

NORTHWESTERN UNIVERSITY

An Electron-Recoil Dark Matter Search Using the Next Generation of SuperCDMS HVeV
Detectors Operated at SNOLAB

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Physics

By

Kyle Kennard

EVANSTON, ILLINOIS

August 2026

© Copyright by Kyle Kennard 2026

All Rights Reserved

ABSTRACT

Over the past several decades, the WIMP (weakly interacting massive particle) has motivated searches for dark matter (DM) in the $\gtrsim$ GeV range. However, as searches continually report null results while constraining the viable parameter space for WIMP models, attention in the field has gradually shifted to lower mass “hidden-sector” DM models where parameter space is still relatively unexplored. The energies imparted to Standard Model (SM) particles by DM in this mass range are eV-scale and thus sensitivity to these models requires excellent energy resolution. For this purpose, the SuperCDMS HVeV program utilizes gram-scale cryogenically cooled silicon crystals equipped with transition-edge sensors (TESs) as phonon detectors to achieve eV-scale resolution. HVeV detectors may be operated with an electric field over the substrate to further amplify ionizing signals via the Neganov-Trofimov-Luke (NTL) effect. HVeV detectors have been utilized a number of times to obtain competitive limits to hidden-sector DM candidates interacting with the SM via electron-recoils. However, despite advances detector resolution, sensitivity across the field remains limited by excesses at low energy of unknown origin, communally referred to as the low energy excess (LEE), which are the dominant feature in the signal region of interest.

In this thesis, I report results from the latest data-taking campaign with HVeV detectors operated at the SNOLAB deep underground facility, which aimed to investigate the LEE. In particular, I focus on observations from a newly fabricated HVeV capable of sub-eV resolution—the current best for HVeV detectors. This improvement in resolution from past HVeV runs enabled observation of eV-scale structures in photon calibration data which revealed evidence of a previously unobserved energy loss mechanism. Using this detector, I characterize non-ionizing components of the LEE for inclusion as a background in an electron-recoil dark matter (ERDM) search. Finally, I conclude with world-leading exclusion limits on a number of ERDM candidates.

ACKNOWLEDGMENTS

Over the past five years—and especially while writing this thesis—it has often occurred to me how lucky I am to have the mentorship and support that I did. Without them, this work would not have been possible. This luck is evidenced, in part, by the long list of people to whom I owe thanks. Throughout writing this thesis, I have continually added to this list and yet, in spite of this, I am sure that I have neglected some.

I first would like to thank my principle advisor, Enectalí “Tali” Figueroa-Feliciano. Tali puts a lot of trust into his students; I recall at one point him telling me to treat his lab as though it were mine. Even though I strongly suspect it was not always the path of least resistance, I was given a lot of freedom to pursue my research interests. I am immensely grateful to have been given the opportunity to be a part of the work presented in this thesis and for that I have him to thank.

I am similarly grateful for Dan Baxter, who, while not officially my advisor until the latter half of my PhD, was always a voice of encouragement, especially when I needed it most. There were a few times throughout my graduate studies when I felt completely overwhelmed or incapable, and, in those times, I could rely on Dan to set my head straight.

I must thank Ziqing Hong who, despite not technically being my advisor, taught me more about HVeV detectors, experimental design, data analysis, hell, maybe all of physics, than anyone. I sent Ziqing, at times, a nearly constant stream of questions, which he would entertain and discuss at length with me. At some point he had joked that, through these conversations, we would reconstruct all of physics from first principles. Despite our best efforts, I don’t think we are quite there yet, but I hope that we can continue to work towards that goal in the coming years.

To Noah Kurinsky and Kelly Stifter, thank you for welcoming me into your research group. I learned a lot in the short time I was at SLAC and am greatly looking forward to working with you

when I arrive there in the coming fall.

I owe a lot to all the students and postdocs in the research group at Northwestern, which has grown considerably from when I started. To Valentina Novati and Benjamin Schmidt, I thank you for not only for teaching me the foundational skills that I needed to carry out this work, but for holding me to a high standard. Though you moved on to bigger and brighter things early on in my graduate studies, I still aspire to your level of quality and precision. I also owe a great deal to Ran Chen and Tom (Runze) Ren, both for showing me the ropes but also for being excellent (if not unrealistic) role models to strive for. To Tom, I hope you find that this work is a worthy continuation of your contributions to SuperCDMS. To Ran, it would be impossible to repay you properly for the countless times you helped me with coding or in the lab so all I can really say is that it meant a lot to me. I didn't fully realize how much I had relied on your help until you graduated. Finally, a huge thank you to everyone in the research group. It was a pleasure to work along side you all!

I have many people to thank at SNOLAB for the positive memories I have there. In particular, thank you to Ruslan Podviianiuk for being an excellent roommate at the SuperCDMS shifter house. I miss the talks that you, me, and Emanuele would have about language, science, and faith. My time at the shifter house was formative and I will always remember it fondly thanks to you. Additionally, I have Matt Stukel, Andy Kubik, and Marco Olivares to thank for not just supporting the work at CUTE, but for somehow managing to make waking up early to go underground fun (although, I will confess, I do not miss waiting in line to shower).

There are many people with whom I have had the good fortune to collaborate on this work and other projects that deserve more than just the thanks I can provide here. First, I thank Mason Buchanan and Julius Viol who collaborated with me on essentially all of the work shown in this thesis. Also, a big thanks to Aviv Simchony, Aditi Pradeep, and Jamie Ryan for their contributions

to the photon calibration shown in Chapter 4 and for their support and insight. I am very much looking forward to a few good years at SLAC with you all! Both the work presented in this thesis and my development as a researcher benefited immensely from discussions with the HVeV Run 4 analysis team, who pioneered many of the methods I used here. In particular, I would like to thank Imran Alkhatib, Enze Zhang, Huanbo Sun, Ata Sattari, and Sudipta Das for the insights they provided. Finally, I would like to thank Miriam Diamond, Wolfgang Rau, Madeleine Zurowsky, Shubham Pandey, and Musaab Al Bakry for their time spent reviewing the analysis and providing feedback.

A special thanks goes to Emanuele Michielin, with whom I worked closely on the majority of this work. From the beginning, Emanuele treated me not just as a student but as a partner. I learned a lot under his supervision, not the least of which was how to make pasta properly (it's a long story, involving a late night between underground shifts at SNOLAB). I am grateful for his tolerance of my constant messages to discuss analysis and even more so for being the voice of reason when I began to panic, as I often do. Most importantly, I have him to thank for showing me that there is no better reason to pursue our work other than because we love it. It was a pleasure to work with you, my friend, and I sincerely hope that we collaborate again in the future.

I would like to thank my family for their constant support. It was hard being away from my siblings and my mother over these past years, especially after a lifetime of being able to rely on their constant presence. To my older sister, Callie, I miss your advice and your candor. I can't wait to come home and go bird-watching with you. More than anything, I miss simply talking with you. To my twin brother, Brandon, before I moved to Chicago, I couldn't imagine a day without you. While I think we handled it well enough, I think it's about time that we go on a hike again. To my mother, who was brave enough to read this thesis (and to suggest edits, which I heavily relied on), there's more that you've done for me than I could possibly hope to acknowledge here. I hope that

this thesis serves as a tribute to your love and support.

And finally, to my girlfriend, Tessa, who I need to thank for, well, everything. For driving with me to Chicago for the first time. For staying on the phone with me for hours while we were apart. For taking a huge leap of faith and moving across the country so that we could be together. For patiently listening while I thought aloud and for offering reassurance in those times when graduate school felt overwhelming. For being my home while we were away from ours. None of this would have been possible without you. I love you, Tessa.

This thesis is dedicated to my twin brother, Brandon, as consolation for the adventures foregone while writing it and as a promise for many more to come.

TABLE OF CONTENTS

Acknowledgments	4
List of Figures	14
List of Tables	19
Chapter 1: Dark Matter: The Missing Mass in Our Universe	20
1.1 Evidence for Dark Matter	20
1.1.1 Galactic Rotation Curves	22
1.1.2 Gravitational Lensing of Galaxy Clusters	23
1.1.3 The Power Spectrum of the Cosmic Microwave Background	25
1.1.4 The Influence of Dark Matter in Structure Formation	28
1.2 What is Dark Matter?	29
1.2.1 WIMP Dark Matter	32
1.2.2 Hidden Sector Dark Matter	37
1.3 This Work: Continuing the Search for Dark Matter at Low Masses	40
Chapter 2: SuperCDMS HVeV Detectors	41

	10
2.1 Measurement Scheme	42
2.1.1 Phonon Sensors	42
2.1.2 SQUID-Based Readout	50
2.2 Detector Response to Particle Interactions	55
2.2.1 Ionization in Silicon	57
2.2.2 Charge-Trapping and Impact Ionization	59
2.3 Signal Models	62
2.3.1 Dark Matter-Electron Scattering	63
2.3.2 Dark Photon Absorption	68
2.3.3 Axion-like Particle Absorption	71
2.4 The Development of HVeV Detectors	72
2.4.1 HVeV Run 3	74
2.4.2 HVeV Run 4	76
Chapter 3: HVeV at CUTE	78
3.1 Experimental Setup	78
3.1.1 The CUTE Facility at SNOLAB	78
3.1.2 The Detector Payload	80
3.1.3 Run Summary	83
3.2 Event Reconstruction	85
3.2.1 Offline Triggering	85

	11
3.2.2 Optimal Filtering	85
Chapter 4: Photon Calibration of NFCSO	91
4.1 Calibration Strategy	94
4.2 Electrical Crosstalk	96
4.2.1 An Independent Measurement of CTA	98
4.3 Fitting and Results	104
4.4 Another Measurement of Energy Loss	107
4.5 Implications of Energy Loss on Previous Calibrations of HVeV Detectors	110
Chapter 5: Background Observations at 0V	114
5.1 The Shape of the 0V Energy Spectrum	114
5.2 The 0V Rate Over Time	119
5.3 Comparison with HV Spectra	121
Chapter 6: The HVeV Run 5 ERDM Search	128
6.1 Data Acquisition and Analysis Strategy	128
6.2 Data Selection	129
6.2.1 Live Time Cuts	129
6.2.1.1 Mean Baseline Cut	129
6.2.1.2 Coincidence Cut	130
6.2.1.3 Burst Event Cut	131

6.2.1.4	Excess Rate Cut	134
6.2.2	Quality Cuts	135
6.2.2.1	χ^2 Cut	135
6.2.2.2	Final Spectrum After Cuts	137
6.3	Limit Setting using the Profile Likelihood Ratio Method	138
6.4	Fitting the Observed Energy Spectrum	140
6.5	Exclusion Limits	148
6.5.1	Final Limits	151
Chapter 7:	Conclusion	153
7.0.1	Calibration and Energy Loss	154
7.0.2	Future Investigations of Excesses at Low Energy	156
7.0.2.1	The Ionizing Background	157
7.0.2.2	The Non-Ionizing Background	159
7.0.3	Outlook	160
References		169
Appendix A:	Burst Events	171
A.1	HVeV Run 3	171
A.2	HVeV Run 4	172
A.3	HVeV Run 5	173

A.4	Triggering Efficiency of Burst Events in Run 5	173
A.5	Discussion	178
Appendix B: The Profile Likelihood Ratio with Degenerate Signal and Background Models		180

LIST OF FIGURES

1.1	Rotation curve of M33.	23
1.2	The Bullet Cluster.	24
1.3	CMB power spectrum.	27
1.4	Dark matter candidates and their corresponding masses.	31
1.5	Comoving number density of freeze-out dark matter.	36
1.6	Feynman diagrams of hidden-sector dark matter interactions with electrons.	38
2.1	SuperCDMS detector schematic.	41
2.2	Diagram of a QET.	43
2.3	Resistance vs. temperature for a superconducting film.	44
2.4	TES bias circuit.	45
2.5	HVeV channel layout and close-up of a QET.	49
2.6	Schematic of a DC SQUID.	50
2.7	IV curve for a Josephson Junction.	51
2.8	Qualitative relationship of screening current in a DC SQUID as a function of applied magnetic flux.	53

2.9	Voltage vs. magnetic flux in a DC SQUID.	54
2.10	Simplified flux-locked loop.	55
2.11	Probability to produce n eh pairs in silicon as function of electron-recoil energy. . .	58
2.12	Measurements of ionization yield from nuclear recoils in silicon.	59
2.13	Illustration of charge-trapping and impact-ionization.	60
2.14	Example of how an energy spectrum varies depending on charge-trapping and impact-ionization.	62
2.15	Differential rate for DM-electron scattering in silicon.	67
2.16	Differential event rate as a function of phonon energy in the detector for DM- electron scattering.	68
2.17	Example signal models for dark photon absorption.	70
2.18	Example signal models for axion-like particle absorption.	73
2.19	HVeV exclusion limits for DMe scattering with a light mediator as of HVeV Run 4.	74
2.20	Detector payloads for HVeV Runs 3 and 4 and their associated energy spectra. . . .	75
2.21	An open LED module from HVeV Run 4.	77
3.1	Diagram of the CUTE facility at SNOLAB.	79
3.2	Cross-section of an SiO ₂ HVeV.	81
3.3	The HVeV Run 5 detector payload.	82
3.4	Current spectrum density of NFCSO.	87
3.5	Signal templates for NFCSO.	89

4.1	Simulation of a charge paths in a surface-trapped event.	93
4.2	Crosstalk trace.	96
4.3	Crosstalk amplitude vs. LED amplitude.	97
4.4	0th peak as a function of λ	102
4.5	Measured λ vs. true λ	103
4.6	Detector response model fits to data taken at CUTE.	105
4.7	0th peak vs. λ including systematics.	106
4.8	Finer structure of eh peaks due to surface-trapping.	107
4.9	Comparison of detector packaging between SLAC and CUTE.	108
4.10	Detector response model fit to SLAC data.	110
4.11	Energy loss measured at CUTE and SLAC.	111
5.1	Mean of 1eh peak over time.	115
5.2	0V spectrum at different times.	117
5.3	Shape parameters for the 0V spectrum plotted against time.	118
5.4	Shape of the 0V spectrum.	119
5.5	Rate of 0V background plotted against time.	120
5.6	Leakage model.	123
5.7	HV data compared to models of an ionizing 0 V background.	124
5.8	Model fit to HV data.	125
5.9	Rate of non-ionizing background in HV data.	126

6.1	Distribution of the mean baseline in an example series.	130
6.2	Example of a coincident event.	131
6.3	Burst cut in open data.	132
6.4	Poisson fits to dt spectra in open data.	133
6.5	The energy spectrum in open data before and after the burst event cut.	134
6.6	χ^2 of injected pulses as a function of OFL.	135
6.7	Efficiency of the χ^2 cut.	136
6.8	The energy spectrum in open data before and after all data selections.	137
6.9	Electron-recoil dark matter signal models.	141
6.10	Calibration uncertainty.	144
6.11	Reference and global fits to open data.	146
6.12	The NLL as a function of the cross-section.	147
6.13	The test statistic plotted with the reference cross-section.	148
6.14	Exclusion limits for DM-electron scattering set using open data.	149
6.15	Exclusion limits set on the kinetic mixing parameter via dark photon absorption, ϵ (left) and for the axio-electric coupling parameter, g_{aee} via ALP absorption.	150
6.16	Final exclusion limits for DM-electron scattering.	151
6.17	Final exclusion limits set on the kinetic mixing parameter via dark photon absorption, ϵ (left) and for the axio-electric coupling parameter, g_{aee} via ALP absorption.	152
6.18	Comparison of final exclusion limits for DM-electron scattering (heavy mediator) to SENSEI and DAMIC-M.	152

7.1	Updated comparison of exclusion limits for DM-electron scattering (light mediator).	153
7.2	Comparison of the HV and 0V spectra.	158
A.1	Example of a burst event from HVeV Run 3.	172
A.2	Rate of burst events in HVeV Run 5.	174
A.3	Contributions of burst events to energy spectra at different voltages.	175
A.4	Integral of saturated pulses as a function of voltage.	176
A.5	Triggering efficiency of secondaries in burst events.	177
B.1	Simulated energy spectrum.	181
B.2	Test statistic for degenerate signal and background models.	183
B.3	Distribution of the test statistic for degenerate signal and background models. . . .	184

LIST OF TABLES

4.1	Summary of fitting parameters in the detector response model.	95
4.2	Results of the detector response model fit to data taken at CUTE.	104
4.3	Calibration results for SLAC and CUTE.	109
4.4	HVeV Run 4 calibration parameters.	112
5.1	Best fit parameters of the normalized rate model (Eqn. 5.3) to 0V data.	121
6.1	Summary of the remaining live time in open data after each data selection.	137
6.2	Nuisance parameters in the profile likelihood.	143

CHAPTER 1

DARK MATTER: THE MISSING MASS IN OUR UNIVERSE

The term “dark matter” originally referred to matter that was just that: dark or non-luminous. In that sense, the concept and discussion of dark matter has been around since the late 1800s, when astronomers began to contemplate the existence of matter in the galaxy that could not be observed with a simple telescope. Somewhat ironically, the first few attempts to estimate the amount of dark matter (see those by Henri Poincaré [1] and Jacobus Kapteyn [2], for example) concluded that its presence was negligible in comparison to the amount of “shining matter” such as stars. However, as technological advances in astronomy enabled precise observations of our Milky Way and beyond, the evidence for a Universe dominated by a mysterious and as-of-yet undiscovered “dark matter” accumulated. In this way, dark matter as we understand it has evolved from being a possible systematic uncertainty on astronomical measurements into what is now one of the most interesting and fundamental questions in physics.

In this chapter, I explain the evidence which has motivated searches for dark matter in spite of decades of null results. In the first section, I explain what we know about dark matter. In the second section, I will highlight some of the leading candidates proposed as possible constituents of dark matter, and in the third section, I will summarize the current efforts to detect them.

1.1 Evidence for Dark Matter

Fritz Zwicky’s observations of the cluster of galaxies called the Coma cluster are often cited as the first indication of the existence of dark matter [3]. In his work [4], Zwicky estimated the mass of the Coma cluster as the product of the number of galaxies and the average mass of a galaxy.

Using the cluster's mass, M , and radius, R , he estimated the gravitational potential energy in the system by taking the cluster to be a sphere with a uniform mass density wherein the gravitational potential energy can be written as shown in Eqn. 1.1.

$$U_g = -\frac{3GM^2}{5R} \quad (1.1)$$

G in Eqn. 1.1 is the gravitational constant. The virial theorem (Eqn. 1.2) provides a relationship between the gravitational potential energy and the kinetic energy, K contained in the cluster:

$$\langle K \rangle = -\frac{1}{2}\langle U_g \rangle \quad (1.2)$$

Eqn. 1.2 can be used to calculate the velocity dispersion of the galaxies contained in the cluster since $K = \frac{M}{2}\langle v^2 \rangle$.

From this line of reasoning, Zwicky found that the velocity dispersion for galaxies in the Coma cluster should be ~ 80 km/s. However, his measurements of the motion of galaxies in the cluster were an order of magnitude larger, which implied that there was a far larger amount of mass contained in the cluster than one might estimate from the luminosity of the cluster. From this, he made the controversial claim that if these observations were confirmed then we would be left with the “surprising result that dark matter is present in much greater amount than luminous matter” [4]. It is important to note that at the time Zwicky made this statement, the term “dark matter” did not refer to the exotic non-baryonic matter to which the term typically refers today. At the time it literally meant matter that was not luminous, hence “dark.” However, even without suggesting the presence of an exotic type of matter, the conclusion Zwicky drew from his observations was surprising, given that previous measurements found that the amount of dark matter contained in local galaxies was small [1]. For decades after publication of Zwicky's observations of the Coma

cluster, the astronomical community was skeptical of his findings and, in particular, his dark matter hypothesis. It wasn't until the 1970s that Zwicky's hypothesis became largely accepted. At that point, numerous observations seemed to suggest dark matter made up a majority of the mass in most galaxies. The following section will summarize the evidence that supported the dark matter hypothesis as well as leading to the more concerning conclusion of its non-baryonic composition. This is by no means an exhaustive summary—for that I would refer the reader to reviews given in Refs. [3], [5] and [6]—but is enough to provide the reader with an understanding for the widespread acceptance of the existence of dark matter (DM).

1.1.1 Galactic Rotation Curves

One of the first hints that the majority of the universe must be composed of an undiscovered and non-baryonic form of dark matter came from observations of the radial velocity of stars in a galaxy as a function of the distance from the galactic center, referred to as a rotation curve. Applying the virial theorem to individual galaxies in a similar manner that Zwicky applied it to the Coma cluster, one can predict the shape of a galaxy's rotation curve based on the mass profile estimated from the luminosity of the galaxy. (See the dashed line in Fig. 1.1.) However, observations consistently found that the rotation curves of galaxies “flattened” for large distances from the center of the galaxy rather than decreasing as would be expected from observations of the stellar disk. (See the points in Fig. 1.1.)

Similar to Zwicky's observations of the Coma cluster, this implies that stars and luminous material in galaxies make up only a small fraction of the total mass.

On its own, the evidence from rotation curves does not require an exotic type of matter. For example, it is possible that the “missing mass” could be explained by gas and dust that may not emit strongly in the visible spectrum. While the mass of hot gas in a galaxy can be estimated via X-ray observations (*e.g.* Ref. [8]), if the gas is sufficiently cool, then it will not emit strongly and




Figure 1.1: The measured rotation curve of M33 is shown by the points. A best-fit model (solid line) consisting of contributions from the stellar disk (short dashed), gas (long dashed), and the halo (dot-dashed). Note the difference between the rotation curve of the stellar disk and the measured data which exhibits a comparatively flat rotation curve. Figure taken from Ref. [7].

will therefore not be visible through direct observation. However, we will show in the next section that observations of galactic clusters via gravitational lensing strongly suggest that gas and dust make up only a small fraction of the mass contained in galaxies, thus leaving the problem of the “missing mass” open.

1.1.2 Gravitational Lensing of Galaxy Clusters

Gravitational lensing is an effect predicted by Einstein’s theory of general relativity (GR). Rather than being a fundamental force as postulated by Newton, gravity in GR is a measure of the curvature of spacetime due to the presence of a massive body—the more massive the body, the more spacetime curves around it. A notable consequence of this is the bending of light around massive bodies, which we refer to as gravitational lensing. Gravitational lensing is a useful tool for measuring the mass of large systems, such as galaxies and galaxy clusters since the lensing does

not depend on the luminosity or dynamics of constituent bodies within the system.¹

The Bullet Cluster is a famous example of how gravitational lensing provided evidence in support of dark matter. Despite its name, the Bullet Cluster is actually the remains of a collision of two galaxy clusters consisting of a smaller “bullet” that passed through a larger cluster. During galaxy cluster collisions, it is unlikely for galaxies themselves or the stars they contain to collide directly. However, the gas coronas surrounding galaxies extend far beyond the radius of the galax-




Figure 1.2: A composite image of the Bullet Cluster. X-ray observations reveal that most of the hot gas that was stripped out of the clusters during the collision lies in the areas highlighted in red. Observations of the mass density via gravitational lensing show that most of the mass is contained in the blue regions, which suggests that only a small portion of the mass contained in the clusters can be attributed to gas. Color intensity is proportional to density. Credit: NASA.

ies themselves and are therefore far more likely to interact. As a result, when the galaxy clusters collide, stars and galaxies continue their paths unimpeded while gas experiences a drag force due

¹For the purposes of this thesis, it is sufficient for the reader to understand that the strength of the lensing effect is proportional to the mass between the observer and a distant object/system, for this allows one to deduce the mass of a system based on observations of the lensing. However, for an excellent introduction to gravitational lensing in the context of DM and without the need for a deep understanding of GR, I refer to the reader to Chapter 2 of Ref. [5].

to interactions with gas coronas from the other cluster. The end result is that gas from both clusters gets largely separated out from the galaxies and stars. Measurements of the mass distribution of the Bullet cluster show that most of the mass in the system is contained in the regions highlighted in blue in Figure 1.2—these are the regions that contain the clusters’ galaxies and stars. X-ray observations reveal that the gas in the clusters is distributed in regions highlighted in red and contains less than 1% of the total mass of the system. This is strong evidence that gas contained in galaxies cannot explain the flat rotation curves discussed in Section 1.1.1. Note as well that the Bullet Cluster is not the only example of this. See, for example, Refs. [9]–[11].

The combination of the observations shown in this section and in Section 1.1.1 can tell us about an important feature of dark matter. In Section 1.1.1, we found rotation curves that implied that mass from dark matter was distributed far past the radius of the galaxy (as defined by the extent of luminous matter). Therefore, similar to the gas coronas contained in the clusters, one might expect dark matter to separate out during cluster collisions. However, observations of the Bullet Cluster and other similar clusters show that most of the mass in the cluster does not separate out as the gas does, which tells us that dark matter is relatively collisionless. From a cosmological viewpoint, this is a significant finding; we will show in the next section that the existence of a non-luminous and (nearly) non-interacting type of matter in the early universe would have an observable effect on structures in the universe today.

1.1.3 The Power Spectrum of the Cosmic Microwave Background

The commonly accepted model for the early universe suggests that all mass and energy was condensed into a singularity before undergoing a period of rapid expansion in a process called inflation. The early universe was incredibly dense and hot, and particles had too much energy to form stable bound states—there were no atoms or molecules but, instead, a hot “soup” of elementary particles. As the universe expanded, the energy it contained was diluted over a larger volume, so

the universe also cooled down. Within the first few minutes, the early universe had cooled enough that electrons and protons could bind to form hydrogen in a time period referred to as the epoch of recombination. During recombination, photons released by the electrons as they fell into a bound state were able to travel unimpeded for the first time; prior to this point, the universe was essentially a plasma, so any radiation was quickly scattered. As a result, we can still observe this light which now forms a nearly uniform background of radiation referred to as the Cosmic Microwave Background (CMB).²

Since most photons in the CMB have traveled without scattering since the epoch of recombination, the CMB is a valuable tool for studying the early universe. The CMB is nearly a perfect blackbody with a mean temperature of ~ 2.7 K, but small anisotropies in the temperature of the CMB can be used to probe the distribution and composition of matter in the early universe. To understand why, let us consider the distribution of matter at very short timescales after the Big Bang.

Let's first imagine that matter in the early universe is first uniformly distributed. Since the universe at this time existed on such small scales, quantum effects were dominant and thus quantum fluctuations would initially create anisotropies in the density of the matter and energy—"clumping" it together in some places. The matter in these clumps would experience two competing forces: gravity would act to pull the matter closer together while radiation in the clump would exert a pressure that would push the matter farther apart. These competing forces created an oscillation similar to a sound wave and resulted in oscillations in the baryon³ density called Baryon Acoustic Oscillations (BAO).

²While the binding of electrons to hydrogen emits photons in the visible and infrared parts of the electromagnetic spectrum, the CMB emits most strongly in microwaves. The reason for this is that as the universe expanded, so too did the wavelength of light in the CMB.

³In general, the contribution from leptons is ignored in the early universe since the stable states for leptons are many orders of magnitude lower in mass compared to stable states of baryons. As a result, the density of matter in the anisotropies discussed will be dominated by baryons.

Now consider if a very weakly interacting dark matter existed in this period in the early universe. This dark matter's reluctance to interact with radiation or other particles would make it much more resistant to the radiation pressure and would therefore “clump” more easily. Dark matter (DM) would then essentially serve to increase the size of the anisotropies in the matter density since its mass contribution would increase the gravitational pull towards an over-fluctuation in density without being strongly affected by the radiation pressure.

During the epoch of recombination, photons released near clumps of matter were red-shifted as a result of escaping the gravity of the higher-density region, thus leading to lower-energy (or colder) areas in the CMB. Therefore, the size of anisotropies in the CMB can tell us about the density of matter in the early universe, which in turn, depends on the amount of DM. Figure 1.3 shows the power spectrum of the CMB measured by *Planck 2018* [12]. The relative height of the




Figure 1.3: The temperature power spectrum of the CMB as measured by *Planck 2018* (points). The best model fit is shown in blue. The x-axis corresponds to the multipole moment of the corresponding spherical harmonic. The y-axis corresponds to the amplitude of the various spherical harmonics of the multipole moment ℓ averaged over all orders. From Ref. [12].

peaks of the model shown in Figure 1.3 depend on features in the early universe such its curva-

ture and the amount of matter. The x-axis can be thought of as the scale of angular anisotropies in the temperature of the CMB; a *larger* value for ℓ corresponds to a *smaller* fluctuation in the temperature. The fact that the first peak occurs at $\ell \sim 200$ tells us about the curvature of the universe; in this case, that the universe is incredibly flat. If the universe were curved, light from hotter areas in the CMB would be bent towards (away) if the curvature were positive (negative) and thus make the hot area appear larger (smaller). Thus, if the peaks were shifted to lower values of ℓ , the universe would be positively curved, and if they were shifted to higher values of ℓ , the universe would be negatively curved. The heights of the two other peaks depend on the matter density in the universe—the more matter there is, the more anisotropies will be introduced due to BAOs. Finally, the relative height of the 2nd and 3rd peaks tell us about the amount of dark matter in the universe. Dark matter clumps like baryonic matter, but is not strongly affected by radiation pressure and therefore will lead to finer scale anisotropies in the CMB. The prominence of the 3rd peak tells us that DM makes up roughly 84% of the total matter in the universe [12]. Thus, not only do astrophysical observations discussed in the previous subsections strongly imply the existence of a non-luminous and weakly interacting matter, cosmological observations find it impossible to explain the structure of the CMB without it.

1.1.4 The Influence of Dark Matter in Structure Formation

The evidence provided by the CMB that DM existed in the early universe and seeded the initial matter anisotropies in the universe suggests that it played a major role in the formation of large structures such as galaxies and galaxy clusters. In fact, hints that DM was needed in order to explain structure formation could be found from early numerical simulations that showed that galaxies were not rotationally stable when just considering a flat disk of stars [13]. It was later demonstrated that, in these simulations, by embedding the galaxy in a massive spherical “halo”, the galaxy would become stable and thus agree with observations [14], once again supporting the

hypothesis that galaxies contained a large amount of non-luminous matter.

On a larger scale, cosmology turned to simulations to understand the structure of the universe given a set of initial conditions, in particular, the scale of density fluctuations (evidenced by measurements of the CMB shown in the last section) and the composition of the universe into matter, radiation, and other fields [15]. While these simulations were not necessarily sensitive to types of interactions between DM and other particles, they did consistently reveal that DM must be mostly non-relativistic or “cold” in order to explain the scale of structures that are observed now (see Ref. [16] for instance). We shall see in Sec. 1.2 how this conclusion influenced models for what DM could be composed of—in particular, this ruled out a popular hypothesis that neutrinos could make up a large portion of DM.

Over the past few decades, cosmological simulations of structure formation since the Big Bang have grown increasingly sophisticated and continued to prove a valuable tool for probing properties of DM. In that time, while it has remained certain that “hot” or relativistic DM can not explain the observed structures in the universe, precision measurements of the CMB by the Dark Energy Spectroscopic Instrument (DESI) [17] and of high-redshift galaxies by the James Web Space Telescope (JWST) [18] have found discrepancies between observation and the cold DM (CDM) model. These new observations have opened up questions regarding the necessity for extensions or alternatives to the CDM model. For this, I refer the reader to Ref. [19] as a detailed discussion of the failures of the CDM model; its proposed alternatives are beyond the scope of this thesis.

1.2 What is Dark Matter?

The previous section established the existence of DM, which seems to make up most of the mass in the universe despite never having been observed directly. An obvious question is “What is it?” While the field has not yet been able to answer that question, from the evidence and observations presented above, we can deduce some characteristics of DM:

- **DM either does not interact or very weakly interacts via the electro-magnetic force.** Since the electro-magnetic force is mediated via photons, the very fact that DM is non-luminous tells us that the degree to which DM interacts electromagnetically is, at most, very small.
- **DM is relatively collisionless.** Observations of mergers of galaxy clusters, such as the Bullet Cluster, show that DM must not have strong interactions with other types of matter or itself—if it did, it would have separated out of the galaxy clusters along with the hot gas.
- **DM is long-lived and has existed since the early universe.** Anisotropies in the early universe cannot be explained at their observed scales without DM, which suggests that DM must have existed in the universe in times prior to recombination.
- **DM is “cold” or non-relativistic.** We know this from comparisons of between observations of large structures in the universe and numerical simulations. Simulations of structure formation since the Big Bang reveal that hot DM would lead to diffuse structures that are incompatible with observations.

The challenge in observing DM directly lies in the fact that we do not know what type of object or particle we are searching for. Our ignorance here is well evidenced by our ignorance in the mass of a single constituent or particle of DM: Lower bounds on the mass of a single DM particle (if it is a particle), m_{DM} , can be obtained from considering a system of DM particles confined to a galactic halo, as is demonstrated in Section 3.5 of Ref. [5]. Given some assumptions regarding the velocity distribution of DM, which is informed by both by measurements of the CMB and numerical simulations discussed in the previous section, one can find that $m_{\text{DM}} \gtrsim 10^{-22} \text{ eV}$.⁴ For

⁴This is the bound one would set when considering bosonic DM. If one were to consider fermionic DM, a more stringent limit can be set by applying the Pauli exclusion principle yielding a lower bound of $\sim 10 \text{ eV}$. (See Section

upper limits of m_{DM} , we do not have to consider DM as a single particle but instead can consider it to be a macroscopic object such as a dim star, asteroid, or even black hole. Weak lensing data from our own galaxy indicate that the upper bound on the mass of individual objects that may contribute to the DM halo would then have to have masses below $\sim 10^{59}$ eV [20].

From the point of view of an experimentalist, this is a huge parameter space. It is impossible to design a single experiment that is sensitive to the entire spectrum of possible DM masses, and therefore a broad portfolio of experiments, each searching in different mass ranges, is required to fully explore the mass range. Optimally, rather than randomly selecting a mass range to search over, we would prioritize mass regions that overlap with the parameter space of well-motivated DM candidates. Fig. 1.4 shows a number of popular candidates and the various mass ranges they could occupy. For a detailed review of these candidates, I refer the reader to reviews in Refs. [21], [22]. For the purposes of this thesis, we will limit our discussion to models with viable parameter space in the eV to GeV range—in particular, weakly-interacting massive particles (WIMPs), hidden sector DM, axions, and axion-like particles (ALPs).




Figure 1.4: A graphic showing DM candidates and models and corresponding mass ranges in which they have viable parameter space. This is a modified version of a graphic taken from Ref. [21].

3.5 of Ref. [5].) Typically, the bosonic bound is referred to since DM has not been confirmed to be either bosonic or fermionic.

1.2.1 WIMP Dark Matter

The WIMP model is motivated by the fact that DM existed in the early universe and assumes that DM was also in thermal equilibrium with SM particles. In this scenario, we assume DM particles, X , can annihilate to produce a pair of SM particles $XX \rightarrow YY$. The reverse process, $YY \rightarrow XX$, in which a pair of SM particles interact to produce DM particles may also occur and, while DM is in thermal equilibrium with SM and the temperature is greater than the DM mass, both sides of the interaction balance each other out so that the populations remain constant.

If we consider massive DM particles such that, in general, $m_X > m_Y$, then as the universe expands and cools, SM particles will eventually not be sufficiently energetic to create DM through interactions, thus shutting off the $YY \rightarrow XX$ side of the interaction. Similarly, as the universe expands, the number density of DM will decrease and will, therefore, be less likely to interact; this shuts off the $XX \rightarrow YY$ side of the interaction, leaving us with a remaining population of DM. This process is called “freeze-out”.

We can solve for the remaining relic density of DM in the case of freeze-out and set the solution equation equal to DM density extracted from CMB measurements to determine under what conditions such a scenario is possible. I begin with the derivation for the rate of change in the co-moving number density of DM—the number density in volume units that remain constant with the expansion of the universe—which closely follows the derivation shown in Ref. [23]. We can start with the Boltzmann equation, which describes the rate of change in the number density of a particle as the difference in the rates in which its annihilated and created.

$$\begin{aligned}
 a^{-3} \frac{d(na^3)}{dt} = & \frac{1}{2(2\pi)^3} \int \frac{d^3p_{X_1}}{E_{X_1}} \frac{1}{2(2\pi)^3} \int \frac{d^3p_{X_2}}{E_{X_2}} \frac{1}{2(2\pi)^3} \int \frac{d^3p_{Y_1}}{E_{Y_1}} \frac{1}{2(2\pi)^3} \int \frac{d^3p_{Y_2}}{E_{Y_2}} \\
 & \times (2\pi)^4 \delta^3(p_{X_1} + p_{X_2} - p_{Y_1} - p_{Y_2}) \delta(E_{X_1} + E_{X_2} - E_{Y_1} - E_{Y_2}) |\mathcal{M}|^2 \\
 & \times (f_{Y_1} f_{Y_2} [1 \pm f_{X_1}] [1 \pm f_{X_2}] - f_{X_1} f_{X_2} [1 \pm f_{Y_1}] [1 \pm f_{Y_2}])
 \end{aligned} \tag{1.3}$$

Eqn. 1.3 is long enough to warrant some explanation. First, the time derivative on the LHS is of the co-moving number density na^3 where a is a scale factor such that the volume of the universe, V divided by a^3 , is constant as the universe expands. The physical number density of the DM particles, n , falls with the expansion of the universe, which is why there is a factor of a^{-3} multiplying the time derivative. The RHS of Eqn. 1.3 contains the interactions between two DM particles $X_{1,2}$ and two SM particles $Y_{1,2}$. The first term on the last line of the RHS is the "creation term" describing the interaction between two SM particles that produces two DM particles—naturally, since this type of interaction increases the number of DM particles, it is positive. The second term is then the annihilation term and contributes negatively to the number of DM particles. These terms depend on the density of states f for each class of particle. The second line contains delta functions that enforce conservation of energy and momentum as well as a scattering amplitude $\mathcal{M}$, which is determined by the specific interaction between DM and SM particles. Since in the early universe, SM and DM are in thermal equilibrium, we assume that $\mathcal{M}$ is the same regardless of whether we are considering $XX \rightarrow YY$ or $YY \rightarrow XX$. Finally, the top line on the RHS contains the integrals over all of phase space for each particle.

In the early universe, particles interact frequently so that the occupancy of each particle, f , is either given by the Bose-Einstein or Fermi-Dirac distributions depending on whether the particle in question is a boson or fermion.⁵

$$f(E, T) = \frac{1}{e^{-(E-\mu)/T} - 1} \quad \text{for bosons} \quad (1.4a)$$

$$f(E, T) = \frac{1}{e^{-(E-\mu)/T} + 1} \quad \text{for fermions} \quad (1.4b)$$

The chemical potential in Eqn. 1.4 is represented by μ . Part of our initial assumption is that

⁵I will generally use natural units to simplify notation. Thus, $c = k_B = 1$ where c and k_B are the speed of light and the Boltzmann constant, respectively.

DM particles are in chemical equilibrium with SM particles such that $\mu_X = \mu_Y$. For our purposes, we will be considering temperatures much lower than $E - \mu$ so that the ± 1 in Eqn. 1.4 can be ignored and the occupancy can be written as

$$f(E, T) \approx e^{-E/T} e^{-\mu/T} \quad (1.5)$$

We can use Eqn. 1.5 to put Eqn. 1.3 in terms of the number densities. For a general particle in thermal equilibrium, the number density can be calculated as the integral of the occupancy over all phase space:

$$n = \int \frac{g}{2\pi^3} f(p) d^3p \quad (1.6)$$

Here, the factor g accounts for the degeneracy of states. Substituting in the approximation in Eqn. 1.5, we have

$$n = \frac{g}{2\pi^3} e^{-\mu_X/T} \int e^{-E/T} d^3p \quad (1.7)$$

While DM and SM particles are in equilibrium, we assume that it takes no additional energy to create or annihilate DM from SM particles and thus we can consider n_{eq} where $\mu_X \sim 0$ (or, equivalently, $\mu_Y \sim 0$). If the temperature at which this is true is high so that $m_X \ll T$, then the energy of DM particles will be related to the momentum such that $E \sim p$ (in natural units). If the temperature is lower, so that we have $m_X \gg T$, then $E \sim m$. Integrating Eqn.1.7 in each limit will give

$$n_{eq} = \begin{cases} g \frac{T^3}{\pi^2} & m_X \ll T \\ g \left(\frac{m_X T^{3/2}}{2\pi} \right) e^{-m_X/T} & m_X \gg T \end{cases} \quad (1.8)$$

where n_{eq} is the number density while SM and DM particles are in equilibrium. In either limit shown in Eqn. 1.8, we can represent the chemical potential as the ratio of the number density to

the number density while in equilibrium:

$$e^{-\mu_X/T} = n/n_{\text{eq}} \quad (1.9)$$

Before substituting the number densities into the Boltzmann equation, we can make one more simplification by defining the thermally-averaged cross-section.

$$\begin{aligned} \langle \sigma_A v \rangle = & \frac{1}{n_{\text{eq}}^2} \frac{1}{2(2\pi)^3} \int \frac{d^3 p_{X_1}}{E_{X_1}} \frac{1}{2(2\pi)^3} \int \frac{d^3 p_{X_2}}{E_{X_2}} \frac{1}{2(2\pi)^3} \int \frac{d^3 p_{Y_1}}{E_{Y_1}} \frac{1}{2(2\pi)^3} \int \frac{d^3 p_{Y_2}}{E_{Y_2}} e^{-(E_{X_1} + E_{X_2})/T} \\ & \times (2\pi)^4 \delta^3(p_{X_1} + p_{X_2} - p_{Y_1} - p_{Y_2}) \delta(E_{X_1} + E_{X_2} - E_{Y_1} - E_{Y_2}) |\mathcal{M}|^2 \end{aligned} \quad (1.10)$$

Here, σ_A is the annihilation cross-section, and v is the velocity of a DM particle. In general, σ_A depends on v hence why we average the product of the two rather than just the cross-section. Re-writing the remaining terms on the RHS of Eqn. 1.3 and expanding out the derivative on the LHS, the Boltzmann equation becomes

$$\frac{dn}{dt} + 3Hn = \langle \sigma_A v \rangle [n_{\text{eq}}^2 - n^2] \quad (1.11)$$

where I have used the definition of the Hubble parameter $H = \frac{\dot{a}}{a}$ when expanding the LHS of Eqn. 1.3. We can use Eqn. 1.11 to solve for the remaining number density of the dark matter when the universe has reached the temperature at which freeze-out occurs, T_f . This happens when the annihilation rate of DM, given by $\Gamma(T) = n(T)\langle \sigma_A v \rangle$, is equal to the Hubble expansion $H(T)$.

Solving Eqn. 1.11 for the DM number density at the time of freeze-out yields a solution that depends on the thermally averaged cross-section as shown in Fig. 1.5. The cross-section that gives the relic density of DM consistent with measurements just so happens to be consistent with the scale of weak interactions. This was seen as significant because around the same time this




Figure 1.5: The co-moving number density of DM as a function of time (here, parameterized by the mass of a DM particle over temperature). The solid line shows the number density if DM remained in thermal equilibrium with the SM particles, while the dashed lines show the remaining relic density at the time of freeze-out, which depends on the thermally-averaged cross-section. Figure taken from Ref. [24].

discovery was made, Supersymmetric theories were proposed as possible solutions to the hierarchy problem and predicted the existence of weakly-interacting particles in the $>\text{GeV}$ mass range, and were seen as excellent DM candidates. This coincidence, referred to as the “WIMP-miracle,” spurred decades of searches for DM particles in the form of WIMPs.

Among the particles that were predicted by Supersymmetry (SUSY), the most motivated DM candidates were perhaps the Supersymmetric partners of neutrinos, called “neutralinos,” due to their being long-lived and electrically-neutral. However, a number of theorists independently found that for masses below $\sim 2 \text{ GeV}$ (referred to as the Lee-Weinberg bound), WIMPs in the form of

heavy neutrino states would not annihilate fast enough to explain the abundance of DM [25], [26]. Decades of searches for WIMPs turned up null results and placed increasingly stringent bounds on the available parameter space; in particular, direct-detection experiments utilizing liquid scintillators as targets have led sensitivities for masses down to a few GeV. Though WIMPs have not yet been entirely ruled out, attention has slowly drifted to DM candidates at lower masses, whose parameter space is comparatively unexplored.

1.2.2 Hidden Sector Dark Matter


Hidden sector models can be seen as a generalization to the WIMP model in that DM is now allowed to interact via “hidden” forces outside of the SM. As a result, production methods other than the freeze-out mechanism can be used to explain the current abundance. With these constraints on the WIMP model removed, hidden sector models evade many bounds on the viable parameter space. Notably, many of these models predict DM candidates with masses down to the eV-keV scale, where the parameter space is still relatively unexplored.

While one may consider hidden sectors such that DM has no predicted interactions with SM particles, symmetries in SM motivate certain “portal” interactions that weakly couple the “visible” sector of the SM to these hidden sectors. In general, there are a number of portals with viable parameter space, most of which are beyond the scope of this thesis. Rather, I will focus on the vector portal, which introduces a new vector force mediator and thus leads to a mixing term in the total Lagrangian given by


$$\mathcal{L} \ni \frac{\epsilon}{2} F_Y^{\mu\nu} F'_{\mu\nu} \quad (1.12)$$

where ϵ is a kinetic mixing parameter determining the strength of the portal interaction, $F_Y^{\mu\nu}$ is the hypercharge field strength, and $F'^{\mu\nu}$ is the field strength of the hidden sector gauge field. The gauge boson introduced in the hidden sector A' can then mediate interactions between a particle

charged under the hidden sector's $U'(1)$ symmetry (which I refer to as χ) and SM particles.



(a) DM-electron scattering.



(b) DM absorption.

Figure 1.6: Feynman diagrams showing DM-electron scattering mediated by an introduced vector or scalar (left) and absorption of DM by an electron (right). Here χ is the particle charged under the introduced $U'(1)$ symmetry and A' , ϕ , and a , are possible mediators that, respectively, take the form of a vector, scalar, or pseudoscalar, depending on the model under consideration. Figures taken from Ref. [21]

From this term, one can predict a number of possible interactions between DM candidates in nuclei and electrons. However, interactions with nuclei are only useful when searching for DM models with higher masses, where experiments searching for WIMPs have already explored most of the viable parameter space. This is due to the fact that as you consider DM candidates with lower masses, the amount of energy imparted to nuclei during an elastic collision with DM shrinks below the thresholds of detectors used in WIMP searches. We can show this simply by considering the kinematics of a two body collision. The maximum energy transferred to a nucleus during an

elastic collision can be written as⁶

$$\Delta E_{\max} = \frac{2\mu^2 v_\chi^2}{m_N} \quad (\text{elastic collisions}) \quad (1.13)$$

where, in this context, μ is the reduced mass of the system (given by $\mu = \frac{m_\chi m_N}{m_\chi + m_N}$ where m_χ is the mass of the DM candidate), m_N is the mass of the nucleus, and v_χ is the initial velocity of the DM candidate. For a small m_χ , Eqn. 1.13 can be approximated as

$$\Delta E_{\max} \approx \frac{4m_\chi}{m_N} \left(\frac{1}{2} m_\chi v_\chi^2 \right) \quad (\text{elastic collisions}) \quad (1.14)$$

where the term in the parentheses is the kinetic energy of the DM before the collision, which is suppressed by a factor proportional to the ratio of the DM candidate mass with the mass of the nuclei. Since nuclei have masses $\gg 1$ GeV, the suppression factor $\frac{4m_\chi}{m_N} \ll 1$.

However, one can instead consider scattering with electrons, in which case, for elastic scatters, the mass of the nucleus in the discussion above would instead be replaced by the mass of an electron and thus greatly increase the suppression factor. For atoms, which have internal excited states, we may also consider inelastic scattering with bound electrons (Fig. 1.6a) or even total absorption (Fig. 1.6b), in which case the maximum energy transferred can be written as

$$\Delta E_{\max} = \frac{1}{2} m_\chi v_\chi^2 \quad (\text{inelastic collisions}) \quad (1.15)$$

To get a sense of the relevant energy scale in Eqn. 1.15, consider a DM candidate with a mass of 5 MeV. For cold DM, typically velocities are dominated by our solar system's motion through the galactic halo and thus v_χ can be taken to be ~ 220 km/s. Using this, we see that, for

⁶Here, I consider the non-relativistic case since both CMB measurements and numerical simulations of structure formation in the universe imply that DM particles have velocities much smaller than the speed of light.

inelastic collisions, we obtain a maximum energy transfer of $\sim 2\text{--}3$ eV. While this is an energy that is inaccessible to detectors that have lead the charge in WIMP searches—primarily liquid scintillator time projection chambers—even these small transfers are enough to cause ionization in common semiconductor materials such as silicon and germanium. This thesis will show results from one such detector architecture which takes advantage of this fact.

1.3 This Work: Continuing the Search for Dark Matter at Low Masses

As experiments close in on the remaining WIMP parameter space, attention has gradually shifted to lower mass candidates where the parameter space remains relatively unexplored. Searching in these low-mass regions comes with its own set of challenges, not the least of which being the need for detectors sensitive to eV-scale energy depositions. This work presents the latest results of a DM search using SuperCDMS HVeV detectors—detectors which exploit eV-scale bandgaps of semiconductors to achieve sensitivity to low mass DM, which I will introduce in Chapter 2. The design and installation of the detector payload is described in Chapter 3. Chapters 4 through 6 show the results from the data analysis including evidence of a previously unobserved mechanism for energy loss in solid-state detectors, observations of low energy backgrounds which are currently limiting the field, and world-leading exclusion limits on electron-recoil dark matter.

CHAPTER 2

SUPERCDMS HVEV DETECTORS

To meet the energy thresholds required to probe low-mass DM, the Super Cryogenic Dark Matter Search (SuperCDMS) collaboration uses cryogenic, phonon-mediated calorimeters. These detectors consist of a semiconductor, crystalline absorber (silicon or germanium) equipped with thin superconducting films (called QETs) that, when operated in their superconducting transition, act as precision thermometers. In the general measurement scheme, shown in Fig. 2.1, the crystal substrate acts as a target for DM particles. A particle interaction with atoms in the crystal lattice will produce phonons that then propagate throughout the crystal until they eventually deposit their energy into the superconducting films, causing their temperature to increase. I will show that this




Figure 2.1: A schematic of a SuperCDMS HV detector. A crystal absorber (either silicon or germanium) is shown in light gray. A DM particle, labeled χ , could then interact with atoms in the crystal lattice and produce an electron-hole pair (charges are labeled e^- and h^+ , respectively) and phonons (shown by the blue arrows). As the charges drift due to the applied voltage bias, labeled ΔV , they produce more phonons via the Neganov-Trofimov-Luke effect. Figure taken from Dr. Noah Kurinsky's thesis [27].

increase in temperature leads to a change in the films' resistance, which can be readout by a circuit

that I will introduce later in the chapter. In addition, we will see that the use of a semiconductor crystal absorber enables the amplification of ionizing signals, thus increasing sensitivity to hidden sector DM candidates.

The SuperCDMS collaboration generally employs two detector architectures: 1) an HV detector, to which Figure 2.1 applies, and which measures the phonon energy deposited in the crystal with high precision, and 2) an iZIP detector which uses a similar scheme but also measures the ionization yield in order to distinguish nuclear recoils (NRs) from electron recoils (ERs) [28]. These detectors are currently being utilized at an underground facility called SNOLAB and are expected to set world-leading exclusion limits for low-mass WIMPs [29]. However, the detectors that are the subject of this thesis are from an R&D program in SuperCDMS and referred to as HVeV detectors. These are smaller versions of the HV detector (HV detectors are ~ 600 g whereas HVeVs are only 0.93 g), and were initially conceived as a more deployable detector that could be used for supporting studies such as detector characterization or studies of charge-transport, but have since become a competitive technology for low-mass DM searches on their own.

In this chapter I will give an overview of the SuperCDMS HVeV program, beginning with a description of measurement scheme and readout circuit and ending with a summary of recent data-taking campaigns using HVeVs that are relevant to contextualize results shown in later chapters.

2.1 Measurement Scheme

2.1.1 Phonon Sensors

Phonons from particle interactions in the crystal substrate are sensed by parallel arrays of QETs [30], which stands for “quasiparticle-trap-assisted transition-edge sensors”—an acronym long enough to comprise a thesis on its own. A QET consists of a thin superconducting film biased in its superconducting transition—called a transition-edge sensor (TES)—surrounded by aluminum (Al) films, which serve to “funnel” energy, transferred in the form of quasiparticles,

towards the TES.




Figure 2.2: A schematic for sensing phonons with a QET. Phonons from interactions in the substrate (shown in blue) travel to the superconducting Al fin (shown in grey). Phonons which are transmitted through the interface interact with Cooper pairs, producing quasiparticles. Quasiparticles then diffuse through the Al fin into the TES where they convert most of their remaining above-gap energy into phonons, heating the TES. Figure taken from Ref. [31]

As shown in Figure 2.2, athermal phonons that enter the Al can interact with Cooper pairs, producing quasiparticles in the substrate. The superconducting bandgap energy for W is less than that of Al, thus as quasiparticles diffuse through the Al into the region where the W and Al overlap, they release energy in the form of phonons, becoming “trapped” as they no longer have the energy required to cross into the bulk Al. The released phonon energy then transmits into the TES, increasing its temperature. At cryogenic temperatures ($\lesssim 100$ mK), the Al bandgap (~ 350 μeV) is sufficiently high such that thermal phonons in the substrate do not have the energy required to break Cooper pairs [31]. In this way, the Al fins serve to increase the energy collection area and thermally decouple the TES from the substrate.

The TES is biased in its superconducting transition (see Figure 2.3) so that a small change in temperature causes a measurable change in resistance. The slope and width of the transition region determine the sensitivity and dynamic range of the TES. The response of the TES can be described by a set of coupled differential equations; one “electrical” equation that describes the TES bias circuit shown in Figure 2.4 and one “thermal” equation that describes the heat flow through the




Figure 2.3: An example of the resistance as a function of temperature for a transition-edge sensor (in this case, a Mo/Cu bilayer). Through application of a current or voltage bias, a TES can be held in the transition region where the superconducting film's resistance changes rapidly with temperature. Figure taken from Ref. [32].

TES.

A quick examination of the simplified bias circuit shown in Figure 2.5b allows us to write the electrical differential equation as

$$L \frac{dI}{dt} = V_{\text{bias}} - IR_L - IR_{\text{TES}}(I, T) \quad (2.1)$$

where R_L is the load resistance and is equal to the sum of the shunt resistance and any parasitic resistance, R_p in the bias circuit. In principle, there can be some parasitic inductance as well in which case we would replace L with the equivalent inductance of the bias circuit.

The thermal differential equation comes from an examination of the heat flow through the TES.




Figure 2.4: (right) A TES bias circuit. In our case, the TES is biased using a current source and a shunt resistor (R_{Sh}) that is connected in parallel with the TES (R_{TES}) and some inductance L which serves as the input for a SQUID-based readout circuit that I will introduce in the next section. (left) The equivalent, thevenized TES bias circuit. Here R_L , called the “load resistance,” represents the sum of the shunt resistance and any parasitic resistance in the circuit. Figures created using www.circuit-diagram.org.

Here, we consider the TES as weakly thermally coupled to some cold bath with temperature $T_b \ll T_c$ (where T_c is the critical temperature of the TES). The time rate of change of the temperature of the TES can then be determined by power flowing into the TES (from the signal, P_s and joule heating from the bias, P_J) subtracted by the power dissipating through the weak thermal link to the bath, P_b , as shown in Eqn. 2.2.

$$C \frac{dT}{dt} = P_s + P_J - P_b \quad (2.2)$$

As shown in Ref. [32], equations 2.1 and 2.2 can be simplified under within the “small signal limit,” wherein we can linearize the equations by expanding around a small change in current and temperature:

$$T \approx T_0 + \delta T \quad (2.3)$$

$$I \approx I_0 + \delta I \quad (2.4)$$

In Eqns. 2.3 and 2.4, the subscript 0 denotes the steady-state values ($P_s = 0$) for the TES. In the small-signal limit, we assume the terms of higher order in δT or δI are negligible. We can use this to expand out the resistance of the TES as well which we write as

$$R_{\text{TES}}(I, T) \approx R_0 + \alpha \frac{R_0}{T_0} \delta T + \beta \frac{R_0}{I_0} \delta I \quad (2.5)$$

where α and β are, respectively, the temperature and current sensitivity of the TES and are defined in Eqns. 2.6 and 2.7.

$$\alpha \equiv \left. \frac{\partial \log(R)}{\partial \log(T)} \right|_{I_0} \quad (2.6)$$

$$\beta \equiv \left. \frac{\partial \log(R)}{\partial \log(I)} \right|_{T_0} \quad (2.7)$$

α is typically given as a good figure of merit of a TES; it corresponds to the slope of the transition region shown in Figure 2.3 and can therefore be understood as a measure of the TESs' sensitivity. We will later see that β would optimally be zero—a higher value for β indicates a larger susceptibility to noise which can negatively impact the device's energy resolution. Using these terms, we can similarly expand out the power from joule-heating in terms of α and β in the small signal limit as

$$P_J \approx P_{J,0} + 2I_0 R_0 \delta I + \alpha \frac{P_{J,0}}{T_0} \delta T + \beta \frac{P_{J,0}}{I_0} \delta I \quad (2.8)$$

where $P_{J,0} = I_0^2 R_0$. Finally, we can write the power dissipating to the bath as

$$P_b \approx P_{b,0} + G \delta T \quad (2.9)$$

where $G \equiv dP_b/dT$.

Substituting Eqns. 2.3–2.9 into the differential Equations 2.1 and 2.2, canceling higher orders of terms with δ , and observing that in a steady state we have $P_{b,0} = P_{J,0}$, we arrive at the following

$$\frac{d\delta I}{dt} = \frac{\delta V}{L} - \frac{R_L + R_0(1 + \beta)}{L}\delta I - \frac{\mathcal{L}G}{I_0 L}\delta T \quad (2.10)$$

$$\frac{d\delta T}{dt} = \frac{\delta P}{C} + \frac{I_0 R_0(1 + \beta)}{C}\delta I - \frac{1 - \mathcal{L}}{\tau}\delta T \quad (2.11)$$

where we have introduced the loop gain $\mathcal{L} = P_{J,0}\alpha/GT_0$ and a time constant $\tau = C/G$ (which is just the thermal decay time constant for the TES in the absence of any signal). Additionally, we have used $\delta V = V - V_{\text{bias}}$ as a small change in the bias voltage resulting from the signal¹ and substituted P for δP both to emphasize that we are in the small signal limit and to be consistent with the formalism in Ref. [32].

There are multiple ways of solving Eqns. 2.10 and 2.11, which are discussed in Ref. [32] though specific solution will depend on the form of the signal power δP . First, by considering the homogeneous solution (where $\delta V = 0$ and $\delta P = 0$), we find that a general solution can be written in terms of exponentials:

$$\delta I(t) = c_1 e^{-t/\tau_+} + c_2 e^{-t/\tau_-} \quad (2.12)$$

$$\delta T(t) = b_1 e^{-t/\tau_+} + b_2 e^{-t/\tau_-} \quad (2.13)$$

In Eqns. 2.12 and 2.13, the constants $c_{1,2}$ and $b_{1,2}$ can be solved for given a set of initial

¹Note that the bias voltage of the circuit shown in Fig. 2.5b can be derived from thevenizing the real bias circuit shown in Fig. 2.5a and does not depend on the resistance of the TES. If it did, we would not be able to assume the change in voltage is small under the small signal limit as V_{bias} would depend on α , which is generally very large for TESs.

conditions. The time constants $\tau_{\pm}$ are the inverse eigenvalues that are obtained via solving the coupled differential Equations 2.10 and 2.11. Their explicit dependence on circuit parameters is shown in Eqn. 2.14.

$$\frac{1}{\tau_{\pm}} = \frac{1}{2\tau_{\text{el}}} + \frac{(1 - \mathcal{L})}{2\tau} \pm \frac{1}{2} \sqrt{\left(\frac{1}{\tau_{\text{el}}} - \frac{1 - \mathcal{L}}{\tau}\right)^2 - 4\frac{R_0\mathcal{L}(2 + \beta)}{L\tau}} \quad (2.14)$$

Here we have defined a the electrical time constant of the RL circuit in the absence of any input power $\tau_{\text{el}} = L/(R_L + R_0(1 + \beta))$. Note that the time constants τ can, in general, be imaginary and thus, the qualitative behavior of the TES response depends on the circuit elements used. Of particular importance is whether or not the real part of τ is positive or negative; if $\text{Re}(\tau) < 0$, then the current through the TES will continually increase with time and thus become unstable. For a more detailed discussion of stable operation of the TES, see the discussion in Section 2.4 of Ref. [32].

To understand the response of the TES resulting from a particle interaction in the substrate, we can consider the case where δP takes the form of a delta function—this would mimic an idealized case where energy is instantaneously deposited in the crystal. In this scenario, we have that, at time $t = 0$, $\delta I = 0$ and $\delta T = \Delta T$. In the limit that the heat capacity of the absorber is insensitive to changes in temperature, we may approximate $\Delta T \simeq E/C$ [33], where E . Using these initial conditions, we can solve for the coefficients in the homogeneous solutions given in Eqns. 2.12 and 2.13. However, since we will only directly measure the current going through the TES, it is sufficient to only solve for the constants in Eqn. 2.12 in which case we arrive at a current response of the form given in Eqn. 2.15.

$$\delta I(t) = A \frac{E}{C} (e^{-t/\tau_+} - e^{-t/\tau_-}) \quad (2.15)$$




Figure 2.5: (right) The top side of an HVeV detector with the inner and outer channel denoted by the outside or inside of the red box, respectively. (left) A zoomed in view of a single QET. The W TES is the center of the Al fins that surround it.

We can see that under these conditions, δI takes the form of a “pulse” with an exponential rise time given by τ_+ and a fall time given by τ_- . The amplitude of the pulse is the product of a constant A , which depends on the circuit elements and time constants and is given explicitly in Ref [32], and the ratio of the energy deposited to the heat capacity of the TES. We may therefore estimate the energy deposited in the absorber during a particle interaction via the height of the observed pulse in the current of the TES.

Note that in Eqn. 2.15, the height of the current pulse is inversely proportional to the heat capacity of the film. From a detector point of view this is significant—to collect more phonons from an interaction in the substrate, one would naively want a large TES. However, by making the TES large, we also increase the heat capacity C , thereby decreasing the amplitude of the current response. This is precisely the genius behind the QET; the Al fins in the QET effectively increase the collection area of the TES without increasing the TESs heat capacity. However, the diffusion length of quasiparticles through superconducting films is finite [34] and, therefore the Al

fins cannot be scaled indefinitely without introducing mechanisms for energy loss.² This can be worked around by utilizing parallel arrays of QETs, thus avoiding excessively large Al fins while still maintaining a high coverage fraction. HVeV devices have two such parallel arrays, which cover an equal amount of area, of QETs arranged in a “bullseye” pattern as shown in Fig. 2.5.

2.1.2 SQUID-Based Readout

To measure the current through the TES in the circuit shown in Fig. 2.4, we use a superconducting quantum interference device (SQUID). There are two types of SQUIDs: an RF (radio frequency) SQUID consisting of a superconducting loop with a single Josephson junction and a DC (direct current) SQUID consisting of a superconducting loop with two Josephson junctions as pictured in Fig. 2.6. Our discussion will focus on the latter, which we use in our readout system. Here, I will present a mostly qualitative explanation of the SQUID’s response and implementation into the readout system and refer the reader to Refs. [35], [36] for an in-depth discussion.




Figure 2.6: A simplified circuit schematic of a DC SQUID, which consists of a loop of superconducting wire containing Josephson junctions (represented by the $\times$ symbols in the diagram).

To understand how the SQUID responds, first let us understand the behavior of the Josephson junctions. Josephson junctions are comprised of two superconducting materials separated by a thin

²This occurs when the Al fins area is large enough such that quasiparticles are likely to recombine in the Al before becoming trapped in the Al-W overlap region, thus preventing any of the energy required to excite them in the first instance from reaching the TES.

insulator. The Cooper pairs in each superconductor form one coherent quantum state with some phase and can tunnel through the potential barrier imposed by the insulator, forming a “supercurrent” of Cooper pairs through the junction. As shown in Ref. [37], the supercurrent, I , is related to the phase difference between the superconductors, δ by

$$I = A \sin(\delta) \quad (2.16)$$

In Eqn. 2.16, A is a constant that depends on the volume of the superconductors, the charge of the Cooper pairs ($2e^-$), and the strength of the coupling between the superconductors. Just as with bulk superconductors, Josephson junctions have a critical current, I_0 , above which there is a non-zero resistance. Thus, for $I > I_0$, the current flows with resistance and a voltage drop can be measured over the junction, which increases with increasing I (see Fig. 2.7). In general, the critical




Figure 2.7: The typical IV behavior of a Josephson Junction. When the current through the junction is less than the critical current, the current flows with zero resistance and thus there is no voltage drop. Above the critical current, current no longer flows without resistance and a voltage drop appears over the junction. For $I \gg I_c$, the voltage drop is given simply by Ohm's Law. Figure taken from Ref. [38].

current for Josephson junctions is much smaller than that of a typical superconductor. Thus, if we assume that the SQUID shown in Fig. 2.6 is symmetric (I_0 is the same for both junctions), then the voltage drop over the SQUID will be zero so long as $I < 2I_0$, where I is the total current through the SQUID. Hereafter, I will refer to this critical current (the critical current in the absence of any applied flux to the SQUID, which we will consider shortly) of the SQUID as a whole as I_c .

SQUIDs can be used to very accurately measure the magnetic flux through the superconducting loop shown in Fig. 2.6 due to the quantization of flux through superconducting loops—the flux through the loop must be an integer multiple of $\Phi_0 = \frac{h}{2e}$ else the requirement that the wave function describing the Cooper pairs in the loop is not single-valued would be violated. Here I will consider a simplified model of the SQUID's response to flux.³ Let us consider the DC SQUID in a magnetic field so that there is some flux Φ passing through the superconducting loop, where $\Phi = n\Phi_0$. Now say the flux through the loop increases slightly. Since the flux is restricted to being an integer multiple of Φ_0 , a screening current I_s will be induced in the loop, generating an opposing magnetic flux so that the net flux through the loop remains equal to Φ . Then, if the bias current running through the SQUID is I , the current through the junctions as a result of the increasing flux becomes

$$i_1 = \frac{I}{2} + I_s \quad (2.17)$$

$$i_2 = \frac{I}{2} - I_s \quad (2.18)$$

where $i_{1,2}$ denotes the current through the individual junctions, indexed by 1 or 2. We can see that the addition of the screening current effectively lowers the critical current of the SQUID—applying a bias current of $I = I_c - I_s$ will cause $i_1 = I_0$ and, thus a voltage drop will be measured over the SQUID. If we continue to increase the applied flux through the loop, the screening current will

³In particular, I will ignore noise sources as well as neglect the response of the Josephson junctions to changing flux in the loop. Even while ignoring these effects, the qualitative behavior of the SQUID is much the same.

continue to increase to oppose the change until the applied flux is equal to $\Phi + \frac{1}{2}\Phi_0$, at which point the resulting screening current reaches a maximum. If the applied flux is increased further, then something interesting happens: the screening current flips direction so that its induced magnetic flux *adds* to the applied flux so that the next flux through the loop becomes $\Phi + \Phi_0$. This is because it now takes less energy for the screening current to increase the next flux to the next allowed state, $\Phi_{\text{net}} = \Phi + \Phi_0$, than it does to continue to oppose the increase in flux and keep $\Phi_{\text{net}} = \Phi$. As the applied flux is increased to $\Phi + \Phi_0$, the screening current will increase back to 0 as shown in Fig. 2.8.




Figure 2.8: A qualitative plot of the screening current generated in the DC SQUID (plotted on the vertical axis) as a function of the increase in applied flux shown on the horizontal axis in units of Φ_0 . The screening current reaches its maximum when the increase in the applied flux is a half-integer multiple of Φ_0 and is zero when the increase in applied flux is an integer multiple of Φ_0 . Figure taken from Ref. [39].

Assuming that the SQUID is held in its resistive state (*i.e.* $I > I_c$), then the changing magnetic flux manifests as a change in the voltage drop over the SQUID. Like the screening current, the voltage is periodic with the applied magnetic flux; the voltage reaches a maximum for $\Phi = (n + 1/2)\Phi_0$ and a minimum when $\Phi = n\Phi_0$ where $n = 0, \pm 1, \pm 2, \dots$. Figure 2.9 shows the voltage response of the SQUID as a function of the flux for a particular bias current (denoted I_B in Fig. 2.9 and is equivalent to I in the discussion above). Therefore, by measuring the voltage output of the SQUID, we can measure the flux going through the SQUID loop. However, as can




Figure 2.9: (left) The IV curves of the SQUID when $\Phi = n\Phi_0$ and when $\Phi = (n + 1/2)\Phi_0$. For a particular bias current, I_B , the voltage varies continuously along the dashed line as the flux is increased from $n\Phi_0$ to $(n + 1/2)\Phi_0$. (right) The voltage vs. applied flux (here plotted in units of Φ_0) for the bias current shown in the left panel.

be seen in the right panel of Fig. 2.9, the voltage drop over the SQUID quickly becomes non-linear with increasing flux. To be sensitive to larger changes in flux, we can linearize the response of the SQUID using a flux-locked loop (FLL), which is shown in Fig. 2.10. The purpose of the FLL is to “lock” the SQUID at a particular point in its V - Φ response (typically we choose the steepest part of the curve for maximum sensitivity). The FLL is a negative-feedback loop wherein the voltage output of the SQUID (V_{sq} in Fig. 2.10) is amplified, converted to a current via a feedback resistor (R_{fb}), and fed through a feedback inductor (L_{fb}), which supplies a magnetic flux, Φ_{fb} , to cancel out the original change in flux. In this way, the net flux through the SQUID remains constant so that the response is linear and the original increase in flux can be read out via the voltage drop over the feedback resistor.

Having shown that SQUIDS are incredibly sensitive flux-to-voltage transducers, let’s now consider how they are implemented to read out signals from the TES. Recall from the previous section that particle interactions in our detectors’ substrate cause a changing current through the detector.




Figure 2.10: A simplified circuit diagram of a flux-locked loop (FLL). The SQUID is biased via current supplied through wire labeled “bias”. The voltage response of the SQUID (V_{sq}) resulting from an applied magnetic flux (Φ_{ext}) is amplified (where G denotes the gain from the amplifier), converted to a current via a large feedback resistor, R_{fb} , and passed through a feedback inductor that supplies a flux, Φ_{fb} , to cancel out Φ_{ext} . Figure taken from Ref. [40].

The change in current results in a change in the flux through the inductor in the bias circuit shown in Figure 2.4. We then couple this inductor to the SQUID so that the flux through the inductor is the Φ_{ext} in Fig. 2.10, which is then read out through the FLL circuit described in this section. The output of the FLL is then digitized and read out at some sampling frequency (in our case, a frequency of 625 kHz, which we then down sample by a factor of 4). The final result is a time series of current values that are ready to be analyzed!

2.2 Detector Response to Particle Interactions

With the method for sensing and measuring energy from particle interactions established, we may now consider what energy we might expect from an interaction. In the majority of cases, a particle interacts with an electron or nucleus in the crystal substrate, which then recoils. The recoiling electron or nucleus can scatter with phonons or produce electron-hole (eh) pairs by exciting electrons to the conduction band. If the detector is operated without a voltage bias (*i.e.* $V = 0$), then any excited electrons first release most of their kinetic energy from scattering with phonons before

recombining with holes in the valence band—these phonons are referred to as “prompt phonons”. I refer to their summed energy as the “prompt energy,” E_p . The energy lost by the freed charges during recombination, in the case of silicon⁴, is also released as phonons and can be written as $E_g n_{eh}$, where E_g is the ionization bandgap and n_{eh} is the number of eh pairs produced. In general, we assume that freed charges recombine on time scales on the order of the prompt phonon arrival time to the TES and thus, at the end of the cascade prompted by the initial interaction (which can involve various recoils, phonon scattering events, and recombination), the total energy deposited into the substrate as phonons is given by Eqn. 2.19.

$$E = E_R = E_p + E_g n_{eh} \quad (V = 0) \quad (2.19)$$

However, as pointed out in Dr. Runze Ren’s thesis (Ref. [41]), it is possible that freed charges recombine on comparatively long time scales, in which case the energy of an event could be as low as the prompt energy produced during the initial recoil, $E \geq E_p$. While we typically assume Eqn. 2.19, we will revisit this assumption in Chapter 4.

This picture remains largely unchanged when we consider a non-zero voltage bias, V , applied across the substrate. The energy from the initial recoil is still distributed between creating phonons or producing electron-hole pairs and the freed charges still release most of their initial kinetic energy via phonon scattering, but now, rather than proceeding to recombine, the applied electric field accelerates the charges across the substrate. As the charges drift through the substrate, they continuously interact with the lattice, releasing their kinetic energy in form of phonons—this is

⁴In general, the energy emitted during recombination can be in the form of a photon or phonons, depending on the nature of the bandgap of the material. Specifically, if the crystal momentum corresponding to the maximum energy of the valence band corresponds to the minimum energy of the conduction band (materials with this property are referred to as “direct bandgap” semiconductors), the energy will most often be emitted in the form of a photon. Otherwise, in order for crystal momentum to be conserved, phonons must mediate the recombination, which is the case for silicon and germanium.

referred to as the Neganov-Trofimov-Luke (NTL) effect [42], [43]. From this, the energy released into the substrate in the form of phonons includes an additional term as compared to Eqn. 2.19:

$$E = E_R + E_{\text{NTL}} = E_p + n_{\text{eh}}(E_g + eV) \quad (2.20)$$

From Eqn. 2.20, we can see that while operating with an HV bias, the energy (in electron-volts) from ionizing signals is amplified by the product of the number of eh pairs produced and the voltage bias. Note that for silicon, $E_g \approx 1.17$ eV at cryogenic temperatures [44] and thus, it is possible to produce a single eh pair with as little as ~ 1.17 eV. Without a voltage bias, this is quite small and difficult to measure. However, we may apply voltage biases exceeding ~ 100 V, and therefore, by applying such a voltage bias, the energy from such an interaction could deposit energies exceeding ~ 101 eV in the substrate! It is precisely this method of amplification that allows SuperCDMS HVeV detectors to be sensitive to the small energy deposits from low mass DM discussed at the end of Chapter 1.

2.2.1 Ionization in Silicon

The amount of amplification one gets from the voltage bias depends on n_{eh} , which in general depends on the energy and type of the interaction. The average number of eh pairs excited in a general interaction can be written as

$$\langle n_{\text{eh}} \rangle = Y \frac{E_R}{\epsilon_{\text{eh}}} \quad (2.21)$$

where E_R is the recoil energy of an interaction, ϵ_{eh} is the average energy required to excite a single eh pair, and Y is a factor sometimes referred to as the “ionization yield” and has a value that depends on whether the interaction is an electron-recoil, (in which $Y = 1$) or a nuclear-recoil, in which (Y depends on E_R and is generally less than 1). Here, ϵ_{eh} can be interpreted as the average amount of energy from an electron-recoil it takes to produce one electron-hole pair and depends

on E_R . For $E_R < E_g$, energy conservation forces ϵ_{eh} to approach ∞ since otherwise electrons would be excited to the conduction band, despite not receiving the full bandgap energy. Similarly, for $E_g \leq E_R < 2E_g$, only a single eh pair can be produced and thus $\epsilon_{\text{eh}} = E_R$. At high energies, $E_R \gg 2E_g$, ϵ_{eh} has been found to be constant with a value between ~ 3.6 – 3.8 eV, depending on the temperature. The behavior of ϵ_{eh} is more complex in the energy region where $E_R \gtrsim 2E_g$. Throughout this thesis, we will reference the phenomenological model constructed in Ref. [44], from which Figure 2.11 (showing the eh pair production probabilities) is taken.




Figure 2.11: The probability to ionize the labeled number of eh pairs in silicon as a function of energy deposited from an electron-recoil. Figure taken from Ref. [44].

For nuclear recoils, the number of eh pairs produced is typically less than that of an equal energy electron-recoil; hence $Y < 1$ for nuclear recoils. For energies above ~ 4 keV the measurements of the nuclear ionization yield have generally been in agreement with a model accepted by Lindhard et. al. [45]. However, below this measurements vastly differ across experiments for nuclear ionization yield [46]. Figure 2.12 shows measurements for the ionization yield in silicon

at energies below ~ 20 keV and is taken from Ref. [46]. The DM models explored in this thesis




Figure 2.12: Measurements of the nuclear ionization yield in silicon. In red are the data points from Ref. [46] and are fit to a power law. Also included in the figure are data from Refs. [47]–[52]. The dashed line shows the Lindhard model with $k = 0.146$ [45]. Figure taken from Ref. [46].

interact via electron-recoil and so I will not discuss nuclear ionization yield further. However, I refer the reader to the thesis of Dr. Runze Ren (Ref. [41]) for recent measurements of the nuclear recoil ionization yield at low energies in silicon.

2.2.2 Charge-Trapping and Impact Ionization

Equation 2.20 gives the energy deposited in the substrate in the idealized case, where charges freed during a particle interaction are drifted via the applied electric field and recombine near the surface. In this case (see the left side of Figure 2.13), the amount of energy released into the substrate via the NTL effect is quantized in integer multiples of the voltage bias (in units of eV). However, defects in the substrate can cause instances in which a non-quantized amount of NTL energy enters the substrate. We consider two distinct cases of this:

1. **Charge-trapping (CT)**: a charge falls into a defect state while drifting across-the crystal, and thus only a fraction of the NTL energy enters the substrate (middle of Fig. 2.13).
2. **Impact-ionization (II)**: a charge drifting through the substrate collides with another charge loosely bound in a defect state and brings it to the conduction band so that more than the expected NTL energy is released in the substrate (right of Fig. 2.13).




Figure 2.13: Illustration of a normal eh pair event (left) compared to an event with charge-trapping (middle) and impact-ionization (right). Diagram taken from Ref. [53].

In either the case of CT or II, there is a single charge released in the bulk of the crystal and thus only part of the NTL energy expected from an eh pair is released. To model these effects, we can consider partitioning the path of a charge in the crystal in small, discrete steps. For each step, there is some probability p_i for CT ($i = \text{CT}$) or II ($i = \text{II}$). Thus, as shown in Ref. [54], after traveling a distance of Δz in n steps, the probability of a charge *not* undergoing CT or II is $(1 - p_i \Delta z / n)^n$. If we assume infinitely small step sizes, we have

$$\bar{C}_i = e^{-p_i \Delta z} = e^{-\Delta z / \tau_i} \quad (2.22)$$

where $\bar{C}_i$ is the probability of a charge *not* undergoing process i ($i = \text{CT}$ or II). In the last equality

of Eqn. 2.22, we have defined $\tau_i = 1/p_i$, the characteristic length a charge travels without undergoing process i . If we consider z to represent the fraction of the crystal thickness traveled through (*i.e.* $z = 0$ and $z = 1$ represent the bottom and top of the crystal), then we can characterize the total probability for a charge to undergo a process i while traveling through the substrate as

$$f_i \equiv 1 - e^{-1/\tau_i} \quad (2.23)$$

Where i in Eqn. 2.23 is CT or II. However, we can not necessarily assume that every instance of charge-trapping or impact-ionization happens with the same probability. For example, in the case of CT, it is possible that the probability for a hole to become trapped differs from that of an electron. In fact, there is some evidence for this shown in Section IV of Ref. [54]. Thus, these two cases of charge-trapping would technically need to be considered separately and are respectively referred to as f_{CTh} and f_{CTe} . The case of II is more complicated as one can consider a case where a freed electron frees an additional electron with some probability f_{IIee} differently from the scenario where the freed electron frees an additional hole (f_{IIeh}). We may similarly consider the case where the starting particle is a hole and we thus have the probabilities f_{IIhe} and f_{IIhh} , depending on whether the charge resulting from II is an electron or a hole (respectively). Rather than account for six different scenarios of CT and II, we will assume that $f_{CTe} = f_{CTh} = f_{CT}$ and that $f_{IIj} = f_{II}$ for all $i, j \in e, h$ for the remainder of this thesis.

Because the energy from the NTL effect is proportional to the work done on the charge by the electric field, which depends on the distance the charge travels through the substrate, events with CT have a lower energy compared to similar events without CT while the opposite is true for events with II. To demonstrate how $f_{CT, II}$ might affect an energy spectrum, I plot an example NTL energy spectrum consisting of poisson distributed eh peaks under different values of $f_{CT, II}$ in Figure 2.14.

I use $V = 50$ V and assume an energy resolution of 2 eV. Thus, rather than a spectrum purely




Figure 2.14: A simulated energy spectrum consisting of poisson distributed eh peaks at a 50 V bias under different probabilities of charge-trapping (right) and impact-ionization (left). For simplicity, only NTL energy is shown ($E_R = 0$ eV). Note that smaller values for f_{II} are considered (compared to those shown for f_{CT}) since previous studies have found f_{II} to be small in most cases [54].

consisting of eh peaks like the ideal case, we expect a non-zero number of events associated with ionizing signals to fall within the inter-peak regions. For a more detailed understanding of how these effects are modeled, I refer the reader to Ref. [54].

2.3 Signal Models

We can now consider the detector's response to dark matter interactions. In particular, I will show the energy spectrum and rate we could expect from DM-electron scattering models, absorption of dark photons, and absorption of axion-like particles. In each case, I will first begin with a derivation of the differential rate (*i.e.* the rate as a function of the recoil energy deposited during an interaction). Next, I will show the expected response of the detector with a voltage bias applied. I will also note that a similar explanation of these signal models can be found in Dr. Enze Zhang's thesis (Ref. [55]), who applied them for HVeV Run 4, the experiment which laid the foundation for much of the work shown in this thesis.

2.3.1 Dark Matter-Electron Scattering

As discussed in Chapter 1, hidden sector dark matter models allow for DM-electron scattering via a hidden mediator. The calculation of the differential rate is complicated by the fact that we are considering electrons in semiconductors, which have complicated energy eigenstates due to the band structure in crystals. However, as I will show here, we can first consider the scattering rate with a free electron. The resulting equation for the differential rate can then be corrected by a form factor, which accounts for the electron being bound in a crystal. For a detailed derivation I refer the reader to Appendix A of Ref. [56], which I abridge here.

The rate of DM interactions which cause electron excitations for a single target can be generally represented as

$$R = \frac{\rho_\chi}{m_\chi} \int g(v) \sigma v_{1 \rightarrow 2} d^3v = \frac{\rho_\chi}{m_\chi} \langle \sigma v \rangle_{1 \rightarrow 2} \quad (\text{Rate for a single electron target}) \quad (2.24)$$

where ρ_χ is the local DM density, m_χ is the mass of a single DM particle, $g(v)$ is the velocity distribution of DM particles, and $\sigma v_{1 \rightarrow 2}$ is the product of the velocity and the cross-section corresponding to the electron transitioning from energy states 1 and 2. Note that Eqn. 2.24 can be understood as the product of the flux of DM particles through a cross-sectional area in which there would be an interaction. However, note that the cross-sectional area, which in this case is the interaction cross-section σ , in general depends on the velocity of the DM particle and is therefore included in the integral. For DM scattering with a free electron, we have

$$\sigma v_{free} = \frac{1}{4E'_\chi E'_e} \int \frac{d^3q}{(2\pi)^3} \frac{d^3k'}{(2\pi)^3} \frac{(2\pi)^4}{4E_\chi E_e} \delta(E_i - E_f) \delta^3(\vec{k}' + \vec{q} - \vec{k}) |\overline{\mathcal{M}_{free}(\vec{q})}|^2 \quad (2.25)$$

where $E_i = E_e + E_\chi$ and $E_f = E'_e + E'_\chi$ are respectively the initial and final energy states of the DM-electron system, $\vec{k}$ is the initial momentum of the electron (such that the initial state of the

electron can be represented as $|e\rangle_i \propto e^{i\vec{k}\cdot\vec{r}}$, and $\vec{q}$ is the momentum transferred to the electron during the interaction. $\mathcal{M}_{\text{free}}(\vec{q})$ is the scattering matrix element computed from the Feynman diagram shown in Figure 1.6 and $|\mathcal{M}_{\text{free}}|^2$ is its absolute square averaged over the initial spins and summed over the final spins. Our first goal is to use the general form of Eqn. 2.25 to calculate $\sigma v_{1\rightarrow 2}$. To do this, first note that $\mathcal{M}_{\text{free}}$ enters Eqn. 2.25 via the calculation of the scattering amplitude with some interaction Hamiltonian H_{int} :

$$\langle \chi_{p-q}, e_f | H_{\text{int}} | \chi_p, e_i \rangle = C \mathcal{M}_{\text{free}} \times (2\pi)^3 \delta^3(\vec{k}' + \vec{q} - \vec{k}) \quad (\text{free electron}) \quad (2.26)$$

The initial and final states of the DM-electron system are represented by $|\chi_p, e_i\rangle$, where p indicates the momentum of the DM particle. C in Eqn. 2.26 is simply a coefficient but does not contain any important dependencies. The term with the delta function comes from the product $\langle e_i | e_f \rangle$ integrated over all k -space. If we instead consider a bound electron, then Ref. [56] shows that we can represent the resulting scattering amplitude involving the bound electron in terms of the free scattering amplitude and an atomic form factor, $f_{1\rightarrow 2}$ which encodes the effect of the bound state wavefunctions (represented in momentum space by $\tilde{\psi}(k)$ in the equation below) on the matrix element.

$$\langle \chi_{p-q}, e_f | H_{\text{int}} | \chi_p, e_i \rangle = C \mathcal{M}_{\text{free}}(\vec{q}) \int \frac{V d^3k}{(2\pi)^3} \tilde{\psi}_f^*(\vec{k} + \vec{q}) \tilde{\psi}_1(\vec{k}) \quad (2.27)$$

$$= C \mathcal{M}_{\text{free}}(\vec{q}) f_{1\rightarrow 2}(\vec{q}) \quad (\text{bound electron}) \quad (2.28)$$

We can further simplify Eqn. 2.25 in the case of scattering with bound electrons by separating out the $\vec{q}$ dependence of $\mathcal{M}_{\text{free}}$. This is done by parameterizing $\mathcal{M}_{\text{free}}(\vec{q})$ in terms of a reference

momentum, which is typically chosen to be $\vec{q}_{\text{ref}} = \alpha m_e$, where α is the fine structure constant and m_e is the mass of an electron.

$$|\mathcal{M}_{\text{free}}(\vec{q})|^2 = |\mathcal{M}_{\text{free}}(\alpha m_e)|^2 \times |F_{\text{DM}}(\vec{q})|^2 \quad (2.29)$$

Here, F_{DM} is the dark matter form factor and encodes the dependence of the scattering amplitude on the momentum transfer, which depends on the DM model under consideration. We will consider the case where the hidden sector mediator is much heavier than the DM particle, in which case $F_{\text{DM}} = 1$, as well as the case where we have a very light mediator, in which case $F_{\text{DM}} \propto 1/q^2$.

We then define

$$\bar{\sigma}_e = \frac{\mu_{\chi e}^2 |\mathcal{M}_{\text{free}}(\alpha m_e)|^2}{16\pi m_\chi^2 m_e^2} \quad (2.30)$$

where $\mu_{\chi e}$ is the reduced mass of the DM-electron system.

Before writing our expression for $\sigma v_{1 \rightarrow 2}$, we can simplify the expression further by writing out the energies in Eqn. 2.25 since we can assume DM is not relativistic.

$$E_i = m_\chi + m_e + \frac{1}{2}m_\chi v^2 + E_{e,1} \quad (2.31)$$

$$E_f = m_\chi + m_e + \frac{|m_\chi v - q|^2}{2m_\chi} + E_{e,2} \quad (2.32)$$

In the equations above, $E_{e,i}$ is the i th energy state of the electron and v is the velocity of the DM particle.

Substituting Eqns. 2.29—2.32 into Eqn. 2.25, we can write the scattering DM-electron cross-section for bound electrons:

$$\sigma_{v_{1 \rightarrow 2}} = \frac{\bar{\sigma}_e}{\mu_{\chi e}^2} \int \frac{d^3 q}{4\pi} \delta(\Delta E_{1 \rightarrow 2} + \frac{q^2}{2m_\chi} - qv \cos(\theta)) \times |F_{\text{DM}}(q)|^2 |f_{1 \rightarrow 2}(\vec{q})|^2 \quad (2.33)$$

In Eqn. 2.33, θ is the angle between the DM particle's initial and final momentum. It is worth noting that Eqn. 2.33 shows that the cross-section for bound electrons can be represented almost entirely in terms of the scattering cross-section for free electrons, while the complexity of the bound states are accounted for by the atomic form factor $f_{1 \rightarrow 2}$. Thus, by considering the electron bound states, 1 and 2, as the energy states in crystal semiconductors, we can calculate the rate of excitations in our detectors. The energy states of a crystal are written as a summation of functions u that are periodic with respect to the reciprocal vectors of the crystal lattice $\vec{G}$:

$$\psi_{i\vec{k}}(\vec{r}) = \frac{1}{\sqrt{V}} \sum_{\vec{G}} u_i(\vec{k} + \vec{G}) e^{i(\vec{k} + \vec{G}) \cdot \vec{r}} \quad (2.34)$$

Here V is the volume of the crystal and $\psi_{i\vec{k}}$ are the eigenstates corresponding to the i th energy level and crystal momentum $\vec{k}$. The crystal form factor f_{crystal} can then be calculated by substituting the wavefunctions in Eqn. 2.34 into the LHS of Eqn. 2.33. The resulting integral over k -space in the resulting expression is f_{crystal} and is generally computed numerically. The corresponding cross-section can be substituted into Eqn. 2.24 for the total rate of electron excitations in the semiconductor crystal. Taking the derivative with respect to energy we obtain the differential rate shown in Eqn. 2.35 (Eqn. A.32 in Ref. [56]).

$$\frac{dR_{\text{crystal}}}{d\ln(E_e)} = \frac{\rho_\chi}{m_\chi} N_{\text{cell}} \bar{\sigma}_e \alpha \frac{m_e^2}{\mu_{\chi e}^2} \int d\ln(q) \left(\frac{E_e}{q} \eta(v_{\min}(q, E_e)) \right) F_{\text{DM}}(q)^2 |f_{\text{crystal}}(q, E_e)|^2 \quad (2.35)$$

The function η is the averaged DM velocity distribution truncated at the minimum velocity required for a transition, $v_{\min}$, which depends on the difference between the energy levels in consideration and the transferred momentum q .

$$\eta(v_{\min}) = \int \frac{d^3v}{v} g_{\chi}(v) \Theta(v - v_{\min}) \quad (2.36)$$

Finally, N_{cell} is simply the number of crystal cells in the target volume and can also be represented in terms of the mass of the target.

I plot the differential rate for a chosen value of $\bar{\sigma}_e$ in Figure 2.15 to show the spectrum of recoil energies expected from DM-electron scattering. The detectors' response to the recoil energies




Figure 2.15: The differential rate for DM-electron scattering in silicon for $F_{\text{DM}} = 1$ (left) and $F_{\text{DM}} \propto 1/q^2$ (right). Different colored lines correspond to different DM particle masses. Here, I assume $\bar{\sigma}_e = 1\text{e-}37 \text{ cm}^2$.

shown in Fig 2.15 can then be determined assuming we know the ionization yield and the detectors' resolution and CT/II probabilities. We will discuss modeling of detector specific parameters (such as $f_{\text{CT, II}}$, resolution, etc.) in Chapter 4. For now, I assume some reasonable values and plot the resulting models in Figure 2.16.




Figure 2.16: The differential rate seen by the detector for DM-electron scattering in silicon for $F_{\text{DM}} = 1$ (left) and $F_{\text{DM}} \propto 1/q^2$ (right). Different colored lines correspond to different DM particle masses. Here, I assume $\bar{\sigma}_e = 1\text{e-}37\text{ cm}^2$, a voltage bias of 100 V, a resolution of 4 eV, as well as some other stand-in values for charge-trapping and impact ionization.

2.3.2 Dark Photon Absorption

Another popular hidden-sector dark matter candidate considered in this thesis is the dark photon—a gauge mediator introduced via a hidden U(1) symmetry that can weakly couple to the SM via the vector portal. In particular, we will consider the absorption of dark photons via the process shown by the Feynman diagram in Figure 1.6. For the derivation of absorption rate of dark photons (and, in the next section, axion-like particles), I will follow the derivation shown in Ref. [57]. We begin with the generic equation for the rate:

$$R = \frac{1}{\rho_{\text{target}}} \frac{\rho_{\chi}}{m_{\chi}} \langle n_e \sigma_{\text{abs}} v \rangle \quad (2.37)$$

Here, ρ_{target} is the density of the target (*e.g.* silicon), σ_{abs} is DM-electron absorption cross-section, and n_e is the number density of electrons. Following Ref. [57], we will first calculate the photon absorption cross-section. Then, by relating this cross-section to the DM absorption cross-section, we can calculate the rate of DM absorption in a target.

Ref. [57] first shows that the velocity averaged photon absorption cross-section is equivalent to the real part of the target's conductivity, σ_1 :

$$\langle n_e \sigma_{\text{abs}} v \rangle_\gamma = \sigma_1(\omega) \quad (2.38)$$

where ω is the frequency (equivalently, in natural units, energy) of the photon. If we then consider a dark photon, then there is a kinetic mixing of the dark photon with the standard model photon via

$$\frac{\epsilon}{2} F_{\mu\nu} F'^{\mu\nu} \in \mathcal{L} \quad (2.39)$$

where $F_{\mu\nu}$ is the SM electromagnetic field strength tensor and $F'_{\mu\nu}$ is the analogous hidden sector field strength tensor. As shown in Ref. [58], this leads to a coupling between the dark photon and the electromagnetic current, J_{EM} :

$$\epsilon e J_{\text{EM}}^\mu A'_\mu \in \mathcal{L} \quad (2.40)$$

Taking into account in-medium effects, one obtains the effective kinetic mixing [57]

$$\epsilon_{\text{eff}}^2 = \frac{\epsilon^2 m_\chi^2}{m_\chi^2 - 2m_\chi \sigma_2 + \sigma_2^2 + \sigma_1^2} \quad (2.41)$$

where m_χ is the mass of the dark photon, and σ_2 is the imaginary part of the conductivity of the target. The term in the Lagrangian shown in Eqn. 2.40, is equivalent to that of the photon absorption cross-section and thus we can see that the matrix elements for the absorption of dark photons by electrons are related to the matrix elements of photons by electrons via

$$|\mathcal{M}|^2 = \epsilon_{eff}^2 |\mathcal{M}_\gamma|^2 \quad (2.42)$$

Therefore, we have that

$$\langle n_e \sigma_{\text{abs}} v \rangle = \epsilon_{\text{eff}}^2 \langle n_e \sigma_{\text{abs}} v \rangle_{\gamma} = \epsilon_{\text{eff}}^2 \sigma_1(\omega = m_{\chi}) \quad (2.43)$$

where in the last equality, I made the dependence of the real part of the conductivity on the energy of the dark photon explicit. Substituting this back into Eqn. 2.37, we have

$$R = \frac{\epsilon_{\text{eff}}^2}{\rho_{\text{target}}} \frac{\rho_{\chi}}{m_{\chi}} \sigma_1(\omega = m_{\chi}) \quad (2.44)$$

Note that while considering a particular dark photon mass, the recoil energy released in the detector is singular and is equal to the energy of the dark photon, m_{χ} , which is generally small—on the scale of a few eV to a few tens of eV. Applying our knowledge of the detector's response to an electron recoil of known energy, we obtain the signal models shown in Figure 2.17.




Figure 2.17: The expected energy spectrum seen by the detector from dark photon absorption. Different color lines correspond to different dark photon masses. I assume a kinetic mixing strength of 5×10^{-13} , a 100 V bias, a 4 eV resolution, and nominal values for $f_{\text{CT},\text{II}}$.

2.3.3 Axion-like Particle Absorption

The last DM candidate that we will consider in this thesis is the axion-like particle (ALP). The ALP is an extension to the axion, a particle postulated as a potential solution to the strong CP problem [59], and appear in many theories proposing extensions to the SM. As psuedoscalars, the ALP field, a , can couple to electrons (with field represented by ψ in Eqn. 2.45) via the following interaction term in the Lagrangian [60]

$$-\frac{\delta_\mu a}{f_a} \bar{\psi} \gamma^\mu \gamma^5 \psi \in \mathcal{L} \quad (2.45)$$

which is similar to the photoelectric interaction term. Here, the parameter f_a characterizes the strength of the interaction.

Similar to our calculation of the rate of dark photon absorption, we can relate the axio-electric cross-section of the interaction shown in Eqn. 2.45 to the cross-section of the photoelectric effect. In general, there are differences between these interactions stemming from the fact that ALPs are pseudoscalars while photons are vectors. Specifically, the wavefunction of absorbed photons in a photoelectric interaction is proportional to $e^{i\vec{k}\cdot\vec{r}}$ where $k = \omega$ while the wavefunction of the absorbed ALP is instead proportional to $e^{im_\chi \vec{v}\cdot\vec{r}}$ where m_χ and $\vec{v}$ are the mass and velocity of the ALP. However, since we are considering non-relativistic DM, $\vec{v}$ can be considered small enough so that the exponential term is close to unity. Similarly, if we consider only energies much smaller than the mass of an electron (which is many orders of magnitude above the energies we consider for the DM search), then the contribution of the exponential in the photon wavefunction is also negligible [60]. Under these assumptions, the cross-section of the axio-electric effect can

be directly related to the cross-section of the photoelectric effect by [60]

$$\sigma_a v \approx \frac{3m_\chi^2}{4\pi\alpha f_a^2} \sigma_{p.e.} \quad (2.46)$$

where σ_a is the ALP absorption cross-section and $\sigma_{p.e.}$ is the photo-electric cross-section.

We define the following coupling constant to characterize the strength of the ALP absorption interaction.

$$g_{aee} = \frac{2m_e}{f_a} \quad (2.47)$$

Substituting this into our generic expression for the DM interaction rate, we arrive at the rate of ALP absorption by electrons

$$R = \frac{\rho_{\text{DM}}}{\rho} \frac{3g_{aee}^2 m_\chi}{16\pi\alpha m_e^2} \sigma_1(m_\chi) \quad (2.48)$$

Just as with dark photons, the expected spectrum of energy delivered to a substrate is a delta function centered on the energy of the ALP (m_χ in the non-relativistic limit). Figure 2.18 shows the expected response of the detector resulting from the absorption of ALPs at various masses.

We can compare the signal models we derived in this section to actual observations from operating HVeV detectors and set exclusion limits on the parameters which characterize the strength of the various interactions—namely $\bar{\sigma}_e$, ϵ , and g_{aee} .

2.4 The Development of HVeV Detectors

So far, in this chapter, we have laid out the operation and response of HVeV detectors, which, owing to their excellent energy resolution and ability to amplify ionizing signals, are capable of probing DM models with masses below the GeV scale. Prior to the data-taking campaign presented in this thesis, HVeVs have been used for dark matter searches four times—we refer to these




Figure 2.18: The expected energy spectrum seen by the detector from ALP absorption. Different color lines correspond to different ALP masses. I assume an axio-electric coupling, g_{aee} , of $5\text{e-}11$, a 100 V bias, a 4 eV resolution, and nominal values for $f_{\text{CT,II}}$.

searches as HVeV Runs 1 through 4, with 4 being the most recent, and, as shown in Figure 2.19⁵, have steadily improved their sensitivity to DM. However, it is important to emphasize that HVeVs are still an emerging technology—thus far, their operation has been in a primarily R&D capacity where much of the data taken is to understand the detectors’ response and operation as opposed to accumulating exposure for DM search experiments. In spite of this, HVeV detectors have recently become promising instruments in the search for sub-GeV dark matter.

To contextualize the work of this thesis, which sought to further improve upon the exclusion limits set by HVeV Run 4, I will briefly describe the progress made by HVeV Run 3 [62] and Run 4 [63]. I refer the reader to Ref. [64] for the foundational results from HVeV Run 2.

⁵I exclude the limits set by HVeV Run 1 as they did not yet account for the effects of charge-trapping and impact-ionization [61].





Figure 2.19: Limits on the DM-electron scattering cross-section for a light mediator ($F_{\text{DM}} = 1/q^2$) set by HVeV Runs 2 (blue), 3 (orange) and 4 (pink).


2.4.1 HVeV Run 3

The detector payload for HVeV Run 3 was operated out of a shallow (~ 100 m) underground facility at Fermilab called NEXUS (Northwestern EXperimental Underground Site). The NEXUS facility is clean room that hosts a *CryoConcept* dilution refrigerator capable of cooling the detector payload to ~ 11 mK [62], [63]. The payload consisted of four $1 \text{ cm} \times 1 \text{ cm} \times 4 \text{ mm}$ HVeVs mounted between printed circuit boards (PCBs) that were secured to copper plates thermalized to the mixing chamber stage of the dilution refrigerator. (See Fig. 2.20a.) Just as in HVeV Run 2 [64], a 1.95 eV laser was used as calibration source. An optical fiber was used to transmit photons to the detector payload (after passing through an infrared filter), allowing the laser to be operated outside the fridge to avoid electrical crosstalk.


A blinded DM search using data taken with a voltage bias of 100 V from a single HVeV detector, which achieved a baseline resolution of ~ 3 eV. To develop data selection criteria, 10% data taken—consisting of ~ 13 g-days—was unblinded and used for analysis. While data from only




(a) HVeV Run 3 payload.



(b) HVeV Run 4 payload.



(c) Energy spectrum from HVeV Run 3.



(d) Energy spectrum from HVeV Run 4.

Figure 2.20: (top) Schematics of the HVeV Run 3 (left) and Run 4 (right) payloads. Note that while the detectors are visible in the Run 3 payload (they are clamped between PCBs that are secured to the vertical copper plates), they are concealed by the light-tight copper housings in the Run 4 payload. (bottom) The energy spectra from HVeV Run 3 (left) and HVeV Run 4 (right). The final spectra used in the DM searches are shown in Refs. [62] and [63] respectively.

one detector was used for setting exclusion limits, the other detectors were used to reject coincident events. The final energy spectrum used in the DM search analysis, consisting of ~ 8 g-days of live time after data selections (13.14 g-days before data selection), is shown in Figure 2.20c.

The exclusion limits set by HVeV Run 3 significantly improved on the exclusion limits set by HVeV Run 2—an improvement that was largely attributed to the increased overburden provided by the rock above the NEXUS facility as compared to the surface facility where HVeV Run 2 was

conducted. However, in spite of this, HVeV Run 3 was limited by a peculiarly high rate of events in the $\geq 2\text{eh}$ peaks. Of the events that contributed to these peaks, most of them were found to be what are now referred to as “bursts”—events characterized by an energetic primary followed by multiple lower-energy secondaries. (See Appendix A for a detailed discussion of burst events). It was later shown that the largest contributor to these backgrounds was likely luminescence of the SiO_2 in the PCBs that were used to mount the detector [65].

2.4.2 *HVeV Run 4*

The discovery of luminescence in the PCBs used in Run 3 led to a complete redesign of the detector housings, the result of which was employed in HVeV Run 4. To avoid luminescent materials, HVeV Run 4 clamped the detectors in copper housings. (See Fig. 2.20b.) A thin PCB was used to facilitate readout, however, the copper housings were designed so that only a small area of the PCB had a line of sight to the detectors. In addition to avoiding luminescence from the PCBs, the housings were designed with interlocking rims to maintain light tightness and thus avoid an increased event rate due to infrared photons from higher temperature stages in the fridge. (That said, to facilitate the readout of the detector, some openings in the housing had to be permitted.)

This redesign of the detector housings also required a redesign of the photon calibration system. As opposed to a laser that was operated out of the fridge, we designed copper modules that hosted LEDs and could stack on top of the detector housings. The “lid” of the modules hosted the PCB that facilitated operation of $\sim 2.05\text{ eV}$ LEDs (see left side of Fig. 2.21). The bottom part of the modules had small ($\sim 0.2\text{ mm}$) holes drilled into their bottom side so that photons emitted from the LED would be collimated onto the detectors. Infrared filters were placed between the LED and the pinhole. The inside of the modules was coated in a black metallic material to absorb any incoming infrared photons, thus lowering the possibility of them transmitting through the filters and being absorbed by the detector.




Figure 2.21: (right) The lid of an LED module. In the center of the lid is the PCB with two LEDs (~ 2.05 eV). (left) The “bottom” piece of the LED module, which contains infrared filters. Small holes (~ 0.2 mm diameter) are drilled in the bottom of the module. When the module is closed, LED photons pass through the infrared filters and are collimated through the drilled holes.

The payload shown in Fig. 2.20b was installed in the NEXUS cryostat and operated at a temperature of ~ 11 mK. Just as in Run 3, a blinded DM search was conducted using data taken from one detector operated with a 100 V bias while the other three detectors included in the payload were used to reject coincident events. 30% of the 10.8 g-days of total exposure was used to develop an analysis.

The result of this redesign can be seen by comparing the spectra shown in Figure 2.20—the rate of events in the eh peaks was drastically reduced, particularly those in the $\geq 2\text{eh}$ peaks, and led to the significant improvement in DM sensitivity exhibited in Fig. 2.19. However, despite a significant reduction in the background rate, Run 4 was still background limited; this time, however, the backgrounds are at lower energies. So, just as HVeV Run 4 improved on the sensitivity of HVeV Run 3 by mitigating high energy backgrounds, the goal of this work was to improve on the sensitivity of Run 4 by mitigating low energy backgrounds. As it turns out, observations of an excess of events at low energy are not unique to HVeV Run 4, but are a common feature observed by essentially every experiment looking for low mass DM; see Ref. [66] for a review of the current observations of this excess. This will be discussed more in Chapter 5.

CHAPTER 3

HVEV AT CUTE

HVeV Run 5 (or, more affectionately, “HVeV at CUTE”) was designed to answer the questions left open by its predecessors. HVeV Runs 1-3 [61], [62], [64] had continually matured our understanding of the detectors’ response and operation until, after mitigating high energy backgrounds observed in HVeV Run 3, HVeV Run 4 [63] reached the obstacle that other experiments searching for DM at low masses were encountering—backgrounds of unknown origin which posed a hard limit on sensitivity [66]. Thus, any improvement in sensitivity will necessarily require a better understanding of these backgrounds.

As such, the initial plan for HVeV Run 5 did not initially include a DM search, though we changed this when we discovered potential to explore new parameter space. We set out to study the low-energy excesses, both ionizing and non-ionizing, which were limiting sensitivities across the field. At the time of writing this thesis, the analysis of the data taken during HVeV Run 5 with the goal of gaining insight into the origin of these backgrounds are in progress or not yet started. The remainder of this thesis will discuss the analysis of HVeV Run 5 thus far including some investigation on the backgrounds at low energy, but primarily focusing on an ERDM search.

3.1 Experimental Setup

3.1.1 The CUTE Facility at SNOLAB

HVeV Run 5 was operated out of the Cryogenic Underground TEsting (CUTE) Facility (leading to its more colloquial nickname “HVeV at CUTE” or “HVeV@CUTE”) at SNOLAB—a deep underground laboratory located in Sudbury, Canada. At ~ 2 km beneath the Earth’s surface, the rock overburden above CUTE provides roughly 6 km m.w.e of shielding, reducing the expected

muon flux by a factor of $\sim 1,000$ compared to expectations at NEXUS [67], [68]. During operation, the cryostat hosting the detector payload is placed in a drywell enveloped by additional layers of shielding (see Fig. 3.1) to further reduce backgrounds from the surrounding environment's backgrounds consisting primarily of radioactive contamination in SNOLAB's cavern walls.




Figure 3.1: A diagram of the CUTE facility at SNOLAB. Figure taken from Ref. [68].

The drywell is seated inside a stainless-steel tank filled with ultra-pure water that provides an additional 1.5 m of shielding around the drywall and another 1 m of shielding below it. Inside the drywell are an two layers of low-activity lead with a combined thickness of ~ 11 cm around the sides and 15 cm on the bottom. Between the two layers of lead is a μ -metal shield, which reduces the magnetic flux within the cryostat; this reduction is necessary for operating the SQUIDs that are used to measure the current through the detectors' TESs. Since the cryostat is placed into the drywell from above, there is a large hole above the cryostat during operation. To offset the lack of shielding in this direction, an 11 cm-thick, copper-plated lead shield is placed in the cryostat above the detector payload. Additionally, a 20 cm-thick layer of polyethylene is placed on top of

the deck to shield the top-side of the cryostat from neutron backgrounds.

Overall, the radiogenic backgrounds at CUTE are expected to be many orders of magnitude less than those at NEXUS for energies $O(>\text{keV})$. However, one of the challenges about looking for eV-scale signals is that the background rates in our energy region of interest ($\lesssim 1 \text{ keV}$); thus, it is not known whether we can expect lower background rates in the HVeV detectors. In fact, comparing event rate measured by HVeVs between CUTE and NEXUS could shed some light on the origin of backgrounds currently limiting the field of experiments looking for low-mass DM signals!

3.1.2 The Detector Payload

The detector payload consisted of six $1 \text{ cm} \times 1 \text{ cm} \times 4 \text{ mm}$ HVeV detectors—the four detectors from HVeV Run 4, which were introduced in Chapter 2 as well as two detectors from a new version of the HVeV design [69]. These detectors use mask designs that were tested during previous HVeV runs (one uses the NFC mask and the other uses an NFH) but they differ from previous HVeVs, such as those used in Run 4, in two key ways:

1. The tungsten TES was fabricated with a lower T_c (estimated to be $\sim 45 \text{ mK}$, compared to $\sim 65 \text{ mK}$, which was the approximate T_c of detectors used in Run 4). Refs. [27], [31] show that the phonon energy resolution, σ_{ph} of QETs scales with the T_c thus, we expect a better energy resolution from these detectors.
2. Between the instrumentation on the 1 cm^2 faces of the detector and the silicon substrate is a 40 nm layer of amorphous SiO_2 (See the graphic included in Figure 3.2.) The primary motivation for this layer was to due to the high rate of 1eh events observed in previous HVeV runs. It is possible that these events could be caused by leakage of single charge-carriers near the surface of the substrate—if that were the case, the charge would propagate throughout

the length of the surface and thus mimic a bulk eh event. A possible origin for these events could be free electrons tunneling from the aluminum in the QETs or the biasing grid into the substrate. The large potential barrier provided by the SiO_2 , however, would make this less likely and could therefore possibly improve the rate of 1eh events.




Figure 3.2: A graphic showing the cross-section of a SiO_2 HVeV detector. The silicon substrate is shown in dark gray surrounded by 40 nm layers of amorphous SiO_2 which are shown in yellow. The QETs deposited on the substrate are shown in light blue and the aluminum biasing grid on the other side of the detector is shown in light grey. Drawing not to scale.

The design of the detector tower built on the success of Run 4 and utilized many of its same components. The same copper detector housings were used to avoid using luminescent material near the detectors. The same LED modules that position an LED to shine on the center of the side of the detector with the biasing grid were used and, just as in Run 4, the detector housings and LED modules were assembled into a tower which is pictured in Figure 3.3. The tower was then screwed into a copper plate (referred to as the “mixing chamber plate”), which was thermalized to the mixing chamber of a CryoConcept dilution refrigerator.

The readout system used to measure and digitize the current through each detector channel was designed with particular care to avoid electrical crosstalk and group loops—both of which were problematic during HVeV Run 4. The system used four shielded looms with 25 wires (12 twisted pairs and a single wire connecting to the stainless-steel mesh shielding) to connect the warm-side




Figure 3.3: **Left:** A CAD model of the detector payload for HVeV@CUTE. The copper housings and LED modules are shown stacked in orange/red, flex cables which facilitate readout and biasing between the mixing chamber plate and the tower are shown in green. In gray is the copper plate which the tower is mounted to. Highlighted in blue is a copper ring which is bolted to the mixing chamber plate. **Right:** A picture of the detector payload for HVeV@CUTE after assembly.

electronics and readout to the tower. Two of the looms contained “housekeeping” lines—lines used for biasing LEDs and applying HV bias to detectors. The other two contained the TES bias and SQUID bias lines respectively. Choosing to separate the LED bias lines from the TES bias lines was a response to the observations from HVeV Run 4 where crosstalk associated with operating the LEDs contributed a significant systematic uncertainty to the calibration [70]. At each thermal stage, a feedthrough board was installed to reduce line of sight for any IR photons that might originate from warmer parts of the fridge and travel to colder stages. We taped the edges of the feedthrough board with copper tape to further prevent any IR leakage. At the mixing chamber stage we added an additional layer of protection for IR by pouring a mixture of black stycast and copper powder to absorb IR that might otherwise pass through the feedthrough.

To measure the currents through each detector channel, two 6-channel SQUID cards were mounted below the mixing chamber plate and thermalized to the cold plate (~ 100 mK). Each SQUID card was protected from external magnetic fields by an aluminum shield that, during oper-

ation, would be superconducting.

The output of the SQUIDs is digitized and read out by a SuperCDMS Detector Control and Readout Card (DCRC) [71], which is also used to control the detector and SQUID biases and facilitates the readout scheme described in Chapter 2. The DCRC samples the digitized current at a frequency of 625 kHz which is then downsampled by a factor of 4 in order to keep the bandwidth below the DCRC’s limit of ~ 7 MB/s. Data is recorded continuously in half-second “traces” and stored for offline triggering and analysis. We further organize data into “series” denoted by the timestamp corresponding to the start of logging. Generally, each series consists no more than 12 hours of live time before any data selection.

3.1.3 Run Summary

HVeV Run 5 consisted of two consecutive fridge cool downs, CUTE Runs 39 and 40 separated by a partial warm up to 20 K and together spanning a period of roughly four and a half months. After reaching the mixing chamber base temperature of ~ 10 mK, we found that 10 of our 12 detector channels were operational including all the channels from the detectors included in HVeV Run 4. However, we did not see the superconducting transition of the outer channel of NFCSO nor from the inner channel of NFHSO. For NFCSO, we found that a damaged trace on the DCRC was preventing us from locking the readout loop. As such, we were able to recover the channel simply by replacing the DCRC. However, the inner channel of NFHSO was not recovered and it was not until the end of HVeV Run 5 that we discovered the problem: a wire bond on the SQUID board had become disconnected.

Following the initial cooldown at the beginning of each fridge run, we frequently measured the bias power and bias current required to operate the detectors at 40% through their superconducting transition. We stopped when we noticed the detectors reaching a stable operating point. However, about three months into CUTE Run 40, we found that the detectors must have continued to slowly

cool. The detectors, as we discovered, had been operating below 20% of their normal resistance. Despite this accident, we saw no degradation in the performance of the devices [69].

Much of the data we took at CUTE has not yet been investigated fully. However, there are future plans to do so. In particular, we took data to test a number of hypotheses regarding the origin of the LEE, to investigate discrepancies in calibrations between HV and 0V operations discovered in HVeV Run 4 [63], [72], and to expand our understanding of charge propagation in solid-state detectors. The majority of these investigations are still open and much of the data we have from HVeV at CUTE has yet to be fully explored. The rest of this chapter will present a blinded ERDM search using NFCSSO, which displayed the best resolution observed from an HVeV detector at the time of writing this thesis and showed potential to uncover new parameter space.

As stated at the beginning of this chapter, we had not originally planned on performing a blinded DM search. However, during the run, we noticed that operating the detectors at lower bias voltages seemed to result in lower rates in the 1eh peak—a critical metric for sensitivity to ERDM since the eh peaks are the primary signal regions of interest for ERDM. We experimented with running the detectors at as low a bias voltage as possible while keeping the 1eh peak in the region above where the LEE began to dominate (<30 eV) and determined the optimal bias to be -35 V. After taking 24 hours of live time with this high voltage, we found that the rate of events in the 1eh peak was lower than had been previously measured for HVeV detectors. We therefore decided to take blinded data for an ERDM search. Thus, starting August 10, 2024 and continuing through August 20, we took blinded data with a high-voltage bias of -35 V. The blinding was performed by saving every 9 out of 10 minutes of data in a directory with reading permissions removed. The remaining 10% of raw data was saved to allow full access. This 10% data was examined purely to see that the operating conditions remained consistent throughout the run. That is, this data was not examined to develop an analysis. For the purpose of analysis, we used the original data taken

before August 10 was used and is henceforth referred to as the “open” data. The rest of this chapter will explain the analysis as well as show preliminary DM exclusion plots for the four DM models whose signal models were discussed in Chapter 2.

3.2 Event Reconstruction

3.2.1 Offline Triggering

Data taken during HVeV Run 5 was saved continuously and stored in half-second long-time streams. A gaussian derivative filter is applied to each trace which is then subjected to a threshold trigger. Wherever the filtered trace exceeds the threshold, a window is drawn between the initial trigger point and the point at which the filtered trace drops below its value at the initial trigger point. If the filtered trace did not drop below the value of the initial trigger point for more than 1024 samples (~ 6.5 ms) after the trigger, the window is then set to 1024 samples after the initial trigger. The trigger point was then adjusted to correspond with the maximum of the filtered trace within the window. In addition to a threshold trigger, we triggered at random point within each half-second trace in order to monitor noise conditions. For each trigger, a 2048 sample (~ 13 ms) window centered at the trigger point was then stored for further analysis.

3.2.2 Optimal Filtering

An optimal filter (OF) was used to calculate the amplitude of triggered pulses. Optimal filtering is a signal processing technique that can be used to extract signal in the presence of noise. The filter itself can be thought of as a flexible band-pass filter which allows frequencies where the signal-to-noise ratio is highest and reduces frequencies where the signal-to-noise is lowest. The OF, Φ in Eqn. 3.1, is calculated in frequency space as the ratio of the complex conjugate of the

signal template, $s(f)$, and the noise CSD, $J(f)$.

$$\Phi = \frac{s^*(f)}{J(f)} \quad (3.1)$$

In the simplest case, where the signal template only depends on the frequency, f , the amplitude returned by the OF can be written as [27]

$$A = \int_{-\infty}^{\infty} \Phi'(f)v(f)df \quad (3.2)$$

where Φ' is the OF normalized by $\int \Phi(f)s(f)df$ and $v(f)$ is the measured signal. For more information regarding optimal filtering, I refer the reader to the excellent summary given in Appendix E of Noah Kurinsky's thesis [27].

Throughout this thesis, I will refer to two quantities returned by the OF, OF0 and OFL, which I define below:

OF0 is the amplitude returned by the OF when the signal is not allowed to move in the time domain. In this case, OF0 is calculated via Eqn. 3.2.

OFL is the amplitude returned by the OF when the signal is allowed to shift in the time domain over a small window limited around the trigger point. The time shift is determined by minimizing the χ^2 (calculated in the frequency domain) with respect to the time shift) and is sometimes referred to as the OFL time.

We calculated the noise CSD, which is used as an input in the OF from random triggers in data. Random triggers from the unblinded portions of each series (including the fully unblinded data) are passed through a set of strict criteria based on their mean, slope, skew, and range (the difference in the minimum and maximum in a given trace). A periodogram is then used to calculate a CSD from each pulse and the result is averaged for each series. Figure 3.4 shows the averaged noise

CSDs from a number of series taken throughout the DM search.




Figure 3.4: Average CSDs calculated from random triggers collected from data series taken with a high voltage bias of -35 V.

For the most part, the CSDs remain relatively constant over time which is a sign that our noise environment does not significantly change throughout the course of data taking. We do, however, see slight deviations in the lower frequency regions below ~ 1 kHz which are most likely caused by contamination from low amplitude pulses. As such, the optimal filter is calculated using the CSD from the series with the lowest mean up to 1 kHz.

The signal template is constructed by averaging events seen in low background data. Ideally, these events would be representative of what we expect for a physics event (as opposed to an event caused by a fluctuation in noise) at an energy within our signal region of interest—for an ERDM search this would be in the region where we would expect to see electron-hole pair events. Thankfully, we know our background energy spectrum has features or peaks that correspond to events in which charges were allowed to propagate through the crystal. Thus, by selecting events in the 1eh peak region in background data, we ensure that we are examining events in an energy region of interest that is relevant to the ERDM search¹. To do this, we construct an initial OF using

¹Optimally, one would use pulses induced by a calibration source to construct the template so that the energy and

a signal template consistent with observations from previous HVeV runs. We do not necessarily expect this template to perfectly capture the shape of our average pulse but it should capture the same general features. Using the resulting OF, we can calculate an OFL spectrum from which we can identify the 1eh peak. Pulses selected from the 1eh peak then pass through additional selection criteria:

1. **Selection based on χ^2 :** With an initial guess of a signal template, we do not expect the χ^2 comparing pulses to the template to follow a χ^2 distribution. However, assuming the template captures the main features of a pulse, the distribution will still be peaked (though perhaps not at 1). To remove outliers, we fit the peak in the distribution to a gaussian (with mean μ and standard deviation σ) and exclude pulses with χ^2 greater than $3\sigma + \mu$.
2. **Selection based on baseline slope:** We exclude events based on the slope of their baseline (defined as the 1024 samples in the analysis window preceding the trigger point) in order to remove pileup events that occur on the falling edge of previous events or events contaminated by low frequency noise fluctuations. This selection doesn't depend on the OF and is therefore unaffected by our choice of template. We exclude events $\pm 3\sigma$ from the mean of the distribution where σ is determined by a gaussian fit.
3. **"dt" selection:** Once again targeting pileup, we define a selection criteria to exclude pulses based on how close in time they are to another pulse which is referred to as "dt". We exclude events which occur within 1024 samples of each other or, equivalently, pulses which occur within the same analysis window. This selection, while not dependent on the OF, does depend on the trigger threshold which can be set arbitrarily in which case this selection is

origin of the events is known. However, we chose not to use LED events for constructing the template because, as will be shown later in the chapter, we observed an electrical crosstalk signal while operating them. To avoid including the effects of the crosstalk in the signal template and possibly degrading our energy resolution, we opted instead to use pulses sampled from low background data.

highly dependent on the processing. To avoid this ambiguity, we consider only pulses with an amplitude 5 times greater than the spread of the amplitude of random triggers—this can be thought of as setting a threshold at 5 times the baseline resolution in units of current.

4. **OFL time selection:** The OFL time selection excludes pulses for which the OF significantly shifted with respect to the initial trigger point. While these pulses are not necessarily poorly reconstructed, they may affect the average since the onset of the pulse in this case may be shifted with respect to a general pulse. As such, we define a selection to exclude pulses with OFL times more than 3σ from the mean of the distribution where the mean and σ are once again determined via a gaussian fit.




Figure 3.5: A comparison of the signal template (black dashed line) constructed from the average of 1eh pulses sampled from data taken with a voltage bias of 50 V and an average of 100 pulses from data taken with a voltage biases of -35 V (red) for the inner channel (top left), outer channel (bottom left) and, the sum of the channels (right).

We construct a signal template for each detector channel as well as for the sum of the response

of the channels for each detector. In the latter case, the traces for each channel of the detector are summed before being included in the average. Figure 3.5 shows the resulting signal template for each channel on NFCSO as well as their sum. Pulses were selected from data taken with a high voltage bias of 50 V as opposed to the -35 V, with which ERDM search data was taken. This was because the -35 V data lacked the higher number of statistics required to average out the high frequency noise present in any given trace. However, it can be seen from the right panel of Fig. 3.5 that there is little difference between the pulse shape in the 1eh peak at -35 V and the 1eh peak at 50 V. The signal model and current CSD are then used to calculate the optimal filter.

CHAPTER 4

PHOTON CALIBRATION OF NFCSO

We calibrated the detectors by shining an LED (605 nm) on the center of the detector face with the biasing grid. The LED photon energy was measured to be $2.0468 \text{ eV} \pm 0.0002 \text{ eV}$ using a spectrum analyzer while the LED was immersed in liquid helium—an energy conveniently chosen as it is capable of producing only a single eh pair [44] as this vastly simplifies the detector response model we employ later in the chapter. As the absorption depth in silicon (and similarly in silicon-dioxide) at this wavelength is on the order of a few micrometers [73], photons that reach the detector surface are very likely to be absorbed¹.

The goal while operating the LEDs was to release only a small number of photons ($\lesssim 10$) at a time so that the detectors' response at low energies could be observed. We did this by sending a combination of a fast (1 us) sinusoidal (500 kHz) voltage pulse and a constant “offset” voltage to the LED. The offset voltage was required to keep the LED warm enough to turn the LED on when the fast signal was sent, but low enough so that the LED turned off immediately after the voltage pulse. An external trigger was set to trip on the rising edge of the voltage pulse sent to the LED so that events in the detector induced by photons could be easily identified and separated from background. We sent the sinusoidal pulse to the LED every 100 ms which was chosen to provide a high rate of calibration events without risking pileup between LED pulses (the average pulse fall times were observed to be $< 1 \text{ ms}$.)

The number of photons that interact with the detector after an LED pulse can be described by

¹It is also possible that photons reaching the detector surface are reflected by the superconducting aluminum which makes up the biasing grid. However, the coverage of the biasing grid is only 5%, and thus we expect $\sim 95\%$ of photons that reach the detector surface to be absorbed.

a Poisson distribution with a mean of λ , which scales with the amplitude of the voltage supplied to the LED. Since the probability of absorption is high and the photon energy is such that only one electron-hole pair can be broken per photon, the number of electron-hole pair peaks we see during HV operation as a result of operating the LEDs will be similarly distributed. Therefore, we theoretically know the energy at each electron-hole pair peak, which is given by Eqn. 4.1:

$$E_{n_{\text{eh}}} = n_{\text{eh}}E_{\gamma} + E_{\text{NTL}} \quad (4.1)$$

where $E_{n_{\text{eh}}}$ is the energy of the n^{th} eh peak, E_{γ} is the photon energy, E_{NTL} is the NTL energy and is given by Eqn. 2.20, and n_{eh} is the number of electron hole pairs produced (or, equivalently, the number of photons) in a given interaction. We can then use the eh peaks as points with known energy that we can use to calibrate the current response of the detector. Optimally, the response of the detector within our energy region of interest is linear. However, we can account for non-linearity by including higher order terms in our calibration. For our purposes, we find it is sufficient to include only a quadratic term and thus we calibrate using Eqn. 4.2 where E is the reconstructed energy of an event in units of eV, $c_{1,2}$ are calibration coefficients obtained through fitting and A_{OF} is the current response of the detector as measured by the optimal filter (either OF0 or OFL, which were defined in Chapter 3).

$$E = c_1 A_{\text{OF}} + c_2 A_{\text{OF}}^2 \quad (4.2)$$

In reality, the energy spectrum resulting from operating the LEDs is far more complicated than the idealization described above. As we already discussed in Chapter 2, free charges produced in an interaction do not always neatly propagate through the crystal but can also be trapped or ionize other charges. Additionally, simulations of charge-transport through crystals have shown another trapping mechanism related to events occurring near the surface, which is the case for low energy

photon events (see Figure 4.1): assuming that free charges produced in an interaction initially have random trajectories then there is a chance that both charges will have initial velocities directed towards the nearest detector surface. If voltage bias is low enough such that the acceleration due to the electric field is not strong enough to reverse the direction of the appropriately charged particle, then both charges will recombine on the surface resulting in an event where little to no NTL energy is released. This complicates the calibration since, for example, a 2eh event where one charge pair is surface-trapped may appear as a 1eh event with a small amount of additional energy from the photon.




Figure 4.1: Illustration of a surface-trapped event (left) compared to a "normal" eh pair events (right) as shown by observing particle tracks in simulations using G4CMP. The path of the electron is shown as the solid blue line while the path of the hole is shown as the dotted red line. Diagram taken from Ref. [54].

In the following section, we will show that we can account for these effects by modeling the LED amplitude spectra are using a detector response model, which includes the effects of charge-trapping, impact-ionization, and surface-trapping. Through this calibration, we find, for the first time, evidence of an energy loss mechanism that wherein the full photon energy does not enter the crystal substrate! In this section, we will describe how these effects are modeled and taken into

account during the calibration of NFCSO. We will end with a discussion of the rather unexpected results in which we find that we do not recover the full photon energy for the majority of events.

4.1 Calibration Strategy

To account for the complicated effects of charge-trapping, impact-ionization, and surface-trapping, we used a full detector response model to fit the uncalibrated spectrum. The model calculates the expected probability distribution for the energy of a single photon event, including effects of charge-trapping, impact-ionization, and surface-trapping. To calculate the probability distribution for the energy in the substrate when multiple photons are absorbed, the single-photon PDF is convolved with itself a number of times equal to the number of photons absorbed. Finally, to account for the detector resolution, the resulting PDF, after convolutions, is convolved with a gaussian with a standard deviation equal to the resolution. For more details, I refer the reader to the Refs. [54], [74]. Thus, using this model, we not only extracted calibration parameters, but the probabilities of charge-trapping, impact-ionization, and surface-trapping, as well as the detector's resolution. As we will see in Section 6.4, these parameters are used to inform signal models in the ERDM search. The full list of parameters included in the model are summarized in Table 4.1. For fitting, we use a Monte-Carlo-Markov Chain (MCMC) to explore the large parameter space.

The parameter space shown in Table 4.1 is not only large, but highly correlated—many parameters have similar or partially degenerate effects on the model. As an example, consider the parameters λ and the surface-trapping probability, p_{ST} : When p_{ST} is increased, the distribution of eh peaks in the detector response model will behave as though λ is decreasing since it becomes less likely that eh pairs produced in an event lead to NTL amplification. To reduce the correlations, we simultaneously fit four LED spectra each taken with different amplitude voltage signals being sent to the LED but otherwise with no changes to detector operation. We expect that these data sets will have different values for λ since the number of photons being emitted by the LED scales with

Fitting Parameters in the Detector Response Model		
Parameter	Description	Shared? (Y/N)
N	Normalization	N
λ	Poisson parameter describing distribution of eh peaks	N
c_1	Linear calibration coefficient	Y
c_2	Quadratic calibration coefficient	Y
p_{CT}	Fractional charge-trapping probability	Y
p_{II}	Fractional impact-ionization probability	Y
p_{ST}	Fractional surface-trapping probability	Y
σ_{ph}	Phonon energy resolution	Y
$E_{\gamma,obs}$	Collected photon energy	Y
CTA	Crosstalk amplitude (see Section 4.2)	Y

Table 4.1: A summary of the fitting parameters included in the detector response model. The first and second columns give their mathematical notation and their descriptions (respectively). The third column indicates whether the parameters are independent in the fit of each data set ("N") or rather optimized for all data sets simultaneously ("Y").

the amplitude of the voltage signal. However, parameters related to charge transport within the detector or the response of the detector should be unaffected by the operation of the LED and are thus constrained to be equal across all data sets included in the fit (as indicated by the "Y" in the third column of Table 4.1). By requiring a single set of detector response parameters to fit multiple data sets simultaneously, we heavily constrain the parameter space and thus reduce the effect of correlations on the uncertainties of our calibration.

The external triggering scheme for LED data allows us to model the entire LED spectrum down to ~ 0 eV which can also help constrain the parameter space. To take advantage of this, we calibrate using OF0 as opposed to OFL. In this case, the external trigger is already well-aligned with the onset of the LED-induced phonon pulse, thus for pulses above the detector's threshold, there is a negligible difference in OF0 and OFL. But for energies near or below the detectors threshold, where fluctuations in the noise are the same order as the amplitude of any pulses, the OFL is unable

to effectively constrain the trigger point and will be heavily biased by noise as a result. OF0 does not suffer this same problem since it makes no attempt to optimize the trigger time. Since the trigger time is already well set by the external trigger, OF0 provides a better amplitude estimation at low energies.

4.2 Electrical Crosstalk

As was the case in HVeV Run 4 (see Chapter 3 of Ref. [70]), we observed a crosstalk signal associated with operating the LEDs, which introduced a systematic uncertainty in our calibration. In order to understand the effect that the crosstalk would have on LED spectra, we measured the response of NFCSO while operating LEDs shining on other detectors in the payload—this way, we could observe the crosstalk in the absence of any actual phonon signal. The shape of the crosstalk signal shown in Figure 4.2 is calculated from the average of 1000 external triggers resulting from operating the LED shining on NFHSO. Regardless of which LED was operated (except for the




Figure 4.2: The average (including 1000 traces) response of NFCSO when operating the LED shining on NFHSO shown in black. The signal template is shown in the dashed blue line for comparison. Both the signal template and the average response have been normalized.

LED focused on NFCSO), we consistently observed a similar response characterized by a fast

O(~ 5 us) “spike” at the trigger point followed by a smaller pulse-like component. We found that the average OF0 measured for events such as these was $> 0 \mu\text{A}$ and, therefore, if the same signal exists while operating the LED shining on NFCSO, the presence of crosstalk could artificially shift the uncalibrated spectrum from LED events. We will define this shift, or the amplitude estimation returned by the OF0 for crosstalk, as the “crosstalk amplitude” or “CTA” for short.

To account for the crosstalk, we would like to include CTA as a fitting parameter in the MCMC. In order to do this, we must first understand whether or not the CTA is constant across all the data sets included in the fitting, which would be the case if CTA does not significantly change with the amplitude of the signal. Assuming that the dependence of CTA on the voltage amplitude does not significantly change from LED to LED, we can measure the mean OF0 on NFCSO when shining LEDs on other detectors. Figure 4.3 shows the measured CTA as a function of the amplitude of the voltage pulse sent to the LED (referred to as “LED Amplitude”). The values of OF0 shown in




Figure 4.3: The CTA measured on NFCSO while shining LEDs on other detectors as a function of LED amplitude. The vertical error bars represent the standard error of the mean.

Figure 4.3 are extracted from a gaussian fit to the OF0 spectrum on NFCSO for external triggers coincident with the LED signal. A line is fit to each set of data gathered by shining on a particular detector where the slope of each line shows how the CTA depends on the voltage amplitude sent to the LED. Figure 4.3 does not show the LED shining on NFC2 due to the odd and unexplained behavior of the LED. We found, for example, that the λ describing the distribution of eh peaks in LED calibration data did not scale in a coherent way with the amplitude of the signal sent to the LED. This was not the case with other detectors and is suspected to be related to troubleshooting with the electronics required to operate the LEDs during the experiment. As such, we exclude it from the discussion. For the remaining detectors, we find that the slope of the line describing CTA vs. LED amplitude are all consistent with 0 nA/V. Taking the weighted average of the slopes yields a value of $-0.27 \text{ nA/V} \pm 0.18 \text{ nA/V}$ which corresponds to a difference of $\sim 0.04 \text{ nA} \pm 0.03 \text{ nA}$ ($\sim 0.01 \text{ eV} \pm 0.01 \text{ eV}$) across the calibration data sets for NFCSO. This small change is deemed negligible and so the CTA is modeled as a parameter that is shared among data sets.

4.2.1 *An Independent Measurement of CTA*

After establishing that the CTA can be treated as constant across all data sets, we attempt to constrain the exact value of CTA with an independent measurement in order to help the MCMC converge. Measuring the value of CTA while shining on NFCSO is more challenging since it is not possible to isolate the crosstalk from genuine phonon signals. However, we can indirectly measure CTA by examining the mean of the 0th eh peak—the peak near $\sim 0 \text{ uA}$ —as a function of the LED amplitude.

In the idealized scenario described at the beginning of this section (where we can neglect charge-trapping, impact-ionization, and surface-trapping), every photon that is absorbed by the detector produces an eh pair and the resulting charges propagate throughout the length of the crystal yielding energy via NTL. Thus, the 0th peak would entirely consist of events where zero photons

were absorbed in the substrate. In such a case, the mean of the 0th peak would be exactly the value of the CTA. However, real data are complicated by charge-trapping and surface-trapping, through which the 0th peak receives contributions from events with non-zero energy. This is because, in the case where either charge-trapping or surface-trapping occurs, the NTL energy in any particular event can be as low as ~ 0 eV (for surface-trapping, this is always true if all free charges are surface-trapped).

Let us consider how these contributions to the 0th peak depend on λ . Starting with surface-trapping, we can show that the mean energy, from events in which all eh pairs that are produced undergo surface-trapping, scales linearly with λ .

Consider first that when the LED is turned on, the probability to get N photons interacting with the substrate is given by a Poisson distribution:

$$P(N) = \frac{\lambda^N e^{-\lambda}}{N!} \quad (4.3)$$

Then, given N photons interacting with the substrate, the probability of m of the resulting eh pairs becoming surface-trapped is given by a binomial distribution:

$$P(m|N) = \binom{N}{m} p_{\text{ST}}^m (1 - p_{\text{ST}})^{N-m} \quad (4.4)$$

We would like to find the distribution of surface trapped events contributing to the 0th peak. These are the events in which all the produced eh pairs are surface-trapped. Thus, we would like to find the distribution of N given that all N eh pairs underwent surface-trapping; when $m = N$. We can calculate this distribution using Bayes' Theorem:

$$P(N|m = N) = \frac{P(m = N|N)P(N)}{P(m = N)} \quad (4.5)$$

$$P(N|m = N) = \frac{p_{ST}^N \frac{\lambda^N e^{-\lambda}}{N!}}{P(m = N)} \quad (4.6)$$

The probability of getting m surface-trapped eh pairs, $P(m)$, is difficult to calculate directly. Instead, however, we can take advantage of the fact that Eqn. 4.6 must equal 1 when summed over all N to calculate $P(m = N)$.

$$\sum_{N=0}^{\infty} P(N|m = N) = \sum_{N=0}^{\infty} \frac{p_{ST}^N \frac{\lambda^N e^{-\lambda}}{N!}}{P(m = N)} = 1 \quad (4.7)$$

$$P(m = N) = \sum_{N=0}^{\infty} \frac{(p_{ST}\lambda)^N e^{-\lambda}}{N!} \quad (4.8)$$

Rearranging terms on the RHS, we get

The sum in Eqn. 4.8 is precisely the Taylor expansion for an exponential with argument λp_{ST} multiplied by $e^{-\lambda}$. Thus, we have that $P(m = N) = e^{-\lambda} e^{\lambda p_{ST}}$. We can then rewrite Eqn. 4.6:

$$P(N|m = N) = \frac{(p_{ST}\lambda)^N}{N!} e^{-\lambda p_{ST}} \quad (4.9)$$

Note that Eqn. 4.9 is just a Poisson distribution with a mean given by λp_{ST} . Our result can be interpreted more generally: for N trials, where N is Poisson distributed with mean λ , and for each trial there is a probability, p , of success, the number of successes is also Poisson distributed with a mean λp . We will use this result later in the section. For now, we observe that for an event with n surface-trapped eh pairs, the energy in the substrate is equal to the energy of the photon, nE_γ . Thus, we can calculate the expectation value for the energy from these events by multiplying the expectation value of N , which we can calculate from Eqn. 4.9, by E_γ :

$$E[E_\gamma N] = E_\gamma \sum_{N=0}^{\infty} N \frac{(p_{ST}\lambda)^N}{N!} e^{-\lambda p_{ST}} \quad (4.10)$$

$$= E_\gamma \lambda p_{ST} \quad (4.11)$$

From Eqn. 4.11, we see that the mean of events from surface-trapping that contribute to the 0th peak scales linearly with λ .

We find that the same is true for charge-trapped events though, rather than go through the process of deriving this, we can make the simple observation that for each charge-trapped eh pair, there is some probability that the charges are trapped early enough so that the NTL energy received is ~ 0 eV, $p_{CT,NTL=0}$. This probability would play the exact same role as p_{ST} in the derivation above thus we see that the mean energy of the 0th peak, $\overline{E_{0th}}$, can be written as a function of λ :

$$\overline{E_{0th}}(\lambda) = E_{CTA} + \lambda E_\gamma (p_{ST} + p_{CT,NTL=0}) \quad (4.12)$$

where E_{CTA} is the crosstalk amplitude in units of eV. From Eqn. 4.12, we see that the mean of the 0th peak follows a linear relationship with λ and that the y-intercept of that relationship is E_{CTA} . Therefore, we can measure E_{CTA} by fitting a line to the measured mean of the 0th peak and λ ! In practice, λ can be varied by changing the amplitude of the voltage signal sent to the LED. We can then measure λ by measuring the relative heights of the eh peaks and fitting the result to a Poisson distribution. The result of a simple measurement of this nature is shown in Fig. 4.4. The reader may note that I distinguish λ_m , the mean number of photons absorbed in the substrate as measured by the relative heights of eh peaks, from λ which can be thought of as the “true” mean number of photons interacting with the substrate. Once again, in an ideal case with no surface-trapping and no charge-trapping, λ_m would be expected to be consistent with λ up to some statistical error. However, surface-trapping, charge-trapping, and impact-ionization can muddy the waters of this and, depending on the strength of these effects, can lead to a systematic difference between λ_m and λ .




Figure 4.4: Data showing the mean of the 0th peak in LED data plotted as a function of the measured λ , λ_m is shown in red where error bars represent statistical uncertainties. A linear fit is shown by the dashed orange line.

From the result given in Eqn. 4.12, we can deduce that, under surface-trapping, the eh peaks follow a Poisson distribution with a mean $\lambda(1 - p_{ST})$ since for every eh pair there is a probability of $(1 - p_{ST})$ that the freed charges won't be surface-trapped. In this case, λ_m and λ only differ by a multiplicative constant, so the line shown in Fig. 4.4 will only change in slope but the y-intercept will stay the same; thus, our estimation of CTA is unaffected.

However, the presence of charge-trapping complicates this somewhat. In the case of surface-trapping, we have been assuming that E_γ is small compared to the NTL energy in which case an event with N total eh pairs broken, m of which are surface-trapped, can be considered an event in the $(N-m)$ peak. For charge-trapping this is not the case—the NTL energy lost from a charge-trapped event varies continuously between 0 eV and $e \times V$ where V is the voltage bias over the detector. In this case, events with N eh pairs may fall into any of the eh peaks between 0eh and N eh. Or, they may not fall into an eh peak at all! In this case, λ and λ_m do not vary by a

simple multiplicative constant, and thus our estimation of CTA shown in Fig. 4.4 may be affected. To estimate the effect this has on the relationship between λ and λ_m we can use data simulated from detector response models with various CT, II and surface-trapping probabilities. From these simulations, we can see how λ_m (referred to as “measured lambda” in Fig. 4.5) differs from λ (referred to as “true lambda” in the aforementioned figure). Fig. 4.5 shows us that our estimate for




Figure 4.5: The ratio of λ_m (referred to as “meas lambda” in the axis labels) to λ (referred to as “true lambda”) in the case of a charge-trapping probability of 0 (left) and 0.4 (right). In both sub-figures, the dashed lines correspond to the value of λ_m/λ when $\lambda = 0.5$ while the dotted points correspond to values measured from simulations. Different colors correspond to different surface-trapping probabilities. It can be seen that as the probability of charge-trapping increases, λ_m/λ begins to depend on λ . Credit: Matthew Wilson.

CTA depends on the charge-trapping probability.² However, simulations also show that this effect is rather small. Thus, in the fitting we constrain the CTA measurement with a prior corresponding to a gaussian with mean given by the y-intercept in the linear fit shown in Fig. 4.4 with a width given by the corresponding uncertainty. We will revisit this assumption in the next sub-section where we discuss the results of the fitting.

²Technically, the CTA estimate also depends on the impact-ionization probability, p_{II} . However, previous observations of the HVeV detector response shown in Ref. [54] found p_{II} to be small and therefore it is most likely negligible in the effects we are considering. Despite this, we still take the impact-ionization is still taken into account when we revisit the CTA estimate in the next section.

4.3 Fitting and Results

Results from the Detector Response Model Fit	
Parameter	Result
c_1	$339.3 \frac{\text{eV}}{\mu\text{A}} \pm 0.2 \frac{\text{eV}}{\mu\text{A}}$
c_2	$55.9 \frac{\text{eV}}{\mu\text{A}^2} \pm 0.4 \frac{\text{eV}}{\mu\text{A}^2}$
p_{CT}	$27.2\% \pm 0.3\%$
p_{II}	$1.2\% \pm 0.4\%$
p_{ST}	$21.9\% \pm 0.6\%$
σ_{ph}	$589 \text{ meV} \pm 5 \text{ meV}$
$E_{\gamma,obs}$	$1.13 \text{ eV} \pm 0.02 \text{ eV}$
CTA	$0.41 \text{ nA} \pm 0.02 \text{ nA}$

Table 4.2: The results of the detector response model fit to data. The values provided are the mean of the marginalized posterior distribution after reaching a stationary distribution. The uncertainties are given by the standard deviation of the marginalized posterior distribution, which were found to be consistent with Gaussian distributions.

Using the detector response model, we define a binned Poisson log-likelihood. To fit the four data sets simultaneously, we will simply add the Poisson log-likelihood calculated for each data set. We then include a prior which constrains the CTA estimate when calculating the log posterior. To fully map this complicated parameter space, I use an MCMC to explore the posterior. The resulting fits to the LED spectra are shown in Fig. 4.6. The values of each parameter included in the fit and their corresponding uncertainty were taken to be the mean and standard deviation (respectively) of the sample of the marginalized posterior after the MCMC chain had reached a stationary distribution. These values are shown in Table 4.2. The first thing to point out is that the fit value for the CTA lies near 1σ from the value provided by the prior and thus could be taken to be in agreement. However, the prior that was provided was based on an estimate of the CTA which did not account for the systematic effect of charge-trapping, impact-ionization, or surface-trapping on λ_m in Fig. 4.4. If we simulate data using the best fit values for the p_{CT} , p_{II} , and p_{ST} , we can




Figure 4.6: The uncalibrated spectra from LED calibration data are shown in black. The best fit detector response models are shown in red.

estimate the value of CTA while accounting for these systematic effects. Figure 4.7 shows the result of the simulation from which we estimate $\text{CTA} = 0.32 \text{ nA} \pm 0.05 \text{ nA}$ corresponding to a $\sim 1.8\sigma$ difference compared to the best fit value, which is still consistent.

The more surprising result from the fitting is that the collected photon energy, $E_{\gamma, \text{obs}}$, is $\sim 0.9 \text{ eV}$ ($\sim 46\sigma$) lower than the measured photon energy emitted by the LED! From the measurement obtained with the spectrum analyzer, we know the LED is emitting photons with an energy of $\sim 2 \text{ eV}$ so the fact that the MCMC prefers a lower value does not suggest that the photon wavelength is different from expected but, rather, that we do not recover the recoil energy that we expect from a photon interaction with the crystal. To understand this effect, let us first understand how the




Figure 4.7: Data showing the mean of the 0th peak in LED data plotted as a function of the measured λ , λ_m is shown in red with a corresponding linear fit shown in orange. The same data but with λ_m , corrected for the effects of charge-trapping, impact-ionization (CTII), and surface-trapping (ST) is shown in blue with a corresponding linear fit shown in cyan.

collected photon energy, $E_{\gamma, \text{obs}}$, is constrained by the model.

Back to our idealized scenario where we have no surface-trapping, we know that eh peaks should appear to be gaussian and that the energy of those peaks will be given by Eqn. 4.1. In this case, there is no way for us to know E_γ or E_{NTL} from the OF spectrum alone—we would need an independent measurement of E_γ . However, as can be seen in Figure 3 of [54], when the surface-trapping probability is non-zero, the eh peaks are not perfectly gaussian but instead have an underlying Poisson distribution consisting of "sub-peaks" which, when calibrated, are separated by $E_{\gamma, \text{obs}}$. Figure 3 of [54] shows that when the detector's energy resolution is $\gtrsim E_{\gamma, \text{obs}}$, the individual sub-peaks are not visible and instead the effects of surface-trapping manifest by both broadening and shifting the eh peaks. However, when the detector resolution is below $E_{\gamma, \text{obs}}$, the underlying Poisson distribution becomes visible and the shape of the eh peaks is strongly dependent on λ and $E_{\gamma, \text{obs}}$ (see Fig. 4.8). Thus, assuming that the value of λ can be constrained by the distribution of the eh peaks in a given data set, the value of $E_{\gamma, \text{obs}}$ can be constrained by fitting the shape of the




Figure 4.8: The 1eh peak from LED calibration data is shown in black along with the fit returned by the MCMC (red) which finds $E_{\gamma,obs} = 1.23$ eV. For comparison, I show the detector response model which has $E_{\gamma,obs} = E_{\gamma}$ (the calibration parameters in this model have been tuned to align the 1eh peak to data though the other parameters are unchanged compared to the fit returned by the MCMC.)

eh peaks. As can be seen in Figure 4.8, the eh peaks in our data have a shape that is incompatible with $E_{\gamma,obs} = E_{\gamma}$.

4.4 Another Measurement of Energy Loss

On its own, the result (shown in the last subsection) is concerning and leads us to ask whether this is a physical effect or whether we neglected some systematic effect which then led to this discrepancy. As is proper in the event of an unexpected result, we attempted to reproduce the results in another measurement. In this case, we performed a separate calibration using an identical detector (though not the same one operated at CUTE) at the SLAC National Accelerator Laboratory. In many ways, the set-up at SLAC was similar to that at CUTE; in both cases the detector was mounted in a copper housing placed in a light-tight can, thermalized to the mixing chamber of a

dilution-unit. Both facilities used a SQUID-based readout which is digitized using a SuperCDMS DCRC.




Figure 4.9: **Left** The SLAC mounting scheme for a single device of the NFC SO design clamped in a copper housing. **Right** The mounting scheme used at CUTE for two devices (NFC SO design on the left and NFHSO on the right) clamped in a similar housing. Both detector housings have an interlocking ridge to create an improved light-tight connection to a lid or other housing (in the case of a stacked payload geometry). The Si detectors shown are 1 cm x 1 cm x 4 mm thick.

The setup at SLAC, like the one at CUTE, used LEDs collimated on the center of the detector face with the Al biasing grid for calibration. At SLAC, however, the LEDs were operated outside of the fridge and connected to an optical fiber which directed photons through the fridge and to the detector [69]. Using this setup, the TES and SQUID readout lines can be separated entirely from the lines used to operate the LEDs, thus avoiding any crosstalk signal. The LED photons used in the calibration at SLAC are 3.4 eV and were operated in a similar fashion to that at CUTE.

Just as with the calibration at CUTE, the analysis of SLAC calibration data used an MCMC to fit the detector response model to the LED calibration spectrum. In this case, only a single data

set was available for the calibration which included far less statistics than what was used at CUTE. However, without crosstalk, the parameter space is less complicated for the model fit with SLAC data. The resulting fit to the LED spectrum is shown in Fig. 4.10. The model parameters returned by the MCMC are shown in Table 4.3, although, here, I have chosen to show the best fit photon energy in terms of the difference from the expectation defined as $E_{\gamma,\text{loss}} = E_{\gamma} - E_{\gamma,\text{obs}}$ in order to better facilitate comparisons between the two measurements.

Calibration Results from CUTE & SLAC		
Parameter	CUTE	SLAC
c_1	$339.3 \frac{\text{eV}}{\mu\text{A}} \pm 0.2 \frac{\text{eV}}{\mu\text{A}}$	$290^{+2}_{-4} \frac{\text{eV}}{\mu\text{A}}$
c_2	$55.9 \frac{\text{eV}}{\mu\text{A}^2} \pm 0.4 \frac{\text{eV}}{\mu\text{A}^2}$	$27^{+5}_{-4} \frac{\text{eV}}{\mu\text{A}^2}$
p_{CT}	$27.2\% \pm 0.3\%$	$10^{+2}_{-3}\%$
p_{II}	$1.2\% \pm 0.4\%$	$7^{+3}_{-2}\%$
p_{ST}	$21.9\% \pm 0.6\%$	$29^{+1}_{-2}\%$
σ_{ph}	$589 \text{ meV} \pm 5 \text{ meV}$	$900^{+22}_{-2} \text{ meV}$
$E_{\gamma,\text{loss}}$	$0.92 \text{ eV} \pm 0.02 \text{ eV}$	$0.8^{+0.2}_{-0.1} \text{ eV}$

Table 4.3: The results of the detector response model fit to data taken at SLAC (HV = 20 V with 3.4 eV photons). To facilitate comparisons, the results from CUTE (HV = -35 V with 2.0468 eV photons) are reprinted here.

The differences in the charge-trapping, impact-ionization, and surface-trapping probabilities shown in Table 4.3 are most likely due to differences in HV bias—SLAC data was taken at 20 V whereas CUTE data was taken at -35 V. The reader may also note that the resolution of the detector operated at SLAC significantly differs from the measurement at CUTE; this was found to most likely be due to vibrational coupling of the pulse-tube to the mixing chamber-stage of the SLAC cryostat [69].

Figure 4.11 shows the marginalized posterior from the MCMC chains used in the calibration for LED spectra taken at CUTE and SLAC. At both facilities, model fits require non-zero values for $E_{\gamma,\text{loss}}$ in order to describe the LED calibration spectra suggesting that, on average, the bandgap




Figure 4.10: LED data taken at SLAC with an HV bias of 20 V using 3.4 eV LED photons is shown in black with corresponding detector response model fit shown in red. The shaded blue region indicates the spread of random triggers consistent with noise to provide a scale of the detector resolution. Note that the “jagged” eh peaks predicted by the detector response model are consistent with the expectation since the detector resolution of 900 meV is much lower than the collected photon energy of ~ 2.6 eV.

energy is not fully recovered from a photon-absorption event. It is also worth noting that the values for $E_{\gamma, \text{loss}}$ between the facilities are in agreement which suggests that the mechanism through which energy is lost does not depend on the wavelength of the photon. At the time of writing this thesis, the origin of this energy loss is unknown and future work is needed to understand its behavior.

4.5 Implications of Energy Loss on Previous Calibrations of HVeV Detectors

Before concluding the chapter, we will note that the energy loss found in the LED calibrations at both CUTE and SLAC offer a possible resolution to previous results from HVeV Run 4, which found the calibration differed depending on whether the detector was operated with or without an HV bias. In the HV calibration, a 100 V bias was applied to the detectors and LED photons (2.05 eV) were collimated on the side of the detector with the biasing grid just as what was done in the HVeV Run 5 photon calibration described earlier in this chapter [63]. The calibration at 0 V used a very active source of gammas, Cs-137, which emits photons at ~ 662 keV. The resulting




Figure 4.11: Samples from the posterior distributions of event energy loss per photon using LED data from SLAC (red) and CUTE (blue). The corresponding shaded regions represent ± 1 standard deviation from the mean. The posterior distribution from CUTE results from the simultaneous fit to four data sets whereas the distribution from SLAC results from a fit to a single data set. The gray band shows the band-gap of silicon [44]. The narrowness of the distribution from CUTE compared to SLAC is a result of using four data sets to constrain the parameter space. The distributions show consistent results between data taken at CUTE and SLAC. Figure taken from Ref. [69].

Compton spectrum is expected to have “steps” corresponding to the binding energy of electrons in various shells of the silicon (*i.e.* ~ 100 eV for the $L_{2,3}$ shells, 150 eV for the L_1 shell, and 1.53 keV for the K shell) [72]. The results of both calibrations for NFC1, the detector used in the ERDM search during HVeV Run 4, are shown in Table 4.4. From both calibrations, we can see that the quadratic term in the calibration function is small suggesting that the detector is in its linear regime up to high energies. We can also see that between the two calibrations, the measured values of the dominant linear calibration term disagree by $\sim 30\%$! This is a surprising result; we do not expect that the detector’s response to a given energy should depend on the voltage bias we apply over the substrate. However, the resolution of NFC1 in HVeV Run 4 was ~ 2 eV—too large to resolve the

NFC1 Calibration Results from HVeV Run 4		
Parameter	HV Calibration	0V Calibration
c_1 [keV/ μ A]	0.74 ± 0.01	1.06 ± 0.05
c_2 [keV/ μ A ²]	0.014 ± 0.006	-0.01 ± 0.03

Table 4.4: Results of the two calibrations performed on NFC1 during HVeV Run 4. We find that dominant term in the calibrations, the linear term c_1 , the two calibration methods have a $\sim 30\%$ disagreement.

underlying structure from surface-trapping in eh peaks in LED spectra which revealed the energy loss mechanism observed in HVeV Run 5. As such, neither calibration in HVeV Run 4 accounted for energy loss. Let us see how these calibrations might change if we assume that for every eh pair produced in the HVeV Run 4 calibrations, about 0.9 eV was lost (the same amount of energy loss that was observed from the calibrations at SLAC and CUTE).

The HV calibration would be relatively unaffected. Each calibration point they had originally was approximately (neglecting surface-trapping and charge-trapping) given by

$$E_{eh} \sim n_\gamma(E_{\gamma,obs} + eV) \quad (4.13)$$

where $V = 100$ V and the original calibration assumed $E_{\gamma,obs} \sim 2.05$ eV. Therefore, most of the energy at each of the eh peaks comes from the NTL energy so if we assume $E_{\gamma,obs} \sim 1.15$ eV instead, we only expect a $\sim 1\%$ difference in the calibration points. Thus, for now, let us consider that the calibration parameters obtained in the HV calibration for HVeV Run 4 remain relatively unchanged by any energy loss from broken eh pairs.

The 0 V calibration, however, could change significantly if there is energy loss for every broken eh pair. On average for every 3.66 eV of photon energy deposited in the detector's substrate, an eh pair is broken [44]. The 0 V calibration uses Compton scatters from high energy photons and so many eh pairs will be broken in any given event. Let's consider that for every eh pair bro-

ken, 0.9 eV is lost. If that is the case, then at each calibration point, the detector only receives $(3.66 \text{ eV} - 0.9 \text{ eV})/3.66 \text{ eV} \sim 75\%$ of the energy expected. Since the quadratic term in the calibration is small, we can essentially interpret this as a correction to the linear term of the calibration. Thus, accounting for this loss, the c_1 term in the 0V calibration becomes $0.75 \cdot 1.06 \text{ keV}/\mu\text{A} \sim 0.80 \text{ keV}/\mu\text{A}$.

While the HV and 0V calibration still disagree under the assumption, the discrepancy is reduced from $\sim 30\%$ to $\sim 8\%$. That said, however enticing this indication is, this is all highly speculative; the exact mechanism by which energy is lost in the calibrations of NFCSO performed at SLAC and CUTE are unknown. Future work is needed to understand the origin of this energy loss and whether it applies more broadly to HVeV detectors or any detectors using similar architectures. In future chapters we will see how this energy loss impacts the ERDM search using data taken at HVeV@CUTE.

CHAPTER 5

BACKGROUND OBSERVATIONS AT 0V

HVeV detectors are uniquely positioned to study excesses at low energy. Most experiments searching for signals below ~ 1 keV are sensitive to either ionizing signals or non-ionizing signals but not both and thus are limited in the phenomena that they can study. HVeVs, on the other hand, are sensitive to both ionizing signals and non-ionizing signals when operated with an HV bias and are thus able to study the excesses reported by detectors sensitive to ionizing events such as SENSEI [75] and DAMIC [76] as well as the excesses reported by detectors sensing phonons such as TESSERACT [77], EDELWEISS [78], CRESST [79]–[81], and SuperCDMS-CPD [82].¹ However, as we will show, there is some difficulty in disentangling the ionizing components of background signals from the non-ionizing components. For this, we rely on comparisons with background spectra taken at 0 V.

In this chapter, we will characterize the components of the 0 V background spectrum recorded by NFCSO. Understanding the background spectrum measured in detectors included in HVeV Run 4 is left for future work. Included in this description is a model for how the 0 V background rate changes with time as well as modeling of the non-ionizing components for use in the ERDM search.

5.1 The Shape of the 0V Energy Spectrum

I examined the energy spectrum of low background data taken at 0 V in the energy region between 7 eV and 200 eV. Theoretically, our detector resolution of 589 meV would allow us to

¹This is by no means an exhaustive list of experiments that have reported observations of low energy excesses. For a more complete summary of observations of the LEE, refer to the review by D. Baxter *et.al.* [66].

probe energies down to ~ 3 eV without needing to worry about a significant contribution from noise. However, we found that the spectrum below ~ 7 eV changed slightly throughout the run and thus was difficult to model. While the reason for this is still not understood, we think that this may be due to changes in the vibration due to re-aligning the fridge during operation. Since I am more interested in modeling backgrounds which seem unaffected by environmental conditions and since 7 eV is still far from the signal region of interest for ERDM models we consider in Chapter 6, I exclude the lower energy part of the spectrum from the remainder of the analysis. The upper bound is set so that I include the energy region of interest for ERDM models with masses up to a few $\sim \text{GeV}$ while not straying far from the highest calibration point provided by the LED calibration (~ 6 eh or ~ 210 eV with a -35 V bias).




Figure 5.1: The mean of the 1eh peak in 60 V plotted as a function of time into CUTE Run 40 (shown in black). The y-axis is converted from units of μA to eV using the calibration obtained in Chapter 4. Errors shown in the y-axis are the standard deviation of the 1eh peak as calculated via a gaussian fit. The vertical blue dotted line denotes the time at which we took calibration data, after which we found the response of the detector to be stable. The red dashed line shows the linear fit to the data shown.

A large number of statistics are required in order to model the underlying distribution of the 0 V background. We will see that for energies below ~ 40 eV, the event rate is high enough that the shape of the distribution can be constrained with only a day or so of live time. However, the event rate steeply drops off for higher energies so that many hours of live time are required to map the underlying distribution. Therefore, I relied on data taken over many months in order to constrain the shape of the spectrum. Over this time, we found that the current response of the detector drifted slightly, as evidenced by the drift in the mean of the 1eh peak observed in background data taken at 60 V shown in Fig. 5.1. Overall, this drift was found to be small compared to the value observed when calibration data was taken—about a 3% difference. To correct for this shift, the calibration of data taken before calibration data (indicated by the blue vertical line in Figure 5.1) is adjusted according to Eqn. 5.1 on a series by series basis.

$$E_{\text{corrected}} = \frac{m * t + b}{m * T_{\text{cal}} + b} E \quad (5.1)$$

where m and b are the slope and y-intercept of the linear fit shown in Figure 5.1 (respectively), t is the time the given data set was taken, T_{cal} is the time at which the calibration data was taken, and E is the uncorrected calibration.

After correcting for the drift in calibration, I first checked that the shape of the spectrum at 0 V remained constant with time. Throughout CUTE Run 40, we took a number of long exposures at 0 V, which I have combined separately and labeled according to the month in which they were taken in Figure 5.2. The shapes of the spectra are well described by exponentials—a feature that is consistent with observations of the LEE across the field (see the recent review by Baxter et al. [66]). In this case, I fit each spectrum to the sum of two exponentials with decay parameters λ_1 and λ_2 and weights given by C and $(1 - C)$, respectively, where C is a constant between 0




Figure 5.2: Calibrated data taken at 0V in each month of CUTE Run 40. The model (Eqn. 5.2) fit to the shape of the energy spectrum is shown in red with corresponding fit parameters provided in the legend. In the lower panel of each figure, the residuals are shown in units of uncertainty where the uncertainty for a bin is taken to be the square root of the number of counts in the bin divided by the live time (minimum of 1 divided by the live time). The green shaded region indicates ± 1 .

and 1 (see Eqn. 5.2). The best fit parameters for each spectrum are determined by minimizing an unbinned negative log-likelihood.

$$f(E|\lambda_1, \lambda_2, C) = \frac{C}{\lambda_1} e^{-E/\lambda_1} + \frac{1-C}{\lambda_2} e^{-E/\lambda_2} \quad (5.2)$$

Figure 5.2 shows the resulting fits to the partitioned data sets and Figure 5.3 shows the shape parameters and their uncertainties.




Figure 5.3: The shape parameters corresponding to the fits shown in Fig. 5.2 plotted with their corresponding errors.

The results shown in Figure 5.3 show that the 0 V background does not vary significantly with time. This provides us with some confidence that regardless of when 0 V data was taken throughout CUTE Run 40, the same underlying distribution was being measured. We can therefore combine all data taken at 0 V to help limit the statistical uncertainty when modeling the higher energy regions of the spectrum. The fit to the combined spectrum, consisting of ~ 200 hours of live time, is shown in Figure 5.4.

Before moving on to discuss the rate of the 0 V background as a function of time, it is worth




Figure 5.4: The combined energy spectrum of data taken at 0V consisting of ~ 200 hours of live time taken throughout CUTE Run 40. The model fit of Eqn. 5.2 obtained by minimizing the unbinned negative log-likelihood is shown in red. The lower panel shows the residuals between the model fit and the data shown in black in units of uncertainty on bin height (taken as the square root of counts in the bin divided by the live time). The green shaded region indicates ± 1 .

comparing the spectrum shown Fig. 5.4 to observations of the low energy excesses reported across the field. The aforementioned review by D. Baxter et al. [66] shows that other experiments using solid-state detectors to measure phonon signals almost all observe a steep rise in event rate “typically observable below ~ 100 - 200 eV,” which is consistent with our observations. Furthermore, the review states that in all observations, the spectrum from the LEE can not be described by a single exponential or power law but generally requires a combination of multiple, which is also consistent with what we observe.

5.2 The 0V Rate Over Time

While the shape of the 0 V spectrum appears to remain constant over time, the rate does not; similar to observations of the LEE reported across the field, we find the rate measured in 0 V data decreases over time. Figure 5.5 shows the rate measured in 0V data between 7–200 eV plotted




Figure 5.5: The event rate integrated between 7-200 eV measured in data taken at 0 V as a function of time into CUTE R40 (black). Vertical error bars account for both statistical errors and systematic errors related to scaling the calibration. Horizontal errors represent the time over which the rate was measured. In blue is the best fit of Eqn. 5.3 to the data. The shaded region shows the time window wherein the blinded DM search data was taken (see Chapter 6).

against time into CUTE Run 40.

We can effectively model this rate of change using a linear combination of two exponentials with the parameterization shown in Eqn. 5.3.

$$R(t) = N \left(\frac{f}{\tau_1} e^{-t/\tau_1} + \frac{1-f}{\tau_2} e^{-t/\tau_2} \right) \quad (5.3)$$

Here $\tau_{1,2}$ are constants with units of time which characterize the rate of decay of the integrated event rate measured in 0 V data, f is a value between 0 and 1 and determines the relative weighting of the decay constants, and N provides the overall normalization. Table 5.1 gives the values of f and $\tau_{1,2}$ that best fit the data shown in Figure 5.5; the model fit is shown as the blue dashed line in the same figure.

Thus, the 0 V background behaves similarly to what is commonly described as the LEE—

Best Fit Parameters of Eqn 5.3 on 0V Data	
Parameter	Value
τ_1 [days]	5 ± 4
τ_2 [days]	61 ± 13
f [unitless]	0.14 ± 0.04

Table 5.1: Best fit parameters of the normalized rate model (Eqn. 5.3) to 0V data.

the event rate steeply increases at lower energies, and the integrated event rate decays with time. Moreover, as we saw in the last section, the fact that the shape of the spectrum does not change significantly over time suggests that the event rate decays similarly over the energy region between 7 eV and 200 eV. This implies that the 0 V spectrum is largely dominated by events with a common origin (at least in the energy region considered so far in this chapter). It is also worth noting that the model fit to the rate seen in 0 V data predicts a somewhat optimistic long-term behavior of the background rate as time goes on; as $t \rightarrow \infty$, $R(t) \rightarrow 0$. This implies that if we operated the detectors for a time much larger than the characteristic decay times given by $\tau_{1,2}$, the background rate in 0 V data would be negligible. However, from Table 5.1, we can see that the rate model is dominated by the longer decay constant, τ_2 , which is on the order of two months, implying that we would have to operate for close to a year to avoid the background completely. Moreover, CRESST operated a silicon device for longer than a year and still found a non-negligible excess [83]. Multiple experiments have also found that the LEE rate “resets” after warming their detector payloads above a few tens of Kelvin [66].

5.3 Comparison with HV Spectra

At this point in the chapter, we have shown that our observations of the 0 V background match the features commonly observed of the LEE. In this subsection, we will demonstrate how our understanding of the 0 V background spectrum is relevant to the HVeV Run 5 ERDM search.

In particular, comparisons between background spectra taken at 0 V and at HV can be used to distinguish non-ionizing and ionizing components of the spectrum. Since we aim to search for DM models that interact via electrons in the substrate, non-ionizing components of the LEE can be modeled as backgrounds in an ERDM search. Thus, by characterizing the non-ionizing components of the spectrum, we can possibly increase our sensitivity to ERDM. We will do this by modeling HV spectra with a non-ionizing and an ionizing component.

For the ionizing component, we can construct a model based on the leakage of charge carriers throughout the substrate. This is motivated by the fact that the $1eh$ peak observed in background data typically has a mean consistent with the NTL energy expected from a single eh pair (where $E_R = 0$ eV). “Leakage” refers to the spontaneous excitation of a charge carrier bound by an impurity in the crystal by the applied electric field. We consider both leakage near the surface or leakage in the bulk. In the former case, the charge is excited near the surface and is drifted throughout the entire thickness of the crystal via the electric field and thus mimics a bulk eh event with 0 recoil energy. In the latter case, the charge is excited in the crystal bulk and thus has an energy distributed between 0 eV and the NTL energy for a single eh pair. For this case, we assume that charges are excited uniformly throughout the crystal. Figure 5.6 shows the expected energy spectrum of both bulk and surface leakage (including effects of charge-trapping and impact-ionization).

To model the non-ionizing component, let us consider the 0V spectrum shown in Fig. 5.4. The parts of this spectrum that are non-ionizing will retain the same shape regardless of the HV bias that is applied. However, we already have reason to believe that the vast majority of the 0 V spectrum within the region of interest is non-ionizing simply by comparing to HV spectra. To demonstrate, let’s consider the possibility that the HV spectrum consists of some electron recoils (ERs). For an ER with initial recoil energy E_R , the average number of eh pairs produced in an interaction is given




Figure 5.6: An example of an energy spectrum (in units of n_{eh}) expected from bulk (surface) leakage shown in red (blue).

by

$$\langle n_{eh} \rangle = Y \frac{E_R}{\epsilon} \quad (5.4)$$

where ϵ is the average energy required to produce an electron-hole pair in silicon and has been measured to be ~ 3.6 eV [44]. Y is the ionization yield which, for ERs, is 1. From Eqn. 5.4 we see that, when considering the lower bound of the energy region of interest, 7 eV, the average ER would produce 2 eh pairs. However, in HV data, we seldom observe events above 2 eh. Figure 5.7 shows what the 0 V spectrum would look like under the application of a 60 V bias if all of the events were ERs. Comparing this extreme scenario with our observations of the spectrum at 60 V, we can place an upper bound on the percentage of the 0 V spectrum which could consist of ERs. In this case, we examine the energy region between 127 eV (corresponding to the 2eh peak plus the minimum recoil energy from an event in our region of interest) and the upper bound of our




Figure 5.7: A comparison of the 0 V background model (blue) with data taken at 60 V. The red dashed line shows the 0 V background under the assumption that the entire 0 V background consists of ERs. In orange is the 0V background under the assumption that only 2.4% of the events are ERs which corresponds to the 90% CL on the percentage of ERs in the 0 V background based on comparisons to ~ 10 hours of data taken at 60 V (black). The shaded region shows the energy range over which the 90% CL was calculated.

region of interest, 200 eV. We find that at 90% confidence, ERs can make up at most $\sim 3\%$ of the 0 V spectrum. This represents an upper bound because we consider no other origins for the events above 2eh. This is significant because from this, we can deduce that the majority of events in the 0 V spectrum are non-ionizing and therefore that a non-ionizing component can be modeled using the shape of the 0 V spectrum.

Using a combination of the models shown in Figure 5.6 and the shape of the 0 V spectrum, we attempt to fit the spectra found in HV data at different times throughout the run. The parameters in the ionizing component (p_{CT} , p_{II} , the relative heights of eh peaks, the fractional contribution

from bulk [as opposed to surface] leakage, and overall normalization) are left unconstrained. The locations of the eh peaks, however, are constrained to be within 5 eV of the expected NTL energy. For the non-ionizing component, the shape is fixed to the model of the 0 V background shown in Figure 5.4 though the normalization is left unconstrained. We specifically chose to fit data taken with a 60 V bias since the voltage bias is high enough that we expect less degeneracy between the ionizing and non-ionizing components of the model.




Figure 5.8: A fit of the detector response model (red) to a background spectrum at 60 V (grey). The non-ionizing (ionizing) component of the model is shown in blue (yellow).

At lower voltages, the ionizing component is "compressed" so that the ionizing background resides in the same region where the 0 V rate is greatest, and thus the model has less power in differentiating its components. An example fit is shown in Figure 5.8.

We consistently find that the ionizing model dominates in the energy region beginning close

to the 1eh peak through the 2eh peak whereas the non-ionizing component dominates for energies much lower than the 1eh peak energy or for energies higher than the 2eh peak energy. In fact, we find that for energies below ~ 30 eV, the non-ionizing component contributes the vast majority of events. We can extract the rate of the non-ionizing component for each data set from the best fit normalization which is shown as a function of time in Figure 5.9. As can be seen from Figure 5.9,




Figure 5.9: The event rate integrated between 10 eV and 200 eV measured in 0V is shown in black. The non-ionizing rate in HV data is calculated from the normalization of the non-ionizing component of the fit to HV spectra (see Figure 5.8) and plotted in red. (Note that the error bars do not account for the systematic uncertainty related to scaling the calibration.) The rate predicted by the rate model with a form given by Eqn. 5.3 and parameters given by Table 5.1 are plotted in blue, with corresponding uncertainties shown. The rate model for 0V data is plotted as the dashed blue line. The shaded region indicates the time window over which the blinded DM search data was taken.

the non-ionizing rate extracted from fits to HV data is in excellent agreement with the rate measured from 0 V data directly. However, this fact should not be too surprising since by comparing the HV and 0 V spectra directly, we could already see that the majority of the 0 V background must not consist of ionizing events. It should be noted that this is also consistent with previous observations of the LEE by EDELWEISS, which found that the LEE was dominated by non-ionizing events down to ~ 25 eV_{ee} [78], and RICOCHET, which similarly found that the LEE was dominated by

non-ionizing events up to energies on the keV scale [84]. Note that this also seems to imply that the overall event rate in the region between 10 eV and 200 eV increases with the application of a high-voltage bias. It is unclear, however, whether the ionizing background is induced by the application of the HV bias or whether there is an excess of ionizing events at energies below our region of interest. The latter hypothesis seems less likely since, as stated in the previous section, we generally find that the 1eh peak observed in background data is consistent with leakage.

In the context of an ERDM search, the implication that the 0 V spectrum primarily consists of non-ionizing events is a boon for sensitivity—events that do not yield ionization can be treated as a background since they do not have the expected phenomenology of ERDM! Therefore, we may use the 0 V background to characterize the contribution of non-ionizing backgrounds to improve our sensitivity to ERDM. This process will be described in the following chapter.

CHAPTER 6

THE HVEV RUN 5 ERDM SEARCH

During the construction and operation of the HVeV Run 5 payload, we had considered the possibility of a DM search, but wanted to prioritize studies of the low energy excesses. In spite of this, while studying the backgrounds in the detectors under different conditions, we observed 1eh rates that were far lower compared to observations in previous HVeV Runs. When we discovered the potential to explore new parameter space, the opportunity was too good to pass up. Thus, in spite of not being optimized for a DM search, we decided to take roughly two weeks of blinded data with which to set exclusion limits on the DM candidates discussed in Chapter 2. This chapter will describe the analysis and the results that followed.

6.1 Data Acquisition and Analysis Strategy

After taking the initial data sets wherein we discovered the low rate in the 1eh peak, we took two weeks of blinded data for the ERDM search. We operated the detectors with a bias of -35 V, far lower than the 100 V bias used for the ERDM searches in the last two HVeV Runs. The reason was that throughout HVeV Run 5, we observed that higher bias voltages seemed to be correlated with higher rates in the 1eh peak.¹ We chose a magnitude of 35 V so that we could keep the bias as low as possible while still positioning the 1eh peak above the region where the LEE dominated. The reason for the negative polarity was purely a matter of circumstance—at the time we were testing the detector response with negative polarities and comparing to data taken with positive polarity.

¹More work is needed to determine whether the apparent decrease in the 1eh rate was truly due to the lower voltage. We have previously observed that the 1eh rate is highly dependent on the biasing history of the detectors and so it is possible that the lower rate was induced by holding the detector at higher voltages for a long period of time rather than simply operating with a lower voltage.

By the time we decided to conduct a DM search, we happened to be biasing with negative polarity and so there was no reason to switch back.

In the two weeks of blinded data, we unblinded one out of every ten minutes of data so that we could monitor operating conditions and check for obvious problems. While these data were not used in setting the final exclusion limits, they were also not used to develop an analysis. Additionally, every 2-3 days, we would pause low background data and take 20 minutes of LED data to confirm that the calibration was stable; these data were also unblinded. To develop the analysis, we used the ~ 24 hours of live time taken at -35 V (henceforth referred to as the “open data”) that we took prior to the decision to conduct a blinded DM search.

6.2 Data Selection

The data selection shown in the following section was designed to exclude events in which the detector’s response differed significantly from the signal template (see Section 3.2) or was inconsistent with the phenomenology expected from DM. These criteria can be separated into two categories: 1) live time cuts in which periods of time are removed from the analysis, and 2) quality cuts in which specific events are removed and the resulting live time is weighted by an overall “efficiency” specific to the nature of the cut.

6.2.1 *Live Time Cuts*

6.2.1.1 Mean Baseline Cut

The “baseline” refers to the detector response in the absence of an energy deposition in the crystal and can be used to monitor changes in temperature or noise performance. As such, we define a criterion based on the mean baseline for each half-second trace. The mean baseline is calculated as the mean of all samples in a given trace that are not included in the analysis window of a triggered pulse. In the absence of any large temperature fluctuations or vibrational coupling to

the detectors, the distribution of mean baselines for a given series can typically be well described as Gaussian (see Fig. 6.1.)




Figure 6.1: An example distribution of mean baselines for a series taken with a -35 V bias. (Note that, in general, the mean baseline does not depend on the HV bias). The Gaussian fit to the distribution is shown in red. The green shaded region corresponds to the region within $\pm 3\sigma$ of the mean, where both the mean and σ are determined via the Gaussian fit.

Our observations of the unblinded data taken throughout the two weeks of exposure showed that operating conditions seemed to remain stable throughout the DM search and thus we expect very little live time to be removed from the mean baseline cut. For example, in the open data set, the mean baseline cut removed less than 1% of the live time.

6.2.1.2 Coincidence Cut

The interaction cross-sections for DM have been bounded such that it is highly improbable for a single DM particle to interact with two different detectors. Therefore, coincident events are not

consistent with the phenomenology expected for DM events and are thus removed. Since NFHSO and NFCSO share an optical cavity and therefore have many open lines of sight between them, we construct a coincidence cut which removes events in the detectors that occur within 0.5 ms of each other. (See Fig. 6.2 for an example of such an event.) We choose not to apply this cut using coincidence with other detectors in the payload because there were no coincident events observed in the open data set, which is consistent with the lack of lines of sight between them. Despite the small amount of live time removed by the cut ($<1\%$), we find that a number of events between ~ 10 -30 eV are removed solely by the coincidence cut.




Figure 6.2: An example of a coincident event found in the open data set. The response of NFCSO (NFHSO) is shown in orange (blue). Credit: Julius Viol

6.2.1.3 Burst Event Cut

The most significant contribution of events that are considered background in the context of an ERDM search are the so-called “burst events”. These events are characterized by a high energy event that typically saturates the detector, called a “primary,” followed by many lower energy

events ($\leq 1\text{h}$) referred to as “secondaries.” Burst events are typically associated with a significantly increased trigger rate, which can persist for as little as a few minutes or up to a few hours (see Fig. 6.3) and, as such, there is little argument against this being the most important data selection criterion with regard to improving sensitivity.




Figure 6.3: The event rate as a function of time in the open data. Regions shaded in red were identified as burst events. Note that a large spike in event rate is often associated with a burst event. The dashed line shows the estimated rate without bursts and the dotted line shows the estimated rate $+ 1\sigma$ where σ is determined via a Poisson fit to the event rate. Credit: Mason Buchanan.

To identify burst events, we first estimate the event rate in their absence. We do this by binning the time between successive triggers (referred to as “dt”) and fitting a Poisson distribution to the resulting histogram. In general, the low dt bins have an abnormally high number of counts due to the presence of bursts and so are excluded from the fitting region (see Fig. 6.4a). We then define the start of a burst event when two criteria are satisfied: 1) We observe a high energy ($>200\text{ eV}$) event and 2) the event rate is higher than 5σ from the average event rate determined by the fit to the dt spectrum shown in Fig. 6.4a. The burst is considered over when the trigger rate falls below one standard deviation above the mean rate as determined by the Poisson fit to the dt spectrum. The




Figure 6.4: **Left** A histogram of dt before applying the burst cut. The presence of bursts inflates the counts in bins with $dt \sim 0$ and are thus excluded when fitting a Poisson to the histogram (shown in red.) **Right** The histogram of dt after applying the first pass of the burst cut using the result of a. Note that the bins with $dt \sim 0$ more closely follow the trend of the Poisson fit shown in red, which demonstrates the effectiveness of the burst cut. Credit: Mason Buchanan.

time windows wherein a burst is identified are then excluded from further analysis. After applying the cut, a second fit to the remaining dt spectrum is performed and is expected to give a lower mean rate which is more representative of the mean rate without contamination from bursts. The results of the second fit to the dt spectrum are then used once more to search for remaining burst events. The effect of the final cut on the open data is shown in Figure 6.5. This cut removes a significant amount of live time—about 55% of the total—but comes with a vast improvement in sensitivity.

Understanding these burst events is of utmost importance for future HVeV runs as they pose a significant hindrance on sensitivity. They have been observed with similar behavior in HVeV Run 4 [63] and have similar characteristics to the burst events seen in HVeV Run 3, though these were found to originate from luminescence in materials close to the detectors [65]. In particular, despite the success of the described cut in removing visible contributions from bursts, there is evidence that the bursts contribute events to the 1eh peak region and below that cannot be removed with 100% efficiency—for more information on this see Appendix A. In light of this, locating the origin of




Figure 6.5: The energy spectrum in open data before (blue) and after (black) the burst cut. The bins are 1 eV wide. Credit: Mason Buchanan.

burst events and developing methods by which they can be prevented is necessary for future DM searches using HVeV detectors.

6.2.1.4 Excess Rate Cut

Following the application of the burst cut, one could in principle more easily spot jumps in the trigger rate that are not initiated by a high energy event. Though we do not observe this in the open data, we define a selection criterion to remove jumps in the trigger rate that would not otherwise be removed with the burst cut described in Section 6.2.1.3.

After applying all other cuts listed above, the average event rate is determined by dividing the remaining number of events by the surviving live time, which we call μ_r . Surviving events are then binned in time in intervals of 5 s. We then define a cut that excludes events with a number of counts for which the probability predicted by a Poisson distribution with mean μ_r is less than or

equal to 10^{-6} .

6.2.2 Quality Cuts

6.2.2.1 χ^2 Cut

The only selection criterion implemented as a quality cut is a selection based on the χ^2 , comparing the signal template to a given pulse in frequency space. In general, the pulse shape can vary depending on the energy deposited in the crystal, hence we implement an energy-dependent threshold cut based on the χ^2 of a pulse, though we did not notice an appreciable change in the shape of the average pulse over the energy region of interest. To determine the threshold above which events will be excluded as a function of energy, we calculate the distribution of χ^2 for pulses that are expected to match the signal template. To do this, the signal template shown in Fig. 3.5 is scaled to various energies and added to noise traces sampled from data. We then apply the OF




Figure 6.6: A 2D histogram of the frequency-based χ^2 vs. OFL for injected pulses. Data shown in red shows the χ^2 located 3σ from the mean as a function of OFL. The dashed line shows the best fit of Eqn. 6.1 to the data shown in red. Credit: Mason Buchanan.

to each injected pulse and plot the distribution of the resulting χ^2 as a function of OFL (see Figure 6.6). To define a cut threshold that is a function of OFL, $T(A_{\text{OFL}})$, we first bin the χ^2 in OFL and calculate the χ^2 at 3σ from the mean (as determined via a Gaussian fit) for each bin. We then fit a line of the form shown in Eqn. 6.1 to the cut thresholds calculated for each bin.

$$T(A_{\text{OFL}}) = a + (A_{\text{OFL}} + b)^c \quad (6.1)$$

The best fit parameters are found to be $a = 0.3$, $b = 0.05$, and $c = 0.03$.

In principle, the efficiency of the χ^2 cut is set by construction—by defining the cut threshold at 3σ from the mean χ^2 at a given energy, we expect an efficiency of $\gtrsim 95\%$. We confirm this to be the case by directly calculating the fraction of injected pulses that pass the cut. We find that the efficiency is consistent with a constant value of $\sim 98.1\%$ across our energy region of interest as shown in Figure 6.7. The efficiency is taken into account through multiplication with the remaining live time.




Figure 6.7: The passage fraction of injected pulses when subjected to the χ^2 cut given by Eqn. 6.1 plotted as a function of OFL. The red line shows the constant value, which best describes the data shown in black and corresponds to a 98.1% efficiency. Credit: Mason Buchanan.

6.2.2.2 Final Spectrum After Cuts

The final spectrum after selections is shown in Figure 6.8 along with the passage rates on open data which are shown in Table 6.1. Most cuts remove very few events in the energy region of interest with the exception of the burst cut, which accounts for nearly all the live time removed.




Figure 6.8: The energy spectrum before (blue) and after (black) all the data selections. The bins are 1 eV wide.

Live Time Remaining After Data Selections	
Cut	% Remaining Live Time
Mean Baseline	99.57
Burst	44.90
Coincidence	99.93
Excess Rate	100.00
Combined	44.69

Table 6.1: Summary of the remaining live time in open data after each data selection.

6.3 Limit Setting using the Profile Likelihood Ratio Method

The exclusion limits for the DM candidates introduced in Chapter 2 are calculated using the profile likelihood ratio (PLR) method. The PLR method relies on the comparison of model fits to the observed energy spectrum where the DM cross-section (or equivalent parameter) is included as a fitting parameter. Specifically, one first conducts the fit with a fixed signal cross-section, which I will refer to as the “reference” or “local” fit. The resulting likelihood is then compared to the globally maximized likelihood obtained by including the signal cross-section as a fitting parameter. Thus, the PLR is defined as

$$\lambda(\mu) = \frac{L(\mu, \hat{\theta})}{L(\hat{\mu}, \hat{\theta})} \quad (6.2)$$

where L is the likelihood of a model to the observed energy spectrum that depends on the signal cross-section μ as well as any other “nuisance” parameters, which are represented by the vector θ . The numerator in Eqn. 6.2 is the conditionally-maximized likelihood assuming the reference cross-section μ . Here, $\hat{\theta}$ are the nuisance parameters that maximize the likelihood under this assumption and are thus the conditionally-maximum likelihood estimators of θ . The denominator in Eqn. 6.2 is the global maximum of the likelihood corresponding to a maximum-likelihood estimator of the cross-section $\hat{\mu}$ and the maximum-likelihood estimators of the nuisance parameters $\hat{\theta}$.

For setting upper bounds on μ , we consider a test statistic of the form shown in Eqn. 6.3 which depends on the PLR [85].

$$q_\mu = \begin{cases} -2\ln(\lambda(\mu)) & \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu \end{cases} \quad (6.3)$$

Smaller values of q_μ indicate that the reference cross-section μ is more compatible with the observed data set. Since we are only interested in setting upper bounds on μ , we set $q_\mu = 0$ when

$\mu \leq \hat{\mu}$. One can calculate the 90% confidence level on the cross-section by calculating the probability distribution of q_μ , $f(q_\mu|\mu')$. However, calculating $f(q_\mu|\mu')$ directly is often prohibitively expensive and requires thousands of model fits to simulated data sets in order to marginalize the distribution over all the nuisance parameters. Thankfully, there are certain conditions under which the distribution can be reliably approximated. It was shown by Wald [86] that when there is only a single parameter of interest (*i.e.* μ is the only parameter we wish to place bounds on), which is normally distributed with a mean $\hat{\mu}$ and standard deviation σ , the test statistic can be written as

$$-2\ln(\lambda(\mu)) = \frac{(\mu - \hat{\mu})^2}{\sigma^2} + O(1/\sqrt{N}) \quad (6.4)$$

where N is the data sample size. Assuming that N is sufficiently large such that the term with $O(1/\sqrt{N})$ can be ignored and that $\hat{\mu}$ is normally distributed, then the distribution $t_\mu = -2\ln(\lambda(\mu))$ can be written as

$$f(t_\mu; \Lambda) = \frac{1}{2\sqrt{2\pi t_\mu}} \left(e^{-\frac{(\sqrt{t_\mu} + \sqrt{\Lambda})^2}{2}} + e^{-\frac{(\sqrt{t_\mu} - \sqrt{\Lambda})^2}{2}} \right) \quad (6.5)$$

which is a non-central χ^2 distribution with non-centrality parameter Λ . A result by Wilks [87] showed that this can be further simplified in the case where the reference cross-section is equal to the maximum likelihood cross-section, in which case t_μ follows a χ^2 distribution with a single degree of freedom. Thus, assuming that $\hat{\mu}$ is normally distributed and the sample size N is sufficiently large, we have

$$f(q_\mu|\hat{\mu}) = \frac{1}{2}\delta(0) + \frac{1}{2}\chi^2(1) \quad (6.6)$$

where $\delta(x)$ is a Dirac-delta function centered at x (and originates from the fact that $q_\mu = 0$ for $\mu < \hat{\mu}$) and $\chi^2(1)$ is a χ^2 distribution with one degree of freedom.

We can rule out a particular cross-section μ at 90% confidence when we have that

$$0.1 = \int_{q_{\mu'}}^{\infty} f(q_{\mu}|\mu) dq_{\mu} \quad (6.7)$$

Thus, if we can assume that $f(q_{\mu}|\mu)$ has the functional form given by Eqn. 6.6, we simply need to find the value of μ' for which Eqn. 6.7 holds. Using the functional form given in Eqn. 6.6, one can find that Eqn. 6.7 holds when $q_{\mu'} \approx 1.64$, thus, the condition above can be equivalently stated as we must find μ' such that $q_{\mu'} \approx 1.64$.

6.4 Fitting the Observed Energy Spectrum

We model our energy spectrum using the combination of three components:

1. **Non-Ionizing Component:** In Chapter 5, we found that the shape of a non-ionizing component in HV data could be modeled after the shape of the background spectrum taken at 0 V. We also showed that we can predict the rate of the non-ionizing component at any given time and therefore can include a prior on the normalization of this component. In the context of an ERDM search, non-ionizing events are inconsistent with signal and therefore including this component improves our sensitivity by limiting the number of events which can be attributed to DM interactions.
2. **Ionizing Component:** The spectrum we observe in open data cannot be described with just the combination of the ERDM signal models and the non-ionizing component. Thus, we include an ionizing component that is constructed based on our knowledge of leakage and has the same form as the model used to fit HV data in Section 5.3. It is important to note that we have no way of constraining the parameters in this model, including the normalization. It is this fact that prevents the ionizing component from being treated as background; in places




Figure 6.9: Signal models for a) DM-electron scattering assuming a heavy mediator, b) DM-electron scattering assuming a light mediator, c) Dark photon absorption, d) Axion-like particle absorption. Different colors correspond to different masses of the DM particle in each panel.

where the signal model is degenerate with the ionizing model, all events are automatically considered consistent with signal in the PLR method. This is demonstrated in Appendix B.

3. **Signal Component:** We model the observed spectra using four different types of signal models corresponding to ERDM candidates described in Chapter 2. We constrain the various nuisance parameters included in the model (see Table 6.2) using the results from the

calibration. In general, signal models for the considered candidates are similar to the spectra expected from LED events in that they are expected to have peaks which correspond to the NTL energy plus the energy imparted by the DM particle during the recoil. For DM-electron scattering models, the recoil energy is given by a spectrum which depends on the DM mass under consideration. For absorption models (*i.e.* dark photons and ALPS), the recoil energy is mono-energetic and corresponds to the energy of the DM candidate. However, unlike the spectra from LEDs, we do not expect a significant number of surface-trapped events since DM particles would not necessarily interact near the surface; the probability of any given DM particle interacting with the crystal is low and so it is more likely that interactions will occur in the bulk, which makes up the majority of the volume in the crystal. Because of this, the reader will notice the lack of a “0th peak” in the models shown in Figure 6.9. For a detailed discussion on ERDM signal models, see Chapter 2.

The nuisance parameters included in the model are summarized in Table 6.2. We constrain the nuisance parameters for the signal and non-ionizing components according to the results from the LED Calibration (see Table 4.2) and 0V background modeling (see Fig. 5.4) respectively. The parameters in the ionizing component are left unconstrained since we have no way to measure the behavior of this component independently. To reduce the large parameter space shown in Table 6.2, we fix the impact-ionization probability in the ionizing component of the model $p_{\text{II,bkg}}$ to 1.2%, the value obtained from the LED calibration since the impact ionization has been previously observed to be small and thus is not expected to have a large effect on the modeling.

It is worth noting that we differentiate between the charge-trapping probabilities for the ionizing component and the signal model, $p_{\text{CT,bkg}}$ and $p_{\text{CT,DM}}$ respectively. This is due to the fact that fits to HV spectra similar to that done in Section 5.3 yield higher charge-trapping probabilities than found from fits to LED spectra. This effect was most prominent at lower voltages and is likely due

Nuisance Parameters		
Component	Parameter	Gaussian Prior
Non-Ionizing	λ_1 [eV]	$\mu = 6.23, \sigma = 0.08$
	λ_2 [eV]	$\mu = 71, \sigma = 11$
	C	$\mu = 0.995, \sigma = 0.001$
	Normalization	$\mu = 677, \sigma = 39$
Ionizing	$p_{\text{CT,bkg}}$	None
	σ_{eh}	None
	f_{bulk}	None
	Normalization	None
Signal	$p_{\text{CT,DM}}$	$\mu = 27.2\%, \sigma = 0.2\%$
	$p_{\text{II,bkg}}$	$\mu = 1.1\%, \sigma = 0.4\%$
	$E_{\text{loss, constant}}$ [eV]	$\mu = 0.92, \sigma = 0.02$
	$E_{\text{loss, fraction}}$ [eV]	$\mu = 0.55, \sigma = 0.01$
	σ_{res} [eV]	$\mu = 0.589, \sigma = 0.005$
	ϵ_{eff}	$\mu = 98.1\%, \sigma = 0.2\%$
All Components	$\sigma_{\text{calibration}}$	$\mu = 0, \sigma = 1$

Table 6.2: Nuisance parameters for each of the three components that are used to model the observed energy spectrum. The values of μ and σ given in the “Gaussian Prior” column correspond to the mean and standard deviation of the prior placed on the parameter. The priors for the signal and non-ionizing components are based on the results of the modeling discussed in Chapters 4 and 5, respectively. The parameters $E_{\text{loss, constant}}$ and $E_{\text{loss, fraction}}$ are only included as fitting parameters depending on the model for energy loss we are considering and are otherwise fixed (see text for details). Finally, a parameter that accounts for the calibration uncertainty is applied to all components of the model and is in units of uncertainty of E in Eqn. 4.2.

to contributions from burst events, which contribute significantly to lower energy regions in the energy spectra in data taken with lower voltage biases. Thus, we do not assume that the ionizing and signal components of the model must have the same charge-trapping probability.




Figure 6.10: The uncertainty on the calibration as a function of energy in the region of interest.

We also include the calibration uncertainty, $\sigma_{\text{calibration}}$ as a fitting parameter which is applied to all three components in the model. This is a unitless parameter that scales the calibration according to Eqn. 6.8.

$$E' = E + \sigma_{\text{calibration}} \sigma_E(E) \quad (6.8)$$

In Eqn. 6.8, E is the calibrated energy in Eqn. 4.2 and $\sigma_E(E)$ is the uncertainty on the calibrated energy resulting from uncertainties in c_1 and c_2 , which depends on the energy E (see Fig. 6.10). E' is the calibrated energy adjusted by $\sigma_{\text{calibration}} \sigma_E$. From Eqn. 6.8, it can be seen that $\sigma_{\text{calibration}}$ can be thought of as the number of uncertainties the best fit calibration is away from the original calibration. Since we interpret the error on our calibration to be Gaussian with a standard deviation of σ_E at each energy E , we constrain the value of $\sigma_{\text{calibration}}$ using a Gaussian prior centered at 0

with a standard deviation of 1.

Finally, we include parameters that account for the energy loss mechanism observed from the calibration described in Chapter 4, which we call $E_{\text{loss, constant}}$ and $E_{\text{loss, fraction}}$. These parameters are only included as fitting parameters situationally, depending on the scenario we are considering. The different scenarios in which the energy loss can manifest are described below.

1. **Constant Energy Loss:** We can consider that the energy loss manifests as a constant loss per eh pair produced. In this case, we include $E_{\text{loss, constant}}$ (which characterizes the amount of energy lost per eh pair produced) as a fitting parameter constrained by the prior shown in Table 6.2 and fix $E_{\text{loss, fraction}}$ to 1.
2. **Fractional Energy Loss:** We can also consider the case that the energy loss is a fractional loss of the initial recoil energy deposited in the crystal. In this case, we include the fitting parameter $E_{\text{loss, fraction}}$, (which characterizes the fraction of the recoil energy that is kept) and fix $E_{\text{loss, constant}}$ to 0.
3. **No Energy Loss:** The simplest case we can consider is one where the energy loss observed in the LED calibration is only associated with events near the surface and thus would not affect DM events occurring in the bulk. In this case, $E_{\text{loss, constant}}$ and $E_{\text{loss, fraction}}$ would be fixed to 0.

By considering each of these cases separately, we will end up with three different exclusion plots corresponding to the different possible ways that the energy loss behaves. More work will be needed in the future to determine which scenario describes the actual physics of the detector.

The unbinned likelihood of the model (the parameters for which are summarized in Table 6.2)

can be written as

$$L(\mu, \theta) = e^{-N} \left[\prod_i^N f_{0Q}(x_i, \theta_{0Q}) + f_{\text{ion}}(x_i, \theta_{\text{ion}}) + f_{\chi}(x_i, \mu, \theta_{\chi}) \right] \prod_j p(\theta_j) \quad (6.9)$$

where f_{0Q} , f_{ion} , and f_{χ} are the non-ionizing, ionizing, and signal components of the model (which respectively depend on the nuisance parameters contained in the vectors θ_{0Q} ; θ_{ion} , and θ_{χ}); $p(\theta_j)$ is the prior on the nuisance parameter θ_j ; x_i are the energies of observed events; and N is the total number of observations. Recall from Section 6.3 that to calculate the test statistic, we maximize this likelihood when assuming a reference cross-section, $L(\mu, \hat{\theta})$, and again including the signal cross-section μ as a fitting parameter, $L(\hat{\mu}, \hat{\theta})$. Figure 6.11 shows examples of both such fits.




Figure 6.11: a) The reference fit to the observed energy spectrum in open data. b) The global fit to the observed energy spectrum in open data. In both panels, the ionizing component is shown as the green dotted line, the non-ionizing component is shown as the dark blue dotted line, the signal component is shown as the light blue dotted line, and the combination of all the components is shown as the red solid line.

We would then like to find the test statistic that corresponds to a 90% CL on μ . Recall that if N is sufficiently large (such that $\frac{1}{\sqrt{N}} \ll 1$) and if $\hat{\mu}$ follows a Gaussian distribution, we can use

the result from Wilks' theorem, which tells us that q_μ is distributed according to Eqn. 6.6. For the first condition, we have ~ 2000 observations in the open data set. Already, this is sufficiently large since $\frac{1}{\sqrt{N}} \simeq 0.02 \ll 1$. Moreover, the blinded data set consists of roughly two weeks of live time—more than ten times the live time included in the open data set—so it is safe to assume that this condition will also hold when we calculate the limit on blinded data. For the second condition, we can estimate the form of the likelihood as a function of μ while the nuisance parameters are held at their maximum-likelihood estimators ($\theta = \hat{\theta}$.) We want to see that in the region near $\mu \sim \hat{\mu}$, the likelihood can be reasonably approximated as Gaussian. Figure 6.12 shows the likelihood of the open data set around $\hat{\mu}$ while $\theta = \hat{\theta}$ where it can be seen that the likelihood is roughly parabolic near $\hat{\mu}$. As such, the conditions for Wilks' theorem are satisfied.




Figure 6.12: An example of a profile of the negative log likelihood (NLL) of the model fit to open data (black) plotted as a function of μ in the region near $\hat{\mu}$ when $\theta = \hat{\theta}$. Note that in the region near $\hat{\mu}$ (where the NLL is minimized), the profile of the NLL can be approximated by the gaussian shown in red.

From Wilks' theorem, we have that $q_{\mu_{90\%}} \simeq 1.64$. Thus, to calculate the 90% CL on the cross-section, $\mu_{90\%}$, we can simply repeat the fits shown in Fig. 6.11 for a number of reference cross-sections until we find one for which the test statistic is equal to 1.64. In practice, this is done

by sampling the test statistic for a number of reference cross-sections and interpolating between the samples as is shown in Fig. 6.13. Our final exclusion limits, presented in Section 6.5, can then be calculated for each DM candidate under consideration by repeating this process for many different DM masses.




Figure 6.13: The test statistic for a DM model sampled at a number of reference cross-sections. The dashed horizontal line corresponds to $q_\mu = 1.64$.

6.5 Exclusion Limits

The results of the PLR method on open data are shown in Figures 6.14 and 6.15. Note that we must consider three cases depending on the nature of the energy loss discovered in Chapter 4 and therefore we obtain three limits: a limit where we assume no energy loss (as would be the case if only surface events were affected and so bulk DM events would be unaffected), one where we assume we lose a constant amount of energy per eh pair produced, and one where we lose a percentage of the recoil energy. In general, in either scenario the considers energy loss, the signal regions of interest appear at lower energies and so are closer to the regions of our spectrum dominated by backgrounds. For this reason, the limits set while considering energy loss are generally worse than the limits set where no energy loss is considered. Despite this, it can be seen in Fig-




Figure 6.14: Exclusion limits calculated on open data for the DM-electron scattering cross-section, $\bar{\sigma}_e$ for $F_{\text{DM}} = 1$ (left) and $F_{\text{DM}} = 1/q^2$ (right).

ure 6.15 that there are some mass points for which considering an energy loss mechanism seems to improve the limits. We found that this is a purely statistical effect—running the limits on many toy data sets show that the limits between the different scenarios of the energy loss are generally compatible for masses where the limit without energy loss does not set the strongest limit. The reason this effect is only observed for absorption models is due to the signal model being concentrated over small energy ranges as compared to the DM-electron scattering models, which are distributed more broadly over the analysis range (see Figure 6.9). Especially when considering higher mass dark photons or ALPs, the signal becomes concentrated in the $\geq 2\text{eh}$ region, where we have few observations. Thus, depending on the specific location of the signal peaks in relation to the few observations above the 1eh peak, the limit we set may fluctuate.

Regardless of the energy loss scenario we consider, we improve on the limits set by Run 4 for low masses by about an order of magnitude. While encouraging, this is not unexpected—the decision to take blinded data for an ERDM search was motivated by the lower rate in the 1eh peak, where most of the signal for lower mass models is concentrated.




Figure 6.15: Exclusion limits calculated on open data for the kinetic mixing parameter, ϵ for dark photon absorption (left) and on the axio-electric coupling parameter, g_{ae} , for ALP absorption (right).

For the higher masses, HVeV Run 4 generally sets better limits, which can be attributed to the higher amount of live time included in their data set. (The final limits of Run 4 were set using ~ 6 g-days while our open data contains ~ 0.4 g-days after cuts.) For these higher mass DM models, most of the signal is concentrated above the 1eh peak where our backgrounds are low and thus by accumulating more exposure, we can set higher limits. This is also generally the region where HVeV detectors are not competitive with other dedicated DM search experiments, notably the CCD experiments like SENSEI and DAMIC, which utilize large arrays of detectors designed for large exposures.

With this in mind, before unblinding the full data set, we had a good idea of where our limits would improve and where they would stay largely unchanged: For low mass models where the signal was concentrated in regions with high backgrounds (near the 1eh peak), we expect that the limit will not improve outside of statistical fluctuations. In the higher mass regions, we are limited by the exposure we have, and thus by including more data, we expect to set a better limit.

6.5.1 Final Limits

The limits calculated with the full unblinded data sets are shown in Figures 6.16 and 6.17. Note that our earlier prediction was largely correct—for lower masses where the limits are limited by backgrounds, we see little to no improvement compared to the open data. For higher mass regions, however, the inclusion of an additional ~ 3 g-days (after data selections were applied) greatly improved the limits so that we now surpass HVeV Run 4 for all masses considered. Note that at the time of writing this thesis, these results are not yet published and should be considered a work in progress.




Figure 6.16: The final exclusion limits on the DM-electron scattering cross-section, $\bar{\sigma}_e$ for $F_{\text{DM}} = 1$ (left) and $F_{\text{DM}} = 1/q^2$ (right).

Importantly, the exclusion limits we set here explore some new parameter space. To demonstrate this, I have plotted the exclusion limits we set on the DM-electron scattering cross-section (considering a heavy mediator) next to the most recent results from SENSEI [88] and DAMIC-M [89], which are the experiments with the best sensitivity in this low mass region (Fig. 6.18).




Figure 6.17: The final exclusion limits on the kinetic mixing parameter, ϵ for dark photon absorption (left) and on the axio-electric coupling parameter, g_{aee} , for ALP absorption (right).




Figure 6.18: A comparison of the final exclusion limits for DM-electron scattering (heavy mediator) to those from SENSEI [88] (red) and DAMIC-M [89] (pink). Note that the limits from SENSEI and DAMIC-M have been digitized and so may slightly differ from the official limits reported in their corresponding references.

CHAPTER 7

CONCLUSION

In this thesis, I have described the latest dark matter result from HVeV detectors. To show this result in the broader context of HVeV runs, I have reproduced Figure 2.19 from Chapter 2, where I first showed the progress of HVeVs, now updated with HVeV Run 5 (Fig. 7.1).




Figure 7.1: The exclusion limits on DM-electron scattering (light mediator) for past HVeV runs now updated to include Run 5. The limit shown from HVeV Run 5 does not consider energy loss. Please note that the Run 5 limit should still be considered a work in progress.

Figure 7.1 paints an encouraging picture. HVeV Run 5, despite not being a dedicated DM

search, continues the history of improvement seen thus far from HVeV detectors. In addition, even in our R&D capacity, we set limits that are complementary to those set by SENSEI and DAMIC-M—experiments that have historically lead the field in terms of sensitivity. Thus, one of the takeaways from this thesis is proof that SuperCDMS HVeV detectors, while young and still developing, are emerging as world-leading tools for searching for low-mass DM.

However, while the DM exclusion limits obtained in this thesis are an excellent benchmark for HVeV detectors, Run 5 was not intended to be a DM search. Rather, the data taken at SNOLAB were intended to be used to investigate the excesses at low energy and further our understanding of the detectors’ response. Dr. Runze Ren—who is, in a way, my predecessor as the “HVeV student” at Northwestern—noted in the conclusion of his thesis [41] two key mysteries that needed further investigation following HVeV Run 4: 1) the discrepancy in the 0 V and HV calibrations and 2) the low-energy excesses in both 0 V and HV spectra. The work presented in this thesis, while far from providing a thorough explanation, represents a step towards understanding these still open questions.

7.0.1 Calibration and Energy Loss

As discussed in Chapter 4, HVeV Run 4 observed a significant discrepancy in the calibration obtained while operating the detectors at HV compared to operating them at 0 V. Runze pointed out that there were two categories of explanations for this discrepancy: either the response of the detector was changing with voltage (as in, the same phonon energy would register a different response depending on the voltage bias) or the estimations of the phonon energy in one or both of the calibrations was incorrect.

The calibration conducted in HVeV Run 5 tentatively favors the latter hypothesis—the enhancement in detector resolution in Run 5 (0.589 eV compared to the ~ 2 eV obtained in Run 4) enabled observations of the finer structure in the photon calibration peaks and ultimately led to the

discovery of a mechanism for energy loss that had previously escaped observation. If this same energy loss was present in HVeV Run 4, it could resolve the discrepancy between the calibrations. In fact, the scale of the observed energy loss (~ 0.9 eV per photon, a value suspiciously close to the ionization bandgap of silicon) hints at a possible explanation that was also proposed by Runze: that charges freed during an interaction may not recombine on sufficiently fast time scales for the phonons produced by recombination to contribute to the energy of the initial event.

However, there is still work to be done in order to understand the nature of the energy loss observed in Run 5 and to determine whether we can assume that the same mechanism was present in Run 4. We could consider, for example, that trap states introduced by the SiO_2 -Si interface were responsible for the loss of energy in which case a different explanation would be required to resolve the discrepancy observed in Run 4. One possible way to test this would be to operate detectors capable of sub-eV resolutions but without the SiO_2 layer and check the structure of the calibration peaks for evidence of the energy loss. The HVeV detectors required for such a test are already fabricated and are currently being operated.

If the energy loss is consistently observed in detectors with and without the SiO_2 layer, then there is additional work required to understand its behavior. Our current ignorance regarding the physics behind the energy loss is evidenced by our inclusion of three different exclusion limits to cover for its dependence on recoil energy or proximity to the detector's surface. In particular, it would be beneficial to understand whether or not the loss is associated with surface effects (in which case we would not need to consider the energy loss when constructing DM signal models) and whether or not the loss in energy scales with the recoil energy. For the latter question, while the consistency in the energy loss observed for the different photon calibrations shown in Figure 4.11 may imply that the energy loss does not depend on the recoil energy at first glance, the large errors on the measurement conducted at SLAC leave room for a possible scaling. Thus, more precise

calibrations are needed to pin down the dependence on recoil energy. Beyond just HVeV detectors, evidence of an energy loss that is not associated with the SiO_2 layer would have profound implications for the field of experiments searching for DM with similar detectors, such as TESSERACT and CRESST.

7.0.2 Future Investigations of Excesses at Low Energy

The investigation of the low energy backgrounds observed in Runs 4 and 5 should be the highest priority. These backgrounds are not just a limitation for HVeV detectors but for the entire field of experiments searching for low mass DM.

The work presented in this thesis demonstrates an effective, but short-term method for reducing the impact of the non-ionizing component of the LEE on sensitivity to ERDM. While this method seems to have pushed us into an exposure-limited regime for the higher masses of each DM candidate examined in Chapter 6, sensitivity at low masses remains fundamentally limited by the rate in the 1eh peak. Similarly, searches for low-mass DM interacting via nuclear recoils are limited by the non-ionizing component of the LEE. Thus, until these excesses at low energy are understood and can be mitigated, sensitivity to low mass DM has a hard limit.

However, the vast majority of the data taken during HVeV Run 5 was taken with the intent of studying how the rate of these excesses is affected by operating conditions. Prior to any dedicated DM search using HVeV detectors that may be conducted in the future, the remaining data should be analyzed for insights into how the LEE can be mitigated in order to maximize gains in sensitivity. Below, I summarize the highest priority studies that may yield insights on the behaviors of both the ionizing and non-ionizing backgrounds using the data taken at SNOLAB.

7.0.2.1 The Ionizing Background

The improvement in the ERDM exclusion limits obtained in Run 5 compared to Run 4 can be attributed to the inclusion of the non-ionizing component of the LEE as well as an overall lower rate in the 1eh peak. However, the reason for this reduction in rate is still unclear. For the detector used in the DM search, NFCSO, it is unclear as to whether the reduction in rate can be attributed to the cleaner environment, the increased overburden, the SiO₂ layer, or a combination of these effects. Moreover, HVeV Run 4 observed a strong dependence of the 1eh rate on the biasing history of the detector—operating the detector at a given voltage appeared to lead to higher rates in the 1eh peak compared to the rate obtained by first “pre-biasing” at a higher voltage for a period of time. We observed a similar effect in Run 5 but also found that lower bias voltages seemed to lead to lower rates in the 1eh peak in general. During Run 5, we had taken many long exposures at a variety of voltage biases and with varying levels of pre-bias so that the dependence of the rate on both the bias voltage and the pre-bias can be better understood, but this has yet to be analyzed in detail. Before designing and operating a dedicated DM search with HVeV detectors, this data ought to be analyzed to determine the optimal biasing conditions.

The largest contribution to the rate in the 1eh peak in HVeV Run 5 came from the burst events described in Section 6.2.1.3. We found evidence that the rate of these burst events seemed to depend on the voltage bias though this was possibly due to a voltage dependence in the efficiency of the burst identification algorithm; see Appendix A for a more in-depth description. These bursts contribute heavily to the event rate for energies at and below the 1eh peak over long time scales, which prevents us from removing their contribution with 100% efficiency. With this in mind, if one compares data taken at -35 V to the data taken at 0 V (in Fig. 7.2, I have chosen two data sets close in time so that the LEE rate is similar between the two), one can see that the two spectra are similar except that the rate in HV is higher for energies between ~ 5 eV and 1eh peak—the same

region where burst events contributed most heavily, as evidenced by the spectrum before and after the burst selection shown in Figure 6.5.




Figure 7.2: The comparison of the 0V (blue) and the HV ($V = -35$ V) spectrum taken from open data (black) after data selections have been applied. The shaded green region gives the energy distribution on noise triggers and is provided to give a scale for the trigger threshold. Both data sets include ~ 12 hours of live time.

It is possible, then, that contributions from the burst events are responsible for the difference between the two spectra in the region below the 1eh peak and perhaps account for the majority of remaining events in the 1eh peak. Thus, future DM searches using HVeV detectors would benefit considerably by identifying and mitigating the source of bursts. In HVeV Run 4, during which these bursts were also observed, it was proposed that the bursts are caused by highly energetic events originating from radioactive elements embedded in the copper housing. The same could be true for HVeV Run 5, however, a much higher rate of bursts was observed prompting the additional hypothesis that perhaps the SiO_2 luminesces as a response to high energy events absorbed near the surface (see Appendix A for a more detailed discussion of either hypothesis). If the latter

hypothesis is correct, sensitivity could be improved by operating detectors without the SiO_2 layer.

7.0.2.2 The Non-Ionizing Background

Chapter 5 of this thesis showed that the background at 0 V observed by NFCSO between 7 eV and 200 eV was

1. Mostly composed of non-ionizing events.
2. Had an event rate that exponentially rose towards lower energies.
3. Had an overall rate which decayed over time.

These three traits are common among the vast majority of limiting backgrounds observed by experiments in the field—the previously referenced review by Baxter, et. al. [66] synthesizes the observations of the LEE for a large sample of experiments in the field. Thus the work in this thesis, beyond its utility for the ERDM search, merely confirms that the LEE observed by our detectors matches the characteristics of the LEEs observed across the field and thus, gives another observation to compare to.

However, there is far more data from HVeV Run 5 that can be used to understand the LEE. First and foremost, future work is needed to characterize the low energy background observed by NFCSO down to trigger threshold, which based on the resolution of 589 meV, should be ~ 3 eV. We neglected this region in the DM search since there is little signal in that region. More interestingly, the shape of the spectrum for energies below ~ 7 eV seemed to change slightly throughout the course of the run, making it more difficult to characterize. The reason for this change is yet to be understood and could perhaps provide hints on the origin of the background. Secondly, we have not formally examined the 0 V rate and spectral shape in NFCSO before the partial warm-up to 20 K on June 15th, 2026, after which the detector's reponse slightly changed. An examination

of how the LEE rate depends on time before and after the partial warm-up would help add to the observations of the LEE rate “resetting” upon warming up (see observations by CRESST-III [83], for instance) and perhaps allow us to better understand the mechanism behind this behavior.

Most importantly, however, are the 0 V observations from the five other detectors included in HVeV Run 5 that have yet to be formally characterized. The payload featured a number of different detector mask designs, which have differing fractions of the detector surface covered by Al and/or W from the QETs. Thus, comparisons in the LEEs observed by each detector could shed light on the role stress from the relaxation of thin-films deposited on crystal substrates in contributing to the excess rate. This is a mechanism suggested as a possible contribution to the LEE [90] and has yet to be formally tested. Along a similar vein, it would be beneficial to formally compare the LEE rate observed by the detectors used in Run 4 and in Run 5 to see whether the excess we observe depends on the overburden or operating history of the detector.

7.0.3 *Outlook*

The SuperCDMS HVeV program has reached an exciting stage of development. Over the past five years, HVeV detectors have matured significantly, positioning them at the forefront of the search for sub-GeV DM. However, like other experiments in the field, we face a fundamental limit on sensitivity due to excesses at low energy with unknown origin. The goal of the work presented in this thesis was to study and mitigate these excesses. The analysis presented herein characterizes the non-ionizing components of the LEE, which we used to increase our sensitivity to ERDM. However, future studies—including analyses of the remaining data from the work underpinning this thesis—should aim to characterize these backgrounds completely. Thus, HVeV detectors are poised to overcome the obstacles facing not only SuperCDMS, but the field as a whole.

REFERENCES

- [1] H. Poincaré, “The milky way and the theory of gases,” *Popular Astronomy*, vol. 14, pp. 475–488, 1906.
- [2] J. C. Kapteyn, “First Attempt at a Theory of the Arrangement and Motion of the Sidereal System,” vol. 55, p. 302, May 1922. DOI: 10.1086/142670.
- [3] G. Bertone and D. Hooper, “History of dark matter,” *Rev. Mod. Phys.*, vol. 90, p. 045002, 4 Oct. 2018. DOI: 10.1103/RevModPhys.90.045002.
- [4] F. Zwicky, “Die Rotverschiebung von extragalaktischen Nebeln,” *Helv. Phys. Acta*, vol. 6, pp. 110–127, 1933. DOI: 10.1007/s10714-008-0707-4.
- [5] J. Cooley and S. Sekula, *The Dark Matter Discoverer’s Guidebook: Exploring WIMPS, Axion-like particles, and the Dark Sector*. YBK Publishers, Inc., 2024.
- [6] E. Oks, “Brief review of recent advances in understanding dark matter and dark energy,” *New Astronomy Reviews*, vol. 93, p. 101632, 2021.
- [7] E. Corbelli and P. Salucci, “The extended rotation curve and the dark matter halo of m33,” *Monthly Notices of the Royal Astronomical Society*, vol. 311, no. 2, pp. 441–447, Jan. 2000. DOI: 10.1046/j.1365-8711.2000.03075.x. eprint: <https://academic.oup.com/mnras/article-pdf/311/2/441/2881340/311-2-441.pdf>.
- [8] J. Rasmussen, J. Sommer-Larsen, K. Pedersen, *et al.*, “Hot gas halos around disk galaxies: Confronting cosmological simulations with observations,” *The Astrophysical Journal*, vol. 697, no. 1, p. 79, Apr. 2009. DOI: 10.1088/0004-637X/697/1/79.
- [9] M. Bradač, S. W. Allen, T. Treu, *et al.*, “Revealing the properties of dark matter in the merging cluster macs j0025.4–1222,” *The Astrophysical Journal*, vol. 687, no. 2, pp. 959–967, 2008.
- [10] W. A. Dawson, D. Wittman, M. J. Jee, *et al.*, “Discovery of a dissociative galaxy cluster merger with large physical separation,” *The Astrophysical Journal Letters*, vol. 747, no. 2, p. L42, 2012.

- [11] F. Gastaldello, M. Limousin, G. Foëx, *et al.*, “Dark matter–baryons separation at the lowest mass scale: The bullet group,” *Monthly Notices of the Royal Astronomical Society: Letters*, vol. 442, no. 1, pp. L76–L80, 2014.
- [12] N. Aghanim, Y. Akrami, M. Ashdown, *et al.*, “Erratum: Planck 2018 results: Vi. cosmological parameters (astronomy and astrophysics (2020) 641 (a6,” *Astronomy & Astrophysics*, vol. 652, pp. 1–3, 2021.
- [13] F. Hohl, “Numerical Experiments with a Disk of Stars,” vol. 168, p. 343, Sep. 1971. DOI: 10.1086/151091.
- [14] J. P. Ostriker and P. J. E. Peebles, “A Numerical Study of the Stability of Flattened Galaxies: or, can Cold Galaxies Survive?,” vol. 186, pp. 467–480, Dec. 1973. DOI: 10.1086/152513.
- [15] E. Bertschinger, “Simulations of structure formation in the universe,” *Annual Review of Astronomy and Astrophysics*, vol. 36, no. 1, pp. 599–654, 1998.
- [16] G. Blumenthal, F. S., J. Primack, and M. Rees, “Formation of galaxies and large-scale structure with cold dark matter,” *Nature*, vol. 311, pp. 517–525, 1984.
- [17] U. Kumar, A. Ajith, and A. Verma, *Evidence for non-cold dark matter from desi dr2 measurements*, 2025. arXiv: 2504.14419 [astro-ph.CO].
- [18] R. P. Gupta, “Jwst early universe observations and cdm cosmology,” *Monthly Notices of the Royal Astronomical Society*, vol. 524, no. 3, pp. 3385–3395, Sep. 2023. DOI: 10.1093/mnras/stad2032. eprint: <https://academic.oup.com/mnras/article-pdf/524/3/3385/51015444/stad2032.pdf>.
- [19] D. H. Weinberg, J. S. Bullock, F. Governato, R. Kuzio de Naray, and A. H. Peter, “Cold dark matter: Controversies on small scales,” *Proceedings of the National Academy of Sciences*, vol. 112, no. 40, pp. 12 249–12 255, 2015.
- [20] J. Yoo, J. Chaname, and A. Gould, “The end of the MACHO era: Limits on halo dark matter from stellar halo wide binaries,” *The Astrophysical Journal*, vol. 601, no. 1, pp. 311–318, 2004.
- [21] M. Battaglieri, Y. Hochberg, A. Villano, *et al.*, “Arxiv: Us cosmic visions: New ideas in dark matter 2017: Community report,” Tech. Rep., 2017.

- [22] A. Arbey and F. Mahmoudi, “Dark matter and the early universe: A review,” *Progress in Particle and Nuclear Physics*, vol. 119, p. 103 865, 2021.
- [23] S. Dodelson, “Beyond equilibrium,” in *Modern Cosmology*. Academic Press, 2003, pp. 58–79.
- [24] E. Kolb, “Ms turner the early universe front,” *Phys*, vol. 69, pp. 1–547, 1990.
- [25] B. W. Lee and S. Weinberg, “Cosmological lower bound on heavy-neutrino masses,” *Phys. Rev. Lett.*, vol. 39, pp. 165–168, 4 Jul. 1977. DOI: 10.1103/PhysRevLett.39.165.
- [26] K. Sato and M. Kobayashi, “Cosmological Constraints on the Mass and the Number of Heavy Lepton Neutrinos,” *Prog. Theor. Phys.*, vol. 58, p. 1775, 1977. DOI: 10.1143/PTP.58.1775.
- [27] N. Kurinsky, *The low-mass limit: Dark matter detectors with eV-scale energy resolution*. Stanford University, 2018.
- [28] M. F. Albakry, I. Alkhatib, D. Amaral, *et al.*, “A strategy for low-mass dark matter searches with cryogenic detectors in the SuperCDMS SNOLAB facility,” *arXiv preprint arXiv:2203.08463*, 2022.
- [29] R. Agnese, A. J. Anderson, T. Aramaki, *et al.*, “Projected sensitivity of the SuperCDMS SNOLAB experiment,” *Phys. Rev. D*, vol. 95, p. 082 002, 8 Apr. 2017. DOI: 10.1103/PhysRevD.95.082002.
- [30] K. D. Irwin, S. W. Nam, B. Cabrera, B. Chugg, and B. A. Young, “A quasiparticle-trap-assisted transition-edge sensor for phonon-mediated particle detection,” *Review of Scientific Instruments*, vol. 66, no. 11, pp. 5322–5326, 1995.
- [31] R. Ren, C. Bathurst, Y. Y. Chang, *et al.*, “Design and characterization of a phonon-mediated cryogenic particle detector with an ev-scale threshold and 100 kev-scale dynamic range,” *Phys. Rev. D*, vol. 104, p. 032 010, 3 Aug. 2021. DOI: 10.1103/PhysRevD.104.032010.
- [32] K. D. Irwin and G. C. Hilton, “Transition-edge sensors,” in *Cryogenic particle detection*, Springer, 2005, pp. 63–150.
- [33] M. A. Lindeman, “Microcalorimetry and the transition -edge sensor,” Ph.D. dissertation, 2000, p. 203, ISBN: 978-0-599-86727-7.

- [34] J. Yen, B. Shank, B. A. Young, *et al.*, “Measurement of quasiparticle transport in aluminum films using tungsten transition-edge sensors,” *Applied Physics Letters*, vol. 105, no. 16, 2014.
- [35] J. Clarke, “Squid fundamentals,” in *SQUID sensors: fundamentals, fabrication and applications*, Springer, 1996, pp. 1–62.
- [36] R. Cantor, “Dc squids: Design, optimization and practical applications,” in *SQUID sensors: fundamentals, fabrication and applications*, Springer, 1996, pp. 179–233.
- [37] R. Feynman, R. Leighton, and M. Sands, *The schrödinger equation in a classical context: A seminar on superconductivity (chapter 21), the feynman lectures on physics, vol. iii, quantum mechanics*, 1995.
- [38] J. F. Annett, *Superconductivity, superfluids and condensates*. Oxford University Press, 2004, vol. 5.
- [39] R. Simon, M. Burns, M. Colclough, G. Zaharchuk, and R. Cantor, *Mr. SQUID user’s guide*, 6.6, Available at <https://starcryo.com/wp-content/themes/education-pro/manuals/MrSQm66.pdf>, STAR Cryoelectronics, LLC., 1996.
- [40] D. Prêle, M. Piat, L. Sipile, and F. Voisin, “Operating point and flux jumps of a squid in flux-locked loop,” *IEEE Transactions on Applied Superconductivity*, vol. 26, no. 2, pp. 1–5, 2016. DOI: 10.1109/TASC.2015.2510606.
- [41] R. Ren, “Ionization yield measurements and low-energy background investigation using a SuperCDMS-HVeV detector,” Ph.D. dissertation, Northwestern University, 2022. DOI: 10.21985/n2-yvdv-2390.
- [42] B. S. Neganov and V. Trofimov, “Colorimetric method measuring ionizing radiation,” *Otkryt. Izobret*, vol. 146, no. 215, p. 53, 1985.
- [43] P. N. Luke, “Voltage-assisted calorimetric ionization detector,” *Journal of Applied Physics*, vol. 64, no. 12, pp. 6858–6860, 1988.
- [44] K. Ramanathan and N. Kurinsky, “Ionization yield in silicon for eV-scale electron-recoil processes,” *Phys. Rev. D*, vol. 102, no. 6, p. 063 026, 2020. DOI: 10.1103/PhysRevD.102.063026.

- [45] J. Lindhard, M. Scharff, and H. E. Schiøtt, *Range concepts and heavy ion ranges*. Munksgaard Copenhagen, 1963, vol. 33.
- [46] M. F. Albakry, I. Alkhatib, D. Alonso, *et al.*, “First measurement of the nuclear-recoil ionization yield in silicon at 100 eV,” *Phys. Rev. Lett.*, vol. 131, p. 091 801, 9 Aug. 2023. DOI: 10.1103/PhysRevLett.131.091801.
- [47] G. Gerbier, E. Lesquoy, J. Rich, *et al.*, “Measurement of the ionization of slow silicon nuclei in silicon for the calibration of a silicon dark-matter detector,” *Phys. Rev. D*, vol. 42, pp. 3211–3214, 9 Nov. 1990. DOI: 10.1103/PhysRevD.42.3211.
- [48] A. T. Lee, B. Cabrera, B. L. Dougherty, M. J. Penn, J. G. Pronko, and S. Tamura, “Measurements of the ballistic-phonon component resulting from nuclear and electron recoils in crystalline silicon,” *Phys. Rev. B*, vol. 54, pp. 3244–3256, 5 Aug. 1996. DOI: 10.1103/PhysRevB.54.3244.
- [49] A. E. Chavarria, J. I. Collar, J. R. Peña, *et al.*, “Measurement of the ionization produced by sub-keV silicon nuclear recoils in a CCD dark matter detector,” *Phys. Rev. D*, vol. 94, p. 082 007, 8 Oct. 2016. DOI: 10.1103/PhysRevD.94.082007.
- [50] F. Izraelevitch, D. Amidei, A. Aprahamian, *et al.*, “A measurement of the ionization efficiency of nuclear recoils in silicon,” *Journal of Instrumentation*, vol. 12, no. 06, P06014–P06014, 2017.
- [51] A. Villano, M. Fritts, N. Mast, *et al.*, “First observation of isolated nuclear recoils following neutron capture for dark matter calibration,” *Physical Review D*, vol. 105, no. 8, p. 083 014, 2022.
- [52] R. Agnese, T. Aralis, T. Aramaki, *et al.*, “Energy loss due to defect formation from ^{206}Pb recoils in SuperCDMS germanium detectors,” *Applied Physics Letters*, vol. 113, no. 9, 2018.
- [53] M. J. Wilson, “A new search for low-mass dark matter and an examination and reduction of the uncertainty due to the photoelectric absorption cross section using a cryogenic silicon detector with single-charge sensitivity,” Ph.D. dissertation, University of Toronto (Canada), 2022.
- [54] M. J. Wilson, A. Zaytsev, B. von Krosigk, *et al.*, “Improved modeling of detector response effects in phonon-based crystal detectors used for dark matter searches,” *Phys. Rev. D*, vol. 109, p. 112 018, 11 Jun. 2024. DOI: 10.1103/PhysRevD.109.112018.

- [55] E. Zhang, “Light dark matter constraints from a supercdms cryogenic silicon detector with single-charge sensitivity,” Ph.D. dissertation, University of Toronto (Canada), 2025.
- [56] R. Essig, M. Fernandez-Serra, J. Mardon, A. Soto, T. Volansky, and T.-T. Yu, “Direct detection of sub-gev dark matter with semiconductor targets,” *Journal of High Energy Physics*, vol. 2016, no. 5, pp. 1–54, 2016.
- [57] Y. Hochberg, T. Lin, and K. M. Zurek, “Absorption of light dark matter in semiconductors,” *Physical Review D*, vol. 95, no. 2, p. 023 013, 2017.
- [58] M. Fabbrichesi, E. Gabrielli, and G. Lanfranchi, *The physics of the dark photon: a primer*. Springer, 2021, pp. 2–5.
- [59] J. E. Kim and G. Carosi, “Axions and the strong CP problem,” *Rev. Mod. Phys.*, vol. 82, pp. 557–601, 1 Mar. 2010. DOI: 10.1103/RevModPhys.82.557.
- [60] M. Pospelov, A. Ritz, and M. Voloshin, “Bosonic super-wimps as kev-scale dark matter,” *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, vol. 78, no. 11, p. 115 012, 2008.
- [61] R. Agnese, T. Aralis, T. Aramaki, *et al.*, “First dark matter constraints from a supercdms single-charge sensitive detector,” *Phys. Rev. Lett.*, vol. 121, p. 051 301, 5 Aug. 2018. DOI: 10.1103/PhysRevLett.121.051301, Erratum published in “Erratum: First dark matter constraints from a supercdms single-charge sensitive detector [phys. rev. lett. 121, 051301 (2018)],” *Phys. Rev. Lett.*, vol. 122, p. 069 901, 6 Feb. 2019. DOI: 10.1103/PhysRevLett.122.069901. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevLett.122.069901>.
- [62] M. F. Albakry, I. Alkhatib, D. Alonso-González, *et al.*, “Light dark matter constraints from SuperCDMS HVeV detectors operated underground with an anticoincidence event selection,” *Phys. Rev. D*, vol. 111, p. 012 006, 1 Jan. 2025. DOI: 10.1103/PhysRevD.111.012006.
- [63] M. F. Albakry, I. Alkhatib, D. Alonso-González, *et al.*, “Search for low-mass electron-recoil dark matter using a single-charge sensitive supercdms-hvev detector,” *Phys. Rev. D*, vol. 113, p. 032 001, 3 Feb. 2026. DOI: 10.1103/51np-6mng.
- [64] D. W. Amaral, T. Aralis, T. Aramaki, *et al.*, “Constraints on low-mass, relic dark matter candidates from a surface-operated supercdms single-charge sensitive detector,” *Phys. Rev. D*, vol. 102, 091101(R), 9 Nov. 2020. DOI: 10.1103/PhysRevD.102.091101.

- [65] M. F. Albakry, I. Alkhatib, D. W. P. Amaral, *et al.*, “Investigating the sources of low-energy events in a supercdms-hvev detector,” *Phys. Rev. D*, vol. 105, p. 112 006, 11 Jun. 2022. DOI: 10.1103/PhysRevD.105.112006.
- [66] D. Baxter, R. Essig, Y. Hochberg, *et al.*, “Low-energy backgrounds in solid-state phonon and charge detectors,” *Annual Review of Nuclear and Particle Science*, vol. 75, no. Volume 75, 2025, pp. 301–326, 2025. DOI: <https://doi.org/10.1146/annurev-nucl-121423-100849>.
- [67] L. M. Garrison, “Measurement of neutron and muon fluxes 100 m underground with the scibath detector,” Ph.D. dissertation, Indiana University, 2014.
- [68] P. Camus, J. Corbett, S. Crawford, *et al.*, “CUTE: A cryogenic underground test facility at SNOLAB,” *Frontiers in Physics*, vol. 11, p. 1 319 879, 2024. DOI: 10.3389/fphy.2023.1319879.
- [69] K. Kennard, A. Pradeep, M. Buchanan, *et al.*, “Performance of a superCDMS HVeV detector with sub-ev energy resolution and single charge-sensitivity,” *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 1091, p. 171 753, 2026. DOI: <https://doi.org/10.1016/j.nima.2026.171753>.
- [70] I. Alkhatib, “Analysis of Signals in Semiconductor Cryogenic Dark Matter Search Detectors and new Constraints on Interactions of Light Dark Matter With Electrons,” Ph.D. dissertation, Toronto U., 2025.
- [71] B. A. Hines *et al.*, “The Cryogenic Dark Matter Search Test Stand Warm Electronics Card,” in *2010 IEEE Nuclear Science Symposium, Medical Imaging Conference, and 17th Room Temperature Semiconductor Detectors Workshop*, 2010, pp. 1392–1395. DOI: 10.1109/NSSMIC.2010.5874000.
- [72] M. F. Albakry, I. Alkhatib, D. Alonso-González, *et al.*, “Low-energy calibration of Super-CDMS HVeV cryogenic silicon calorimeters using compton steps,” *Phys. Rev. D*, vol. 112, p. 092 014, 9 Nov. 2025. DOI: 10.1103/jj7w-gkgg.
- [73] M. A. Green, “Self-consistent optical parameters of intrinsic silicon at 300 K including temperature coefficients,” *Solar Energy Materials and Solar Cells*, vol. 92, no. 11, pp. 1305–1310, 2008.

- [74] F. Ponce, C. Stanford, S. Yellin, *et al.*, “Measuring the impact ionization and charge trapping probabilities in SuperCDMS HVeV phonon sensing detectors,” *Phys. Rev. D*, vol. 101, 031101(R), 3 Feb. 2020. DOI: 10.1103/PhysRevD.101.031101.
- [75] O. Abramoff, L. Barak, I. M. Bloch, *et al.*, “SENSEI: Direct-detection constraints on sub-gev dark matter from a shallow underground run using a prototype skipper CCD,” *Phys. Rev. Lett.*, vol. 122, p. 161 801, 16 Apr. 2019. DOI: 10.1103/PhysRevLett.122.161801.
- [76] A. Aguilar-Arevalo, D. Amidei, D. Baxter, *et al.*, “Constraints on light dark matter particles interacting with electrons from DAMIC at SNOLAB,” *Phys. Rev. Lett.*, vol. 123, p. 181 802, 18 Oct. 2019. DOI: 10.1103/PhysRevLett.123.181802.
- [77] R. Anthony-Petersen, C. L. Chang, Y.-Y. Chang, *et al.*, “Low energy backgrounds and excess noise in a two-channel low-threshold calorimeter,” *Applied Physics Letters*, vol. 126, no. 10, p. 102 601, Mar. 2025. DOI: 10.1063/5.0247343. eprint: https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/5.0247343/20436183/102601_1_5.0247343.pdf.
- [78] Q. Arnaud, E. Armengaud, C. Augier, *et al.*, “First germanium-based constraints on sub-mev dark matter with the EDELWEISS experiment,” *Phys. Rev. Lett.*, vol. 125, p. 141 301, 14 Oct. 2020. DOI: 10.1103/PhysRevLett.125.141301.
- [79] G. Angloher, M. Bauer, I. Bavykina, *et al.*, “Results from 730 kg days of the CRESST-II dark matter search,” *The European Physical Journal C*, vol. 72, no. 4, p. 1971, 2012.
- [80] A. H. Abdelhameed, G. Angloher, P. Bauer, *et al.*, “First results from the CRESST-III low-mass dark matter program,” *Phys. Rev. D*, vol. 100, p. 102 002, 10 Nov. 2019. DOI: 10.1103/PhysRevD.100.102002.
- [81] G. Angloher, S. Banik, G. Benato, *et al.*, “Doublets detectors to investigate the CRESST low energy background: Results from above-ground prototypes,” *The European Physical Journal C*, vol. 84, no. 10, p. 1001, 2024.
- [82] I. Alkhatib, D. W. P. Amaral, T. Aralis, *et al.*, “Light dark matter search with a high-resolution athermal phonon detector operated above ground,” *Phys. Rev. Lett.*, vol. 127, p. 061 801, 6 Aug. 2021. DOI: 10.1103/PhysRevLett.127.061801.
- [83] G. Angloher, S. Banik, G. Benato, *et al.*, “Latest observations on the low energy excess in CRESST-III,” *SciPost Physics Proceedings*, no. 12, p. 013, 2023.

- [84] C. Augier, G. Baluieu, V. Belov, *et al.*, “First demonstration of 30 eVee ionization energy resolution with ricochet germanium cryogenic bolometers,” *The European Physical Journal C*, vol. 84, no. 2, p. 186, 2024. DOI: <https://doi.org/10.1140/epjc/s10052-024-12433-1>.
- [85] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, “Asymptotic formulae for likelihood-based tests of new physics,” *The European Physical Journal C*, vol. 71, no. 2, p. 1554, 2011.
- [86] A. Wald, “Tests of statistical hypotheses concerning several parameters when the number of observations is large,” *Transactions of the American Mathematical Society*, vol. 54, no. 3, pp. 426–482, 1943.
- [87] S. S. Wilks, “The large-sample distribution of the likelihood ratio for testing composite hypotheses,” *The Annals of Mathematical Statistics*, vol. 9, no. 1, pp. 60–62, 1938.
- [88] I. M. Bloch, A. M. Botti, M. Cababie, *et al.*, “SENSEI at SNOLAB: Single-electron event rate and implications for dark matter,” *Phys. Rev. Lett.*, vol. 134, p. 161 002, 16 Apr. 2025. DOI: [10.1103/PhysRevLett.134.161002](https://doi.org/10.1103/PhysRevLett.134.161002).
- [89] K. Aggarwal, I. Arnquist, N. Avalos, *et al.*, “Probing benchmark models of hidden-sector dark matter with DAMIC-M,” *Phys. Rev. Lett.*, vol. 135, p. 071 002, 7 Aug. 2025. DOI: [10.1103/2tcc-bqck](https://doi.org/10.1103/2tcc-bqck).
- [90] R. K. Romani, “Aluminum relaxation as the source of excess low energy events in low threshold calorimeters,” *Journal of Applied Physics*, vol. 136, no. 12, 2024.

**AN ELECTRON-RECOIL DARK MATTER SEARCH USING THE NEXT
GENERATION OF SUPERCDMS HVEV DETECTORS OPERATED AT SNOLAB**

Approved by:

Enectali Figueroa-Feliciano
Principle Advisor
Northwestern University

Daniel Baxter
Co-Advisor
*Northwestern University & Fermilab National
Accelerator Laboratory*

Eric Dahl
Committee Member
Northwestern University

Date Approved: , 2026

APPENDIX A

BURST EVENTS

Burst events, characterized by energetic primaries followed by multiple lower energy secondaries, have been observed in a number of HVeV runs beginning in HVeV Run 3. In every HVeV run since, burst events have been one of the main limitations on sensitivity due to their association with lasting increases in event rate. Thus, identifying, characterizing, and mitigating burst events in future HVeV runs is of utmost importance. This appendix compiles the observations of burst events from HVeV Runs 3, 4, and 5.

A.1 HVeV Run 3

HVeV Run 3 observed a large population of events consisting of multiple pulses occurring within quick succession, which the analysts referred to as burst events. These events could be readily identified by a jump in event rate and were characterized by an energetic primary followed by secondary events with energies distributed mostly among the 1st, 2nd, and 3rd eh peaks [65]. Events such as these were found in multiple detectors and thus were believed to be caused by an external background. Evidently, as shown in Ref. [65], the distribution in the energy and arrival time of secondaries in these bursts matched expectations from luminescence in SiO_2 —the primary component of the PCB boards that were used to mount the detectors. For this reason, future HVeV payloads avoided using scintillating materials near the detectors and replaced the mounting scheme that used PCBs with the copper housings that were used in HVeV Runs 4 and 5. The massive reduction in the 2eh and 3eh rate observed in HVeV Run 4 compared to that of Run 3 was attributed to this redesign of the detector housing.




Figure A.1: An example of a (filtered) burst event observed in Run 3 shown in blue (orange). Vertical dashed lines correspond to points at which the filtered trace exceeds the trigger threshold. Figure taken from Ref. [65].

A.2 HVeV Run 4

The burst events observed during HVeV Run 4 differed slightly from those seen in HVeV Run 3. In general, the primary event was at a far higher energy ($\gtrsim 1$ keV) and usually saturated the detector. Moreover, the burst events in HVeV Run 4 were associated with increases in the event rate lasting for up to 10 hours after the primary event. The events were generally not associated with coincidence in other detectors and usually occurred at a rate of ~ 1 -2 per day. Interestingly, Dr. Imran Alkhatib points out in his thesis that the events were not observed after the copper lids that enclosed the detectors' housings were switched [70]. Because of this fact, it was proposed that the events may originate from radioactive impurities in the copper, specifically $^{210}_{82}\text{Pb}$, which produces alpha particles in its decay chain. The alpha particle, if ejected from the copper lid and subsequently implanting near the detector surface, would initially deposit a large amount of energy in the detector and possibly dislocate atoms in the crystal lattice. Thus, the increased event rate following the initial collision could be explained by the relaxation of the lattice. While not tested

explicitly, this hypothesis was generally accepted as the likely cause of the burst events.

A.3 HVeV Run 5

In general, the burst events observed by NFCSO in HVeV Run 5 exhibit similar behaviors to those seen in Run 4; the primary event generally saturates the detector, and the energies of secondaries are typically at or below the 1eH peak energy. We also observe an increased event rate following a burst event though the time scale over which the rate is increased was seen to be anywhere between a few minutes to several hours in Run 5, whereas the time scale was on the order of ~ 10 hours in Run 4.

The frequency of the burst events in NFCSO, however, differed significantly from the observations of HVeV Run 4 and seemed to depend on the voltage bias. While operating with voltage biases $\gtrsim 60$ V, the rate of burst events was roughly consistent with HVeV Run 4—we observed usually 1-2 per day. When operating with voltages $\lesssim 20$ V, bursts were never observed. However, for voltages between 20-60 V, the rate of bursts increased significantly—at 40 V, the rate was more than 50 times the rate observed in HVeV Run 4. While this appears to suggest that the mechanism behind the burst events depends on the voltage bias, and would therefore disfavor the hypothesis suggested during Run 4, I will present evidence in this section to show that the apparent dependence is related to our ability to identify bursts as opposed to a physical relationship. Additionally, our results suggest that the contribution from burst events to the event rate cannot be removed with near 100% efficiency and therefore burst events pose a significant limitation on DM sensitivity to HVeV Run 5 and future HVeV runs.

A.4 Triggering Efficiency of Burst Events in Run 5

In this section, I will show how we conclude that burst events can not be removed with 100% efficiency. Here, I will formally define burst events as events with energies $\gtrsim 300$ eV and with




Figure A.2: The observed rate of saturated (burst) events plotted against the HV bias applied to the crystal is shown in black (red). Vertical error bars represent the statistical uncertainty of the measured rates from Poisson statistics. Here, saturated events are defined as events with energies exceeding ~ 300 eV. Burst events are defined as saturated events in which at least three additional pulses are observed in the 260 ms time window following the saturated event.

at least three pulses occurring within 260 ms after the initial event. The choice to examine a window of 260 ms was somewhat arbitrary but was chosen to be long enough for the detector to return to its operating point. To test our ability to identify bursts as a function of voltage bias, I examined the triggering efficiency of secondaries at a number of voltage biases. The ability to trigger on secondaries is directly related to the ability to identify burst events since both in the analysis shown here and in the burst event cut defined in Section 6.2.1.3, burst events are identified by the existence of secondaries following a saturated event. I found that the triggering efficiency of the secondaries was strongly affected by both the energy of the secondaries and the duration for which the detector is saturated after the primary event, both of which depend on the voltage bias as described below.

1. **Energy Distribution of Secondaries:** We found that for voltage biases ≥ 50 V, the energies of secondaries associated with bursts were generally distributed at the 1eh peak. Secondaries with energies above or below the 1eh peak were found to occur at rates roughly consistent with the levels of charge-trapping and impact-ionization measured from calibration data. An example of the spectral contribution from bursts at 60 V can be seen in Fig. A.3a. However,




Figure A.3: The energy spectrum in units of n_{eh} before (blue) and after (black) applying the burst cut described in Section 6.2.1.3 for data taken at 60 V (left) and at -35 V (right). The difference in the distributions shown in each panel can be interpreted as the contribution from secondaries from burst events.

for voltages below ~ 40 V, the distribution of secondaries appreciably changed; rather than being concentrated at the 1eh peak, the energy distribution of secondaries formed a broad “shoulder” ranging between 0 eV and the 1eh peak energy as shown in Fig. A.3b. Notably, the rate at which events appeared below the 1eh peak was too high to be consistent with charge-trapped 1eh peak events. Possible reasons for this difference will be discussed later in this appendix.

2. **Time Saturated after Primary:** The time for which the detector is saturated after a primary event scales with the energy of the event. As such, it is not surprising that for higher voltages,

primaries in burst events generally saturate the detector for longer times and thus have larger integrals (see Fig. A.4). This is relevant for the trigger efficiency of secondaries since events that occur while the detector is saturated yield no change in the measured current and thus are not triggered.




Figure A.4: The mean integral of saturated events plotted as a function of HV bias. Vertical error bars represent the standard deviation of integrals for saturated events.

To estimate how the effects described above affect the triggering efficiency of secondaries I simulate bursts by adding the pulse template to saturated events sampled from data for a number of voltage biases. The energy of the simulated secondaries is randomly sampled from a distribution between 0 and $1 n_{eh}$, where the probability to get an energy below $1 n_{eh}$ is based on the charge-trapping probability calculated by fitting a detector response model to LED data taken at the corresponding voltage bias. To capture the effects of saturation on the triggering efficiency, the amplitude of the simulated secondaries is reduced by a factor, α , which depends on the mean baseline, $\bar{I}$, of the secondary and is given by Eqn. A.1.

$$\alpha = \frac{\sqrt{\bar{I}_{\max} - \bar{I}}}{\sqrt{\bar{I}_{\max}}} \quad (\text{A.1})$$

The functional form of Eqn. A.1 was determined through an empirical fit to the maximum energy of secondaries as a function of the mean baseline for burst events observed in data. The maximum baseline at which the detector is completely saturated, $\bar{I}_{\max}$, is also extracted from data.

For simulated burst events, the triggering efficiency is calculated via the ratio of triggered secondaries to simulated secondaries as a function of time after the primary event (see Fig. A.5a). The triggering efficiency is then averaged over the first 5000 samples (~ 32 ms) (see Fig. A.5b). The averaged triggering efficiency for secondaries can be thought of as a proxy for the efficiency in identifying bursts since burst events are identified by the proximity of pulses occurring after saturated events.




Figure A.5: **Left:** The triggering efficiency of secondaries as a function of time (1 sample $\sim 6.4 \mu\text{s}$). Different lines correspond to different voltage biases. **Right:** The average triggering efficiency for secondaries calculated over the first 5000 samples after the primary event. This can be interpreted as the efficiency of identifying bursts.

The limited triggering efficiency for secondaries at low voltages is due to the low energy of secondaries, whereas at high voltages it is caused by the longer lasting saturation. This leads to the dependence shown in Fig. A.5b which displays a strikingly similar dependence on voltage as

the observed rate of burst events shown in Fig. 6.3. This suggests that the cause of the voltage dependence of the observed burst event rate is due to the efficiency in the identification of bursts depending on voltage. Additionally, Fig. A.5b tells us that we likely fail to identify a significant fraction of burst events due to our inability to reliably trigger on secondaries occurring soon after primaries. Since burst events are associated with long-lasting increases in event rate, it is possible that a significant fraction of the events observed in HV data are associated with unidentified burst events. The inability to trigger on secondaries for higher voltages may also explain the qualitative difference in the contribution of bursts to the energy spectra at high voltages compared to low voltages as shown in Figure A.5. If at higher voltages, the detector is saturated for longer times after primary events, then it is possible that low-energy pulses occurring soon after the primary are lost. As we go to lower voltages, near $\sim 30\text{-}40$ V, however, the detector is saturated for less time and thus it is possible to trigger on these lower energy events with a higher efficiency, thus giving rise to the increased event rate for energies below the 1eh peak.

A.5 Discussion

The observations on the efficiency of identifying bursts have interesting implications regarding the bursts observed in Run 4. For example, it is possible that the bursts observed by NFCSO share an origin with those observed in Run 4. In that case, the reason for the far lower rate of burst events observed in Run 4 detectors may be due to the fact that Run 4 operated with a higher voltage bias (100 V) and have a worse energy resolution. Based on the study shown in this appendix, both the high voltage bias and the worse resolution would contribute to a lower efficiency in identifying burst events, leading to a lower observed rate. If this is the case, then it's possible that many of the 1eh events observed by Run 4 are actually associated with unidentified bursts and may actually explain the higher overall rate in the 1eh peak compared to NFCSO. On the other hand, it is possible that the bursts in Run 5 have a different origin than those in Run 4. For example, it may

be that the burst events observed by NFCSO are associated with luminescence in the SiO_2 layer induced by energetic events. At the time of writing this thesis, there are ongoing tests of HVeV detectors capable of sub-eV resolution but fabricated without the SiO_2 layer, which may help prove or disprove this hypothesis. In any case, further work is needed to determine the origin of these burst events. Identifying and mitigating this background would be a boon for sensitivity to ERDM and thus should be a top priority for future HVeV runs.

APPENDIX B

THE PROFILE LIKELIHOOD RATIO WITH DEGENERATE SIGNAL AND BACKGROUND MODELS

The models to the observed energy spectra shown in Chapter 6 include a background ionizing component that is degenerate with the signal component (*i.e.*, the models share features—the 1eV peak being the most prominent—that often overlap.) Under these conditions, one might be concerned that the exclusion limit will be overly optimistic due to some events which could be consistent with signal in the spectrum being attributed to the ionizing component. In this appendix, I will demonstrate that this is not the case by showing that as the signal and background models approach 100% degeneracy (*i.e.*, the background and signal models are functionally identical), the exclusion limit calculated by the PLR is equivalent to that of Poisson counting under the signal-only hypothesis.

To demonstrate this, I simulate 2000 events that are sampled from a DM signal model with a known cross-section, μ_{true} , which results in the energy spectrum shown in Figure B.1. We can calculate the 90% cross-section, $\mu_{90\%}$, under a signal-only hypothesis from Poisson statistics: interpreting the number of “observations” as the mean of a Poisson distribution, then the 90th percentile of that distribution is ~ 2057 events. Equivalently, this means that in 90% of experiments under these conditions, we would observe fewer than 2057 events, which corresponds to $\mu_{90\%} = 3.22 \times 10^{-31} \text{ cm}^2$.

We now want to see if we would obtain this same limit using the PLR method when using a “background” model that occupies the same parameter space as the signal model. As such, we will consider a background model that has the same functional form as the signal model but with parameters that are allowed to fluctuate independently of the parameters included in the signal




Figure B.1: The simulated energy spectrum sampled from a DM signal model with a known cross-section of $\mu_{\text{true}} \simeq 3.13 \times 10^{-31} \text{ cm}^2$ is shown in gray. The DM signal model is shown in red.

model. Importantly, the normalization of the background model is unconstrained; constraining the normalization would imply that we understand our backgrounds well enough to predict the number of events that are not attributable to signal.

First, let us consider the case where the reference cross-section is less than the true cross-section, $\mu < \mu_{\text{true}}$. When this is the case, the local fit will find that the number of signal events N_{signal} is less than the total number of observations, N . Thus, we can expect that the likelihood in the reference fit will be maximized when the normalization of the background model corresponds to a number of background events N_{bkg} such that $N_{\text{signal}} + N_{\text{bkg}} = N$. Similarly, in the global fit, the model is free to choose how many events to attribute to signal or background with no cost to the likelihood and so will choose some normalizations for the signal and background such that the total number of observations predicted is consistent with observations: $N'_{\text{signal}} + N'_{\text{bkg}} = N$. In other words, the likelihood is only sensitive to the total number of counts—it cannot distinguish between

the signal and background. We want to then check that the limit that is set is representative of all the observations being counted as signal. Thus, we have that $L(\mu, \hat{\theta}) = L(\hat{\mu}, \hat{\theta})$ when $\mu < \mu_{\text{true}}$. Or, equivalently

$$q_\mu = 0, \quad \mu < \mu_{\text{true}} \quad (\text{degenerate signal and bkg}) \quad (\text{B.1})$$

Since here we consider μ_{true} to be the cross-section that is consistent with all the observations originating from signal, this is completely consistent with the condition that $q_\mu = 0$ when $\hat{\mu} > \mu$ in our definition of the test statistic given in Eqn. 6.3.

When $\mu > \mu_{\text{true}}$, the reference fit begins to predict a number of counts larger than the number of observations. For the local fit, the likelihood will be maximized when $N_{\text{bkg}} = 0$ but we still have $N_{\text{signal}} > N$. The global likelihood, however, will be free to lower the total amount of signal so that $N'_{\text{signal}} + N'_{\text{bkg}} = N$. Thus, we will find that $L(\mu, \hat{\theta}) < L(\hat{\mu}, \hat{\theta})$, in which case $q_\mu > 0$. The value of q_μ will continue to increase for larger values of μ . This is exactly what we observe to be the case when we calculate the test statistic for a number of reference cross-sections, as shown in Fig. B.2. Again, this behavior is consistent with our definition of the test statistic given by Eqn. 6.3.

Now that we have established the behavior of q_μ , we would like to see that we can still apply Wilks' theorem. Recall that for Wilks' theorem to hold, we require, 1) the number of observations is large and 2) $\hat{\mu}$ is Gaussian distributed. In this model, the first condition is satisfied by construction—we simulate a spectrum of 2000 events which is a sufficiently large sample size. This is also true for our real data as the open data contains a number of observations on the order of thousands. The second condition, however, we have already shown is not generally held. Because the signal and background are completely degenerate, the likelihood can be maximized for $\hat{\mu}$ anywhere between 0 and μ_{true} with an appropriate choice for N_{bkg} , thus $\hat{\mu}$ is not Gaussian distributed.

Despite this, the *total number* of observations, $N = N_{\text{sig}} + N_{\text{bkg}}$ that maximizes the likelihood




Figure B.2: A plot of q_μ as a function of the reference cross-section. For each reference cross-section, ~ 50 toy data sets are simulated with the test statistic calculated for each one. The 1- and 2- σ bands indicated by the green and yellow shaded regions correspond to the mean q_μ plus 1 or 2 standard deviations.

is Gaussian distributed in the limit that the number of observations is large.¹. Thus, we may still apply Wilk's theorem to the *total number* of observations rather than specifically the signal. This is equivalent to the case of Poisson counting, where we do not attempt to model the background or signal at all, but instead, calculate a limit simply using the total number of observations. The PLR ends up doing this “automatically”—since the likelihood is only sensitive to the total number of observations, then $L(\mu, \theta) = L(N, \theta')$ where θ' is θ without N_{bkg} since it is not included in N .

To show that this is the case, applying the PLR method to the simulated data set yields a $\mu_{90\%} = 3.23 \times 10^{-31} \text{ cm}^2$ (Figure B.2) which is in excellent agreement with the exclusion limit expected from Poisson counting. Thus, we have shown that the exclusion limits calculated by the PLR method using degenerate signal and background models are consistent with the result from

¹In general, the total number of observations would be Poisson distributed. However, if the number of observations is large enough, then the Poisson distribution can be well approximated by a Gaussian




Figure B.3: The distribution of $q_{\mu=\mu_{\text{true}}}$ calculated from ~ 100 toy experiments assuming degenerate signal and background models is shown in red. The distribution expected from Wilk's theorem (Eqn. 6.6) is shown as the dashed black line.

Poisson counting under the signal-only hypothesis. In the context of the ERDM search shown in Chapter 6, we conclude that degeneracy in the signal and background models does not violate the assumptions implicit in the PLR method.

Before concluding, one final remark I will make is that sensitivity in the case of degenerate signals and backgrounds will always be worse compared to sensitivity in the case where you can differentiate the two. To make this explicit, note that in the example considered in this appendix, I considered μ_{true} as the cross-section that yields $N_{\text{sig}} = N$ where N is the number of observations in the simulated data. However, I could have just as easily considered the case where μ_{true} yields $N_{\text{sig}} < N$ and assumed that the remaining counts were from background. If this were the case, obviously the limit set by the PLR does not change—the likelihood cannot differentiate signal and background and is only sensitive to the total number of observations. Thus, in the presence of a degenerate background, the exclusion limit we set will be higher than the “true” 90% confidence

level.