing of experience and best practices for the tonne-scale design of nEXO.

Chapter 6 - Simulating the nEXO Light Detection System

In this chapter, I review the previous stage of nEXO light simulations, work carried out primarily by Ako Jamil at Yale University. As one of the main developers of nEXO's Chroma simulation package, I also led and carried out a light simulation workshop at McGill University in 2024 to concentrate efforts necessary for upgraded simulations and framework. I led and carried out a major refactoring of the codebase to make our wrapper framework, `chroma-simulation`, more flexible for use by collaborators.

The upgraded framework included substantial contributions from TRIUMF PhD student Alex Li (who implemented the 2D SiPM PDE look-up-table and did preliminary testing). Using Alex's updated SiPM model, based on data shared from TRIUMF collaborators (pre-publication) for their detailed full wavelength domain measurement campaign, we made substantial updates to all optical surfaces in the nEXO detector simulation. During the period of my thesis work, I also supervised several undergraduate student projects with various contributions to components of the Chroma simulation used for the work presented in this chapter.

This includes work on Cherenkov implementation by Sandhya Rottoo, and integration-testing and initial LoLX Chroma simulation testing by Simon Lavoie.

The updated simulation presented is my work, including all analysis of simulation output, and the development of the SiPM response model simulation. The SiPM response model was supported by input data measurements done by collaborators within nEXO.

Chapter 7 - Discussion and Conclusions

The final chapter is a summary of the work presented in my thesis, with my interpretations and potential directions for future work.

Contents

I	Introduction	1
1	Introduction	2
1.1	A Brief History of Neutrinos	4
1.2	The Standard Model of Particle Physics	7
1.3	Neutrinos	9
1.3.1	Neutrino Mass Theories	12
1.4	Neutrinoless Double Beta Decay	13
1.4.1	Beta and Double Beta Decay	14
1.4.2	$0\nu\beta\beta$ decay process	16
1.4.3	Experimental Signatures of $0\nu\beta\beta$	17
1.4.4	Experimental Limits and Interpretations	18
1.5	Experimental Considerations for $0\nu\beta\beta$ Searches	22
1.5.1	Background Signals	22
1.5.2	Background Mitigation	23
II	Characterizing SiPM Performance in Liquid Xenon	25
2	The Light Only Liquid Xenon Experiment	26

2.1	Motivation	26
2.1.1	Xenon Scintillation	27
2.1.2	Silicon Photomultipliers	34
2.2	LoLX Support Infrastructure	39
2.2.1	Gas Handling and Recovery System	39
2.2.2	Cryogenics System and Upgrades	42
2.3	LoLX Detectors	47
2.3.1	Phase-1 Detector	47
2.3.2	LoLX Phase-2 Detector	51
2.4	LoLX Data Acquisition and Slow Control	54
2.4.1	Slow Control with MIDAS	54
2.4.2	Data Acquisition System	56
2.5	Analysis and Simulation Software	58
2.5.1	Online/Offline Analysis Software	58
2.5.2	LoLX Geant4 Simulations	66
2.5.3	GPU-based Light Transport Simulations	70
3	Measuring SiPM ExCT	74
3.1	Motivation	74
3.1.1	SiPM Optical Crosstalk	75
3.2	Measurement Technique	77
3.2.1	Analysis Criteria	79
3.2.2	SiPM ExCT Correlation Probability	83
3.3	Simulating SiPM ExCT	86
3.3.1	Simulating SiPM Surface Optics	92
3.3.2	Simulation ExCT Analysis	95
3.4	Modelling ExCT Coincidence Probability in LoLX	97

3.5	Results	100
3.5.1	Systematic Study	102
3.5.2	ExCT probabilities and secondary emission yields	105
3.5.3	Discussion	108
4	LoLX Phase-2 Commissioning	111
4.1	LXe Commissioning Runs	112
4.1.1	Run Overviews	112
4.1.2	Results	115
4.2	Simulation Comparison	123
4.2.1	Geant4 Input	123
4.2.2	Chroma Simulations in LoLX phase-2	125
4.3	Calibration Sources and Future Runs	134
4.3.1	Hafnium-181 Internal Calibration Source	135
4.4	Conclusion and Next Steps	139
III	Simulating the nEXO Light Detection System	142
5	The nEXO Experiment	143
5.1	nEXO Overview	145
5.1.1	Outer Detector	145
5.1.2	Inner Detector	146
5.2	nEXO Multiparameter Data Analysis	153
5.3	nEXO Sensitivity to Neutrinoless Double Beta Decay	156
6	nEXO Light Simulations	160
6.1	nEXO Light Simulations	161

6.2	Simulation Upgrades and Validation Testing	164
6.2.1	Refactored chroma-simulation Codebase	164
6.2.2	Updated Simulation Optical Parameters	168
6.3	Simulating SiPM ExCT in nEXO	176
6.3.1	ExCT Implementation in chroma-simulation	176
6.3.2	Validation Testing and Spatial Dependencies	179
6.3.3	ExCT Induced Correlated Noise in nEXO	184
6.3.4	nEXO Light Collection Efficiency Map	187
6.4	SiPM Response Modelling	191
6.5	Revisiting nEXO Energy Resolution	197
6.6	Conclusions	203

IV Discussion and Conclusions 208

7 Discussion and Conclusions 209

List of Figures

1.1	Standard Model of Elementary Particles.	9
1.2	Illustration of forbidden and allowed decays for ^{136}Xe and neighbouring isotopes.	15
1.3	Illustration of the experimental signatures of standard and neutrinoless double beta decay in ^{136}Xe	18
1.4	Feynman diagram of $0\nu\beta\beta$ through light Majorana neutrino exchange. . . .	19
1.5	Lobster plot showing available phase-space for $0\nu\beta\beta$ in effective Majorana mass vs lightest neutrino mass.	21
2.1	Illustration of liquid xenon scintillation production and decay pathways. . . .	28
2.2	NEST LXe scintillation yields.	30
2.3	Illustration of energy loss pathways in LXe.	32
2.4	Cherenkov cone visualization.	33
2.5	Illustration of the SiPM operational principle.	35
2.6	Illustration of SiPM correlated noise pathways with expected SiPM waveform readout signals.	36
2.7	Photo of the LoLX experiment at McGill University from November 2024. . .	39
2.8	Illustration of xenon flow into LoLX in a simplified version of the LoLX XGHS P and ID	41

2.9	Phase diagram of xenon.	42
2.10	LXe vapour pressure vs temperature.	43
2.11	MIDAS Slow control temperature readout of the LoLX phase-1 October 2021 LXe run.	45
2.12	Temperature and pressure conditions of the LoLX phase-2 cryostat with up- graded cooling system during January 2025 commissioning run.	46
2.13	The LoLX Phase-1 detector during test assembly.	47
2.14	Decay pathway for ^{90}Sr	48
2.15	Transmission of LoLX-1 filters.	49
2.16	Illustration of LoLX detector unwrapped.	50
2.17	LoLX Phase-2 detector during assembly.	53
2.18	Illustration of LoLX phase-2 cooling system and slow control readout variables.	55
2.19	Overview illustration of the LoLX phase-1 and phase-2 DAQ data flow.	56
2.20	MERCI LoLX Analysis flow overview.	60
2.21	Comparison of LoLX Phase-1 DAQ digitizer and Phase-2 digitizer.	61
2.22	LoLX phase-2 waveform analysis overview.	64
2.23	LoLX phase-2 Geant4 simulated geometry.	69
2.24	Ray-tracing rendering of LoLX phase-2 detector in the Chroma simulation detector viewer.	71
3.1	Visualization of primary scintillation and ExCT signatures in an inverted LoLX phase-1 detector.	78
3.2	Illustration of SiPM ExCT search method with candidate pulses.	79
3.3	Occupancy of all SiPMs within LoLX for 4 V OV data collections.	82
3.4	SPE charge distribution and selection window for an unfiltered SiPM channel in LoLX.	83
3.5	5 V OV dataset ExCT delta-time distribution, with linear fit overlaid.	84

3.6	Wavelength dependence of LoLX phase-1 ExCT simulation parameters. . . .	87
3.7	Illustration of the Geant4 SiPM ExCT simulation method with labelled steps.	89
3.8	Transmission probability for secondary photons within Si to enter into LXe. .	90
3.9	Secondary photon distribution for ExCT into LXe from Geant4 simulation. .	91
3.10	Cross-section illustration of HPK SiPM surface for optical simulation. . . .	93
3.11	Comparison of high N summation versus approximate relation for ExCT prob- abilities.	98
3.12	Plot of unperturbed channel rate ratio (R_R) versus measured rate ratios ($\widetilde{R}_R$) for a range of N_{Xe} values.	99
3.13	4 V OV dataset delta-time distribution, with linear fit overlaid.	101
3.14	LoLX SiPM ExCT comparison of effective ExCT model and data.	108
4.1	LoLX phase-2 detector offline event-display example.	115
4.2	Example of a selection of FBK SiPM channel reverse IV curves at room tem- perature and in-LXe at 170 K.	116
4.3	FBK channel SiPM ID#7 SPE calibration plots.	118
4.4	LoLX phase-2 light asymmetry distributions for ^{133}Ba 4 V OV data runs. . .	121
4.5	Energy spectrum comparison of external ^{133}Ba source data and background- only data from 4 V January 2025 LoLX phase-2 commissioning.	122
4.6	2D light channel signal projections of external ^{133}Ba source data with LoLX.	123
4.7	Total event energy of ^{133}Ba decay events that deposited energy in the LoLX Geant4 LXe scoring volume.	124
4.8	External ^{133}Ba decay events energy deposition positions in LoLX phase-2 detector scoring volume.	125
4.9	Light collection efficiency map for scintillation events in LoLX phase-2 de- tector.	127

4.10	Pure detected photon and estimated PE spectra comparison for simulated external ^{133}Ba source events.	128
4.11	Comparison of simulated and measured external ^{133}Ba source energy spectrum, in total detected PE.	129
4.12	Profile distribution of simulated trigger efficiency for LoLX phase-2 commissioning data with external ^{133}Ba source events.	130
4.13	2D light channel simulation projections for external ^{133}Ba source in LoLX phase-2.	131
4.14	Comparison of simulated and measured LXe time profiles in LoLX phase-2 for external ^{133}Ba source calibrations.	133
4.15	Decay schematic for ^{181}Hf	136
4.16	Sealed hafnium source and witness samples before irradiation at McMaster in March 2025.	137
4.17	HPGe spectrum of ^{181}Hf wire sample overlaid with HPGe background, displayed in MCA channel number.	138
5.1	Sketch of nEXO pre-CDR detector concept.	145
5.2	Geant4 2/3 projection visualization of nEXO TPC	147
5.3	Illustration of LXe TPC operational principles.	148
5.4	TPC Rotated energy plot from nEXO simulations.	150
5.5	Overview of nEXO light collection system components.	152
5.6	nEXO Multi-parameter analysis illustrated overview.	155
5.7	nEXO sensitivity to neutrinoless double beta decay projection.	157
5.8	nEXO sensitivity lobster plot.	158
6.1	Simulated photon transport efficiency map of nEXO TPC.	162

6.2	Schematic diagram outlining the modular components of the updated chroma-simulation architecture.	166
6.3	LXe optical simulation input data.	169
6.4	Total reflectivity of $\text{MgF}_2\text{-Al-MgF}_2$ coatings modelled using POCAL [177].	171
6.5	Simplified Chroma simulation test geometry.	172
6.6	Geometry simulation testing results for nEXO anode tile.	173
6.7	Full wavelength-dependent gold anode surface refractive indices.	174
6.8	COMSOL nEXO electric field simulation projections.	176
6.9	3D visualization of simulated nEXO SiPM channel positions.	178
6.10	ExCT emission profile simulation validation.	179
6.11	ExCT detection hit pattern in unwrapped nEXO TPC barrel.	180
6.12	Prompt photon travel time distribution for SiPM ExCT in nEXO.	181
6.13	Comparison of ExCT photon travel time for varying emission positions in nEXO.	182
6.14	Height and wavelength dependence on ExCT transport in nEXO.	183
6.15	nEXO ExCT collection efficiency map in Z vs wavelength.	185
6.16	Profile histogram projections of nEXO ExCT collection efficiency map. . . .	185
6.17	Correlation probability for HPK SiPM ExCT vs overvoltage in nEXO. . . .	186
6.18	nEXO 2D Light collection efficiency maps for FBK and HPK configurations.	188
6.19	Projections in R and Z of nEXO light collection efficiency maps.	188
6.20	Absorbed photon positions in R-Z projection for nEXO.	189
6.21	SRML core simulation overview flowchart.	193
6.22	Excess charge factor distributions for varying numbers of avalanches per event.	195
6.23	Simulated ECF and CAF results from SRML for FBK and HPK SiPMs. . . .	196
6.24	Comparison of reconstructed and simulated photons for $0\nu\beta\beta$ events in nEXO TPC FV for the HPK configuration.	199

6.25	Comparison of reconstructed and simulated photons for $0\nu\beta\beta$ events in nEXO	
	TPC FV for the FBK configuration.	200
6.26	Model of nEXO total energy resolution at $Q_{\beta\beta}$	202

List of Tables

1.1	Summary of selected isotopes considered for $0\nu\beta\beta$ searches.	16
1.2	Summary of selected isotopes for $0\nu\beta\beta$ and their current and next-generation experiments and sensitivities.	16
2.1	List of step-level Geant4 output variables for Chroma simulations.	70
3.1	Systematic uncertainties considered, based on data and simulations.	102
3.2	Input parameters for ExCT model with 1-sigma uncertainties considered in the analysis.	105
3.3	External cross-talk probabilities from the time correlation analysis, and the resulting mean number of photons emitted per avalanche within silicon and into LXe.	108
6.1	Reproduced table from [52], nominal simulated input optical parameters and values.	163
6.2	High-level comparison of chroma-simulation architectures before and after updates.	167
6.3	VUV reflective coating stack design for nEXO.	170
6.4	Comparison of SRML results with measured values.	197
6.5	Contributions of simulated nEXO light channel energy resolutions at $Q_{\beta\beta}$, for both SiPM vendor configurations at 3 V OV.	201

Acronyms

- **3P2S**: Three Paralell Two Series
- **ADC**: Analog to Digital Conversion
- **AP**: AfterPulse
- **BP**: BandPass filter
- **BDV**: BreakDown Voltage
- **CAF**: Correlated Avalanche Fluctuation
- **CE**: Collection Efficiency
- **DCR**: Dark Count Rate
- **DNN**: Deep Neural Network
- **DeCT**: Delayed optical CrossTalk
- **DiCT**: Direct optical CrossTalk
- **EC**: Excess Charge
- **ECF**: Excess Charge Factor
- **EXO**: Enriched Xenon Observatory

- **ExCT**: External Crosstalk
- **FBK**: Fondazione Bruno Kessler
- **FFT**: Fast Fourier Transform
- **FOV**: Field Of View
- **FSR**: Field Shaping Ring
- **FV**: Fiducial Volume
- **GEANT**: GEometry ANd Tracking
- **GPU**: Graphics Processing Unit
- **GSPS**: GigaSamples Per Second
- **HV**: High Voltage
- **HPK**: Hamamatsu Photonics K.K
- **ID**: Inner Detector
- **LED**: Light Emitting Diode
- **LoLX**: Light only Liquid Xenon
- **LXe**: Liquid Xenon
- **MERCI**: Modular onlinE multi-thReaded C++ based waveform toolkIt
- **MC**: Monte Carlo
- **MPPC**: Multi Photon Pixel Counter
- **OD**: Outer Detector

- **PD**: Photodiode
- **PMT**: PhotoMultiplier Tube
- **PPS**: Pulses Per Second
- **PDE**: Photon Detection Efficiency
- **PDF**: Profile Density Function
- **QE**: Quantum Efficiency
- **RMS**: Root Mean Square
- **RFI**: Radio-frequency interference
- **RT**: Room Temperature
- **SiPM**: Silicon PhotoMultiplier
- **SM**: Standard Model
- **SNR**: Signal to Noise Ratio
- **SPE**: Single PhotoElectron
- **SRML**: SiPM Response Model Library
- **SNO**: Sudbury Neutrino Observatory
- **TDC**: Time to Digital Converter
- **TPC**: Time Projection Chamber
- **ToT**: Time over Threshold
- **WIPP**: Waste Isolation Pilot Plant

Part I

Introduction

1

Introduction

The science of fundamental physics has continued to push detector technologies to evolve and grow, to reach the precision necessary for the discovery of physics beyond our current understanding of the Universe. At the so-called “intensity frontier” of the search for new physics, we push the sensitivity of instruments to the ultimate limit, searching for extremely rare events, with rates as low as a single event detection per year in a tonne-scale experiment. To achieve this, we need ultra-low background detectors, coupled with precise modelling and accurate characterizations. The proposed nEXO experiment is one such detector, and aims to reach world-leading sensitivity for the rare hypothetical nuclear decay mode, neutrinoless double beta decay, for half-lives up to $\sim 10^{28}$ y. nEXO will employ liquid xenon as both detector media, and target, enriched to 90 % in the double beta decaying isotope ^{136}Xe . It will

measure both the scintillation light and ionization electrons produced from incident radiation in liquid xenon, and combine the two signals to achieve its ultimate performance goals. To reach such unprecedented sensitivities, it is critical to develop a detailed understanding of the experiment using validated simulation tools that can be bench-marked against prototypes and existing detectors. Here we reach the goals of this thesis: to develop a detailed understanding of light transport for an ultra-low background liquid xenon detector that employs state-of-the-art single-photon light sensors (SiPMs), and to characterize the performance of these light sensors in a liquid xenon experiment. In order to achieve these goals, we must develop more sophisticated analysis and simulation tools. To this end, we employed the R&D experiment “Light only liquid xenon” (LoLX), located at McGill University, to perform the first in-liquid characterization of the correlated detection of secondary light emission from the light sensors being considered for use in the nEXO experiment. To go beyond the limitations of the earlier design, we upgraded the LoLX experiment in every aspect: cryogenics and control, detector design and data acquisition, as well as the tools used for analysis and simulations. These upgrades will enable the critical measurements needed for a discovery-oriented experiment like nEXO. By combining upgrades to simulations with our improved understanding of light transport and detection with SiPMs in liquid xenon from LoLX, we can, for the first time, fully characterize the light detection system of the proposed nEXO experiment.

This thesis is separated into 3 main parts. First, an introduction to neutrino physics and neutrinoless double beta decay, its motivation, and detection considerations. This is followed by a three-chapter part on our R&D experiment LoLX: an overview of the experiment, including all relevant systems and tools upgraded as part of this work, the measurement of SiPM secondary light emission in liquid xenon, and the preliminary commissioning of the upgraded LoLX experiment. The final part is dedicated to nEXO, with an introduction to the experiment and its current design, followed by an overview of the upgraded optical simulations and updated estimates of the performance of the light detection system.

1.1 A Brief History of Neutrinos

Neutrinos were first evoked to explain the reported continuous beta radiation energy spectrum, initially observed in 1911 by Lise Meitner and Otto Hahn [1], later confirmed through calorimetry in 1927 by Ellis and Wooster [2]. Beta radiation, in contrast to alpha radiation, that was observed to be emitted in discrete energy lines, followed a continuum. This apparent “missing” energy created a minor crisis in physics, calling into question a fundamental law, the conservation of energy. This worrying observation remained unexplained until 1930, when a letter from Wolfgang Pauli¹ proposed the existence of a light neutral spin 1/2 particle (which he called neutrons) [3]. Shortly after, the neutrino received its name from Enrico Fermi², who in 1933 derived the first theoretical framework for beta radiation. The energy in this decay was not truly missing, only carried away by the neutrino, that, at the time, was undetectable. With the keen postulations of Fermi and Pauli, and supporting work by experimentalists, the crisis in physics was resolved, re-establishing the law of conservation of energy.

Fermi’s theory led to many more developments, including the proposition of a so-called “double” beta decay by Maria Göppert-Mayer in 1935 [4]. In Göppert-Mayer’s theory of double beta decay, she considered the consequences of a nucleus with forbidden beta decay transitions, possible in the so-called “even-even” nuclei with equal numbers of neutrons and protons. These isotopes were expected to be more strongly bound than neighbouring odd-odd nuclei. Shortly afterwards, while developing an extension of Dirac’s groundbreaking work on anti-electrons, called “positrons” [5], Ettore Majorana developed a theory for real-valued solutions to a modified version of Dirac’s theory [6]. This insight posited the existence of neutral fermions, which are their own anti-particle, in contrast to the Dirac anti-particle formulation. Neutrinos are an example of neutral fermions. We now refer to these hypothetical particles as

¹Pauli’s letter was sent to scientists at the Gauverein meeting in Tübingen, and was famously addressed “Dear Radioactive Ladies and Gentlemen”.

²In the meantime, what we now call the neutron was discovered, hence neutrino or “little neutral one” from the Italian.

“Majorana” fermions, in honour of Ettore’s contributions. The connection to Mayer’s double beta decay formulation was made in 1939 by Wendel Furry, by demonstrating that if neutrinos are Majorana particles, double beta decay could occur without the emission of neutrinos [7]. This would violate the conservation of a quantity known to be preserved in all observed interactions, called lepton number, and provide an experimental test of Majorana neutrino nature. We refer to this hypothetical decay mode as “neutrinoless double beta decay”, and searching for it experimentally is the motivation of this thesis work. The work led by these pioneers blazed the trail for modern experimental searches into neutrino properties, however, concrete estimates of the rate of these hypothetical decay events were not possible at the time, and the prospect of ever detecting neutrinos at all was bleak [8].

The 1930’s were a very active time for research into the fundamentals of particle physics and its mathematical framework in quantum mechanics. It would take another two decades before we obtained the first direct evidence of the existence of neutrinos, through their capture in the inverse process of beta decay. The experiment, aptly known as “Project Poltergeist”³, led by Clyde Cowan and Frederick Reines in 1956, was the first to confirm the existence of neutrinos. The experiment used a large tank of proton-rich liquid scintillator [9], loaded with cadmium to detect anti-neutrinos from a nearby nuclear reactor. The experiment depended on the delayed-coincidence signature from the inverse-beta decay capture of neutrinos, producing a prompt signal from the emitted positron annihilation, with the delayed signal from the neutron being captured on the added cadmium. This was a major milestone in the story of the neutrino, providing the first direct experimental evidence of its existence, confirming the theory of inverse-beta decay, and establishing a technological strategy for their detection using massive volumes of sensitive media ⁴.

Following from the first detection by Reines and Cowan, there was critical work into

³Owing to the neutrinos “ghostly” weakly-interacting nature.

⁴Frederick Reines was awarded the Nobel Prize in Physics in 1995 for the detection of the neutrino. Clyde Cowan had passed away in 1974 and was therefore ineligible [10].

the foundation of neutrino interactions by Chien-Shiung Wu in 1957, establishing the chiral nature of neutrinos [11]. Afterwards, experimentalists turned their attention to the most abundant source of neutrinos expected, the Sun. Neutrinos were estimated to be produced in massive quantities from nuclear fusion reactions in the core. The Homestake experiment, led by Ray Davis⁵ and John Bahcall, began in 1965 and ran from 1970 – 1994, counting neutrino captures on chlorine (in the form of a giant tank of dry-cleaning fluid) by measuring the gaseous argon produced by the capture reaction [13]⁶. The initial results confirmed that nuclear fusion reactions in the core of the sun drove energy production. The second main result of this experiment was puzzling, as the measured flux of neutrinos from the sun was significantly lower than predicted by Bahcall, roughly 1/3 of the expected signal [14]. This would come to be known as the “solar-neutrino problem”, and its resolution would require another two decades of experimentation.

Early interpretations of the discrepancy focused on uncertainties in solar models or experimental systematics. However, Bruno Pontecorvo had earlier proposed that neutrinos might undergo flavour transformations during propagation, which would result in a fraction of solar neutrinos being undetectable by experiments such as Homestake [15]. This idea gained support as improved solar models continued to predict higher fluxes than observed, and independent experiments consistently confirmed the deficit [16]. Definitive evidence for neutrino flavour change, which would come to be known as neutrino oscillation, was found at the end of the 1990s. The Super-Kamiokande experiment observed a deficit of atmospheric neutrinos passing through the Earth, dependent on zenith angle, providing the first compelling demonstration of neutrino oscillations [17]. The Sudbury Neutrino Observatory (SNO) directly measured both the electron-neutrino component (the signal measured by Homestake),

⁵Ray Davis was awarded the Nobel Prize in Physics in 2002, shared with Masatoshi Koshiba and Riccardo Giacconi, for pioneering contributions to astrophysics, in particular for the detection of cosmic neutrinos [12].

⁶This radio-chemical detection method was originally proposed by Bruno Pontecorvo in 1946, while working at the Canadian Nuclear Laboratories. Pontecorvo’s prescience was firmly established through his 1956 prediction of the oscillating neutrino hypothesis. This would come to explain the results of the Homestake experiment, that wouldn’t begin until 10 years later.

and the total neutrino flux from the Sun, demonstrating that the missing solar neutrinos were present in other flavours after transforming states [18]. Together, these results resolved the solar-neutrino problem and established that neutrinos have non-zero mass, the first confirmed experimental evidence for physics beyond the Standard Model [19]⁷.

While incomplete, we have covered the history of the neutrino, from its conception, and the first discovery, to its confounding nature that still eludes us. Much has been omitted, and we recommend the reader consider the third edition of Kai Zuber’s seminal text “Neutrino Physics” for a more complete review [21]. We will now move onto a high-level overview of the Standard Model of particle physics. This will be followed by a description of neutrino theory, especially those points relevant to Majorana neutrinos and their mass generation mechanisms.

1.2 The Standard Model of Particle Physics

The Standard Model (SM) of particle physics is a crowning achievement of 20th century physics. The model is based in quantum field theory and describes all elementary particles and three of the four fundamental forces of nature, excluding gravity. Over decades of experimentation, the model has been tuned to exceptional precision, and provides a framework for understanding a broad range of phenomena from radioactive decays to collisions in particle colliders [22]. Importantly, the model predicted the existence of several particles, most notably the Higgs boson, which was discovered in 2012 [23].

Symmetry plays an essential role in modern particle physics theory, building on the contributions of Emmy Noether. In 1918, Noether proved that every continuous and differentiable symmetry of the action of a physical system, has a corresponding conservation law [24]. For example, a physical process that produces the same outcomes regardless of choice of location, is said to be invariant under translation in space, and corresponds to the conservation of linear

⁷Takaaki Kajita and Arthur B. McDonald were awarded the Nobel Prize in Physics in 2015 for the discovery of neutrino oscillations, which demonstrated that neutrinos have mass [20].

momentum.

Going from fundamental symmetries at a global level to local transformations gives additional constraints on a physical theory. Maintaining invariance under these local symmetries requires the introduction of gauge fields, which manifest physically as the interactions between particles. This provides a framework in which forces arise from symmetry principles. The high energy gauge symmetry group of the SM is:

$$SU(3)_C \times SU(2)_L \times U(1)_Y, \quad (1.1)$$

and represents the allowed interactions between fundamental particles in the SM. Each term corresponds to a fundamental force, with $SU(3)_C$ representing the strong nuclear interaction, and $SU(2)_L \times U(1)_Y$ representing the electroweak interaction, which at low energies breaks into electromagnetism and the weak nuclear interaction through a process called electro-weak symmetry breaking [25]. Matter fields are given conserved quantities (quantum numbers), which under these symmetries, determine how they interact with force carriers. In this structure, the observed patterns of particles and interactions emerge from the principles of symmetry, rather than being imposed ad-hoc.

The fundamental constituents of matter in the SM are fermions, which obey the Pauli exclusion principle and have half-integer spins. They are organized into three generations, each containing two quarks and two leptons. In contrast, bosons are integer-spin particles that mediate the fundamental forces: the photon, W and Z bosons, gluons, and the Higgs boson. The distinction between fermions and bosons comes from the allowed quantum statistics and interactions in the theory, and their roles are illustrated in Figure 1.1. In the SM we have two subclasses of fermions, one that feels the strong nuclear interaction, quarks, and one that does not, leptons. There are charged leptons (electron, muon, tau), that interact electromagnetically and weakly, and neutral leptons (neutrinos), which only interact weakly. In SM interactions, the total number of leptons does not change, an observation known as the conservation of lepton number. For example, in ordinary beta decay ($n \rightarrow p + \beta^- + \bar{\nu}_e$), the creation of an

electron ($L = +1$) is accompanied by an antineutrino ($L = -1$), keeping the total lepton number zero.




Figure 1.1: Overview figure highlighting all elementary particles in the SM, reproduced from [26]. Particles are categorized into matter (fermions) and interactions (bosons), and further separated by matter generation.

While the symmetry structure of the SM successfully describes a broad range of experimental data, it predicts neutrinos to be massless, which is in direct conflict with the observation of neutrino oscillations [18]. This discrepancy provides one of the clearest indications of physics beyond the SM.

1.3 Neutrinos

Neutrinos exist in three flavour states; electron, muon, and tau, corresponding to the three mass generations of fermions shown in Figure 1.1. We have introduced the impact of neutrino

oscillations in the historical context of the solar neutrino problem, where electron flavoured neutrinos from nuclear fusion reactions in the sun are transformed into muon and tau flavours on their way to earth, causing a deficit in the observed neutrino flux. This happens because neutrinos are created and captured through weak interactions as well-defined flavour eigenstates corresponding to the lepton involved in the process. Neutrino flavour eigenstates are superpositions of mass eigenstates with differing masses (ν_1, ν_2, ν_3). Meanwhile, as neutrinos travel through space their individual mass eigenstates evolve differently in time due to the different masses, causing a corresponding evolution in the flavour state of the travelling neutrino. This phenomenon is what we call “neutrino oscillation”, and the mixing between neutrino mass states and flavour states is represented by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix [27]. The mixing can be written as,

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = U_{\text{PMNS}} \cdot \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix}, \quad (1.2)$$

$$U_{\text{PMNS}} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{cp}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{cp}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{cp}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{cp}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{cp}} & c_{23}c_{13} \end{pmatrix} \cdot P_{M/D}, \quad (1.3)$$

where, c_{ij} is a shorthand for $\cos \theta_{ij}$, s_{ij} is a shorthand for $\sin \theta_{ij}$, with $i, j \in \{1, 2, 3\}$, the neutrino mass eigenstates. δ_{cp} is a complex phase term that allows for CP violation in neutrino oscillations [28], as a non-trivial value of δ_{cp} results in differing oscillation probabilities for neutrinos and antineutrinos due to the conjugate phases acquired by their respective amplitudes. The matrix $P_{M/D}$, is the identity matrix in the case of Dirac neutrinos, or in the case of Majorana neutrinos:

$$P_M = \begin{pmatrix} e^{\frac{i\alpha_1}{2}} & 0 & 0 \\ 0 & e^{\frac{i\alpha_2}{2}} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (1.4)$$

where $\alpha_{1/2}$, are the two Majorana phases. These Majorana phases are additional CP-violating terms and do not appear in neutrino oscillation observables, because the oscillation probabilities depend on combinations of the terms where the phases cancel. They will play a role in lepton-number violating processes such as $0\nu\beta\beta$.

To better understand this, we will relate the probability that a neutrino ($\nu_\alpha = \nu_{e/\mu/\tau}$) produced at time $t = 0$, is observed in a different flavour state ν_β at a later time t . First we write the neutrino state at $t = 0$:

$$|\nu(t=0)\rangle = |\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle, \quad (1.5)$$

where, $U_{\alpha i}^*$ is the conjugate of the PMNS matrix element at flavour α and mass-state i . Recalling that mass eigenstates evolve in time through the Hamiltonian operator, $e^{-iE_i t}$ [28], at later times this becomes:

$$|\nu_i(t)\rangle = e^{-iE_i t} |\nu_i(t=0)\rangle, \quad (1.6)$$

for each mass eigenstate, ‘ i ’. Coming together, the probability $P_{\alpha \rightarrow \beta}$ is:

$$P_{\alpha \rightarrow \beta} = |\langle \nu_\beta | \nu_\alpha(t) \rangle|^2 = \left| \sum_i \sum_j U_{\alpha i}^* U_{\beta j} \langle \nu_j | \nu_i(t) \rangle \right|^2. \quad (1.7)$$

If we take the relativistic limit, the distance travelled by the neutrino is $L \approx ct$, where c is the speed of light, and $p_i \approx p_j \equiv p \simeq E$. Then the energy of a neutrino mass state is then $E_i = \sqrt{p^2 + m_i^2} \approx p + m_i^2/2E$. Using this approximation, we can re-write Equation 1.7 as:

$$P_{\alpha \rightarrow \beta} = \sum_{ij} U_{\beta i} U_{\alpha i}^* U_{\beta j}^* U_{\alpha j} e^{-i\Delta m_{ij}^2 L/E}, \quad (1.8)$$

where, $\Delta m_{ij}^2 = m_j^2 - m_i^2$ are the mass-squared differences between mass eigenstates m_i , and m_j [29]. From this representation, we can draw some inferences. The first is that at least two neutrinos mass states, m_i , must be non-zero, otherwise the mass-squared differences vanish. The second is that the oscillation probability is proportional to L/E , a fact exploited by modern oscillation experiments to sample different parts of the probability spectrum [30]. The third inference is critical for our case with $0\nu\beta\beta$, oscillation experiments are insensitive to Majorana phases, and only measure the mass-squared differences between neutrino mass

eigenstates. Therefore, in oscillation measurements, we are insensitive to the absolute scale of neutrino masses. Because of this, the ordering of mass states is not currently known and allows two configurations, referred to as hierarchies. The first, called the “Normal Ordering” (NO), assumes $m_1 < m_2 < m_3$, and the second, called the “Inverted Ordering” (IO) assumes $m_3 < m_1 < m_2$. The mass-squared differences are determined by measuring neutrino rates from specific sources, and are typically referred to by those sources. Δm_{21}^2 is called the “solar” mass splitting, and Δm_{31}^2 is the “atmospheric” mass splitting. Both are measured simultaneously with the mixing angles θ_{ij} of the PMNS matrix. The observed value of Δm_{21}^2 is 2 orders of magnitude smaller than Δm_{31}^2 , indicating that m_1 and m_2 are close in value compared to m_3 , and motivating the two possible configurations [31]. δ_{cp} is constrained by comparing the measured probabilities for neutrino to anti-neutrino oscillations [32].

1.3.1 Neutrino Mass Theories

With non-zero neutrino masses established by neutrino oscillation experiments, we can consider how they may obtain this mass as extensions to the SM. Here we consider two options, a Dirac mass, and a Majorana mass.

To generate a Dirac mass, analogous to other fermions, we require a Yukawa coupling between the left and right-handed neutrino states. This requires us to invoke the existence of right-handed neutrinos and left-handed anti-neutrinos, which interact through the Higgs mechanism. In order to preserve the tiny masses of left-handed neutrinos, we require a vanishingly small Yukawa coupling constant [33].

In contrast, if neutrinos are Majorana particles, we generate their mass through a coupling term between neutrino and anti-neutrino states. For a left-handed neutrino field, we can include the Lagrangian [33]:

$$\mathcal{L}_{Maj.} = -\frac{1}{2}m_M\bar{\nu}_L^c\nu_L + h.c. , \quad (1.9)$$

where the superscript c represents a charge conjugate state, m_M is the Majorana mass, and $h.c.$

is shorthand for Hermitian conjugate terms. Notably, this interaction violates lepton number by 2 units, as $0\nu\beta\beta$ produces two electrons ($\Delta L = +2$) from an initial state with $L = 0$. In this formalism, we have no need to introduce a right-handed neutrino.

To explain the observed tiny neutrino masses⁸ we can combine these two mass terms in what is known as the type-I see-saw mechanism [35], [36]. Because the right-handed neutrinos are singlets under the SM gauge group, they can carry a Majorana mass m_M that is not tied to the electroweak scale and can be very large [37]. Combined with a Dirac mass m_D from the usual Higgs Yukawa coupling, this generates naturally small left-handed neutrino masses for an order-unity Yukawa coupling. For $m_D \ll m_M$, this is approximately:

$$m_1 \approx \frac{m_D^2}{m_M}, \quad m_2 \approx m_M \quad (1.10)$$

for two mass eigenstates, m_1 , associated with a light neutrino that interacts with SM particles, and m_2 , associated with a heavy right-handed Majorana neutrino that does not interact and is referred to as sterile. With a large m_M , this naturally produces very small observable neutrino masses without fine-tuning [28], [35].

With our understanding of neutrino masses, we will now move onto an overview of beta decay, double beta decay, and $0\nu\beta\beta$. This is followed by essential experimental considerations, including the potential detection mechanisms and main background sources.

1.4 Neutrinoless Double Beta Decay

Neutrinoless double beta decay ($0\nu\beta\beta$), is a rare hypothetical nuclear decay mode, only observable in a handful of candidate isotopes where double beta decay is accessible. In this section we will differentiate beta decay, double beta decay, and the hypothetical neutrinoless variant.

⁸So far we only have limits on neutrino masses, which constrain them below 0.45 eV, we can also get a rough estimate by taking the square root of the mass-squared differences observed in oscillations, which are of similar scale [34].

1.4.1 Beta and Double Beta Decay

In beta decay, an unstable atom decays radioactively, emitting a beta particle and a neutrino.

This can be described as:

$$\beta^- : n \rightarrow p + \beta^- + \bar{\nu}_e, \quad \text{or} \quad (1.11)$$

$$\beta^+ : p \rightarrow n + \beta^+ + \nu_e, \quad (1.12)$$

where n, p are neutron or proton, respectively, inside a nucleus. As a result of the decay process, the parent atom is transmuted into the daughter isotope. The atomic number Z increases by one in beta-minus decay and decreases by one in beta-plus decay, while the nucleon number A is unchanged. In both cases, we conserve lepton number through the emission of a corresponding neutrino. The distribution of beta energies observed in this decay mode is a broad continuum, as the neutrino carries away the balance of the available decay energy, and is practically undetectable.

In isotopes with even numbers of protons and neutrons, we observe a stronger nuclear binding due to the spin-pairing of nucleons [38], where nucleons of the same type couple their spins antiparallel to form $J = 0$ pairs, lowering the nuclear energy and increasing binding energy relative to odd-mass or odd-odd nuclei. Isotopes with odd numbers of both protons and neutrons are referred to as odd-odd isotopes, and have reduced binding energy due to two unpaired nucleons. For some even-even isotopes, such as ^{136}Xe , ordinary beta decay is energetically forbidden because the neighbouring odd-odd isobar (^{136}Cs) is higher in rest mass. This is illustrated in Figure 1.2 for the $A = 136$ isobars. The difference in mass-energy between the parent isotope and the daughter isotope is known as the Q -value. In beta and double beta decay, the end-point energy of the continuum of measured beta energies is equal to the Q -value, less m_ν . Being a second-order weak process, $2\nu\beta\beta$ decay has an extremely low rate, with half-lives ranging from 10^{18} – 10^{21} yr [39].

There are only 35 known isotopes that match the conditions necessary for double beta




Figure 1.2: Illustration of forbidden and allowed decays for ^{136}Xe and neighbouring isotopes.

decay to happen, i.e. where single beta decay is energetically forbidden, of which only 14 have been measured [40]. A list of double beta decay isotopes relevant for $0\nu\beta\beta$ is shown in Table 1.1, with experimentally relevant parameters. The corresponding current and next-generation experiments and their projected half-life sensitivities are summarized in Table 1.2.

The natural abundance of an isotope is an important consideration in terms of experiment cost and viability, however, it is not completely indicative. We must also include the base material cost, enrichment process considerations, alongside detection technology. For example, with its high Q-value, and the low cost of natural calcium, calcium-48 is an excellent choice, with established technology in calcium fluoride scintillating crystals. However, this is offset by the very low abundance and the difficulty and cost associated with enriching. Similar cost-benefit analysis can be carried out for all the isotopes listed in Table 1.1, including additional

Table 1.1: Summary of selected isotopes considered for $0\nu\beta\beta$ searches, highlighting Q -values, natural isotopic abundance from [41], measured $2\nu\beta\beta$ half-lives from [42], and current best $0\nu\beta\beta$ half-life limits at 90 % C.L.

Isotope	Natural Abundance (%)	$Q_{\beta\beta}$ (MeV)	$2\nu\beta\beta$ half-life (10^{20} yr)	Best $T_{1/2}^{0\nu}$ limit (yr, 90 % C.L.)
^{48}Ca	0.187	4.263	0.53	$> 5.6 \times 10^{22}$ [43]
^{76}Ge	7.8	2.039	18.8	$> 1.9 \times 10^{26}$ [44]
^{100}Mo	9.8	3.035	0.071	$> 1.8 \times 10^{24}$ [45]
^{130}Te	34.08	2.527	7.91	$> 3.5 \times 10^{25}$ [46]
^{136}Xe	8.9	2.459	21.4	$> 3.8 \times 10^{26}$ [47]

Table 1.2: Summary of selected $0\nu\beta\beta$ candidate isotopes with current experiments and next generation experiments and their projected $0\nu\beta\beta$ half-life sensitivities.

Isotope	Current Generation Experiment	Next Generation Experiment	Next-Gen Proj. $T_{1/2}^{0\nu}$ Sens. (yr, 90 % C.L.)
^{76}Ge	LEGEND-200 [44]	LEGEND-1000 [48]	$> 1.6 \times 10^{28}$
^{100}Mo	CUPID-Mo [45]	CUPID-1T [49]	$> 10^{27}$
^{130}Te	CUORE [46]	SNO+ [50] ⁹	$> 1.8 \times 10^{26}$
^{136}Xe	EXO-200 [51]	nEXO [52]	$> 1.35 \times 10^{28}$
	NEXT-100 [53]	NEXT [54]	$> 10^{27}$
	KamLAND-Zen [47]	KamLAND2-Zen [55]	$> 2.0 \times 10^{27}$

considerations on background mitigation techniques, detector technology scalability, nuclear matrix element and phase-space factor magnitudes, and more [56].

1.4.2 $0\nu\beta\beta$ decay process

In the case of $0\nu\beta\beta$, in contrast to standard double beta decay, we have the emission of two beta particles without the corresponding neutrinos. This process would violate the conservation of lepton number and is possible only if neutrinos are Majorana fermions [40]. We can

⁹SNO+ is an existing neutrino observatory at SNOLAB, which in 2025 completed refilling with liquid scintillator and plans to dissolve tellurium salts to study $0\nu\beta\beta$ [57].

compare double beta decay and $0\nu\beta\beta$ as:

$$2\nu\beta^-\beta^- : (A, Z) \rightarrow (A, Z + 2) + 2\beta^- + 2\bar{\nu}_e, \quad (1.13)$$

$$0\nu\beta^-\beta^- : (A, Z) \rightarrow (A, Z + 2) + 2\beta^-, \quad (1.14)$$

where the parent isotope (A, Z) , is transmuted into a daughter isotope 2 units of Z away, with the emission of two betas and two electron anti-neutrinos. As in ordinary beta decay, where both electron and positron emissions are possible, double beta-plus and neutrinoless double-beta plus decay is also considered, as well as double electron-capture and neutrinoless double electron capture [58], [59].

1.4.3 Experimental Signatures of $0\nu\beta\beta$

There are two essential characteristics of $0\nu\beta\beta$ that differentiate it from the ordinary double beta decay mode, and enable detection. During ordinary double beta decay, momentum is shared between the two emitted betas, two anti-neutrinos, and the recoiling nucleus, producing a broad energy spectrum and distribution of emitted particle angles. In $0\nu\beta\beta$ decay at rest, the two electrons carry away the full Q-value, since the recoiling nucleus absorbs momentum while taking up negligible energy. Their relative angle is not fixed by momentum conservation but is set by the weak-interaction dynamics, which mildly favour back-to-back emission [60]. The mean energy of the outgoing electrons is roughly half the Q-value, with a mean opening angle (θ , between electrons) of around 100° – 110° , while the $\cos\theta$ distribution peaks at back-to-back emission. This signature is illustrated in Figure 1.3, with a conceptual sketch of the expected measured beta energy distribution. In the case of $0\nu\beta\beta$, the smoking-gun signature is a small peak-like excess above background signals at the Q-value of the decay. The width of the peak is determined by the resolution of the instrument, typically $\mathcal{O}(\%)$ [39], and with higher resolutions the signal could be unresolvable below the end-point of the much more frequent ordinary double beta decay mode.

A detector that can image particle tracks directly could also take advantage of the opening-




Figure 1.3: Illustration of the experimental signatures of standard and neutrinoless double beta decay. Only electron energies are measured in the detector, resulting in a broad energy continuum for standard double beta decay up to the end-point energy or Q-value. For $0\nu\beta\beta$ we expect a sharp line at the Q-value, broadened by the energy resolution of the instrument. The relative heights of the signals are purely for illustration, and the expected $0\nu\beta\beta$ intensity is greatly exaggerated here.

angle distribution of the two beta tracks alongside the energy measurement to provide additional discrimination power against the ordinary double beta decay mode [60].

1.4.4 Experimental Limits and Interpretations

At the time of writing, $0\nu\beta\beta$ remains undetected, despite decades of experiments [39]. Without a clear signal, experiments searching for $0\nu\beta\beta$ place limits on the minimum half-life of $0\nu\beta\beta$, inferred from non-observation of the signal. The current world-leading limit is from the KamLAND-Zen experiment in Japan, which uses a liquid scintillator detector with

a low-background interior region loaded with ~ 785 kg of ^{136}Xe . The current limit is $T_{1/2}^{0\nu} > 3.8 \times 10^{26}$ yr at a 90 % C.L, from their full-dataset [47].

In order to compare experiments and isotopes, the standard convention is to convert an experimental limit on the half-life to the effective Majorana mass, $m_{\beta\beta}$. The simplest model that is widely used to perform this calculation is called “light Majorana neutrino exchange” [7]. In the case of light neutrino exchange, we assume Majorana neutrinos and that no other lepton-number violating processes contribute to the effective neutrino mass [40]. This process is illustrated in Figure 1.4, reproduced from [61].




Figure 1.4: Feynman diagram of $0\nu\beta\beta$ through light Majorana neutrino exchange, figure reproduced from [61].

The effective Majorana mass, which enters the $0\nu\beta\beta$ half-life expression, is a coherent sum of the previously mentioned PMNS matrix, $U_{e,i}$, with the neutrino mass eigenstates, m_i [40],

$$\langle m_{\beta\beta} \rangle = \left| \sum_i^3 U_{e,i}^2 m_i \right|, \quad (1.15)$$

where, because the $U_{e,i}$ terms are squared amplitudes rather than the $|U_{e,i}|$ in neutrino oscillation, interference between terms is allowed, and a vanishing $m_{\beta\beta}$ is possible. While many other theories describing $0\nu\beta\beta$ are possible, it has been postulated that any process creating $0\nu\beta\beta$ necessarily implies Majorana neutrinos, known as the “black-box” or Schechter-Valle

theorem [62]. Inversely, any observation of $0\nu\beta\beta$, necessarily implies Majorana neutrinos, regardless of the underlying mechanisms, and correspondingly, lepton number violation.

The half-life for $0\nu\beta\beta$, under light neutrino exchange can be expressed as:

$$(T_{1/2}^{0\nu})^{-1} = \langle m_{\beta\beta} \rangle^2 G^{0\nu} |\mathcal{M}^{0\nu}|^2, \quad (1.16)$$

where, $G^{0\nu}$ is the phase-space factor that depends on the Z transition and the Q-value, calculated analytically with high precision [63]. $\mathcal{M}^{0\nu}$ is the nuclear matrix element (NME) and quantifies the overlap between initial and final nuclear states that represents the many-body nuclear physics involved $0\nu\beta\beta$ [39]. The NME is not accessible in experiments and must be calculated using nuclear modelling frameworks, which have substantial disagreements depending on model and method [39]. This represents the largest source of variance between limits on $m_{\beta\beta}$ from experimental half-life limits [39], [52].

Experimentally, the half-life enters into the equation through an estimate of the number of decays expected in a detector, $N_{0\nu\beta\beta}$. If we assume the exposure time is much shorter than the decay lifetime, so that the number of decays grows linearly with t , the number of expected decays is:

$$N_{0\nu\beta\beta} = \ln(2) \frac{\epsilon m N_A}{M_{\beta\beta}} \frac{t}{T_{1/2}^{0\nu}}, \quad (1.17)$$

where, t is the exposure time, ϵ is the enrichment fraction (what % of mass is composed of the $\beta\beta$ isotope), m is the sensitive mass being exposed, N_A is Avogadro's constant, and $M_{\beta\beta}$ is the molar mass of the $\beta\beta$ isotope. We can compare the current best limit on effective Majorana mass, and available phase-space for $0\nu\beta\beta$ as shown in Figure 1.5, known colloquially as the ‘‘Lobster’’ plot. The effective Majorana mass is plotted against the lightest neutrino mass (m_1 in NO, m_3 in IO), and the bands are computed using measured PMNS matrix elements with the currently unconstrained Majorana phases, as the lightest neutrino mass is varied in either hierarchy. The width of the KamLAND-Zen limit band represents the combined uncertainty from the experiment and NME model variances [47]. Uncertainties on individual PMNS

matrix elements are omitted, but are typically at the level of a few percent for the mixing angles [28].




Figure 1.5: Lobster plot showing available phase-space for $0\nu\beta\beta$ in effective Majorana mass vs the lightest neutrino mass. The inverted mass ordering (IO) is shown as a purple band, and normal mass ordering (NO) in red. Overlaid are the current best limits for $m_{\beta\beta}$ from KamLAND-Zen [47], and the best limit for m_β from KATRIN [34] for comparison. The width of the bands comes from the unknown Majorana phases (Equation 1.4). Note, m_β cannot be directly translated to m_{lightest} without knowledge of the mass orderings, and is overlaid purely to compare $0\nu\beta\beta$ constraints with the best limit from direct neutrino mass measurements.

Future $0\nu\beta\beta$ experiments will push the sensitivity to $m_{\beta\beta}$ another order of magnitude, aiming to cover the entire IO phase space indicated in Figure 1.5, down to $\sim 15\text{meV}$ [52], [64], [65]. This will require innovation in signal extraction techniques, advances in background reduction, and most significantly, much more isotope exposure. In the following section, we will review these details and the varied considerations necessary to achieve this order of magnitude increase in sensitivity to $0\nu\beta\beta$.

1.5 Experimental Considerations for $0\nu\beta\beta$ Searches

When searching for $0\nu\beta\beta$, we must be aware of the key limitations in obtaining our desired experimental signal. In Section 1.4.3, we introduced the key parameters of interest for identifying $0\nu\beta\beta$ from conventional $2\nu\beta\beta$. We will address the two key limitations, the extremely low rate of events, and the abundance of diverse background signals.

To put this into our context, with a sensitive mass of 3.2 tonnes of xenon, enriched to 90 % in ^{136}Xe , we should expect ~ 1 event per year of live exposure at a half-life of 10^{28} yr. At the same time, we should expect a continuous ~ 10 mHz rate of ordinary double beta decay events.

1.5.1 Background Signals

Background events can be divided by origin into internal and external backgrounds. Internal backgrounds are intrinsic to the detector itself, with ordinary two-neutrino double beta decay being the most important example for $0\nu\beta\beta$ searches. External backgrounds originate outside the detector volume but produce signals within.

The dominant sources of external backgrounds are primordial uranium and thorium decay chains, present at varying concentrations in most materials, along with radiogenic isotopes such as ^{40}K and ^{60}Co commonly found in detector components [56]. Of particular concern is ^{214}Bi , which has one of the highest energy beta-decays in the uranium chain at a Q-value of 3.27 MeV, overlapping with the Q-values of most $0\nu\beta\beta$ candidate isotopes (Table 1.1). Radon contributes to both categories: as an external background through decay-daughter deposition onto detector materials during fabrication, and as an internal background when carried into detector volumes through gas handling systems.

Cosmogenic backgrounds arise from cosmic-ray muons and their secondaries. On Earth's surface, muon fluxes are untenable for rare-event searches, and so experiments are operated in deep underground laboratories where muon rates are suppressed by $\mathcal{O}(10^6)$ [66]. Residual

high-energy muons can still activate detector materials through spallation, producing long-lived radioactive isotopes. Additional external backgrounds include gammas from surrounding rock and detector infrastructure [67], and neutrons from spontaneous fission and (α, n) reactions in the cavern walls, which produce both prompt capture signals and long-lived activated backgrounds.

1.5.2 Background Mitigation

Background mitigation techniques are divided into passive and active categories. Passive mitigation prevents background signals from entering the detector or signal region, while active mitigation uses supplementary information to discriminate against background events.

Passive mitigation is primarily achieved through radiopurity screening, where every detector component must satisfy strict radioactivity requirements, often at the part-per-billion level and beyond, using techniques including surface and underground HPGe counting, neutron activation analysis, mass spectrometry, and coincidence gamma counting [68], [69]. Shielding with high-Z materials such as lead attenuates gamma radiation, while hydrogen-rich materials such as water or polyethylene moderate neutrons [40]. Radon emanation screening and active radon removal suppress radon and its progeny [70]. In $0\nu\beta\beta$ searches, energy resolution itself acts as a passive mitigation tool: sharper resolution reduces overlap between the signal peak and the tail of the two-neutrino double beta decay spectrum, enabling “cut-and-count” analyses over a narrow signal window [39].

Active mitigation techniques exploit additional detector information to veto or reject background events. Passive water shields can be instrumented as muon Cherenkov vetoes, flagging events coincident with through-going muons [71]. Position reconstruction enables fiducialization, removing surface-proximate events where backgrounds are elevated. Particle identification exploits differing signal properties between event types; in liquid xenon, the ionization-to-scintillation ratio differs between signal-like β and background-like α events

[72], and topological reconstruction or inter-event timing coincidences provide further handles [40], [60].

A uniquely powerful active technique available in xenon-based $0\nu\beta\beta$ searches is barium tagging, where following a double-beta decay event the barium daughter ion is individually identified within the detector volume. High-specificity tagging provides a direct signature of double-beta decay, rejecting all non-double-beta backgrounds at the single-event level [73], [74]. This approach is enabled by the homogeneous liquid or gaseous xenon target, which permits extraction and identification of the stable barium daughter ion. This capability is unavailable in solid-state detectors where the daughter isotope remains embedded in the matrix.

In this chapter, we have introduced the historical and theoretical motivation for $0\nu\beta\beta$, outlined the characteristics of the expected signal, and explored the experimental considerations for building a rare-event search detector. $0\nu\beta\beta$ remains one of the priority experimental searches for physics beyond the standard model, highly motivated by the non-zero mass of neutrinos confirmed in oscillation experiments [75]. If discovered, $0\nu\beta\beta$ is an unambiguous indication of lepton number violation, and would be the first observed process where matter is created in excess of anti-matter, pointing towards a potential solution to the observed baryon asymmetry problem [40].

In the following part, we will transition to detector characterization work in the context of a next-generation xenon-based $0\nu\beta\beta$ search experiment, such as nEXO. This work is split into three chapters, with the first introducing our prototype detector at McGill University and upgrades carried out as part of this thesis work. The second chapter will cover a novel measurement technique and first detailed characterization of secondary light emission from solid-state photosensors in liquid xenon. In the third chapter, we will review commissioning measurements carried out with our upgraded experiment and outline future work.

Part II

**Characterizing SiPM Performance in
Liquid Xenon**

2

The Light Only Liquid Xenon Experiment

2.1 Motivation

The Light Only Liquid Xenon experiment (LoLX) is a small-scale modular liquid xenon (LXe) detector located at McGill University. The LoLX experiment has three main physics goals; developing a detailed understanding of SiPM performance in LXe (1), study and characterize LXe light emission properties (2), and improve the modelling of light transport and detector response in LXe detectors (3). In this chapter, we will give a detailed introduction to xenon scintillation, SiPM operation and their key performance metrics, and considerations for light

transport within a LXe detector. This information will be helpful for the subsequent overview of the LoLX detector, its design considerations, and supporting infrastructure. We will then conclude with an introduction to the analysis and simulation tools used for the LoLX experiment, developed as a part of this work.

2.1.1 Xenon Scintillation

A material that produces light proportionally to the energy deposited into it by the passage of ionizing radiation is known as a “scintillator” and the light produced is called “scintillation”. Scintillation is a subclass of luminescence, the phenomenon of characteristic light emission of a material through the absorption of energy. All noble gases exhibit scintillation, and liquid noble gases are used in many applications as scintillators for calorimetry [76], positron emission tomography (PET)[77], [78], and tracking detectors [79]. In its liquid state, xenon scintillation light has a mean emission wavelength of ~ 175 nm, in the vacuum ultraviolet (VUV) range of light [80]. The high energy of the VUV photons restricts the choice of compatible detector support materials significantly because of strong attenuation and very low reflectivities (for example in air, with the linear attenuation coefficient $\mu = 10^{-2} \text{ cm}^{-1}$). Fortunately, LXe is transparent to its own scintillation light, thanks to the molecular processes involved in the scintillation light production.

When ionizing radiation passes through LXe, collision and ionization occur with xenon atoms. The relative fraction of ionization to collisional excitation varies significantly by the linear energy transfer (LET) of the incident particle [81]. Xenon scintillation is produced by the decay of xenon excimers, molecules bound only in their excited state that dissociate back into two free atoms as they emit a photon. These excimers (Xe_2^*) can be formed through both excitation and ionization interactions, producing varying populations in either singlet or triplet meta-stable states [82]. An example of xenon scintillation and the pathways for xenon excimer formation are illustrated in Figure 2.1.

The overall LXe scintillation emission time profile depends strongly on the applied electric field, influencing the relative populations of singlet and triplet excimer states [83], through the recombination channel.




Figure 2.1: Illustration of liquid xenon scintillation production and decay pathways. Incident radiation ionizes and excites xenon atoms which produces direct excimers, or recombination excimers that decay back to ground state xenon, emitting vacuum ultra-violet (VUV) light. The two decay channels (singlet/triplet spin configurations) have different half-lives, corresponding to meta-stable states of the xenon excimer molecule. The measured scintillation time profiles will also depend on applied electric fields by drifting away electrons that would have combined and populated the recombination branch.

A key feature of LXe detectors is the anti-correlation between charge and light signal fluctuations [84]. When an electric field is applied across the LXe volume, ionization electrons can be drifted to charge collection detectors, while scintillation light from direct excitation and recombination is still collected. By measuring both the charge and the light signal simultaneously, energy resolutions as good as $\sim 1 \%$ (σ/μ) are achievable due to the suppressed

variance from recombination effects [85].

The Noble Element Simulation Technique or “NEST” is a tool used by the LXe detector community for calculations of light and charge yields from interactions with LXe [82]. The NEST model is empirical and informed by a large variety of LXe detector measurements across a range of scales and applied electric fields. At zero-field there are a limited number of measurements [86], leading to larger uncertainties in the estimated light yield and its resolution. The NEST software is available for direct use, integration into particle transport simulations like Geant4, and through a python library binding, `nestpy`. The estimated light yields at zero-field from NEST model 2.0 through `nestpy`, are shown in Figure 2.2, which is approximated in the NEST model using 10 V/cm. Figure 2.2 includes the response for beta-like (including gamma), nuclear recoils, and alpha particle radiation. The range of the nuclear recoil response is limited to the range of accessible nuclear recoil energies from calibration neutron sources.

Noble liquid scintillators are exceptionally bright, with light production yields of ~ 40 ph per keV of energy deposited. This light level is much brighter than plastic scintillators at ~ 1 -10 ph/keV [87], and nearly as bright as the best crystal scintillators, such as thallium doped sodium iodide (NaI:Tl) at ~ 60 ph/keV [88]. Noble liquids also have $\mathcal{O}(\text{ns})$ scintillation response times [89], appropriate for fast timing applications, and are flexible in application owing to its liquid state.




Figure 2.2: LXe scintillation yields, in photons per keV of deposited energy of the incident radiation, versus incident particle kinetic energy. Data is shown from NEST at an applied electric field of 10 V/cm (the approximation for zero-field in the model), using nestpy (2.0). Three different LET regimes are shown through different incident radiation types, β , recoiling nuclei (NR or nuclear recoil), and α .

Not all energy deposited into LXe will produce a measurable signal, and there are several pathways of energy loss that are possible through differing mechanisms. Some of these losses are avoidable, while others are not, and the main mechanisms are illustrated in Figure 2.3. Nuclear quenching, also known as Lindhard quenching [90], is the energy lost during primary collisions of the recoiling xenon nuclei with nearby xenon nuclei, heating the xenon. Biexcitonic quenching, where excited atoms in a dense track region collide with one another, resulting in ionization and a net loss of one quanta (a quanta is either an ionization electron or scintillation photon here) for every collision of two excited atoms. The amount of biexcitonic quenching depends on the LET of the primary particle and structure of its track through the material. While biexcitonic and nuclear quenching are unavoidable loss mechanisms, the remaining mechanisms outlined in Figure 2.3 are dependent on the character and quantity of non-xenon impurities in the liquid. Electronegative impurities are contaminants that exhibit a high electron affinity, or electronegativity, and act as trapping sites for ionization electrons.

Their presence leads to more variance in recombination fluctuations as impurity levels vary, and to a loss of measurable charge and light signal. The two main electronegative impurities found within LXe detectors are water and oxygen, commonly out-gassed from materials within the detector volume [91]. Photoabsorption impurities absorb scintillation photons while they are travelling through the medium, and the impact of this channel depends on concentration and detector geometry. Water and oxygen will also act as photoabsorption impurities. Penning quenching, where an excited xenon atom ionizes an impurity molecule, and collisional dissociation, where molecular impurities are broken apart by excited xenon dimers, both require energy level compatibility with the impurity [92]. This criterion is not met for any other noble gases, which are commonly found in xenon, while it is possible for molecular gases such as CO_2 [93]. These pathways are essential to understand for diagnosing LXe detector performance, with each pathway having its own energy, particle, and concentration dependence. Excellent xenon purity is required to reduce the avoidable energy loss channels, and will constrain material selection during detector design.

Energy Loss Mechanisms in LXe




Figure 2.3: Illustration of energy loss pathways in LXe. Where nuclear and biexcitonic quenching represent unavoidable losses, all other energy losses are correlated with the xenon purity level. Energy loss due to material heating is represented by the fire icon.

Operationally, LXe detectors require continuous purification to remain a high quality scintillation media. This is typically achieved through the use of a heated purifier or “getter”, which employs a large surface area material with irreversible chemical bindings to most chemical impurities, such as zirconium [94].

To monitor the detector purity level, electron lifetime measurements are routinely carried out in LXe TPCs. The electron lifetime is very sensitive to electronegative impurities, where electron loss from attachment on impurities attenuates the signal. High lifetimes of up to $\mathcal{O}(10 \text{ ms})$ are attainable with continuous purification [95]. One draw-back of this technique is that the lifetime measurement cannot separate specific contaminants and is difficult to translate to concentrations of impurities [96]. However, it is a representative measurement

that is directly correlated to detector light and charge yields, and can be compared between experiments.

In addition to scintillation light, Cherenkov radiation can be produced in a LXe detector from both β and γ interactions (γ 's scatter or photoabsorb and produce electrons). Cherenkov radiation is the optical equivalent of a shock-wave, and is produced when a charged particle is travelling faster than the speed of light in a dielectric medium ($v > c/n_{media}$). The passage of the charged particle excites the material and creates a polarization gradient as it travels, exciting the material which subsequently relaxes and releases this energy as light, known as “Cherenkov Radiation”, named after its discoverer Pavel Cherenkov [97].

Cherenkov emission is characteristically prompt, emitted very quickly during the passage of the incident charged particle, broad-band in energy with a peak at high energies, and produced in a conical wavefront. An illustration of the Cherenkov cone is shown in Figure 2.4. The opening angle of the cone is related to the refractive index of the medium, n , and the velocity of the charged particle (represented by β the dimensionless velocity factor), is given as, $\cos \theta = \frac{1}{n\beta}$. The full description of Cherenkov radiation was given by physicists Ilya Frank and Igor Tamm, and is represented by the eponymous relation, the “Frank-Tamm Formula” (in wavelength units) [98], [99]:

$$\frac{d^2N}{dx d\lambda} = \frac{2\pi\alpha}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)} \right). \quad (2.1)$$




Figure 2.4: Illustration of a Cherenkov cone from incident particle moving with relativistic velocity β in a dielectric medium.

2.1.2 Silicon Photomultipliers

The light sensors employed in LoLX and planned for the nEXO experiments are silicon photomultipliers (SiPMs). SiPMs are single photon sensitive light detectors, used in many fields of applications [100]. They have many advantages over the traditional technology for single-photon counting, photomultiplier tubes (PMTs). In the context of our cryogenic low background experiment beneficial features of SiPMs include: their intrinsically low concentrations of primordial uranium and thorium, high detection efficiencies, and low power requirements [101]. The operational principle of SiPMs is the avalanche, or Geiger, process, where a single photon can liberate $\mathcal{O}(10^6)$ charge carriers. Avalanches occur within the thousands of individual diodes that make up the SiPM, known as single-avalanche photo-diodes (SPADs).

Each SPAD operates as a binary counter, and has a characteristic reset time, which can be understood as an RC time constant of the diode circuit with the SPAD as a capacitor, alongside the quenching resistor. Individual SPADs are P-N junctions with a reverse bias voltage applied until a breakdown condition is reached. This point is aptly called the “breakdown voltage”, and is illustrated in Figure 2.5. The strength of the avalanche signal is directly proportional to the applied voltage across the device, up to the saturation voltage. Once the breakdown voltage has been reached, the sensor no longer responds linearly to incident light, and a single photon is sufficient to create an avalanche of charge carriers.

The sensitive region for photon detection is called the depletion region, at the interface between the n-doped and p-doped layers of the SPAD. The location of the depletion region is a key consideration in SPAD design for different regimes of photon wavelengths, as penetration depths vary significantly [102]. This operational mode is called “Geiger” mode, and the applied voltage in excess of the breakdown voltage is known as “overvoltage”. Once Geiger mode has been reached, the SiPM operates as a proportional counter, where a fraction of the incident light on the device will trigger avalanches in individual SPADs, and the number

of incident photons may be determined by estimating the number of fired SPADs from the resulting signal. SiPMs are designed such that each SPAD has the same gain, and so by measuring the current of the SiPM signal, we may “count” the number of SPADs that have fired. This allows us to perform a proportional measurement of the incident light on the sensor. The estimated number of fired SPADs is often referred to as N_{av} , or sometimes N_{PE} for anachronous reasons ¹.




Figure 2.5: Illustration of the SiPM operational principle. As photodiodes, SiPMs respond linearly to light when biased with low applied voltage, gain increases significantly as the bias voltage increases until the breakdown voltage is reached.

The responsivity of SiPMs to incident light is characterized by the “Photon detection efficiency” (PDE). The PDE is quantified as the probability of the SiPM to trigger an avalanche when a photon of a given wavelength is incident upon it. The PDE of a SiPM can be described functionally as depending on several key factors, while more detailed parameterizations are

¹PE or Photoelectrons, are the quanta counted by PMTs, where an incident photon on the photocathode of the PMT liberates a photoelectron, triggering the dynode avalanche process. It is common for this jargon to be applied to SiPMs, which as solid-state sensors, can have avalanches triggered by holes or electrons in the silicon lattice.

possible [103]. A common approach is to break the PDE into transmission effects, geometric effects, and intrinsic responsivity:

$$PDE(\lambda, \theta, t, V, n) = T(n, t, \theta) \times FF \times P_{av}(\lambda, V, t), \quad (2.2)$$

where the first of three main contributions is the transportation probability, T , which depends on the refractive index of the external media, n , the temperature t , and the incident angle θ . The second term FF is the geometric fill factor, related to the SPAD packing efficiency within the SiPM. The third term is the avalanche initiation probability, P_{av} , which incorporates the quantum efficiency, depending on the incident wavelength of light λ , the applied electric field strength, proportional to reverse bias voltage V , and the operational temperature t .

A key consideration when employing SiPMs as light sensors is the presence of correlated noise. Correlated SiPM noise is a source of excess signal, that is proportional to the number of triggered SPADs, but does not originate from initial photons. There are several classes of correlated SiPM noise, and we will discuss the four main types as illustrated in Figure 2.6.




Figure 2.6: Illustration of SiPM correlated noise pathways with expected SiPM waveform readout signals. External crosstalk can be considered a subclass of optical crosstalk, where the correlated hit SPAD is in a SiPM in a different device.

- (a) Afterpulsing (AP) - Charge carriers liberated in the primary avalanche (blue) are trapped by defects and impurities within the silicon lattice, and released with characteristic de-

cay times with lifetimes $\mathcal{O}(1\text{-}100\text{ ns})$ [102]. These subsequent correlated avalanches occur within the same SPAD, and often before the SPAD has fully recovered from the primary avalanche. This leads to a correlation between the time since primary avalanche (Δ_t), and the amount of charge contributed by the afterpulse. As Δ_t approaches the SPAD reset time, the excess charge contributed by the afterpulse tends towards the single avalanche charge.

- (b) Direct Optical Cross-talk (DiCT) - Optical photons created in the avalanche transmit through the bulk silicon to a neighbouring SPAD, inducing an additional avalanche. These photons have been shown to be in the near-infra-red (NIR) regime [102]. DiCT will predominantly be localized to the neighbouring SPADs, with a photon transit time of $\mathcal{O}(\text{ps})$, and appears as two SPADs firing in coincidence.
- (c) Delayed Optical Cross-talk (DeCT) - Differing slightly from DiCT, DeCT is produced by charge carriers liberated near the depletion region of the SiPM which randomly walk back into the depletion region, triggering an avalanche [104]. The characteristic of this correlated noise source is a primary avalanche, with a delayed single avalanche equivalent signal. This differs from AP, which has a delayed avalanche within the same SPAD, resulting in a less-than-single avalanche signal. The delay time structure of DeCT is defined by diffusion of the charge carriers within a sensitive region around the depletion region.
- (d) External Cross-talk (ExCT) - In a small fraction of cases, the cross-talk photons can leave the device entirely and be detected by a distant SiPM [105]. This source of noise is of particular importance for detector designs employing large surface coverage, with SiPMs facing one-another. This is characterized by a primary avalanche on one SiPM, with a time coincident signal present on another sensor. Developing an understanding of SiPM ExCT in a LXe detector is the subject of the subsequent chapter (Chapter 3),

and will be discussed in detail therein.

In addition to correlated noise, SiPMs experience uncorrelated noise, known as “dark noise” or “dark current”. They typically exhibit very large rates of dark noise, up to MHz/mm² at room temperature [106], and can be a limiting factor for SiPM operation in some applications (low energy threshold experiments for example). Dark noise in SiPMs is largely driven by thermally generated charge carriers, as such, a significant Boltzmann suppression occurs as the devices are cooled and operated at lower temperatures. Below 100 K, thermal dark noise becomes secondary to quantum tunnelling noise [107]. At LXe operating temperatures of ~ 170 K, the dark noise is $\mathcal{O}(\text{Hz/mm}^2)$ [101], and is a negligible consideration for many analyses.

Photomultiplier Tubes

When light is incident on a photomultiplier tube (PMT), electrons are ejected from the thin photocathode through the photoelectric effect, and subsequently amplified by successive dynode impacts within the evacuated tube to produce an avalanche of $\mathcal{O}(10^6)$ electrons. In the context of LXe detectors for rare-event searches, PMTs carry two main drawbacks relative to SiPMs: a higher baseline radioactivity from primordial uranium and thorium within the PMT glass [108], and higher voltage and power requirements leading to increased heat dissipation in the cryogenic fluid. Both PMTs and SiPMs have comparable efficiencies for LXe scintillation light ($\sim 20\%$), fast intrinsic timing resolution, and signal linearity, though SiPMs offer superior single-photon resolution at low light levels while PMTs maintain linearity over a larger dynamic range. In LoLX phase-2, a VUV-sensitive PMT is included alongside the SiPM arrays to provide an independent cross-check of the light detection and a broader dynamic range for calibration data-taking.

2.2 LoLX Support Infrastructure

Operating a liquid xenon detector requires significant supporting apparatuses, equipment and monitoring capabilities. In particular, for the LoLX experiment, we have two additional requirements: running autonomously, and running remotely. The two main categories of supporting infrastructure for the LoLX experiment are the gas handling and recovery system, and the cryogenics and detector monitoring system.




Figure 2.7: Photo of the LoLX experiment at McGill University from November 2024. The LoLX gas handling system valve board is visible on the left, and the cryostat and DAQ are on the right.

2.2.1 Gas Handling and Recovery System

The LoLX gas handling and recovery system, referred to internally as the “XGHS”, is shown on the left side of Figure 2.7. The XGHS is used to transfer gaseous xenon from the storage

bottles, installed behind the valve board, to the LoLX cryostat during operations, with flow control and gas-phase purification, and to ultimately recover the xenon in the storage bottles at the end of a LXe campaign.

The XGHS is controlled by a custom LabView program, developed by former McGill Postdoctoral Fellow Chris Chambers, who led the design of the XGHS. The LabView control program operates the pneumatic valves for automated and remotely controllable recovery of xenon. While recovery is automated and remote controllable, xenon deployment requires in-person monitoring and valve control. Deployments require pre-heating of the purifier and opening of the manual bottle valves. A simplified illustration of the XGHS Piping and Instrumentation Diagram (P&ID), with key components highlighted for a conceptual understanding, is shown in Figure 2.8. This illustration excludes most monitoring sensors and valve control. Highlighted in purple is the flow of xenon into LoLX during liquefaction, with recovery flow highlighted in orange. All xenon gas flow is controlled by thermal gradients in the gas system, using a technique known as “cryopumping”. This method is simple and ensures no contamination in the GXe from oils or other contaminants commonly found in commercial pumps. There are two main storage bottles employed in the XGHS, and one emergency recovery bottle to provide additional capture capacity in the event of a critical system failure during operation.




Figure 2.8: Illustration of xenon flow into LoLX in a simplified version of the LoLX XGHS P and ID. Gas flow during deployment (purple) is a single pass-through of the purifier, no recirculation is currently employed. During recovery, gas flow (orange) is directed to the storage bottles which are cryopumped by a dewar of LN_2 as shown.

The key XGHS components, shown in Figure 2.8, are:

- Heated Purifier - Heated zirconium-alloy getter purifier. Removes non-noble gas impurities through chemical bonds and large surface area (Model MonoTorr PS3-MT3).
- Cold Purifier - In-line purifier. Removes residual gasses after heated getter, also acts as a particulate filter (Model SAES 902).
- Mass Flow Controller (MFC) - Controls and monitors xenon gas flow. Typically up to ~ 5 standard litres per minute (slpm) (Model MKS GM50A).
- Residual Gas Analyzer (RGA) - Vacuum compatible probe with electronics control unit. Analyzes gas composition through mass spectroscopy using a quadrupole probe and ion detector (Model SRS RGA200).

2.2.2 Cryogenics System and Upgrades

To condense xenon into LXe at near atmospheric pressures (0.8-1.2 atm), the stable temperature range is 166-172 K, requiring precise control over thermal conditions. This is clearly illustrated in the xenon phase diagram Figure 2.9, from [109], with optimal conditions away from the triple point of ~ 161 K at ~ 82 kPa. An important consideration when operating the LXe system is the vapour pressure, which determines when condensation will begin. The LoLX cryostat operates by condensing GXe into LXe on the interior walls of the cryostat, cooled by the cryogenics system. In order to reach stable conditions within the LXe cell, the cryogenic system must provide stable cooling power in steady-state conditions, in addition to sufficient power during condensation.




Figure 2.9: Phase diagram of xenon from [109]. LoLX operates near 1 atm of pressure during LXe campaigns.

During condensation, heat is expelled due to the phase transition. This is the inverse of the enthalpy of vaporization, called the enthalpy of condensation. For xenon, the enthalpy of condensation is 12.6 kJ/mol (at boiling point or 165 K) [110]. More conveniently we can write

this as ~ 96 kJ/kg of xenon. In LoLX, we operate with approximately 5 kg of xenon, meaning we require a ~ 500 kJ condensation fee, in addition to the steady-state heat load of the system. Each filling cycle takes on average ~ 6 hours to complete, translating to an additional ~ 22 W of cooling power required for a filling cycle. This is an important consideration in the design of the upgraded system, which also requires an estimated ~ 20 W of steady-state cooling power. The estimated steady-state cooling power comes from early estimates of the heat load in LoLX during initial commissioning, from simulations, and conservative estimations.




Figure 2.10: LXe vapour pressure vs temperature. Stable liquid states are above the curve, while stable vapour states are below.

The earlier iteration of the LoLX cooling system had several shortcomings which ultimately lead to its failure. The previous cooling system employed an open-loop LN₂ circuit, which was controlled by a custom program using a solenoid valve that ran with a duty cycle tuned to the desired cryostat settings. The LN₂ flowed from the supply dewar through the solenoid valve and through a copper block at the base of the cryostat. The block had a serpen-

tine path to increase surface area, and the LN_2 was expected to vaporize completely in this block, extracting the heat from LoLX, to vent out of the cryostat as GN_2 .

Two main issues in this system prevented stable operating conditions, leading to its eventual failure and replacement. The first critical issue was the variation of LN_2 flow rate as the supply dewar was drained. Using a duty-cycle control system, the system's cooling rate decreased as the supply decreased. This was partially mitigated by applying a constant pressure to the supply dewar with an external GN_2 bottle and regulator. However this introduced additional challenges. By pressurizing with GN_2 , when the supply dewar was empty and needed to be replaced, there would be a brief period of cooling power loss, and the system instability was exacerbated by the introduction of warm GN_2 . The cooling power loss driven by supply dewar replacements created instability in the cryostat environment, sometimes leading to LXe freezing or rapid boiling. Each dewar replacement would introduce several hours of dead time for data taking.

The second major shortcoming in the earlier iteration of the cooling system was a severe underestimation of the heat transfer coefficient between the serpentine block and LN_2 . The expectation that all LN_2 would boil in the block was never met, and at high LN_2 flow rates there was significant liquid flow out of the gas venting port. It is likely that film boiling [111], driven by the large temperature difference between LN_2 at 77 K and the serpentine block at ~ 180 K, was the culprit. The non-evaporated LN_2 introduced significant thermal strain to the welded gas ports through the insulation vacuum into the room temperature exhaust line. This was mitigated moderately through insulation and the addition of a heat exchanger at the exhaust to warm up the cold exiting GN_2 using the warm water supply from the building. Eventually, the stress was too much and the connection snapped off on the room side, and a cryogenic redesign was initiated. While the cooling system challenges necessitated significant work, the redesign was also motivated by the need to perform longer duration LXe campaigns for extended data taking periods.



Figure 2.11: MIDAS Slow control temperature readout of the LoLX phase-1 October 2021 LXe run. Overlaid highlights show pre-run cooldown, LXe filling and data-taking periods. Extensive instability is clear during the filling cycle, with additional spikes during LN₂ re-fill cycles, shown as dead-time windows.

During the October 2021 LXe campaign, shown in MIDAS readout in Figure 2.11, ~ 40 hrs was the duration of pre-run cooldown, filling and re-stabilization, which resulted in ~ 24 hrs of stable conditions for data-taking, with an additional ~ 20 hrs of dead-time. Ignoring the pre-run duration, this results in $\sim 55\%$ up-time.

The upgraded cooling system replaced the open loop LN₂ and serpentine block with a SHI CH-104 cryocooler and helium compressor. The SHI CH-104 cryocooler is a single-stage Gifford-McMahon cryocooler, designed for LN₂ operations. The cryocooler is specified for a 42 W heat load at 77 K and 60 Hz. During operation in LoLX, the cold head is stable at ~ 110 K, with a heat load capacity of ~ 60 W. The upgraded cooling system was outfitted with 6×100 W cartridge heaters, and 4 temperature sensors (PT-110, Resistive Temperature Detector or RTD) for monitoring. The cartridge heaters and RTDs are read out using a Lakeshore model 350 cryogenic temperature controller. The cryostat is now controlled using the built-in

Proportional-Integral-Derivative (PID) functions of the Lakeshore controller, configured by a custom control program over the ethernet interface. An illustration of the upgraded cryogenics and controls is available in the upcoming Section 2.4.




Figure 2.12: Temperature and pressure conditions of the LoLX phase-2 cryostat with upgraded cooling system during January 2025 commissioning run. The inset plot shows the $\mathcal{O}(\text{mK})$ stability on the PID control sensor over a 24 h period of data-taking.

We can compare the last LoLX phase-1 run in Figure 2.11 with the most recent LXe run with the upgraded cryostat shown in Figure 2.12. While operation of the upgraded cryostat is vastly improved over the previous version, the operating conditions and pre-run cooldown have also significantly improved. Comparing this to our last LXe run with the previous cryostat, we have $\sim 100\%$ up-time for data-taking once pre-run filling and stabilization is achieved. Taking ~ 24 hrs for initial cooldown and filling compared to ~ 40 h from the last phase-1 run.

2.3 LoLX Detectors

2.3.1 Phase-1 Detector

During the first phase of LoLX operations from 2019-2022, the first LoLX phase detector was operated in the LXe cryostat. The phase-1 detector shown in Figure 2.13 during a test assembly, was fabricated out of a 3D printed resin cage with 24 Hamamatsu Photonics K.K (HPK) VUV4-Q SiPM packages, for 96 individual SiPMs. The majority of the SiPM surface area in LoLX phase-1 was covered in optical filters, with 22/24 packages covered by a long-pass (LP) optical filter with a cut-off of 220 nm. One of the remaining packages was covered by a bandpass (BP) optical filter with a band centered on LXe scintillation at ~ 175 nm. The last remaining packages was left unfiltered and is referred to as the “bare” SiPM package. A ^{90}Sr needle source (β) was installed from the bottom of the detector pointing into the middle region and used as the calibration source.




Figure 2.13: The LoLX Phase-1 detector during test assembly, photo by Thomas McElroy.

The decay pathway of ^{90}Sr is shown in Figure 2.14, ^{90}Sr decays to ^{90}Y , which has a rela-

tively short half-life of 64 hours compared to the 28-year half-life of ^{90}Sr . As the half-life of the parent isotope is much longer than the progeny (^{90}Y), we can safely assume the source has reached secular equilibrium. ^{90}Sr was selected because of its simple decay path which covers the range of β energies of interest for a $0\nu\beta\beta$ search in LXe, from a few hundred keV to ~ 1 MeV. LoLX phase-1 also included a UV optical fibre, with an exposed end inserted into the top of the cage. The UV optical fibre is connected through the top of the Conflat feedthrough to a PicoQuant-PDL 800-D laser with a pulsed diode head ($\lambda \sim 405$ nm).




Figure 2.14: Decay pathway for ^{90}Sr from [112].

The main physics goals of the LoLX phase-1 experiment can be summarized as follows:

- Characterize and improve the understanding of SiPM performance in a LXe environment.
- Validate photon transport simulations through data-simulation comparisons.
- Investigate the potential in separating Cherenkov and scintillation light in a LXe detector, towards improved background discrimination for future $0\nu\beta\beta$ LXe detectors, among other applications [113].

The investigation into LXe scintillation and Cherenkov light separation motivated the detector design. By covering most of the photosensor surface area with optical filters to select

for > 220 nm light we isolate this contribution to specific channels. This increases our target area for the relatively weak Cherenkov emission, removing most of the dominant VUV emission from LXe. In LoLX phase-1 events, the expected ratio of LXe scintillation to Cherenkov was ~ 150 to 1. The optical properties of the BP and LP filters used in LoLX phase-1 are shown in Figure 2.15, overlaid with the spectra of the principal light emissions observed. LXe scintillation is centred at 175 nm with a full-width at half-maximum (FWHM) of 10 nm [114]. This is selected in LoLX-1 with the BP filter with a peak transmission of ~ 20 %. The expected Cherenkov distribution from a 2 MeV β travelling through LXe (normalized to 1) is selected for by the LP filters, with a flat transmission of ~ 90 % above 220 nm. The detector, illustrated unwrapped to highlight the filtering, is shown in Figure 2.16.




Figure 2.15: Transmission of LoLX-1 filters, overlaid with light emission spectrums of Cherenkov, LXe Scintillation and SiPM external crosstalk (ExCT).

In addition to these two expected light production mechanisms, we also observed the cor-

related light production from SiPM ExCT. The wavelength distribution shown in Figure 2.15 originates from the ex-situ measurements by our collaborators at TRIUMF [115]. We also observed significant luminescence from the 3D printed cage of LoLX during LXe operations induced by the VUV scintillation photons. The luminescence was observed to exhibit a short time constant ($\mathcal{O}(\text{ns})$) fluorescence component and a longer lived $\tau \sim 1 \mu\text{s}$ component. The spectrum shown in Figure 2.15 was measured ex-situ with a 265 nm exciting light source at room temperature, taken by former undergraduate summer intern Guanqiao Li, under supervision by myself and Professor Thomas Brunner.




Figure 2.16: Illustration of LoLX detector unwrapped to show positions of optical filters, reproduced from [116].

The addition of a significant visible component from fluorescence of the cage material made resolving the subdominant Cherenkov signal in LoLX phase-1 intractable. Another point of uncertainty in LoLX phase-1 data was the observation of a significantly lower than

expected SiPM efficiency in LXe. As reported in [117], when comparing simulations of ^{90}Sr events to the reconstructed energy, agreement was found with PDE $\sim 12\%$ in the VUV band.

The final challenge encountered in LoLX phase-1 was the precise modelling of the energy distribution of the ^{90}Sr source. Due to the degeneracy between light production, transport, and detection effects, the lower light yield seen in the LoLX data could be due to impurities in the LXe, or to absorptive build-up on the SiPM windows, such as water outgassed from surrounding components. Simulation and data energy spectral shapes were not in agreement even after extensive tuning [117], which we suspected were due to the idealized geometry of the ^{90}Sr in the Geant4 simulations. During decommissioning our suspicions were confirmed with microscope images of the needles source showing significant cracking and flakes in the plated coating.

2.3.2 LoLX Phase-2 Detector

In the second phase detector for LoLX (LoLX phase-2), we have endeavoured to account for all the main challenges in the phase-1 detector design. The LoLX phase-2 detector employs 40 HPK VUV-4 SiPMs, as well as 40 Fondazione Bruno Kessler (FBK) HD3 VUV sensitive SiPMs, the two candidate SiPMs for the nEXO experiment [101]. In addition to the FBK SiPMs for cross-comparison, a 1" VUV sensitive PMT (Hamamatsu R8520-406) is installed in the first stage of LoLX phase-2. The heat load from the PMT was considered negligible, and no special measures were taken to mitigate it. The addition of this calibrated light sensor allows for relative measurements of the efficiency of the SiPMs employed in LXe conditions.

We can summarize the challenges of phase-1 and the interventions taken in phase-2 as follows:

1. LoLX Detector Cage Fluorescence - Limiting non-VUV sensitivity.

- Replaced 3D printed cage with open non-polymer based PCB to inhibit any fluorescence from traditional FR4-type boards, non-utilized surface have an exposed

copper layer for enhanced reflection of VUV light.

2. Low light yield in LXe - Isolate the possibility of degraded HPK SiPM performance in LXe.
 - Two VUV SiPM types employed, alongside a calibrated reference PMT with a known VUV sensitivity.
3. Uncertain source geometry - Difficulties in energy reconstruction and simulation data comparison due to the peculiarities of the source surface.
 - During commissioning, we employed external gamma sources for calibration. In parallel, internal gaseous calibration sources and neutron activated sources are being developed.




Figure 2.17: Upside-down photo of LoLX Phase-2 detector during assembly. SiPMs are tiled with one HPK VUV-4-Q and 4 FBK HD3's in a 2×2 array for each row, resulting in 16 HPK and 16 FBK SiPMs per tile face. The PMT, shown at the bottom replaces the top tile of the cube, and the laser calibration fibre passes through the spaces between tiles.

The LoLX phase-2 detector design is highly modular, and each tile may be outfitted with optical filters to enable wavelength selection as in LoLX phase-1. However, during commissioning runs covered in this work, no optical filters were used. Not shown in Figure 2.17 is the optical fibre connected to the pulsed diode laser used for SiPM and PMT calibration.

2.4 LoLX Data Acquisition and Slow Control

In this section, we review the functional aspects of the LoLX experiment, including the data acquisition (DAQ) system and the slow control system for experiment monitoring. The LoLX slow control and DAQ systems are built with the MIDAS framework [118], developed by TRIUMF in Vancouver, Canada, and PSI in Würenlingen, Switzerland. MIDAS, or the Maximum Integrated Data Acquisition System, is a general purpose control software designed for particle physics experiments [118].

2.4.1 Slow Control with MIDAS

The LoLX slow control system has several components, some of which interface with the XGHS described previously. Here, the “slow” in slow control refers to any components of the detector monitoring and control system which are not time-critical and are run at a lower priority. In our case, this means all pressure and temperature monitoring, cryogenic control settings, and (planned) liquid level sensors.

An illustration of LoLX highlighting temperature sensors and heater positions is shown in Figure 2.18. All temperature sensors and the four pressure sensors (not shown) are readout through a custom MIDAS front-end (FE) program. There are high pressure (Pirani) and vacuum pressure gauges (Cold-cathode) in the vacuum and xenon-wetted spaces. The custom slow control FEs are written in Python or C++, and communicate with the sensors directly, or through single-board computers which are reading serial data from the sensors. Remote controllable and online monitoring tools are essential for longer duration LXe campaigns with LoLX, allowing us to extend the physics reach of the experiment.

The main slow-control FEs including a brief description are:

- Lakeshore Control - Communicates with the Lakeshore 350 control board, sets PID settings and controls heater outputs. Reads back RTD temperature values.

Legend





Figure 2.18: Illustration of LoLX phase-2 cooling system and slow control readout variables. Heaters distributed throughout the cooling system are controlled by PID using a Lakeshore 350 temperature controller to maintain stable LXe conditions. The main readout temperature sensor for PID is l-RTD4.

- Cryocontrol FE - Communicates with sensor readout LabJack™ and RaspberryPI™. Reads pressure sensor values, thermocouples and external RTD-2.
- LLSensor - Communicates with liquid level sensor Arduino via a serial interface, and reads back capacitance values to convert into liquid level with a configured offset value.
- nanoPI_0# - Six identical frontends that communicate with nanoPI preamplifier and voltage supply boards. Reads back voltage and current on SiPM channels, and sets bias voltages.

2.4.2 Data Acquisition System

The LoLX DAQ has undergone two major stages, called phase-1 and phase-2. The phase-1 DAQ was in use from the initial commissioning stages in 2019 until early 2022. The phase-2 DAQ was commissioned in 2022 and has been in use since. The two DAQ systems and the key components are outlined in Figure 2.19. The major difference between the two systems was the upgrade from the CAEN V1740 digitizer, which was limited to 62.5 Mega Samples per Second or MSPS, to the WaveDAQ or WDAQ system from the MEG-II experiment [118]. This upgrade enables much finer timing, with up to 80x faster digitization rates (5 gigasamples per second).




Figure 2.19: Overview illustration of the LoLX phase-1 and phase-2 DAQ data flow. Both DAQ systems use the MIDAS framework, while phase-1 used a commercial CAEN V1740 digitizer, the phase-2 DAQ was loaned to us through a partnership with the MEG-II collaboration.

The high-speed phase-2 WDAQ enables fast-timing studies of LXe scintillation and the investigation into Cherenkov separation using photon arrival times. While the WDAQ may operate at much faster rates, each waveform is limited to 1024 entries, trading off a shorter

event window for faster timing. Both the WDAQ and V1740 are 2 V dynamic range digitizers with 12 bit resolution.

There are a variety of trigger settings that are employed during operation, we highlight the primary triggers used with LoLX. Trigger thresholds are configured for each channel individually using the estimated SPE amplitudes.

- External Trigger — Trigger supplied by external signal from waveform generator, the same signal drives the laser for calibration.
- SiPM Majority Trigger — Trigger supplied when a given majority level (number of channels where the trigger condition is met) is reached on the SiPM channels.
- Phase-1 Group Triggers — Triggers on each filter group individually, typically with majority $\equiv 1$ (Unfiltered, VUV Filtered, VIS filtered).
- Phase-2 PMT Trigger — Triggers on the PMT signal when threshold is met (Typically set between 2.5-4.5 photon equivalent amplitude).

The DAQ was created with the MIDAS framework, and consists of several key custom and built-in programs for operation. A brief description of the key DAQ programming and their function is given.

- mlogger — MIDAS Logging program, manages writing/reading from online data-base (ODB) and writing of output data files.
- mserver — MIDAS Server program, controls remote access to main DAQ and manages client programs.
- mhttp — MIDAS interactive web client program, creates HTTP server which enables interactive MIDAS experimental control.

- Digitizer FE - Communicates data and settings to front-end electronics, including desired trigger and digitizer settings from ODB. This program creates MIDAS databanks from the input waveforms to be stored in binary data files when triggered events are returned.
- MERCI Online — Online version of analysis software, reads data stream in real-time to enable live data quality checks.
- Trigger Control — Python scripting client to configure DAQ, used to streamline trigger setting application to the large number of ODB locations, to avoid human input errors.

In addition to the MIDAS DAQ software described here, and the custom analysis code described in the subsequent section, there is an additional tool used during DAQ operations. In phase-2 of LoLX, the MEG-II WDAQ was shipped with an inspection and calibration utility, “wavedaq-scope” (WDS) . WDS has firmware upgrade options, overriding control for DAQ boards during troubleshooting, timing and voltage calibration tools, and board synchronization. Before each data-collection period, the timing, voltage, and synchronization between the two digitizer boards in LoLX is carried out.

2.5 Analysis and Simulation Software

The following section provides an overview of our LoLX data analysis and simulation software programs. The online/offline analysis software outlined here was a substantial effort, with significant contributions from TRIUMF collaborators, as explained in the Statement of Contributions section.

2.5.1 Online/Offline Analysis Software

The analysis framework used by the LoLX experiment is called MERCI (Modular online multi-threaded C++ based waveform toolkit). The analyzer takes as input the output of the

MIDAS DAQ system, MIDAS files (.mid), and produces .root output files using the ROOT C framework. MERCI is developed around the analyzer framework using an optimized C++ methodology with Flow Queue and Flow Events. Each module included a Flow Analyzer method, which takes in the Flow Queue with Flow Events and can modify or append the flow objects. This allows for significant performance improvements through multi-threading, where each module runs independently.

For example, once Event#0 has been unpacked by the unpacker module, it will be passed to the baseline module to analyze the ADC Flow event, and simultaneously Event#1 will begin unpacking. Through this infrastructure, the entire analysis pipeline may be as fast as its slowest module, rather than the sum of the modules in series.

An overview of the LoLX MERCI analysis flow is illustrated with a flowchart in Figure 2.20, with key components illustrated. There are two operational modes for MERCI: online and offline.

In the offline MERCI analysis, the program takes in the MIDAS output files and performs the binary decoding, waveform correction, pulse-finding and high level analysis methods before writing to a ROOT TTree within an output ROOT file. In online MERCI, the program is connected to the MIDAS mserver and is run as a worker program under the central DAQ server. Running in online mode requires the MERCI program to be installed on a MIDAS DAQ server, while the offline MERCI program can be executed with an independent analyzer installation that is fetched by CMake during installation.

While running in online mode, a ROOT THttp server is created in place of the TreeWriter module, and the MIDAS data buffer is sampled periodically. The sampled binary data events are sent through the MERCI analysis flow and are analyzed identically to the offline events. In the online THttp server, all waveforms and a large selection of analysis plots are available, created by the Histogram module, to allow for live data quality checks by LoLX data taking shifters. One significant advantage of the MERCI online/offline implementation is running




Figure 2.20: MERCI LoLX Analysis flow overview. Flow analysis modules are shown with brief descriptions of their function, flow event objects are shown above in coloured arrows. Each subsequent module has access to previous flow events added to the flow queue for a given event.

the same offline analysis pipeline during live data analysis for quality checks.

SiPM Pulsefinding

The most critical part of our analysis pipeline is the extraction of photon information from the raw waveforms recorded by the digitizer. All quantifiable information about the event is encoded in the SiPM pulses present within the waveforms. An example of SiPM pulse waveforms comparing the phase-1 digitizer (CAEN V1740) and the phase-2 WDAQ is shown in Figure 2.21, where the WDAQ is operated at its nominal rate of 1 giga-samples per second (GSPS).

As seen in Figure 2.21, there is a significant improvement in timing and shape information in the phase-2 digitizer (WDAQ) waveform. This improvement comes with a downside of higher input bandwidth, leading to a lower signal-to-noise (SNR) as a result of the increased fast baseline noise. This is a key consideration in the change between pulse-finding techniques




Figure 2.21: Comparison of LoLX Phase-1 DAQ digitizer (top) and Phase-2 DAQ digitizer (bottom). Both SiPM waveforms show single pulses with ~ 2 PE, by pulse height.

used for LoLX phase-1 and LoLX phase-2 analysis.

In LoLX phase-1, pulse-finding was implemented using an algorithmic approach to waveform analysis. Each waveform is analyzed individually, traversing from the beginning of the waveform to the end. The LoLX phase-1 pulse-finding algorithm is an extension of the algorithm developed by myself from another project [119]. Pulses are identified using rolling thresholds in amplitude, and by threshold in the first derivative of the waveform (numerically calculated using Newton's method). This effectively selects pulses in a high SNR regime, as in LoLX phase-1. The first derivative is used to identify compound pulses, which contain “sub-peaks”, i.e. subsequent photon induced avalanches within the same SiPM, at later times.

Before pulse-finding can occur, the waveforms are first corrected for baseline shifts and distortion. Zeroing baselines for waveform analysis is standard, however, the distortion correction requires a more thorough approach. The distortion is induced by the amplifier circuit in the readout chain, and manifests as a positive overshoot after a large signal in the SiPM. Correction is necessary for accurate photon information reconstruction from the waveform,

and follows the same procedure described in [119].

After baseline reconstruction, phase-2 analysis proceeds to pulse-finding as outlined in Figure 2.20. The LoLX phase-2 SiPM pulse-finding analysis differs from the traditional approach used in phase-1 by employing matched filtering for pulse identification, with custom SiPM pulse templates.

Matched Filter SiPM Pulsefinding

The “Matched Filter” is a linear transformation to the waveform, that with perfect template matching, produces the maximum SNR in the output filtered waveform [120]. As the literature of matched filtering techniques is extensive, we will restrict ourselves to describing the necessary details for SiPM pulsefinding.

To reach the performance requirements for LoLX, where for each data-taking campaign we collect $O(10^9)$ waveforms, we take advantage of Fast Fourier Transform techniques (FFTs) to accelerate the analysis [121]. The templates used are a key consideration, and for LoLX phase-2 we have three different sensor types, each with unique pulse-shapes. The templates were optimized by manually fitting a sample of single and 2-PE pulse waveforms for each SiPM vendor (FBK and HPK) to the Single-Avalanche-Response-Function (SARF) developed by TRIUMF’s photodetector development group. A simple Gaussian template was selected for the PMT signal. The SARF model has the form of a doubly-exponentiated Gaussian function. Each exponential component has its own independent time constant, related to the fast and slow discharge of avalanches from SiPMs [122].

The SARF Function can be written as:

$$V(t) = A \cdot \left(\frac{(1-k)}{2 \cdot \tau_s} \cdot V_s(t) + \frac{k}{2\tau_l} \cdot V_l(t) \right), \quad (2.3)$$

where A is the amplitude of the template, while V_s and V_l are defined as:

$$V_{s/l}(t) = \exp \left(\frac{\sigma^2}{2\tau_{s/l}^2} - \frac{t-t_0}{\tau_{s/l}} \right) \cdot \operatorname{erfc} \left[\frac{\sigma}{\sqrt{2}\tau_{s/l}} - \frac{t-t_0}{\sqrt{2}\sigma} \right], \quad (2.4)$$

with short and long exponential time constants τ_s and τ_l respectively, gaussianity parameter

σ , initial time t_0 , and k , which represents the sharing fraction between short and long time constants.

A modified version of SARF was used to model the HPK pulse-shape, which as seen in Figure 2.21, has a rounded-top that the SARF could not reproduce. The modified SARF function used for the HPK SiPM includes a sigmoid term, allowing for a gentle turning-on of $t > t_0$ components using an effective time constant for the sigmoid (ϵ):

$$V(t) = A \cdot \left(\frac{(1-k)}{2 \cdot \tau_s} \cdot V_s(t) + \frac{k}{2\tau_l} \cdot V_l(t) \right) \cdot \left(\frac{1}{1 + \exp\left(\frac{t-t_0}{\epsilon}\right)} \right). \quad (2.5)$$

By taking the FFT of the waveform and of the template waveform, we can take advantage of the Cross-Correlation Theorem [123] to reduce the computational complexity from $O(N^2)$ to $O(N \log_2 N)$ [124]. The cross-correlation of $f(t)$ and $g(t)$, given here as $h(t)$, may be written as:

$$h(t) = (f * g)(t) = \int_{-\infty}^{\infty} f(t) g(t - \tau) d\tau. \quad (2.6)$$

Using the cross-correlation theorem, we relate the complex conjugates of the FFT of $f(t)$, and $g(t)$, here $\hat{f}(\xi)$ and $\hat{g}(\xi)$, using:

$$\hat{h}(\xi) = \hat{f}(\xi) \cdot \hat{g}(\xi). \quad (2.7)$$

Then by applying the inverse FFT to $\hat{h}(\xi)$ we recover the cross-correlation $h(t)$, which results in a correlation waveform that is shifted by $+t_0$ relative to the start of the raw signal waveform. This is illustrated in the top two panels of Figure 2.22.

While the correlation waveform contains valuable information about the probability of finding a template-matched signal at a given location, the photon information in the waveform still needs to be reconstructed accurately. The reconstruction or “pulse extraction” happens using information combined from the baseline corrected waveform, and the correlation waveforms. The analysis procedure is described at a high level as:




Figure 2.22: LoLX phase-2 wave form analysis overview outlined. 1) Example of a baseline corrected signal waveform, and pulse finder results shown in pink with pulse template overlaid. Pulsefinder results have reconstructed pulse timing shown via box edges, and reconstructed amplitude in box height. 2) Shifted and SNR normalized Correlation Waveform (nCWF), 3) first derivative of nCWF (nCWF') with relative and absolute thresholds overlaid, 4) Second derivative of nCWF (nCWF'').

1. Correlation waveform preparation

- Template waveform setup - Using user-defined parameter values, define the template function and create the template waveform for each device.
- Matched Filter with template in Fourier space - Take the FFT of both Template and Baseline Corrected waveforms and perform matched filtering using cross-correlation theorem.
- Correcting Correlation waveform - Take the Inverse FFT of the matched filter results to create the raw correlation waveform, then cut and shift waveform to recover the original timing, and align with estimated t_0 of the Baseline Corrected

waveform.

2. Peak finding setup with Correlation waveform

- Low variance region identification - Using the baseline corrected waveform, we identify a low variance region by measuring baseline RMS with a rolling window. Low variance regions are assumed to be pulse-free and are used for normalization.
- Normalize Correlation waveform to SNR - Using the low variance region, take the offset and RMS of the baseline of the Correlation waveform and then shift and normalize the Correlated waveform (nCWF).
 - This is done to allow for inter-waveform data cleaning cuts with a normalized scaling by using the “significance” of a pulse as the criteria.
- First derivative of Correlation waveform - Take the numerical first derivative of the normalized Correlation waveform, used to evaluate the absolute slope in significance units (nCWF').
- Second derivative of Correlation waveform - Again take the numerical derivative, of the first derivative in the previous step, to constrain the polarity of pulse peaks (nCWF''). As the SiPM produces negative polarity peaks, positive polarity peaks are identified as noise.

3. Pulse region identification and peak extraction.

- Peak locating - Using a rolling absolute (significance) and a relative threshold (to baseline RMS), identify peak locations in the first derivative of the normalized Correlation waveform.
- Peak selection - Check second derivative for peak polarity, SiPM peaks in LoLX have negative polarity, while radio-frequency interference (RFI) pickup may appear in either polarity, also find the pulse peak location from inflection point.

- Pulse end determination - Find the end of the pulse region using proximity to pre-pulse baseline of the normalized Correlation waveform.

4. Pulse analysis for sub-peak extraction

- True Minima - Using the baseline corrected input waveform and pulse region information, evaluate the true minimum for pulse amplitude and peak time.
- Charge Integration - Integrate the baseline corrected waveform within the pulse region using trapezoidal integration, also evaluate the “gate” charge in a sub-region².
- Sub-peak Extraction - Using the second derivative of the normalized Correlation waveform, identify sub-peak candidates and zero-crossings positions within the pulse-region using a relative threshold (in nCWF'' RMS units).
 - Within the main pulse region, significant spikes in nCWF'' followed by zero-crossings are evidence for subsequent pulses piling up in the primary pulse region.

2.5.2 LoLX Geant4 Simulations

The standard tool for particle physics radiation transport simulations is Geant4 (formerly an acronym for GEometry ANd Tracking or GEANT) [125], developed at CERN by the Geant collaboration over the past five decades. For phase-1 of LoLX, we employed a fully detailed Geant4 simulation with a detailed optical transport simulation. As part of the work in the following chapter on characterizing SiPM ExCT in LoLX, the default Geant4 optical transport physics was modified to enable the addition of a custom physics process for SiPM ExCT. The details of the ExCT simulation are described in Chapter 3. The Geant4 simulation framework is a complete toolkit for the radiation transportation simulations relevant to LoLX. It

²Gate integration for improved charge estimation is explained in Chapter 4.

allows for implementation of detector geometry, custom radiation particle sources, extensive physics processes employing catalogued tables of interaction cross-sections, and energy loss and transport with customized scoring for tracking of analysis relevant parameters. Simulation analysts prepare “macro” files (.mac) which use standard and custom methods to send instructions to the Geant4 run manager. This includes toggles for physics lists and methods, selection of simulation sources, and output control.

The phase-1 LoLX Geant4 simulation, `lolx-sim`, was initially developed by postdoctoral fellow Thomas McElroy, and subsequently developed by myself and fellow PhD Student Bindiya Chana [117]. `lolx-sim` was built with the Geant4 framework, based on the advanced examples for standalone simulations.

In `lolx-sim`, we have modelled the phase-1 LoLX detector in detail using `CADMesh` [126] to handle the import of the complex cage geometry shown in Figure 2.13. The remainder of the geometry (cryostat and support components) is modelled with standard solids in the Geant4 framework. All relevant optical surfaces have been modelled with built-in and custom optical physics. One major component of this work was the implementation of dichroic surfaces by Bindiya Chana [117], which correctly models the optics of the longpass filters.

As part of our work in modelling SiPM ExCT we have implemented a fully detailed SiPM model for the HPK VUV4 devices used in LoLX phase-1. The ^{90}Sr calibration source is modelled and the full decay spectrum has been implemented as a custom particle source.

The physics list employed in LoLX phase-1 simulation, `lolx-sim`, is as follows:

- `G4EmStandardPhysics` - Standard electromagnetic physics (γ , β , ions, hadrons) up to 100 TeV.
- `G4EmLowEPPPhysics` - Low energy electromagnetic physics (γ , β , ions, hadrons) from 10 eV to 1 GeV, including polarized processes and atomic de-excitation.
- `G4DecayPhysics` - Decay of particles in flight and at rest.

- G4RadioactiveDecayPhysics - Decay in flight and at rest of radioactive nuclei ($\alpha, \beta, \gamma, EC$).
- G4OpticalPhysics - Optical interactions for $\lambda > 10$ nm photon energies, including Rayleigh scattering, absorption and boundary effects (Default G4OpBoundary and G4Scintillation are replaced by custom versions).
- NESTScintillation - Replacement of G4Scintillation with NEST toolkit Geant4 bindings using a modified NEST v2.2.4 model.
- LOLXMunPhysics - Custom ModularPhysicsList including all relevant muon interactions.
- LOLXExtendedOpticalPhysics - Custom ModularPhysicsList which includes customized LoLXOpBoundaryPhysics and the ExternalCrossTalkPhysics methods.

Hardonic physics processes are not enabled for the majority of particles, except muons, which can reach relevant energies.

In phase-2 of LoLX we have moved to a serial heterogenous simulation strategy. The simulation is broken into two stages, a traditional CPU-based Geant4 simulation for particle transport followed by a light transport simulation using a GPU accelerated light transport simulation framework known as Chroma [127]. In the first stage of the simulation, called G4LoLX, we have modelled the LoLX cryostat with an oversized custom sensitive scoring volume representing the LoLX phase-2 cubic detector geometry. The phase-2 visualization is shown in cross-section in Figure 2.23.




Figure 2.23: LoLX phase-2 Geant4 simulated geometry, G4LoLX. The sensitive scoring volume is highlighted in magenta within the LXe volume, within the surrounding xenon volume.

Departing from the design of the phase-1 LoLX simulations, we only model the radiation transport and energy loss with G4LoLX, while optical transport and production processes are handled by a subsequent stage of the simulation. This next stage produces a custom simulation input table, necessary for the second stage of the simulation, by G4LoLX in a custom ROOT TTree. The output of the G4LoLX stage includes the following necessary information for the following simulation stages, for each “step” in the simulation, we write the variables shown

in Table 2.1.

Table 2.1: List of step-level Geant4 output variables for Chroma simulations.

- Particle Type	- Starting energy at this step
- Particle ID	- Deposited energy at this step
- Track ID	- Remaining energy at this step
- Parent Track ID	- Position (X,Y,Z) at the end of the step
- Momentum ($\vec{x}, \vec{y}, \vec{z}$) at the end of the step	- Time at this step

This set of information is required for sorting tracks based on parent/daughter structure, producing scintillation light with the NEST model, and Cherenkov radiation using a custom implementation of G4Cerenkov in the GPU simulation.

Primary particle and radiation sources are modelled using the standard GPS or “General Particle Source” tool, however custom source implementation is possible and may be required for future simulation development.

2.5.3 GPU-based Light Transport Simulations

The major change in simulation strategy between LoLX phase-1 and 2 is the transition to heterogeneous computing with the application of GPU accelerated photon transport. The reason for the significant change is to align the LoLX simulation and data comparisons with the nEXO simulation framework, to enable cross-validation. For LoLX we use Geant4 for particle transport simulations, and Chroma [127] for photon transport simulations. Chroma is a GPU accelerated photon transportation framework developed by Stan Seibert [127], and is also in use by the nEXO collaboration. Simulations in phase-2 of LoLX use G4LoLX for the initial radiation transport and energy loss calculations, and the second stage of the simulations occur within the Chroma simulation framework. The second stage includes light production, transportation and detection, and detailed detector modelling. The phase-2 detector is shown in Figure 2.24 using the built-in visualizer tool from the Chroma framework.




Figure 2.24: Ray-tracing rendering of LoLX phase-2 detector in the Chroma simulation detector viewer.

Chroma is an open-source high performance optical photon simulation tool developed for particle physics detectors. It uses a surface-based geometry modelling paradigm, distinct from the solid-based paradigm used by Geant4 [127]. The surface-based geometry uses triangle meshes in the form of STL (Standard triangle language) files to represent bodies in the simulation. One key advantage to this methodology is the simplification of geometry design and maintenance for experimentalists. As STL files can be readily prepared by CAD softwares, existing models of the detector used in design and production may be exported to mesh files for direct application in the Chroma simulation [127].

In contrast, the solid-based geometry paradigm used by Geant4 requires strict hierarchy of mother-daughter volumes which is achieved by Geant4 with a tree structure of solids. In the solid-based design one must define the geometry of the detector using primitive solid shapes such as: spheres, cylinders, boxes, and some polyhedra. This implementation may lead to performance trade-offs, where sophisticated models required for some geometries lead to complex implementation and lower performance in tracking due to the denser geometry tree

traversal. With a triangular mesh, Chroma can apply optical surface properties to each triangle individually for fine-grained control, however, it is typical to implement a common optical surface for each mesh.

Beyond implementation simplicity and performance differences with a surface-based geometry implementation, the main advantage of Chroma over traditional tools is the use of GPU acceleration for photon transport. Each signal event ($0\nu\beta\beta$) in nEXO produces $\sim 70\text{k}$ photons. This requirement introduces computational roadblocks quickly as several decades of simulated detector exposure are required for some analyses. The speed of the simulation is a combination of the efficiency of algorithm implementation, software design choices, and hardware. Driven by advances in graphics in video game industries, and general purpose GPU computing from Artificial Intelligence (AI) industries, GPU hardware has grown exponentially in capability. The Chroma simulation toolkit takes advantages of some of these developments, by implementing the fully detailed optical physics core code in CUDA-C using NVIDIA GPUs, with its own implementation of photon ray casting. This development was several years ahead of the development of dedicated ray casting hardware and API implementations in modern GPUs, so further performance upgrades are still accessible. However, as photon transportation is a highly parallelizable process, even without dedicated hardware, Chroma reports a $200\times$ speed-up over traditional CPU photon transport for an equivalent simulation in Geant4 [127].

The Chroma source code [127], has been integrated into a complete simulation package for the nEXO experiment, which we call “chroma-simulation” to differentiate. In the full simulation package, we have developed a separate detector simulation for the LoLX phase-2 detector. This development required some significant upgrades to the core code of chroma-simulation, in order to generalize detector building, event and photon production, and event variable saving. These changes were necessary to expand the chroma-simulation tool beyond nEXO specifications, to be generalized for any LXe detector within the nEXO collabo-

ration network, and are described in Chapter 6. The Chroma simulation toolkit also includes a simplified DAQ and scoring implementation, where meshes are assigned as “detectors” with a wavelength dependent detection efficiency. We have recently expanded on this capability to include 2D wavelength and incident angle detection efficiencies for surfaces.

With our overview of LoLX and the upgraded components for the second phase of the experiment completed, we move onto the measurement of SiPM ExCT in LXe carried out with LoLX phase-1. In Chapter 3 we present an expanded reproduction of the recent publication in European Physical Journal C [116]. This will be followed by the initial commissioning of the LoLX phase-2 experiment.

3

Measuring SiPM ExCT

3.1 Motivation

For experiments using SiPMs for light detection, understanding and characterizing all sources of correlated noise is essential. This is especially true for next generation detectors looking for evidence of physics beyond the Standard Model, such as the nEXO experiment. In order to make precision measurements beyond the state of the art, all sources of variances in detector response should be well understood, and, if possible, reduced to negligible levels.

In this chapter we discuss the measurement, simulation, and modelling of SiPM optical crosstalk between distant sensors. This correlated noise source, introduced in the previous chapter as “SiPM external crosstalk” (ExCT), is of particular importance due to its geometry

dependence and the gap in the current literature of understanding this phenomenon. These measurements were carried out with the LoLX phase-1 detector system in October 2021 with analysis and modelling conducted over the following years.

3.1.1 SiPM Optical Crosstalk

As discussed in Chapter 2, SiPMs detect light by inducing an avalanche of charge carriers in a Gieger-mode SPAD, with an avalanche releasing between 10^5 - 10^7 electrons. This avalanche process also produces secondary emission of photons, known as optical cross-talk (OCT). The emission of secondary photons is correlated with several factors, such as the electric field, impurity concentration, doping profiles, and the geometry of the SPAD [128].

These secondary photons can induce several types of correlated noise within SiPMs. This includes internal cross-talk, optically induced after-pulsing, and external crosstalk. Understanding the fluctuations in these sources of noise and their impact on future particle physics experiments that employ SiPMs is an ongoing field of study. For the case of SiPM ExCT in particular, experimental measurements have been carried out in a variety of settings, which we can classify into two broad categories: DC measurements of SiPM light emission characteristics, and idealized geometry analogues for SiPM ExCT.

Briefly, in DC-mode measurements of light emission characteristics from SiPMs, such as those done by our collaborators at TRIUMF in [105], [115], the emitted light from SiPMs is collected by microscopes with cameras and spectrometers that integrate the signal from many avalanches. These measurements are performed in vacuum at operational temperatures, and employ detailed modelling to estimate the light transfer function from the SiPM to the narrow field-of-view (FOV) light integration optics. These measurements are essential for understanding the fundamental processes involved in secondary light emission from SiPMs, but translating this information to detector-specific noise estimations is challenging.

In contrast to DC-mode measurements, idealized geometry analogues for SiPM ExCT

studies use small-scale geometries, often in operational conditions (LXe or LAr), with SiPMs directly facing each other [129]–[131]. In this measurement methodology the test detectors are sensitive to the optics and conditions expected in full-scale detectors, however, they are limited by the narrow FOV of the mirrored SiPM geometry. Moreover, the measurement techniques used by these apparatuses rely on sweeping over OV's on one SiPM tile, while using the opposite facing SiPMs as detectors, to work out the expected “cross-talk” probabilities.

For our measurement of SiPM ExCT in LXe, we aimed to develop simulation and analytical models of cross-talk for use in other detector geometries. To accomplish this, we developed a novel measurement technique using timing coincidences between SiPM pulses to estimate the levels of ExCT observed in the LoLX detector. By using the LoLX phase-1 detector described in the previous chapter, we cover a much larger FOV and solid-angle than those used previous measurements, with a geometry similar to the large cylindrical barrel of the nEXO photodetector system [52]. Using our simulated model of SiPM ExCT we are also able to better understand the effect of optical surfaces in the detector on SiPM ExCT, as photon transport considerations differ greatly between VUV and NIR for many materials. In developing these models, we move away from describing SiPM ExCT as something with a distinct “probability” specific to a SiPM model, into a more broadly applicable model of SiPM secondary photon emission, measured in avalanche photon yields. By describing SiPM ExCT in this way, we hope to be more precise in our estimations of the impact of SiPM ExCT, and its peculiar environmental and geometric relationships. Every measurement of SiPM ExCT will depend on the optical media in which the SiPM is contained, and the geometry in which they are configured. To go beyond this, our simulation and analytical models may be applied to arbitrary detector systems employing SiPMs, to answer specific questions about resolution, trigger thresholds, and other confounding detector response effects to which SiPM ExCT contributes.

3.2 Measurement Technique

The measurement of SiPM ExCT in LXe was performed using the LoLX phase-1 detector and cryostat, described in Section 2.3. LoLX phase-1 used a ^{90}Sr needle source to induce scintillation events within the detector volume, Figure 3.1 shows an illustration of the inverted LoLX phase-1 detector with SiPM ExCT. This figure highlights SiPM ExCT signatures within LoLX, under consideration in our study. At step 1) the β from a ^{90}Sr decay induces primary LXe scintillation light, in step 2) the primary or “seed” photons travel to SiPMs on the walls of LoLX, inducing an avalanche in the unfiltered or “bare” SiPM. In step 3) the seeded avalanche from 2) produces NIR secondary photons that travel to a nearby LP filtered SiPM and induces a correlated avalanche, which is referred to as ExCT. SiPM ExCT produced in the unfiltered SiPMs, that is detected by any surrounding LP filtered SiPM channel, is called “forward” ExCT. Similarly, in step 4) secondary photons produced by seeded avalanches in the LP filtered channels, that are then detected by the unfiltered SiPMs are called “backward” ExCT. This distinction is necessary for our analysis, and the classification is motivated by the analysis strategy employed here. The LP filtered channels make up the majority of the light sensitive surface area within LoLX (22/24 packages), and filter the majority of scintillation light. As such, they are the ideal search channel for the ExCT analysis.




Figure 3.1: Visualization of primary scintillation and ExCT signatures in an inverted LoLX phase-1 detector. 1) Primary scintillation light is produced by ^{90}Sr β emission, 2) Photons travel across detector to SiPMs and “seed” ExCT emission, 3) NIR photon emission from the avalanche escapes the unfiltered device and travels to a nearby LP filtered SiPM. 4) Seed photons detected by a LP SiPM create subsequent emission of NIR secondary photons, which are detected by the unfiltered SiPM. ExCT coincidences from 3) are classified as “forward” and those from 4) as “backward” correlations. The illustration shown here is inverted relative to the installed orientation, and is reproduced from [116].

The experimental signature of ExCT we are looking for in LoLX is a single photon equivalent pulse in the LP channel group and the unfiltered groups in time coincidence. When two pulses are observed in coincidence, the direction is ambiguous, and can either be forward (unfiltered $\rightarrow$ LP) or backward (LP $\rightarrow$ unfiltered). The expected secondary emission photon travel time within LoLX is ~ 200 ps, nearly three orders of magnitude lower than the sampling rate of the LoLX phase-1 DAQ. The timed analysis approach is illustrated in Figure 3.2. After the initial burst of scintillation light, we look late in the event waveform for delayed pulses. These “seed” photons, induce the primary SPE pulse on the source SiPM and are primarily from the fluorescence of the cage, late after-pulsing and stray scintillation light. When a candidate pulse is identified in the unfiltered channels, all LP channels are searched for

single PE pulses in a coincidence-candidate window surrounding the initial source candidate pulse on the unfiltered channel.




Figure 3.2: Illustration of SiPM ExCT search method with candidate pulses. Coincidences between two pulses in the search window are recorded as the time difference between the source candidate in the unfiltered channel, and the detection candidate in the LP channel, this process is repeated for all pairs of channels. Figure reproduced from [116].

The LP $\longleftrightarrow$ unfiltered pairing was selected for the large sensitive surface area of the LP channels, which maximizes the statistics for “successful” ExCT events (meaning a NIR photon produced in one avalanche successfully triggers an avalanche in another sensor), and has a high solid angle coverage to the unfiltered package. The filtering of the LP filtered SiPM group reduces the contribution of uncorrelated background pulses.

3.2.1 Analysis Criteria

The data in this analysis was collected in October and November of 2021, with applied OV's of 4 V and 5 V. The waveforms collected for each channel are processed through MERCI and pulse-finding is carried out, saving the pulse and sub-pulse information, for each run individ-

ually. To apply single-photon equivalent (SPE) thresholds for the analysis, SPE calibration is carried out for each run individually, setting analysis thresholds on a run-by-run basis. The run specific SPE and analysis cuts corrects for minor temperature induced gain variations from run to run.

There are two main stages of analysis selection criteria that are applied. The first stage is applied at event selection, which removes confounding events from the analysis dataset. The second stage is at the pulse level, using reconstructed pulse information for ExCT candidate pulse pair selection. This data was taken for >3 photon equivalent ADC threshold on any of the 4 channels in the unfiltered group, to prevent an excess of noise triggers present at the single photon equivalent level.

The first stage of event selection criteria is as follows:

- No channels in the event may saturate the digitizers dynamic range.
 - This cut removes high energy events, such as muons passing through LoLX.
- All channels must have stable pre-trigger baselines.
 - This cut removes events that follow saturated or clipped preceding events.
- An event must have only one trigger candidate.
 - This cut removes accidental pileup and events with significant coherent RFI noise pickup. Trigger candidacy is based on the run DAQ settings.
- An event must have a well-defined start time.
 - This cut removes events with significant pre-trigger light, which may be in other channel groups.

The second stage of analysis criteria is to select candidate pulse pairings between the unfiltered and LP channel groups. The criteria for candidate pulses is as follows:

- Candidate pulses must occur within the low-occupancy search window. This window is defined by the post-trigger time window with less than 0.2 pulses on average per event. The window is shown in Figure 3.3 as a highlighted red band.
- Candidate pulses are SPE-like from charge-based estimation. This cut ensures that pulses produce approximately constant secondary photon emissions, known to be gain dependent [105].
- Candidate pulses have no additional sub-peaks. In combination with the previous criteria, this ensures that delayed secondary photons from after-pulsing do not contribute to the signal.

The low-occupancy search window was defined by first measuring the total average rate of pulses on all channels within LoLX, as a function of time since the beginning of the event. This is shown in Figure 3.3 for all 4 V OV runs in the dataset. The occupancy distribution has 16 ns bins, set by the sampling rate of the digitizer in phase-1 of LoLX.




Figure 3.3: Occupancy of all SiPMs within LoLX for 4 V OV data collections. Low-occupancy search window is highlighted in red and spans [800,1500] ns from the start time (t_0) of the event. The broad dip in the distribution between 200-600 ns is an artifact of the pulse-finding conventions used. Pulses that follow the primary scintillation burst will be classified as sub-peaks, until the signal reaches baseline. The occupancy distribution uses the leading edge of primary rather than sub-peaks. Figure reproduced from [116].

The low-occupancy window spans [800,1500] ns after the initial time of the event. An additional search window, the candidate search window, is defined as ± 296 ns relative to the pulse time of an initial candidate ExCT pulse in the unfiltered channel group. This window is much larger than the $\mathcal{O}(10 \text{ ns})$ timing resolution of LoLX phase-1, and will include both random uncorrelated backgrounds and ExCT coincidences.

The SPE pulse selection criterion is based on the estimated pulse charge, from waveform integration. An example of an unfiltered SiPM SPE distribution is shown in Figure 3.4, with a highlighted SPE selection region. The selection region is the one-sided 95 % confidence interval for the Gaussian fit to the 1-PE peak of the SPE distribution. This region was selected to capture the majority of variance in the 1-PE peak, while removing any events with after-

pulsing, which may not have been flagged from sub-peak information.




Figure 3.4: SPE charge distribution and selection window for an unfiltered SiPM channel in LoLX. The highlighted SPE pulse selection region is a one-sided 95 % confidence interval of the fitted SPE Gaussian model.

3.2.2 SiPM ExCT Correlation Probability

Through this analysis, we aimed to estimate the SiPM ExCT correlation probability in LoLX. That is, the probability, given a single-photon avalanche on any unfiltered SiPM, of finding a time-correlated SPE pulse on any LP filtered SiPM above the uncorrelated background. This was carried out following the analysis criteria in the previous subsection.

When an SPE pulse that passes selection cuts is found in the low-occupancy search window in any unfiltered SiPM channel we search all LP filtered channels for correlated pulses in the candidate search window fixed about the leading-edge pulse time of the unfiltered candidate pulse. When a correlated SPE pulse is found, the time difference to the unfiltered pulse, Δt , is calculated and stored in a histogram for post-analysis. We sum the paired pulse time differences for all unfiltered and LP channels into a single histogram of Δt for each run indi-

vidually. The summed Δt distribution for all 5 V OV runs is shown in Figure 3.5. The time bins are 16 ns wide, set by the digitization rate.




Figure 3.5: 5 V OV dataset delta-time distribution, with linear fit overlaid. The fit function shown is the linear background component of the combined signal+background fit, with residuals showing random fluctuations in background regions, and a signal excess in the correlation region about $\Delta t = 0$ as expected from SiPM ExCT. Figure reproduced from [116].

The distribution in Figure 3.5 has two main features, a peak centred at $\Delta t = 0$, and a flat falling background. The peak at zero is the expected signal of correlated pulses that we attribute to SiPM ExCT, and the falling background is composed of randomly correlated background pulses from accidental coincidences between the unfiltered and LP filtered channels.

The background components of Figure 3.5 are fit to a first-order polynomial. The result of this fit is used as the constrained prior for a combined fit of the linear background and a Gaussian signal peak with a fixed mean of zero. The linear background fit shown in Figure 3.5 is the linear component of the combined fit, to emphasize the signal from ExCT in residuals. The width of the signal peak is influenced by the digitizer timing, DAQ jitter, and pulse-finding time resolution. The residuals of the linear background and correlation data are shown in the bottom panel, normalized to the data rate.

Using the data from Figure 3.5, we calculate the probability for a correlated pulse pair between the LP and unfiltered channels. Critically, this differs from the true ExCT probability in either direction, as this distribution includes both “forward” and “backward” correlations. The correlation probability is expressed by:

$$P_{corr,B} = \frac{N_{correlated} - N_{background}}{N_B}, \quad (3.1)$$

where the denominator N_B is the total number of unfiltered (bare) candidate pulses found, and the numerator is the excess of correlated pulses found above the background. The first numerator term, $N_{correlated}$ is the sum of the three signal bins about $\Delta t = 0$, from $[-24 \text{ ns}, 24 \text{ ns}]$. The second numerator term, $N_{background}$ is estimated by evaluating the linear component of the combined fit at the bin centres of the three signal bins.

The term $P_{corr,B}$ represents the probability of observing a time-coincident SPE pulse in the LP filtered SiPMs, given an SPE pulse in an unfiltered SiPM. This observable includes contributions from both directions of time correlations, forward and backward, as illustrated in Figure 3.1. Because of the prompt nature of ExCT and the considerably slower timing resolution in LoLX phase-1, it is not possible to differentiate between forward and backward correlations with timing information alone. One important consideration is the normalization of $P_{corr,B}$, which includes only the trials for forward correlations, and does not include “failed” backward correlated pulses. Failed backward correlated pulses are SPE pulses in the LP filtered SiPM group, that do not “successfully” trigger a correlated coincident pulse in an unfiltered SiPM. This is a key contribution to extracting the so-called “true” ExCT probabilities between channel groups. Since we include only “successful” trials in the numerator from the backward correlations, we account for this in our effective model.

In the following subsection we will describe the SiPM ExCT Geant4 simulations and simulation analysis technique. This will be followed by a description of our effective model for SiPM ExCT, in Section 3.4, followed by results and discussion.

3.3 Simulating SiPM ExCT

The LoLX phase-1 simulations, introduced in Section 2.5, are Geant4 based with NEST integration for LXe scintillation light production. The optical transport and photon detection components of the simulation are the key inputs in this study, and all relevant optical surfaces have been defined across the full wavelength range under consideration (150—1000 nm). These optical data are summarized in Figure 3.6.




Figure 3.6: Wavelength dependence of LoLX phase-1 ExCT simulation parameters. Top panel: Distributions of optical filter transmission and expected ExCT wavelength distribution from [115], also shown is the measured room temperature (RT) emission spectrum from fluorescence of the 3D printed LoLX cage. Bottom panel: Summary of SiPM reflectivity and detection efficiency data used in this study. Three SiPM reflectivity curves are shown; in purple is the default reflectivity calculated by Geant4 Fresnel relations using SiO_2 refractive index data, in brown is the effective Geant4 reflectivity used in nominal simulations, and in grey is the upper limit on SiPM reflectivity from LXe measurement matching, overlaid with the reported value at LXe scintillation wavelengths from [132]. Figure reproduced from [116].

We implemented the SiPM ExCT process as a custom physics process in Geant4. This was combined with customized versions of the built-in optical transportation codes, and functions similarly to a fluorescence simulation. In our simulation framework and modelling, we assume every avalanche in a SiPM produces some amount of secondary photons. From those avalanches, light may be emitted into the surrounding environment. The simulation process and input information can be summarized as:

- During initialization, the wavelength dependent emission spectrum shown in Figure 3.6, from [115], is loaded into the custom physics process. Look-up tables are read-in for all optical surfaces.
- Photons that pass through the optical boundary of a SiPM detector surface trigger the modified G4OpBoundaryProcess. This only triggers for photons that are “detected”, after efficiency checks.
- After the successful detection of a seed photon, secondary photons are produced within the bulk silicon at the detection vertex. The emission intensity (in photons/avalanche) is a tuneable simulation parameter, and each avalanche samples from a Poisson distribution with the input intensity as the mean. The secondary photons are emitted isotropically in the outgoing hemisphere (w.r.t the surface normal) with a random polarization.
- Before exiting the SiPM, photons must pass a transmission check. The transmission probability from bulk Si to LXe was calculated by collaborators at TRIUMF using an ANSYS Lumerical™ far-field Maxwell’s equation solver. This was done to accurately model the microstructures within a SPAD, at a level of detail that is computationally infeasible within a full-scale Geant4 simulation. The transmission is both direction and wavelength dependent due to dispersion, and is shown in Figure 3.8.
- Photons that pass the transmission check are manually refracted and placed into LXe

just in-front of the SiPM surface. This was done to avoid built-in optical photon transportation through the SiPM volume, handled already by the transmission check. The emitted secondary photons are flagged as ExCT and are added to the stack of secondary particles for standard transport through the detector, producing secondary and even tertiary avalanches. The outgoing photon distribution after transmission checks, with refraction into LXe is shown in Figure 3.9 for reference¹.

These steps are illustrated in Figure 3.7 for the reader's understanding.




Figure 3.7: Illustration of the Geant4 SiPM ExCT simulation method with labelled steps. In step 1) a photon is detected at a given location, 2) a number of photons are produced isotropically, photons in the lower hemisphere (as defined by the surface normal) are killed. In step 3) photons must pass a transmission check for emission into LXe from $\text{Si} \rightarrow \text{SiO}_2 \rightarrow \text{LXe}$. In step 4) photons surviving the transmission check are refracted into LXe manually according to Snell's law using input optical data. The minor x/y displacement from within Si is negligible and not included. The resulting photon distribution is shown in Figure 3.9. Figure reproduced from [116].

¹A high density version of this emission map is stored locally for applications in our Chroma simulations, which will be described in the subsequent chapters.




Figure 3.8: Transmission probability for secondary photons within Si to enter into LXe. The data used for this map comes from our collaborators at TRIUMF and was created as part of a complimentary study of SiPM ExCT in vacuum [115]. Photons in the ExCT simulation are subject to a transmission probability check, with a probability determined by their incident angle and wavelength.

The 2D distribution of transmission probabilities for secondary photon emission highlights the critical angle across the wavelength range ($\sim 20^\circ$), where total internal reflection dominates. Totally internally reflected photons are killed in simulation to avoid double counting direct cross-talk contributions, which are applied at the end of the event during photon hit summation.




Figure 3.9: Secondary photon distribution for ExCT into LXe from Geant4 simulation. This distribution is normalized to the total number of emitted photons in a forward hemisphere, and is the resulting distribution from the ExCT Geant4 simulation process. Figure reproduced from [116].

The emitted photon distribution in Figure 3.9 has a broad emission profile, centred at 45° with respect to the SiPM surface normal. This differs significantly from an isotropic profile, due to the strong refractive index change from silicon into LXe, and the effect of total internal reflection.

The ExCT simulation is directly embedded as a physical process within our `l0lx-sim` Geant4 simulation and is used in conjunction with the existing optical photon processes. SiPM ExCT simulations can be enabled by standard Geant4 macro commands, implemented through custom messenger classes. Additionally, macro commands allow for toggling which SiPMs have the process enabled, the mean emission intensity for ExCT avalanches, and alternative emission profiles. The ExCT process can be enabled for nominal physics simulations, and ExCT correlated hits are flagged and collected separately from primary photon detections to aid analyses. In the following subsection we will discuss the detailed SiPM optics simula-

tion, which was developed in collaboration with TRIUMF PhD candidate Austin de St. Croix as a necessary extension to the simplified Geant4 optics for simulation accuracy.

3.3.1 Simulating SiPM Surface Optics

The optical simulation for photons that reach SiPM surfaces submerged in LXe is a key consideration in this analysis. The SiPM PDE and reflectivity are coupled directly to the ExCT measurement described here, so we must employ a more sophisticated surface simulation. This method provides the expected coupling to the LXe medium using available optical data (n, k , for wavelength and angular dependence). This depends on using additional directional surfaces to scale the total reflectivity to align with existing literature data for vacuum and LXe measurements of SiPMs surfaces [132], [133]. A directional surface in Geant4 defines different properties for incoming and outgoing optical photons. We use a so-called “internal” SiPM PDE, defined as the net PDE after factoring out optical transport effects, to agree with the HPK datasheet [134]. This simulation method, with optical surfaces, is illustrated in Figure 3.10 reproduced from [116].




Figure 3.10: Cross-section illustration of HPK SiPM surface for optical simulation. A degenerate (overlapping) volume of LXe, labelled “LXe Buffer” is defined above the SiPM surface to circumvent limitations in Geant4 optics. This volume has directional surface properties, diffusely reflecting photons on the ingoing surface, and absorbing excess reflected photons on the outgoing surface. The absorption (A) is used to match Geant4 optics to measured reflectivities. Overall SiPM efficiency is matched using an internal PDE that depends on transmission into the SiPM. Figure reproduced from [116].

Before photons reach the physical SiPM, incoming photons are reflected diffusely from the first directional surface. The surface, referred to as the “LXe” buffer, is a degenerate volume of LXe containing the SiPMs, within the mother volume of LXe. Incoming photons are reflected off the top with a wavelength independent $R \approx 18\%$ [133]. We attribute the diffuse reflectivity to SiPM surface microstructures (resistors, trenches, etc...), which are not reflected in the simplified Geant4 SiPM model. The SiPMs are simulated as a flat stack of silicon-SiO₂-medium, where in our circumstances the “medium” is LXe, while reference measurements are primarily within vacuum. Fresnel ray optics are applied at each interface, with input n and k data from [135], [136]. Some datasets are incomplete across the entire wavelength range, so a combination of measurements is used where necessary. The total combined specular reflectivity of the stack is captured in R_{Fr} and is applied at the SiO₂ surface as shown in Figure 3.10.

Using ray optics neglects some possible thin-film effects, while the shape of R_{Fr} (see Figure 3.6) follows that of vacuum reflectivity, with a scaling factor. The higher specular reflectivity can be partially attributed to the lack of insensitive, non-specular SiPM surfaces. We reduce the overall specular reflectivity in simulations using another directional optical surface. This surface is the outgoing LXe buffer top surface, shown in Figure 3.10. The overall specular reflectivity is reduced by adding an absorption probability, A , to this outgoing directional surface. The effective parameter A is tuned to set the overall reflectivity of the SiPM, and is physically related to absorption on microstructure, surface absorption on SiO₂ and non-sensitive surfaces.

The transmission probability for an incident photon is proportional to the diffuse and specular reflectivity as $(1-R_d)(1-R_{Fr})$. The PDE data from Hamamatsu, PDE_{HPK} , was normalized by the transmission probability to produce the “internal efficiency”, $iPDE$, with R_{Fr} from vacuum data with normal incidence to match the experimental conditions of the measurement. The internal efficiency is not angular dependent.

We have defined $iPDE$ as:

$$iPDE = \frac{PDE_{HPK}}{(1 - R_D)(1 - R_{Fr})}. \quad (3.2)$$

The effective specular reflectivity is prepared similarly. For nominal reflectivity, the absorption term is tuned such that the effective specular reflectivity, $R_{Fr}(1-R_D)(1-A)$, is equal to the vacuum reflectivity from [133] (R_{lit} , using R_{Fr} from vacuum data). The nominal absorption probability is given for the range of data in [133] (120—280 nm), with A fixed to 0.45 for longer wavelengths,

$$A_{nom} = 1 - \frac{R_{lit}}{R_{Fr}(1 - R_d)}. \quad (3.3)$$

In nominal simulations in the ExCT wavelength band (600—1000 nm), this results in a specular reflectivity that is fairly flat, with $\sim 10\%$ in vacuum and $\sim 15\%$ in LXe. This was taken as the nominal value for simulations, and as the lower bound on SiPM reflectivity when eval-

uating simulation systematic uncertainties.

We also estimated a systematic upper-limit for reflectivity in the ExCT wavelength band. For this, we estimated that the diffuse reflectivity is attributed entirely to SiPM microstructure (proportional to 1 minus the fill factor, FF), and the active area was assumed to be purely specular with R_{Fr} . We used these assumptions to rescale the effective specular reflectivity ($R_{Fr}(1-R_D)(1-A)$) and rearranged for A .

This resulted in an absorption probability, A , of 0.27. When taken together, this resulted in a specular reflectivity in LXe at LXe scintillation wavelengths of $\sim 33\%$. We see that this estimate is in good agreement with the values of specular reflectivity in LXe of 28-34 % from [133]. The nominal, upper bound systematic, [132] and Geant4 Fresnel reflectivity curves are shown in the bottom panel of Figure 3.6.

Using this methodology, we determine an upper limit on LXe reflectivity that agrees with measurements at LXe wavelengths [132], and a nominal reflectivity that agrees with [133] in vacuum. In all cases, the simulated PDE matches the data-sheets provided by HPK at normal incidence. We do not include any shadowing or angular dependence effects from SiPM surface structure, however the shadowing of the geometry in LoLX and VUV-Q-4 packaging restricts viewing angles to $\sim 65^\circ$ (with respect to surface normal), which we expect to be dominant over surface structures.

3.3.2 Simulation ExCT Analysis

With a detailed simulation model in hand, we can move onto simulations of SiPM ExCT in the LoLX detector. The primary goal of these simulations was to estimate the combined transportation and detection probability of SiPM secondary photon emission in the LoLX phase-1 detector geometry, with conservative uncertainties that reflect the input optical data.

To differentiate between the forward and backward ExCT processes (see Figure 3.1), two distinct simulations are required. For each of the simulations, we simulated mono-energetic

beta decays at the centre of the detector volume to seed ExCT processes.

In the first simulation, SiPM ExCT emission is artificially restricted to the unfiltered SiPMs, which produces a map of the combined probability from unfiltered→LP channels. Conversely, in the second simulation, emission is restricted to the LP filtered SiPMs, producing a map of the combined probability from LP→unfiltered channels.

We refer to these combined transportation and detection probabilities as $P_{\rightarrow}$ and $P_{\leftarrow}$, for the forward and backward processes, respectively. They are single-valued parameters, as the entire ExCT emission spectrum (Figure 3.6, top-panel) is sampled for each photon from a simulated avalanche. Since we have the Monte-Carlo (MC) truth information about the creator process of every detected photon, the combined ExCT transport and detection probabilities are the ratio of ExCT photons in the detector set of channels, to the number of emitted ExCT photons on the emission source channels, which is recorded separately for each simulation, represented as:

$$P_{\rightarrow} = \frac{1}{N_{\gamma}} \sum_{i \in LP} N_{ExCT,i} \quad P_{\leftarrow} = \frac{1}{N_{\gamma}} \sum_{i \in Bare} N_{ExCT,i}, \quad (3.4)$$

where N_{γ} is the total number of emitted ExCT photons in the simulation. We sum over channels, i , in the bare and LP filtered groups for the number of detected ExCT hits (N_{ExCT}) individually. The combined transport and detection probabilities in either direction, $P_{\rightarrow}$ and $P_{\leftarrow}$, are averaged over wavelength and individual channels. These are highly detector-specific quantities, and they are not sensitive to the energy or incident particle type simulated. Since the backward simulation allows for ExCT emission from all LP channels, it will contain the contribution of second and higher order ExCT. Second order ExCT in the backward direction occurs when one LP ExCT detection produces more secondary emission, leading to another LP ExCT detection. At sufficiently high emission intensities this can induce a runaway effect, however, in real detector systems SiPMs are never operated with the overvoltage required to reach such intense emissions.

We explicitly do not include SiPM DiCT in these simulations. DiCT enhances ExCT

by producing additional SPAD avalanches, leading to more secondary photon emissions en masse. This omission is supported by the analysis method described in this chapter, where only SPE avalanches are selected, and the model introduced in the following section takes this into account.

3.4 Modelling ExCT Coincidence Probability in LoLX

In order to estimate the number of emitted secondary photons per avalanche from SiPMs in the LXe environment, an effective modelling framework was developed. The model encodes the forward and backward contributions to the correlated probability, $P_{corr,B}$, measured from LoLX data, and ultimately estimate the number of emitted photons, N_{Xe} . This parameter is especially important, as it is detector independent, and can be used to inform correlated noise studies for other detector configurations.

The strategy is then, to derive an analytical expression for $P_{corr,B}$, that depends on the emission yield, N_{Xe} , the simulated transport and detection efficiencies in both directions, $P_{\rightarrow}$ and $P_{\leftarrow}$, alongside other detector specific parameters. The final relationship depends on the rate of candidate pulses in the LP and unfiltered channel groups, also a function of N_{Xe} , which we represent analytically. This produces a pair of nested relations, both dependent on N_{Xe} , which are solved numerically.

First, we define the probability for a single successful correlated SiPM ExCT hit in the forward direction, P_{Forw} , from secondary photon emission. This is defined as a weighted sum over the Poisson distribution for an avalanche generating a mean number of photons, N_{Xe} , and each term (photon) has a binomial probability of being successfully transported and detected:

$$P_{Forw} = \sum_{N_{\gamma}=1}^{\infty} \frac{N_{Xe}^{N_{\gamma}} e^{-N_{Xe}}}{N_{\gamma}!} \cdot N_{\gamma} P_{\rightarrow} (1 - P_{\rightarrow})^{N_{\gamma}-1} \approx N_{Xe} P_{\rightarrow}. \quad (3.5)$$

For the backward direction, P_{Back} , we replace the simulated probability to the equivalent term, $P_{\leftarrow}$ for the resulting expression, $P_{Back} \approx N_{Xe} P_{\leftarrow}$. The simplifying approximation here is

correct to within $\sim 3\%$ for the low values of N_{Xe} expected (Figure 3.11).




Figure 3.11: Comparison of full summation versus approximate relation for ExCT probabilities. Using the simulated forward combined transport and detection probability of 4.1 %, at the estimated N_{Xe} level around ~ 0.5 photons, produces a relative difference of $\sim 3\%$ between the two relations.

The relationship for $P_{corr,B}$ in Equation 3.1, is normalized to the number of all candidate SPE pulses in the unfiltered channels only, while the numerator includes all successful ExCT coincidences. In this case, we are explicitly accounting for all trials in the forward (unfiltered $\rightarrow$ LP) direction, but only successful trials in the backward (LP $\rightarrow$ unfiltered) direction. To account for this effect, we have introduced two additional terms. $N_{1,LP}$ and $N_{1,B}$, the number of seed SPE pulses in the unfiltered (bare) and LP filtered channels without ExCT contributions, represented as:

$$P_{corr,B} = \frac{(P_{Forw}N_{1,B} + P_{Back}N_{1,LP})(1 - P_{DiCT})}{N_{1,B} + P_{Back}N_{1,LP}(1 - P_{DiCT})}. \quad (3.6)$$

This relationship is simplified by introducing the ratio of $N_{1,LP}$ and $N_{1,B}$, as a ratio of rates, $R_R \equiv R_{1,LP}/R_{1,B} = N_{1,LP}/N_{1,B}$. However, the observed rate ratio is influenced by the presence of ExCT between the LP and unfiltered groups. We can correct for this by representing

the observed rate ratio as a function of the unperturbed ratio, with the confounding effect of the ExCT process. This is:

$$\widetilde{R}_R = \frac{R_{1,LP} + R_{1,B}P_{Forw}(1 - P_{DiCT})}{R_{1,B} + R_{1,LP}P_{Back}(1 - P_{DiCT})}, \quad (3.7)$$

and if we rearrange and use our simplifying approximation from Equation 3.5, we arrive at the relationship:

$$R_R = \frac{\widetilde{R}_R - P_{Forw}(1 - P_{DiCT})}{1 - \widetilde{R}_R P_{Back}(1 - P_{DiCT})} \approx \frac{\widetilde{R}_R - N_{Xe}P_{\rightarrow}(1 - P_{DiCT})}{1 - \widetilde{R}_R N_{Xe}P_{\leftarrow}(1 - P_{DiCT})}. \quad (3.8)$$

While the relationship for $P_{corr,B}$ is simplified to:

$$P_{corr,B} = \frac{(P_{Forw} + P_{Back}R_R)}{\frac{1}{(1-P_{DiCT})} + P_{Back}R_R} \approx \frac{P_{\rightarrow} + R_R P_{\leftarrow}}{\frac{1}{N_{Xe}(1-P_{DiCT})} + R_R P_{\leftarrow}}. \quad (3.9)$$

Equations 3.8 and 3.9 form a pair of relations, both functions of N_{Xe} that are solved simultaneously using numerical methods. The effect of ExCT on the measured rate ratio is shown in Figure 3.12, for a range of N_{Xe} values.




Figure 3.12: Plot of unperturbed channel rate (R_R) ratio versus measured rate ratios ($\widetilde{R}_R$) for a range of N_{Xe} values. Increasing the emission of secondary photons increases the measured ratio of SPE pulse rates between LP and unfiltered channels, due to the larger sensitive area of the LP group.

Using simulations, we estimated the directional transport and detection probabilities, $P_{\rightarrow}$

and $P_{\leftarrow}$. The observed rate ratio, $\widetilde{R}_R$, is measured from data, and P_{DiCT} from literature [101]. The relations described here are solved numerically to determine the value of N_{Xe} , that is in agreement with the observed $P_{corr,B}$. This estimation method can be applied to arbitrary detector systems, given that estimates of $P_{\rightarrow}$ and $P_{\leftarrow}$ are available.

3.5 Results

In this section, we will present the results of the data and simulation analysis, and their comparison with the effective model of ExCT described in the previous section. Afterwards, we will discuss the main systematic considerations for this analysis, and their impact on our main observables, the correlated ExCT probability, $P_{corr,B}$, and the estimate of the secondary photon emission yield in LXe, N_{Xe} .

We estimated $P_{corr,B}$ for both OV datasets using Equation 3.1, with the fitting procedure described earlier. The excess in the Δt distribution, about $\Delta t = 0$, above the combinatoric background of uncorrelated pulses is ascribed to SiPM ExCT between the unfiltered and LP channel groups, consisting of “forward” and “backward” correlations (see Figure 3.1). The excess is normalized to the number of pulses in the unfiltered channels, N_B . For the 5 V dataset, the correlated probability is, $P_{corr,B} = 5.2\%$, for 4 V OV, this probability is, $P_{corr,B} = 3.8\%$. The Δt distribution for the 4 V dataset is shown in Figure 3.13, for comparison to Figure 3.5.




Figure 3.13: 4 V OV dataset delta-time distribution, with linear fit overlaid. The fit function shown is the linear background component of the combined signal+background fit, with residuals showing random fluctuations in background regions, and a signal excess in the correlation region about $\Delta t = 0$ as expected from SiPM ExCT.

From the functional dependence in Equation 3.9 we can see that $P_{corr,B}$, is not directly proportional to ExCT between the two groups. For large R_R , corresponding to a very high SPE rate in the LP channels compared to the unfiltered, this will exceed 1 even for small values of N_{Xe} . The linear falling background in Figure 3.13 and Figure 3.5 originates from stray LXe scintillation light, very late afterpulsing, dark noise, and predominantly the slow luminescence of the 3D printed cage structure. Measurements of the cage luminescence under UV excitation at room temperature were carried out ex-situ at McGill in 2022, the emission spectrum is shown in Figure 3.6. In-situ measurements of the photoluminescent lifetimes from the cage showed a dominant prompt component that could not be resolved from the prompt LXe emission, and a delayed component with a lifetime of 1026 ± 88 ns [117]. We expect that the low intensity delayed photoluminescence constitutes the majority of seed photons in the ExCT measurement.

Table 3.1: Systematic uncertainties considered, based on data and simulations. Simulation uncertainties are evaluated by comparison to a reference simulation with nominal inputs, and the combined 4 V dataset for data-driven uncertainties. The symbol ‘–’ indicates negligible uncertainties. First block corresponds to simulation uncertainties on $P_{\rightarrow/\leftarrow}$, the second block corresponds to the data analysis systematics on $P_{corr,B}$

Name	Nominal Value(s)	Variation	Relative Uncertainty (%) on $P_{\rightarrow/\leftarrow}$ or $P_{corr,B}$
NIR SiPM PDE	~5% (See Figure 3.6)	+10%, -30% relative	+9, -29 ($\rightarrow$ / $\leftarrow$)
NIR LP Transmittance	~90% (See Figure 3.6)	$\pm 10\%$ relative	± 4.4 ($\rightarrow$), ± 0.5 ($\leftarrow$)
Cage Reflectivity	~95% (Varies, calculated by Geant4 from $n(\lambda)$)	Set to 50% for all λ (-45 % absolute)	-30 ($\rightarrow$), -40 ($\leftarrow$)
SiPM Reflectivity	~15% (See Figure 3.6)	Upper limit from vacuum data (-5% absolute)	-3 ($\rightarrow$), – ($\leftarrow$)
Total Simulation Uncertainty			+11, -40 ($\rightarrow$) +9, -49 ($\leftarrow$)
Pile-up Algorithm Threshold	>3 photon equivalent	± 1 photon equivalent	+12, -7 ($P_{corr,B}$)
Low Occupancy Window Selection	[800,1500] ns	± 200 ns	–
Δt fit range	[-200,250] ns	± 80 ns	–
Total Data Uncertainty			+12, -7

3.5.1 Systematic Study

We investigated the main systematic uncertainties in the measurement of LXe secondary photon emission yield, N_{Xe} . All systematics considered, including negligible contributions, are summarized in Table 3.1. Systematics on the measurement from LoLX data, impacting $P_{corr,B}$, were subdominant compared to the simulation uncertainties, impacting $P_{\rightarrow}$ and $P_{\leftarrow}$. We also include uncertainties from literature values used in evaluation of the effective model described in the previous section, summarized in Table 3.2. Our primary sources of systematic uncertainties are due to the lack of measurements of detector optical properties, due

to different measurement conditions, or incident wavelength ranges. Simulated systematics are evaluated by re-running the simulations and corresponding analysis described previously, with the described systematic variations, with one systematic per simulation set. The corresponding variations to $P_{\rightarrow}$ and $P_{\leftarrow}$ are captured as asymmetric relative uncertainties with respect to the nominal simulations.

The first key systematic is the simulated SiPM PDE in the NIR range, for which supplementary measurements, beyond the Hamamatsu datasheet, are sparse. We varied the NIR PDE by a relative +10 % and -30 % with respect to the datasheet values. We ran a larger negative variation on the PDE to conservatively capture a lower than expected PDE observed in the Hamamatsu SiPMs in LoLX in the LXe scintillation band [117]. At the applied OV's used in this measurement, the SiPM PDE is approaching saturation, and the nominal 4 V OV SiPM responsivity in the Hamamatsu datasheets is used for all simulations. The SiPM PDE variation is expected to cover any potential increases in SiPM PDE at 5 V compared to the 4 V datasheet values.

The second key systematic originates from the custom 3D printed cage of the LoLX phase-1 detector. The cage is made of a proprietary plastic blend, Formlabs Durable Resin [137], and is stated to be acrylate based. In the simulations of this work, the 3D printed cage is modelled with a refractive index from acrylic plastic, which Geant4 uses to compute reflectivity and transmittance based on the Fresnel relations. To conservatively account for the acrylic approximation in optical response, we vary the absolute reflectivity of the cage by 50 %. This is the largest individual contribution to the total simulation systematic uncertainty. This systematic is simulated by overriding the Fresnel calculations with a fixed 50 % reflectivity of the cage surface at all wavelengths.

The last key simulation systematic relates to the reflectivity of the SiPM surfaces. Using the model described in Section 4.1, we have defined the upper and lower limits on the SiPM reflectivity. Summarizing Section 4.1, the nominal values used in simulations for $P_{\rightarrow}$,

and $P_{\leftarrow}$, representing the lower-limit on SiPM reflectivity reproduces the measured vacuum reflectivities. The upper limit represents a scenario where the entire photosensitive area of the SiPM has a reflectivity derived from Fresnel relations using literature values of the SiO_2 refractive index. This gives agreement with measurements of UV reflectivity in LXe [132]. The nominal and variational reflectivities are shown in Figure 3.6 for reference. Finally, a systematic on the NIR transmittance of the LP filters was investigated by applying a relative variation of $\pm 10\%$ to the nominal data-sheet value of $\sim 90\%$ in the NIR band. The effect of this variation is subdominant compared to the cage reflectivity and SiPM PDE.

Systematic effects from the event selection and candidate selection were considered for their impact on the observable, $P_{\text{corr},B}$. The analysis depends on the low-occupancy assumption stated previously, and so there must be a sufficiently low probability of observing two uncorrelated pulses in accidental coincidence. This is needed in order to resolve the small contribution of correlated pulses from ExCT. This correlation is restricted to the small window about $\Delta t = 0$, encompassing $[-24 \text{ ns}, 24 \text{ ns}]$. We have optimized the event and candidate selection to remove any sources of truly coincident light late in the event window. One potential source of coincidence is pile-up, from the ^{90}Sr source, which had an activity of $\sim 330 \text{ Bq}$ at the time of data taking. If we assume secular equilibrium between ^{90}Sr and ^{90}Y , we can estimate the pile-up probability in the $3 \mu\text{s}$ DAQ window. From this we expect pile-up in ~ 1 in 1000 events. This contribution is small compared to the dominant source of false late time coincidences from radio frequency interference (RFI). RFI was present in significant quantities for the entirety of the data collection period used in this analysis, and was approximately coherent across all channels when observed. The most troublesome RFI spikes were $\sim \text{SPE}$ level wavelets that occurred randomly through the event window, and when observed in the low-occupancy search window, can contribute to false ExCT coincidences. To prevent this contribution, a pile-up rejection algorithm was implemented. The algorithm searches for trigger candidates in the summed and aligned waveform from all channels by evaluating

the waveform against the trigger condition and identifying crossing points. Coherent RFI pickup across multiple channels late in the event window was strongly suppressed by this algorithm, as events with more than one trigger candidate are excluded from the event selection for this analysis. To estimate a systematic effect from missed false coincidences, we varied the threshold of the pile-up algorithm by ± 1 SPE equivalent ADC units. This resulted in a ~ 10 % change to the measured $P_{corr,B}$ values. We include this systematic as an asymmetric uncertainty in the final results.

Lastly, we investigated several other sources of systematic uncertainty for their impact on the observables reported here. The linear uncorrelated background fitting range was varied by up to 100 ns with no effect. Other negligible sources of systematic uncertainty include: the location and length of the low-occupancy window, the energy of preceding trigger events, and temperature variations within a run. The fit range of the Gaussian peak was varied by up to 80 ns with miniscule impact, subdominant by a factor of 10 to the main sources of systematic uncertainty and is not considered in the final evaluation.

3.5.2 ExCT probabilities and secondary emission yields

Table 3.2: Input parameters for ExCT model with 1-sigma uncertainties considered in the analysis. $\widetilde{R}_R$ is the ratio of the rates of candidate pulses between the LP and unfiltered channels. P_{DiCT} is the probability of internal cross talk. $P_{\rightarrow}$ and $P_{\leftarrow}$ are the forward (bare $\rightarrow$ LP) and backward (bare $\leftarrow$ LP) transport efficiencies, respectively.

Input Parameter	4 V Value	5 V Value	Source
$\widetilde{R}_R$	6.3 ± 0.9	7.7 ± 1.2	Measured
P_{DiCT}	5.0 ± 0.5 %	7.0 ± 0.7 %	[101]
$P_{\rightarrow}$	$4.1^{+0.5}_{-1.7}$ %	Same as 4 V	Evaluated from simulation
$P_{\leftarrow}$	$0.6^{+0.1}_{-0.3}$ %	Same as 4 V	Evaluated from simulation

To estimate the key observable of this work, the secondary emission yield in LXe (N_{Xe}), we combined the experimentally measured parameters, simulation inputs, and literature inputs, detailed in Table 3.2, and Table 3.2. Following the methodology described in Section 3.4,

we solve Equation 3.9, and Equation 3.8, together numerically to estimate the value of N_{Xe} that agrees with the measured $P_{corr,B}$. The effective model of $P_{corr,B}$ and the measured values for the 4 V and 5 V OV datasets are overlaid in Figure 3.14. The intersection of the central value of the experimentally observed value of $P_{corr,B}$ and the model curve at the corresponding OV represents the best-fit value of N_{Xe} for each measurement, respectively. The observed secondary emission yield in LXe, N_{Xe} , which is the average number of photons emitted into LXe per avalanche, is estimated to be 0.51 for 4 V and 0.64 for 5 V OV. Summaries of measured values and their respective uncertainties are included in Table 3.3 for reference.

The total uncertainties (combined statistical and systematic) on data and the effective models are shown as shaded bands in Figure 3.14, representing the $1-\sigma$ uncertainty regions. All non-negligible systematic sources of uncertainty were added in quadrature, with total uncertainties and individual contributions to $P_{\rightarrow}$ and $P_{\leftarrow}$ summarized in Table 3.1. The uncertainty on $\widetilde{R}_R$ was determined from the variance observed across individual runs, and the uncertainty on the literature input, P_{DiCT} , was taken from the reported measured uncertainties in [101].

We also estimated the combined uncertainty arising from the effective model and data-driven uncertainties including statistical uncertainty. This was done numerically by adding in quadrature the observed differences from evaluating the effective model at the upper and lower bounds of the measured data values, together with the differences from evaluating the central data values with the upper and lower bounds of the effective model uncertainty bands.

Earlier in this chapter, we discussed the difference between the measured correlation probability, $P_{corr,B}$, and the so-called “true” ExCT probabilities in either direction, P_{Forw} , and P_{Back} . Since the simulated transport and detection probabilities, $P_{\rightarrow}$, and $P_{\leftarrow}$, preserve the angular and wavelength character of the ExCT process, we use these to estimate the true probabilities for ExCT in LoLX. This is represented by Equation 3.5 in the previous section, when we substitute our measured values of N_{Xe} , and $P_{\rightarrow}$, the forward probability is obtained. In plain terms, this is the probability that an avalanche in the unfiltered SiPM group will cause

a correlated ExCT secondary avalanche in any LP filtered SiPM. This value is 2.1 % and 2.6 % at 4 V and 5 V OV respectively. As the simulations of $P_{\rightarrow}$ and $P_{\leftarrow}$ are highly detector geometry dependent, these estimates are specific to the LoLX phase-1 detector in LXe, while the measurement of N_{Xe} may be used to inform other experimental settings and apparatuses.

Lastly, we estimated the secondary photon yield within the bulk silicon of the SiPM. As stated previously, the vast majority of secondary photons are not expected to leave the SiPM due to the large refractive index mismatch (~ 3.78 vs ~ 1.46 for Si and SiO_2 , at 700 nm for reference) resulting in mostly internal reflections. The fraction of internally reflected photons under the isotropic emission approximation is 94.7 % in the outgoing hemisphere, or 97.4 % in 4π , evaluated using the simulation method described in Section 3.3, averaged over wavelength and emission angles. To estimate this so-called “internal secondary photon yield” we divide N_{Xe} , the observed estimate of secondary photons in LXe, by the simulation-driven estimate of the escape fraction in 4π . This results in an internal silicon secondary photon yield of 20 and 25 photons/avalanche on average for 4 V and 5 V OV, respectively. Just as in the secondary LXe photon yield, the secondary internal yield may be used for modelling the effect of SiPM ExCT in different media, such as liquid argon, for example.




Figure 3.14: LoLX SiPM ExCT comparison of effective ExCT model and data. Using the intersection of the data bands (green/blue) with the model of ExCT at the respective over-voltages (pink/orange), the LXe secondary photon emission yield, N_{Xe} is estimated. Shaded regions represent the $1\text{-}\sigma$ confidence intervals, with statistical and systematic uncertainties added in quadrature. Figure reproduced from [116].

Table 3.3: External cross-talk probabilities from the time correlation analysis, and the resulting mean number of photons emitted per avalanche within silicon and into LXe.

Overvoltage	Observable $P_{corr,B}$ (%)	ExCT probability (P_{Forw}) (%)	Emission within Si ($\gamma/\text{Av.}$)	Emission into LXe, N_{Xe} ($\gamma/\text{Av.}$)
4 V	3.9 ± 0.2 (stat.) $^{+0.5}_{-0.3}$ (sys.)	$2.1^{+1.2}_{-1.3}$	20^{+11}_{-9}	$0.5^{+0.3}_{-0.2}$
5 V	5.2 ± 0.2 (stat.) $^{+0.6}_{-0.4}$ (sys.)	2.6 ± 1.4	25^{+12}_{-9}	$0.6^{+0.3}_{-0.2}$

3.5.3 Discussion

We have measured ExCT between SiPMs submerged in LXe using the LoLX phase-1 detector. These measurements were done at two operating voltages, under nominal detector conditions, with a standard triggering scheme. The measurement methodology developed here is novel

and can be replicated by other experiments aiming to characterize SiPM ExCT in their systems. To better understand the impact of SiPM ExCT, and the effect of optical transport effects on it, we developed a dedicated Geant4 model of ExCT emission using custom physics processes. To separate the two contributing components of SiPM ExCT correlations measured using our technique, so-called “forward” and “backward” correlations, we developed a model of our expected signal using a combination of measured, simulated and literature data inputs. By combining these contributions, we estimate the secondary photon emission yield in LXe, N_{Xe} , a detector independent parameter that can be compared to other measurements. The measured emission intensities, reported in Table 3.3, show the expected gain-dependent relationship of SiPM ExCT. The measured intensity increased by ~ 1.25 , when increasing the overvoltage on all SiPMs from 4 V to 5 V, a 25 % increase.

The measured secondary emission intensity in LXe, and the corresponding estimate internal emission intensity (N_{Si}) measured here are in tension with the intensities reported by Hamamatsu for VUV4 SiPMs [115]. In [115], the reported yield at 4 V OV is 48.5 ± 9.5 ph/avalanche, with an estimated 1.55 ± 0.62 photons emitted into LXe. Compared to 20_{-9}^{+11} in silicon, and $0.51_{-0.24}^{+0.29}$ in LXe, a factor of ~ 3 difference. The two main differences between our measurement and the TRIUMF measurement are angular acceptance and measurement conditions. As introduced at the start of the chapter, LoLX phase-1 is sampling a much larger solid angle than the limited FOV microscope measurement (NA=0.45) in [115], and importantly, it is a measurement in vacuum at cryogenic temperatures and not in LXe. When correcting for optical effects, we assume a perfectly flat surface of the SiPM, ignoring surface and microstructures within the SPAD, and any absorption of photons during emission. These differences may lead to an over estimation of the large-angle emission profile here, when extrapolating from ex-situ measurements. Another source of discrepancy could arise from the assumption of isotropic emission profile in silicon used in both works. The LoLX phase-1 detector is sensitive to a much broader angular acceptance, and is natively sensitive to the

effects of being submerged in LXe, further efforts are required to isolate sources of tension. Follow-up measurements will constrain the angular emission profile using angular correlations between different SiPMs in the LoLX phase-2 detector, with larger statistics. The data taken for this measurement was from the last operational run of LoLX phase-1, and was statistically limited when investigating angularly resolved SiPM ExCT measurements. The main systematic limitations of this analysis originated from the lack of measurements of the necessary optical properties in LXe operating conditions, motivating a strategic need for in-LXe optical property measurements of detector-relevant materials.

In Chapter 6, we evaluate the impact of ExCT on the nEXO detector, building on the simulations developed here. Using our measurements of SiPM ExCT in LXe, we find that SiPM ExCT constitutes a percent-level and non-negligible contribution to correlated noise in a detector with photocoverage comparable to that of nEXO. We then examine the impact of SiPM ExCT on the energy resolution of nEXO, and contribution to other analyses.

4

LoLX Phase-2 Commissioning

The LoLX phase-2 detector is a significant change from the phase-1 detector employed for the measurement of SiPM ExCT in LXe. With this upgraded detector design, we addressed a number of challenging design elements of the phase-1 detector. By removing superfluous materials, selecting non-polymer based components for the structure, and including the ability to place optical filters over each tile/face, we have reduced significantly the presence of non-VUV fluorescence light compared to the first phase of the experiment. With a reference PMT installed and two varieties of VUV sensitive SiPMs, we are able to isolate and validate the SiPM relative efficiency to VUV light through cross-referencing. With substantial upgrades to detector monitoring, cryogenics, and autonomous operations we have prepared the system for future long-term behaviour monitoring, an essential validation step before deploying $\sim 46k$

SiPMs in the nEXO detector for 10 years of operation.

These upgrades reflect the primary goals of the second phase of LoLX experimentation: detailed characterization of SiPM performance and light transport in LXe, long-term performance and efficiency studies in LXe for nEXO photodetector candidate SiPMs, and the measurement and characterization of Cherenkov radiation in LXe for future background rejection capability studies, with additional applications in medical physics [113]. This chapter describes the commissioning data runs taken to-date, with preliminary MC and data analysis comparisons for the January 2025 commissioning run. We describe the external calibration procedure and its limitations, and conclude with an overview of an internal calibration source using ^{181}Hf produced through neutron activation of natural hafnium that has been prepared for future LoLX runs.

4.1 LXe Commissioning Runs

4.1.1 Run Overviews

The LoLX phase-2 detector was installed in the upgraded LoLX cryostat in June 2023, after assembly and preliminary bench testing of SiPMs in May 2023. The first commissioning run occurred in August 2023, with initial cooldown on August 15th, and LXe recovery on August 19th.

In the initial LXe commissioning run, we encountered significant challenges originating from the custom UHV coaxial cables. All cables from the tiles to the feedthrough of the cryostat were custom-made at TRIUMF and McGill, with only crimped connections. Several channels which behaved normally at room temperature lost connections after cooldown, due to poor contact between intermediate connections.

During the first commissioning run, data was taken with a fixed PMT voltage of 750 V, and SiPMs were operated at 3 V, 4 V and 5 V OV. Before external source calibration data-taking,

internal calibration is carried out with a pulse diode laser. Laser illumination is used to provide sufficient photocurrent for breakdown voltage measurements from reverse IV curves, and low-intensity laser calibration is used to perform SPE calibration of the SiPMs and PMTs. Detector calibration was performed with three different γ -sources, placed on the outside of the cryostat, at approximately the same height as the LoLX phase-2 detector. These sources were selected for their intermediate energies, as < 100 keV sources have insufficient energy to penetrate the stainless steel, copper and LXe layers, and > 1 MeV sources experience much higher rates of Compton scattering in LXe, complicating event reconstruction considerably. The three sources used were ^{133}Ba , ^{22}Na , and ^{137}Cs , with dominant γ -line energies of approximately 356 keV, 511 keV, and 661 keV, respectively. Most data were taken while triggering on a PMT signal with > 3 photons equivalent in amplitude, with some sample datasets taken with global OR trigger on SiPM signals with varying amplitude thresholds. Several hundred thousand events were taken for each source, and longer background-only data were taken overnight to reach equivalent event counts for background subtraction.

The second commissioning run took place in January 2025, after interventions and modifications during October and November 2024. During this intervention period, we repaired several connectors, and replaced a non-functioning HPK VUV-4-Q package with a windowless replacement from a separate production batch. With the comparison to a windowless device we aimed to decouple the in-LXe performance differences for the HPK SiPMs from any batch-specific deficiencies, or unexpected absorption of VUV light from outgassing components freezing out on the quartz windows. Additionally, a prototype LXe capacitive level sensor was installed in the cryostat, alongside a LXe purity monitor using electron drift lifetimes [138]. For the second commissioning run, we installed a remote controllable high voltage PMT power supply to allow for autonomous PMT control, which was limited in the first commissioning run due to the use of a bench top analog HV supply.

In the second run we prioritized ^{133}Ba with 4 V OV to increase statistics significantly

over the data taken during the first commissioning run. During runs with a ^{133}Ba calibration source, approximately $\sim 30\text{M}$ events were taken over four days, compared to the $\text{O}(500\text{k})$ in the first commissioning run. The focus on barium was motivated by the detector performance evaluated from the first commissioning run, where the ^{22}Na , and ^{137}Cs runs had a large proportion of saturated and clipped events, making event reconstruction challenging. In the January 2025 run, we had a total of twelve unplugged SiPM channels. Of these, five were non-functioning FBK SiPMs, two were non-functioning HPK SiPMs, and the remaining five were unplugged due to a non-functioning SiPM present in the same summation group for readout. We unplugged functioning channels in the same summation group as non-functioning channels because of excess noise arising from electrical cross-talk. Of the non-functioning channels, only one FBK SiPM response was consistent with a completely dead device. The remaining six non-functioning channels all showed signatures of poor connectivity from the readout cabling chain. These signatures include an unreasonable photovoltaic voltage for SiPMs, a significantly lower readout capacitance than expected, and DC current in reverse IV scans, consistent with an open electrical circuit before the SiPM device.

As described in Section 2.5, the analysis techniques used for pulse-finding and event reconstruction were upgraded substantially over the previous implementations used in LoLX phase-1. Using the commissioning data, we will show the results of this upgraded analysis for SiPM calibration, and its impact on energy reconstruction for simulation-data comparisons with our combined Geant4 $\rightarrow$ Chroma simulation pipeline.

To further support analysis efforts, we created an offline event display for LoLX phase-2 data inspection. An example event is shown in Figure 4.1, from a 4 V ^{133}Ba calibration run. This tool is essential for standardizing data quality checks, where questionable sub-populations of events can be flagged by event and run number for manual post-inspection. Problematic events can include clipped waveforms from saturated detector events, significant pre-trigger artifacts from preceding events, and large light asymmetries in SiPM response

from surface-events among other causes.




Figure 4.1: LoLX phase-2 detector offline event-display example event. Left: Unwrapped LoLX phase-2 detector cube with FBK channels highlighted in red, and unplugged (non-functional) channels hatched out. The estimated PE level for each channel is overlaid in text, with a single-colour saturating scaling for relative PE scale within an event. Right: Normalized stacked raw waveform display plot (polarity flipped).

4.1.2 Results

SiPM Calibration

The first stage of data-taking for LoLX phase-2 is for SiPM calibration, needed to balance the gains of the sensors before physics data taking. Since the detection efficiency and readout signal size are both strongly coupled to the applied OV, it is essential to match OVs across all channels for data collection. This is a required step to preserve the symmetry of the detector response. We begin by taking laser calibration runs for SiPM breakdown voltage with high illumination to produce sufficient current while devices are at LXe temperatures.

We reverse bias every SiPM individually at fixed voltages, while reading back the current from the built-in current readout capability of the front-end nanoPI board's. This is done on the board level, where each board has 16 SiPMs from a single vendor type (from two faces of

the detector). We do not mix SiPM types on readout boards, because the maximum voltage difference between channels for a board is 4 V, while HPK and FBK SiPMs have breakdown voltages that differ by more than that. The reverse-IV curves for a selection of FBK SiPMs on readout board #03 are shown in Figure 4.2. We overlay the curves at R.T for comparison, as well as the fitted breakdown voltage (BDV) for each channel at LXe temperatures. We see that in future calibration runs, we must take finer voltage steps during the breakdown region to capture the full shape, which will improve the stability of the breakdown voltage estimates.




Figure 4.2: Example of a selection of FBK SiPM channel reverse IV curves at room temperature (RT) in red and in-LXe at 170 K in blue. Overlaid are the calculated breakdown voltages for in-LXe measurements using a pulse laser diode. We observed significant scatter in estimated BDV compared to the underlying IV curves, signalling a need for finer voltage increments in the breakdown voltage region for future measurements.

Once we have determined the breakdown voltages for every SiPM, we must calibrate the gain of the devices at the applied overvoltage. This is required to convert from raw estimates of pulse charge and amplitude from pulse finding into estimates of PE, and ultimately detected photons. To calibrate the SPE amplitude and charge for the devices, we take additional low-

intensity laser calibration runs where each channel should receive a few photons on average per event. Due to the geometry of the detector, the bottom tile SiPMs receive much more illumination, and secondary very-low illumination runs are needed to accurately estimate their SPE gains. An example of the SPE calibration data is shown in Figure 4.3, where we highlight our pulse charge reconstruction estimates.

We encountered a new challenge in LoLX phase-2 detector operations, a result of the lower single-PE SNR described in Section 4.2. The estimate of an individual pulse’s charge, or integral, is highly dependent on the estimated end-position of the pulse. With the observed baseline voltage RMS of $\sim 3\text{-}4$ mV we reach the “noise-floor” before we reach the end of the pulse. This makes estimates of the pulse-end especially difficult, more-so for FBK SiPMs which have a strongly bi-exponential decay shape at cryogenic temperatures (see Figure 2.22 for an example).

To mitigate this effect, we take advantage of our knowledge of the underlying SiPM SPE pulse-shape from template fitting, and we construct a secondary charge estimate to complement the so-called “raw charge”. To estimate SPE-level pulse charges with poor information about the end of the integration window, we instead integrate in a “gate” that begins at the pulse start and ends at a fixed time relative to the pulse start. For FBK the window is $\delta t = 60$ ns. We then normalize our clipped estimate of the pulse charge by the fraction of the pulse template in the gate region, to estimate the full pulse charge. This charge estimate is only used for SPE pulse charge calibration, since for physics data most signal pulses are much larger than the SPE level and the pulse end is clear relative to the peak amplitude and baseline noise. This gate correction only works for pulses coherent in time, and is therefore not appropriate for normalizing pulses from LXe signals, where the photon arrival time is exponential

in nature. The relationship for the two pulse charge estimators are:

$$Q_{raw} = \sum_{i=t_{start}}^{t_{end}} V_i \Delta t, \quad (4.1)$$

$$Q_{gate} = \frac{\sum_{t_{start}}^{t_{start} + \delta t_{gate}} V_i \Delta t}{N_{template}}, \quad (4.2)$$

where V_i , is the signal in waveform bin ‘ i ’, Δt is the bin width, and $N_{template}$, is the normalization factor from integrating the SiPM pulse template (T_{SiPM}) in the gate region and to infinity:

$$N_{template} = \frac{\int_{t_{start}}^{t_{start} + \delta t_{gate}} T_{SiPM}(t) dt}{\int_{t_{start}}^{\infty} T_{SiPM}(t) dt}. \quad (4.3)$$




Figure 4.3: FBK channel SiPM ID#7 SPE calibration plots. Left: Gated pulse charge (blue) and raw charge estimates (orange). Right: Pulse amplitude vs gate charge estimate. We see substantial recovery of pulse charge resolution when correcting gate charge estimates with a template normalization factor.

We see in Figure 4.3 the improvement in pulse charge resolution and accuracy of SPE charge estimates by switching to the gate integration. When we look at the 2D projection versus pulse amplitude (right panel of Figure 4.3), we clearly see the 1 PE, and 2 PE signatures, while the 3 PE signal is smeared due to the presence of afterpulsing and delayed cross-talk,

which are not included in the template for normalization.

For run preparation, we fit 1-D Gaussian functions to every channel for SPE charge and amplitude, and where fits fail to converge because of poor charge resolution or sub-optimal laser illumination we estimate the SPE charge and amplitude based on the 2D distribution, and the mean values from successful fits. For the January 2025 run, we observed 3-5 % relative variance in the HPK SiPM gain, while the FBK channels had up to 10 % relative variance between channels. This points towards the need for longer calibration runs, ideally with improved baseline noise and finer binning in reverse-IV curves for more consistent overvoltages, as these effects are coupled.

The PMT calibration is done using the baseline corrected waveforms collected during low-illumination SiPM calibration data-taking, where the number of PMT pulses on per event average are much less than 1. For each trigger, we integrate in a fixed window near the trigger, where we expect the PMT signals to appear. This produces a large “pedestal”, when there is no signal, and a small SPE signature from ~ 1 SPE PMT pulses. To subtract the “pedestal”, i.e. the mean baseline integration when there are no PMT pulses, we use the same method for a run where the PMT was turned off. This gives us estimates of the PMT SPE amplitude for trigger configuration and charge for signal integration in data analysis.

All SiPM and PMT gains for SPE amplitude and charge are loaded into the configuration for the online/offline MERCI analysis. The estimated charge and amplitude signals are normalized to SPE estimates to provide estimates of light signal intensity using either pulse charge or pulse amplitude.

External Barium-133 Source Data

In the January 2025 commissioning run, our largest dataset was for the external ^{133}Ba calibration run, taken with a PMT trigger, equivalent to approximately 8 photons in amplitude. The data shown here is for the 4 V OV runs, runs 31 - 43.

We begin by looking at the position dependence of interactions in the LoLX phase-2 de-

tector cube. We can project the position dependence by leveraging the cubic symmetry of our detector, i.e. by summing the total signal in each SiPM tile, and the PMT. By taking the weighted mean of the tiles relative to their opposite face we get an estimator for how “symmetric” the light distribution inside the detector is. We define this tile-tile asymmetry as:

$$A_{i-j} = \frac{w_i Q_i - w_j Q_j}{w_i Q_i + w_j Q_j}, \quad (4.4)$$

where $w_{i/j}$ is the weight for a direction pair (e.g. N-S), based on sensitive area for SiPM tiles (taking into account unplugged channels), relative sensitivity and effective area for the PMT in top-bottom calculation. $Q_{i/j}$ is the sum of the signal reconstructed for channels on that face. We have 4 cardinal directions in the X-Y plane (N-S,E-W), and the Z-direction is referred to as “top-bottom”. The distribution of events along the three main axes of the cube are shown in projections in Figure 4.4. We highlight the Top-Bottom asymmetry individually, which shows two main populations of events, one concentrated near the PMT surface, and one near the bottom of the detector volume. The estimates of asymmetry are degenerate with the overall light production and collection efficiency in the detector, and only serve as a crude position estimate for events.




Figure 4.4: LoLX phase-2 light asymmetry distributions for ^{133}Ba 4 V OV data runs. We see two main populations of events, a large excess at high and low positions along the height of the detector (Z-axis or top-bottom axis). These are fairly aligned in X-Y direction, as seen in the left panel, which shows an extended region for events from the North-East corner of the detector.

With the rough position estimation from asymmetry, we identify a set of data selection cuts along asymmetry directions, Top-Bottom, North-South and East-West. We apply the following selection criteria: $A_{T-B} < 0.3$, $A_{N-S} > 0.1$, and $A_{E-W} < -0.1$. These cuts should help isolate events from the external source, from scatters, and backgrounds from directly in-front of the PMT.

Next, we look at the total signal sensor contributions to the total signal per event in the energy distributions shown in Figure 4.5, and compare these to the distributions of rough position-selected cuts. We see very broad features with no clearly identifiable γ -like peak features. This was initially unexpected as preliminary simulations of LoLX phase-2 indicated a higher resolution was possible. We will revisit this point when we compare this data with detailed simulations in the following section. An additional observation present in Figure 4.5, but not in the previous August 2023 dataset, is the lack of a low-energy peak in the distribution. This is likely due to the combination of the higher relative trigger threshold used in 2025, and the $\mathcal{O}(10\times)$ higher activity barium source used for calibration, leading to some DAQ buffer overflow.




Figure 4.5: Energy spectrum comparison of external ^{133}Ba source data and background-only data from the 4 V OV January 2025 LoLX phase-2 commissioning. Top: Live-time normalized rate of events per 10 PE bin, for source and background data with and without positioning cuts. Bottom: Background subtracted source data compared with background subtracted position-cut data. We see a very poor resolution in total signal integration, which was unexpected.

To compare the sharing of the total signal between the three types of light sensors, including FBK SiPMs, HPK SiPMs, and the PMT, and to assess their relative performance, we projected the channels against each other in 2D distributions and compared them to simulated estimates. These projections are shown in Figure 4.6, for the position-cut dataset for external ^{133}Ba source calibration data at 4 V OV. We use the position-cut data set to remove spurious and background-like contributions to the distribution for clarity. However, by restricting to the bottom-half of the detector, we are introducing a bias in light collection efficiency that will reduce the overall PMT signal. Because the SiPMs are uniformly distributed, they do not experience the same level of bias in collection efficiency.




Figure 4.6: 2D light channel signal projections of external ^{133}Ba source data with LoLX, with position-cuts applied, for all sensors. Left: Total SiPM detected PE, summing all sensors for both vendors vs total PMT detected PE from position selected events. Right: Comparison of summed HPK and summed FBK SiPM signals.

As we see clearly in the right panel of Figure 4.6, the HPK SiPMs significantly underperform compared to the FBK devices. The suspected cause of this discrepancy is discussed in the recently submitted manuscript [139] by collaborator Alex Li. For completeness, we briefly comment on it later in the comparison with simulations. In the left panel of Figure 4.6, the light collection in the LoLX phase-2 data shows that the PMT and SiPM signals are anti-correlated due to the detector geometry: when the PMT receives the most light, the SiPMs receive less, and vice versa.

4.2 Simulation Comparison

4.2.1 Geant4 Input

We simulated the ^{133}Ba decays using the G4LOLX simulation described in Section 2.5, to compare to the data taken during the second commissioning run of LoLX phase-2. The simulated position is approximately located where the source was positioned during data-taking, and the Global Particle Source (GPS) was used to simulate decay γ s with a piece-wise user-defined energy histogram, with relative intensities scaled to the data from online IAEA NDS tables

[140].

The resulting total event energy deposition and position of energy depositions in the LXe scoring volume are shown in Figures 4.7 and 4.8, respectively. The skewed position distribution for energy depositions is consistent with the observed signal asymmetry shown in Figure 4.4. The position distribution appears clipped due to the limited extent of the LXe scoring volume used in G4L0LX. The simulated energy deposition positions, shown in Figure 4.8, are displayed in the Chroma frame of reference, which has a -283.0 cm Z-position offset with respect to the G4L0LX positions.




Figure 4.7: Total event energy of ^{133}Ba decay events that deposited energy in the LoLX Geant4 LXe scoring volume. The dominant γ lines are apparent, with the most dominant line at ~ 356 keV.




Figure 4.8: External ^{133}Ba decay events energy deposition positions in LoLX phase-2 detector scoring volume, in the LoLX phase-2 Chroma frame of reference. The source was simulated at (7.5 cm, 0 cm, 30 cm) with respect to the G4LoLX geometry, and we see significant source-side bias in X-position of deposits, and a milder bias towards lower Z positions.

We can see clearly that the source energy deposition positions are biased towards the lower Z and +X direction, similar to what we saw previously with the barium source data asymmetry projections in Figure 4.4. We can deduce that the relative orientation of the LoLX phase-2 detector in the cryostat is likely rotated by $\sim 45^\circ$, compared to its position in the simulation. This may lead to some differences in the propagation of light for data-simulation comparisons, and motivates the need for a specific clocking to be used during installation so that the simulated source positions are known a priori.

4.2.2 Chroma Simulations in LoLX phase-2

LoLX Phase-2 Light Map

To better understand the propagation of light within the LoLX phase-2 detector, we can measure the total light collection efficiency as a function of simulated position. This is shown for 20k “scintillation photon-bombs”, where we use the NEST model in Chroma to produce LXe scintillation-like photons with fluctuations driven by the chosen source (356 keV β s here).

The light collection efficiency is the ratio of detected to created photons per event, averaging over multiple events for each position voxel (volume pixel). This is shown as the X vs Z position, with projections averaging over either X or Y in Figure 4.9. The projections are displayed as profile histograms, where the mean value for each projection is averaged over X or Y versus Z, and the uncertainty on each point represents the $1\text{-}\sigma$ standard deviation of the estimates. For the sake of comparison, we include the projections and 2D lightmap for a more realistic detector, removing detected photons from the list of non-functioning SiPMs in the calculations of efficiency. These channels are fairly evenly distributed in the detector, as shown in Figure 4.1, and when averaging over the entire detector, we move from an average light collection efficiency in the ideal case of $\sim 7.6\%$ to $\sim 6.6\%$ in the realistic case, a $\sim 13\%$ decrease in overall light collection efficiency.




Figure 4.9: Light collection efficiency map for scintillation events in LoLX phase-2 detector, with idealized and estimated response side-by-side. Left panels: “Perfect” detector, with all working SiPM channels, and projections using profile histograms of X/Y versus Z position. Right panels: “Realistic” detector, with unplugged channels from January 2025 LXe run excluded from detected photon counts. Simulated SiPM PDE is for 3 V OV operation for both device types.

Barium-133 Response

To compare the simulated Geant4→Chroma barium events to our data, we must first apply some minor corrections. The simulated photon detection efficiency is for nominal 3 V OV data-taking, which requires us to scale up the collected light by a small amount. The scaling is derived from [103], a $\sim 7.5\%$ relative increase in efficiency for VUV photons driven by the increase in avalanche initiation probability from 3→4 V. Additionally, we remove detected

photons from any channels that were unplugged in the January 2025 commissioning run, for comparison to the source data from the previous section. Lastly, we smear the perfect detected photon count by the SiPM response expectation from [101]. The smearing is done by randomly sampling a normal distribution with parameters from [101] for mean excess charge and excess charge variance. For each detected photon count in a given SiPM group we multiply by the randomly sampled noise factor from the prescribed distribution for each group individually. Total smeared counts and PMT detected photon counts are summed at the end to produce the orange curve in Figure 4.10, which we compare to the pure detected photon distribution and the PMT detected photon distribution. We can also do an approximate re-scaling of the simulated data to better match what we observed in LoLX phase-2 January 2025 commissioning runs, shown in Figure 4.5.




Figure 4.10: Pure detected photon and estimated PE spectra comparison for simulated external ^{133}Ba source events, normalized. The detected “PE” spectrum includes smearing from the expected excess charge distribution from [101] for FBK and HPK SiPMs. The re-scaled distribution has been downsampled by the observed data-simulation yield difference, with $f = 0.7$. SiPM channels which were unplugged in the January 2025 run are excluded from the summation here.

As seen in Figure 4.5, while the initial detected photon distribution still contains some peak-like features, once we include the effect of re-scaling to a lower overall light yield, and introduce the variance from SiPM correlated noise, there is a significant loss in resolution. This effect was initially unexpected and we have been pursuing a neural-net based position reconstruction algorithm to recover resolution in combination with a detailed lightmap like the one shown in Figure 4.9.

We can also compare the simulation and the data energy spectrum qualitatively, by overlaying and normalizing the data and simulation to the same amplitudes. This is shown in Figure 4.11 with the non-cut and not-background subtracted ^{133}Ba source data overlaid with the smeared and re-scaled distributions from simulation.




Figure 4.11: Comparison of simulated and measured external ^{133}Ba source energy spectrum, in total detected PE. By applying an overall down-scaling of the detected PE in simulation with a factor, $f = 0.7$, we get a fair shape agreement between data and simulation.

While the overall shape of the down-scaled and smeared simulation energy spectrum agrees roughly with the measured data, there are a number of differences. The first is the narrower peak, and broader tail of the energy spectrum above 1500 PE. This is likely an ar-

tifact of clipped pulses, since we observe for a significant fraction of events at full-energy that one or more sensors have clipped the dynamic range of the readout electronics. This contributes to the broadening in the distribution of high-energy events. The second significant difference is the lack of the low-energy peak in the data. While some of this loss can be attributed to overflow in the DAQ buffer (as the majority of events expected are high energy, as seen in Figure 4.7), we can model our detector’s trigger efficiency for a more quantitative comparison. This is shown in Figure 4.12, using the simulated external barium source dataset shown here. Barium calibration data was taken with a -25 mV threshold on the PMT signal, approximately 8 photons equivalent in amplitude.




Figure 4.12: Profile distribution of simulated trigger efficiency for LoLX phase-2 commissioning data with external ^{133}Ba source events. An event is considered “triggered” when 8 photons are detected in the PMT. Each point represents the mean trigger efficiency for events with the corresponding total simulated light level, an analogue for energy, while the error-bars are the $1\text{-}\sigma$ standard deviations of the trigger efficiency distribution for each bin.

As we see in the profile distribution shown in Figure 4.12, there is a very large variance in trigger efficiency even for high light levels in the detector, due to the geometric effects in light collection as shown in Figure 4.9. The trigger efficiency seems to stabilize above ~ 2000 created photons, and begins again increasing at ~ 8000 photons. For events in the bottom half of the detector, which make up most of our simulated vertices, we expect a light collection

efficiency of $\sim 6\%$. Converting to these values to detected photons, we estimate the trigger threshold to be approximately 200 PE for ^{133}Ba . The efficiency increases above roughly 500 PE, which corresponds to the low-energy cutoff observed in simulations, as shown in Figure 4.12.

We also project the simulated results across channels in 2D to compare to the distributions from data measurements, shown in Figure 4.6. The simulated external ^{133}Ba dataset used here, with the data driven rescaling, but without ECF smearing, is shown in Figure 4.13.




Figure 4.13: 2D light channel simulation projections for external ^{133}Ba source in LoLX phase-2. The overall light-yield has been down-scaled by $f = 0.7$ to match the overall yield in data. Left: Detected PMT photons vs detected SiPM photons, without noise. Right: Simulated HPK channel vs FBK channel response, with January 2025 run channels omitted from photon summation in the respective groups.

By comparing the simulation projection in Figure 4.13 with the data in Figure 4.6 we can make some additional observations. We see that the overall shape in SiPM vs PMT signals is reproduced, although the branch extending to high PMT signals occurs at a higher SiPM light level. Comparing the HPK and FBK projections shows that the HPK devices underperform relative to simulation expectations. Because the total light yield was scaled by matching the total PE, this procedure does not account for the HPK underestimation, leading to the branch at high-PMT in the simulated SiPM vs PMT projection, to be shifted relative to the data. The simulations could be better described by a specific down-scaling of the HPK simulated per-

formance, with an absolute yield difference from the PMT-only energy spectrum. However, this would be contingent on correcting for PMT clipping to accurately estimate the end-point in energy.

Finally, we can compare the timing response of the detector to the expected LXe timing pulse-shape from the simulated data that uses the NEST model. This is done by summing the time-aligned baseline corrected PMT waveforms, we write to the output .root file in full for one barium calibration dataset, run# 43. Because the PMT signal is very prompt, without the long decay tails in the SiPM waveforms, we can estimate the pulse-shape by summing many PMT waveforms [141]. The average PMT waveform is compared to the simulated photon travel time data in Figure 4.14. For a more accurate comparison, we also overlay a simulated spectrum that includes a Gaussian smearing with width $\sigma = 7$ ns, tuned to match the rise-time of the data profile, which includes contributions from DAQ jitter, PMT response width, and event misalignment. The simulated spectrum comes from Chroma simulation results, and was shifted to match the trigger offset in data.




Figure 4.14: Comparison of simulated and measured LXe time profiles in LoLX phase-2 for external ^{133}Ba source calibrations. The simulated time profile without SiPM correlated noise is shown in blue, and a Gaussian smeared simulation ($\sigma = 7$ ns), is overlaid in orange for a more realistic comparison with the summed PMT waveform from data. The residual panel shows the relative residuals for data versus the smeared simulation profile. We observe an intermediate time excess and late time deficit when comparing the shapes of the simulated and measured time profiles for LXe scintillation.

The observed deficit of late time LXe photons, shown in Figure 4.14, is understood to be an effect of scintillation quenching from impurities in the liquid [142]. We have independently measured $\mathcal{O}(100$ ppm) nitrogen contamination in the GXe during inspection with a residual gas analyzer, probing Xe gas samples of the LoLX gas handling system. We have also observed an overall quenching of the total light yield of the data when comparing to simulations, which is only partially explained by a lower than expected HPK performance in LXe compared to vacuum measurements. These observations motivated the inclusion of the purity monitor in the January 2025 commissioning run.

Preliminary modelling of the PMT LXe profile, shown in Figure 4.14, indicated an effective triplet life-time lower than expected from simulations and reported measurements [143].

This effect is consistent with nitrogen-driven quenching through collisional de-excitation, shown earlier in Figure 2.3. Longer lived triplet excimers, formed primarily through the ionization channel, will disassociate through collisions, and begin forming new excimers with nearby neutral xenon. This populates singlet states primarily by forming through excitation channel, and promptly de-exciting through the singlet decay mode ($\tau_s \sim 4$ ns) [142]. We consider this effect together with an unknown level of electronegative impurities in the bulk LXe. These impurities predominantly quench the late-light by absorbing ionization electrons before dimers can form. As a result, a significant reduction in the overall light yield is expected, consistent with the scales observed in LoLX phase-2 commissioning data. Preliminary modelling of the measured pulse-shape indicated fair agreement with a modified 2-exponential model that included an intermediate “recombination” term with a negative contribution, as described in [142], with input from [144]. This negative contribution could be interpreted as a quenching term from collisional de-excitation, which has an interaction rate roughly the same order as free recombination ($\text{Xe}_2^+ + e^- \rightarrow \text{Xe}^{**} + \text{Xe}$). However, more work and a fully constrained model are needed as the individual fit contributions are highly coupled, and the measured fit χ^2 are incompatible with a good fit.

4.3 Calibration Sources and Future Runs

From our commissioning data runs, we have seen the limitations of using an external γ -source for energy calibrations of LoLX phase-2. In order to improve the precision of detector response calibration for improved data-simulation comparisons, and to enhance energy resolution for future studies of LXe scintillation response, we must move to an internal calibration source. We have the following requirements for an internal source: low enough mean energies to avoid saturation and clipping artifacts, a point-like geometry for reconstruction simplicity and to avoid shadowing effects, and suitability for a cryogenic noble liquid environment. While the first requirement can be addressed partially from more advanced triggering

by requiring single-photon sensitivity, we are limited by the minimum gain required to have sufficiently high SNR for SPE pulse identification. The second requirement is motivated by experience with LoLX phase-1 where an extended ^{90}Sr needle source was used, introducing both shadowing effects and artifacts from local energy loss in the needle volume. The final requirement is partially motivated by nominal LXe detector component requirements already discussed, with additional consideration from the observed flaking of the ^{90}Sr coating used in LoLX phase-1 operations.

4.3.1 Hafnium-181 Internal Calibration Source

^{181}Hf was chosen as a calibration source for several reasons: it is readily available, its production is relatively straightforward, and it satisfies the experimental constraints outlined previously. During private communications with Professor Dominic Ryan (McGill), we were recommended to use natural hafnium activated through thermal neutron activation to produce ^{181}Hf . This occurs through the reaction $^{180}\text{Hf} + n \rightarrow ^{181}\text{Hf}$ and was carried out at the MacMaster University research reactor. Activated hafnium is an attractive source for our needs, the absence of long-lived activation isotopes other than ^{181}Hf , and an energy range for β and γ emissions that is well matched with the needs for LoLX phase-2. ^{181}Hf decays via β^- decay. The dominant branch ($\sim 93\%$ B.R.) has an endpoint energy $Q = 404\text{ keV}$ and a mean emission energy of 120.7 keV [145].

The full decay scheme of ^{181}Hf is shown in Figure 4.15 reproduced from [146]. There are several dominant γ lines useful for LoLX phase-2 calibration, including the two main lines, from the primary β^- channel, of 133 keV followed by a delayed 482 keV emission with a decay constant of $\tau \sim 1.8 \times 10^1\text{ }\mu\text{s}$. In LoLX phase-2 we expect to be able to correlate these delayed emissions through coincidence tagging. With a decay half-life of 42.4 d , the quantity of ^{181}Hf in our activated source will decay over time, but is easily “refreshed” with subsequent activations.




Figure 4.15: Decay schematic for ^{181}Hf reproduced from [146].

As an activated all-metal source, ^{181}Hf is suitable for use in a LXe cryostat. For our preliminary source preparation we purchased >95 % purity natural hafnium wire with $D = 0.5$ mm from Goodfellow advanced materials, and prepared three samples for activation. For activation the reactor used was the McMaster Reactor Rabbit activation system [147], a fast pneumatic deployment system for small-sample short (< 1 h) activations at the 5 MW reactor facility. Samples must be < 5 cm in length and fit within a 1.8 cm diameter polyethylene sample tubes. The three samples we activated are pictured in Figure 4.16 in a sealed quartz ampule¹. The three samples were: a ~ 3 cm hafnium wire, a ~ 3 cm zirconium rod witness sample, and the prepared source sample, a ~ 3 cm zirconium rod also from Goodfellow, with a spot-welded hafnium “nub”. We attempted to spot-weld hafnium to its sister metal, zirconium, which is a widely used material in nuclear reactors due to its low neutron absorption [148]. After production, each sample was cleaned by hand with degreaser, sonicated for 30 minutes in acetone, then ethanol, and finally rinsed in ethanol then dried. All three samples were then sealed in a quartz ampule packed with quartz wool, following guidance from the McMaster reactor manager. Unfortunately, the spot-welded hafnium “nub” delaminated dur-

¹Many thanks to Dominic Ryan for his assistance in preparing and sealing the samples.

ing transport, leaving the activated hafnium wire as our source for upcoming runs.



Figure 4.16: Sealed hafnium source and witness samples before irradiation at McMaster in March 2025.

Activation estimates were carried out using the online NIST neutron activation analysis estimator [149], to provide a guide for selecting activation times. The source was irradiated for 60 s in a $\sim 5 \times 10^{12}$ n/cm²/s thermal neutron flux at one of the Rabbit sites without a cadmium foil-lined container. This was estimated to produce a total γ activity of ~ 175 kBq after 24 h of cooling for all materials in the sample container, which was very close to the 2-day post-irradiation measured total activity of ~ 170 kBq². The wire sample will still have an estimated ~ 440 Bq of activity remaining after one year, suitable for calibration purposes in LoLX phase-2 where the nominal DAQ event rate maxes out at ~ 400 Hz.

After activation on May 13th 2025, we recieved the samples back the following week and proceeded with unpacking and cursory inspection using the departmental high purity germanium (HPGe) gamma spectroscopy setup available in the undergraduate physics labs. This work was carried out by undergraduate summer research intern Ethan Coulthard, under the

²Private communication, McMaster Nuclear Reactor facility manager, 2024.

supervision of Thomas Brunner and myself. The HPGe inspections were carried out for all samples, with the wire sample and background spectrum shown in Figure 4.17, reproduced with data shared by Ethan. We do not observe any higher energy γ lines, our primary concern for activated isotope contamination due to saturation effects in LoLX phase-2.




Figure 4.17: HPGe spectrum of ^{181}Hf wire sample overlaid with HPGe background, displayed in MCA channel number. Both samples are normalized to their rates in the shaded normalization region for visualization and qualitative inspection. Only hafnium emission lines are observed in the activated ^{181}Hf sample. The highest energy feature represents the transition of β -decay daughter, $^{181}\text{Ta}^*$ from $+5/2 \rightarrow +7/2$ ground state of ~ 482 keV. HPGe data was taken by undergraduate research intern Ethan Coulthard.

Future LoLX phase-2 calibration will be performed using the installed ^{181}Hf source, which will be mounted inside the LoLX detector volume in a source holder designed to expose the hafnium wire tip while shielding the rest of the wire. This source will be used to better estimate the intrinsic energy resolution of LoLX phase-2, constrain LXe scintillation leakage through optical filters that will be installed, and improve modelling constraints for optical transport and SiPM detection efficiency measurements.

4.4 Conclusion and Next Steps

We showed and discussed the results from the second commissioning run of LoLX phase-2, with a comparison to simulations for key parameters. We demonstrated that the SiPM SPE charge calibration can be recovered in a low SNR regime by employing pulse template fitting to inform charge corrections. We observed that the performance of the LoLX phase-2 detector has broad agreement with simulations in measured energy response shape, if not scale, and in timing and positioning response through weighted charge averaging. More modelling work is needed to prepare for future runs with LoLX phase-2, which will include an internal calibration ^{181}Hf source. The inclusion of the hafnium calibration source is expected to address our key challenge points from these commissioning runs. By using a low-energy internal source, we aim to improve detector energy resolution, and fully constrain source position knowledge for more precise simulation-data validations in the future.

Through three independent measurements, we observed phenomena that could contribute to the observed deficit of scintillation light in LoLX. These include N_2 detected in the GXe processing system by a residual gas analyzer, a light yield deficit relative to simulations that cannot be attributed to sensor performance, and a marked reduction of triplet light in the LXe scintillation profile. While none of these alone are conclusive, we have also installed an electron drift lifetime monitor to establish a baseline LXe purity measurement, sensitive to all electronegative impurities except nitrogen [95]. The design is based on a version developed by nEXO collaborators at the University of Massachusetts [138], using a xenon flashlamp to produce photoelectrons on a gold-coated photocathode. During the January 2025 commissioning run, we did not observe any signals in the purity monitor when submerged in LXe, while operation was nominal during GXe testing. During GXe measurements, we saw significant purity degradation from outgassing, during the room temperature calibration measurements. Follow-up testing and the inclusion of a higher power flashlamp are being carried out to better

understand the source of signal loss when submerged in LXe. Independently, we have begun the design of a xenon recirculation system, expected to significantly improve the purity of the xenon for operations. Improved purity is a requirement for any future measurements of xenon scintillation properties or timing behaviour.

We omitted the discussion on the LXe responsivity of the SiPMs relative to the VUV PMT, which is the subject of a recently submitted manuscript lead by collaborator Alex Li [139]. In brief, these characterizations, which required the development of light-only position reconstruction for improved detector resolution, show that by modelling the surface features and shadowing of light from oblique angles on the HPK SiPMs, we can reproduce the measured relative performance of FBK and HPK in LXe. However, this work still requires the inclusion of a global light-yield quenching to match the absolute light levels observed in LoLX phase-2 data.

With a much more stable system, a better understood detector, and improved source and hardware calibration for future runs, LoLX is poised to perform more detailed and critical detector R&D for LXe based detectors, such as nEXO, XLZD and PIONEER [150], [151], and to further our understanding of xenon light production and transport through dedicated studies. These studies include: long-term stability performance of SiPMs in LXe for weeks-month-long operations, angle resolved measurements of SiPM ExCT, and SiPM ExCT from FBK SiPMs in LXe, measurements of zero-field light yield in LXe for β , α , and γ radiation, and the characterization of Cherenkov production from low energy β 's, alongside other fast-timing measurements such as the LXe scintillation time profile at zero-field.

Future work could also include a physics-informed model of light quenching from impurities in LXe in order to improve on the poorly constrained data shown from our commissioning run. This would be significantly improved by dedicated timing resolution and response measurements of the PMT, afterpulsing constrains for the PMT from dark and cold laser calibration data, combined with benchmark purity estimates from electron drift lifetime measure-

ments. Since N_2 contamination is known to affect the measured timing response, but is only marginally sensitive in electron drift lifetime monitors, characterizing in detail the effect of nitrogen contamination in LXe would provide valuable insight to its impact on LXe detector performance.

Shortly after the January 2025 commissioning runs explored here, we experienced a critical component failure during gas transfer operations. The high-pressure regulator used to separate the gas storage bottles from the low-pressure deployment side of the system formed a leak. Dedicated characterization and preliminary purification through xenon distillation is underway to return the xenon to a sufficiently pure state. While this additional purification and characterization continue, we prepare for future runs, including the development of a re-circulation system and the installation of upgraded readout electronics boards to improve baseline noise and readout SNR.

Part III

Simulating the nEXO Light Detection System

The nEXO Experiment

The nEXO experiment is a next-generation $0\nu\beta\beta$ search experiment that builds on the foundation of the successful predecessor experiment EXO-200 [152]. The Enriched Xenon Observatory (EXO-200), ran from 2011 until 2018 [152]. EXO-200 performed the first measurement of the two neutrino decay mode, $2\nu\beta\beta$ in ^{136}Xe [153], and set world leading limits in $0\nu\beta\beta$, with a final half-life limit of $\tau_{1/2} > 3.5 \times 10^{25}$ yr [152]. The EXO-200 experiment employed a low-background LXe time projection chamber (TPC) with ~ 200 kg of LXe, and combined charge and light readouts with an energy resolution of ~ 1.2 % of the FWHM at the ^{136}Xe endpoint energy [152]. The $0\nu\beta\beta$ half-life limit in ^{136}Xe from EXO-200 was recently surpassed by the KAMLAND-Zen experiment [47].

The nEXO experiment aims to improve significantly on the EXO-200 results, with an up-

graded detector design, a greatly increased xenon mass, and a reduced background level. To meet these ambitious goals, the nEXO design has been developed over the past decade [154], taking advantage of technological advances in the field, such as VUV-sensitive SiPMs [101]. nEXO seeks to constrain the entire inverted mass hierarchy with an ultimate projected sensitivity to $0\nu\beta\beta$ of $\tau_{1/2} > 1.3 \times 10^{28}$ yr at the 90 % confidence level after ten years of running [52].

The nEXO design is composed of two detector systems, the outer detector (OD) a 1.5 kt Water-Cherenkov muon-veto and radiation shield, and the inner detector, a 5 t LXe TPC (~ 3.7 t active). Time projection chambers are an excellent technology choice for $0\nu\beta\beta$ searches, driven by two main criteria: energy resolution and background suppression. The large monolithic LXe volume surrounded by light sensors along the barrel, and charge collection tiles at the anode top-plane allows for fiducialization of an analysis volume at the innermost part of the detector. This fiducial volume will have excellent signal purity, i.e. with very few contributions from background signals. In order to reach its sensitivity goals required for $0\nu\beta\beta$ discovery, we must address the main challenges of the EXO-200 experiment, and increase the isotope mass significantly. To reduce the cosmogenic activation of detector components, especially ^{137}Xe , nEXO will be located at SNOLAB, the deepest underground science site in North America [155]. SNOLAB, host of the eponymous SNO experiment, is located in Sudbury, Ontario, and boasts a dense overburden of Canadian Shield rock with a water equivalent depth of ~ 6000 m [155]. This shielding is more than three times the overburden of EXO-200, which had a water equivalent depth of ~ 1600 m [156].

In this chapter, we will review the nEXO detector systems in Section 5.1, the main data analysis methodologies for the $0\nu\beta\beta$ search in Section 5.2, and nEXO's sensitivity to $0\nu\beta\beta$ in Section 5.3.

5.1 nEXO Overview

The nEXO experiment will be composed of a 5 t LXe TPC inner detector, nested within the outer detector, a 1.5 kt water Cherenkov veto, planned to be housed within the SNOLAB Cryopit. The cross-section of the conceptual nEXO design is shown in Figure 5.1.




Figure 5.1: Sketch of the nEXO detector concept from the 2018 Pre-conceptual design report [154]. The illustration shows the LXe TPC located inside the insulated cryostat filled with HFE-7000 refrigerant, all enclosed within the outer detector water tank, filled with 1.5 kt of ultrapure water. A cartoon person is included in the access tunnel for scale.

5.1.1 Outer Detector

To reach the sensitivities required for nEXO, we employ the techniques described briefly in Section 1.5.2. The nEXO outer detector (OD) serves the dual purpose of passive background shielding and active muon veto.

The outer detector is a 1.5 kt water-Cherenkov detector instrumented with 125 PMTs. The large water volume provides passive shielding against the γ and neutron radiation flux from the Cryopit walls, with hydrogen-rich water being particularly effective for neutron moderation. The OD dimensions are constrained by the SNOLAB Cryopit geometry, and meet the minimum shielding requirements for γ radiation [71].

Beyond passive shielding, the primary objective of the OD is to tag all through-going muons with efficiency $>95\%$. In particular, muons passing within 3.5 m of the cryostat vessel produce the vast majority of ^{137}Xe atoms expected in the LXe TPC [71]. For muons passing within 4 m of the TPC, a conservative veto window of 25 minutes is applied to effectively remove ^{137}Xe decays ($\tau_{1/2} \sim 4$ min), adding only 1.3 days of dead-time over the ten-year experiment lifetime [71]. Additional synergies are possible by combining muon and TPC signals to tag ^{137}Xe de-excitation radiation, potentially reducing the required veto window length.

5.1.2 Inner Detector

The heart of the nEXO experiment is the inner detector, the LXe TPC containing xenon enriched to 90 % of the $0\nu\beta\beta$ isotope ^{136}Xe . The inner detector is a 1.3 m right-cylindrical copper vessel. The vessel will be made of electroformed copper, formed underground at SNOLAB to reduce cosmogenically activated backgrounds [52]. A rendering of the nEXO inner detector with TPC is shown in Figure 5.2 reproduced from [52].




Figure 5.2: Visualization of the Geant4 simulated nEXO TPC from [52]. This 2/3 projection shows the main components of the TPC, including charge and light detection locations and E-field shaping rings.

nEXO’s detector key components, highlighted in Figure 5.2, include the charge collection system, the light collection system and the electric field shaping rings.

Liquid Xenon Time Projection Chambers

The operational principle of a LXe TPC is illustrated in Figure 5.3. When incident radiation (for example $2\nu\beta\beta$) interacts with LXe is within the TPC, the ionization electrons and scintillation photons are produced in an anti-correlated manner [157]. Additionally, the short γ -attenuation length of LXe (~ 8.5 cm at 2.5 MeV), induces a self-shielding effect. The centre of nEXO will be ~ 8 radiation lengths away for the critical background γ ’s near the $0\nu\beta\beta$ region of interest.




Figure 5.3: Illustration of LXe TPC operational principles. When radiation interacts with LXe, ionization electrons are drifted to the charge collection tiles on the anode in a 2D grid. The drift field is modified by the field shaping rings (orange) to produce a linear field within the fiducial volume. Scintillation photons are detected by SiPMs surrounding the barrel. The interaction 3D position is estimated by combining the (x,y) position of the interaction estimated from the 2D charge grid signal, with the z position, estimated using the time difference between scintillation and charge signals.

The ionization electrons are drifted by the electric field produced by the potential difference between the anode and cathode planes, accented by the field shaping rings along the length of the cylindrical chamber. The field cage is configured so that a static electric field is present throughout the bulk of the detector volume, essential for event position reconstruction. During the drift of ionization electrons to the charge collection tiles on the anode top-plane, scintillation photons travel promptly from the point of origin of the interaction, to be detected by the SiPMs along the barrel of the chamber. By combining the detected charge and light sig-

nals, the estimated energy of the interaction is reconstructed. nEXO’s energy resolution target at $Q_{\beta\beta}$ is $<1\%$ (σ/μ), while 0.67% has been demonstrated in a similar configuration [158]. The reconstructed 3D position of the interaction from the TPC is a key input in background discrimination. This is because the anticipated signal ($0\nu\beta\beta$) would be distributed volumetrically within the bulk LXe, and most backgrounds are concentrated at the detector surfaces [52].

The (x,y) position is easily extracted from the charge collection pads, which are oriented in a 2D grid. Lastly, the z-position can be estimated from the time difference between the scintillation and charge signals. The scintillation signal is produced promptly, and is transported in $\mathcal{O}(\text{ns})$, while the charge signal takes $\mathcal{O}(100\ \mu\text{s})$ to arrive. When light and charge information is combined, the 3D position vertex can be estimated within $\sim 1\text{ cm}^3$ [159].

With the 3D position vertex of the interaction site in-hand, we can define one of our key background discrimination variables, standoff distance. We define standoff distance as the Euclidean distance between the interaction site, and the nearest detector surface.

By combining the light and charge signals together, we can extract an energy estimator, referred to as “rotated energy”. The linear combination is selected to maximize linear response for β -like events, which incidentally provides effective discrimination against α -like events due to their significantly different ionization fractions [160]. Figure 5.4 shows the rotated energy spectrum in charge vs light space for simulated backgrounds from [52]. The rotated energy can be represented by Equation 5.1, where the scintillation and charge constants (c_P , c_Q) are selected for optimal combined energy resolution:

$$E_R = c_P \tilde{N}_P + c_Q \tilde{N}_Q, \quad (5.1)$$

where, $\tilde{N}_P$ and $\tilde{N}_Q$ are the reconstructed light and charge signals from the event, correcting for detector non-uniformities using a lightmap and expected charge attenuation loss, respectively. In this representation, c_P and c_Q are obtained from calibration and encode the detector

response to photons and charge per unit energy.




Figure 5.4: 2D plot of charge vs light energy from nEXO simulations. Rotated energy axis overlaid. By inspecting the reconstructed light ($\tilde{N}_P$) and charge ($\tilde{N}_Q$) signals of simulated ^{232}Th background decays from the TPC vessel. Figure reproduced from [52].

nEXO Charge Collection System

The charge collection system consists of two main components, the E-field shaping cage and charge collection tiles. The nEXO field shaping cage has been designed to provide a stable 400 V cm^{-1} electric field region within the fiducial volume of the LXe TPC, containing $\sim 3300 \text{ kg}$ of LXe with the remaining LXe in the so-called “skin” region. The TPC will have a $\sim 1.2 \text{ m}$ drift volume, and the fiducial volume is set by a minimum of 2 cm of standoff distance for uniform fields from simulations [52]. At this field strength, the electron drift velocity is $\sim 0.2 \text{ cm}/\mu\text{s}$. The field cage will consist of 58 field shaping rings (FSRs) of 113.4 cm inner diameter, supported by UltemTM (PEI) rods and fused silica spacers. The cathode will

be coated with a reflective MgF_2 /aluminum stack to enhance light collection efficiency.

The drifted charge will be collected at the anode by an arrangement of segmented charge readout tiles. Each tile will be $\sim 10 \times 10 \text{ cm}^2$, made of fused silica with gold and SiO_2 coatings. They will have 30 ‘X’ and 30 or ‘Y’ strips, arranged diagonally with $\sim 3 \text{ mm}$ pads, with strips separated by thick layers of SiO_2 at crossing points to electrically isolate channels [161]. The geometry was chosen to optimize charge collection, while minimizing charge build-up and capacitance. The final design and arrangement of the charge collection system is being finalized, but will consist of ~ 120 tiles, covering the majority of the anode.

nEXO Light Collection Subsystem

The nEXO scintillation light collection subsystem plays a central role in background reduction and improving energy resolution. The main design goal of the nEXO light collection system is to detect more than 3 % of the total light produced from scintillation within the fiducial xenon volume [154]. There are two additional requirements: the photodetectors should not introduce more than 1 % of the total background in the region of interest, and the total detector energy resolution (σ_E/μ_E) should be less than 1 % [162].

Additionally, nEXO’s photodetectors must be suitable for long-term operation in LXe at cryogenic temperatures, and in order to reach the energy resolution and light collection targets, must have low correlated noise and high detection efficiency for LXe scintillation. The baseline light system design includes five integration tiers are illustrated Figure 5.5 reproduced from [162] described according to label:




Figure 5.5: Overview of nEXO light collection system components from [162]. a) Basic unit of $1 \times 1 \text{ cm}^2$ SiPM, b) Image of a proto-type sub-array with 6 SiPMs, each sub-array is read out as a single channel c) Back-view of prototype tile-module, each tile-module has 16 sub-arrays. d) SiPM stave, the largest unit in nEXO PD system, each stave has 20 tile-modules. e) Full light collection barrel of nEXO with 24 staves for a total of 46080 SiPMs and 7680 readout channels.

- a) SiPMs — The $1 \times 1 \text{ cm}$ photosensor is the base unit of the light collection system. Two main vendors are being considered for nEXO: FBK HD3 SiPMs and HPK VUV4 multi-pixel photon counters (MPPCs).
- b) SiPM sub-array — 6 SiPMs are arranged in a “3P2S” configuration (3-parallel and 2 series) to minimize noise and maximize sensor granularity while optimizing for readout capacitance.
- c) SiPM Tile — 16 sub-arrays are readout by a single tile module, equipped with an FeASIC for fast analogue analysis, performing shaping and amplification. SiPMs on the tile module are connected to the FeASIC through a fused silica interposer.
- d) Photodetector Stave — 20 tiles are arranged in a 2×10 configuration, each stave is equipped with an ADCTX ASIC connected to the tile modules by radiopure Kapton cabling [163], placed at the top of the stave to reduce LXe boiling in critical areas. This ASIC performs analogue-to-digital conversion for each tile and passes along the

waveform snippets to the DAQ system.

- e) The full light collection barrel with 24 staves, for a total of 46080 SiPMs and 7680 readout channels.

Both FBK and HPK SiPM vendors are still under consideration for the nEXO photodetector system (PD), and as mentioned in Chapter 2, LXe testing is underway to help inform vendor selection. We have some additional considerations for vendor selection which include: dark noise rates, long term performance stability in LXe, and the capacity for through-hole vias. The high granularity of the light collection system will aid with background reduction by enabling hit-pattern based analyses, especially helpful for surface and skin-volume backgrounds.

Detailed optical characterization of in-situ components is required in order to produce accurate simulations of the nEXO light collection system, and extensive R&D has been underway to support these efforts [101], [132], [164]. In the following chapter, we will discuss the past and current state of nEXO’s optical simulations, and work underway to improve modelling and understanding of the essential scintillation light collection system.

5.2 nEXO Multiparameter Data Analysis

nEXO builds on the multi-parameter profile likelihood developed for the EXO-200 experiment [152]. By employing profile likelihoods, rather than using a simple cut-and-count method, nEXO will optimize its sensitivity to $0\nu\beta\beta$, while minimizing the contribution from background fluctuations [52]. The profile likelihood analysis uses input profile density functions (PDFs) weighted proportional to the background contribution exposures in the analysis dataset (total activity and acceptance dependent). These background PDFs are 3-dimensional (energy, standoff distance, and event topology) and bounded by the first-pass analysis cuts with non-uniform binning. Input PDFs are required for each key background contributions

[52]. The fits are completed using MC techniques, minimizing the negative-log-likelihood (NLL) to estimate the best fit PDF weights for each toy dataset.

Each input PDF is generated by high statistics Geant4 simulations of background radiation from nEXO detector components, with estimated activities from extensive radio-assay campaigns [52], [165]. All key detector components have been assayed to determine their intrinsic levels of U/Th contamination, and the measured activities are used to scale the background PDFs by detector component mass in the profile likelihood analysis.

There are three high-level analysis variables that serve as the input parameters for the multi-parameter likelihood analysis. These three variables are selected for their discrimination power, and relevance to the $0\nu\beta\beta$ search, and are illustrated in Figure 5.6. In the top insets of Figure 5.6, the categories of “background-like” and “signal-like” are included for illustrative purposes, as the profile likelihood uses the entire input PDF for fitting and estimation.




Figure 5.6: nEXO’s multi-parameter likelihood analysis illustrated overview. Simulated background radiation spectra (1D projections) are overlaid with illustrations describing each parameter. The simulated $0\nu\beta\beta$ assumes a half-life of 7.4×10^{27} y. Left: Energy near $Q_{\beta\beta}$ with orange signal and grey background. Middle: Standoff distance in mm, with background-like components peaking near zero. Right: Machine learning estimator to identify charge clusters, $0\nu\beta\beta$ -like at 1, and background-like at 0.

The first parameter is the reconstructed energy of the event. This has clear importance as we are searching for the evidence of $0\nu\beta\beta$, which appears as an excess bump at $Q_{\beta\beta}$. This figure emphasizes the importance of energy resolution for nEXO, which is essential for constraining the leakage of $2\nu\beta\beta$ into the $0\nu\beta\beta$ region of interest.

The second parameter was introduced in the previous section, the “standoff” distance. As seen in the simulated distributions in the bottom panel (middle) of Figure 5.6 the standoff distance distributions are significantly different between signal and background. The double beta decay signal is uniform per unit volume, whereas most backgrounds originate from outside the detector and concentrate at small standoff distances. Note that a volume-uniform signal is not flat in standoff distance, since the volume at a given standoff is set by the detector geometry.

The final input parameter is related to the “topology” of each event’s charge readout sig-

nal. As Compton scattering dominates for γ 's at ~ 2.5 MeV, background-like events tend to produce multiple spatially separated energy deposition sites that sum to the full energy. In contrast to this, the $0\nu\beta\beta$ signal is expected to be localized to a single-site energy depositions. We can take advantage of this by employing a Deep Neural Network (DNN) to analyze the charge readout signals for clusters. Events with single-site charge clusters are scored as $\beta\beta$ -like approaching one, and multi-site events are scored approaching zero [52]. This technique was successfully applied for the EXO-200 full data-set analysis [152], and has been further developed, leveraging developments in machine learning analysis techniques.

5.3 nEXO Sensitivity to Neutrinoless Double Beta Decay

The nEXO experiment is projected to be sensitive to $0\nu\beta\beta$ decay in ^{136}Xe . The following section outlines the expected sensitivity based on background estimates, with the discovery potential under the inverted mass ordering discussed in Chapter 1.

The key background components for the $0\nu\beta\beta$ search with nEXO are radon-induced, cosmogenic, and γ -like backgrounds. Radon daughter ^{214}Bi is of particular concern, being a β -emitter with a Q -value extending beyond $Q_{\beta\beta}$ (~ 3.2 MeV [166]), and is expected to accumulate on the cathode and field shaping rings from drifting radon daughter ions. Intrinsic γ -radioactivity from U/Th in detector materials makes up another substantial background contribution, targeted through the use of electroformed copper for the TPC vessel [52] and a low-radioactivity nickel cryostat [167]. Cosmogenic ^{137}Xe , with an endpoint of ~ 4 MeV [168] overlapping the $0\nu\beta\beta$ region, is the primary target of the OD muon veto described in Section 5.1.1.

The key metric used to evaluate the design and material selection for nEXO is the half-life sensitivity for $0\nu\beta\beta$, and the corresponding discovery potential. This is evaluated using a frequentist binned profile likelihood method. Simulated datasets are generated from the background distributions and signal model for a given half-life, then fitted with the same

models to determine the confidence intervals. These datasets are generated with Poisson fluctuations, informed by the evaluated uncertainties on each component. While some non-Poissonian fluctuations may be introduced from cosmogenic background sources, these are considered negligible here due to their small contribution to the overall background.

Using this approach, we can estimate the sensitivity of nEXO to $0\nu\beta\beta$ and evaluate it along experimentally relevant variables such as the live-time of the experiment. Shown in Figure 5.7, reproduced from [52] is the sensitivity projection versus detector live-time over the expected duration of the nEXO program.




Figure 5.7: nEXO Sensitivity to $0\nu\beta\beta$ versus detector live-time [52]. Both 90 % C.L sensitivities at $\tau_{1/2} = 1.35 \times 10^{28}$ yr and the 3- σ discovery level at $\tau_{1/2} = 7.45 \times 10^{27}$ yr are overlaid with the final sensitivity from EXO-200 [152].

We can translate our sensitivities into an isotope-agnostic metric, the effective Majorana mass $m_{\beta\beta}$, introduced in Chapter 1, allowing for direct comparison between $0\nu\beta\beta$ searches across isotopes. The nEXO projection is compared with that from EXO-200 in Figure 5.8, alongside the allowed ranges of $m_{\beta\beta}$ derived from neutrino oscillation measurements [169],

shown as solid contours for unconstrained Majorana phases and dashed contours otherwise.

There is large variation in technique and results for evaluating the nuclear matrix elements (NME) for $0\nu\beta\beta$, required to estimate the half-life from the measured rate (or limit). This variance is reflected in the horizontal bands in Figure 5.8, where the solid band represents $\sim 68\%$ of the NME models considered.




Figure 5.8: Plot of 90 % CL exclusion sensitivity for effective Majorana mass $\langle m_{\beta\beta} \rangle$ versus the lightest neutrino mass in the normal and inverted hierarchy, figure reproduced from [52]. The width of the nEXO horizontal bands is dependent on the variance from nuclear matrix element calculations.

nEXO is being designed with a projected half-life sensitivity at 90 % CL of 1.35×10^{28} y. As a discovery-oriented experiment, nEXO requires precise characterization of detector response and rigorous background mitigation, both of which are central to the work presented in this thesis. In the next chapter, we will explore how our work with the LoLX experiment plays a key role within nEXO through optical simulations of the light collection subsystem. We will use the models of SiPM ExCT developed in this work to estimate the currently un-

constrained impact of SiPM ExCT on nEXO's light channel energy resolution, and how this correlated noise may affect other physics searches beyond $0\nu\beta\beta$.

6

nEXO Light Simulations

The successful detection of $0\nu\beta\beta$ with nEXO depends on the achievable energy resolution of the experiment. Therefore, the collection of scintillation light must be optimized to reach this goal. This is achieved in two ways: maximizing the scintillation collection efficiency and minimizing the variance introduced by correlated noise during the detection process.

The simulation of scintillation light within the nEXO LXe TPC is the topic of this chapter. We will outline the past and current simulation strategies for light collection with nEXO in Section 6.1, followed by a revisit of previous assumptions relating to the TPC design, including updating and extending the range of optical parameters in Section 6.2. In Section 6.3, we will outline the implementation, simulation, and results of SiPM ExCT within nEXO, and in Section 6.4 we will describe the SiPM correlated noise simulation model, and combining all

of these efforts together, we will give an updated estimate of the nEXO light channel energy resolution including SiPM ExCT in Section 6.5. Finally, we conclude with a discussion of the current reach of these simulations and the work required to further optimize and improve the modelling of light transport in the nEXO TPC.

6.1 nEXO Light Simulations

nEXO simulations are carried out using the Geant4 framework for particle transport and rely on dedicated optical simulations performed with Chroma. To estimate the scintillation signal for an energy deposit within the nEXO TPC, the NEST LXe model is first called to produce an estimate of the number of scintillation photons produced for the energy deposited in the surrounding region. The average photon transport efficiency (PTE) for that region is retrieved from a look-up-table (LUT) called the “light map”. By combining the PTE with the mean photodetection efficiency (PDE) of the SiPMs at VUV wavelengths (175 nm), and including the estimated excess charge from SiPM correlated noise (quantified as the excess charge, EC), we arrive at an estimate of the total photodetector signal for that event. This relationship can be expressed as:

$$\widehat{Q_{Scint.}} = N_\gamma \times \text{PTE}(R, Z) \times \text{PDE} \times (1 + \text{EC}), \quad (6.1)$$

where the estimated reconstructed light channel energy, $\widehat{Q_{Scint.}}$, depends on N_γ , the sum of photons produced in this event ($N_\gamma = \sum E_{dep.} \times LY(E, |\mathbf{E}|)$, where the NEST light yield (LY), carries units of photons/keV, and depends on the energy and electric field strength at that position). The lightmap used in the most recent nEXO $0\nu\beta\beta$ sensitivity estimate [52], is shown in Figure 6.1. This lightmap is the result of 5.0×10^{11} photons simulated uniformly within the nEXO TPC.




Figure 6.1: Simulated light collection efficiency map of nEXO TPC from [52]. Photon transport efficiency, or PTE, does not include the detection effects, but does include reflections from SiPM surfaces, which is factored from the PDE for reconstruction.

The production of the lightmap in Chroma offers a significant computational advantage, with speed-ups of up to $\sim 200\times$ expected by using the GPU-based software over the CPU-based Geant4. The map will critically depend on the simulation input optical parameters, and the level of physical detail included in the meshed model used for the simulation. In Table 6.1, reproduced from [52], are the nominal input optical simulation parameters used for nEXO sensitivity estimates, and to produce Figure 6.1. The lightmap is a key observable for the nEXO photodetector system, and can be monitored and estimated with experimental data by dissolving radioactive calibration material into the TPC [170]. Both ^{222}Rn and ^{127}Xe are being considered for this calibration source, which would be circulated into the LXe purification system periodically for continuous monitoring of the detector light collection efficiency [170]. The lightmap production and initial Chroma simulation development was carried out

Table 6.1: Reproduced table from [52], nominal simulated input optical parameters and values.

Parameter	Value (unit)
LXe Refractive Index	1.69
LXe Absorption Length	20 m
LXe Scattering Length	30 cm
TPC Reflectivity (specular)	0 %
Cathode Reflectivity (specular)	80 %
Anode Reflectivity (specular)	20 %
Field shaping ring Reflectivity (specular)	80 %
Stave and FSR Support Reflectivity (specular)	0 %
SiPM Reflectivity (specular)	25-28 % (θ dependent)

by former Yale PhD student Ako Jamil [171], whose work focused on variations of the nEXO design and the corresponding constraints on optical parameters required to reach the sensitivity goals. This work also explored the optimization variations on design aspects, including: reflective coatings for interior components, photodetector coverage and channelling, and geometry variations.

Using the simulation method described in this section, nEXO is able to make accurate and fast estimates of the light channel energy contribution, without requiring the computationally expensive simulation of optical photons in Geant4. However, by replacing the full optical simulation with the lightmap LUT, we lose significant information about the events, and their correlations with the charge signals. Additionally, the conservative and slightly ad-hoc optical input parameter values do not reflect the full picture of light transportation within nEXO. In the following section, we will revisit the key input parameters for the nEXO Chroma simulations, and expand their range beyond 175 nm to the full range of wavelengths expected for light in nEXO. While the lightmap in Figure 6.1 uses uniform photon levels sampled uniformly in the detector, the actual scintillation light production will depend on the specific electric field strength at that location. By combining simulations of the TPC field with Chroma simulations of light transport, we can investigate the variance in reconstructed energy estimates for varying

positions within the TPC.

This is especially important for the so-called “skin” region of the TPC, the area between the FSRs and the SiPM staves, where the field changes the most rapidly. This region is also of significant importance for calibration near $Q_{\beta\beta}$. Energy calibration of the TPC near $Q_{\beta\beta}$ is an important performance metric for the detector, and ties directly into the experimental sensitivity to $0\nu\beta\beta$. This will be done with dedicated regular thorium γ -source calibrations from the exterior of the inner detector, and due to the short attenuation length for γ ’s in LXe, a large fraction of these calibration events will deposit energy in the skin region. SiPM saturation and exposure effects from the high light levels are under investigation, and by upgrading our simulation tools we will be able to model these effects in appropriate detail.

6.2 Simulation Upgrades and Validation Testing

To increase confidence in our simulation results, we validate the input data and simulation tests. The ultimate test for any simulation is a direct comparison with an equivalent measurement in an experimental setup. In Chapter 4, we have demonstrated broad agreement for light transport simulations with the LoLX phase-2 commissioning data, showcasing the physical accuracy of the Chroma simulation toolkit. These essential simulation-data validation tests require robust simulation tools which are flexible for a variety of experimental apparatuses within the nEXO collaboration. Spurred by this need, and the desire for simulation cross-validation with the LoLX phase-2 detector and nEXO, we have dedicated substantial efforts in upgrading the `chroma-simulation` framework used by nEXO and collaborators.

6.2.1 Refactored `chroma-simulation` Codebase

Building on the work by former nEXO PhD Student Ako Jamil, who lead the initial development work for `chroma-simulation`, we have further developed the package to generalize beyond nEXO-specific applications. These developments focused on flexibility, abstraction,

and standardizing input and output data streams. Code base refactoring is a substantial undertaking and should only be carried out once the potential benefits are clearly recognized. The primary goals in our refactoring were: future upgradability, experiment genericism, and simulation standardization.

The pre-existing `chroma-simulation` codebase made extensive use of global hard-coded variables specific to the nEXO experiment, and with the exception of the geometry building utility, was written in a procedural and monolithic fashion. The primary simulation script was extended ad-hoc with if-else statements for new simulation methods, with minimal modularity. We have replaced this script-based implementation with a modular class-based design, standardizing concepts for I/O, position generators, and photon generators. Both previous and updated `chroma-simulation`'s make use of the human-readable data serialization file format, YAML¹. The YAML files are used for simulation specifications, geometry details, and optical data inputs. Simulation specifications include the type of position generator to use, the photon production/emission choice, number of events and photons per event, and the output file components and format. Users have the option to select which branches of the simulation results they would like: photon level, detector level, or high level. Photon level information includes the details for every photon simulated and is not often required for analysis. Detector level information includes information for detected photons and their channel distributions. High level information includes summary variables describing an entire event.

An overview schematic outlining the upgraded architecture is shown in Figure 6.2. Additionally, a high-level comparison between the original and updated implementations is outlined in Table 6.2.

¹YAML originally stood for “Yet another markup language”, but was rebranded as the recursive acronym “Yaml ain’t a markup language”, to emphasize the data oriented nature.




Figure 6.2: Schematic diagram outlining the modular components of the updated *chroma-simulation* architecture. Users run the primary simulation executable with a simulation specification YAML table. Modular classes generalize the concepts of: the geometry of the simulation, input events, output events, position generators and photon generators. The main controller delegates tasks to each module, transferring data back and forth to the central class. The modular class-based architecture increases the software complexity and number of files, in exchange for abstraction that enables flexible and extensible simulations.

The photon generation mechanisms have been extended in the *Photons.py* module, using *nestpy* for scintillation light production, and a custom re-implementation of the *Geant4* Cherenkov radiation methods into dedicated python functions. The scintillation and Cherenkov implementation may be used as standalone simulations (assuming minimum ionizing particles for Cherenkov production) or in conjunction with an input *Geant4* simulation. With expanded position generation and input file reading, we use the *upROOT* package to read output *Geant4* files with a specified tabular format (outlined in Table 2.1). Each primary and secondary particle track are sorted, with locational energy depositions and momentum information provided to *Photons.py* for (optional) concurrent scintillation and Cherenkov generation. Producing a large quantity ($O(10^7)$) of events in *Geant4* with secondary pho-

Table 6.2: High-level comparison of chroma-simulation architecture before and after refactoring, emphasizing improvements in use and simulation fidelity.

Characteristic	Original Implementation	Upgraded Implementation	Outcomes
Code organization	Long, nested scripts with ad-hoc logic	Modular, class-based architecture	Improved readability and maintainability
Event/position generation	Monolithic routines	Dedicated generator classes	Easier to extend and test individual components
Input+Output handling	Manual parameter parsing	Dedicated Reader & Writer classes for standardized YAML parsing and output preparation	Reliable, reproducible configuration
Simulation control	Central script orchestrating all steps	Controller class delegating to modules	Clear separation of responsibilities, simplified upgrades
Extensibility	Difficult to modify without breaking code	Well-defined interfaces between modules	Facilitates upgrades & future development

ton and ionization electron production disabled is very computationally efficient, decoupling the more expensive photon transport simulation. The input Geant4 simulation file can be sampled repeatedly, with indexed or random events while fine-tuning simulation parameters. This is a significant advantage during the preliminary stage of simulation and development, as program start-up times for Geant4 are much longer.

Beyond the core simulation code upgrades, we have also created a number of helper scripts to support specific simulations. One such helper script is a tool for converting .csv optical data tables to the required .yaml format to help avoid user input errors. Other helper scripts include: example G4 $\rightarrow$ Chroma table generation scripts, E-field interpolator production scripts, benchmarking and simulation validation scripts, and more. We have also prepared the groundwork for continuous integration/deployment and validation for nEXO optical simulations, which will be implemented when simulations, and detector design are in their final stage.

6.2.2 Updated Simulation Optical Parameters

To simulate SiPM ExCT within nEXO (among other things), we must first extend the simple single-valued optical simulation parameters used in [52], and described previously, to the full wavelength range.

LXe Refractive Index and Scattering Length

One of the key inputs for the optical simulation in Chroma are the optical properties of the bulk LXe that surrounds all detectors components. Within the Chroma framework, this is specified as the “outside” material for all imported meshes of detector components. In this role, photons cast from within a component mesh into the bulk LXe undergo refraction governed by the Fresnel equations, based on the wavelength-dependent refractive index curve applied to the LXe material. For accurate simulations, we require a wavelength-dependent refractive index curve in the range 160–1000 nm. Another important input is the attenuation length, which is represented in Chroma with two components, a Rayleigh scattering length, and an absorptive attenuation length. In Rayleigh scattering, light scatters elastically, changing directions and polarization but preserving energy. Conversely, absorptive attenuation in LXe is determined by the purity level achieved, as the scintillation light is absorbed by impurities in the liquid. We use the nominal 20 m absorption length from earlier nEXO studies for all wavelengths.

The refractive index of LXe at scintillation wavelengths has been measured several times, with varying results. A value of $n = 1.69 \pm 0.02$ was measured in 2004 by Solovov et al. [172], at 170 K with a simultaneous measurement of the attenuation length of 369 ± 18 mm. Previous measurements at the triple point, with $\lambda = 180$ nm, indicated a lower value of $n \approx 1.56$ [173].

To harmonize the refractive index over the entire wavelength range, we stitched together two datasets from empirical modelling of refractive index, based on the approach from [174], with LXe refractive index data from [175]. The LXe refractive index optical input data curve


shown in Figure 6.3a is a combination of the 178 K fit data from [174], with the Sellmeier equation,

$$n^2 = a_0 + \frac{a_{UV}\lambda^2}{\lambda^2 - \lambda_{UV}^2} + \frac{a_{IR}\lambda^2}{\lambda^2 - \lambda_{IR}^2}, \quad (6.2)$$


used to extrapolate beyond the maximum reported range (~ 650 nm), where a_0 , a_{UV} , a_{IR} , λ_{UV} , and λ_{IR} are material-dependent Sellmeier coefficients describing the contributions of UV and IR resonances to the refractive index dispersion. We use the fitted values from Table 2 of [174] as inputs to the Sellmeier equation, but opt to smooth over the estimated IR resonance, which has not been measured to-date.

The produced LXe refractive index curve, shown in Figure 6.3a, is used to generate the Rayleigh scattering length distribution shown in Figure 6.3b. The thermodynamic parameters used in the model: density (ρ), and isothermal compressibility (κ_T), were not included in [174] and were instead determined using the reported Rayleigh scattering length at scintillation wavelengths. The Rayleigh scattering length according to [174] is calculated as:

$$l^{-1} = \frac{16\pi^3}{6\lambda^4} \left[kT\rho\kappa_T \left(\frac{(n^2 - 1)(n^2 + 2)}{3} \right)^2 \right]. \quad (6.3)$$



(a) Refractive index of LXe with inset showing VUV range zoomed-in.



(b) Rayleigh Scattering Length of LXe with inset showing the VUV range zoomed-in.

Figure 6.3: LXe optical simulation input data distributions.

Table 6.3: VUV Reflective coating material stack design, used in POCAR thin-film simulations to determine optical behaviour of VUV reflective coating.


Material	Extent
LXe	Exterior
MgF ₂	100 nm
Aluminum	100 nm
MgF ₂	20 nm
Copper	Interior

Cathode and FSR VUV Reflective Coating


One major design choice, already established in [52], is to coat the TPC cathode and FSR surfaces with a custom reflective coating, optimized for VUV light. The nEXO requirement of 80 % reflectivity for this surface is based on the targeted light-channel energy resolution of the inner detector ($< 1 \%$ FWHM/ μ). The reflective coating planned for nEXO is a stack of MgF₂-Al-MgF₂ deposited on the copper surfaces of the TPC. The thicknesses of each layer are given in Table 6.3. Preliminary validation measurements of the planned coatings have already been carried out by collaborators [176], with total reflectivities in GN₂ in the range of 84–92 % at VUV wavelengths for different substrate surface treatments. The initial measurements show the viability of the reflective coating for increased light collection in nEXO at VUV wavelengths. However, for the optical simulations we require a full wavelength-dependent description. As Chroma does not currently support thin-film effects, we use a standalone thin-film simulation tool, POCAL (Python Optical Coating Analysis Library [177]), to estimate the full wavelength range (160–1000 nm) total reflectivities of the coating stack. POCAL offers a flexible python implementation of the thin-film simulation needed to estimate the total reflectivity of the coatings.

Using input refractive index curves for MgF₂ [178], Al [179], Cu [180] and LXe, we run a wavelength sweep from 160 nm to 1000 nm. Chroma does not currently support 2D LUTs for optical properties with angular and wavelength dependence, so we must produce a single

wavelength-dependent curve that is representative. The input curves are shown in Figure 6.4a, for a variety of polarizations and incident angles, with significant variations. We average the different polarization curves in Figure 6.4a, assuming random polarization of light from LXe scintillation, and combine the incident angle curves with a $\sin \theta$ weighting, which accounts for the solid-angle distribution of isotropic light impacting the surface. The resulting average polarization and incident angle distribution is shown in Figure 6.4b, with the inset plot showing the values in the VUV range. This procedure is intended as a place-holder until witness samples of the final nEXO coatings are available for measurements, and broadly agrees with the values from [176] in the VUV range. The measurements in [176] report significant variations between specular vs diffusive reflectivity, depending on surface preparation. For nEXO simulations we assume only specular reflectivity and provide a complementary absorption probability curve to prevent transmission into the coating volume.



(a) MgF_2 total reflectivity for varying incident angles and polarizations produced using POCAL [177].



(b) Average polarization total reflectivity of $\text{MgF}_2\text{-Al-MgF}_2$ coatings, normalized by isotropic incident angle probabilities at surface.

Figure 6.4: Total reflectivity of $\text{MgF}_2\text{-Al-MgF}_2$ coatings modelled using POCAL [177].

Extending Optical Parameters & Test Environment

The final design of the nEXO experiment is still under development, including components that have significant impact for light simulations. One of these components is the charge read-out tiles that are suspended below the anode plane. The current design, outlined in Chapter 5,

has a surface that is primarily gold, with SiO_2 spacers and substrate. We aim to maintain high performance in the optical simulations with Chroma, where ray casting computational time scales with the number of triangles in the geometry, requiring careful management of mesh file complexity. Modelling the $\mathcal{O}(10\text{ }\mu\text{m})$ details of each anode tile would increase simulation time substantially². To circumvent this constraint, we studied the optical transport properties of the anode tile in a stand-alone test environment. The test environment is shown in a CAD rendering in Figure 6.5 containing only three components. The outer sphere serves as the “detector”, collecting all photons with 100 % efficiency. Inside centred at (0,0,0) cm is an anode tile, with all gold pad coatings, and all SiO_2 treated as separate optical surfaces. These three components were loaded into Chroma in a “Testing” experimental configuration, and two different test simulations were carried out.




Figure 6.5: Simplified Chroma simulation test and validation geometry. With the anode tile centred at the origin, light sources are generated within the spherical detector volume that records the position, incident angle, and time of each photon reaching the surface of the inner shell. This geometry was used to validate Chroma simulation physics, and characterize optical transport properties of materials.

The first simulation was to determine if the initial non-angular dependent parametrization (wavelength-dependent or otherwise) was sufficiently detailed. In this simulation, a “laser”

²How much is hard to say, during testing of the geometry with full anode tiles, the program crashed my workstation during geometry building after consuming +200 GB of VRAM.

photon source was used, beaming a ~ 1 cm spot at the centre of the anode tile from varying positions on the sphere for a representative set of wavelengths. This allows us to estimate the angular dependence of the tile reflectivity, with a large enough spot to average over any surface feature effects. For these tests, we use the complex refractive index data for gold [180] over the wavelength range 160–1000 nm, while SiO_2 is modelled as a dielectric material with real-valued refractive indices. The results of these simulations for three representative wavelengths are shown in Figure 6.6. We can see that a large variation for different wavelengths is apparent, necessitating a full wavelength description of the anode’s reflectivity. We also see good agreement with the theoretical expectation from the Fresnel relations at 175 nm, which also agrees with the nominal value of 20 % at normal incidence.






Figure 6.6: Geometry Chroma simulation testing for a nEXO anode tile model. Incident angles are with respect to the surface normal of the anode tile. Blue simulation data-points are for the simplified anode design and single-valued 20 % reflectivity implemented in [52]. Yellow points for 175 nm incident light are overlaid with the expected results from Fresnel relations using the same input data, showing good agreement for the simulation. A large variation in expected reflectivity across incident angles and wavelengths is apparent.

The second simulation test was to estimate the effective surface coverage for the gold and SiO₂ materials individually. This was done using an isotropic photon source (photon-bombs) with wavelength of 175 nm, positioned distant from the anode tile to sample a large range of incident angles, similar to the conditions within nEXO. From these tests the effective surface coverage for gold was $\sim 95\%$, and $\sim 5\%$ coverage for SiO₂. Based on these two test simulations, we conclude two things: the fully complex Frensel representation for gold is required to capture the full behaviour of the anode tile optics, and the small effective coverage of SiO₂ has a negligible contribution to the surface properties at this stage. A full treatment of the optical properties of the anode will require samples to be measured with broad spectrum light sources to cover the range of wavelengths within nEXO, ideally within LXe in operational conditions. This will also require a completed implementation of 2D LUTs for optical properties in Chroma.

Lastly, the input to Chroma for optical simulations is the full wavelength range of gold complex refractive indices from [180]. These distributions, while sparse, are shown in Figure 6.7, with VUV data range shown in the inset. Chroma will linearly interpolate between points for optical parameters, but does not extrapolate.



(a) Gold Refractive Index - Real.



(b) Gold Refractive Index - Imaginary.

Figure 6.7: Full wavelength-dependent gold anode surface refractive indices.


In addition to extending the optical properties for the anode, we also relax the assumption that the copper components within the inner detector are fully absorptive. The reflectivity data

for copper was taken from vacuum measurements [181], [182], stitched together for a complete dataset. We also translated the SiPM optical and PDE implementation from Section 3.3 into Chroma. This method, described in detail in [103], relies on recent Chroma developments by nEXO and LoLX collaborator, Alex Li, who implemented the 2D LUT (θ, λ) efficiency data loader. This technique provides the expected SiPM detection efficiency, and reflectivities, for the full range of photon wavelengths.

Incorporating E-field Simulations in chroma-simulation

For completeness of our optical simulations, we would like to incorporate the position dependent electric field strength within nEXO. As discussed earlier in Chapter 2, the scintillation light yield depends strongly on the local electric field strength. Within the fiducial volume (FV) of the nEXO TPC, the field varies only marginally by design, but changes quickly outside the FSRs. To model this effect, COMSOL Multiphysics[®] simulations of the TPC with a 1/4 azimuthal projection were performed by collaborators at Pacific Northwest National Laboratory (PNNL). A high density map of the field strength and direction at 3D positions within the TPC was used to create the two projections shown in Figure 6.8. The field data is normalized to 400 V/cm, and displayed on a power-scale to emphasize the homogeneity within the TPC FV regions ($R < 0.56$ m) and, the exaggerated field strength outside the FSRs. While the field map has a high density of points near the rings, there are voids in the interior of the xenon volume which require us to smooth and interpolate the data for use in optical simulations. A smoothed version of Figure 6.8b, averaged over ϕ and restricted to the FV, was used for subsequent simulations of light production variability within the nEXO TPC in Chroma.

³E-field simulation data provided by Luca Pagnini (PNNL), private communication.



(a) Normalized E-field strength in the nEXO TPC displayed on a power-scale in 1/4 polar projection. We see that beyond the FSRs ($R > 0.57$ m), in the “skin” region, the field strength varies significantly, reaching $O(10^4$ V/cm).

(b) Cylindrical projection of nEXO TPC E-field, averaged over ϕ . Below the cathode plane, and outside the FSRs, the so-called “skin” region, the field increases and fluctuates significantly.

Figure 6.8: nEXO 1/4 TPC E-field simulation in COMSOL Multiphysics[®] shown in radial and cylindrical projections. ³

6.3 Simulating SiPM ExCT in nEXO

With upgraded simulations, and a fully detailed optical input dataset, we can now prepare for simulations of SiPM ExCT in the nEXO detector using Chroma. These simulations will be used to evaluate the average amount of excess charge induced by ExCT in the nEXO inner detector. We also investigated the spatial dependence on the transport probability of NIR secondary photons in nEXO.

6.3.1 ExCT Implementation in chroma-simulation

To simulate ExCT with Chroma, we need a dedicated position generator to create event vertices, and a dedicated photon generator to produce the anisotropic light emission. As outlined in Chapter 3, we make use a high-density map of in-LXe secondary photon emission probabilities from our simulations of SiPM ExCT in LoLX, shown in Figure 3.9, to generate the emitted photon distribution. Since we require the position vertex generation to work for any

detector configured in `chroma-simulation`, we take a generic approach for vertex generation. Added as an additional method to `Generator.py`, we first define a case-switch method for the physical extent of different SiPM board designs. For nEXO, the base simulation photodetector unit is a (20×20) cm² “tile” (modelled as a perfectly flat rectangular prism in simulation). This defines a plane of “avalanche positions” which is randomly sampled for vertices based on the requested number of avalanches in the ExCT simulation. For each simulation, the user may request a single photodetector channel for emission based on its unique ID. We then sample from the prescribed plane to create vertices. Afterwards, we must translate and rotate these positions into the global coordinate system set by the detector geometry. This is done by flagging all meshes that are configured as “detectors”, and using the triangle-level information to identify the centre and surface normal of the detector. The results of this procedure are illustrated in Figure 6.9, for all nEXO `chroma-simulation` channels. Once the global detector centre and surface normal are defined for the requested channel, we use Euler’s method to rotate the generated vertices into the global coordinate system, and encode the channel’s surface normal in metadata for use by the photon generation algorithm.




Figure 6.9: 3D visualization of all nEXO chroma-simulation channels, with centres and surface normals illustrated. Each SiPM channel is a “tile” in the simulation, and may be requested individually for ExCT simulations. ExCT is generated about the centre positions (blue) with emission profile dependent on the direction of the surface normal for that channel (red).

The ExCT photon generator was created as an additional option in the photon generator module class as `ExtCrosstalk`. The input map that was converted into a 3D data-table $(\theta, \lambda, P_{emit})$ is first converted into a 1D cumulative array ordered by probability. Next, the cumulative array is converted into a cubic spline interpolator for sampling. Every simulation request specifies the number of photons to be emitted per “avalanche” vertex, N_γ , and the algorithm samples the input ExCT map interpolator N_γ times. The output of this algorithm is compared to the input data and shown in Figure 6.10.




Figure 6.10: ExCT emission profile input data, sampling method, and validation test. Left: Input data for emission profile of secondary photons into LXe from [116]. Middle: Results of random sampling method implemented, normalized. Right: Ratio of Sampled/Input data for each bin in the left and middle panels, values are distributed primarily around 1, indicating good agreement for the implemented sampling routine. The edge effects visible at high and low emission angles are artifacts of the extremely low emission probabilities in those regions.

6.3.2 Validation Testing and Spatial Dependencies

Initial validation testing for the implemented ExCT simulation was carried out with the optional “photon-tracking” functionality of the chroma-simulation package. With photon-tracking enabled, every step for every photon is saved during ray-casting propagation, which quickly becomes untenable at large photons numbers. Running ExCT simulations with photon-tracking readout enabled allows us to validate the emission directions to independently validate the anisotropic emission profile implementation. Additionally, we can propagate a large quantity of photons ($\mathcal{O}(10\text{ M})$) to produce an “ExCT lightmap”, which is unique to a single SiPM tile in nEXO. Each tile will produce its own unique lightmap, but owing to the azimuthally symmetric design of nEXO, we should expect significant symmetries for tiles in any row of the detector. An example ExCT lightmap is shown in Figure 6.11, from the easterly tile on the middle band of the nEXO barrel. We see several interesting features in Figure 6.11. First an intense spot centred about the origin position, driven by reflections of the secondary NIR photons off the field cage back onto the emissive and neighbouring tiles.

Secondly, we can see two large features from direct emission detection, centred around the $\sim 45^\circ$ mode of the emission distribution in Figure 6.10, and several low-intensity spots from shadowing by field cage components.




Figure 6.11: ExCT hit pattern in nEXO, projected in z - ϕ from mid-plane SiPM tile. Average detected photon position from ExCT on tile ID#584, easterly (+X) position in the centre-plane of the stave barrel. The labelled origin is the position of the first sampled photon location on the emissive tile.

A second validation step for the ExCT generation simulation is through inspection of photon travel times through the TPC. Due to the long scattering length ($\mathcal{O}(100 \text{ m})$) in the NIR band for LXe, shown in Figure 6.3b, we expect a “beat” pattern in the arrival time distribution, representing the travel time across the TPC barrel. For reference, the round-trip time for a 650 nm photon across the nEXO TPC barrel with a diameter of $\sim 1.2 \text{ m}$ is $\sim 5.5 \text{ ns}$. We observe this effect clearly in the simulation. In Figure 6.12, a zoomed-in travel time distribution for ExCT within nEXO is shown. We also observe a unique timing feature arising from the prompt and uninterrupted light propagation from SiPM ExCT. For event vertices within $\sim 120 \text{ mm}$ of a stave, VUV scintillation light will travel to the nearest tiles, producing secondary NIR photons which will propagate uninterrupted to the far side of the TPC, arriving

before the initial VUV photons. This is mainly driven by the significant difference in refractive index between the NIR secondary photons and VUV scintillation light (~ 1.35 vs ~ 1.69 respectively). This unique feature could potentially be taken advantage of for analyses that require very low thresholds, such as in DM detectors [130], but is inaccessible in the baseline nEXO electronics readout design, with sampling times longer than the separation.




Figure 6.12: Prompt photon travel time distribution for ExCT photons in nEXO. The peak at ~ 1 ns corresponds to reflections by field cage components directly in-front of the SiPM tile, and subsequent ~ 6 ns and ~ 12 ns peaks align with the round-trip time for NIR photons across the nEXO TPC diameter.

By comparing the photon travel time and mean detection efficiencies for SiPMs across the azimuth and height of the nEXO TPC, we can look for potential symmetries. In Figure 6.13, the photon travel time for detected photons from different tile positions are overlaid, with representative samples. The chosen positions are the four cardinal directions around the centre

plane of the nEXO staves, and 4 additional positions across the height of the northernmost stave.




Figure 6.13: 2-panel comparison of detected ExCT photon travel time distributions in nEXO TPC. Left: Comparison of travel times for detected photons at four positions around the azimuthal axis at the centre-plane of the nEXO TPC. Right: Comparison of travel times for detected photons at five positions along a single stave. No significant travel time differences are present, and a minor difference is observed only for the lowest tile, indicating a minor Z-dependence.

In order to model the ExCT noise contribution for nEXO without a dedicated wavelength-dependent transport map for each SiPM position, we investigate the Z and wavelength dependence of the NIR ExCT transport. In these simulations, the internal PDE for all wavelengths are set to unity, capturing only the transport effects of the detector (including SiPM optics). The results are shown below in Figure 6.14, for five tile positions along the northern SiPM stave. We can observe the expected minor variation in transport efficiency, with a larger variance near the bottom of the TPC due to the large expected variation in the VUV reflective coating (shown in Figure 6.4b). We also observe a $\sim 2\text{-}3\%$ absolute variation in wavelength averaged ExCT transport probabilities along the height of the stave. The observed U-shape is likely driven by the higher NIR reflectivities near the cathode and anode, while the anode is

only moderately reflective in the VUV. It is highly reflective in the NIR as seen in Figure 6.6. Additionally, PTE variation for tiles about the azimuth in-plane are negligible, cementing the expected azimuthal symmetry in the cylindrical TPC.




Figure 6.14: Simulating wavelength and Z -dependence of SiPM ExCT transport in nEXO. Left: ExCT PTE versus emission wavelength for five positions along a single stave in the nEXO TPC. Right: Bar plot of wavelength averaged PTE (%) versus position. A mild wavelength dependence in transport efficiency is observed, and a non-monotonic change in average PTE versus Z position is observed, with an absolute variation of $\sim 2\%$.

We have demonstrated and validated SiPM ExCT simulations for the nEXO experiment, informed by our work for the LoLX experiment. We can make the following conclusions relevant for including ExCT in a SiPM noise model: 1) SiPM ExCT transport is azimuthally symmetric in nEXO. 2) A wavelength-dependent efficiency is necessary to capture the dispersion expected from simulations. 3) A minor non-symmetric Z -dependence is present, ne-

cessitating a position-dependent parametrization of the transport efficiency. While there is significant dispersion in the detected positions of ExCT emission, as seen in Figure 6.11, we opt to omit this in modelling the excess charge from ExCT. This is because of computational constraints, with hundreds of tiles, each with their own wavelength-dependent transport efficiencies, adding in the second order effect of hit photon locations is untenable. This simplification will only mildly affect second-order ExCT, or Seed Photon→ExCT→ExCT events, which are at the $\mathcal{O}(0.01 \%)$ contribution level.

6.3.3 ExCT Induced Correlated Noise in nEXO

To evaluate the impact of ExCT in nEXO we have parametrized the average transport and detection efficiency (Collection Efficiency or CE) over avalanche Z-position and secondary photon emission wavelengths. The average CE as a function of height and wavelength are shown in Figure 6.15. For application in subsequent simulations, we apply a Gaussian kernel smoothing, to produce a version of this map as an interpolator for subsequent SiPM response simulations, described in the following section. This is also shown in Figure 6.16 as profile histogram projections, averaging over Z-position or wavelength. In these projections, we see the expected behaviour seen from earlier, with a bimodal distribution of collection efficiency as a function of height in the TPC, we also observe that this dependence is much weaker for NIR photons above ~ 800 nm.




Figure 6.15: nEXO SiPM ExCT average collection efficiency in Z and λ . This distribution is smoothed with Gaussian kernel smoothing to produce a 2D `RegularGridInterpolator` for ExCT CE sampling in subsequent analysis.




Figure 6.16: Profile histogram projections of nEXO SiPM ExCT collection efficiency. Left: Avalanche Z -position dependent CE, averaged over wavelengths. Right: Emission wavelength-dependent CE averaged over Z -position. Uncertainty bands are $1\text{-}\sigma$ standard deviation from simulated datasets.

We can combine the map of collection efficiency versus wavelength, with the emission profile of ExCT secondary photons in LXe from Figure 3.6, and the expected emission yield from Chapter 3 to estimate the SiPM ExCT “Correlation probability”. The correlation probability is the probability that, given a SiPM avalanche in nEXO (dependent on OV), we detect an emitted correlated photon anywhere else in the detector. We can model the emission

strength as a linear response with gain, and so overvoltage, using our two different data-points from [116]. This gives us an expected yield of ~ 0.13 photons per volt of OV. Putting this together with the Z-position averaged emission collection efficiency from Figure 6.16 (assuming uniform avalanches as a function of Z) we get the expected correlation efficiency as function of overvoltage in nEXO, shown in Figure 6.17.




Figure 6.17: Expected correlation probability per avalanche for HPK SiPM ExCT in nEXO as a function of SiPM OV. Y-axis represents the mean probability for a SiPM avalanche anywhere in nEXO to produce a correlated avalanche anywhere else. Uncertainty on data points is the relative uncertainty on the expected yield per volt of OV from [116].

As we can see in Figure 6.17, the level of ExCT correlation probability is quite low overall, resulting from low emission, and only moderate collection efficiency. The overall level of correlation probability is roughly 1/4 of that from [116], reflecting the presence of much more obstruction in the near-field of view of the SiPMs (in particular from FSRs). This places SiPM ExCT at approximately 1/10th of the direct optical cross-talk contribution, similar to measurements for LAr-based DM experiments [129], where external-to-internal crosstalk acceptance ratios of 9–15 % are obtained across various detector geometries and estimation

methods (estimated from Table 1 of that work).

6.3.4 nEXO Light Collection Efficiency Map

With our expanded and upgraded optical model for light transport in nEXO, we revisit the lightmap from Figure 6.1. As we now include all major effects impacting detection in the optical simulation with Chroma, we can not trivially factor in or out a single-valued PDE estimate, as it has both strong angular and wavelength dependencies. In this case, just as we did for SiPM ExCT emission, we combine the transport and detection effects together into the parameter “Collection Efficiency”, which represents the average probability that a photon from a specific generation source, for example SiPM ExCT, is transported and detected, as a function of position within the detector. This collection efficiency lightmap is required to reconstruct the estimated number of scintillation photons produced using the estimated number of detected photons in a given event. Since the total light collected in an event is degenerate between its emission intensity and its position within the nEXO TPC, this is an essential step to avoid biased energy estimates. We use the nominal simulated SiPM response of 3 V overvoltage input data for both HPK and FBK configurations in the Chroma simulations.

To produce the collection efficiency lightmap, we use the stand-alone NEST generator in Chroma to produce “scintillation photon-bombs”, from 1.2 MeV β ’s. This produces events with the light yield variance expected from real detector events, and a timing and wavelength distribution for LXe scintillation. We perform a biased position sampling, oversampling in the low radius region to have sufficient samples for the entire fiducial volume. The whole simulation has $O(1 \times 10^5)$ positions sampled within the TPC FV, each producing $O(1.0 \times 10^4)$ photons, giving an estimated $\sim 2 \times 10^9$ photons in a lightmap. This procedure is repeated for both SiPM vendor configurations. The collection efficiency maps in 2D, restricted within the FV, are shown in Figure 6.18, with projections in R and Z shown in Figure 6.19. The projection averages were restricted within the FV for this study. The 2D collection lightmaps

are generated first as 2D heatmaps, then smoothed with a narrow Gaussian kernel smoothing to produce a `RegularGridInterpolator` for post-analysis evaluations.




Figure 6.18: nEXO Light collection efficiency maps for FBK and HPK SiPM configurations. Simulated scintillation photon bombs are generated through the nEXO FV, and efficiency is determined by the ratio of collected (transported + detected) to produced photons per event, for each location. With sufficient statistics, a 2D heatmap is produced, then Gaussian kernel smoothed to produce the lightmaps shown here.




Figure 6.19: R and Z projections of the average light collection efficiency in nEXO for β scintillation at 1.2 MeV, for FBK and HPK SiPM configurations at 3 V OV.

As shown in Figure 6.19, the nEXO FV target for light collection of efficiency of 3 % is not met at 3 V OV for either SiPM configuration. We carried out exploratory simulations to better understand the significantly lower estimated light collection efficiency in the updated

simulations compared to earlier projections. We identified two main contributors to transport light loss, surface absorption on the gold-coated charge collection tile at the anode surface, and bulk attenuation in LXe. A side-by-side comparison of absorption locations for bulk and surface absorbed processes is shown in Figure 6.20 for the FBK configuration, with the updated simulations. The effect of increased light loss during transport is combined with the updates to the SiPM response using 2D LUT's for more accurate in-LXe response, resulting in lower effective PDE than expected from normal incidence vacuum testing [139].




Figure 6.20: Absorbed photon positions in R-Z with the updated nEXO simulations for the FBK SiPM configuration. Left: Heat map of bulk absorption flagged photon locations. Right: Heat map of surface absorption flagged photon locations. On the left we see a uniform distribution from LXe absorption with a small excess at the field rings due to the shorter attenuation length of the SiO_2 spacers. On the right we see the excess of surface absorption at the anode tiles at high Z. The total fraction of created photons absorbed by the bulk process is $\sim 78\%$, and surface absorption is $\sim 20\%$.

These are the first full-wavelength simulations of the nEXO light detection system and require independent validation. A potential explanation for the increased bulk absorption fraction is the enhanced reflectivities in the updated model. In earlier simulations, non-SiPM surfaces were treated as perfectly absorbing, meaning photons were killed on first contact with any non-SiPM surface, reducing the impact of bulk absorption as a loss channel for photons. Collection efficiency was estimated using an average PDE correction based on normal-incidence measurements in vacuum, which underestimates the angular dependence

of reflectivity at the SiPM surface. Introducing physical reflectivities, including at the SiPM surface itself, increases the mean photon path length substantially. Now a significant fraction of photons that would otherwise have been counted as lost to non-SiPM surfaces instead survive additional reflections only to be bulk-absorbed in LXe before reaching a SiPM. We can consider this together with the now wavelength-dependent treatment of Rayleigh scattering length, which is ~ 35 cm at 177 nm, and follows a λ^4 scaling. Previous estimates assumed mono-chromaticity for LXe scintillation, while we now sample the emission spectrum for each photon [183], leading to ~ 25 % of emitted light experiencing < 30 cm scattering lengths. The consequence of a short scattering length compared to the $\mathcal{O}(\text{m})$ detector size leads to scintillation photons “random walking” out to the surrounding surfaces. The gain in collection efficiency from reflective surfaces is now partially offset by enhanced bulk attenuation, an effect suppressed in the absorbing-surface approximation. The exploratory simulations suggest targeted R&D is necessary to validate simulations with independent testing setups, ideally in LXe, and emphasize the need for robust full wavelength-dependent treatment of optical surfaces. One key input is the LXe absorption length which is experimentally challenging due to the much shorter scattering length, and strong purity dependence [184].

We also investigated the impact of the TPC E-field variance on scintillation production. Using the 2D field-map interpolator and the `Chroma nestpy` scintillation model, we simulated ~ 1 MeV beta energies and compared the produced distributions of created photons to a static field estimate of 400 V/cm. With ~ 10 k positions sampled throughout the TPC FV, we see no significant variance added from the varying E-field as expected in the baseline design. Employing a 2-sided KS-test, we saw a maximum difference of ~ 0.02 , corresponding to a p-value of 0.77 for the sampled distributions to come from the same prior. While this is true within the FV where the E-field only varies by ~ 15 V/cm at most, the ability to simulate skin region events, where E-field variances are much larger, is now possible. This will be used for detailed studies of SiPM saturation during external gamma calibration in the future.

6.4 SiPM Response Modelling

With our parametrized SiPM ExCT model prepared, we have all the ingredients required to estimate the impact of SiPM correlated noise on nEXO light channel resolution. We would also like to be able to analyze the contribution of each noise individually. This is essential for dedicated studies where photon timing information is required. Additionally, we also require the SiPM response to be easily integrated with the Chroma simulations, and with the Python-based detector response framework used by nEXO called `Strax`. To achieve all of these aims, we developed a modular Python-based SiPM response model library (SRML), which can easily be installed through Python package managers such as `pip`.

Created using test-driven development techniques, the SiPM response model library, or SRML, has several key components:

- `model.py` — Central SRML class implementation, with I/O, configuration and input data parsing, and core simulation loop code.
- `noise.py` — Module with each individual noise source modelled as individual functions with the required input information outlined.
- `event.py` — Custom `DataClass` objects for “Hit” (detected photon), “Channel” (collections of “Hits”), and “Event” (collections of “Channels”).
- `efficiency.py` — Models for SiPM efficiency simulations.
- `data/` and `config/` — Directories for storing configuration information about detector being simulated, and input data tables (Yaml).

An overview of the process used for nEXO SiPM response modelling is shown in Figure 6.21, with key stages highlighted. The SRML is used to simulate stochastic SiPM correlated noise on an avalanche-by-avalanche basis, and also produces the expected dark noise

contribution if the event window length is provided. This model works under the assumption that the SiPMs are in “photon-counting” mode, far away from saturation effects. Each detected photon is first converted into a “Hit” object for simulation in the loop, where each noise source is generated independently based on the provided input data. The input data for nEXO SRML simulations was taken from [101]. OV dependent probabilities for DiOCT and correlated delayed avalanches (a combination of DeOCT and AP), were scraped from the distributions in Figures 11 and 13 from [101], with exact values for SRML calculated with piece-wise spline interpolation. Figure 11 from [101] gives the number of correlated delayed avalanches per primary avalanche in $1\ \mu\text{s}$, which must be converted into a probability of at least one correlated delayed avalanche, using $P_{1,CDA} = 1 - e^{-N_{CDA}}$, where N_{CDA} is the reported mean number of correlated delayed avalanches. The inverse relationship between AP avalanche charge and time was scraped from Figure 3 from [101], combined with the approximate exponential decay time derived from the excess charge and the underlying charge-vs-time relationship. This is an important consideration, as the excess charge added per avalanche for AP is typically < 1 , unlike the other noise sources which produce a full gain avalanche. In addition to proportional charge, subsequent avalanches triggered by AP parent hits have an effective probability that scales with the avalanche gain linearly. For example, an AP with 0.2 charge units, will reduce the nominal probability for DiOCT from 5 % $\rightarrow$ 1 %. Dark noise is added in the primary hit stage, assuming a $3\ \mu\text{s}$ event window and scaling the rates from [101] in Hz mm^{-2} to the expected SiPM surface area of nEXO, $4.608\ \text{m}^2$, adding only a few dark avalanches per event on average.




Figure 6.21: SiPM Response simulation overview flowchart for nEXO simulations. nEXO simulations include PDE as part of the optical transport stage, and down sampling by efficiency checks is omitted here. Each detected photon is converted into a primary “hit” and structured for response simulation. Once all noise types have been checked for a given primary hit, the randomly generated noise hits are added to the stack of hits and the original primary hit is flagged as complete. The simulation proceeds until all hits in the stack have noise generated, allowing for run-away under some noise conditions.

To model ExCT in the SRML, we generate an emission wavelength sampler using the inverse transform technique [185], with input data from Figure 3.6, and sample detections using the ExCT light collection efficiency map from the previous section. Due to the lack of in-LXe measurements of FBK SiPM ExCT, we make the simplifying assumption that the yield and emission profiles are the same as HPK. This is likely an underestimate of the emission yield, which in vacuum measurements from [115] shows a $\sim 25\%$ higher yield for FBK compared to HPK, while the absolute estimated yields are in tension with our results from LoLX [116]. This underestimated yield is covered by the overestimate of light collection efficiency from using the CE map from the HPK configuration, as the FBK SiPMs are only $\sim 60\text{--}80\%$ (λ -dependent) as efficient at detecting the NIR SiPM ExCT emission. This will serve as a

preliminary conservative estimate of the impact of SiPM ExCT in nEXO for an FBK vendor selection, until follow-up measurements of SiPM ExCT from FBK are completed with LoLX phase-2 as described in Chapter 4.

At the time of writing, the SRML contains only pseudocode for running a subsequent frontend electronics response and noise simulation. The front-end electronics of the nEXO photodetector readout system is still under design, so we omit some important effects at this stage. An ideal simulation would include the single avalanche charge fluctuations, a simulated SiPM waveform with readout noise and simulated front-end shaping effects, and digitization effects. Future work on nEXO light channel simulations will replace the pseudocode in the SRML with a simulated waveform generator using the measured noise and front-end electronics response. For the purposes of this analysis, we represent each primary and correlated hit as a single unit of charge (1.0), while after-pulses may have less than one unit of charge due to SPAD resetting effects.

For comparison to our input data, we define two quantities. The excess charge factor (ECF) is defined as $ECF = 1 + \hat{\Lambda}$, where $\hat{\Lambda}$ is the mean excess charge from correlated avalanches in $1 \mu\text{s}$ per primary avalanche. The correlated avalanche fluctuation (CAF) is then defined as $CAF = \sigma_{\Lambda}/ECF$, where σ_{Λ} is the RMS of the excess charge distribution Λ per avalanche.

From the SRML we can make additional observations about the expected noise contributions to nEXO events, on a per-event basis, by increasing the number of primary hits per event. Since we model noise on a per-hit basis, including both probabilistic contributions (like early successes/failures in discrete Bernoulli trials) and secondary effects such as crosstalk-induced excess hits. At low numbers of detected hits per event, this results in broad, bi-modal distributions of the per-event excess charge, reflecting the stochastic nature of individual noise realizations. As the number of hits per event increases, these independent contributions average out, and the variance of the per-event excess charge decreases. This is consistent with

the expected $1/N$ scaling from the central limit theorem (CLT) effect. This behaviour is captured in Figure 6.22 and Figure 6.23, showing the evolution of the mean of the excess charge distribution with hit counts, and in the changing shapes of the per-event charge distributions. This is a key observation made possible by performing MC for the noise response using our SRML, where we know the true number of primary avalanches per event. The distributions below are simulated for nEXO, including the expected dark-noise contribution in the $3\mu\text{s}$ event window, this leads to much higher CAF and variance for low N_{hit} trials, as dark hits may produce correlated noise hits, but are excluded from the primary hit count. This normalization choice was motivated by the need to estimate the nEXO-specific single avalanche detection fluctuation levels, useful for low-threshold and no-threshold analyses. This normalization explains the higher mean ECF for low N_{hit} for HPK despite lower mean noise at high numbers of avalanches per event, as the DCR at LXe temperature for HPK is more than $2\times$ that of FBK.




Figure 6.22: FBK 3 V ECF distributions from SRML for varying number of primary hits per event. As we increase the number of primary detections (avalanches), the per-event variance of the ECF distribution decreases according to the CLT.




Figure 6.23: Simulated SiPM noise mean ECF and CAF versus number of primary avalanches per event, for 1000 toy events per configuration. Top: Mean ECF vs number of primary hits per event for FBK and HPK SiPMs simulated with SRML, plotted error bars are $1-\sigma$ standard deviations of the ECF distribution for each. Bottom: Correlated avalanche fluctuations. As we increase the number of primary avalanches (hits) per event, we experience a CLT suppression of the variance per event.

The presence of bimodal noise distributions has been observed in other SiPM applications, when considering the branching effects of the Poisson processes, where correlated noise with low probabilities can produce a so-called “heavy-tail”, when early successful trials compound to produce more noise [186]. This effect could have been excluded by analysis selection conditions in the input data sources, especially if measurements were carried out at different incident light levels.

We compare the simulated results from SRML with the measured values from [101] in Table 6.4. We see some differences when considering mean excess charge, and the variance. When comparing CAF, we observe that SRML underestimates substantially compared to [101], which is expected. This is partly because the results from [101] are from a full electronics readout including shaping and digitization, and are more representative of the true

Table 6.4: Comparison of SRML simulated excess charge factor (ECF) and correlated avalanche fluctuations (CAF) with SiPM ExCT % contribution to ECF highlighted, with [101] measurement results, all values are for 3 V OV. The estimates for SRML are the tangential estimates for high N_{hit} , relevant for the large number of avalanches expected in $0\nu\beta\beta$ events in nEXO.

SiPM Type	ECF	CAF	ExCT Contrib. (%)
SRML			
HPK VUV4	1.05	0.05	8.1
FBK HD3	1.18	0.16	2.6
G2022 - [101]			
HPK VUV4	1.06 ± 0.02	0.24 ± 0.02	—
FBK HD3	1.23 ± 0.06	0.42 ± 0.07	—

variance expected from SiPM correlated noise, but is largely driven by the high N_{hit} limit taken here. The results reported here for CAF can be interpreted as an estimate of the lower limit on variance, driven purely from the correlated noise sources’ stochastic nature, when approaching the high N_{hit} limit, the expected conditions for $0\nu\beta\beta$ events in nEXO.

6.5 Revisiting nEXO Energy Resolution

With our updated optical model of the nEXO TPC, and detailed SiPM correlated noise simulation, we are prepared to make an updated estimate of the nEXO light channel energy resolution. We restrict this analysis to the most meaningful energy value for the experiment, the ^{136}Xe $Q_{\beta\beta}$, an end-point energy of 2458 keV. The MC methodology is given as follows:

1. Using the nEXO Geant4 simulation toolkit, `nexo-offline`, we produce 30k events using the $0\nu\beta\beta$ generator to simulate the expected particle transport physics, including interactions with LXe that may produce Bremsstrahlung photons and other nuisance effects. These simulations are run with no light or charge transport, and we record all position dependent energy deposition information as outlined in Section 2.5, with the addition of the nEXO specific NEST⁴ photon production estimate for each energy

⁴nEXO NEST is modified from the public package to use the EXO-200 W-value estimate [160], changing the scintillation light yield by +18 %.

deposit.

2. After converting the `nexo-offline` ROOT file to the required input shape for Chroma, we run the full optical simulation for each input Geant4 $0\nu\beta\beta$ event. This is repeated for both Hamamatsu and FBK SiPM configurations.
3. Read in event-by-event information from optical simulation results, including: travel time for detected photons, detected photon Z position, origin (X,Y,Z) position of $0\nu\beta\beta$, and the MC truth number of simulated photons.
4. Run SiPM response simulation for each event, looping over detected primary hits to add correlated and dark noise.
5. Load the high statistics $0\nu\beta\beta$ simulated dataset from the SRML results, and for each event reconstruct the estimated number of scintillation photons produced.

To reconstruct the estimated number of produced photons $\widehat{N_{Scint.}}$, we combine the total charge from primary photons and correlated hits, and the MC truth origin position to evaluate the CE light map as:

$$\widehat{N_{Scint.}} = \frac{N_{Ph,Q}}{CE(R, Z)} \times \frac{1.0}{ECF}, \quad (6.4)$$

where, $\widehat{N_{Scint.}}$ is the estimated number of produced scintillation photons for the event from reconstruction, $N_{Ph,Q}$ is the total charge of the event, from primary photons plus excess charge from the SRML simulation, $CE(R, Z)$ is the collection efficiency lightmap value at event position (R, Z) , and lastly ECF is the mean of the excess charge distribution for the simulated $0\nu\beta\beta$, in the high N_{hit} limit.

We can directly compare the intrinsic fluctuations for scintillation production, as modelled by the nEXO NEST simulation, including the variance from some non-local energy depositions, with the expected reconstructed charge distribution after correcting for transport, detection and excess charge from correlated avalanches by fitting the simulated data set for

each SiPM vendor configuration. This comparison is shown in Figures 6.24 and 6.25. We observe the expected $\sim 5\%$ resolution in light-only scintillation production from the nEXO NEST model.




Figure 6.24: Comparison of reconstructed and simulated photons for $0\nu\beta\beta$ events in nEXO TPC FV for the HPK configuration. We fit each distribution to a Gaussian signal model, with parameter estimates and reduced χ^2 shown in the legend. There is a small negative bias of $\sim 0.3\%$ in the mean of the reconstructed photon count vs MC truth mean. The source of the bias is likely from low statistics in the light-map.

We compare the simulated MC truth, the reconstructed photon count from total charge (including SRML) and reconstructed photon count from primary VUV photon detections. With the information from SRML, we can independently evaluate the contribution of each stage: NEST MC truth (the variance from LXe microphysics), primary hits only reconstruction (dominated by Poisson statistics from re-scaling detected photons), and total charge reconstruction (includes primary hits plus SRML noise). We estimate the contribution to resolution at each stage by subtracting in quadrature the variance from the previous stage. This is repeated for both configurations, where we use the NEST MC truth mean value of 94630 photons as the mean for resolution estimates. This was chosen to avoid the observed bias from the low-statistics lightmap, assuming a well-calibrated lightmap has no inherent bias and will recover




Figure 6.25: Comparison of reconstructed and simulated photons for $0\nu\beta\beta$ events in nEXO TPC FV for the FBK configuration. We fit each distribution to a Gaussian signal model, with parameter estimates and reduced χ^2 shown in the legend. There is a small positive bias of $\sim 1.7\%$ in the mean of the reconstructed photon count vs MC truth mean. The source of the bias is likely from low statistics in the light-map.

the mean of the MC prediction, but not the variance. This comparison is shown in Table 6.5. The contribution column is the additional contribution in resolution from the added variance from the previous stage.

Due to the limited statistics in the simulated CE lightmap, we are restricted in the conclusions we can draw. We see from Table 6.5 that the additional resolution contribution introduced from SiPM correlated noise is $\sim 0.5\%$ for HPK and $\sim 1\%$ for FBK. While the dominant source of additional variance is due to primary hit reconstruction, at $\sim 4\%$ for both, increasing the σ/E resolution in the light-only channel from $\sim 5\%$ from the NEST MC predictions, to $\sim 6.4\%$ when reconstructing. This additional variance is dominated by Poisson counting statistics, as we only collect $\sim 1.5\%$ of the total scintillation light, with a contribution from the coarse voxels in the simulated lightmap used here.

In order to estimate the impact on overall detector energy resolution, and the OV dependence, we can use the energy resolution model from [101] (equation 22 therein), using input

Table 6.5: Contributions of simulated nEXO light channel energy resolutions at $Q_{\beta\beta}$, for both SiPM vendor configurations at 3 V OV. Comparison of the contribution to the resolution in σ/μ , using the mean of the MC Truth fit, ~ 94630 photons, assuming lightmap reconstruction bias is negligible. Contribution column represents the additional resolution in % added from the previous stage, starting from MC.

Stage	$\sigma_{N_{Scint.}}$ (ph)	Resolution (%)	Contribution (%)
NEST MC	4745 ± 96	5.01 ± 0.1	—
FBK Configuration			
Primary Hits	6069 ± 45	6.41 ± 0.05	~ 4
Total Charge	6149 ± 43	6.49 ± 0.05	~ 1
HPK Configuration			
Primary Hits	6093 ± 35	6.43 ± 0.03	~ 4
Total Charge	6113 ± 37	6.46 ± 0.04	~ 0.5

parameters from table 4 of that work. This model assumes perfect anti-correlation between charge and light signals, includes contributions from SiPM correlated noise, lightmap reconstruction uncertainty, and collection statistics using the estimated mean light collection efficiency in the nEXO TPC. The charge channel contribution includes estimated readout noise variance and charge loss from impurities. The energy resolution is given by:

$$\frac{\sigma_n}{\langle n \rangle} = \frac{\sqrt{\left(\frac{(1-\varepsilon_p)n_p}{\varepsilon_p} + \frac{n_p}{\varepsilon_p} \cdot \frac{\sigma_\Lambda^2}{(1+\langle \Lambda \rangle)^2} + n_p^2 \sigma_{lm}^2 \right) + \left(\frac{n_q t}{\tau} + \frac{\sigma_{q,noise}^2}{\varepsilon_q^2} \right)}}{\langle n \rangle}, \quad (6.5)$$

where ε_p is the mean light collection efficiency, n_p is the estimated number of scintillation photons produced at $Q_{\beta\beta}$, σ_{lm} is the lightmap reconstruction uncertainty, τ is the estimated electron lifetime, n_q is the estimated number of ionization electrons produced at $Q_{\beta\beta}$, t is the mean drift time, ε_q is the estimated mean charge collection efficiency, and $\sigma_{q,noise}$ is the expected readout charge noise in units of electrons. The energy resolution $\sigma_n/\langle n \rangle$, with mean total quanta produced $\langle n \rangle$, is expressed in % units.

We use the mean light collection efficiency shown Figure 6.18, averaged over R and Z, with the expected relative increase/decrease in detection efficiency for a specific OV. The change in efficiency is derived from the changing avalanche initiation probability from electron-driven avalanches in VUV from [103], and varies from 0.8—1.15 from 2—5 V with

1.0 at 3 V OV. For comparison, we model three sets of FBK/HPK energy resolution curves: first using the updated optical MC presented here with CAF from SRML, secondly the results from [101], using CAF input data scraped from the figures for each OV and SiPM vendor, and lastly a hybrid version using the SRML estimated CAFs (including SiPM ExCT, omitted from [101]) but with the nominal light collection efficiency assumed in [101]. This comparison is shown in Figure 6.26, where all sets of curves use the assumed σ_{lm} of 0.5 %.




Figure 6.26: Model of nEXO total energy resolution at $Q_{\beta\beta}$, comparing three different scenarios. SRML estimated CAF with updated MC presented in this work, reproduction of model presented in [101], and SRML estimated CAF with previous estimate of light collection efficiency. The shaded regions in the updated MC represent the $\pm 1\sigma$ variance in light collection efficiency across the fiducial volume. All three sets of curves use $\sigma_{lm} = 0.5\%$, and the last two scenarios overlap with the estimated HPK resolution from [101] (G2022).

We can observe in Figure 6.26 that the targeted energy resolution of 1 % σ/E at $Q_{\beta\beta}$ is not met in the current estimations. As noted earlier, this is driven mainly by the substantially lower overall light collection efficiency predicted from the full wavelength-dependent treatment described in this chapter. These results represent the first full wavelength-dependent

optical simulation of nEXO. While extensive cross-checks have been performed, independent validations and experimental measurements are essential next steps.

6.6 Conclusions

In this chapter, we have outlined the substantial efforts in refactoring and expanding the utility and applicability of the nEXO light simulation framework, `Chroma`, and its corresponding nEXO collaboration packaged implementation `chroma-simulation`. With these upgrades in-hand, we have the ability to simulate more realistically, and more broadly than ever before, enabling future studies using nEXO light channel information. These upgrades included expanded optical simulation input data and validation efforts, and essentially a more detailed and upgraded simulated optics model for SiPMs based on recent efforts [103]. Our primary goal in these efforts was to expand the ability of `Chroma` to the LoLX phase-2 experiment, in order to better understand the impact of SiPM ExCT on the nEXO detector energy resolution.

The impact of SiPM ExCT on nEXO was evaluated here for the first time, and in agreement with other works, is a non-negligible but sub-dominant component of SiPM correlated noise. When combined with updated optical simulation lightmaps, referred to here as collection efficiency or CE lightmaps, we have estimated the contribution of SiPM correlated noise to the resolution of the nEXO light readout system. We have also revisited the design requirements for total light collection efficiency of 3 %. We observed that at the nominal 3 V OV this target is not met in the full wavelength-dependent optical simulation model presented here. When comparing the collection efficiency for the FBK and HPK detector configurations, we see that FBK outperforms HPK. The mean light collection efficiency of 1.03 ± 0.39 % for HPK and 1.37 ± 0.52 % for FBK, with uncertainties from the $1\text{-}\sigma$ standard deviation for events distributed throughout the fiducial volume. This performance deficit is complemented by a much lower correlated noise level for the HPK devices, ultimately producing similar estimates for light channel energy resolution at $Q_{\beta\beta}$ as shown in Table 6.5. We highlighted

the need for independent validation of these first results, and emphasized the impact of increased bulk absorption due to the increased optical path lengths from Rayleigh scattering and enhanced reflectivities.

Future measurements with the $1 \times 1 \text{ cm}^2$ devices planned for nEXO will need to be carried out in LXe to validate their ultimate performance, a role that LoLX is poised to fulfill. This is especially important for HPK devices, which, as indicated in Chapter 4, exhibit $\sim 15\text{-}20\%$ lower PDE (relatively) in LXe than expected from the optical models used here. It remains to be shown if this decrease is as apparent in a nEXO-like geometry. This is because we ascribed most of the reduction in PDE observed in LoLX phase-2 to shadowing from surface features on the SiPM due to the large amount of oblique light in the small cubic detector geometry. This analysis and the accompanying LoLX phase-2 simulation comparison is available as a preprint [139] currently under review.

As is often the case for simulation studies, we are limited significantly by the quality and applicability of input data. In many cases we make assumptions, some minor and others more significant, about the input data distributions, and the fidelity of the simulated detector model. Some key assumptions for the optical simulation itself include: the perfect flatness and uniformity of the SiPM, tile and stave surfaces, the uniform responsivity of all SiPMs in nEXO, and the absence of surface effects, such as thin-film interference, on light transport. We have also omitted the tile-level fill factors in photon collection efficiency calculations, which is not currently known, but will reduce overall collection efficiency proportionally. While our input data has been described in detail, we restate the importance of having subsequent in-LXe measurements of any critical detector components, as expectations from vacuum measurements may be broken when including the temperature and optical medium change. Some of the key components for nEXO that should be measured in detail include: MgF_2 -Al VUV reflective coating witness samples from the cathode and FSRs, the ultimately selected SiPMs themselves along with their mounting tiles (all studies so-far are for 6 mm SiPMs from HPK

and FBK, while nEXO versions will be 1 cm devices), FSR spacers and resistive coatings, and the anode charge collection tiles. Many of these limitations and approximations will be updated readily once nEXO is online and taking data, using effective parametrizations, ad-hoc applications of surface absorption, and effective reflectivities. The charge collection tiles are also a key component that need to be measured in-LXe to inform accurate simulations, and represent an opportunity for enhanced light collection efficiency with increased reflectivity. With these additional inputs good agreement between data and simulation is expected to be achieved. These efforts would be substantially eased by significant investments in testing and characterization capabilities for important components in-LXe and at operating temperatures if liquid measurements are untenable.

While the simulated statistics considered in the MC estimate of nEXO light channel resolution here are limited, we can make general statements about the relative contribution of SiPM correlated noise fluctuations on energy measurements from an integrating detector such as nEXO. The CE lightmaps used here will ultimately be replaced by much higher statistics maps from dedicated data calibration runs [170], and introduce an artificial variance due to the limited number of voxels from limited statistics. Nevertheless, we see that the reconstruction adds $\sim 4\%$, and SiPM correlated noise, including ExCT, adds $0.5\text{--}1\%$ depending on vendor, at 3 V OV. In Table 6.5 we compared the relative effect of reconstructing only primary detections versus all hits including all SiPM correlated noise, on the light channel energy resolution. Employing the model developed by nEXO collaborators for an earlier estimate of nEXO detector response in [101], we estimated the overall nEXO detector energy resolution for two different scenarios, as a function of SiPM OV. The first scenario used the estimated variance in SiPM correlated noise (CAF) from the SRML developed in this work, with the updated estimate of the mean light collection efficiency in nEXO. The final scenario used the SRML estimated CAF for each configuration, but the nominal light collection efficiency assumed in [101]. We see an improvement due to the improved understanding of SiPM correlated noise

fluctuations expected at $Q_{\beta\beta}$, based on the SRML estimates, when comparing to the results from [101]. Beyond the observed excess absorption in bulk LXe and at the anode surface, an additional source of the lower light collection efficiency estimated in this work is the improved SiPM surface optics and detection model employed here, based on the results from [103]. This 2D LUT implementation handles the full wavelength and incident angle dependent behaviour expected for the SiPMs in LXe conditions, and is significantly different from the expectations based on vacuum measurements. More validation and future measurements of in-LXe SiPM performance is needed to have full confidence in the MC strategy presented here, while work is underway to improve our understanding [139], strong MC agreement with LoLX data is a necessary next-step.

The simulation framework developed here has already been used in a recent nEXO publication estimating the sensitivity of the detector to solar neutrinos and fermionic dark matter [187]. For this work, using the models described in this chapter, I prepared input timing PDFs for ExCT transport in nEXO, and expected transport and detection information for the lead author’s Glenn Richardson and Brian Leonardo, and advised on simulation inputs and modelling the detector response. For this analysis, which uses the delayed coincidence of de-excitation γ ’s from previously unknown isomeric states of ^{136}Cs , the distribution of detected photon timing is a critical input. By accurately measuring the time and energy of delayed energy emission in nEXO, it is possible to separate the primary and delayed events to perform particle identification. The impact of SiPM ExCT on timing resolution was studied and determined to be subdominant to the effect of the readout timing resolution. The planned shaping and integrating amplifier readout limits nEXO timing resolution to $\mathcal{O}(10\text{ ns})$, while the low intensity of SiPM ExCT introduces a negligible additional smearing on photon arrival times. While SiPM ExCT is subdominant when considering the analysis in [187], it signals the importance of having a detailed and flexible simulation toolkit for expanding the physics reach of the detector, enabling new and unique analyses.

In the future, expanded and further validated Chroma simulations will be used for nEXO and other experiments in the noble liquid detector ecosystem. This has been spurred by a recent cross-collaboration workshop on light transport simulations, where we renewed the interest in, and importance of open-source accurate simulation tools for the use by the community. Future upgrades to the SRML will include the effects of SiPM saturation, readout noise, simulated waveform response and other effects neglected here. This is critical for a rare-event search experiment like nEXO, which aims to discover new physics. The gold standard in detector response modelling is to simulate the detector in as much detail as computationally feasible, making no sacrifices on accuracy, which produces a simulated event-by-event response that can be analyzed in the exact same way as the data from the real experiment. This would allow for critical evaluation of the readout efficiencies, including pulse-finder losses, reconstruction errors, signal distortions, and more, by comparing the reconstructed results from simulated waveforms, where ground-truth information is known, with real detector data, where it is not.

In the final chapter of this work, we proceed to a detailed discussion of the entire thesis and how each part fits together, outline the remaining challenges in modelling a system as complex and detailed as the nEXO detector, and highlight the steps that can be taken in the future to improve upon the work presented here.

Part IV

Discussion and Conclusions

Discussion and Conclusions

Neutrinos are at the forefront of fundamental physics research, as one of the few challenge points of the otherwise extremely successful Standard Model of particle physics. The origin of the neutrino mass is not currently known, and is not included in the SM. Majorana neutrinos are a natural and elegant extension of the Standard Model, capable of explaining the origin of the neutrino mass mechanism and the several orders of magnitude separation between neutrino masses and those of other fermions. Furthermore, the lepton number violation and CP-violation inherent to Majorana neutrinos can satisfy the Sakharov conditions required to explain the observed matter-antimatter asymmetry of the Universe [188]. An observation of the hypothetical decay mode, neutrinoless double beta decay, is an unambiguous signature for Majorana neutrinos, and would be the first observation of lepton number violate and sub-

sequent breaking of a global conservation law of the SM, while also sensitive to information about the underlying mechanism. The search for $0\nu\beta\beta$ has been ongoing for several decades, and has matured beyond earlier experiments which probed initially for $0\nu\beta\beta$ and measured the $2\nu\beta\beta$ half-lives. The current stage of experimentation is aiming for coverage of the full range of $m_{\beta\beta}$ values under the inverted mass ordering. Achieving this requires half-life sensitivities of $> 10^{28}$ y, which in turn demands tonne-scale isotope masses in next-generation experiments.

The nEXO experiment is an experiment highly optimized to maximize the discovery potential for $0\nu\beta\beta$, with advanced design and a robust R&D program for radioactive background control. nEXO builds on the success of the EXO-200 experiment, and will employ a monolithic single-phase LXe TPC, capable of robust background suppression through self-shielding and targeted analysis. The light detection system of nEXO will use SiPMs with a large photocoverage, 4.6 m² of surface area, readout as 7680 channels. In order to mitigate all conceivable design risks, and to improve detector understanding, we need a complete simulation toolkit that has been benchmarked against real detector data through the use of validation experiments. The LoLX experiment, located at McGill University, is one such prototype detector that has helped improve our understanding of light transport and detection of LXe scintillation. The LoLX experiment has gone through two experimental phases to-date, with upgrades to the cryogenic cooling system, control processes, DAQ and data analysis, and simulation frameworks, between phases, all of which was a substantial part of this thesis work.

In the first phase of LoLX operations, we were limited by several challenging design aspects, including detector material choices, the cooling system and control, and analysis toolkits. The fluorescence of the 3D printed cage material represented a limitation in our physics reach in the first phase, but resulted in a novel method for understanding SiPM ExCT in the LXe detector. By using the late fluorescence as “seeds” for ExCT, after scintillation light has

dissipated, we measured SiPM ExCT correlation between sensors in LoLX. By combining this with a detailed custom Geant4 simulation of SiPM ExCT in LoLX, we were able to model SiPM ExCT and parameterize its impact in an experiment-agnostic fashion for the first time. We were able to confirm the expected gain-dependent behaviour of SiPM ExCT, and estimate the number of emitted photons per avalanche in LXe to be $0.5^{+0.3}_{-0.2}$ and $0.6^{+0.3}_{-0.2}$ at 4 V and 5 V of OV, respectively, for VUV4-Q HPK devices. This model and simulation represent a step forward in our understanding of SiPM ExCT and highlights the dependence on robust optical modelling across the entire wavelength range expected inside the detector. We also made our simulation model available as an open source Geant4 example project for community use [189]. This is important because the modelling of SiPM ExCT was a missing ingredient in many experiments considering SiPMs as light sensors, including nEXO [101].

We invested substantial resources into upgrading the structure and organization of the chroma-simulation framework to support prototype characterization comparisons with LoLX and collaborators within nEXO. With our updated and more flexible simulation framework we set to updating the previous conservative and ad-hoc estimates of optical properties in the nEXO TPC, for a more realistic estimate of light transport for the scintillation readout channel of the detector. The optical property inputs still required assumptions and modelling due to an absence of applicable input data in the literature, highlighting the need for experimental validation of these key input parameters. With updated simulations, and an upgraded PDE model for SiPM candidates in nEXO, and a validated ExCT simulation model from LoLX, we were prepared to update performance estimates of the light detection system of nEXO. We investigated the impact of SiPM ExCT in the nEXO experiment, relevant for future noble liquid detector experiments that use SiPMs. These investigations showed that SiPM ExCT has interesting timing and hit pattern properties, owing to the very long scattering length in LXe at NIR wavelengths, which could be exploited by low-threshold experiments for ExCT suppression. From the simulations of SiPM ExCT transport in nEXO we extracted a parame-

terized light transport efficiency for SiPM ExCT as a necessary ingredient in estimating the level of correlated noise production for physics analyses.

We developed a SiPM Response Model Library (SRML), which uses MC techniques to iteratively add all sources of SiPM noise onto the primary detected photon signal from scintillation. Using data from earlier nEXO publications for both HPK and FBK candidate SiPMs, with our simulations of SiPM ExCT transport in nEXO, we were able to determine that SiPM ExCT is a non-negligible but subdominant component of excess charge at nominal operating voltages at the 5–10 % level of total excess charge. With the detailed MC-based SRML for SiPM noise we were able to investigate its impact on the energy resolution in the light channel for nEXO. We observed a central-limit-theorem suppression of the variance in SiPM correlated noise when going beyond the conditions of benchtop measurements with single SiPMs, to the conditions for $0 \leq \beta \leq 1$ in nEXO. The expected reduction in correlated avalanche fluctuations (CAF) was $2 \times$ to $5 \times$ versus expectations from single-avalanche signals in benchtop measurements, with the caveat that this reduction does not include effects of SiPM gain and electronics and reconstruction. The updated optical simulations for nEXO, including the more realistic optical parameter estimates and upgraded and validated SiPM PDE model from LoLX and TRIUMF, led to a notable reduction in the expected light collection efficiency for VUV photons in nEXO compared to earlier estimates. At nominal operating voltages, we found that the FBK SiPMs will outperform the HPK devices. We found that SiPM correlated noise, even with the inclusion of ExCT, is a subdominant component of the resolution in the light channel, dominated by the effects of photon counting statistics. Using the $\sigma_n / \langle n \rangle$ energy resolution model from [101], we estimated the detector energy resolution to be $1.35^{+0.30}_{-0.16}$ and $1.51^{+0.35}_{-0.19}$ at 3 V OV using our updated optical MC for HPK and FBK SiPMs respectively. We compared this to the estimated 0.73 ± 0.02 % and 0.76 ± 0.01 at 3 V OV for FBK and HPK respectively reported in [101]. The SRML developed in this work was used for the estimate of CAF included in the resolution, and the uncertainty is derived from the $\pm 1\sigma$ variance in

light collection efficiency across the fiducial region of the TPC in the updated optical MC. Under this model of resolution, the dominant source of difference is the updated estimate of mean light collection efficiency within the TPC, which points to opportunities for optimization of the detector design through improved photon collection rather than correlated noise suppression. Exploratory studies indicated that the majority of the non-detected photons were absorbed by the LXe, and surrounding surfaces, enhanced over earlier estimates because of the more robust and detailed optical treatment used here. The literature on LXe absorption length is limited, and available estimates indicate a strong purity dependence [184]. There is an opportunity for future large-scale LXe experiments like nEXO and XLZD [150], to study this effect in dedicated test setups that may be able to perform this challenging measurement, requiring separating the independent contributions from scattering and absorption. One such apparatus is the Xenoscope in Zurich [190], which has purity monitoring capabilities and a 2.6 m LXe column, that may be modified for measuring the achievable absorption lengths in-LXe as a function of xenon purity.

For the second phase of the LoLX experiment, we redesigned the detector to address key points of challenge in the first phase. These upgrades included diversifying the light sensors to include both of nEXO’s candidate SiPMs, and the inclusion of a VUV PMT. Additionally, we invested in upgrading and improving the cryogenic cooling system, the DAQ, and downstream analysis. The upgraded cooling system makes use of a GM cryocooler, with custom copper coupling and temperature control through PID with dedicated controller hardware. By migrating from the previous open-loop LN₂ system to our upgraded cryocooler system, we have improved data-taking up-time by ~ 100 %, and reduced cool-down and filling times by ~ 40 %, while greatly simplifying control and remote monitoring requirements. The design and implementation of the improved cooling system was a major undertaking, done in parallel with the ExCT analysis from phase-1, led by myself. With improved conditions for data-taking, and upgraded DAQ with faster timing, the upgraded analysis software MERCI,

was developed. MERCI developed as a collaborative effort led by myself and TRIUMF scientist Andrea Capra, with support from LoLX and TRIUMF collaborators, enables our online data analysis, and high efficiency modular offline analysis for SiPM waveform data. For LoLX phase-2 we further evolved this analysis toolkit with advanced matched-filter based pulse finding for the lower SNR regime and flexible configurations for high and low light intensity data-taking.

We carried out commissioning data-taking campaigns with the upgraded detector systems, using external gamma calibration sources in August of 2023 and January of 2025. With our upgraded analysis tools, we were able to carry out preliminary data-simulation comparisons for external ^{133}Ba calibration data taken in January 2025. These comparisons show broad agreement in timing, position reconstruction, and light transport, key performance metrics for the GPU accelerated photon transport simulations. While the absolute light-yield observed is significantly lower ($\sim 30\%$) than expectations for pure LXe, we observed several effects which point towards impurity-driven quenching of the overall light production and transport. The energy resolution achieved with the LoLX phase-2 experiment was far below expectations, due to poor containment of light and attenuation and scattering of the external gammas. In order to address this poor performance, we designed and prepared an internal calibration source using neutron activated natural hafnium to produce a ^{181}Hf needle source. Combined with planned updates to the gas handling system to allow for xenon re-circulation, for improved purity, and a dedicated purity monitor, we aim to provide key measurements of LXe light yield, production characteristics, including Cherenkov light generation, and efficiency of SiPMs in LXe. The first stages of this work advanced, with a relative comparison of the FBK and HPK SiPMs submitted recently to the journal Nuclear Instruments and Methods in Physics Research A. This analysis highlighted the context-dependence of SiPM performance, where vacuum measurements used to extrapolate to LXe require reinterpretation due to the optical coupling between media and silicon, and with an angular dependence that will result

in different effective efficiencies depending on the specific detector geometry.

From our work with the LoLX experiment, and with detailed simulations of the proposed nEXO experiment, we have made substantial progress in modelling the light transport and detection process for future noble liquid detectors. The `chroma-simulation` framework upgrades developed here will be made available as an open-source GPU-accelerated photon transport package, and have already contributed to additional nEXO physics sensitivity analyses [187]. We identified a clear need for a cryogenic in-liquid optical characterization facility capable of angle-resolved measurements of reflectivity, transmittance, and PDE for future detector components. This is especially important for validating simulation studies, and improving simulation tools for more accurate future projections of detector performance.

Unambiguous discovery of $0\nu\beta\beta$ requires a well-understood detector with low backgrounds, and reproduction of any signal across multiple isotopes will be required to constrain the underlying mechanisms given the variance in nuclear matrix element calculations [40]. The proposed nEXO experiment is positioned to either observe $0\nu\beta\beta$ or exclude the inverted neutrino mass hierarchy for Majorana neutrinos. The discovery of Majorana neutrinos remains one of the most promising avenues for understanding physics beyond the Standard Model, and $0\nu\beta\beta$ provides the ideal unambiguous signature for lepton number violation.

Through detailed analysis, careful planning, and prototyping, we can build detectors with unprecedented sensitivity and precision, extending experimental reach to probe the limits of the Standard Model. With diligence, such efforts may bring us closer to a deeper understanding of the Universe and its most fundamental properties.

Bibliography

- [1] H.-E. Mahnke, “Lise Meitner, β -decay and non-radiative electromagnetic transitions,” *Notes and Records: the Royal Society Journal of the History of Science*, volume 76, number 1, pages 107–116, Sep. 2020, ISSN: 0035-9149. DOI: 10.1098/rsnr.2020.0036.
- [2] C. D. Ellis and W. A. Wooster, “The Average Energy of Disintegration of Radium E,” *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character*, volume 117, number 776, pages 109–123, 1927, ISSN: 0950-1207. JSTOR: 94774.
- [3] W. Pauli, “Dear radioactive ladies and gentlemen,” *Physics Today*, volume 31N9, page 27, 1978.
- [4] M. Goeppert-Mayer, “Double Beta-Disintegration,” *Physical Review*, volume 48, number 6, pages 512–516, Sep. 1935. DOI: 10.1103/PhysRev.48.512.
- [5] P. a. M. Dirac, “Discussion of the infinite distribution of electrons in the theory of the positron,” *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 30, number 2, pages 150–163, Apr. 1934, ISSN: 1469-8064, 0305-0041. DOI: 10.1017/S030500410001656X.

- [6] E. Majorana, “Teoria simmetrica dell’elettrone e del positrone,” *Nuovo Cim*, volume 14, number 4, pages 171–184, Apr. 1937, ISSN: 1827-6121. DOI: 10.1007/BF02961314.
- [7] W. H. Furry, “On Transition Probabilities in Double Beta-Disintegration,” *Physical Review*, volume 56, number 12, pages 1184–1193, Dec. 1939. DOI: 10.1103/PhysRev.56.1184.
- [8] H. Bethe and R. Peierls, “The ‘neutrino’,” *Nature*, volume 133, page 532, 1934. DOI: 10.1038/133532a0.
- [9] C. L. Cowan *et al.*, “Detection of the Free Neutrino: A Confirmation,” *Science*, volume 124, number 3212, pages 103–104, Jul. 1956. DOI: 10.1126/science.124.3212.103.
- [10] The Royal Swedish Academy of Sciences, *The Nobel Prize in Physics 1995*, Nobel Prize Outreach AB, 1995. [Online]. Available: <https://www.nobelprize.org/prizes/physics/1995/summary/>.
- [11] C. S. Wu *et al.*, “Experimental Test of Parity Conservation in Beta Decay,” *Physical Review*, volume 105, number 4, pages 1413–1415, Feb. 1957. DOI: 10.1103/PhysRev.105.1413.
- [12] The Royal Swedish Academy of Sciences, *The Nobel Prize in Physics 2002*, Nobel Prize Outreach AB, 2002. [Online]. Available: <https://www.nobelprize.org/prizes/physics/2002summary/>.
- [13] J. N. Bahcall and R. Davis, “Solar Neutrinos: A Scientific Puzzle,” *Science*, volume 191, number 4224, pages 264–267, Jan. 1976. DOI: 10.1126/science.191.4224.264.

- [14] R. Davis, “A review of the Homestake solar neutrino experiment,” in *Proceedings of the 21st International Cosmic Ray Conference*, volume 7, Adelaide, Australia, 1990, page 155. [Online]. Available: <https://adsabs.harvard.edu/full/1990ICRC...7..155D> (visited on 12/11/2025).
- [15] B. Pontecorvo, “Mesonium and Antimesonium,” *Soviet Physics Journal of Experimental and Theoretical Physics*, volume 6, pages 429–431, 1958.
- [16] K. S. Hirata *et al.*, “Results from one thousand days of real-time, directional solar-neutrino data,” *Physical Review Letters*, volume 65, number 11, pages 1297–1300, Sep. 1990. DOI: 10.1103/PhysRevLett.65.1297.
- [17] Super-Kamiokande Collaboration, “Evidence for Oscillation of Atmospheric Neutrinos,” *Physical Review Letters*, volume 81, number 8, pages 1562–1567, Aug. 1998. DOI: 10.1103/PhysRevLett.81.1562.
- [18] SNO Collaboration, “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory,” *Physical Review Letters*, volume 89, number 1, page 011 301, Jun. 2002. DOI: 10.1103/PhysRevLett.89.011301.
- [19] B. Kayser, *Neutrino Mass, Mixing, and Flavor Change*, Nov. 2002. DOI: 10.48550/arXiv.hep-ph/0211134. arXiv: hep-ph/0211134. [Online]. Available: <http://arxiv.org/abs/hep-ph/0211134>.
- [20] The Royal Swedish Academy of Sciences, *The Nobel Prize in Physics 2015*, Nobel Prize Outreach AB, 2015. [Online]. Available: <https://www.nobelprize.org/prizes/physics/2015summary/>.
- [21] K. Zuber, *Neutrino Physics*, 3rd edition. Boca Raton: CRC Press, May 2020, ISBN: 978-1-315-19561-2. DOI: 10.1201/9781315195612.

- [22] D. Griffiths, *Introduction to Elementary Particles*. New York, USA: John Wiley & Sons, 1987.
- [23] G. Aad *et al.*, “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC,” *Physics Letters B*, volume 716, number 1, pages 1–29, Sep. 2012, ISSN: 0370-2693. DOI: 10.1016/j.physletb.2012.08.020.
- [24] E. Noether, “Invariante variationsprobleme,” *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch Physikalische Klasse*, volume 1918, pages 235–257, 1918.
- [25] A. Djouadi, “The Anatomy of electro-weak symmetry breaking. I: The Higgs boson in the Standard Model,” *Physics Reports*, volume 457, pages 1–216, 2008. DOI: 10.1016/j.physrep.2007.10.004.
- [26] Wikimedia Commons. “Standard Model of elementary particle.svg,” Wikimedia Foundation. (2025), [Online]. Available: https://commons.wikimedia.org/w/index.php?title=File:Standard_Model_of_Elementary_Particles.svg&oldid=1062088276 (visited on 12/13/2025).
- [27] V. Gribov and B. Pontecorvo, “Neutrino astronomy and lepton charge,” *Physics Letters B*, volume 28, number 7, pages 493–496, Jan. 1969, ISSN: 0370-2693. DOI: 10.1016/0370-2693(69)90525-5.
- [28] P. D. Group, “Review of particle physics - Chapter 14,” *Progress of Theoretical and Experimental Physics*, volume 2024, number 8, page 083C01, 2024. DOI: 10.1093/ptep/ptae089.
- [29] C. Giganti, S. Lavignac, and M. Zito, “Neutrino oscillations: The rise of the PMNS paradigm,” *Progress in Particle and Nuclear Physics*, volume 98, pages 1–54, Jan. 2018, ISSN: 0146-6410. DOI: 10.1016/j.ppnp.2017.10.001.

- [30] K. Abe *et al.*, “The T2K experiment,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 659, number 1, pages 106–135, Dec. 2011, ISSN: 0168-9002. DOI: 10.1016/j.nima.2011.06.067.
- [31] P. F. de Salas *et al.*, “2020 global reassessment of the neutrino oscillation picture,” *Journal of High Energy Physics*, volume 2021, number 2, page 71, Feb. 2021, ISSN: 1029-8479. DOI: 10.1007/JHEP02(2021)071.
- [32] T2K Collaboration, “Search for CP violation in neutrino and antineutrino oscillations by the t2k experiment with 2.2×10^{21} protons on target,” *Physical Review Letters*, volume 121, number 17, page 171 802, Oct. 2018. DOI: 10.1103/PhysRevLett.121.171802.
- [33] G. Altarelli and F. Feruglio, “Models of neutrino masses and mixings,” *New Journal of Physics*, volume 6, number 1, page 106, Aug. 2004, ISSN: 1367-2630. DOI: 10.1088/1367-2630/6/1/106.
- [34] KATRIN Collaboration, “Direct neutrino-mass measurement based on 259 days of KATRIN data,” *Science*, volume 388, number 6743, pages 180–185, Apr. 2025. DOI: 10.1126/science.adq9592.
- [35] P. Minkowski, “ $\mu \rightarrow e \gamma$ at a rate of one out of 10^9 muon decays?” *Physics Letters B*, volume 67, number 4, pages 421–428, Apr. 1977, ISSN: 0370-2693. DOI: 10.1016/0370-2693(77)90435-X.
- [36] R. N. Mohapatra and G. Senjanović, “Neutrino mass and spontaneous parity nonconservation,” *Physical Review Letters*, volume 44, page 912, 1980. DOI: 10.1103/PhysRevLett.44.912.

- [37] R. N. Mohapatra and A. Y. Smirnov, “Neutrino mass and new physics,” *Annual Review of Nuclear and Particle Science*, volume 56, page 569, 2006, arXiv:hep-ph/0603118. DOI: 10.1146/annurev.nucl.56.080805.140534.
- [38] K. S. Krane, *Introductory Nuclear Physics*. New York: John Wiley & Sons, 1987, ISBN: 978-0-471-80553-3.
- [39] J. J. Gómez-Cadenas *et al.*, “The search for neutrinoless double-beta decay,” *La Rivista del Nuovo Cimento*, volume 46, number 10, pages 619–692, Oct. 2023, ISSN: 1826-9850. DOI: 10.1007/s40766-023-00049-2.
- [40] M. J. Dolinski, A. W. P. Poon, and W. Rodejohann, “Neutrinoless Double-Beta Decay: Status and Prospects,” *Annual Review of Nuclear and Particle Science*, volume 69, number Volume 69, 2019, pages 219–251, Oct. 2019, ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev-nucl-101918-023407.
- [41] M. Berglund and M. E. Wieser, “Isotopic compositions of the elements 2009 (IUPAC technical report),” *Pure and Applied Chemistry*, volume 83, number 2, pages 397–410, 2011. DOI: 10.1351/PAC-REP-10-06-02.
- [42] A. Barabash, “Precise Half-Life Values for Two-Neutrino Double- β Decay: 2020 Review,” *Universe*, volume 6, number 10, page 159, 2020. DOI: 10.3390/universe6100159.
- [43] S. Umehara *et al.*, “Search for Neutrino-less Double Beta Decay with CANDLES,” *Physics Procedia*, 13th International Conference on Topics in Astroparticle and Underground Physics, TAUP 2013, volume 61, pages 283–288, Jan. 2015, ISSN: 1875-3892. DOI: 10.1016/j.phpro.2014.12.046.
- [44] G. Saleh, “First results from LEGEND-200 on the search for neutrinoless double beta decay,” *Il Nuovo Cimento C*, volume 49, number 1–2, pages 1–6, Jan. 9, 2026, ISSN: 03905551, 03905551. DOI: 10.1393/ncc/i2026-26005-8.

- [45] C. Augier *et al.*, “Final results on the $0\nu\beta\beta$ decay half-life limit of ^{100}Mo from the CUPID-Mo experiment,” *European Physical Journal C*, volume 82, page 1099, 2022. DOI: 10.1140/epjc/s10052-022-10942-5. arXiv: 2202.081ba.
- [46] CUORE Collaboration, “Constraints on lepton number violation with the 2 tonne · year CUORE dataset,” *Science*, volume 390, number 6777, pages 1029–1032, Dec. 4, 2025. DOI: 10.1126/science.adp6474.
- [47] S. Abe *et al.*, *Search for Majorana Neutrinos with the Complete KamLAND-Zen Dataset*, Jun. 2024. DOI: 10.48550/arXiv.2406.11438. arXiv: 2406.11438 [hep-ex]. [Online]. Available: <http://arxiv.org/abs/2406.11438>.
- [48] N. Abgrall *et al.*, *LEGEND-1000 preconceptual design report*, 2021. arXiv: 2107.11462 [physics.ins-det].
- [49] D. Trotta, “The CUPID neutrinoless double-beta decay experiment,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 1066, page 169 657, Sep. 2024, ISSN: 0168-9002. DOI: 10.1016/j.nima.2024.169657.
- [50] T. Kroupová, “Improving the sensitivity to neutrinoless double beta decay in SNO+,” PhD Thesis, University of Oxford, 2020. [Online]. Available: <https://snoplus.phy.queensu.ca/theses/tkroupova.pdf>.
- [51] G. Anton *et al.*, “Search for Neutrinoless Double- β Decay with the Complete EXO-200 Dataset,” *Physical Review Letters*, volume 123, number 16, page 161 802, Oct. 2019. DOI: 10.1103/PhysRevLett.123.161802.
- [52] G. Adhikari, S. Al Kharusi, E. Angelico, *et al.*, “nEXO: Neutrinoless double beta decay search beyond 10^{28} year half-life sensitivity,” *Journal of Physics G: Nuclear and Particle Physics*, volume 49, number 1, page 015 104, Dec. 2021, ISSN: 1361-6471. DOI: 10.1088/1361-6471/ac3631.

- [53] C. Adams *et al.*, “The NEXT-100 Detector,” *The European Physical Journal C*, volume 86, number 2, page 114, Feb. 5, 2026, ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-025-14951-y.
- [54] C. Adams *et al.*, “Sensitivity of a tonne-scale NEXT detector for neutrinoless double-beta decay searches,” *Journal of High Energy Physics*, volume 08, page 164, 2021. DOI: 10.1007/JHEP08(2021)164. arXiv: 2005.06467.
- [55] C. Grant, *Results from KamLAND-Zen and SNO+*, Talk presented at Neutrino 2020 (neutrino2020), Zenodo, Jul. 2020. DOI: 10.5281/zenodo.4142683. [Online]. Available: <https://doi.org/10.5281/zenodo.4142683>.
- [56] A. Giuliani and A. Poves, “Neutrinoless Double-Beta Decay,” *Advances in High Energy Physics*, volume 2012, number 1, page 857 016, 2012, ISSN: 1687-7365. DOI: 10.1155/2012/857016.
- [57] D. J. Auty *et al.*, “A method to load tellurium in liquid scintillator for the study of neutrinoless double beta decay,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 1051, page 168 204, Jun. 1, 2023, ISSN: 0168-9002. DOI: 10.1016/j.nima.2023.168204.
- [58] K. Blaum *et al.*, “Neutrinoless double-electron capture,” *Review of Modern Physics*, volume 92, number 4, page 045 007, Dec. 2020. DOI: 10.1103/RevModPhys.92.045007.
- [59] J. Barea, J. Kotila, and F. Iachello, “Neutrinoless double-positron decay and positron-emitting electron capture in the interacting boson model,” *Physical Review C*, volume 87, number 5, page 057 301, May 2013. DOI: 10.1103/PhysRevC.87.057301.

- [60] NEXT Collaboration, *Reconstructing neutrinoless double beta decay event kinematics in a xenon gas detector with vertex tagging*, Jun. 2025. DOI: 10.48550/arXiv.2502.10198. arXiv: 2502.10198 [hep-ex]. [Online]. Available: <http://arxiv.org/abs/2502.10198>.
- [61] W. Commons, *File:Doublebetadecayfeynman.Svg — Wikimedia Commons, the Free Media Repository*, 2025. [Online]. Available: https://upload.wikimedia.org/wikipedia/commons/3/34/Double_beta_decay_feynman.svg.
- [62] J. Schechter and J. W. F. Valle, “Neutrinoless double- β decay in $SU(2) \times U(1)$ theories,” *Physical Review D*, volume 25, number 11, pages 2951–2954, Jun. 1982. DOI: 10.1103/PhysRevD.25.2951.
- [63] J. Kotila and F. Iachello, “Phase-space factors for double- β decay,” *Physical Review C*, volume 85, number 3, page 034316, Mar. 2012. DOI: 10.1103/PhysRevC.85.034316.
- [64] LEGEND Collaboration, *LEGEND-1000 Preconceptual Design Report*, Jul. 2021. DOI: 10.48550/arXiv.2107.11462. arXiv: 2107.11462 [physics]. [Online]. Available: <http://arxiv.org/abs/2107.11462>.
- [65] A. Armato *et al.*, *Toward CUPID-IT*, Apr. 2022. DOI: 10.48550/arXiv.2203.08386. arXiv: 2203.08386 [nucl-ex]. [Online]. Available: <http://arxiv.org/abs/2203.08386>.
- [66] D. Mei and A. Hime, “Muon-induced background study for underground laboratories,” *Physical Review D*, volume 73, number 5, page 053004, Mar. 2006. DOI: 10.1103/PhysRevD.73.053004.
- [67] V. V. Golovko *et al.*, “Ambient Dose and Dose Rate Measurement in SNOLAB Underground Laboratory at Sudbury, Ontario, Canada,” *Sensors (Basel)*, volume 23, number 4, page 1945, Feb. 2023, ISSN: 1424-8220. DOI: 10.3390/s23041945.

- [68] S. Cebrian, *Low radioactivity techniques for Large TPCs in rare event searches*, Mar. 2024. DOI: 10.48550/arXiv.2403.18143. arXiv: 2403.18143 [physics]. [Online]. Available: <http://arxiv.org/abs/2403.18143>.
- [69] G. Heusser, “Low-Radioactivity Background Techniques,” *Annual Review of Nuclear and Particle Science*, volume 45, number Volume 45, pages 543–590, Dec. 1995, ISSN: 0163-8998, 1545-4134. DOI: 10.1146/annurev.ns.45.120195.002551.
- [70] K. Pushkin *et al.*, “Study of radon reduction in gases for rare event search experiments,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 903, pages 267–276, Sep. 2018, ISSN: 0168-9002. DOI: 10.1016/j.nima.2018.06.076.
- [71] S. Al Kharusi, “Development of a water Cherenkov muon veto for the nEXO neutrinoless double beta decay experiment,” Ph.D. dissertation, McGill University, Montreal Canada, 2024. [Online]. Available: <https://escholarship.mcgill.ca/concern/theses/3j333812w>.
- [72] L. Baudis, “Dual-phase xenon time projection chambers for rare-event searches,” *Philos Trans A Math Phys Eng Sci*, volume 382, number 2266, page 20230083, Dec. 2023, ISSN: 1364-503X. DOI: 10.1098/rsta.2023.0083.
- [73] H. Rasiwala *et al.*, “‘Searching for a needle in a haystack;’ A Ba-tagging approach for an upgraded nEXO experiment,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, volume 541, pages 298–300, Aug. 2023, ISSN: 0168-583X. DOI: 10.1016/j.nimb.2023.05.047.
- [74] N. K. Byrnes *et al.*, “Fluorescence imaging of individual ions and molecules in pressurized noble gases for barium tagging in ^{136}Xe ,” *Nature Communications*, volume 15,

- number 1, page 10 595, Dec. 2024, ISSN: 2041-1723. DOI: 10.1038/s41467-024-54872-0.
- [75] Subatomic Physics Long-Range Plan Committee, “Canadian subatomic physics long-range plan 2022–2026,” Canadian Institute of Nuclear Physics and Institute of Particle Physics, 2022. [Online]. Available: <https://subatomicphysics.ca/lrp-2022-2026/> (visited on 07/12/2026).
 - [76] A. M. Baldini, V. Baranov, M. Biasotti, *et al.*, “The Search for $\mu^+ \rightarrow e^+ \gamma$ with 10^{-14} Sensitivity: The Upgrade of the MEG Experiment,” *Symmetry*, volume 13, number 9, 2021, ISSN: 2073-8994. DOI: 10.3390/sym13091591.
 - [77] C. Romo-Luque, “PETALO: Time-of-Flight PET with liquid xenon,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 958, page 162 397, 2020, ISSN: 0168-9002. DOI: 10.1016/j.nima.2019.162397.
 - [78] L. Gallego Manzano, J. M. Abaline, S. Acounis, *et al.*, “XEMIS2: A liquid xenon detector for small animal medical imaging,” *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 912, pages 329–332, 2018, ISSN: 01689002. DOI: 10.1016/j.nima.2017.12.022.
 - [79] B. Abi *et al.*, *Deep Underground Neutrino Experiment (DUNE), Far Detector Technical Design Report, Volume II: DUNE Physics*, Mar. 2020. DOI: 10.48550/arXiv.2002.03005. arXiv: 2002.03005 [hep-ex]. [Online]. Available: <http://arxiv.org/abs/2002.03005>.
 - [80] J. Jortner *et al.*, “Localized Excitations in Condensed Ne, Ar, Kr, and Xe,” *The Journal of Chemical Physics*, volume 42, number 12, pages 4250–4253, Jun. 1965, ISSN: 0021-9606. DOI: 10.1063/1.1695927.

- [81] A. Hitachi *et al.*, “Effect of ionization density on the time dependence of luminescence from liquid argon and xenon,” *Physical Review B*, volume 27, number 9, pages 5279–5285, May 1983. DOI: 10.1103/PhysRevB.27.5279.
- [82] M. Szydagis *et al.*, “NEST: A comprehensive model for scintillation yield in liquid xenon,” *Journal of Instrumentation*, volume 6, number 10, P10002, Oct. 2011, ISSN: 1748-0221. DOI: 10.1088/1748-0221/6/10/P10002.
- [83] D. Cichon *et al.*, “Scintillation decay-time constants for alpha particles and electrons in liquid xenon,” *Review of Scientific Instruments*, volume 93, number 11, page 113 302, Nov. 2022, ISSN: 0034-6748, 1089-7623. DOI: 10.1063/5.0087216.
- [84] E. Conti, R. DeVoe, G. Gratta, *et al.*, “Correlated fluctuations between luminescence and ionization in liquid xenon,” *Physical Review B*, volume 68, number 5, page 054 201, 2003. DOI: 10.1103/PhysRevB.68.054201.
- [85] C. A. O. Henriques *et al.*, *Unraveling xenon primary scintillation yield for cutting-edge rare event experiments*, Oct. 2023. DOI: 10.48550/arXiv.2309.14202. arXiv: 2309.14202 [astro-ph]. [Online]. Available: <http://arxiv.org/abs/2309.14202>.
- [86] M. Szydagis *et al.*, “A review of NEST models for liquid xenon and an exhaustive comparison with other approaches,” *Frontiers in Detector Science and Technology*, volume 2, Jan. 2025, ISSN: 2813-8031. DOI: 10.3389/fdest.2024.1480975.
- [87] T. J. Hajagos *et al.*, “High-Z Sensitized Plastic Scintillators: A Review,” *Advanced Materials*, volume 30, number 27, page 1 706 956, 2018, ISSN: 1521-4095. DOI: 10.1002/adma.201706956.
- [88] N. T. Luan *et al.*, *Absolute light yield measurement of NaI:Tl crystals for dark matter search*, Apr. 2024. DOI: 10.48550/arXiv.2404.18551. arXiv: 2404.18551 [hep-ex]. [Online]. Available: <http://arxiv.org/abs/2404.18551>.

- [89] LUX Collaboration, “Liquid xenon scintillation measurements and pulse shape discrimination in the LUX dark matter detector,” *Physical Review D*, volume 97, number 11, page 112 002, Jun. 2018. DOI: 10.1103/PhysRevD.97.112002.
- [90] A. Hitachi, “Luminescence Response and Quenching Models for Heavy Ions of 0.5 keV to 1 GeV/n in Liquid Argon and Xenon,” *Instruments*, volume 5, number 1, 2021, ISSN: 2410-390X. DOI: 10.3390/instruments5010005.
- [91] G. J. Boyle *et al.*, *Review of the experimental and theoretical landscape of electron transport in noble liquids*, Apr. 2025. DOI: 10.48550/arXiv.2504.16338. arXiv: 2504.16338 [physics]. [Online]. Available: <http://arxiv.org/abs/2504.16338>.
- [92] Ö. Şahin *et al.*, “Penning transfer in argon-based gas mixtures,” *Journal of Instrumentation*, volume 5, P05002, 2010. DOI: 10.1088/1748-0221/5/05/P05002.
- [93] F. M. Benoit and A. G. Harrison, “Predictive value of proton affinity. Ionization energy correlations involving oxygenated molecules,” *Journal of the American Chemical Society*, volume 99, number 12, pages 3980–3984, Jun. 1977, ISSN: 0002-7863. DOI: 10.1021/ja00454a015.
- [94] A. Dobi *et al.*, “Study of a zirconium getter for purification of xenon gas,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 620, number 2, pages 594–598, Aug. 2010, ISSN: 0168-9002. DOI: 10.1016/j.nima.2010.03.151.
- [95] G. Plante *et al.*, “Liquid-phase purification for multi-tonne xenon detectors,” *European Physical Journal C*, volume 82, number 10, page 860, Oct. 2022, ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-022-10832-w.

- [96] E. Aprile, R. Mukherjee, and M. Suzuki, “Measurements of the lifetime of conduction electrons in liquid xenon,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 300, number 2, pages 343–350, Jan. 1991, ISSN: 0168-9002. DOI: 10.1016/0168-9002(91)90446-W.
- [97] P. A. Cherenkov, “Visible luminescence of pure liquids under the influence of γ -radiation,” *Proceedings of the USSR Academy of Sciences*, volume 2, number 8, pages 451–454, 1934. DOI: 10.3367/UFNr.0093.196710n.0385.
- [98] K. Nakamura and Particle Data Group, “Review of Particle Physics,” *Journal of Physics G: Nuclear and Particle Physics*, volume 37, number 7A, page 075 021, Jul. 2010, ISSN: 0954-3899. DOI: 10.1088/0954-3899/37/7A/075021.
- [99] I. Tamm, “Radiation Emitted by Uniformly Moving Electrons,” in *Selected Papers*, I. E. Tamm *et al.*, Eds., Berlin, Heidelberg: Springer, 1991, pages 37–53, ISBN: 978-3-642-74626-0. DOI: 10.1007/978-3-642-74626-0_3. [Online]. Available: https://doi.org/10.1007/978-3-642-74626-0_3 (visited on 07/08/2025).
- [100] F. Acerbi and S. Gundacker, “Understanding and simulating SiPMs,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 926, pages 16–35, 2019, ISSN: 0168-9002. DOI: 10.1016/j.nima.2018.11.118.
- [101] G. Gallina, Y. Guan, F. Retiere, *et al.*, “Performance of novel VUV-sensitive Silicon Photo-Multipliers for nEXO,” *The European Physical Journal C*, volume 82, number 12, Dec. 2022, ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-022-11072-8.
- [102] G. Gallina *et al.*, “Characterization of SiPM Avalanche Triggering Probabilities,” *IEEE Transactions on Electron Devices*, volume 66, number 10, pages 4228–4234, Oct. 2019, ISSN: 1557-9646. DOI: 10.1109/TED.2019.2935690.

- [103] A. de St Croix *et al.*, “Demonstrating a broadband photon detection efficiency model on VUV sensitive silicon photomultipliers,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 1089, page 171 526, 2026, ISSN: 0168-9002. DOI: 10 . 1016/ j . nima.2026.171526.
- [104] H. Tajima, A. Okumura, and K. Furuta, “Studies of propagation mechanism of optical crosstalk in silicon photomultipliers,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 1049, page 168 029, Apr. 2023, ISSN: 0168-9002. DOI: 10 . 1016/ j . nima.2023.168029.
- [105] J. B. McLaughlin *et al.*, “Characterisation of SiPM Photon Emission in the Dark,” *Sensors*, volume 21, number 17, 2021, ISSN: 1424-8220. DOI: 10 . 3390/s21175947.
- [106] M. Lucchini, Y. Musienko, and A. Heering, “Experimental method to monitor temperature stability of SiPMs operating in conditions of extremely high dark count rate,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 977, page 164 300, Oct. 2020, ISSN: 0168-9002. DOI: 10 . 1016/ j . nima.2020.164300.
- [107] J. Zhang *et al.*, “Scintillation detectors with silicon photomultiplier readout in a dilution refrigerator at temperatures down to 0.2 K,” *Journal of Instrumentation*, volume 17, number 06, P06024, Jun. 2022, ISSN: 1748-0221. DOI: 10 . 1088 / 1748-0221/17/06/P06024.
- [108] XENON Collaboration, “Lowering the radioactivity of the photomultiplier tubes for the XENON1T dark matter experiment,” *European Physical Journal C*, volume 75, number 11, page 546, Nov. 2015, ISSN: 1434-6052. DOI: 10 . 1140/epjc/s10052-015-3657-5.

- [109] A. Cherubini and A. Bifone, “Hyperpolarised xenon in biology,” *Progress in Nuclear Magnetic Resonance Spectroscopy*, volume 42, number 1, pages 1–30, Jun. 2003, ISSN: 0079-6565. DOI: 10.1016/S0079-6565(02)00052-3.
- [110] Y. Zhang, J. R. G. Evans, and S. Yang, “Corrected Values for Boiling Points and Enthalpies of Vaporization of Elements in Handbooks,” *Journal of Chemical & Engineering Data*, volume 56, number 2, pages 328–337, Feb. 2011, ISSN: 0021-9568. DOI: 10.1021/je1011086.
- [111] M. Cha *et al.*, “Cryogenic fracturing for reservoir stimulation – Laboratory studies,” *Journal of Petroleum Science and Engineering*, volume 124, pages 436–450, Dec. 2014, ISSN: 0920-4105. DOI: 10.1016/j.petrol.2014.09.003.
- [112] S. K. Basu and E. A. Mccutchan, “Nuclear Data Sheets for A = 90,” *Nuclear Data Sheets*, volume 165, pages 1–329, 2020, ISSN: 0090-3752. DOI: 10.1016/j.nds.2020.04.001.
- [113] N. V. H. Viet *et al.*, “Simulation Study of Photon-to-Digital Converter (PDC) Timing Specifications for LoLX Experiment,” *IEEE Transactions on Nuclear Science*, volume 72, number 5, pages 1680–1685, May 2025, ISSN: 1558-1578. DOI: 10.1109/TNS.2025.3546965.
- [114] K. Fujii *et al.*, “High-accuracy measurement of the emission spectrum of liquid xenon in the vacuum ultraviolet region,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 795, pages 293–297, 2015, ISSN: 0168-9002. DOI: 10.1016/j.nima.2015.05.065.
- [115] K. Raymond, F. Retière, H. Lewis, *et al.*, *Stimulated Secondary Emission of Single Photon Avalanche Diodes*, 2024. DOI: 10.48550/arXiv.2402.09634.

- [116] D. Gallacher *et al.*, “Characterization of external cross-talk from silicon photomultipliers in a liquid xenon detector,” *European Physical Journal C*, volume 85, number 6, page 692, Jun. 2025, ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-025-14381-w.
- [117] B. Chana, “Characterization and operation of vacuum ultraviolet-sensitive silicon photomultipliers toward the nEXO experiment,” Ph.D. dissertation, Carleton University, 2023. DOI: 10.22215/etd/2023-15767. [Online]. Available: <https://repository.library.carleton.ca/concern/etds/kd17cv066>.
- [118] L. Galli *et al.*, “WaveDAQ: An highly integrated trigger and data acquisition system,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 936, pages 399–400, 2019, ISSN: 0168-9002. DOI: 10.1016/j.nima.2018.07.067.
- [119] D. Gallacher, “Comissioning and Calibration of a Small-Scale Modular Liquid Argon Detector, Argon-1, for Novel Studies in Background Rejection Techniques Towards Next-Generation Dark Matter Detectors,” M.S. thesis, Carleton University, 2021. HDL: 20.500.14718/42441. [Online]. Available: <https://hdl.handle.net/20.500.14718/42441> (visited on 07/08/2025).
- [120] G. L. Turin, “An introduction to matched filters,” *IRE Transactions on Information Theory*, volume 6, number 3, pages 311–329, 1960. DOI: 10.1109/TIT.1960.1057571.
- [121] P. Duhamel and M. Vetterli, “Fast fourier transforms: A tutorial review and a state of the art,” *Signal Processing*, volume 19, number 4, pages 259–299, Apr. 1990, ISSN: 0165-1684. DOI: 10.1016/0165-1684(90)90158-U.
- [122] F. Corsi *et al.*, “Modelling a silicon photomultiplier (SiPM) as a signal source for optimum front-end design,” *Nuclear Instruments and Methods in Physics Research*

- Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Frontier Detectors for Frontier Physics, volume 572, number 1, pages 416–418, Mar. 2007, ISSN: 0168-9002. DOI: 10.1016/j.nima.2006.10.219.
- [123] E. W. Weisstein, *Cross-Correlation theorem*, Text, 2025. [Online]. Available: <https://mathworld.wolfram.com/Cross-CorrelationTheorem.html>.
 - [124] C. Moler and S. Eddins, *Faster finite Fourier transforms MATLAB*, 2001. [Online]. Available: <https://www.mathworks.com/company/technical-articles/faster-finite-fourier-transforms-matlab.html>.
 - [125] S. Agostinelli, J. Allison, K. Amako, *et al.*, “Geant4—a simulation toolkit,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 506, number 3, pages 250–303, 2003, ISSN: 0168-9002. DOI: 10.1016/S0168-9002(03)01368-8.
 - [126] C. M. Poole *et al.*, “A CAD Interface for GEANT4,” *Australasian Physical & Engineering Science in Medicine*, Sep. 2012. DOI: 10.1007/s13246-012-0159-8.
 - [127] A. LaTorre and S. Seibert, *Fast Optical Monte Carlo simulation with Surface-Based geometries*, Chroma software whitepaper, 2011. [Online]. Available: <https://github.com/BenLand100/chroma/blob/master/doc/source/chroma.pdf>.
 - [128] L. Liu *et al.*, “2D microspatial distribution uniformity of photon detection efficiency and crosstalk probability of multi-pixel photon counters,” *Quantum Electronics*, volume 50, number 2, page 197, Feb. 2020, ISSN: 1063-7818. DOI: 10.1070/QEL17141.
 - [129] M. G. Boulay, V. Camillo, N. Canci, *et al.*, “SiPM cross-talk in liquid argon detectors,” *Frontiers in Physics*, volume 11, May 2023, ISSN: 2296-424X. DOI: 10.3389/fphy.2023.1181400.

- [130] R. Gibbons, H. Chen, S. J. Haselschwardt, *et al.*, “Why would you put a flashlight in a dark matter detector?” *Journal of Instrumentation*, volume 19, number 01, P01013, 2024. DOI: 10.1088/1748-0221/19/01/P01013.
- [131] S. Lee, W.-S. Choong, and J. W. Cates, “Characterization of external optical crosstalk reduction for SiPM-based scintillation detectors with an optical bandpass filter,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 1061, page 169 101, 2024, ISSN: 0168-9002. DOI: 10.1016/j.nima.2024.169101.
- [132] M. Wagenpfeil *et al.*, “Reflectivity of VUV-sensitive silicon photomultipliers in liquid Xenon,” *Journal of Instrumentation*, volume 16, number 08, P08002, Aug. 2021. DOI: 10.1088/1748-0221/16/08/P08002.
- [133] P. Lv, G. F. Cao, L. J. Wen, *et al.*, “Reflectance of Silicon Photomultipliers at Vacuum Ultraviolet Wavelengths,” *IEEE Transactions on Nuclear Science*, volume 67, number 12, pages 2501–2510, 2020. DOI: 10.1109/TNS.2020.3035172.
- [134] Hamamatsu Photonics K.K., *Hamamatsu VUV4 MPPC datasheet (S13370 series)*, Product datasheet, 2024. [Online]. Available: [https://hamamatsu-su/files/uploads/pdf/3_mppc/s13370_vuv4-mppc_b_\(1\).pdf](https://hamamatsu-su/files/uploads/pdf/3_mppc/s13370_vuv4-mppc_b_(1).pdf).
- [135] D. E. Aspnes and A. A. Studna, “Dielectric functions and optical parameters of Si, Ge, GaP, GaAs, GaSb, InP, InAs, and InSb from 1.5 to 6.0 eV,” *Physical Review B*, volume 27, number 2, pages 985–1009, Jan. 1983. DOI: 10.1103/PhysRevB.27.985.
- [136] D. T. Pierce and W. E. Spicer, “Electronic Structure of Amorphous Si from Photoemission and Optical Studies,” *Physical Review B*, volume 5, number 8, pages 3017–3029, Apr. 1972. DOI: 10.1103/PhysRevB.5.3017.

- [137] Formlabs, *Durable Resin datasheet*, Product datasheet, 2024. [Online]. Available: <https://formlabs.com/store/materials/durable-resin/>.
- [138] O. Njoya *et al.*, “Measurements of electron transport in liquid and gas Xenon using a laser-driven photocathode,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 972, page 163 965, Aug. 2020, ISSN: 0168-9002. DOI: 10 . 1016 / j . nima . 2020 . 163965.
- [139] X. Li *et al.*, *In-situ performance of FBK VUV-HD3 and HPK VUV4 SiPMs in the LoLX liquid xenon detector*, Oct. 2025. DOI: 10 . 48550 / arXiv . 2510 . 15270. arXiv: 2510 . 15270 [physics.ins-det]. [Online]. Available: <https://arxiv.org/abs/2510.15270v1>.
- [140] Y. Khazov, A. Rodionov, and F. G. Kondev, “Nuclear Data Sheets for A = 133,” *Nuclear Data Sheets*, volume 112, number 4, pages 855–1113, Apr. 2011, ISSN: 0090-3752. DOI: 10 . 1016 / j . nds . 2011 . 03 . 001.
- [141] J. Li *et al.*, “Waveform simulation in PandaX-4T,” *Chinese Physics C*, volume 48, number 7, page 073 001, Jul. 2024, ISSN: 1674-1137. DOI: 10 . 1088 / 1674 - 1137 / ad380f.
- [142] R. Acciarri *et al.*, “Effects of Nitrogen contamination in liquid Argon,” *Journal of Instrumentation*, volume 5, number 06, P06003–P06003, Jun. 2010, ISSN: 1748-0221. DOI: 10 . 1088 / 1748 - 0221 / 5 / 06 / P06003. arXiv: 0804 . 1217 [nucl-ex].
- [143] I. Murayama and S. Nakamura, “Time profile of the scintillation from liquid and gaseous xenon,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 763, pages 533–537, Nov. 2014, ISSN: 0168-9002. DOI: 10 . 1016 / j . nima . 2014 . 07 . 003.

- [144] G. Ribitzki *et al.*, “Electron densities and temperatures in a xenon afterglow with heavy-ion excitation,” *Physical Review E*, volume 50, number 5, pages 3973–3979, Nov. 1994. DOI: 10.1103/PhysRevE.50.3973.
- [145] S. Wu, “Nuclear data sheets for $A = 181$,” *Nuclear Data Sheets*, volume 106, pages 367–600, Nov. 2005. DOI: 10.1016/j.nds.2005.11.001.
- [146] International Atomic Energy Agency, *Livechart - table of nuclides: Nuclear structure and decay data*, IAEA Nuclear Data Section online database, 2025. [Online]. Available: <https://www-nds.iaea.org/relnsd/vcharthtml/VChartHTML.html#lastnuc=181Hf>.
- [147] McMaster Nuclear Reactor, *Core irradiations*, McMaster University, Nuclear Operations and Facilities, 2025. [Online]. Available: <https://nuclear.mcmaster.ca/products-services/core-irradiations/>.
- [148] P. M. Prajapati *et al.*, “Measurement of Neutron-Induced Reaction Cross Sections in Zirconium Isotopes at Thermal, 2.45 MeV and 9.85 MeV Energies,” *Nuclear Science and Engineering*, volume 171, number 1, pages 78–84, May 2012, ISSN: 0029-5639. DOI: 10.13182/NSE11-02.
- [149] P. Kienzle, *Neutron activation and scattering calculator*, NIST Center for Neutron Research, Oct. 2025. [Online]. Available: <https://www.ncnr.nist.gov/resources/activation/>.
- [150] XLZD Collaboration, “Neutrinoless double beta decay sensitivity of the XLZD rare event observatory,” *Journal of Physics G: Nuclear Particle Physics*, volume 52, number 4, page 045 102, Apr. 2025, ISSN: 0954-3899. DOI: 10.1088/1361-6471/adb900.

- [151] PIONEER Collaboration, *PIONEER: Studies of Rare Pion Decays*, Mar. 2022. DOI: 10.48550/arXiv.2203.01981. arXiv: 2203.01981 [hep-ex]. [Online]. Available: <http://arxiv.org/abs/2203.01981>.
- [152] EXO-200 Collaboration, “Search for Neutrinoless Double- β Decay with the Complete EXO-200 Dataset,” *Physical Review Letters*, volume 123, number 16, page 161 802, Oct. 2019. DOI: 10.1103/PhysRevLett.123.161802.
- [153] EXO Collaboration, “Observation of two-neutrino double-beta decay in ^{136}Xe with the EXO-200 detector,” *Physical Review Letters*, volume 107, number 21, page 212 501, Nov. 2011. DOI: 10.1103/PhysRevLett.107.212501.
- [154] nEXO Collaboration, *nEXO Pre-Conceptual Design Report*, Aug. 2018. DOI: 10.48550/arXiv.1805.11142. arXiv: 1805.11142 [physics]. [Online]. Available: <http://arxiv.org/abs/1805.11142>.
- [155] N. J. T. Smith, “The SNOLAB deep underground facility,” *European Physical Journal Plus*, volume 127, number 9, page 108, Sep. 2012, ISSN: 2190-5444. DOI: 10.1140/epjp/i2012-12108-9.
- [156] EXO Collaboration, “The EXO-200 detector, part II: Auxiliary systems,” *Journal of Instrumentation*, volume 17, number 02, P02015, Feb. 2022, ISSN: 1748-0221. DOI: 10.1088/1748-0221/17/02/P02015.
- [157] E. Aprile and T. Doke, “Liquid xenon detectors for particle physics and astrophysics,” *Review of Modern Physics*, volume 82, number 3, pages 2053–2097, Jul. 29, 2010. DOI: 10.1103/RevModPhys.82.2053.
- [158] G. Pereira *et al.*, “Energy resolution of the LZ detector for high-energy electronic recoils,” *Journal of Instrumentation*, volume 18, number 04, page C04007, Apr. 2023, ISSN: 1748-0221. DOI: 10.1088/1748-0221/18/04/C04007.

- [159] Z. Li *et al.*, “Simulation of charge readout with segmented tiles in nEXO,” *Journal of Instrumentation*, volume 14, number 09, P09020–P09020, Sep. 2019, ISSN: 1748-0221. DOI: 10.1088/1748-0221/14/09/P09020. arXiv: 1907.07512 [physics].
- [160] EXO-200 Collaboration, “Measurement of the scintillation and ionization response of liquid xenon at MeV energies in the EXO-200 experiment,” *Physical Review C*, volume 101, number 6, page 065 501, Jun. 2020. DOI: 10.1103/PhysRevC.101.065501.
- [161] M. Jewell *et al.*, “Characterization of an Ionization Readout Tile for nEXO,” *Journal of Instrumentation*, volume 13, number 01, P01006, Jan. 2018, ISSN: 1748-0221. DOI: 10.1088/1748-0221/13/01/P01006.
- [162] M. Watts and the nEXO collaboration, “nEXO photon detection system and read-out electronics,” *Journal of Instrumentation*, volume 20, number 06, page C06009, Jun. 2025, ISSN: 1748-0221. DOI: 10.1088/1748-0221/20/06/c06009.
- [163] I. J. Arnquist *et al.*, “Ultra-low radioactivity flexible printed cables,” *European Physical Journal Techniques and Instrumentation*, volume 10, number 1, page 17, Dec. 2023, ISSN: 2195-7045. DOI: 10.1140/epjti/s40485-023-00104-6.
- [164] G. Gallina *et al.*, “Characterization of the Hamamatsu VUV4 MPPCs for nEXO,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 940, pages 371–379, Oct. 2019, ISSN: 0168-9002. DOI: 10.1016/j.nima.2019.05.096.
- [165] R. H. M. Tsang *et al.*, “An integrated online radioassay data storage and analytics tool for nEXO,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 1055, page 168 477, Oct. 2023, ISSN: 0168-9002. DOI: 10.1016/j.nima.2023.168477.

- [166] S. Zhu and E. A. McCutchan, “Nuclear Data Sheets for A=214,” *Nuclear Data Sheets*, volume 175, pages 1–149, Jul. 2021, ISSN: 0090-3752. DOI: 10.1016/j.nds.2021.06.001.
- [167] T. J. Roosendaal *et al.*, *Ultra-pure Nickel for Structural Components of Low Radioactivity Instruments*, Aug. 2025. DOI: 10.48550/arXiv.2508.08230. arXiv: 2508.08230 [nucl-ex]. [Online]. Available: <http://arxiv.org/abs/2508.08230>.
- [168] E. Browne and J. K. Tuli, “Nuclear Data Sheets for A = 137,” *Nuclear Data Sheets*, volume 108, number 10, pages 2173–2318, Oct. 2007, ISSN: 0090-3752. DOI: 10.1016/j.nds.2007.09.002.
- [169] I. Esteban *et al.*, “The fate of hints: Updated global analysis of three-flavor neutrino oscillations,” *Journal of High Energy Physics*, volume 2020, number 9, page 178, Sep. 2020, ISSN: 1029-8479. DOI: 10.1007/JHEP09(2020)178.
- [170] nEXO Collaboration, “Development of a ^{127}Xe calibration source for nEXO,” *Journal of Instrumentation*, volume 17, number 07, P07028, Jul. 2022, ISSN: 1748-0221. DOI: 10.1088/1748-0221/17/07/P07028.
- [171] A. Jamil, “Rare Event Searches in Liquid Xenon with EXO-200, nEXO and Beyond,” Ph.D. dissertation, Yale University, 2022. [Online]. Available: https://bpb-us-w2.wpmucdn.com/campuspress.yale.edu/dist/0/1594/files/2022/10/Jamil_Ako_DISSERTATION.pdf.
- [172] V. N. Solovov *et al.*, “Measurement of the refractive index and attenuation length of liquid xenon for its scintillation light,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 516, number 2, pages 462–474, Jan. 2004, ISSN: 0168-9002. DOI: 10.1016/j.nima.2003.08.117.

- [173] L. M. Barkov *et al.*, “Measurement of the refractive index of liquid xenon for intrinsic scintillation light,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, Proceedings of the Sixth International Conference on Instrumentation for Experiments at e⁺ E- Colliders, volume 379, number 3, pages 482–483, Sep. 1996, ISSN: 0168-9002. DOI: 10.1016/0168-9002(99)00518-5.
- [174] E. Grace *et al.*, “Index of refraction, Rayleigh scattering length, and Sellmeier coefficients in solid and liquid argon and xenon,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 867, pages 204–208, Sep. 2017, ISSN: 0168-9002. DOI: 10.1016/j.nima.2017.06.031.
- [175] A. Sinnock and B. Smith, “Refractive indices of the condensed inert gases,” *Physical Review*, volume 181, number 3, pages 1297–1307, 1969, ISSN: 0031-899X. DOI: 10.1103/PhysRev.181.1297.
- [176] D. Bajpai *et al.*, “Validation of the VUV-reflective coating for next-generation liquid xenon detectors,” *Journal of Instrumentation*, volume 19, number 06, P06008, Jun. 2024, ISSN: 1748-0221. DOI: 10.1088/1748-0221/19/06/P06008.
- [177] T. Fontanot *et al.*, “POCAL: A Python-based library to perform optical coating analysis and design,” *Optics Continuum*, volume 2, number 4, pages 810–824, Apr. 2023, ISSN: 2770-0208. DOI: 10.1364/OPTCON.484972.
- [178] H. H. Li, “Refractive index of alkaline earth halides and its wavelength and temperature derivatives,” *Journal of Physical and Chemical Reference Data*, volume 9, number 1, pages 161–290, Jan. 1980, ISSN: 0047-2689. DOI: 10.1063/1.555616.

- [179] A. D. Rakić, “Algorithm for the determination of intrinsic optical constants of metal films: Application to aluminum,” *Applied Optics*, volume 34, number 22, pages 4755–4767, Aug. 1995, ISSN: 2155-3165. DOI: 10.1364/AO.34.004755.
- [180] W. S. M. Werner, K. Glantschnig, and C. Ambrosch-Draxl, “Optical Constants and Inelastic Electron-Scattering Data for 17 Elemental Metals,” *Journal of Physical and Chemical Reference Data*, volume 38, number 4, pages 1013–1092, Dec. 2009, ISSN: 0047-2689. DOI: 10.1063/1.3243762.
- [181] C. N. Panagopoulos, A. Markaki, and E. Hontzopoulos, “Excimer laser treatment of copper-coated mild steel,” *Journal of Materials Science*, volume 32, number 6, pages 1425–1430, Mar. 1997, ISSN: 1573-4803. DOI: 10.1023/A:1018589514262.
- [182] C. Liu *et al.*, “Vacuum ultraviolet high-reflectance aluminum mirrors on copper substrate for application in noble liquid time projection chamber,” *Vacuum*, volume 197, page 110 806, Mar. 2022, ISSN: 0042-207X. DOI: 10.1016/j.vacuum.2021.110806.
- [183] K. Fujii *et al.*, “High-accuracy measurement of the emission spectrum of liquid xenon in the vacuum ultraviolet region,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 795, pages 293–297, 2015, ISSN: 0168-9002. DOI: <https://doi.org/10.1016/j.nima.2015.05.065>.
- [184] A. Baldini *et al.*, “Absorption of scintillation light in a 100l liquid xenon γ -ray detector and expected detector performance,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, volume 545, number 3, pages 753–764, 2005, ISSN: 0168-9002. DOI: 10.1016/j.nima.2005.02.029.

- [185] S. Navas *et al.*, “Review of particle physics,” *Physical Review D*, volume 110, page 030001, 2024, Monte Carlo techniques review, Sec. 42. DOI: 10 . 1103 / PhysRevD . 110 . 030001.
- [186] S. Vinogradov, “Analytical models of probability distribution and excess noise factor of solid state photomultiplier signals with crosstalk,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, New Developments in Photodetection NDIP11, volume 695, pages 247–251, Dec. 2012, ISSN: 0168-9002. DOI: 10 . 1016 / j . nima . 2011 . 11 . 086.
- [187] G. Richardson *et al.*, *Sensitivity of nEXO to ^{136}Xe Charged-Current Interactions: Background-free Searches for Solar Neutrinos and Fermionic Dark Matter*, Jun. 2025. DOI: 10 . 48550 / arXiv . 2506 . 22586. arXiv: 2506 . 22586 [nucl-ex].
- [188] S. Davidson, E. Nardi, and Y. Nir, “Leptogenesis,” *Physics Reports*, volume 466, number 4, pages 105–177, 2008, ISSN: 0370-1573. DOI: 10 . 1016 / j . physrep . 2008 . 06 . 002.
- [189] D. Gallacher, *G4-underground-eXT*, GitHub repository, 2024. [Online]. Available: <https://github.com/dgallacher1/G4-underground-eXT>.
- [190] L. Baudis *et al.*, “Electron transport measurements in liquid xenon with Xenoscope, a large-scale DARWIN demonstrator,” *The European Physical Journal C*, volume 83, number 8, page 717, Aug. 10, 2023, ISSN: 1434-6052. DOI: 10 . 1140 / epjc / s10052 - 023 - 11823 - 1.