Physics Theses and Dissertations   Physics

A Likelihood Search for Low-mass Dark Matter via Inelastic Scattering in SuperCDMS

Daniel Jardin
Southern Methodist University, djardin@smu.edu

Follow this and additional works at: https://scholar.smu.edu/hum_sci_physics_etds

Part of the Astrophysics and Astronomy Commons

A LIKELIHOOD SEARCH FOR LOW-MASS DARK MATTER

VIA INELASTIC SCATTERING IN SUPERCDMS

Approved by:

_____

Dr. Jodi Cooley
Associate Professor of Experimental
Physics

_____

Dr. Stephen Sekula
Associate Professor of Experimental
Physics

_____

Dr. Fredrick Olness
Professor of Theoretical Physics

_____

Dr. Tarek Saab
Associate Professor, University of Florida

A LIKELIHOOD SEARCH FOR LOW-MASS DARK MATTER

VIA INELASTIC SCATTERING IN SUPERCDMS

A Dissertation Presented to the Graduate Faculty of the

Dedman College

Southern Methodist University

in

Partial Fulfillment of the Requirements

for the degree of

Doctor of Philosophy

with a

Major in Physics

by

Daniel Jardin

B.S., Physics, Syracuse University

August 6, 2019

iv

for alternative dark matter led by Belina von Krosigk, my fellow co-Chair Matt Wilson, and myself.

I thank Richard Schnee, my undergraduate research advisor and SuperCDMS collaborator for first introducing me to the world of radon studies. I clearly remember the day you subtly urged me to apply to SMU. Texas felt like a foreign land at the time, but that turned out to be the beginning of a very fulfilling journey. I dove deeper into radon measurements thanks to Ray Bunker, who was an accommodating host during my stint at PNNL. You taught me that having a knack for attention to detail can be crafted into an art form when it comes to writing research logs and grant proposals. I'm glad I get to pass on my knowledge of the radon decay chain to Taylor Wallace at SMU. Taylor, you have a tremendous work ethic and mentoring you through our various radon plate-out studies has been a rewarding experience.

It has been a wonderful six years of navigating the challenging world of physics with my friends at SMU, especially Ryan Staten, Eric Godat, and Andrew Turvey. I don't know that I would have made it through the gauntlet of classes if I didn't have you next to me, banging your heads against the whiteboard (literally and figuratively) until three in the morning. Also a special thanks to Matthew Feickert for always being a brain to pick about new coding techniques and best practices. Your expert development of the SMU thesis template allowed me to focus on the contents instead of fighting with LaTeX for (quite as many) hours. I will miss our late night / early morning Commiseration Committee meetings.

Nobody helped preserve my sanity more than my love, Megan Giehl. You understand me better than anybody I've ever known. No matter how late I had to work, you were always there for me with open arms and a warm meal ready. I'm so glad that I helped you through your one year of graduate school so that you could return the favor for my six. Let's call it even.

Finally, a warm thank you to my family that has ceaselessly supported me throughout the highs and lows of graduate school. Despite not understanding a word I said, you always showed a genuine interest in my work and I would repay you by diving even deeper into the gory details. I always told my parents that being in school for nearly thirty years was worth it if I could blow their minds with new ideas and concepts. To them, I say enjoy these 234 pages.

Jardin, Daniel   B.S., Physics, Syracuse University

A Likelihood Search for Low-mass Dark Matter
via Inelastic Scattering in SuperCDMS

   Advisor: Dr. Jodi Cooley

Doctor of Philosophy degree conferred August 6, 2019

Dissertation completed July 15, 2019

An abundance of evidence suggests that most of the Universe is composed of nonluminous matter. This "dark matter" is believed to be a new elementary particle and experiments around the world are attempting to directly detect rare collisions with terrestrial detectors.

The properties of dark matter have yet to be identified, thus efforts are ongoing to explore a range of possible masses and interaction cross-sections. For the latter, experiments can increase exposure by scaling up the detector mass and operating for a longer time. To search for dark matter with less mass than a nucleon, new technologies and analysis techniques need to be developed to be sensitive to energy deposits less than a few keV.

SuperCDMS is a direct detection experiment that measures ionization and phonon energy in cryogenic germanium crystal detectors. A special mode of operating the SuperCDMS detectors, called CDMSlite, amplifies the ionization signal via phonon creation. This amplification leads to a lower energy threshold, which provides sensitivity to smaller dark matter masses.

Typically, direct detection experiments assume dark matter scatters elastically off the nuclei in the detector. In this thesis, I will highlight the most recent CDMSlite elastic dark

matter search. Then I will describe how inelastic dark matter collisions can manifest in the detector and be useful for extending experimental sensitivity to lower dark matter masses. Finally, I will present a re-analysis of CDMSlite data using a profile likelihood to search for low-mass dark matter through two inelastic scattering channels: Bremsstrahlung radiation, and the Migdal Effect.

TABLE OF CONTENTS

LIST OF FIGURES

xviii

LIST OF TABLES

My grandparents, Homer and Murph Ouellette, instilled a sense of intellectual curiosity in me from a very young age. While they were not able to see me complete my degree, I know they were both extremely proud of me. I dedicate this in their honor.

PREFACE

The following are some useful concepts that should aid in the understanding of this thesis.

0.1. Natural Units

I will not always be explicit throughout this thesis, so it is important to understand the system of units that are typically used in high-energy physics. Physicists prefer to use units of measurement that are based on fundamental constants [1]. However, to simplify equations and avoid carrying these fundamental constants through mathematical proofs, they are set to unitless values of 1. The trade-off of making the math easier is that units then get confusing and everything is essentially in terms of energy. For example, a constant particularly relevant to this thesis is the speed of light, c. If $c = 1$, then mass and energy have the same unit (think of Einstein's famous $E = mc^2$). In order to convert back to the more common SI units, those physical constants need to be added back in. Therefore, even though mass is written as a unit of energy such as electron-volt (eV), it is actually eV/c^2, whether that factor of c is explicitly written or not.

0.2. Energy Units

In SuperCDMS, there are multiple "energy equivalent" units used. They will typically be indicated by a subscript added to the regular energy unit (in these examples I will use keV), although best practice is not always followed and context clues are needed. The most common unit I will use is keV_{ee}, which means "electron-equivalent" energy, or the energy assuming the interaction is an electron recoil (ER). In other words, the electron-equivalent

1

energy is the number of charges collected times the average energy needed to produce an electron-hole pair, which is 3 eV for germanium and 3.8 eV for silicon. If no subscript is included then it is most likely "ee".

0.2.1. NTL Gain

The second most common unit is keV_t, which means the "total" energy collected and is larger than the energy of the interaction. This is used when referring to the phonon energy, which is the sum of the original recoil energy and the Neganov-Trofimov-Luke (NTL) gain,

$$E_t = E_r + N_{eh}eV_b \,, \tag{0.1}$$

where E_t is the total energy measured with the gain, E_r is the initial recoil energy, N_{eh} is the number of electron-hole pairs created given E_r, e is the electron charge, and V_b is the bias voltage across the detector. From Eq. 0.1, the unit conversion between keV_{ee} to keV_t is

$$keV_t = keV_{ee}(\frac{\Delta V_b}{3} + 1) = keV_{ee}(g(V_b) + 1) \,, \tag{0.2}$$

where $g(V_b)$ is the NTL gain that takes advantage of the bias voltage across the crystal to amplify the ionization signal. Physically, as the liberated electrons and holes are drifted through the crystal by the electric field, they are quickly accelerated to their terminal velocity. Any additional work done by the electric field is then converted to phonons where the number of phonons created is proportional to the drifting electrons [2, 3]. A cartoon of this process is depicted in Fig. 0.1. For stronger voltage biases, these phonons dominate the number of phonons created in the original recoil, so this effectively converts the phonon collectors into ionization electrodes.
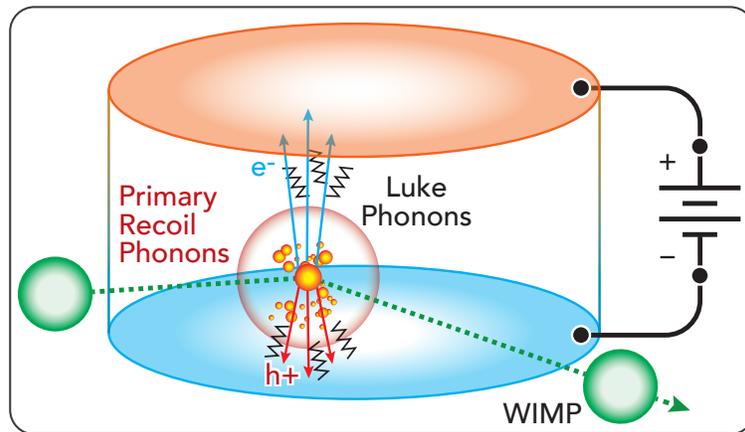
Figure 0.1: Cartoon showing a WIMP recoil creating primary phonons and electron-hole pairs in a crystal detector. As the electrons and holes drift through the electric field, additional ("Luke") phonons are created. [4].

0.2.2. Lindhard Model

In typical dark matter searches one needs to assume the interaction was a nuclear recoil (NR), which often means converting from keV_{ee} to keV_{nr} using

$$E_r\,[keV_{nr}] = E_r\,[keV_{ee}] \left(\frac{1 + g(V_b)}{1 + Y(E_{r,nr})g(V_b)} \right) . \tag{0.3}$$

The unknown variable $Y(E_{r,nr})$ is the ionization yield. When experiments take advantage of the NTL amplification with a higher bias voltage, they sacrifice the ionization electrodes and the ability to measure the yield for ER and NR events. Therefore, a yield model needs to be assumed and the most common is the Lindhard Model [5], where the yield is defined as

$$Y(E_{r,nr}) = \frac{k \times g(\epsilon)}{1 + k \times g(\epsilon)}, \tag{0.4}$$

where Z is the atomic number of the detector material, $\epsilon = 11.5 E_{nr}(keV)Z^{-7/3}$, and $g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon$ is determined by empirically fitting neutron scattering measurements of the relative efficiency between ER and NR events [6]. In the standard Lindhard Model, $k = 0.157$ for germanium and $k = 0.146$ for silicon.

The prediction of this model has been measured at high energies, but as more measurements are being made it appears the model overestimates the ionization yield in silicon at energies below a few keV as shown in Fig. 0.2 [7]. A similar array of measurements exists for germanium (Fig. 0.3), except there are no reliable measurements below 1 keV, which means that any analysis in this range needs to extrapolate [8].

All of this introduces a systematic uncertainty because different values for k still describe the experimental data. This uncertainty will be included as a nuisance parameter in the likelihood function in Chapter 4. Fortunately for the analysis presented in this thesis, Chapter 5 is searching for dark matter signals that manifest as ER signatures in the detector, so this extra systematic does not need to be included.
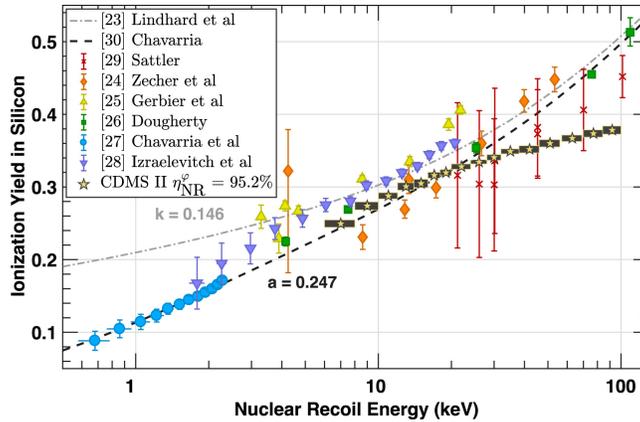
Figure 0.2: This plot shows the rough agreement above 4 keV between a large number of measurements and the predicted ionization yield in silicon from the Lindhard Model. In 2016, DAMIC was able to accurately measure the ionization yield at low energy, strongly disagreeing with the Lindhard prediction [9].
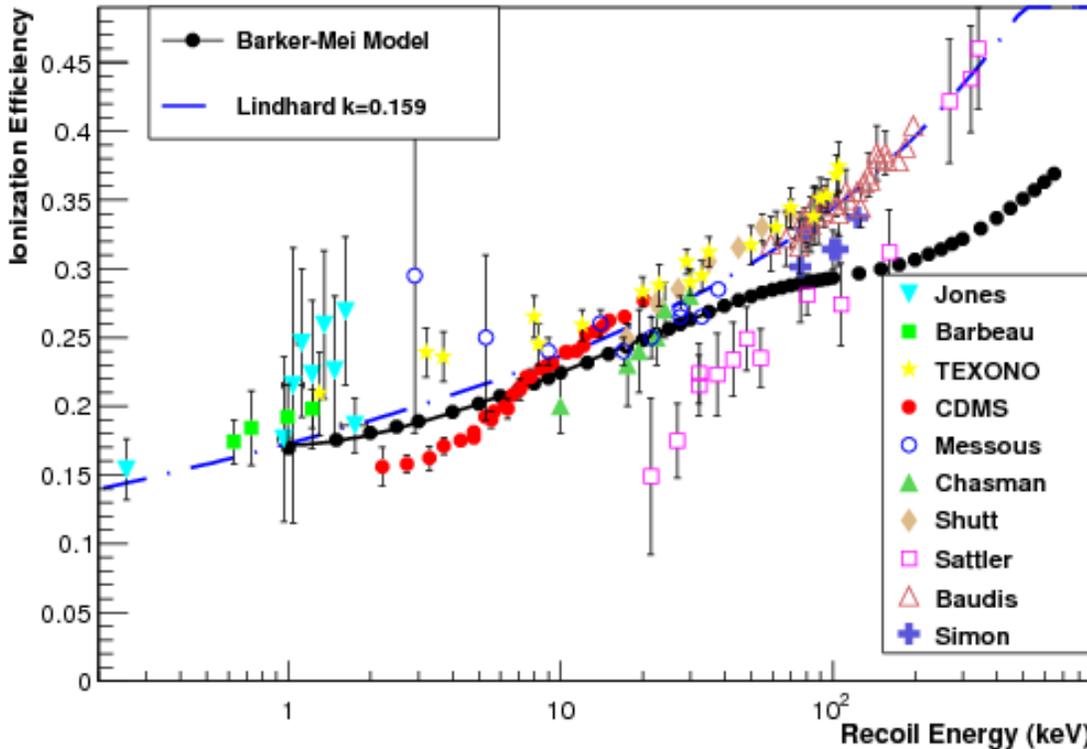
Figure 0.3: This plot shows the rough agreement between a large number of measurements and the predicted ionization yield in germanium from the Lindhard Model [8]. However, the data do not validate the model below a few keV where low-mass dark matter analyses are searching.

5

0.3. Cross-Section

In scattering theory, the number of events is determined from the product of luminosity and probability of interaction. For classical interactions, this probability is determined by the geometry of the particles, but actions-at-a-distance such as electromagnetism can complicate this picture. Therefore, physicists broadened the definition of cross-section from a slice of an object to be the area in which two particles must meet in order to interact. By defining the cross-section in this way, the interaction rate is independent of the intensity and luminosity of the incident particle sources [10].

The unit of cross-section is an area, where cm^2 will be the unit commonly used in this thesis. For a spin-independent dark matter direct detection experiment, introduced in Section 1.4.3 and discussed throughout this thesis, the cross-section referenced most often will define the interaction probability between a dark matter particle and a nucleon. There is also a small discussion in Section 5.1.1 about the cross-section between a photon and an electron in regards to the photoelectric effect.

DARK MATTER INTRODUCTION

To this day, the mystery of dark matter remains one of the great unanswered questions of science. Understanding its nature would provide information about the formation of the Universe, which is why thousands of scientists across the globe are dedicating their careers to the potential identification of this elusive entity.

1.1. Missing Mass

The story of dark matter begins in 1932 when Jan Oort studied the motion of stars in the Milky Way. He measured the redshift to determine how fast stars are moving and found that, for the amount of luminous matter observed, there is not enough mass to keep the stars gravitationally bound in orbit. He concluded there must be some additional nonluminous matter, and postulated that this could be in the form of stars that are dim or hidden behind other objects [11].

One year after Oort, Fritz Zwicky published a similar result by using the Virial Theorem to calculate the mass of the Coma Cluster. The Virial Theorem states that the average kinetic energy of a particle in a gravitationally-bound system is equal to negative one half the gravitational potential of that particle,

$$\langle K \rangle = -\frac{1}{2} \langle V \rangle \, , \tag{1.1}$$

where the average kinetic energy for a particle with mass m_p and velocity v_p is

$$\langle K \rangle = \frac{1}{2} m_p v_p^2 \, , \tag{1.2}$$

and the gravitational potential energy is

$$\langle V \rangle = -\frac{GM_{tot}m_p}{R_{tot}}, \tag{1.3}$$

where G is the gravitational constant, R_{tot} is the radius of the system, and M_{tot} is the total mass enclosed by R_{tot}. Substituting these into the Virial Theorem gives the relation

$$v_p^2 = \frac{GM_{tot}}{R_{tot}}. \tag{1.4}$$

Zwicky measured the redshift of the galaxies in the Coma cluster to calculate their velocities. The radius of the cluster can be calculated by measuring the angular size and estimating the distance from the luminosity. Then Eq. 1.1 is solved to obtain the total mass of the system. In doing this, Zwicky found that there was significantly more gravitating mass than he predicted from the luminosity alone. He was the first to hypothesize that the missing mass could be a new type of matter, and he is credited with coining the phrase "dark matter," or "dunkle materie" as he wrote in German. He proposed extending his study to individual galaxies, but he did not have the instrumentation required to measure the redshift at great distances [12].

Forty years later, Vera Rubin measured the rotational velocity of stars in individual galaxies as a function of their radius. An example of a "rotation curve" is shown in Fig. 1.1. This method did not require an assumption of the mass-to-light ratio, and was therefore more robust than previous estimates of nonluminous mass. It is clear that the data become constant as a function of radius, but the expectation for a disk of stars is that the rotational velocity should fall as $1/\sqrt{r}$ according to Eq. 1.4. Adding a gas component was not enough to account for the difference either, but the inclusion of a hypothetical spherical halo of uniformly distributed nonluminous mass brought the models into agreement with the data [13].

8
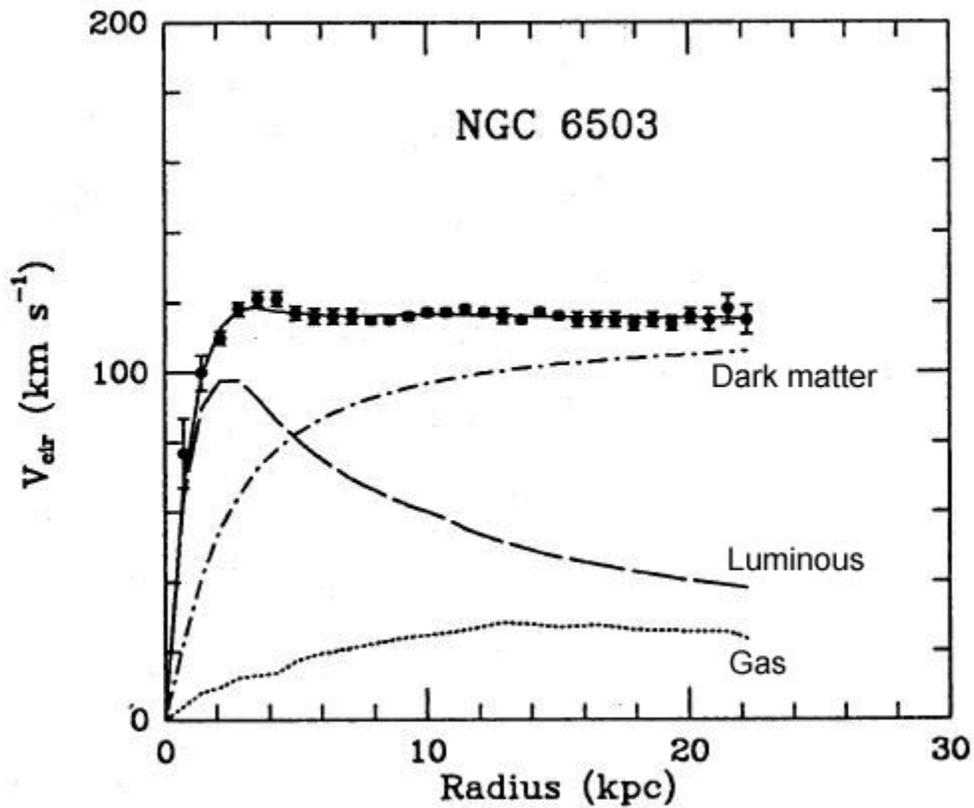
Figure 1.1: Shown is the rotation curve for NGC 6503 [14]. The black points are the measured velocity plotted as a function of radius. It can be seen that the disk and gas models would not sum to match the data shape, but adding a halo of hypothetical mass to those models results in the line that closely follows the black data points.

1.1.1. The Unknown Universe

Besides influencing the motion of individual galaxies or clusters, high-resolution n-body simulations have shown that additional mass is required to explain the large-scale structure of the Universe. In this model, shown in Fig. 1.2, the dark matter halos merge hierarchically to form more massive systems. The gravitational self-attraction of dark matter causes it to collect into long fibrous structures and areas of higher density are believed to initiate the formation of stars and galaxies [15, 16].
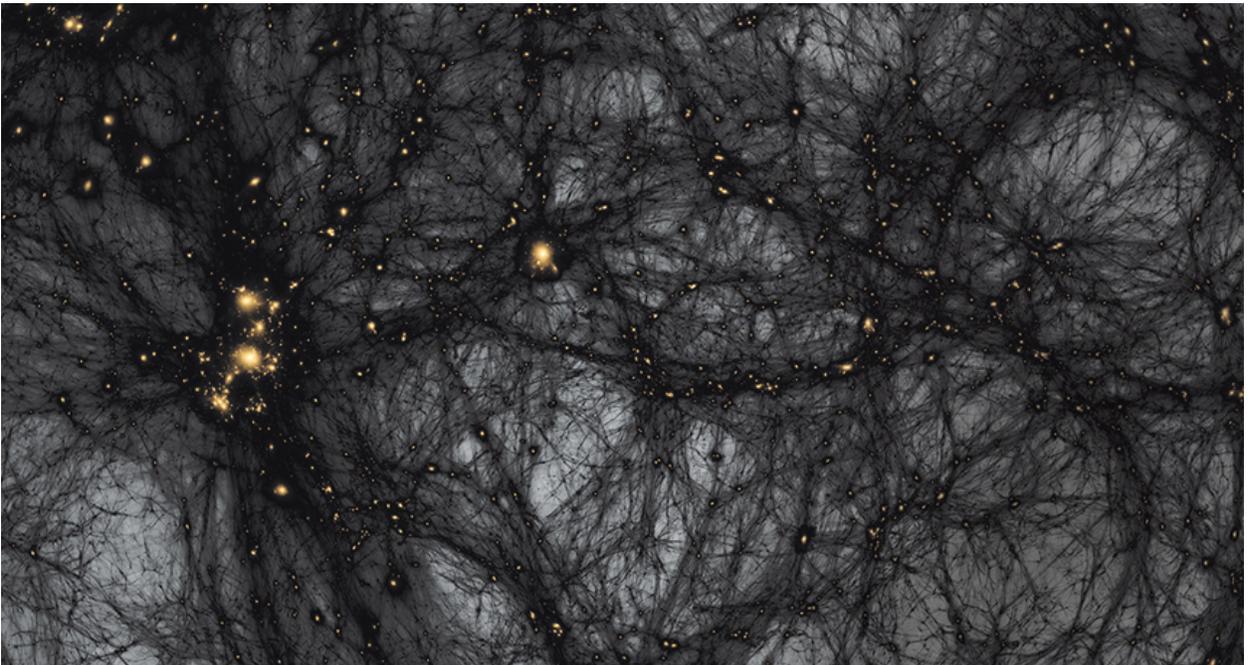
Figure 1.2: Visualization of an n-body simulation of large-scale structure formation in the Universe [16]. The black fibers are collections of dark matter halos, and higher densities of dark matter initiate the formation of stars and galaxies, indicated with yellow.

From the large-scale structure, it is natural to wonder how much mass in the Universe is actually composed of "unseen" matter. This information can be obtained from measurements of the Cosmic Microwave Background (CMB). The CMB is radiation that was emitted shortly after the Big Bang, when the Universe existed as a hot, dense plasma of nucleons. As the Universe expanded, this plasma cooled and the nuclei and electrons condensed primarily into hydrogen and helium. This is known as the epoch of recombination. During this time,

the Universe became transparent to the radiation originally present in the plasma. This radiation continues to travel nearly isotropically through space and can be measured to this day.

The discovery of the CMB was actually a fortuitous accident when radio astronomers Penzias and Wilson were working on a new type of horn antenna in 1964. In order to make accurate measurements of the Milky Way, they were trying to account for all possible background sources and found there was a 3 K excess they could not explain [17]. As it turned out, this was the first measurement of the CMB that was a prediction of the Big Bang Theory [18, 19].

To first order the CMB radiation is isotropic, but it carries information about the space it traverses in the form of polarization, gravitational lensing, and changes in the energy density. This manifests as small anisotropies on the level of 10^{-5}. Taking the harmonic transpose of these anisotropies provides a frequency power spectrum of the different sizes of features. Since the power spectrum from a hot gas is well understood, the location of peaks and relative amplitudes in the spectrum can be used to extract the information carried by the radiation.

The most recent measurement of the CMB was performed in 2013 by the Planck experiment. A map of the anisotropies is shown in Fig. 1.3. Analysis of the power spectrum indicates that all of the elements of the Periodic Table, all the particles of the Standard Model, and all known matter makes up only 4.9% of the Universe. Dark matter is determined to be five times more prevalent and compose 26.8% of the Universe. The remaining 68.3% is attributed to dark energy, which is thought to combat the attraction of gravity and cause the Universe to expand at an accelerated rate [20, 21].
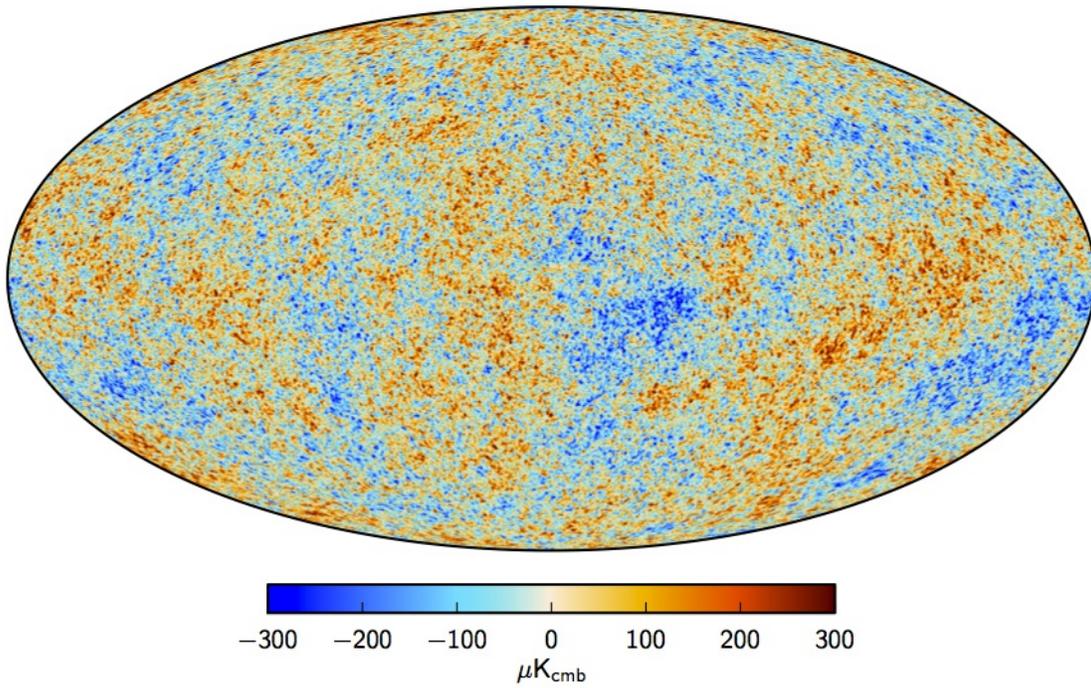
Figure 1.3: Map of the CMB, as measured by the Planck experiment [20, 21]. The color pattern indicates the location of anisotropic fluctuations of the energy density.

1.2. Evidence for New Matter

There are many proposed explanations for the astronomical mass discrepancies, but they generally fall into one of two categories. One very reasonable idea is that our current theory of gravity is flawed at large distances [22]. This is analogous to Newtonian mechanics breaking down at small scales, where the necessity for quantum mechanics arises. The other idea is that there is a new type of matter, as Zwicky first proposed. The observation of the Bullet Cluster (and other astronomical phenomena in this same category) provides a crucial piece of evidence to break the stalemate between the categories.

In order to understand the interpretation of the Bullet Cluster, however, one needs to understand gravitational lensing. Gravitational lensing is a phenomenon in which light is bent around massive objects due to the warping of spacetime, which manifests as a physical force of attraction around those objects. In the most extreme case, objects can be observed despite being hidden directly behind another object. Figure 1.4a depicts an example of gravitational lensing where the light from a galaxy bends around a galaxy cluster. One effect of this is that the light gets bent around all sides, which can result in multiple images or an annulus called an Einstein Ring as shown in Fig. 1.4b. The curvature of the ring can be used to calculate the mass of the object responsible for the lensing. Lensing is more direct for mapping the location of massive nonluminous objects than inferring their location from the relative motion of nearby objects.

The Bullet Cluster shown in Fig. 1.5 provides a unique view into the behavior of dark matter interactions [23, 24]. Perhaps a misnomer, the Bullet Cluster is actually a collision of two smaller galaxy clusters. The stars, observed as visible light, interact gravitationally and are relatively unaffected by the collision because of the astronomical distances between them. Most of the baryonic matter in the clusters is gaseous and interacts electromagnetically. This stronger interaction causes the gas to slow down in the collision, as well as heat up and emit x-rays that are shown in pink. The blue coloring is a gravitational lensing map that shows
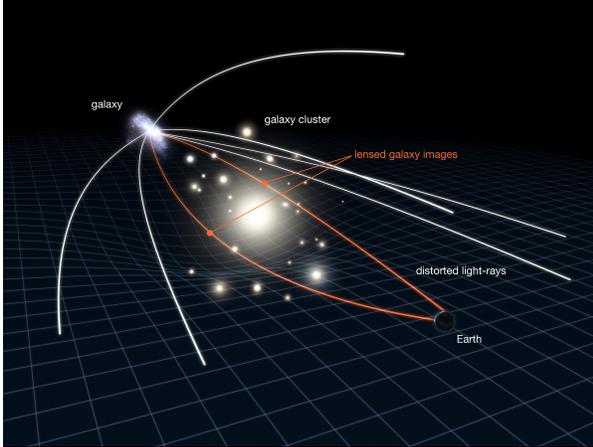
(a) Cartoon depicting the gravitational lensing of light from a galaxy around a massive galaxy cluster. The cluster distorts the fabric of spacetime, causing the light radiating from the galaxy to bend inwards. The red lines indicate light that has been bent enough that it reaches Earth on the opposite side of the cluster. Image: NASA/ESA
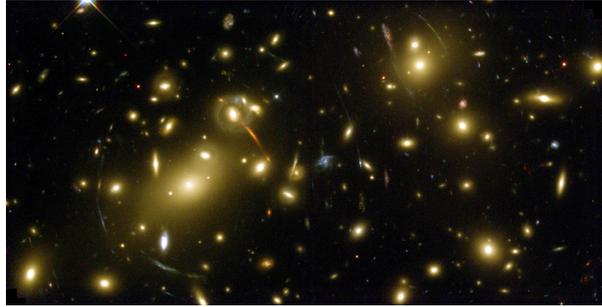
(b) Optical image of an Einstein Ring. The multiple images of the hidden light source are curved so much they nearly form a perfect circle. The curvature of the lensed image can be used to calculate the mass of the massive object responsible for the lensing. Image: NASA/ESA

Figure 1.4: Description and example of gravitational lensing that is useful for mapping higher concentrations of mass in the Universe.

the strongest concentration of mass. If the mass discrepancies were explained by a modified theory of gravity then the lensing would be strongest near the baryonic matter, which is primarily the gas. The fact that there is clear separation between the pink and blue is extremely difficult to explain using modified gravity, which implies that dark matter is a new type of matter that can be separated from baryonic matter. The observation of this collision can be taken a step farther and used to calculate a limit on the rate of dark matter interaction [23]. Other examples of colliding clusters have been observed, although the Bullet Cluster remains the most well-known. [1]

In addition to the observation of the Bullet Cluster, the power spectrum of the CMB cannot be explained by current theories of modified gravity.

---

[1]After the discovery of the Bullet Cluster in 2006 came the Trainwreck Cluster [25] and the Musketball Cluster [26]. Then they stopped receiving creative names.
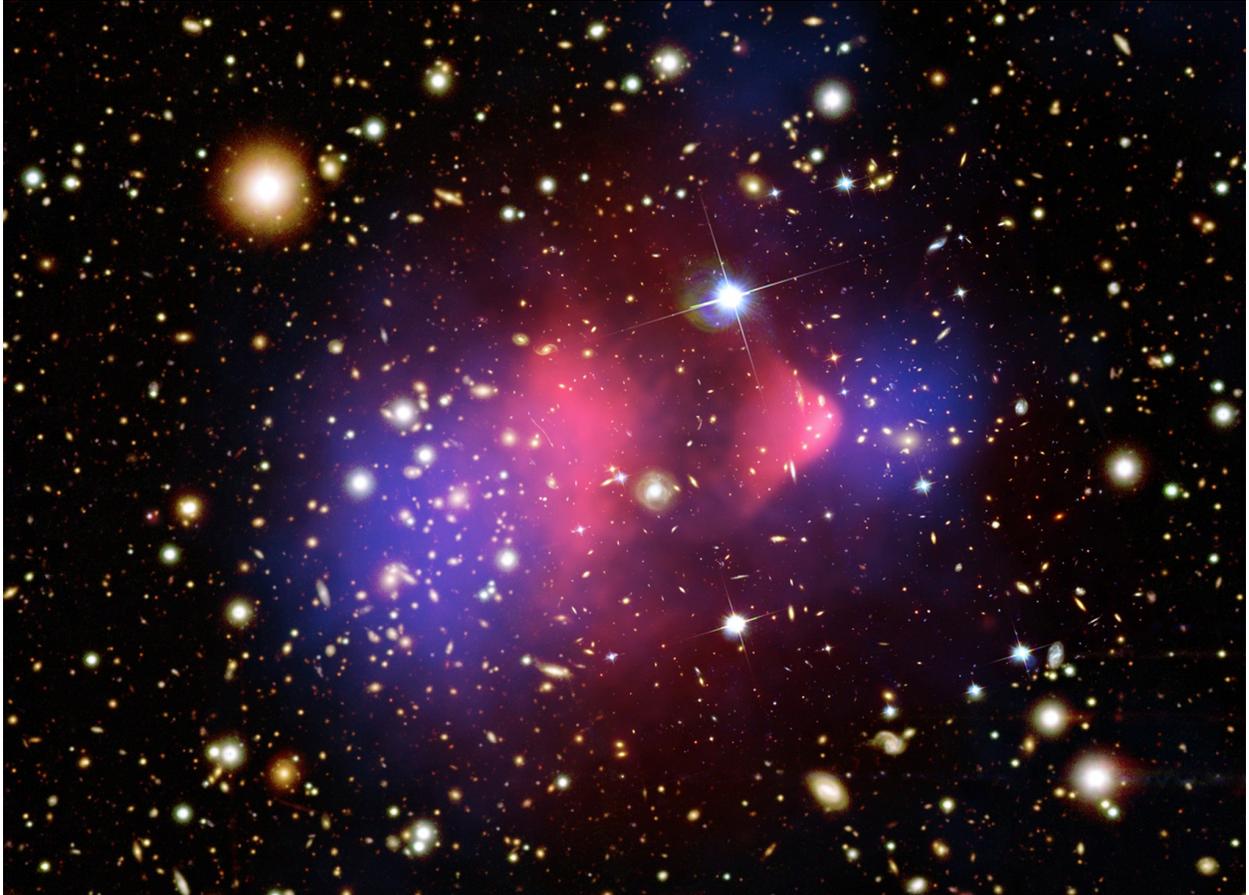
Figure 1.5: Composite image of the Bullet Cluster, which is actually a collision of two clusters [24]. The pink is an x-ray image that shows radiation from the gas that heats up in the collision. It is projected over an optical image that shows the stars in the colliding clusters. The blue is a gravitational lensing map that shows the distribution of mass. The separation of the pink and blue is difficult to explain with a modified theory of gravity, suggesting that dark matter is a new type of matter.

1.3. Weakly-Interacting Massive Particle

If dark matter is indeed a new type of particle, then it needs to have the following properties to be considered a valid candidate.

Electrically neutral. If dark matter had significant electric charge then it would interact electromagnetically and almost certainly would have been detected already.

Non-relativistic. "Cold" dark matter that moves much slower than the speed of light is preferred over "hot" (relativistic) dark matter because it lends to hierarchical structure formation where smaller objects merge into larger objects [27].

Provide the correct relic density. We know how much dark matter there is in the Universe from measurements of the CMB, so any theory needs to match that estimate.

Compatible with constraints on interaction rate. Limits have been placed on the strength of dark matter interactions from observations of collisions such as the Bullet Cluster.

Experimentally detectable. In order to be a valid hypothesis, it needs to be experimentally validated. If dark matter cannot be physically detected, then its existence cannot be proven.

Consistent with current experimental results. Experiments have been searching for dark matter for years and have excluded a lot of possible parameter space. The proposed candidate must lie in unexplored parameter space or else it has already been experimentally ruled out.

For many years, the leading candidate for dark matter has been the Weakly-Interacting Massive Particle (WIMP) [28]. WIMP refers specifically to a theoretical dark matter particle that was in equilibrium in the early Universe. As the Universe expanded and cooled, the primordial WIMPs fell out of thermal equilibrium until they no longer self-annihilated into

Standard Model particles, "freezing out" to the current density of dark matter. This model is clarified in a review published by Bertone, Hooper, and Silk, "A particle species in the early Universe has to interact sufficiently or it will fall out of local thermodynamic equilibrium. Roughly speaking, when its interaction rate drops below the expansion rate of the Universe, the equilibrium can no longer be maintained and the particle is said to be decoupled" [29]. Additionally, the relic density can be used to calculate the self-annihilation cross-section if the WIMP mass is known or assumed.

An added benefit of the theorized WIMP is known as the "WIMP Miracle." This is the coincidence that a new particle with a mass around 100 GeV that interacts via the weak force gives the correct relic density. Physicists appreciate when independent ideas unify, so it would be serendipitous for weakly-interacting massive particles and weak interactions in the Standard Model to be due to the same underlying interaction. Figure 1.6 shows the number density of comoving space as a function of time [30], which means the number of particles in a volume of space that increases with the expansion of the Universe. Notice that a large range of WIMP masses does not make as big of an impact on the relic density as the strength of the interaction.

1.4. Dark Matter Detection Methods

Since the nature of dark matter has yet to reveal itself, scientists all over the world are taking a three-pronged approach to searching for it. Figure 1.7 is not to be interpreted literally as a Feynman diagram, but it shows the relationship between the three methods. Each of them are explained in further detail below.

1.4.1. Production

Among many other measurements, particle colliders such as the Large Hadron Collider (LHC) are searching for the production of dark matter. Since dark matter interactions are
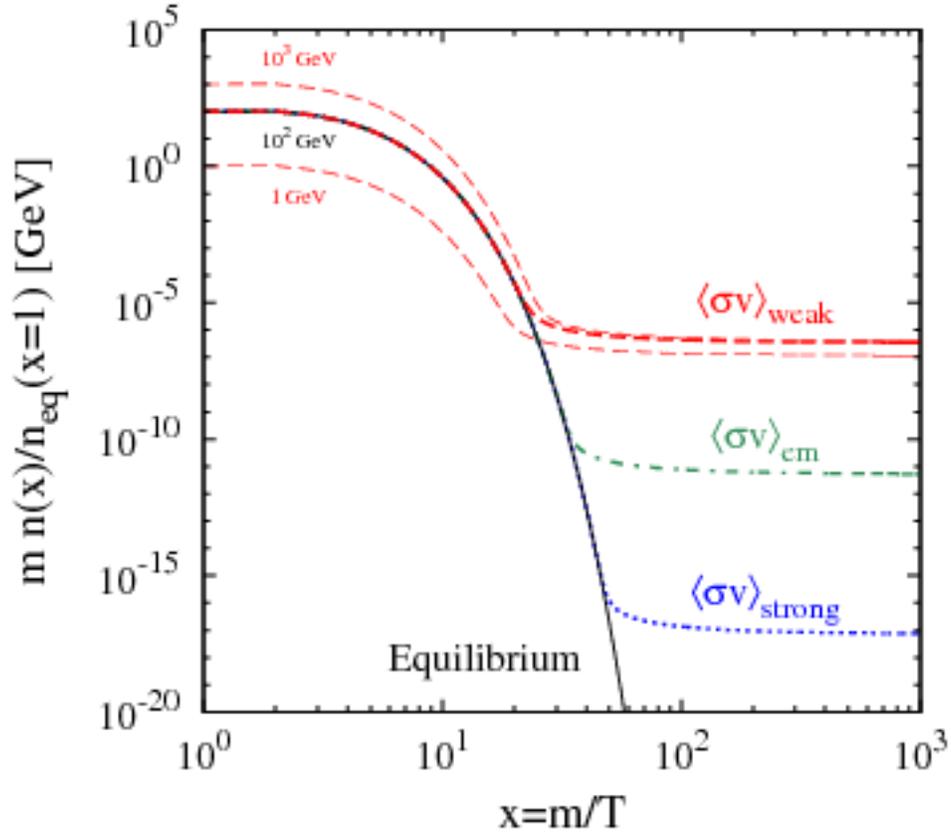
Figure 1.6: Demonstration of the theory of WIMP freeze-out by plotting the number density of dark matter as a function of time. [30]. Time increases to the right, which is inversely proportional to the temperature of the expanding Universe. In the earliest times, the density was a flat line indicating constant density. After that, the Universe expands and WIMP annihilation occurs more than creation. The WIMP density drops until the Universe is large enough that the WIMP self-interaction rate drops below the expansion rate of the Universe. It is at this point that WIMPs "freeze out" and remain at relic density.
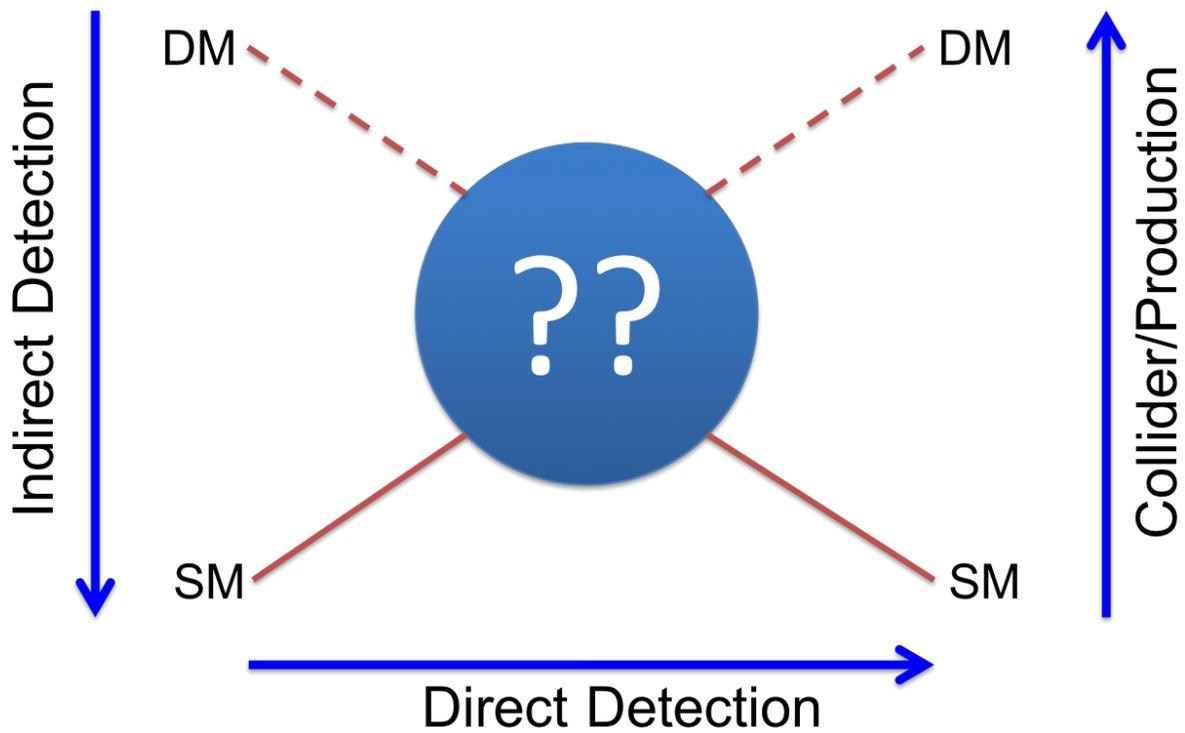
Figure 1.7: Production searches are based on colliding baryonic matter and looking for an imbalance of transverse momentum. Indirect searches attempt to detect byproducts of dark matter annihilation or decay. Direct detection experiments look for specific energy signatures from a dark matter particle colliding with the detector.

so rare, there is no direct signal then deposited in the experiment. Instead, the existence of dark matter must be inferred from the appearance of missing energy and momentum from what was known to go into the initial collision. Particles in a collider travel down a beam pipe where all momentum is in the direction of the beam, so the transverse momentum after the collision must sum to zero. Therefore, the transverse momentum of all leptons and hadrons produced in the collision can be added and any imbalance indicates the direction that an "invisible" particle must have carried some momentum away. However, this method of searching for dark matter is complicated by the creation of neutrinos that will also carry momentum away undetected and detector inefficiency. A schematic of the momentum imbalance is depicted in Fig. 1.8 [31].
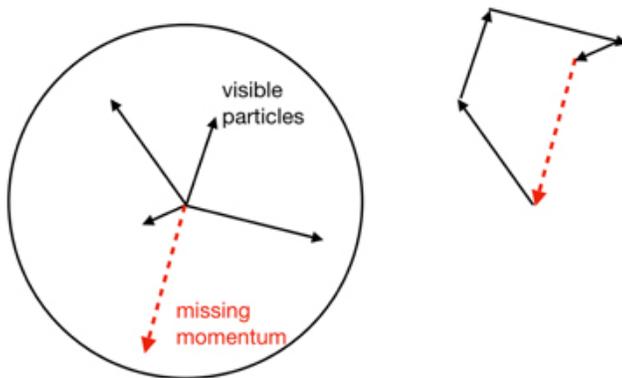
Figure 1.8: Schematic of the inferred missing momentum that is presumably carried away from the detector by dark matter [31]. From this perspective, the beam line is perpendicular to the page, at the center of the circle. The black arrows are transverse momentum vectors that are visible in the detector and the red vector is the missing momentum required to bring the sum to zero. The inset shows the graphical method of tip-to-tail vector addition.

1.4.2. Indirect Detection

Indirect dark matter searches look for byproducts from dark matter interactions. This could include annihilation from self-interactions, or decay of dark matter from an excited state. Dark matter can exist in higher densities in regions of the Universe where baryonic matter is more dense, such as the center of a galaxy. Therefore, indirect searches focus

on these areas where dark matter has a higher probability of interacting with itself. If the theory of freeze-out is correct, then dark matter can self-annihilate and produce a variety of byproducts such as neutrinos, photons, or other particles in the Standard Model. Different types of telescopes are used in an effort to observe higher concentrations of these products to triangulate the location of the dark matter [32].

1.4.3. Direct Detection

Since the Earth is orbiting around the Sun and the Sun is orbiting around the center of the Milky Way, the Earth is constantly passing through the hypothetical dark matter halo that engulfs the entire Milky Way galaxy and extends far beyond its bounds. From the reference frame of the Earth, the dark matter particles that compose this halo are constantly passing through us; this is called the "WIMP wind." In general since they have such a low interaction rate, as determined from collisions like the Bullet Cluster, they should mostly pass through the Earth without scattering. However, there is a rare chance that the incoming particles can collide with an atom and transfer energy. Electrically neutral dark matter and neutrons are examples of hypothetical and well-known particles, respectively, that scatter off of the massive nucleus of the target material in what is called a nuclear recoil (NR) event. Meanwhile, photons and electrons and most other backgrounds that couple electromagnetically will scatter off the electron cloud, leading to an electron recoil (ER) event. A cartoon of the two recoils is depicted in Fig. 1.9. The goal of direct detection experiments is to build ultra-sensitive terrestrial detectors that are capable of measuring both kinds of energy transfers.

The energy from these interactions is typically read out in three different ways. The first is ionization due to the collection of electrons that were freed during the interaction. The second is scintillation, which is the collection of photons emitted during the interaction. The third type of energy is in the form of vibration or heat. This manifests as quantized packets,
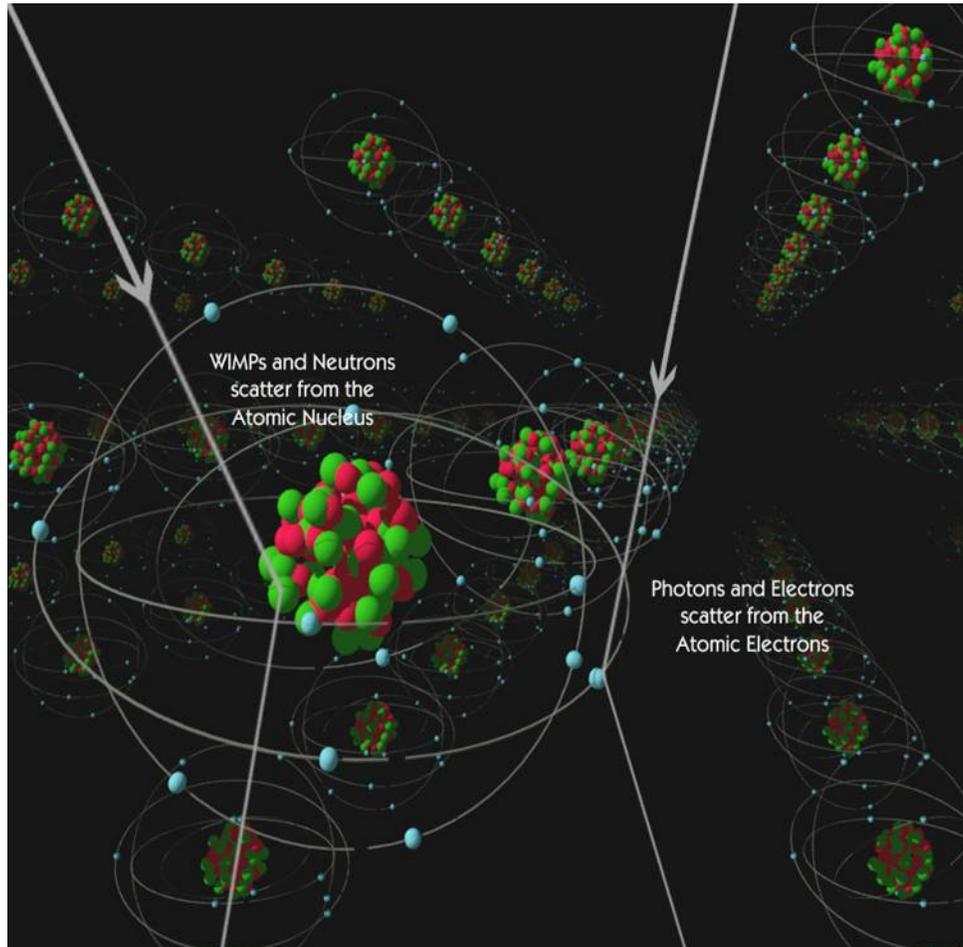
Figure 1.9: Cartoon of the different atomic interactions, depending on the incoming particle. Neutral dark matter is expected to collide with the larger nucleus of one of the atoms in the target material; this is known as a nuclear recoil event. Charged particles that interact electromagnetically will scatter off of the electron cloud; this is known as an electron recoil event. Credit: M. Attisha.

called phonons, that travel through matter like ripples on a pond. They propagate outward from the point of the interaction until they are collected or reflected.

Experiments selectively choose the materials used to construct detectors because the choice then dictates the possible signals they can read out. Most experiments measure energy in two ways because the ratio allows for particle identification, as different types of interactions will lead to different amounts of each signal. This is a useful tool for discriminating between background events and signal events, which is critical when searching for rare phenomena such as dark matter collisions. Figure 1.10 shows the different technologies used by direct detection experiments. For the sake of complementarity in the community, it is also good to have a wide variety of direct detection methods [33].
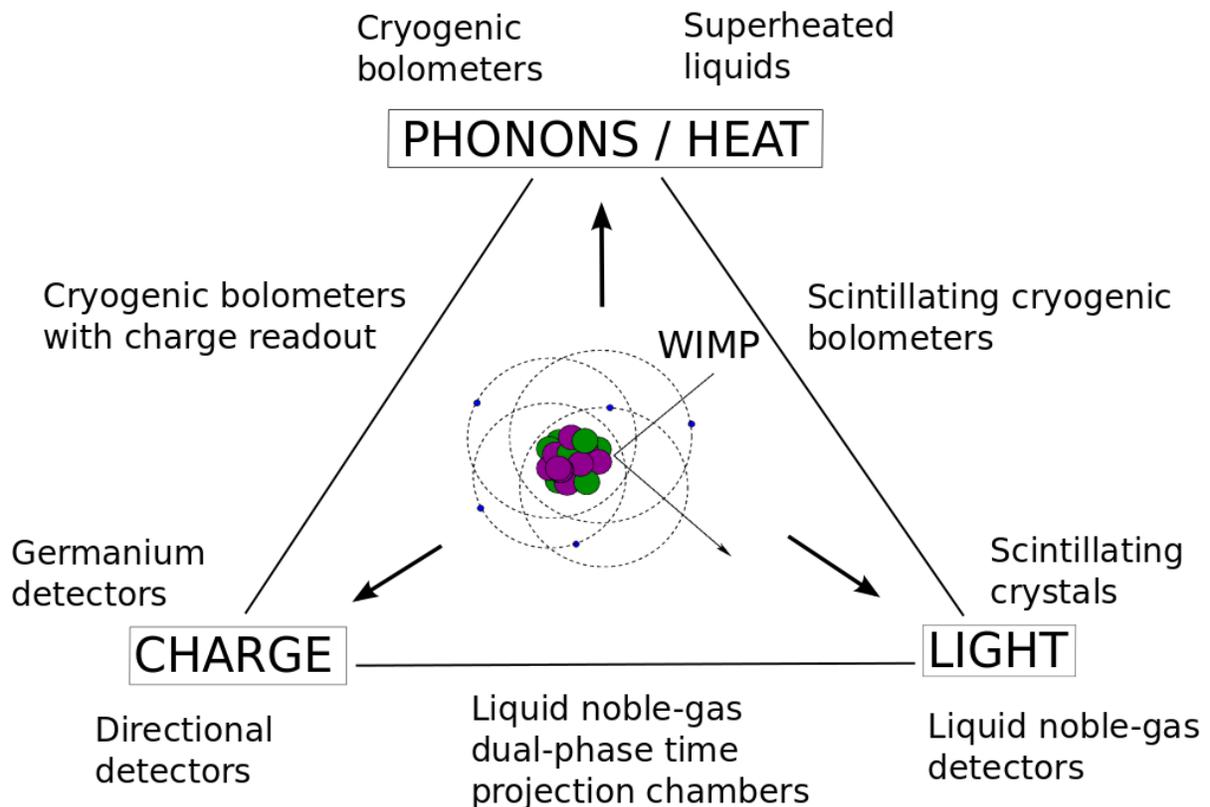
Figure 1.10: Triple Venn diagram of energy signatures caused by the collision of dark matter and the technologies that are used to measure those signals [33].

As can be seen in Fig. 1.10, there are two main categories of detector technologies around the outside of the triangle: cryogenic bolometers and noble liquid detectors. Noble liquid experiments have high scintillation and ionization yield in a large homogeneous target. This type of detector has the advantage of comparatively easy scalability to quickly increase mass and exposure, which allows them to search for dark matter with even smaller interaction cross-sections. The downside of this detector is the light yield is not well understood for small recoil energies, so the detectable energy threshold is a few keV that corresponds to a dark matter mass minimum of roughly 10 GeV [34].

The other main category of direct detection technologies is a cryogenic bolometer, which usually comes in the form of a solid crystal. The detector is operated at tens of mK to reduce the thermal motion of the crystal and allow the measurement of phonons created from small recoil energies. Semiconductor crystals have the advantage of supporting large bias voltages to amplify small signals and provide sensitivity to dark matter masses below 10 GeV. However, it is difficult to readily increase the mass of crystal-based experiments because this requires fabricating more detectors, which is a complex and expensive process [35].

1.4.3.1. WIMP Rate

The expected rate of WIMP collisions in a terrestrial detector can be calculated using information from particle, nuclear, and astrophysics [6]. The number of expected collisions over some energy range from E_1 to E_2 is

$$N = \int_{E_1}^{E_2} \frac{dN}{dE} dE = \int_{E_1}^{E_2} \frac{dR}{dE_r} \varepsilon(E) dE_r \,, \tag{1.5}$$

where E_r is the recoil energy, and ε is the exposure of the experiment, which is the product of the detector mass and the livetime of operation. The differential rate is defined as

$$\frac{dR}{dE_r} = N_t \frac{\rho_\chi}{m_\chi} \int_{v_{min}}^{v_{max}} \frac{vf(\mathbf{v}, \mathbf{v}_E)}{k} \frac{d\sigma}{dE_r} d^3\mathbf{v}\,, \tag{1.6}$$

where N_t is the number of atoms in the target mass, ρ_χ is the local dark matter mass density (0.3 GeV/cm^3), and m_χ is the dark matter mass.

The dark matter velocity distribution $f(v)$ is typically assumed to be the Maxwell-Boltzmann distribution. However, dark matter physicists are using primordial stars in the GAIA all-sky survey to trace the motion of dark matter and empirically determine the velocity distribution [36]. The velocity distribution gets integrated from v_{min} to v_{max}, respectively defined in Eq. 1.7 and Eq. 1.8. The lower velocity cutoff v_{min} is the velocity that kinematically produces the smallest detectable recoil energy in the detector,

$$v_{min} = \sqrt{\frac{m_N E_r}{2\mu_N^2}}\,, \tag{1.7}$$

where m_N is the mass of the nucleus, and μ_N is the reduced mass of the dark matter and the target nucleus. The upper velocity cutoff v_{max} is the highest relative velocity of a dark matter particle that is gravitationally bound by the galaxy,

$$v_{max} = v_{esc} + v_E\,, \tag{1.8}$$

where v_{esc} is the escape velocity of the Milky Way (544 km/s), and v_E is the mean orbital velocity of the Earth (232 km/s). The variable k in Eq. 1.6 is a normalization for the truncated velocity distribution.

The differential cross-section in Eq. 1.6 is where assumptions about the dark matter interaction are contained. This will be revisited, but for now the signal is presumed to be

an elastic nuclear recoil with a differential cross-section,

$$\frac{d\sigma}{dE_R} = \frac{m_N}{2\mu_N^2 v^2}[\sigma_{SI}F_{SI}^2(E_r) + \sigma_{SD}F_{SD}^2(E_r)]\,, \tag{1.9}$$

where F is the model-dependent nuclear form factor that encodes the dependence on the momentum transfer of interactions with the nucleus. The spin-independent cross-section σ_{SI} is defined as

$$\sigma_{SI} = \sigma_0 \left(\frac{\mu_N}{\mu_n}A\right)^2\,, \tag{1.10}$$

where μ_n is the reduced mass of the dark matter with a nucleon, A is the atomic mass, and σ_0 is the cross-section of the dark matter with a nucleon that is independent of target material. Along with the dark matter mass m_χ, σ_0 is the variable that direct detection experiments are probing.

The spin-dependent cross-section σ_{SD} is not explored in this document, but the contribution to the differential cross-section is proportional to the spins of the proton and neutron. The spin dependence is defined as

$$\sigma_{SD} \; \alpha \; (a_p\langle S_p\rangle + a_n\langle S_n\rangle)^2\,, \tag{1.11}$$

where a_p and a_n are model parameters for the proton and neutron interactions, respectively [6].

1.5. Low-mass Dark Matter

The current state of direct detection WIMP searches is shown in Fig. 1.11, demonstrating the sheer number of direct detection experiments around the world. Displayed are limits where the dark matter properties above each curve have been excluded. The canonical 100 GeV WIMP from the "WIMP Miracle" has been excluded down to 10^{-46} cm^2. Despite

the null results, direct detection experiments continue to push to explore this parameter space for the elusive dark matter particle.

The yellow area looming at the bottom of Fig. 1.11a is what is commonly called the "neutrino floor." This region is where direct detection experiments expect to detect cosmogenic neutrinos [37, 38]. Due to the neutral charge and small interaction rate, neutrinos are a very difficult background source to contend with because they are expected to mimic dark matter collisions in direct detection experiments. There are two main ideas proposed to mitigate the neutrino background. The first idea is to use directional information to reject events coming from known neutrino sources [39]. The other idea is to take advantage of the complementarity of different detector materials by measuring the background with a detector that is more sensitive, and using that to estimate the rate in a detector that is less sensitive [40]. Neutrinos have recently been observed coherently scattering off the nucleus of a detector, with the intention of better understanding the interaction process [41].
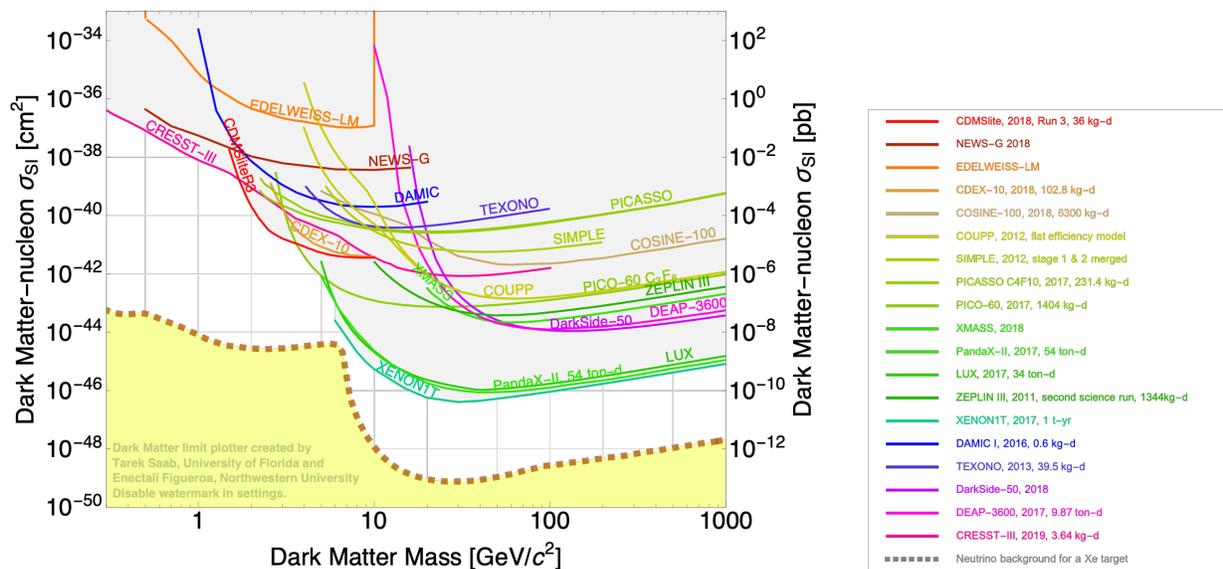
(a) The current state of WIMP-search limits, plotted as DM-nucleon cross-section versus dark matter mass [40, 42–53].

(b) Legend for Fig. 1.11a.

Figure 1.11: This is the current state of direct detection WIMP searches.

A selection of projected sensitivities is shown as dashed lines in Fig. 1.12 along with the published results that compose the current world-leading limit across the range of masses. This shows that the noble liquid experiments plan to continue scaling up to larger exposures until they eventually reach the neutrino floor. Since masses > 10 GeV are so well constrained, other detector technologies have begun to explore smaller dark matter masses (e.g. SuperCDMS, CRESST, NEWS-G, DAMIC, etc.). Due to the kinematics of the interaction, low-mass dark matter deposits very small amounts of energy in the detector, which is a technical challenge to measure. Experimentalists are working diligently to understand and reduce background noise in order to gain sensitivity to these small signals, but as seen in Fig. 1.12, conventional direct detection methods are not sensitive below 100 MeV. There is a lot of unexplored parameter space even below 1 GeV ("sub-GeV").
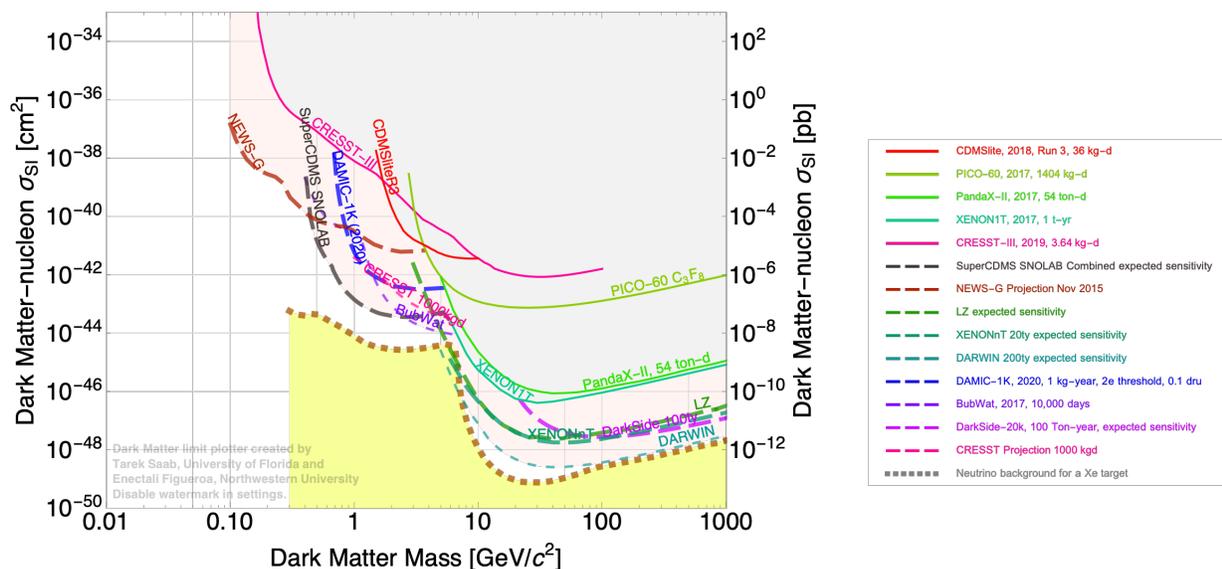
(a) Figure of the world-leading limits for a range of possible dark matter masses, and a selection of projected sensitivities [40, 46, 49, 53–56].

(b) Legend for Fig. 1.12a.

Figure 1.12: The current world-leading limits displayed with a selection of projected sensitivities.

The classical WIMP is not strongly motivated in the low-mass region, but there have been many alternative dark matter candidates proposed [57]. A few of the more popular sub-GeV mechanisms are summarized below.

1.5.1. Freeze-In

Section 1.3 introduced a heavy dark matter particle that self-annihilated into Standard Model particles in the early Universe. In the case of low-mass dark matter, the opposite can be true where the relic abundance was "frozen in" through annihilation or decay of Standard Model particles into dark matter. Similar to the freeze-out model, and compared in Fig. 1.13, this process of freeze-in ceases when the Universe expands faster than the rate of interaction. Freeze-in only works for sub-GeV dark matter because energy conservation prevents a Standard Model particle from decaying into a heavier particle. [57, 58].

1.5.2. Asymmetric Dark Matter

Instead of obtaining the correct relic abundance from a thermal freeze-out or freeze-in process, asymmetric dark matter is the idea that the current dark matter density was set by a primordial asymmetry. In this paradigm, the dark matter and baryon abundances are related because they are in chemical equilibrium in the early Universe. As with the asymmetry of baryonic matter, it is believed there was a symmetric component that was annihilated, leaving behind the current density. While the interaction process that created the original asymmetry may not be detectable, it is possible that the annihilation process might be, assuming a coupling to the Standard Model. [57, 59–61].
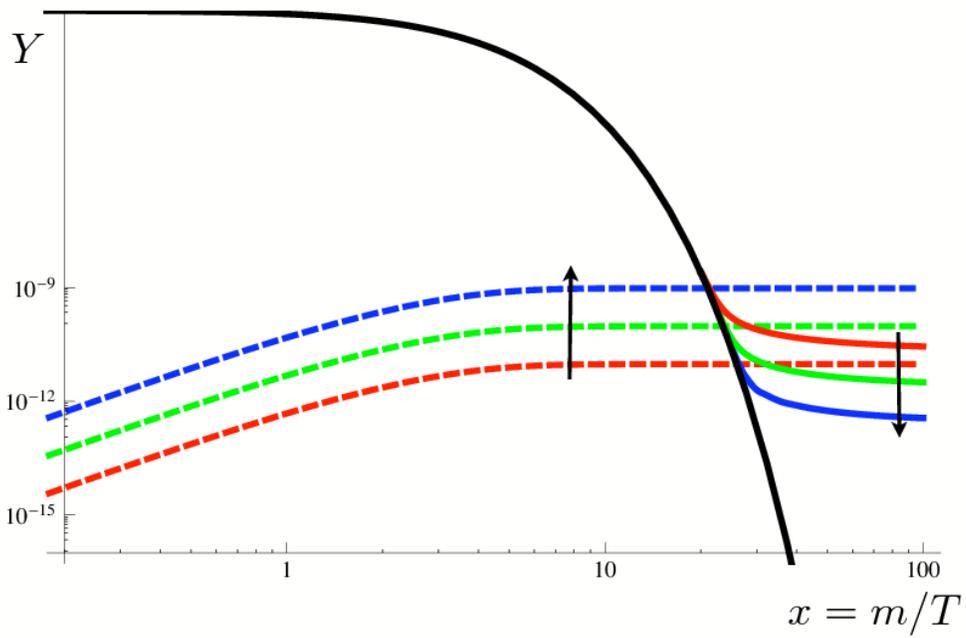
Figure 1.13: Plot comparing freeze-out (solid lines) with freeze-in (dashed lines) [58]. The arrows indicate the direction of increasing strength of interaction.

1.5.3. Hidden Sector

Since dark matter has not been detected yet through interactions with ordinary matter, it is natural to wonder if it does not interact via the known forces in the Standard Model at all, but perhaps only via a new force that is yet to be discovered. The hidden sector is a hypothetical set of new particles and forces that mirrors the Standard Model and includes dark matter. While it is possible that hidden-sector dark matter does not interact with the Standard Model, a hypothetical portal could play a key role in determining the relic abundance. This new force could be responsible for annihilating the symmetric component in the asymmetric dark matter described in Section 1.5.2 or mediating the production of dark matter from Standard Model particles as described in Section 1.5.1. [57, 62, 63].

The detection of hidden-sector dark matter depends on the type of new force, and how it couples dark matter to the Standard model, and the properties of the dark matter candidate. Some theories introduce a new symmetry that leads to a new gauge boson known as a dark photon. The dark photon can theoretically mix with a Standard Model photon. Therefore, a coupling exists that could be observed in a detector [64]. The observable signature would be a peak in the energy spectrum at the energy of the dark photon.

1.5.4. DM-electron Scattering

For a low-mass dark matter interaction, the kinetic energy is transferred much more efficiently to an electron than to the nucleus of the target atom [65]. However, the typical recoil energy of an electron in the outer shell of an atom is a few eV, which is currently below the energy threshold of the direct detection experiments in Fig. 1.11. Higher energy recoils are possible, but their rate is suppressed by the atomic form factor [57]. Therefore, new technology will need to be developed in order to take advantage of this process [66]. Also,

since this process probes a different type of interaction, it cannot be directly compared to the limits in Fig. 1.11 and Fig. 1.12, where the y-axis is DM-nucleon interaction cross-section.

1.5.5. Inelastic Scattering

One of the assumptions of the canonical WIMP described in Section 1.3 is that it scatters elastically off of the target nucleus. Assuming the dark matter particle transfers 100% of its momentum to the target mass, then the energy of the recoiling nucleus is

$$E_r = \frac{|q|^2}{2m_N} = \frac{m_\chi^2 v^2}{2m_N}, \tag{1.12}$$

where m_N is the mass of the nucleus, m_χ is the dark matter mass, v is the velocity of the dark matter, and $q = m_\chi v$ is the momentum of the incoming dark matter particle.

Equation 1.12 shows that the recoil energy is proportional to the square of the dark matter mass, so the elastic recoil energy of low-mass dark matter is extremely difficult to detect. Until direct detection technology is able to sufficiently reduce background noise and energy thresholds, low-mass searches can benefit from analysis techniques and theoretical models of different dark matter interactions.

One such mechanism that has been explored recently is the inelastic scattering of dark matter with a nucleus and how that signal would manifest in a conventional direct detection experiment. Unlike elastic collisions, kinetic energy does not have to be conserved, which allows for more exotic interactions. For example, an additional particle can be created in a 2→3 body scattering process where the maximum possible energy of any particle emitted is the full energy of the system. Considering the detector to be at rest in the laboratory frame, the energy of the system is simply the kinetic energy of the incoming dark matter particle,

$$E_\chi = \frac{1}{2}m_\chi v^2. \tag{1.13}$$

Combining Eq. 1.12 and Eq. 1.13 gives

$$E_r = E_\chi \frac{m_\chi}{m_N} \, . \tag{1.14}$$

By making the assumption that we are limited to the low-mass dark matter regime, then $m_\chi \ll m_N$ and therefore $E_r \ll E_\chi$. In other words, for a low-mass dark matter particle with some given kinetic energy, an inelastic scatter will lead to a larger detectable energy than an elastic scatter. This is because the elastic scatter is suppressed by the relatively large mass of the nucleus, but the inelastic scatter avoids this via the example $2\rightarrow3$ body interaction. This larger detectable energy provides the sensitivity to search for dark matter masses that would otherwise lead to nuclear recoil energies that are below the energy threshold of the detectors.

As indicated by the title, this thesis will describe how inelastic scattering was used to search for low-mass dark matter in SuperCDMS data.

CHAPTER 2

RADON

One of the most troublesome background sources that dark matter direct detection experiments contend with is radon. Radon is a noble gas that is produced from the decay of trace amounts of natural uranium and thorium that can be found in the Earth's crust [67]. These parent isotopes have extremely long half-lives, on the order of a billion years, so a constant supply of radon exists for the lifetime of the Earth or Sun. In its natural state, radon is an inert gas that can diffuse through any barrier that is not airtight. Since it is created in the Earth's crust, it often seeps through cracks in the foundations of houses, hence why it is typically found in higher concentrations in basements [67].

The most detrimental radon isotope for direct detection dark matter experiments is ^{222}Rn, often referred to as simply "radon." Not only is this the most stable radon isotope, but the decay chain in Table 2.1 is supported by the long half-life of one of the decay daughters, ^{210}Pb. This prevents any radioactive contamination from decaying away to a stable isotope within the lifetime of a typical experiment.

As radon decays, its progeny can plate-out (or "stick") onto surfaces such as experimental components and contaminate them with radioactivity. Due to momentum conservation, decays cause the daughter to recoil in the opposite direction. In the case of high-energy alphas, if the parent isotope was plated out, then this recoil could implant the daughter below the surface of the material. Worse, if the recoiling atom is directed towards the detector, then it could cause a nuclear recoil event that mimics the expected dark matter interaction. The most problematic case of this is the ^{210}Po alpha decay to a recoiling ^{206}Pb because it is steadily sourced by ^{210}Pb. Due to effects such as surface roughness and electrostatic attraction,

34

Element	Atomic Mass	Half-life	Decay Mode	Energy [MeV]	Reference
Ra	226	1600 y	alpha	4.784	[68]
Rn	222	3.8235 d	alpha	5.489	[69]
Po	218	3.098 m	alpha	6.002	[70]
Pb	214	26.8 m	beta	1.019 (endpoint)	[70]
Bi	214	19.9 m	beta	3.27 (endpoint)	[70]
Po	214	163.6 μs	alpha	7.686	[71]
Pb	210	22.2 y	beta	0.063 (endpoint)	[71]
Bi	210	5.01 d	beta	1.162 (endpoint)	[71]
Po	210	138.376 d	alpha	5.304	[72]
Pb	206	stable	N/A	N/A	

Table 2.1: Details of the ^{222}Rn decay chain beginning with its parent isotope, ^{226}Ra, and ending with the stable isotope, ^{206}Pb. Notice the long half-life of ^{210}Pb that sustains this decay chain.

surface contamination cannot simply be wiped off; it needs to be acid-etched or machined away [73]. Given the difficulty of removing it, the best way to minimize backgrounds from radon is to understand and reduce the amount of initial exposure.

2.1. Radon Plate-out

The first step to limiting radon plate-out is to minimize the exposure to ambient radon in the room. That is why most of the equipment that is vital for direct detection experiments is fabricated and assembled in cleanrooms; even dust contains radioactive isotopes. Specialty systems can be built that scrub radon from the air that feeds into these cleanrooms to create low-radon environments [74, 75]. Instrumented experimental components can be stored in purge boxes, also called purge cabinets, where boil-off nitrogen gas is flowed in to maintain a slight overpressure and prevent regular air from entering the box. The nitrogen flowing in is free of radon because the boiling point of nitrogen is 77 K and the melting point of radon is 202 K, so all the radon is expected to freeze solid in the liquid nitrogen.

2.1.1. Dead Air Model

The most simple model for describing the plate-out of radon has been called the "Dead Air Model" [1]. The basic assumption of this model is that all of the radon in a contained space will eventually plate out onto every surface within that space. According to the Dead Air Model, the activity of a sample is given by

$$A_{sample} = C_{Rn} \times V_{space} \times f_{area} \,, \tag{2.1}$$

where V_{space} is the volume of the space of interest whether it is a laboratory or a purge box, C_{Rn} is the concentration of radon in that space, and f_{area} is the fractional surface area of the sample with respect to the total surface area of the space. Another way of writing the Dead Air Model is to introduce an effective plate-out height h_{eff} and surface area of the sample S_{sample}:

$$A_{sample} = C_{Rn} \times S_{sample} \times h_{eff} \,. \tag{2.2}$$

In the context of the Dead Air Model, h_{eff} can be considered as the height of a column of air over the sample that is needed to achieve the level of activity A_{sample} if all of the radioactive isotopes in that air plated onto the sample. This imaginary column of air is similar to how the pressure at sea level is the weight of the atmosphere in that column divided by the cross-sectional area of the column. Of course, this is a simplified picture that gets more complicated with fluid dynamics, such as laminar flow and turbulence. However, effective plate-out height is a good way of reporting plate-out results because it is independent of the sample, concentration of radon, or choice of model.

---

[1] Richard Schnee claims he invented this name

2.1.2. Jacobi Model

While it makes for a quick first-order approximation, there are many environmental factors that are not accounted for in the Dead Air Model. A more complex and accurate model is the parametric "Jacobi Model" that is used to account for the number of radon and progeny atoms in a room. The Jacobi Model was originally devised in 1972 from a medical physics perspective. An example of the original intention of the model is to determine the amount of activity inhabitants are exposed to in a house, but it can also be used to calculate an effective plate-out height [76]. This model only deals with probabilities that are averaged for all atoms in the room; it does not track individual atoms and their interactions [77]. The model is also founded on the assumption that the original concentration is uniformly distributed throughout the room [78]. This is a safe assumption for radon because it diffuses much quicker than its half-life of 3.82 days.

While it is true the Jacobi Model is much more complex than the Dead Air Model, it is important to note that this also is not a perfect model. There are still a large number of factors not included in the model such as geometry of the space and objects within the volume, turbulence, aerosol concentration and size, air viscosity, pressure, temperature, humidity, and electrostatics. This is by no means an exhaustive list, and there are unaccounted correlations between these parameters, making it difficult to estimate how these might effect the model and the parameters that are included.

The following terms have very specific meanings in the Jacobi Model.

aerosol: a colloidal suspension of particles dispersed in air or another gas

attached: progeny that have adhered to an aerosol particle

unattached: progeny that are not attached to an aerosol particle (can be free floating or clustered together)

deposition: attachment of particles to surfaces

detachment: the ability of progeny to separate themselves from a surface or aerosol particle (transition from attached or deposited to unattached)

The first input is the decay constant for each isotope λ_i, where the specifier i will identify the progeny number after radon ($i = 0$). The decay constant is defined as

$$\lambda_i = \frac{\ln(2)}{T_{\frac{1}{2},i}}, \tag{2.3}$$

where $T_{\frac{1}{2},i}$ is the known half-life. Another often-known input is the ventilation rate λ_v of the space. This is typically controlled mechanically via air handlers or the volume flow rate, in the case of a purge cabinet.

The first input that is not well-defined is the rate of attachment to aerosol particles λ_a. The attachment rate is defined as the probability of progeny interacting with an aerosol particle times the probability that they then adhere to the particle. The probability that they meet is dependent on the aerosol concentration and the size of the particles [77]. In this description, the attachment rates are assumed to be the same for all progeny, although the model could be expanded by substituting $\lambda_{a,i}$ for λ_a.

A natural progression from the attachment is to introduce the two types of detachment: from an aerosol particle and from a surface after deposition. Within the Jacobi Model there is a parameter that allows for an atom to detach from an aerosol particle when the attached parent undergoes an alpha decay that has enough energy to detach the subsequent daughter. Starting with the assumption that all ^{222}Rn atoms are unattached means there are two alpha decays that can lead to detachment: ^{218}Po decaying to ^{214}Pb and ^{214}Po decaying to ^{210}Pb. The probability of detachment r commonly used is 80%, but this can depend on the aerosol diameter, shape, and chemical structure as well as the energy of the alpha decay [77]. This

picture complicates further when considering that aerosol particles can clump together and trap the parent isotope inside a cluster. Somewhat surprisingly, the Jacobi Model does not consider atoms deposited on a surface to be able to detach. The reason for this is that the progeny have a short recoil range in air of $\sim 80~\mu$m, thus they induce an opposite charge and are re-attracted to the surface they just detached from [77]. Based on the direction of the recoil, it is also possible for an atom to be driven deeper into the surface, which would further reduce its chance to detach in subsequent decays. This process is known as "implantation" and will be revisited at the end of this section.

The deposition rates will be discussed later, but putting everything together gives the following system of equations [79]:

$$
\begin{aligned}
\lambda_0 C_0 &= (\lambda_1 + \lambda_v + \lambda_a + \lambda_d^u)C_1^u \,, \\
\lambda_a C_1^u &= (\lambda_1 + \lambda_v + \lambda_d^a)C_1^a \,, \\
\lambda_1 C_1^u + r\lambda_1 C_1^a &= (\lambda_2 + \lambda_v + \lambda_a + \lambda_d^u)C_2^u \,, \\
\lambda_a C_2^u + (1-r)\lambda_1 C_1^a &= (\lambda_2 + \lambda_v + \lambda_d^a)C_2^a \,, \\
\lambda_2 C_2^u &= (\lambda_3 + \lambda_v + \lambda_a + \lambda_d^u)C_3^u \,, \\
\lambda_2 C_2^a + \lambda_a C_3^u &= (\lambda_3 + \lambda_v + \lambda_d^a)C_3^a \,.
\end{aligned}
\tag{2.4}
$$

The best way to interpret these equations is in terms of conservation of number of atoms where each equation is a different step in the model. The left side of each equation specifies the concentration of a population of atoms coming into that step, and the right side of each equation defines the processes those atoms can undergo. Figure 2.1 is a visual representation of the model from the decay of ^{222}Rn to the decay of ^{210}Pb. In fact, Eq. 2.4 only goes as far as the third isotope in this decay chain, which is ^{214}Bi. They do not include the short-lived daughter, ^{214}Po, because it is assumed to decay before it has time to attach, detach, or deposit itself. However, it has been included in the graphic for the sake of completeness.

The model does not need to be extended beyond ^{210}Pb because its half-life is so long that it will eventually deposit on a surface before it has a chance to decay.
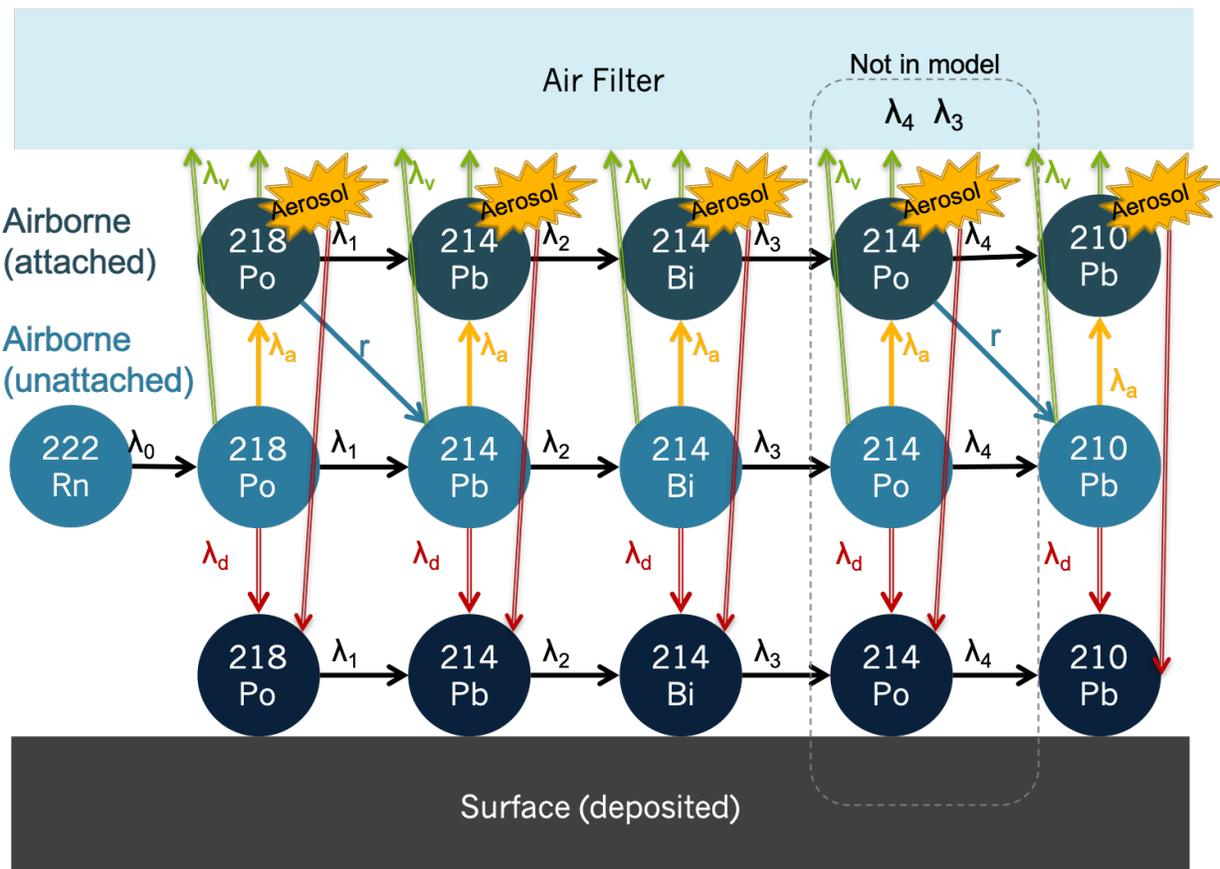
Figure 2.1: This is a visual depiction of the parametric Jacobi Model equations. Each row of atoms corresponds to the deposited, unattached, or attached populations. The red arrows indicate the atom depositing on a surface, and the green arrows indicate the ventilation rate to the air filter at the top. There are two blue diagonal arrows that represent the ability of an attached atom to detach via an alpha decay, although the ^{214}Po is not included in the equations. The last two columns are extensions to the long-lived ^{210}Pb.

The deposition rate for unattached atoms is fairly well-understood. It is defined as

$$\lambda_d^u = vS/V \, , \tag{2.5}$$

where S is the total surface area in the space including any objects, V is the volume of the space, and v is the diffusion velocity of the unattached progeny [78]. There are many

mechanisms that can drive this velocity, a few of which are: Brownian diffusion, turbulent diffusion, gravitational settling, and turbophoresis, which is the tendency for the atoms to migrate in the direction of decreasing turbulence [77].

The formalism described so far has been "textbook" with the particular isotopes of interest substituted. However, the model can be modified to be more applicable for dark matter detection. In this case, only plate-out in laboratory environments with excellent ventilation and air filtration are considered, such as cleanrooms or purge cabinets. Therefore, it can be assumed that all aerosol particles and attached atoms will be filtered away [78]. This effectively sets $\lambda_a = 0$ and $\lambda_d^a = 0$, which greatly simplifies Eq. 2.4 down to:

$$
\begin{aligned}
\lambda_0 C_0 &= (\lambda_1 + \lambda_v + \lambda_d^u)C_1^u \,, \\
\lambda_1 C_1^u &= (\lambda_2 + \lambda_v + \lambda_d^u)C_2^u \,, \\
\lambda_2 C_2^u &= (\lambda_3 + \lambda_v + \lambda_d^u)C_3^u \,.
\end{aligned}
\tag{2.6}
$$

A graphic depiction of these equations is shown in Fig. 2.2.

Having simplified the model, there are a few key outputs that can be calculated. One obvious output is the concentration of radon and each of the progeny at any given time. However, as stated earlier, the effective plate-out height is an excellent way of reporting normalized results that are directly comparable between experiments or models. Equation 2.7 shows how to calculate the effective plate-out height using parameters in the Jacobi Model [78],

$$
h_{\text{eff}} = \frac{\lambda_d}{\lambda_d + \lambda_v} \frac{V_{space}}{S_{space}} \,.
\tag{2.7}
$$

Predictably, the Jacobi Model formula for h_{eff} reduces to the Dead Air Model in the case where there is no ventilation ($\lambda_v = 0$).

Another output from the Jacobi Model is that the total deposition can be found by summing the deposited populations for each progeny. However, the model assumes that all
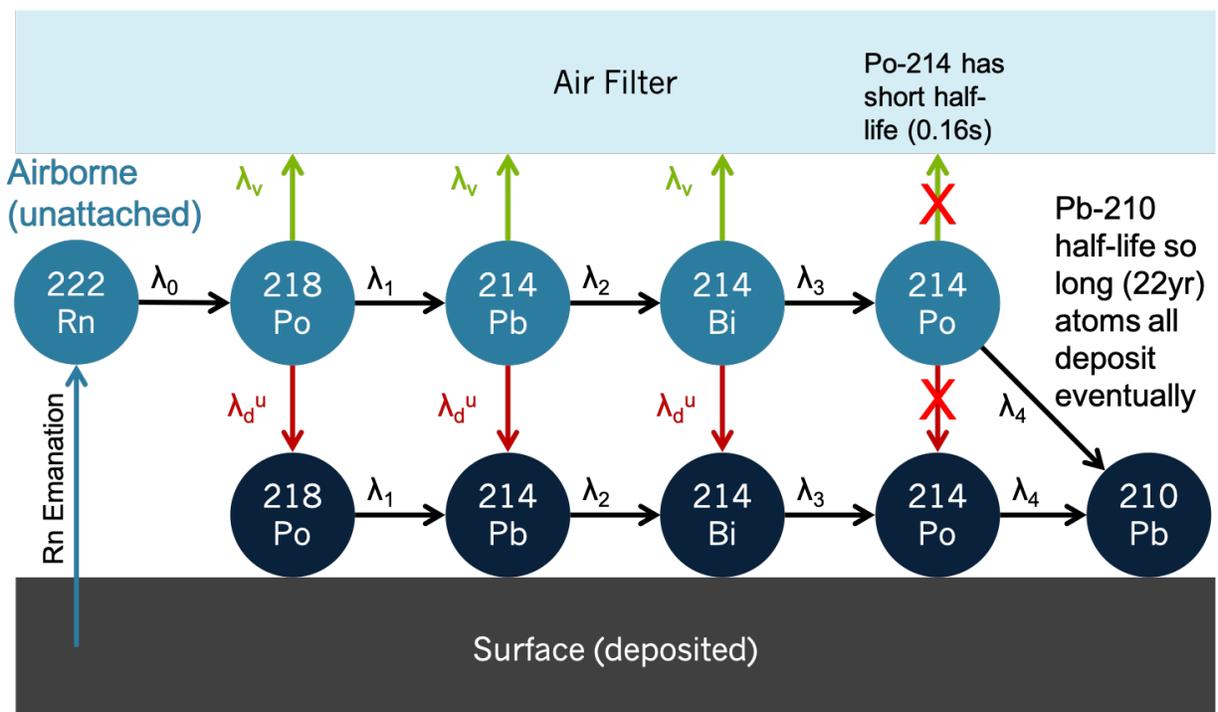
Figure 2.2: Visual depiction of the simplified Jacobi Model equations under the assumption that all aerosol particles and attached atoms have been filtered away. This is a much simpler picture than Fig. 2.1. An additional feature that has been included is the arrow labeled "Rn Emanation" that points to the initial population of ^{222}Rn from the surface. Radon emanation is described in great detail in Appendix B, but this arrow emphasizes that the Jacobi Model is not sensitive to where the original radon population originates. However, in a realistic picture, radon could be introduced from anything in the space.

of the progeny are sitting on the surface, which is not fully accurate. Depending on the direction of the decay, an atom deposited on a surface can recoil further into the material resulting in a daughter that is implanted. An example of how the Jacobi Model could be expanded to include implantation is illustrated in Fig. 2.3.
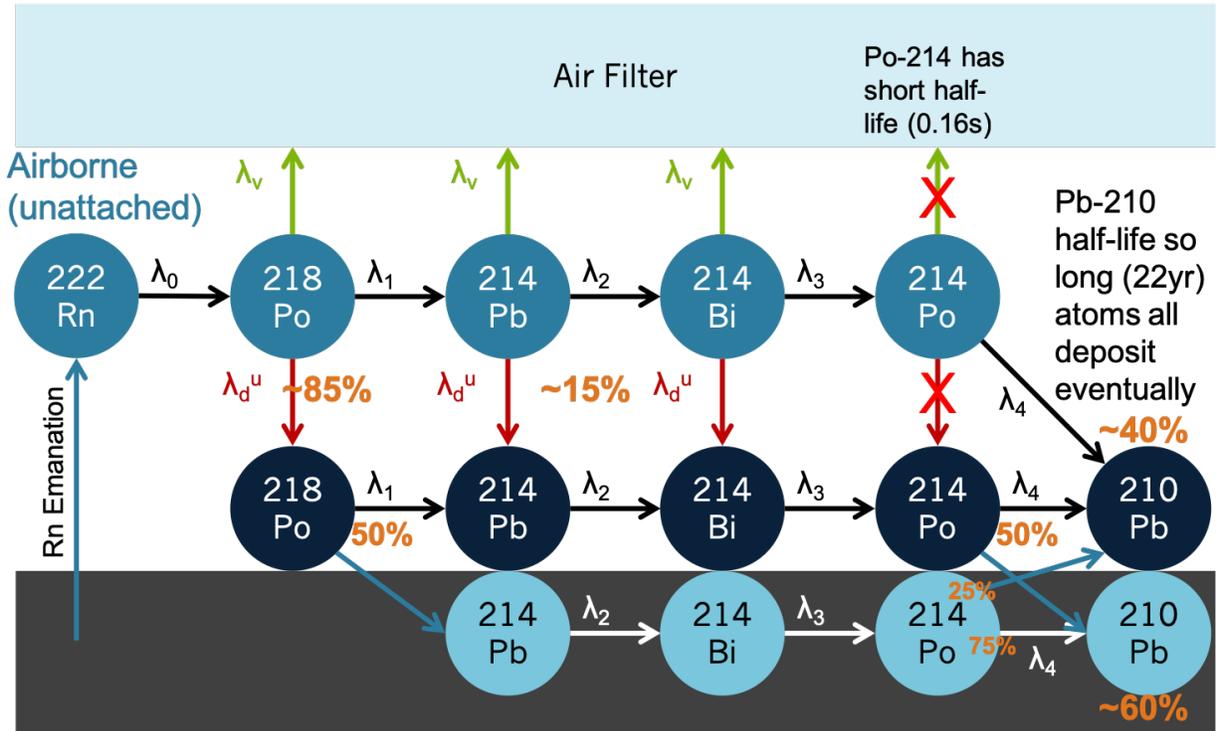
Figure 2.3: Visual depiction of the simplified Jacobi Model with implantation for alpha decays that can implant atoms below the surface of the material.

There are three estimates that go into this implantation example. The first is that 85% of the deposited atoms are ^{218}Po, and the other 15% of atoms deposit some time later. This is because the half-life of ^{222}Rn is so much longer than its progeny (until ^{210}Pb) that the atoms have a much higher probability of depositing on a surface. Another estimate is that there is a 50% chance that the alpha decay points down. In reality, the decay is actually isotropic, but this assumption makes the calculation simpler. Related to that angular assumption, it is estimated that farther down the decay chain, ^{214}Po has a 75% chance of remaining implanted. The ^{214}Po alpha decay has more energy than the ^{218}Po decay, but there is a small range of angles back towards the surface for the implanted atom to be ejected. These

implantation and ejection probabilities can be added to the Jacobi Model similarly to the detachment probability parameter r. Stepping through the estimated probabilities in this example shows how the Jacobi Model can be used to predict the ratio of implanted atoms versus those deposited on the surface [78].

2.1.3. Plate-out Tests

In order to systematically test the Jacobi Model, a controlled radon environment was created inside an exposure chamber at PNNL. A purge box, shown in Fig. 2.4, was modified to have extra ports so that the flow in and out could be controlled, and a pressure gauge was added to monitor the pressure inside the box. It was important to maintain overpressure inside the box so that the laboratory air did not leak in and provide an unknown supply of radon. The known supply of radon came from a calibrated flow-through radon source to control exactly how much radon was allowed into the purge box. Unfortunately, laboratory regulations limited the maximum activity allowed for an unsealed source to < 100 Bq, which ultimately hindered the ability to produce consistent measurements. Since the decays inside the source are constant, the activity obtained is inversely proportional to the flow through the source. Therefore, in order to maximize the radon exposure, the lowest possible flow rate was used, which was approximately 0.1 lpm. However, too little flow lead to problems maintaining overpressure in the purge box. It also took a long time for the radon to reach equilibrium in the modest volume of 0.151 m^3 due to the low flow rate. All of this led to huge variations in measured plate-out activities to the point it was impossible to draw any conclusions.

Since university regulations are not as strict as the national laboratories, this experiment could be re-created at SMU with a 35 kBq source. The stronger source allows a higher flow rate of 4 lpm, which provides more consistent results. This controlled and reproducible radon environment permits systematic tests of some of the Jacobi Model inputs, as well
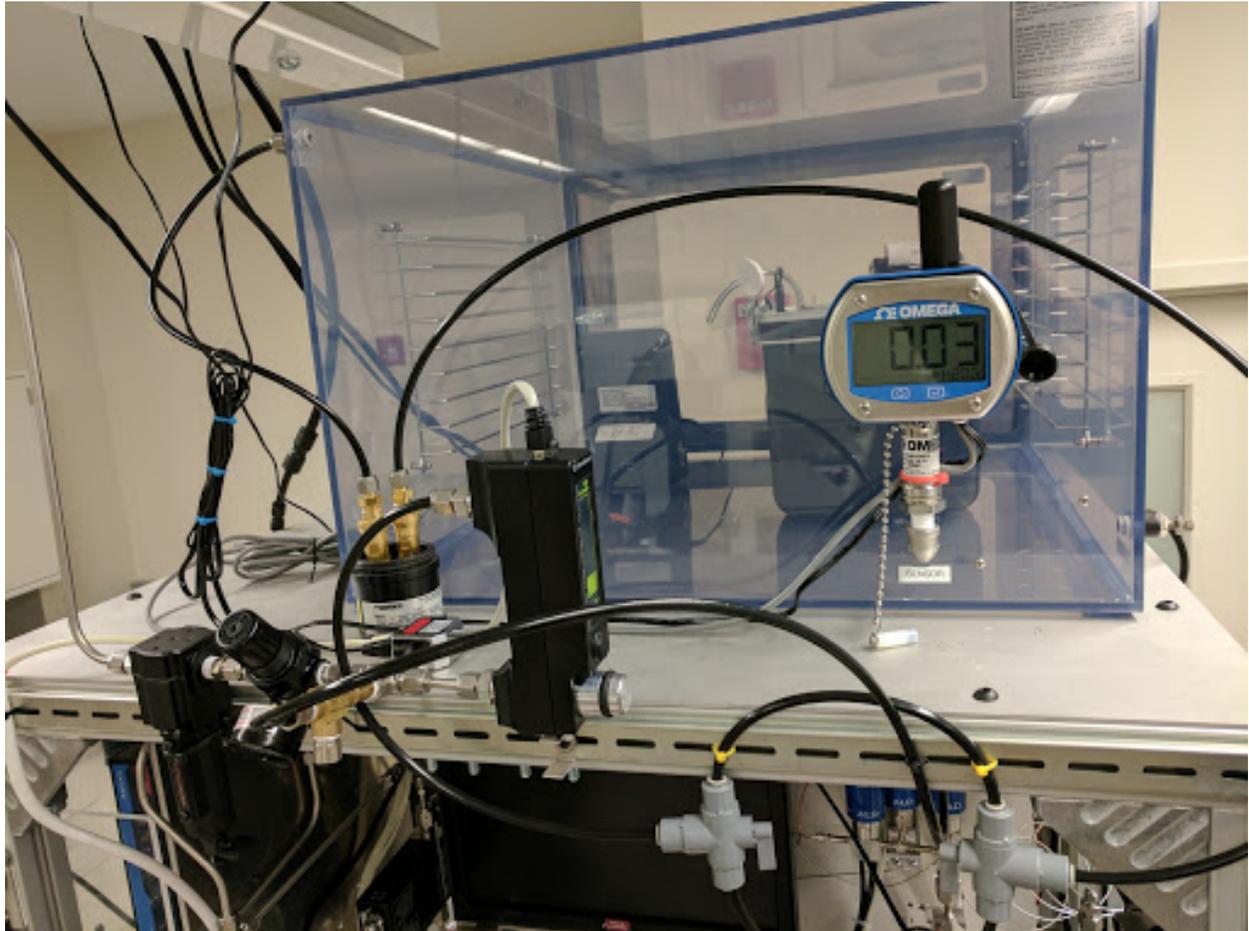
Figure 2.4: Photo of the modified purge box at PNNL that was used as a radon exposure chamber. The black digital meter in the bottom center controlled the flow rate of boil-off nitrogen into the box. The black canister with the brass ports (behind the flowmeter) is the Pylon radon flow-through source. The digital gauge is reading the differential pressure between the purge box and the room. Inside the box are a Dylos particle counter (left) and a Durridge Rad7 radon monitor (right, behind pressure gauge).

as factors that could influence those inputs. Some of the things being tested at SMU are: flow rate, static charge, and different sample materials. A few factors that are not currently being tested, but could be in the future, are: temperature, humidity, surface roughness, and fractional surface area.

In order to systematically determine the effect that different environmental factors have on the plate-out of radon and its progeny, a radon exposure system was built at SMU. Depicted in Fig. 2.5 is the system with a custom acrylic exposure chamber in the shape of a toroid for circulating radon-doped gas during exposures. A bevy of tests have already been performed with this setup. For example, it was recently concluded that circulating the air in the exposure chamber increases the number of atoms plated out on a copper sample, yet the plate-out rate does not depend on the airspeed. The next plan is to switch the conductive copper plate for insulating polyethylene or acrylic samples. Plastics are typically less predictable because they can build up large amounts of static charge that could attract or repel charge radon daughters. An electrostatic detector was purchased in an attempt to measure this charge and determine how charge influences the plate-out rate [80].

This research will provide an understanding of different factors that influence plate-out, ideally leading to practical solutions that can be implemented towards reducing the amount of radon surface contamination and improving sensitivity to dark matter. This is future work beyond the scope of this thesis.

2.2. Material Assay

In order to keep the radon background low, direct detection experiments need to be constructed from radiopure materials, which requires ultra-sensitive measurements to choose the purest materials. There are many techniques for assaying a sample, but they largely fall into two categories: surface screening and bulk assays. Surface screening is conducted by measuring the daughter particles from the decay of surface contamination of a sample placed
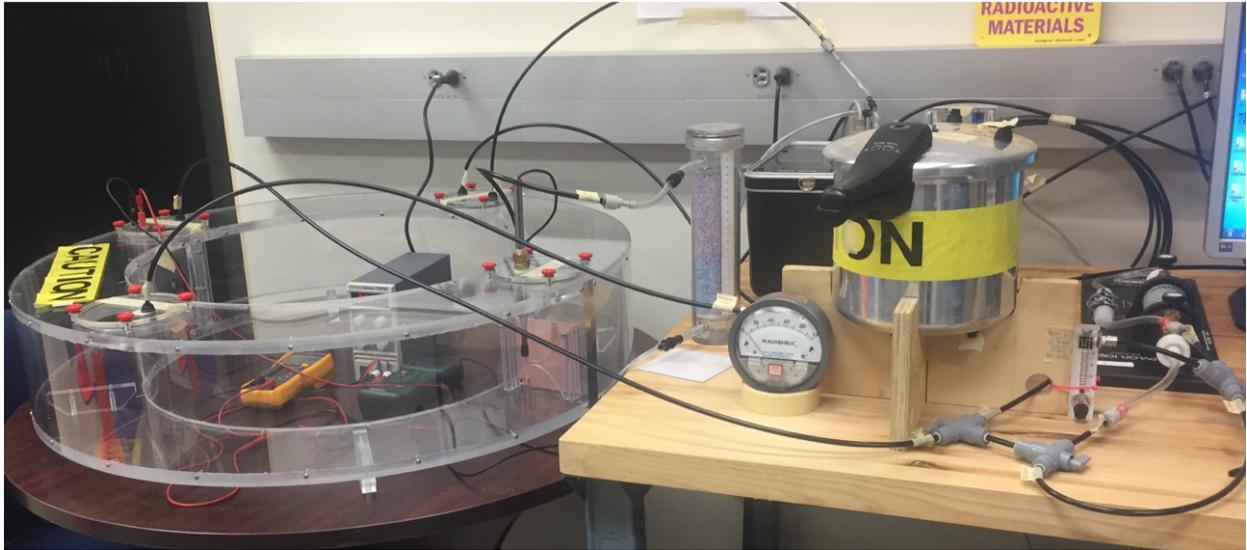
Figure 2.5: Photo of the custom toroidal radon exposure chamber at SMU. The black device to the far right is the Pylon radon flow-through source, with an analog flowmeter in front of it to control the volume flow rate through the source. The white gauge is showing the differential pressure in the chamber is slightly above zero. The clear acrylic chamber is used for radon exposures, and the square copper sample can be see on the right side of the chamber. There is a repurposed computer fan in the left side of the chamber to circulate airflow, and an anemometer is directly in front of the sample to measure that speed. The black box in the back is a Durridge Rad7 radon monitor. This entire setup is designed to be easily switched over to the modified pressure cooker in the middle of the picture for exposures inside a conducting chamber.

in a detector. This is easiest with high-energy alpha decays because they have known energies and appear as mono-energetic peaks above backgrounds, although surface screening can also be done with beta or gamma decays. Measuring the radioactivity from the bulk of a material can be done in multiple ways. Sometimes it consists of sampling a piece of the bulk and measuring it, while sometimes it consists of waiting for the bulk activity to work its way out of the sample to be measured.

2.2.1. Surface Screening

Materials used in the construction of an experiment can adsorb radioactive contamination at any phase of the fabrication, shipping, handling, or assembly. That is why screening is important for testing manufacturing and cleaning procedures. As noted above, alpha decays are clear indicators of surface contamination, and the most sensitive alpha counter that is commercially available is the XIA UltraLo-1800. One of these instruments is located in the Laboratory for Ultrapure Materials and Isotope and Neutron Assessment (LUMINA) at SMU. A full description of the XIA alpha counter and an attempt to improve it using pulse shape discrimination can be found in Appendix A.

2.2.2. Bulk Assay

With improving technology and increased sensitivity, direct detection experiments will require assays of ultra-low radon levels, sometimes needing to detect as few as tens of atoms in equilibrium. Radon emanation is one of the most sensitive and robust ways of making these measurements. While it is possible to infer radon activity from a measurement of bulk uranium content as is done with Inductively-Coupled Plasma Mass Spectrometers (ICP-MS), there are factors that can complicate this prediction [81]. While not an exhaustive list, some of the factors are: uranium not being uniformly distributed in the bulk, unknown diffusion constants, and important physics at material boundaries. It is much more straightforward to

measure the radon itself, but cutting-edge experiments require lower bulk radon levels than commercial devices are capable of measuring. Therefore, the need arises to build a custom system to measure radon emanating from materials. Appendix B details the commissioning of a new radon emanation system built at PNNL that achieves high sensitivity through use of custom high-efficiency ultra-low-background proportional counters [82]. Initial tests demonstrated that it is as sensitive as the best radon emanation systems across the country, and there is a clear path for improvement with a less contaminated detector and running underground to avoid cosmic ray backgrounds.

CRYOGENIC DARK MATTER SEARCH

The Cryogenic Dark Matter Search (CDMS) is an experimental program that is searching for the direct interaction of dark matter with crystal detectors operated at cryogenic temperatures.

The first stage (CDMS I) was located at Stanford Underground Facility (SUF) in Palo Alto, CA. The experiment was operated in an underground tunnel that was an extension of a beam dump in End Station III on the Stanford campus. The reason for being underground is so that the layers of rock in the Earth can shield incoming cosmic rays. The tunnel at SUF is 10.6 m underground, which equates to 17 mwe of overburden. Though not frequently used outside the realm of underground science, "mwe" stands for "meters of water equivalent". Based on the depth of the laboratory and the composition of the Earth and rock above it, mwe is useful for directly comparing the amount of overburden for underground laboratories by converting to a depth of water that would provide an equivalent amount of shielding. Operating at SUF with 17 mwe overburden is expected to reduce the cosmic muon flux by a factor of 5. CDMS I was operated at SUF until 2002 when detectors designed for CDMS II were tested in the existing infrastructure [83, 84]. In 2003, the CDMS II experiment was built and began operation in Soudan, MN [85, 86].

3.1. SuperCDMS Soudan

The SuperCDMS experiment was planned as an upgrade to CDMS II in 2011 in order to deploy a new type of detector inside the existing shielding [87]. SuperCDMS Soudan is the most recent generation of CDMS experiments and operated until it was decommissioned in 2015.

3.1.1. Soudan Shielding

Located 2341 feet below the surface of Soudan, MN is an old iron mine that ceased operation in 1962. In the 1980s, level no. 27 of that mine was converted into laboratory space called the Soudan Underground Laboratory (SUL). The overburden at SUL is 2090 mwe, which reduces the cosmic muon flux by a factor of 5×10^4. Because of the low flux, the laboratory has hosted many low-background experiments, including the Main Injector Neutrino Oscillation Search (MINOS) [88] and SuperCDMS [86].

Within SUL, the SuperCDMS experiment was wrapped in many layers of shielding, shown in Fig. 3.1, to protect the detectors from external sources of background and noise. This included the room that housed the whole experiment and shielded radio frequency electromagnetic radiation. The outermost layer of shielding was an active muon veto constructed from scintillating panels to reject coincident events that made it though the overburden [89]. Within the muon veto was a thick layer of high-density polyethylene (HDPE) to attenuate incoming neutrons from the cavern that can mimic a dark matter signal. Inside the HDPE were two layers of lead shielding to block beta radiation ("betas") and penetrating gamma radiation. Lead has a radioisotope, ^{210}Pb, with a 22.2 year half-life that emits its own radiation as it decays, so there was a thin layer of "ancient lead" that has been shielded on the ocean floor for several half-lives; it is used to block decays from the exterior young lead shielding. The final layer was a thin layer of HDPE to attenuate any neutrons created by the

ancient lead decays. There was also a mu-metal shield to block magnetic flux and a radon purge barrier flushed with boil-off nitrogen to prevent high-radon laboratory air from leaking into the cryostat. At the center of all the shielding was a cryostat composed of nested copper cylinders that housed the actual detectors [90].
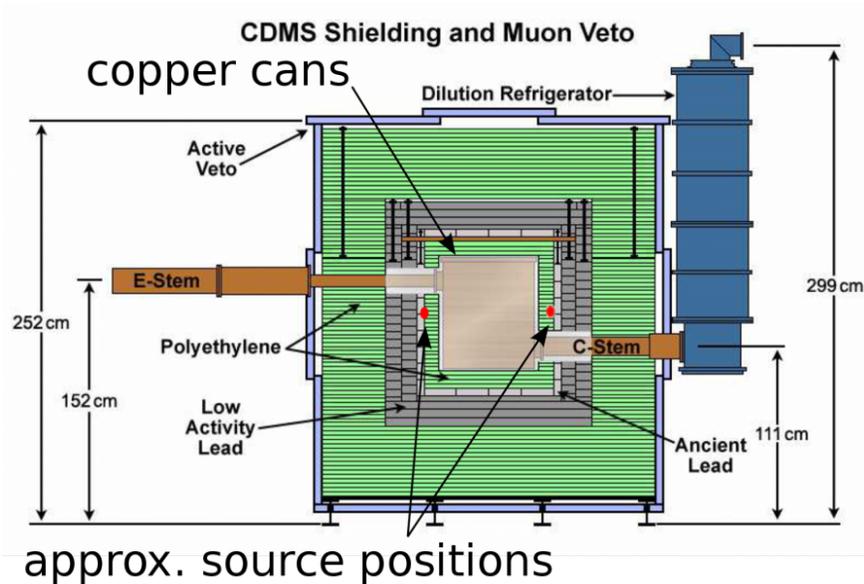
Figure 3.1: Side view of the shielding layers around the Soudan experiment. [90]. Green layers are HDPE, dark gray layers are lead, and the light gray layer is ancient lead.

3.1.2. iZIP Detectors

The 15 new detectors deployed in SuperCDMS Soudan were called interleaved Z-sensitive Ionization and Phonon detectors (iZIPs). The iZIPs were cylindrical germanium crystals that were 1 inch tall, 3 inches in diameter (exactly the size of a regulation hockey puck), and weighed 0.6 kg each. The detectors were arranged in copper towers that held three detectors each, and there were five towers installed in the cryostat. The detector and the installed configuration are shown in Fig. 3.2.

There are two interchangeable nomenclatures used for identifying the detectors. The first is the tower address (e.g., T2Z1), composed of the tower number (1-5) and iZIP within that tower (1-3, from top to bottom). A simpler, but less descriptive, system is to assign each
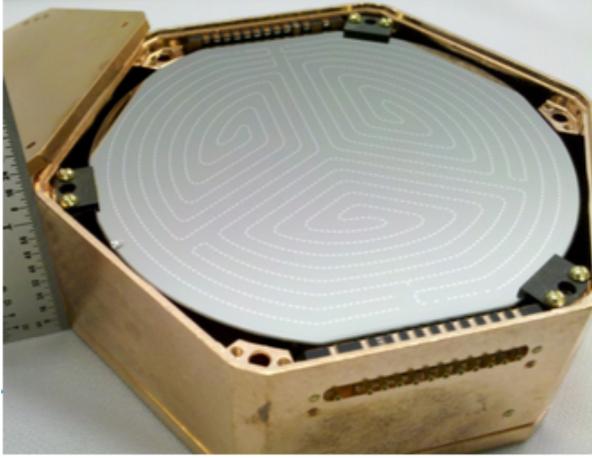
(a) Picture of an iZIP detector in its copper housing. The phonon sensor pattern and partitioning can be seen on the face of the crystal. The charge electrodes are not resolved in this picture, but they are interleaved between the phonon sensors.
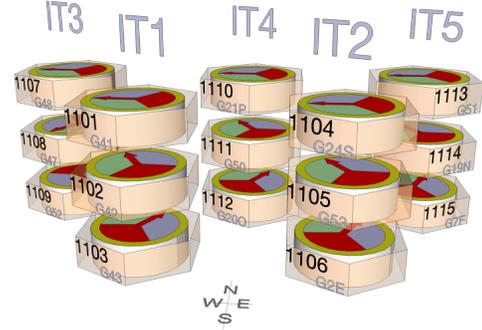
(b) Visualization showing the arrangement and orientation of all 15 detectors.

Figure 3.2: Details of the iZIP detectors installed in SuperCDMS Soudan.

detector a number 1-15 (e.g.iZIP 4). Table 3.1 shows the mapping between these frequently used systems.

Det:	iZIP 1	iZIP 2	iZIP 3	iZIP 4	iZIP 5	iZIP 6	iZIP 7
Add:	T1Z1	T1Z2	T1Z3	T2Z1	T2Z2	T2Z3	T3Z1

iZIP 8	iZIP 9	iZIP 10	iZIP 11	iZIP 12	iZIP 13	iZIP 14	iZIP 15
T3Z2	T3Z3	T4Z1	T4Z2	T4Z3	T5Z1	T5Z2	T5Z3

Table 3.1: Naming conventions for the detectors installed in SuperCDMS Soudan.

3.1.2.1. Energy Collection

The iZIP detectors were designed to simultaneously collect ionization (charge) and phonon energy. The ionization was collected by electrodes photolithographically fabricated on the flat faces of the crystals. The ionization readout was segmented into a circular inner charge channel (denoted "qi") and an outer charge ring (denoted "qo") for radial fiducialization.

In order to detect vibrations in the crystal lattice, the iZIP detectors needed to be cooled to cryogenic temperatures to reduce thermal noise. As the phonons propagate outward from the recoil event, they eventually deposit their energy in aluminum fins that are fabricated on the face of the crystal. This energy breaks Cooper pairs in the superconducting aluminum into quasiparticles that are collected by a tungsten Transition-Edge Sensor (TES), causing the TES to heat up. The change in temperature causes a change in resistance across the TES that is detected by a Superconducting QUantum Interference Device (SQUID). This whole sensor is called a Quasiparticle-assisted Electrothermal feedback Transition Edge Sensor (QET) [87].

TESs are made from materials that have a critical temperature T_c below which the material is superconducting. Above T_c, the material is non-superconducting, or "normal," and the TES is said to be "in transition" when near T_c. Figure 3.3 shows that the change in resistance is quite large for a small change in temperature near the transition point. SuperCDMS detectors use this as a way to amplify small deposits of energy by operating in this transition regime.

The phonon sensors are segmented into 4 channels of equal area: one outer ring (pA) and three inner wedges (pB, pC, and pD). This layout is repeated for both the top face (S1) and bottom face (S2) of the cylindrical detector, with a relative rotation of 180 degrees between S1 and S2. Figure 3.4 shows the alignment of all channels on the detector. This partitioning provides rough position information for each event.

Simultaneously reading out the ionization and phonon signals allows for background rejection and approximate particle identification in the iZIPs. For a given event energy, electron recoils will produce approximately three times as many phonons as nuclear recoils. Therefore, the ratio of phonon energy to charge energy, defined as yield, can be used to discriminate between signal and background events.
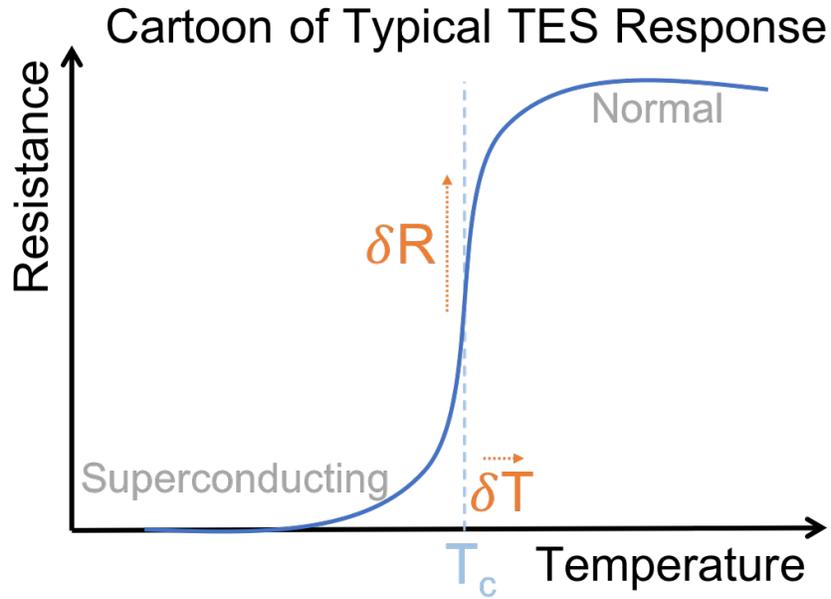
Figure 3.3: This is a cartoon of a typical TES resistance curve. Around T_c the TES is said to be in transition between normal and superconducting. In this regime, any small increase in temperature results in a large change in resistance that allows for detecting small energy deposits.
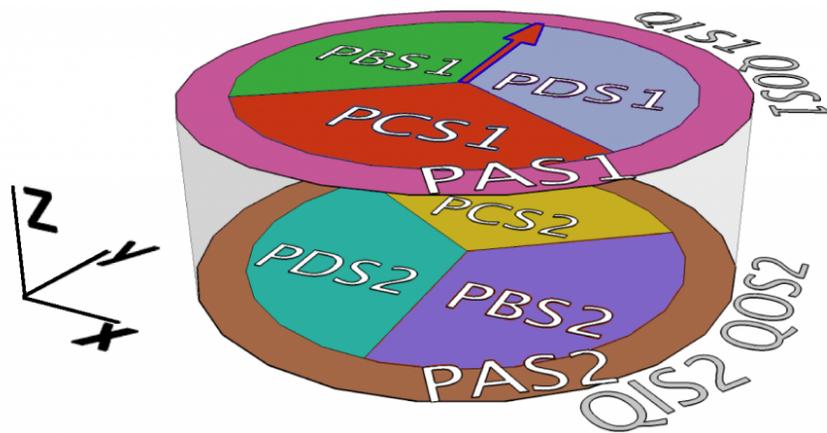
Figure 3.4: Visualization of the 12 channels on each iZIP detector.

The charge electrodes are biased to opposite voltage on the top and bottom face, nominally ± 2 V, and the phonon sensors are grounded. The clever part of the design of the iZIPs is that the charge electrodes and phonon TESs are interleaved to create an electric field that is uniform in the bulk of the detector, but scalloped near the flat faces of each crystal, as shown in Fig. 3.5.
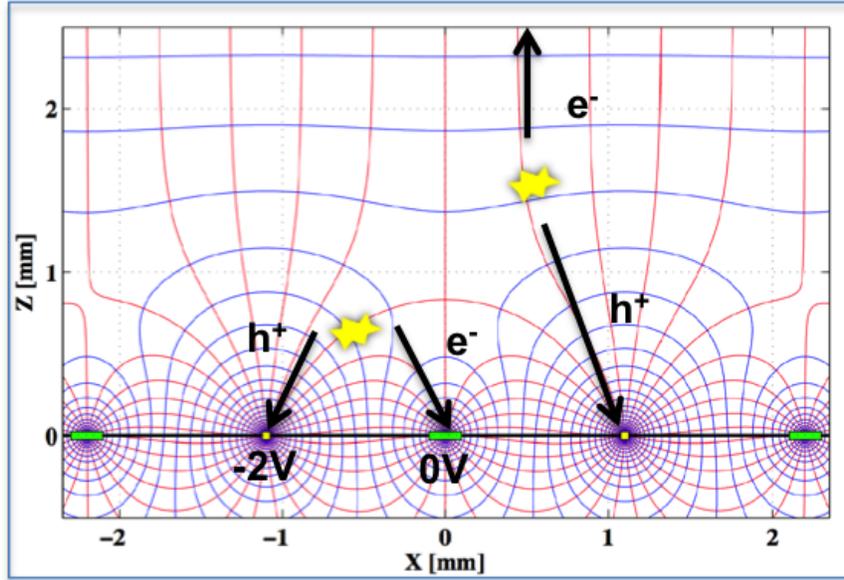
Figure 3.5: The scalloped electric field inside the iZIP detectors. An event that occurs in the bulk of the crystal will experience a uniform electric field, and the electrons and holes will drift to opposite faces. An event that occurs in the scalloped field near one of the faces will only be detected on that face.

Events that occur in the bulk of the crystal will experience the uniform electric field, and the electrons and holes will drift to opposite faces of the detector. However, the scalloped shape of the electric field leads to asymmetrical energy collection for events near the surfaces, which will be dominated by non-penetrating backgrounds where energy is collected entirely on a single face. The asymmetry allows for a clearly defined fiducial volume in the z-direction by selecting events with a signal on both faces, as shown in Fig. 3.6. Since WIMPs have a very small interaction cross-section, they are expected to interact uniformly throughout the

detector. Focusing dark matter searches on the bulk of the crystal maximizes the signal-to-background ratio [91].
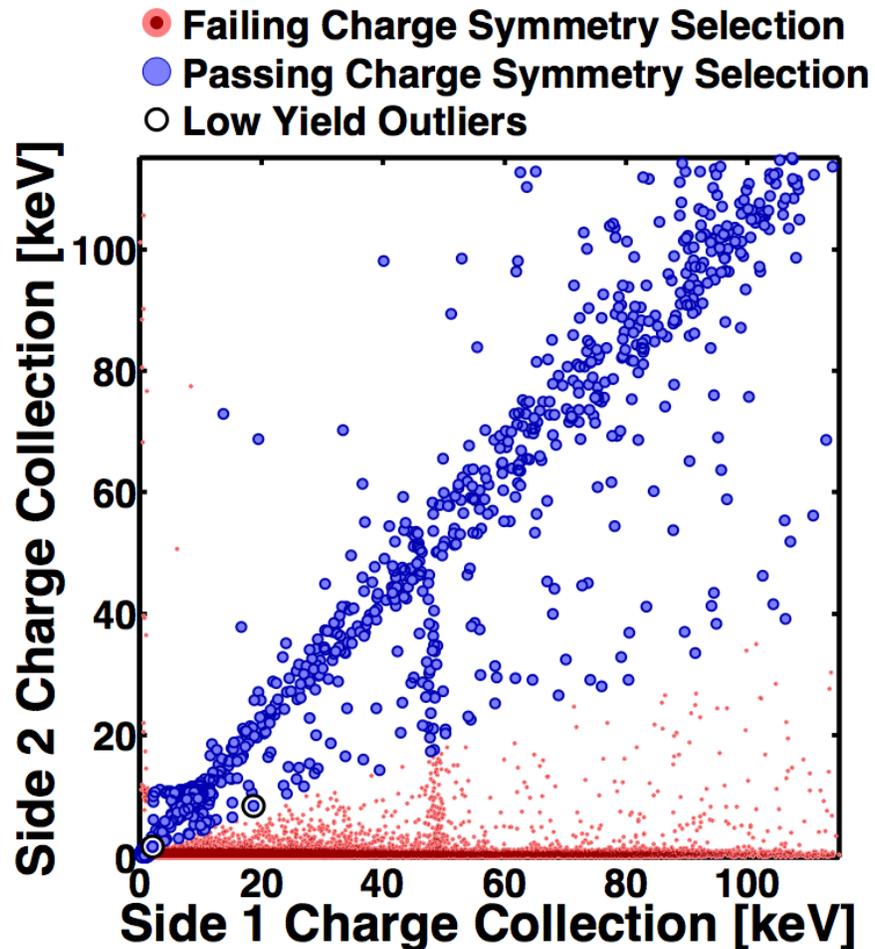
Figure 3.6: Demonstration of the capability of the iZIP to reject surface events [91]. The data points are from a ^{210}Pb calibration source that emits beta and gamma particles. The gammas penetrate to the bulk and are selected as blue points along the diagonal, where betas are surface events and compose the red points that fail the symmetry cut.

3.1.3. Data and Features

There were three types of data collected with SuperCDMS: barium calibration data, californium calibration data, and "low-bg" data. There are a number of energy features within these datasets that are useful for calibrating the energy scale, calculating cut efficiency, and measuring energy resolution.

3.1.3.1. Barium Data

Barium calibration data was collected while a ^{133}Ba calibration source was deployed inside the layers of lead shielding; the source locations are indicated by the red dots in Fig. 3.1. This source is a high-rate gamma-emitter that was used to understand the detector response to electron recoils. Data were collected at regular intervals to make sure the detector response was not changing throughout the run.

The ^{133}Ba spectrum is dominated by a distinct peak at 356 keV, seen as the high-energy spike in the uncalibrated blue spectrum in Fig. 4.1. Additionally, there are much smaller peaks at 276 keV, 303 keV, 383 keV, as well as a dull "hump" from backscattering that rises above a continuum of events that have lost energy via Compton scattering. High-energy gammas from the ^{133}Ba source are useful for energy calibration because they have enough energy to penetrate the inner layers of shielding and deposit energy throughout the crystals. They can also penetrate the crystal, so events are uniformly distributed throughout the bulk.

3.1.3.2. Californium Data

Californium calibration data was periodically collected while a ^{252}Cf neutron calibration source was deployed inside the cryostat. This neutron source served as a proxy for dark matter nuclear recoils.

When exposed to ^{252}Cf source, the stable ^{70}Ge in the crystal atoms can capture neutrons and form heavier germanium isotopes. These heavier isotopes are unstable and decay through various channels. For example, ^{73}Ge decays via a prompt 66.7 keV gamma ray [92]. Due to the prompt emission, this peak can be found in the californium calibration data and was used to measure the iZIP resolution discussed in Appendix C.

Another mechanism for the decay of unstable germanium isotopes is through electron capture. When ^{71}Ge captures an electron, a cascade of energy in the form of gammas and auger electrons is emitted, where the energy from the decay is determined by the binding energy of the shell that captured the electron. The inner K-shell has the highest probability of capture and results in the emission of 10.37 keV. Each successive shell has a lower probability and emits less energy: the L-shell emits 1.3 keV, the M-shell emits 160 eV, etc. [93]. The activated germanium has a half-life of approximately 11 days, so the data collected immediately after a ^{252}Cf exposure will give the strongest electron capture peaks.

This process of activating the germanium is a convenient trick to essentially turn the detectors into their own calibration sources. These electron capture peaks will prove to be important in Chapter 4, Appendix C, and Appendix D.

3.1.3.3. Low-background Data

Low-background (low-bg) data is essentially the rest of the data collected by the experiment where there are no calibration sources present. This is the data that ultimately gets used for dark matter searches, although the calibration datasets are often used for defining cuts, calibrating the energy scale, etc.

3.1.3.4. Data Processing

SuperCDMS uses an optimal filter (OF) to extract useful information from the raw pulses collected. The OF requires a template to be input that it then shifts and stretches to fit the pulse to extract information such as the energy of the pulse. There is also a version called the non-stationary optimal filter (NF) that allows the fitter to shift the template in time to find the delay between the global trigger and the start of the pulse. Different templates can

also be fed in for different event classes and the change in χ^2 from the goodness of fit for each template can be used to classify the event type [94].

Here is a quick introduction to the nomenclature used for identifying energy quantities in the data:

S1: side one, top of the detector

S2: side two, bottom of the detector

qi: charge pulse in the inner electrode

qo: charge pulse in the outer electrode

pA: phonon pulse in the outer ring channel

pB: phonon pulse in the inner wedge channel

pC: phonon pulse in the inner wedge channel

pD: phonon pulse in the inner wedge channel

OF: optimal filter, used to extract information from the raw pulse

NF: non-stationary optimal filter, same as OF but allowed to shift in time

For example, the variable $ptNF$ is the energy extracted by the NF from the total phonon pulse of all channels co-added,

$$ptNF = NF(pt) = NF(pA + pB + pC + pD).$$ (3.1)

This is not to be confused with $psumNF$ that is the sum of the individual pulse energies extracted by the NF,

$$psumNF = NF(pA) + NF(pB) + NF(pC) + NF(pD).$$ (3.2)

3.1.3.5. Energy Resolution

A critical part of understanding the behavior and response of each detector is to measure its charge and phonon energy resolution, and to measure how that resolution changes with energy. Appendix C describes a detailed study that finds the charge and phonon resolution for all fifteen of the Soudan iZIPs.

3.2. SuperCDMS SNOLAB

SuperCDMS was chosen as one of three second generation dark matter experiments to be supported by the United States Department of Energy. SuperCDMS was charged with searching for low-mass dark matter (<10 GeV), LZ was charged with searching for high-mass dark matter (>10 GeV), and ADMX was charged with searching for axions.

Construction has already started on the next generation of SuperCDMS, commissioning will begin in 2020, and the first science run with the full complement of detectors will begin in 2021.

3.2.1. Shielding and Infrastructure

SNOLAB was chosen to be the location of the next stage of SuperCDMS because of the reduced environmental background levels. The laboratory is located in a nickel mine 2 km below Sudbury, Ontario. This converts to 6090 mwe, which reduces the cosmic muon flux by a factor of $> 5 \times 10^7$. The entire underground SNOLAB facility also operates as a class 2000 cleanroom to prevent dust contamination that carries radioactive isotopes, especially from the radon chain [95, 96].

Most of the shielding around SuperCDMS SNOLAB will be similar to the layers that were used at Soudan, except that the outer neutron shield will be made from water tanks instead of HDPE. The other major change for the SNOLAB experiment is that the whole
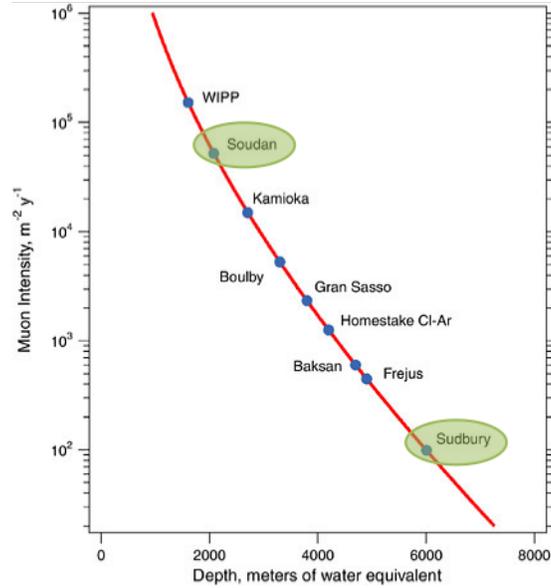
Figure 3.7: Plotting muon intensity versus depth shows that the laboratory under Sudbury has > 2.5 orders of magnitude fewer cosmic muons than the laboratory under Soudan. [97].

experiment will be constructed on a seismic platform to isolate it from laboratory vibrations; this is because parts of the mine are still active. Figure 3.8 shows the planned SuperCDMS shielding.
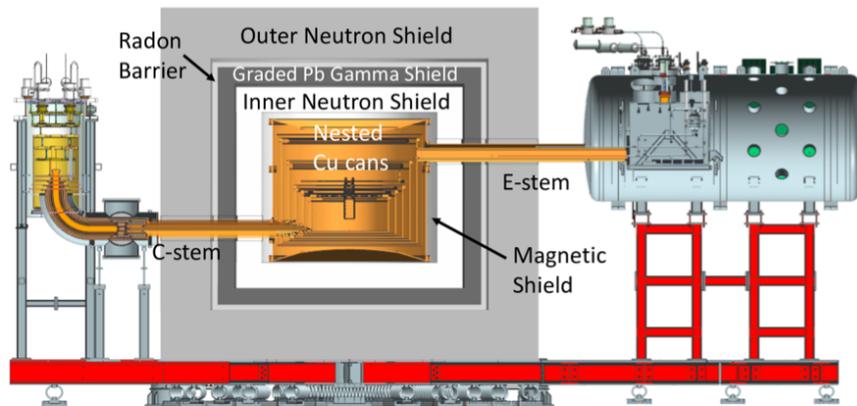
Figure 3.8: Side view of the layers of shielding planned for SuperCDMS [98]. The new additions from the Soudan shielding is an outer water tank and a seismic platform.

Inside the many layers of shielding are nested copper "cans" that make up the cryostat (also known as the "SNOBOX"), which is larger than the one at Soudan. The SNOBOX is capable of holding 31 towers, which was calculated to be the mass necessary for reaching

62

the neutrino floor. However, for the first stage there will only be four towers installed, with six detectors in each. Depending on science goals and funding, more towers can be added in future upgrades [98].

3.2.2. Detectors

There will be two types of detectors deployed at SNOLAB, and there will be germanium and silicon versions of each, for a total of four detector variations. The SNOLAB detectors will be bigger than the Soudan iZIPs for increased exposure. The new detectors will be 100 mm in diameter, 33.3 mm thick, and 1.39(0.61) kg for germanium (silicon) [98].

3.2.2.1. iZIP

Following the success of the iZIPs at Soudan, a new version of iZIP has been developed. The detector will still function the same way as the Soudan iZIPs with an electric field that is uniform in the bulk of the crystal and scalloped near the faces. The SNOLAB iZIPs will also still collect charge and phonon energy, meaning they will maintain the ability to discriminate between electron recoils and nuclear recoils. This discrimination will be necessary for understanding the background rates in the SNOLAB experiment and creating models for the high voltage detectors. There will still be an inner and outer charge channel per face, but there will now be six phonon channels per face. Figure 3.9 shows a picture of the new 100 mm iZIP and the phonon channel mask.

There will be 10 germanium iZIPs installed in the first stage, which is the most of any detector variation. This is because the collaboration has demonstrated a strong ability to understand this type of detector and produce interesting science results from operation at Soudan. The plan is to bias them with 6 V, which will lead to an ionization resolution of 100 eV and a phonon resolution of 50 eV [98].
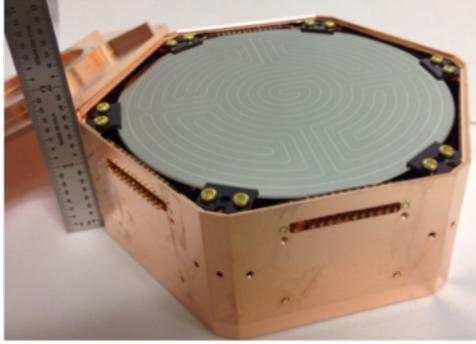
(a) Photo of a SNOLAB iZIP detector in its copper housing. There will be germanium and silicon versions of this detector.
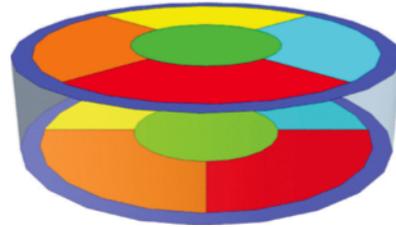
(b) Visualization of the channel layout for the SNOLAB iZIP [98].

Figure 3.9: View of the new iZIP designed for operation at SNOLAB.

There will also be 2 silicon iZIPs installed because the lighter nucleus is more sensitive to low-mass dark matter. This low mass reach is not expected to be better than the HV detectors described below, which is why there are only 2 detectors to begin with. The plan is to bias them with 8 V, which will lead to an ionization resolution of 110 eV and a phonon resolution of 25 eV [98].

3.2.2.2. High Voltage

A new type of high voltage (HV) detector has been developed for SNOLAB that is billed as the successor of the world-leading CDMSlite program. Since the charge electrodes were not used in CDMSlite, these new HV detectors will not have any charge electrodes patterned on the face. Instead, they will have a much higher coverage of TESs that will significantly improve the phonon collection efficiency. As with the redesigned iZIPs, there will be six phonon channels per face. Figure 3.10 shows a picture of the new HV detector and the phonon channel mask. An extra ring was added around the outside to provide better position information for high-radius events [99].

A significant improvement over the CDMSlite operation is that the HV detectors will have their own electronics that are capable of biasing both faces. This will shrink the region

(a) Photo of a SNOLAB HV detector in its copper housing. The face of the detector looks different than the iZIP detectors because there are no charge electrodes interleaved between the phonon TESs. There will be germanium and silicon versions of this detector.
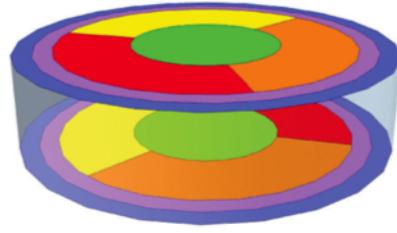
(b) Visualization of the channel layout for the SNOLAB HV detector [98].

Figure 3.10: View of the new HV detector designed for operation at SNOLAB.

at high-radius where events do not experience the full NTL amplification and increase the fiducial volume.

There will be 8 germanium HV detectors installed in the first stage. The plan is to bias them with 100 V, which will lead to a phonon resolution of 10 eV. There will also be 4 silicon HV detectors installed. The plan is to bias them with 100 V, which will lead to a phonon resolution of 5 eV [98].

By combining the lighter silicon nucleus with the HV operation, the silicon HV detectors will have the lowest energy threshold of the four new detector types and therefore be able to search for the lowest dark matter masses. However, calibrating the energy scale will be a serious challenge because silicon does not have activation lines like germanium. Since low-energy sources cannot penetrate the many layers of shielding, only high-energy sources can be used for external calibration. However, the high energy events will saturate the HV detectors. Appendix D describes an algorithm designed to extract the energy of saturated events so that high energy sources can be used to calibrate the energy scale.

3.2.3. Projected Sensitivity

Projected sensitivities have been calculated for each of the SuperCDMS SNOLAB detector variations and are shown in Fig. 3.11. The curves assume 5 years of operation, and the energy threshold is conservatively defined to be $7\sigma_p$ where σ_p is the phonon energy resolution [98].
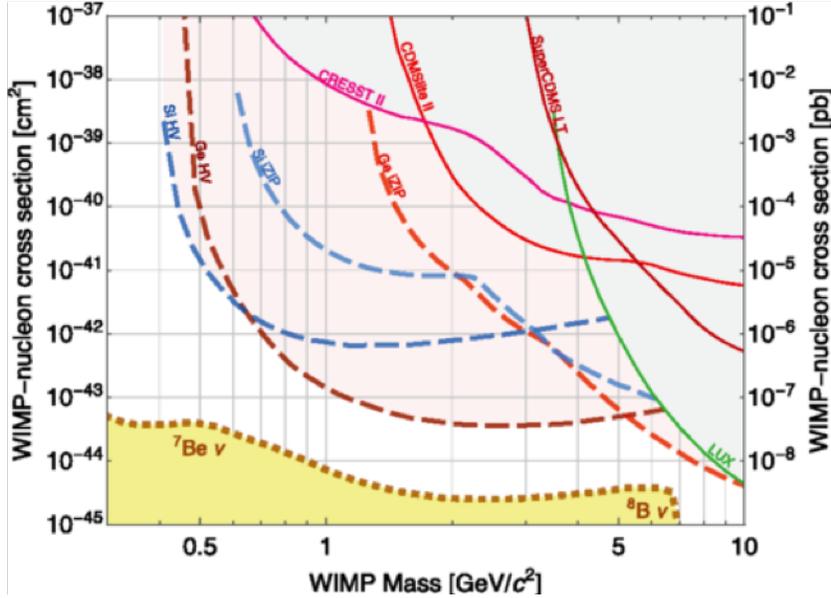
Figure 3.11: Projected sensitivity for each of the detector types in SuperCDMS SNOLAB compared to the current status of dark matter direct detection experiments [98]. The solid lines are the current world-leading exclusion limits except SuperCDMS LT was included for a low-mass comparison. The dashed lines are the projected SuperCDMS SNOLAB sensitivities, and the yellow band at the bottom is the neutrino floor where experiments expect to detect cosmogenic neutrino backgrounds.

CDMSLITE

In order to push the iZIP sensitivity to lower thresholds and correspondingly lower dark matter masses, a new mode of operation was devised called the "low ionization threshold experiment," or CDMSlite for short [100]. In this mode, a significantly higher voltage bias was applied across the crystal to take advantage of the NTL gain described in **??** . The existing electronics were repurposed to apply the new high voltage, but they were not capable of providing a high voltage bias to both sides of the detector, so one side was grounded and the other was biased to high voltage. There were two downsides to this choice. One is that only collecting data on one side negates the advantage of the electric field shape in the iZIP and the symmetric readout, described in Section 3.1.2.2. This makes it very difficult to define a fiducial volume. The other disadvantage is that it led to a distorted electric field with a large region at high radius where events do not experience the full NTL gain, shown in Fig. 4.1.
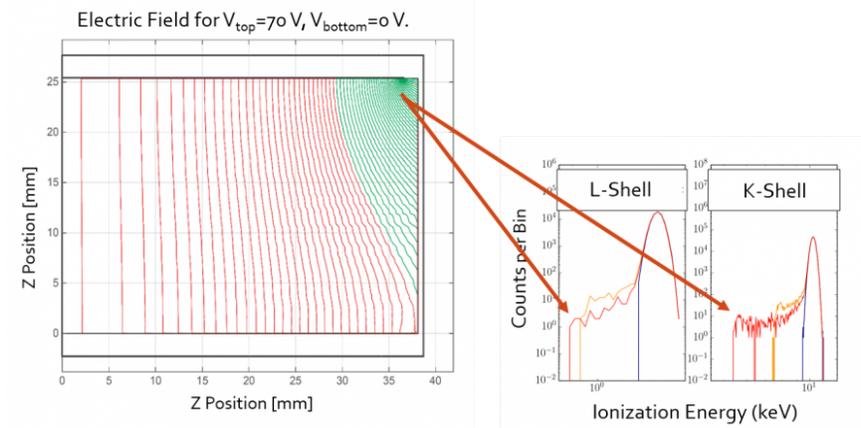
Figure 4.1: Electric field configuration for the one-sided bias of CDMSlite [93]. Events that occur in the red region will experience the full NTL gain, whereas events that occur in the green region will experience some diminished gain from a weaker electric field. These reduced gain events create low-energy tails that can leak down into the signal region.

While the CDMSlite mode of operation results in excellent resolution and a low energy threshold, it is relegated to being a phonon-only analysis. The electronics can still read out the charge electrodes, but the NTL gain amplifies the ionization signal by creating more phonons. Therefore, the un-amplified charge signals for a low-mass dark matter interaction were too small to distinguish from the electronic noise. This means that the CDMSlite analyses contend with higher background rates because they cannot calculate the yield quantity used for particle identification.

There have now been three runs in CDMSlite mode, which will be referred to as CDMSlite R1, R2, and R3.

4.1. The First CDMSlite Run

The first run was largely a proof of concept to modify the existing hardware and bias an iZIP to the highest possible voltage. Each detector was determined to have an upper limit to the voltage that can be applied before dark current began to leak across the semiconductor crystal and dominate the signal. Detector T5Z2 was chosen for the first run because it had one of the best baseline noise resolutions, a low background rate, and a lower affinity for leakage current [100]. Figure 4.2 shows that the signal-to-noise ratio begins dropping around 70 V [101] for this detector. That detector also had the advantage of not having any known shorts on the phonon channels; the list of shorted channels on all detectors can be found in Appendix C.

The detector collected an exposure of 6.2 kg-d and an analysis threshold of 170 eV_{ee} resulted in the final exclusion curve shown in Fig. 4.3. At the time, this was a world-leading limit for WIMP masses below 6 GeV [100, 101].
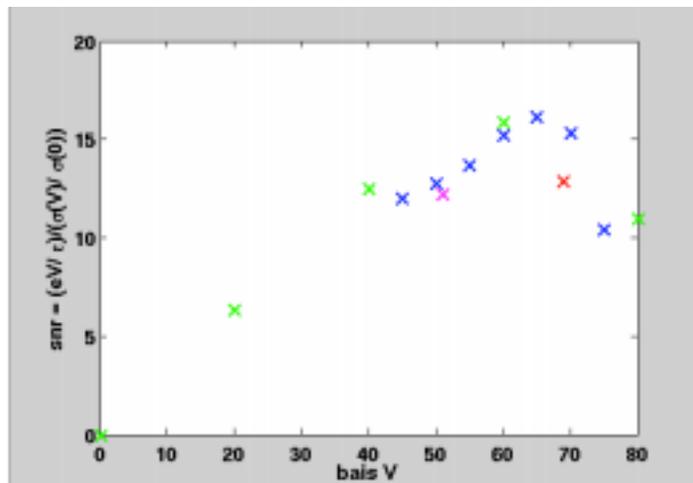
Figure 4.2: A plot of the signal-to-noise ratio running an iZIP in CDMSlite mode [101]. Here the ratio is calculated between the energy resolution of a pulse from barium calibration data and a randomly sampled noise trace. Because of the NTL gain, a higher bias was preferred until the voltage got high enough to make current leak across the semiconductor crystal. The optimum voltage was found to be $60 - 70$ V.
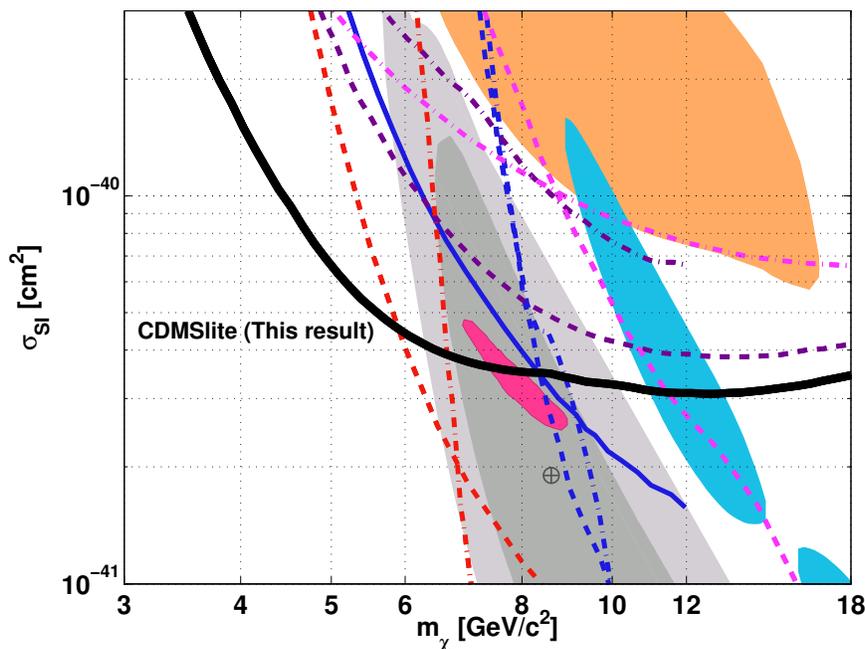
Figure 4.3: The black curve is the CDMSlite R1 WIMP exclusion limit [100]. The shaded blobs are contours where experiments saw an excess of events above expected backgrounds. These regions could be interpreted as WIMP signals, but there is clear tension with other results that have excluded that parameter space. The other curves are grouped by technology: blue are germanium bolometers, purple are point-contact germanium detectors, red is liquid xenon, and magenta are other technologies.

4.2. The Second CDMSlite Run

With a successful first run that led to a world-leading limit below 6 GeV/cm^2, the goal of the second CDMSlite run (R2) was to learn more about the behavior of the detector operating with a high voltage bias and publish the strongest possible result. CDMSlite R2 used the same detector (T5Z2) as R1 since its response to high voltage was better understood. However, R2 ran much longer for a total exposure of 70.10 kg-d.

4.2.1. Data Reduction

By definition, all of the SuperCDMS cuts on the data are named by what they return, hence cGoodEv returns good events. In order to get the opposite, one simply has to negate the cut (e.g. !cGoodEv_v53). The cut names are also appended with processing tags, such as "_v53," "_133," and "_CDMSliteR2."

List of cuts used to remove undesired classes of events:

cHighVoltage_CDMSliteR2: ensures the power supply was set to the correct bias voltage

cPMultTight_blind_v53: removes multiple scatter events because dark matter interacts so rarely that it is not expected to scatter more than once

cNuMI_CDMSliteR2: removes time blocks when the Neutrino at the Main Injector (NuMI) beam from Fermilab is on

cVTStrict_133: removes coincident events with the external muon veto

cRandom_133: selects randomly triggered events to monitor noise environment

List of quality cuts used to clean up the data:

cGlitch_trig_CDMSliteR2: remove electronic misfirings, aka "glitches," whose pulses are un-physically fast and occur in multiple detectors

cPstd_CDMSliteR2: standard deviation cut on the prepulse noise checks that the phonon pulses start from baseline noise levels

cQChiSq_CDMSliteR2: charge chi-squared cut removes glitch triggers that occur in all detectors without issuing a charge trigger

cBadSeries_CDMSliteR2: removes bad series that experienced large trigger bursts, series that did not experience the correct biasing scheme, or series with abnormal chi-squared behavior

cBadGPStime_CDMSliteR2: time cut to remove events where timestamps did not agree with GPS time

cBadTempCurrentFit_CDMSliteR2: removes events whose energy cannot be corrected because they do not have correct base temperature, current through the HV power supply, or bad template fits in the OF

List of cuts based on phonon chi-squared values:

cPChiSq_CDMSliteR2: removes pile-up events where there are multiple pulses in the pulse window

cLFnoise_CDMSliteR2: removes low-frequency noise (LFN) events that can mimic a low-mass dark matter pulse by using pulse shape discrimination (PSD)

cGlitch_chisq_CDMSliteR2: removes some glitches that get past the previous glitch cuts by using PSD

List of physics cuts:

cRad2T_CDMSliteR2: uses ratio of 2-template OF (fast and slow) to remove high-radius events that do not experience the full NTL amplification due to reduced electric field

Fiducialization was a significant improvement for R2 over the R1 analysis. A 2-template OF was developed that fits a typical "slow" template to the pulse to get information from the homogeneously distributed phonon energy. The residual pulse from subtracting the slow template from the pulse is fit again with a "fast" template that is strongest for the phonon channel nearest to the event location. The ratio of the amplitudes from the fast and slow template fits is used to define a radial parameter that can be cut [4].

Arguably the most impactful lesson from R1 was that low frequency noise limited the threshold. It was believed that this noise was caused by the mechanical action of the cryocooler pump, so an accelerometer was installed in order to detect the motion and reject coinciding events. CDMSlite R2 was naturally split into two data-taking periods (R2a and R2b) by cryocooler maintenance. The trigger threshold was lowered between the two periods and the addition of the noise monitor allowed for the analysis thresholds to be lowered from 75 eV$_{ee}$ for R2a to 56 eV$_{ee}$ for R2b [4].

4.2.2. Optimal Interval

The R2 limit, as well as R1, was calculated using the Optimum Interval (OI) method developed by Steve Yellin [102, 103]. The only pieces required for setting a limit were the final spectrum of events after all cuts, the shape of the expected signal multiplied by the efficiency of those cuts, and the total exposure of the final dataset. The OI is not equipped to handle multiple datasets, so R2a and R2b were combined. Some of the analysis pieces were developed for the full dataset, such as the resolution model.

The strength of the OI technique is setting a limit in the presence of unknown backgrounds. However, the OI is incapable of being used to discover dark matter; it can only be used to calculate an exclusion limit.

4.2.3. Published Limit

The published exclusion limit from the R2 WIMP search is shown in Fig. 4.4, which was a world-leading limit for WIMP masses between 1.6 GeV/cm^2 and 5.5 GeV/cm^2 when submitted for publication [4, 93]. An exploration of the systematic effects throughout the analysis and assumptions in the WIMP model has been published for both R1 and R2 [104].
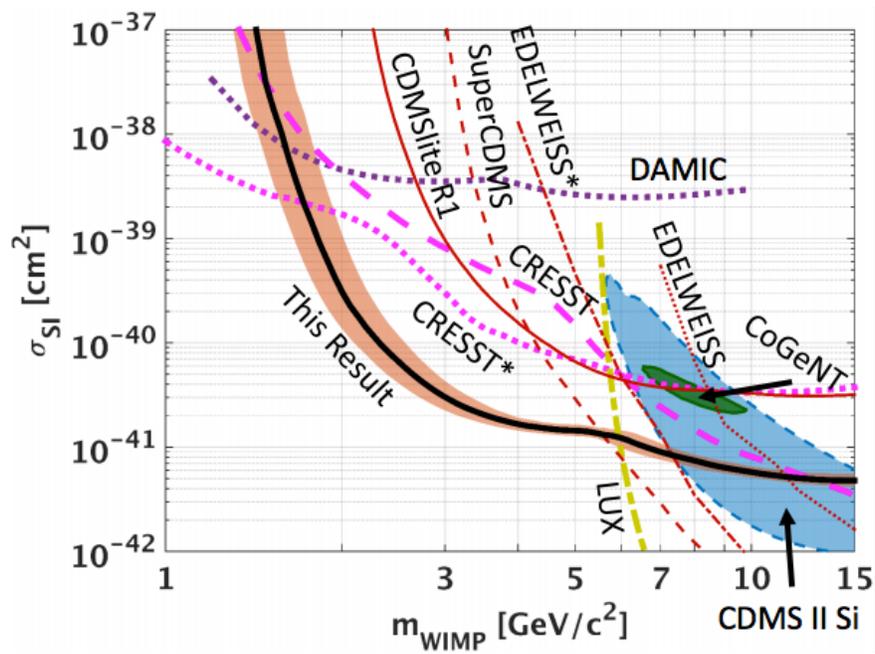
Figure 4.4: The black curve is the CDMSlite R2 WIMP exclusion limit, with the uncertainty band shown in orange [4]. The solid red line is the previous CDMSlite result, which had been surpassed by the CRESST results in pink and the DAMIC results in purple. The limits with an asterisk, dotted CRESST and dot-dashed EDELWEISS, were published around the same time as this result and were still under review when this plot was produced. The shaded contours are where experiments saw an excess of events above expected backgrounds, which were originally interpreted as WIMP signals and have since been excluded.

4.2.4. Likelihood Analysis with R2 Data

The OI method used for R1 and R2 is a very useful tool for quickly calculating an exclusion limit, especially if the dataset is not fully understood. However, if all of the backgrounds are understood and can be modeled, then more powerful limit setting techniques can be used to discriminate between backgrounds and the desired signal (often mistakenly referred to as "background subtraction"). An unbinned profile likelihood ratio is the technique chosen for the inelastic dark matter search in this thesis because it has the potential to claim a signal discovery. Even if there is no excess of events to claim a discovery, a likelihood exclusion limit will be stronger than an OI limit because of the background discrimination.

Since the R2 analysis was optimized for an OI limit, a significant amount of work was needed to create the R2 models for use in a likelihood analysis.

4.2.5. Background Modeling

The first pieces needed for calculating an unbinned profile likelihood limit are probability density functions (PDFs) that model the spectral shape for each source of undesired events in the experiment, called backgrounds, that can mask or mimic the desired dark matter signal. The backgrounds modeled for this analysis are neutron activation by the ^{252}Cf calibration source, cosmogenic activation of the crystal, Compton scattering from radioactive materials in the experiment, and ^{210}Pb contamination on or near the detectors.

The most dominant background source is comes from activation of the germanium crystal via the neutron calibration source. When stable ^{70}Ge in the crystal captures a neutron, it becomes unstable ^{71}Ge that decays via electron capture. The electron capture creates a cascade of energy in the form of gammas and Auger electrons, where the energy emitted from the decay is equal to the binding energy of the shell that captured the electron. The outer K-shell has the highest probability of capturing an electron at 87.57% and results

in the emission of 10.37 keV. Each successive shell has a lower probability and emits less energy: the L-shell captures an electron 10.53% of the time and emits 1.3 keV, the M-shell captures an electron 1.78% of the time and emits 160 eV, etc. [93]. This intrinsic background dominates all other background rates by an order of magnitude. However, the distinct peaks have such high event counts that they actually become useful when calibrating the energy scale of the detectors, calculating the efficiency of cuts, and measuring the energy resolution of the detectors. This background is modeled with a Gaussian distribution for each electron shell peak. The amplitudes of these peaks are set relative to the K-shell peak as determined by the probabilities of each shell capturing an electron. There is one overall normalization parameter in the likelihood function for the ^{71}Ge background.

Before the detectors are brought underground where the overburden shields them from cosmic rays, the crystals are fabricated on the surface of the Earth and transported to the mine. During this time, the interaction of cosmic rays can knock nucleons from the germanium atoms in the crystal and create radioisotopes. Because of the way they are created, these isotopes are called "cosmogenics." One of the more problematic cosmogenics that can be ejected is ^{3}H, commonly known as tritium, which provides a constant source of betas in the detectors because of its long half-life of 12.32 years. The tritium background can be modeled with a standard beta emission spectrum with an endpoint energy of 18.6 keV [105, 106].

Tritium is the dominant source of cosmogenic background in the CDMSlite data, but other cosmogenic isotopes are created as well. As cosmic rays eject nucleons, unstable residual nuclei can remain. Those residual nuclei are considered as backgrounds if they have a half-life that is long enough that they will not decay away before data taking begins, but also short enough that the activity is comparable to other backgrounds. The isotopes modeled in this analysis are ^{68}Ga, ^{65}Zn, and ^{55}Fe [105]. Other isotopes considered were ^{57}Co, ^{54}Mn, and ^{49}Vn, but fits to the R2 spectrum showed that they contributed less than

75

one expected event and could be neglected. Each of the modeled isotopes decay via electron capture in the same way as the activated germanium. As with ^{71}Ge, they all have K-, L-, and M-shells that dictate the energy of the decay. Contributions from each of these peaks are modeled with a Gaussian distribution with fixed relative amplitudes with a single normalization parameter in the likelihood.

In addition to cosmogenic backgrounds, there are also "radiogenic" backgrounds created by the decay of radioisotopes near the detector, such as the radon contamination described in Chapter 2. Typically decays have higher energy and would not be considered in the low energy analysis region, but high energy photons can lead to a Compton scatter where a fraction of the incoming energy is deposited in the detector. To first order, this creates a flat background throughout the analysis region. On closer inspection, there are "steps" at low energy where the scattered photon has enough energy to overcome the binding energy of an electron shell and there is a small jump in rate. This Compton background is modeled as a flat contribution with an error function for each shell to model the steps [107].

The shape of the background models is what eventually determined the energy range of the analysis. The challenge is breaking the degeneracy in the likelihood fit between models that are flat or nearly flat in the energy spectrum. Without external measurements to constrain the backgrounds, the easiest way of breaking this degeneracy was to extend the fit to higher energy. The idea was to extend the analysis range above the endpoint energy of the tritium spectrum, where the only background is the Compton model. This sets the Compton normalization parameter across all energies. Similarly, tritium is the only other background model above the ^{71}Ge K-shell, so the normalization parameter was determined by this part of the spectrum. The other backgrounds described in this section were Gaussian peaks located at specific energies, so there was no degeneracy to worry about with them. Therefore, the final analysis range used was $0 - 25$ keV.

All of the activation and cosmogenic background models described above are distributed nearly uniformly throughout the detector volume, so the efficiency of the physics cuts that were developed for a uniform WIMP signal can also be applied to those models. This allows for the backgrounds to be modeled by analytic functions. However, another background class exists for radon plate-out on the crystal detectors or the surrounding copper housing. Although these events can be implanted below the surface of the crystal, they are predominantly classified as surface events. To understand how these surface events manifest in the detectors, a simulation of a detector tower was contaminated with ^{214}Pb from the ^{222}Rn decay chain. By contaminating with ^{214}Pb, subsequent alpha decays in the simulation can implant the long-lived ^{210}Pb, mimicking the physical contamination [107].

A simulated map of the voltage inside the detector was used to model the detector response to the simulated events. It was determined that there were two main systematics to account for: variability in the map and uncertainty in the location of a voltage-based cut [107]. The output from the simulation was a binned model that can potentially introduce a lot of discontinuities in an otherwise unbinned likelihood; this can cause issues when maximizing the likelihood function.

A Kernel Density Estimator (KDE) was used to smooth the binned models by assigning a non-negative function, or "kernel," to each discrete data point and summing the kernels to get the continuous distribution [108, 109]. Because the kernels are typically summed up to the boundary, KDEs typically struggle to maintain correct boundary conditions. The most robust way to avoid this side effect is to mirror the data at the boundary and then truncate the model after smoothing. This process was repeated for both the median models and the calculated systematics. Then the systematic errors were added in quadrature with the statistical error on the median to obtain a single smooth uncertainty band.

The key to using a KDE is choosing an appropriate bandwidth because a bandwidth that is too small will create artificial features from statistical fluctuations and a bandwidth that is too large will mask the structure of the data. This effect is demonstrated in Fig. 4.5, where the blue curve is the resulting model from different choices of bandwidth [110].
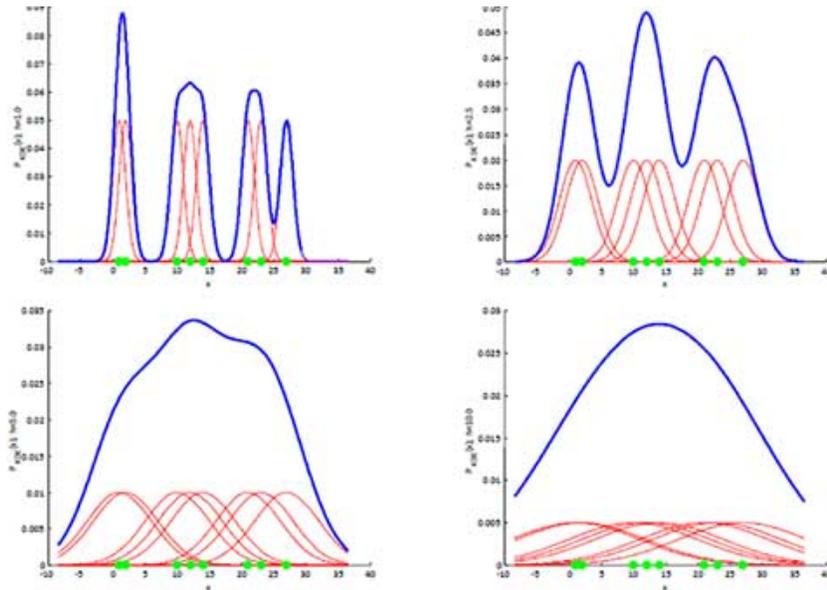
Figure 4.5: Demonstration of the importance of choosing an appropriate bandwidth when using a Kernel Density Estimator [110]. The green data points are the same in each plot, with a red kernel drawn for each point. The blue curve is the sum of the red kernels. The only difference between the four plots is the bandwidth of the red Gaussian kernels, which illustrates how different bandwidths can lead to drastically different distributions.

The problem with smoothing a simulated spectrum is that it can be difficult to determine which features are fluctuations that should be smoothed and which are real features that should be preserved. The first method used to estimate an appropriate bandwidth is known as Scott's Rule of Thumb,

$$w = n^{-1/(d+4)} \, , \tag{4.1}$$

where n is the number of data points and d is the number of dimensions [111]. However, this approach produced a bandwidth that was too small and the features were too sharply resolved. In the end, an adaptive bandwidth was used that did a better job of smoothing flat

sections of the model but preserved the shape where the model changed more sharply [112]. The smoothed models and systematic errors are shown in Fig. 4.6.
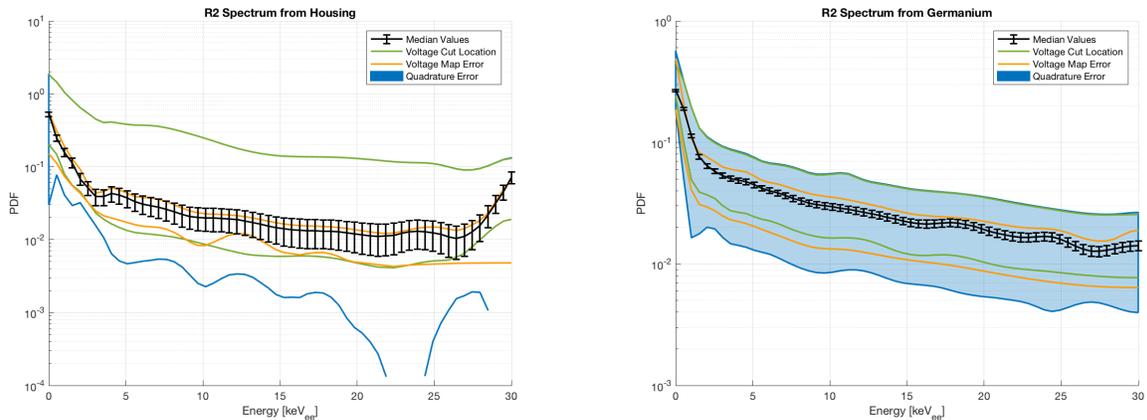
(a) There is a blue curve under the upper green curve because the green error completely dominates the orange error and the black error in the quadrature sum.

(b) Preserving the large spike at low energy was an area of focus because KDEs typically struggle to preserve sharp changes at the boundaries.

Figure 4.6: This figure contains the smoothed surface background models from the KDE for (a) the copper housing and (b) the germanium crystal. The orange curves represent the systematic error on the voltage map used in the simulation and the green lines represent the systematic error of the voltage cut location. The black median curve has been scaled to a pdf with unit area and is shown with statistical error bars. The green and orange lines have been scaled to maintain the ratio with the black curve, and the blue curves are the quadrature sum of the statistical and systematic errors.

The results of the simulation were normalized using an independent measurement of the rate of alpha events in each of the iZIP detectors [113]. The iZIP data contains more information about the position of an event than CDMSlite mode. Thus, the measured alphas could be attributed to the simulated source. For example, the measured rate of alphas on the top surface of the detector are divided by two to normalize the simulated alphas on the germanium. The factor of two comes from the CDMSlite R2 detector being in the middle of the tower. As such, its face is directly adjacent to the bottom face of the detector in the top of the tower. Then the germanium rate is scaled by the ratio of the area between the face and the sidewall to determine the number of sidewall events that originated from

the germanium. Subtracting this from the total measured sidewall rate gives the number of events from the copper housing that faces the sidewall [107].

For the $0 - 25$ keV energy range of this analysis, the normalization procedure results in $N_{SH} = 12.1 \pm 1.0$ events from the sidewall housing (SH), and $N_{GZ} = 10.9 \pm 1.7$ events from the germanium iZIP (GZ), with a correlation coefficient of $\rho_{SH,GZ} = -0.38$ between them. Correcting those numbers for the efficiency of the cuts used in the simulation gives the predicted number of counts to be $N_{SH} = 24.0 \pm 2.0$ and $N_{GZ} = 21.6 \pm 3.4$ events. The uncertainty on these normalizations can be multiplied through the correlation matrix between the sources to give the covariance matrix:

$$\begin{bmatrix} 4.00 & -2.584 \\ -2.584 & 11.56 \end{bmatrix} \tag{4.2}$$

4.2.5.2. Morphing the Surface Backgrounds

As seen in Fig. 4.6, the uncertainty band does not have the same shape as the median curve. Thus, the likelihood needs to account for a shifting shape that is correlated with the normalization uncertainty regarding these surface backgrounds. In order to implement this correlated shift in the likelihood, a fit parameter was added to "morph" the model up or down and correspondingly change the shape. A significant benefit of this approach is that only one parameter is needed in the likelihood for each surface background source. The morphed background model can be calculated as

$$\rho(E, m) = \begin{cases} \rho_{med}(E) + m \times (\rho_{upper}(E) - \rho_{med}(E)) & m > 0 \\ \rho_{med}(E) + m \times (\rho_{med}(E) - \rho_{lower}(E)) & m < 0 \end{cases}, \tag{4.3}$$

where m is the morphing parameter that is allowed to float in the likelihood. The PDFs in Fig. 4.6 have been scaled to be event densities by multiplying them by the predicted normalization: ρ_{med} is the median curve, ρ_{upper} is the upper bound of the 1σ uncertainty band, and ρ_{lower} is the lower bound of the 1σ uncertainty band. Equation 4.3 effectively defines the morphing parameter in sigma-space, where $m = 0$ will reproduce ρ_{med}, $m = -1$ will reproduce ρ_{lower}, and $m = 1$ will reproduce ρ_{upper}. Figure 4.7 demonstrates the performance of the morphing parameters by linearly sampling m from -1 to 1.
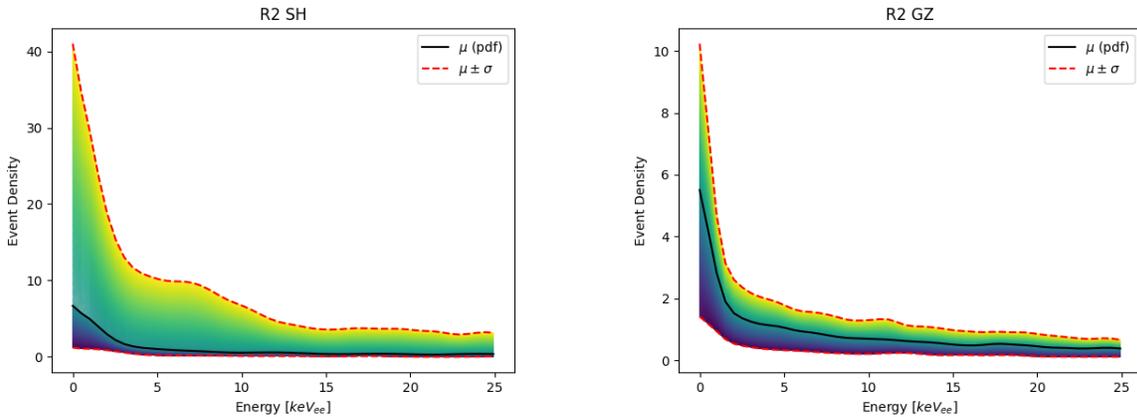
(a) Event density model of surface events from the sidewall housing.

(b) Event density model of surface events from the germanium.

Figure 4.7: Linear sampling of the morphing parameter that controls the model for surface evens from (a) the sidewall housing and (b) the germanium crystal. There are 500 sampled lines on each of these plots with purple curves corresponding to $m = -1$ and yellow curves corresponding to $m = 1$.

The sampling in Fig. 4.7 was linearly varied, but the morphing parameters will be constrained with a Gaussian distribution with mean $\mu = 0$ and width $\sigma = 1$. This will force the likelihood to preferentially fit the median curve or pay a statistical penalty to push the morphing parameter away from the median. In order to calculate the relationship between the Gaussian constraint and the normalization, a one-dimensional Gaussian is randomly sampled. For each sampling, the model is morphed according to the morphing parameter in Eq. 4.3 and the number of events is calculated using Eq. 4.4. Histograms of the number of

events from each morphed model are shown in Fig. 4.8.

$$N_x = \int \rho_x(E)dE\,. \tag{4.4}$$
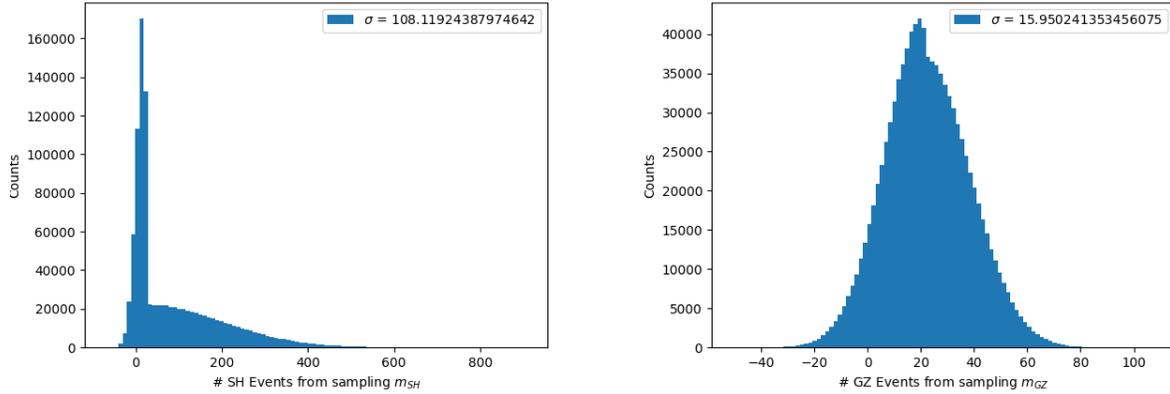
(a) The Gaussian sampling for the sidewall housing produces an asymmetric distribution because the median curve is not equidistant from the upper and lower uncertainty. Said another way, the sampled curves in Fig. 4.7 are much closer together below the median.

(b) The Gaussian sampling for the germanium produces the expected Gaussian shape because the median is roughly centered in the uncertainty band.

Figure 4.8: Histograms of the number of events from Gaussian sampling of morphing parameters.

Randomly sampling the Gaussian distribution provides an estimate of the uncertainty on the number of events for each surface background source. The covariance matrix due to the morphing parameter uncertainty is:

$$\begin{bmatrix} 11494.07 & 1506.58 \\ 1506.58 & 226.47 \end{bmatrix} \tag{4.5}$$

The matrix in Eq. 4.2 can now be added to the matrix in Eq. 4.5 to give the total covariance on the number of surface background events for each source:

$$\mathbf{V} = \begin{bmatrix} 11498.07 & 1503.99 \\ 1503.99 & 238.03 \end{bmatrix}. \tag{4.6}$$

The final step is to translate the correlated uncertainty from the number of surface background events to uncertainty on the morphing parameters so that it can be added to the likelihood function as a multivariate Gaussian constraint. This is done by plotting the relationship between morphing parameter and number of events as defined in Eq. 4.4. Figure 4.9 shows this relationship using the data from the linearly sampled morphing parameter in Fig. 4.7.
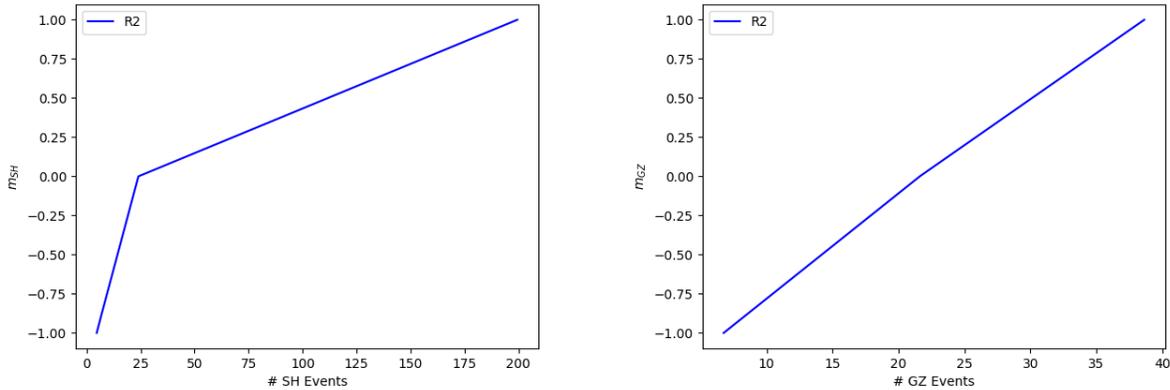
(a) The sharp angle at $m_{SH} = 0$ is related to the asymmetry of the histogram in Fig. 4.8a. This is showing that changing the morphing parameter from -1 to 0 does not result in a large change of number of events.

(b) The dependence on number of events versus morphing parameter has a kink at $m = 0$ that is barely perceptible. This is because of the equidistance of the median from either bound of the uncertainty band.

Figure 4.9: The relationship between morphing parameter and resulting number of events is shown for each surface background source. The independent variable is plotted on the y-axis for these plots to show to the slopes needed in Eq. 4.7.

Following standard error propagation rules, the covariance on the morphing parameters can be calculated according to

$$\mathbf{M}_{ij} = \sum_{k,l}^{2} \frac{\partial m_i}{\partial N_k} \frac{\partial m_j}{\partial N_l} \big|_{x=\mu} \times \mathbf{V}_{kl} \, , \tag{4.7}$$

where $m_1 = m_{SH}$, $m_2 = m_{GZ}$, $N_1 = N_{SH}$, and $N_2 = N_{GZ}$. The resulting matrix in Eq. 4.8 is the covariance matrix between the morphing parameters that will go into the Gaussian constraint in the likelihood to constrain the surface backgrounds:

$$\mathbf{M} = \begin{bmatrix} 11498.07 & 1503.99 \\ 1503.99 & 238.03 \end{bmatrix} . \tag{4.8}$$

A demonstration of the Gaussian constraint is shown in Fig. 4.10, where the blue lines are drawn by randomly sampling the covariance matrix in Eq. 4.8.
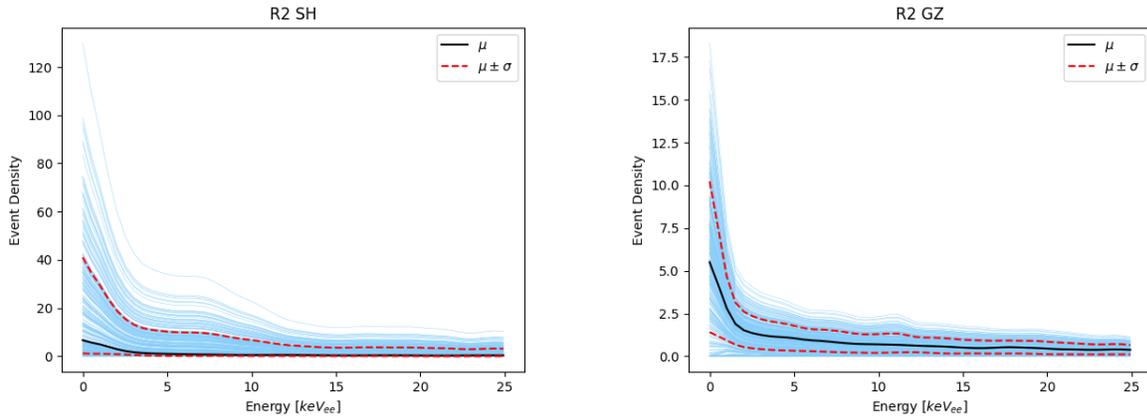
(a) Randomly sampled surface background model for the sidewall housing in R2.

(b) Randomly sampled surface background model for the germanium iZIP in R2.

Figure 4.10: The black curve is the mean model, and the red curves are the 1σ uncertainty band. The blue curves are created by randomly sampling the covariance matrix in the multivariate Gaussian constraint.

4.2.6. Efficiency Model

In order to match the data, all of the models are multiplied by the efficiency of the cuts used to obtain the final data spectrum. Descriptions of these cuts can be found in the long paper about CDMSlite R1 and R2 [104]. The resulting efficiency curve for CDMSlite R2 is shown in Fig. 4.11. The lowest energy portion of the efficiency curve is best described by an error function, however, the R2a and R2b had different threshold energies. When the efficiency from each sub-run is added, it creates the "double-hump" feature seen near 0.1 keV. In order to extend the efficiency curve above 2 keV, the 10.37 keV k-shell peak was used to calculate the efficiency of the cuts and the curve was linearly interpolated between the two measurements. It was originally assumed that the efficiency would be flat above the k-shell, but saturation effects were observed at energies above 17 keV. Thus, the barium calibration data was used to calculate a linear correction in this high-energy regime [105].
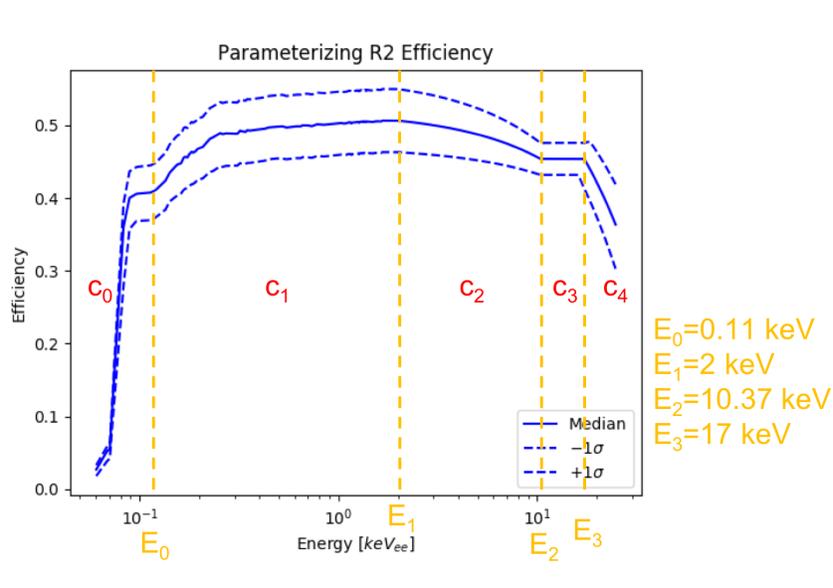
Figure 4.11: Annotated plot of the combined CDMSlite R2 efficiency curve. The solid blue line is the median efficiency curve and the dashed blue lines indicate the 1σ uncertainty on the efficiency. The region labeled c_0 is the expected shape for the low energy efficiency, and region c_1 is the "double-hump" from adding the efficiency from R2a and R2b. Region c_2 is the linear interpolation between the 2 keV analysis range and the 10.37 keV k-shell peak, and c_3 is the presumed flat efficiency. Region c_4 is where saturation was observed and the efficiency was linearly corrected by comparing to barium calibration data.

Due to the inherent uncertainty, the efficiency curve was included in the likelihood function as nuisance parameters. In order to do that, the data-driven curve needed to be modeled so the likelihood can control the parameters of that model; this is known as "parameterization."

The first attempt at modeling the efficiency curve used a piecewise function according to the different efficiency regions described in Fig. 4.11. The difficulty was finding a function that closely matched the shape in region c_1, which ultimately was an exponential function. When parameterizing a model, one needs to be aware of the number of parameters being added to the likelihood function because more parameters will make it more difficult for the fit to converge on the global minimum. By forcing the piecewise function in Eq. 4.9 to be continuous and avoid any discontinuities, the fewest free parameters needed to accurately describe the curve was eight, colored in orange. The piecewise function was defined as

$$Eff_{R2}(E) = \begin{cases} f_0 = Erf(E, \mu, \sigma, A_1) & E < E_0 \\ f_1 = A_2(1 - exp(-\frac{E-E_0}{\tau}) + f_0(E_0) & E_0 \leq E < E_1 \\ f_2 = m_1(E - E_1) + f_1(E_1) & E_1 \leq E < E_2 \\ f_3 = m_2(E - E_2) + f_2(E_2) & E_2 \leq E < E_3 \\ f_4 = m_3(E - E_3) + f_3(E_3) & E_3 \leq E \end{cases} \quad (4.9)$$

Besides the large number of parameters, another drawback of this piecewise approach was the addition of a Gaussian constraint to the likelihood function that required a covariance matrix to encode the uncertainty and correlations between the parameters. The off-diagonal correlations between different functional pieces mixed units and did not make physical sense. Instead, the efficiency curve was handled in the same way as the surface backgrounds, with a single morphing parameter as defined in Eq. 4.3. The resulting morphed efficiency curve

is shown in Fig. 4.12, where the morphing parameter is linearly sampled between $m = -1$ and $m = 1$.
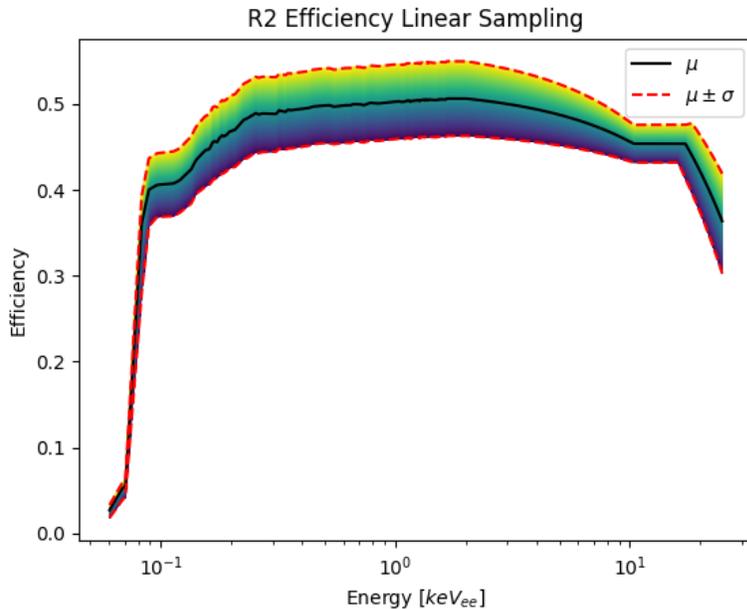
Figure 4.12: Linearly sampled morphing parameter. The purple lines correspond to $m = -1$, the yellow lines correspond to $m = 1$, and $m = 0$ will return the median black curve.

Since there is just one morphing parameter for R2, a one-dimensional Gaussian distribution is used to constrain the efficiency model in the likelihood. This creates a penalty for moving further away from the median, yet the parameter is not constrained to $|m| \leq 1$. Figure 4.13a shows the distribution from sampling this constraint many times. A strange side-effect of the morphing procedure shows up in the blue curves near 17 keV. A zoomed view of this region in Fig. 4.13b shows that the upper (lower) uncertainty band actually has a very small bump (dip) before the efficiency drops, but the morphing parameter exaggerates all features in the data for values of $|m| \geq 1$.

Although the morphing parameter allows for shape and normalization uncertainties, the normalization of the efficiency curve is not a significant factor in the likelihood minimization because all of the derived models except for the signal are converted to PDFs by normalizing
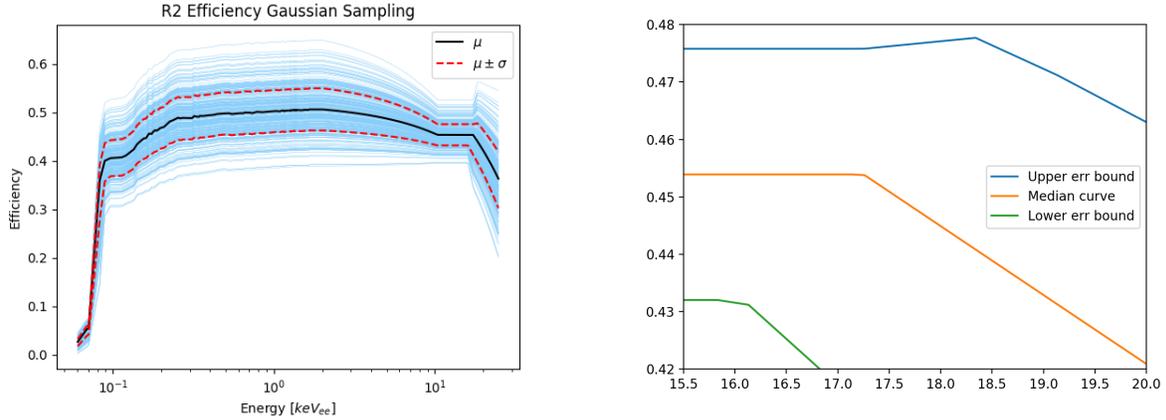
(a) Randomly sampled efficiency model for R2.

(b) Zoomed view of the efficiency curve for R2 shows a small bump near 17 keV. The morphing parameter exaggerates this feature for values of $|m| \geq 1$, which creates the sharp feature seen in Fig. 4.13a.

Figure 4.13: Exploring the Gaussian constraint for the morphed R2 efficiency model.

them to unit-area after they have been multiplied by the efficiency. Therefore, the likelihood function developed to model the data has these shapes constrained by independent measurements and the normalization of the efficiency.

4.2.7. Resolution Model

The background models and signal model are then smeared by the resolution of the detector to match the data. Thus, a parameterized model of the resolution was included in the likelihood function as nuisance parameters. Parameterizing the resolution was more straightforward than the efficiency because the functional form of the detector resolution as a function of energy has been well established by the SuperCDMS collaboration. A full description of the resolution model can be found in Appendix C, in particular the function in Eq. 3.1.

As part of the R2 analysis, the resolution of the baseline noise at 0 keV and the germanium activation peaks at 160 eV, 1.3 keV, and 10.37 keV were measured to span the analysis energy

region of interest. Fitting the resolution model to those measurements provided the best-fit parameters, uncertainty on those parameters, and the correlations between them. Those fits are shown in Fig. 4.14 and the results are summarized in Table 4.1 [4].
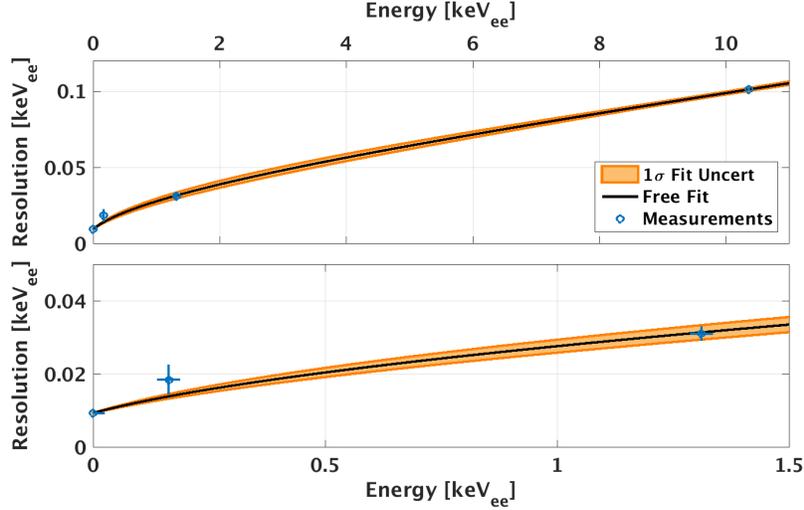
Figure 4.14: The blue points are the measured resolutions of the baseline noise and germanium activation peaks. The black curve is the best fit of the resolution model in Eq. 3.1, and the orange band is the uncertainty of that fit. The top plot shows the measured peaks up to 11 keV, while the lower plot is zoomed in below 1.5 keV.

From these results, it is straightforward to write down the correlation matrix and multiply it through by the uncertainties of each of the model parameters. The resulting covariance matrix is in Eq. 4.10 and will be used to constrain the resolution model in the likelihood function:

$$\mathbf{R} = \begin{bmatrix} 1.21 \times 10^{-8} & -2.29 \times 10^{-10} & 1.73 \times 10^{-9} \\ -2.29 \times 10^{-10} & 1.21 \times 10^{-8} & -1.02 \times 10^{-7} \\ 1.73 \times 10^{-9} & -1.02 \times 10^{-7} & 8.84 \times 10^{-7}. \end{bmatrix} \tag{4.10}$$

A demonstration of the Gaussian constraint is shown in Fig. 4.15, where the blue lines are drawn by randomly sampling the covariance matrix in Eq. 4.10.

Peak	Energy [keVee]	Resolution [eVee]	Res. Err. [eVee]
Baseline	0	9.25	0.11
M-shell	0.16	18.6	4.2
L-shell	1.3	31	2
K-shell	10.37	101	1

Run	A [eV]	B [eV]	C
R2	9.26 ± 0.11	0.64 ± 0.11	$(5.68 \pm 0.94) \times 10^{-3}$

Run	ρ_{AB}	ρ_{BC}	ρ_{CA}
R2	-0.0189	-0.984	0.0167

Table 4.1: Table showing the measured resolutions of known features in the data (top), resolution model parameters from fitting to the measured widths (middle), and correlations between the model parameters (bottom).
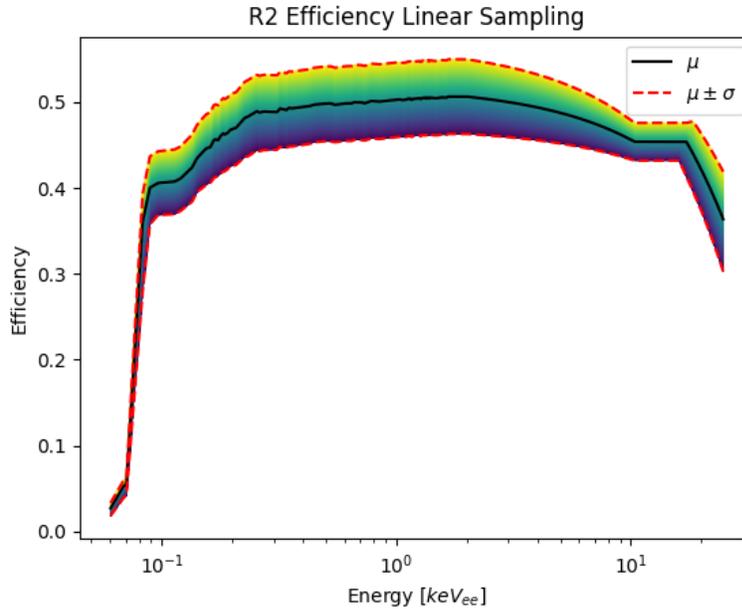
Figure 4.15: Demonstration of Gaussian constraints for the R2 resolution model. The black curve is the mean model, and the red curves are the 1σ uncertainty band. The blue curves are created by randomly sampling the covariance matrix in the multivariate Gaussian constraint. The blue curves account for correlations between the resolution parameters, which is why they do not fully sample the uncertainty band between the red curves.

4.3. The Third CDMSlite Run

After R2, the goal of the third CDMSlite run (R3) was to expand the CDMSlite program and develop analysis tools that could be used when the next generation experiment at SNOLAB starts taking data. Searching for low-mass dark matter is the explicit design goal of SuperCDMS compared to other experiments, so it has been important to extend our knowledge and experience during the CDMSlite program to ensure the success of the SuperCDMS SNOLAB HV program.

The first challenge was to prove that another detector could successfully be operated in CDMSlite mode, so the R3 data was collected with T2Z1 [54]. In addition to demonstrating the ability to yet again adapt the electronics, this provided an opportunity to compare the different detectors and learn more about their behavior in high voltage mode.

As with R2, the R3 dataset was split into R3a and R3b due to a changing noise environment that prompted different cut efficiencies and resolutions for each of the designated periods. The two periods were kept separate in the analysis, but the signal and background models were normalized for the combined dataset. The livetime after all cuts was 31.50 days for R3a and 29.39 days for R3b, with a combined total exposure of 36.92 kg-d and an energy threshold of 70 eV.

4.3.1. Data Reduction

Each CDMSlite analysis builds upon its predecessor. As such, all of the cuts developed for R2 were adapted to the R3 data. The following list contains cuts that were developed in addition to those in Section 4.2.1.

List of cuts not applied to R2:

cIsSquarePulse_CDMSliteR3: square glitch events are removed with a delta chi-squared cut that compares the fit of a normal template to a glitch template in the OF

cIsLFnoisedX2_CDMSliteR3: another delta chi-squared cut that compares the fit of a normal pulse to a LFN template in the OF

bifurcated LFN: cut assisted by a boosted decision tree (BDT) to ensure negligible LFN leakage events

The radial cut was changed from strictly empirical data to use simulated pulses in an attempt to more accurately track the true radius of observed events. The fluctuations seen in the efficiency curve in Section 4.3.5 are caused by the limited statistics of those simulated events when placing a cut on the calculated radial parameter [54].

In order to perform a likelihood analysis, all backgrounds in the data need to be modeled. As a corollary, if a model is not known for a particular background source, then a strict cut needs to be defined such that <1 event is expected in the final data spectrum. That is why a strict BDT-based bifurcated analysis cut was developed to remove LFN from the R3 data [94]. In order to trust the R2 data in a likelihood analysis, a similar BDT cut will likely need to be developed.

4.3.2. Analysis Improvements

CDMSlite R3 was the first CDMSlite analysis to be performed on blinded data. Previous SuperCDMS analyses with the iZIPs used a "black box" blinding scheme that removed all events within a certain range of recoil energy and yield where the dark matter signal was expected to be. However, as mentioned in Section 3.1, there was no yield quantity in CDMSlite to hide the nuclear recoil events so a new scheme was needed. It was determined

that adding an unknown number of fake signal events to the data was the best approach, which is a process called "salting." The entire analysis region, $0 - 2$ keV for R3, needed to be salted but salt was not added above 5 keV. This allowed for a background-only fit to the spectrum above 5 keV to determine the rates for each background. Those rates were then extrapolated to predict the number of each background source in the analysis region, which was necessary prior to unblinding the data.

CDMSlite R3 was also the first CDMSlite analysis to use a profile likelihood instead of the OI. In fact, most of the pieces described in Section 4.2 were originally developed by the R3 analysis team. In addition to the background models, the likelihood function also included nuisance parameters for the efficiency, resolution, and Lindhard uncertainty [5, 94].

4.3.3. Published Limit

There was no excess of events observed above the expected background, so the profile likelihood was used to calculate an exclusion limit [94]. The published limit from the R3 WIMP search is shown in Fig. 4.16, and the accompanying paper contains a detailed description of the complete analysis [54].

Most of the models and other likelihood components were developed as part of the R3 WIMP search analysis, but the inelastic dark matter analysis in this thesis used a different energy range so all of the models had to be reevaluated. The rest of this section will mirror the corresponding subsections in Section 4.2 to obtain the likelihood function components based on the R3 data.

4.3.4. Background Modeling

All of the R2 background models described in Section 4.2.5 were also included in the R3 analysis. Those backgrounds are neutron captures from the ^{252}Cf calibration source that
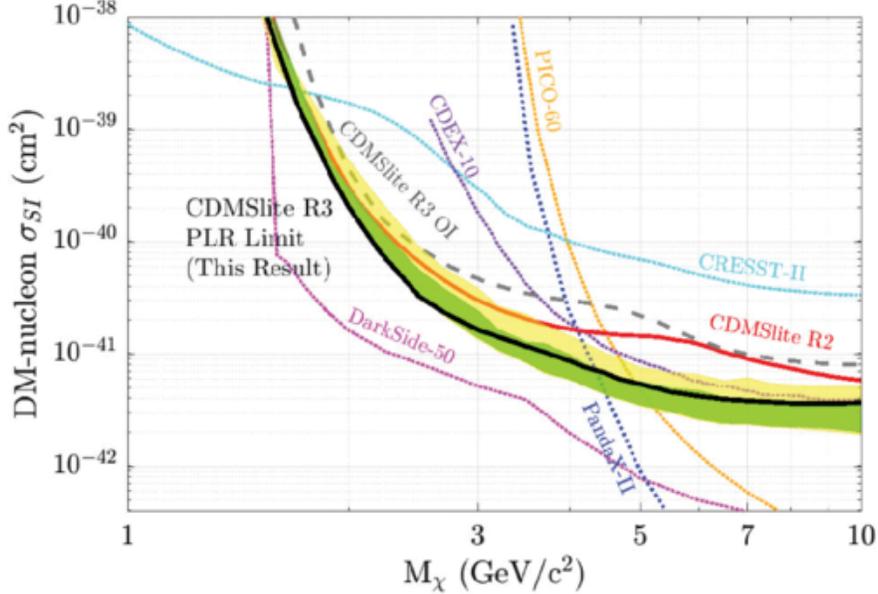
Figure 4.16: The black curve is the CDMSlite R3 WIMP exclusion limit, with 1σ and 2σ uncertainty bands shown in green and yellow, respectively [54]. The solid red line is the previous CDMSlite result, which had been surpassed by the R3 limit and the Darkside-50 limit in magenta [114]. Although R3 is about half the exposure of R2, the profile likelihood method sets a stronger limit than the OI method because it uses information about the backgrounds such as shape and normalization constraints.

activate the germanium crystal, cosmogenic activation of the crystal, Compton scattering from high energy photons produced by radioactive decays, and ^{210}Pb contamination on or near the detectors [54, 107].

4.3.4.1. Surface Background Models

This section follows exactly the same procedure as detailed in Section 4.2.5.1 for obtaining the surface background models. The only difference is that T2Z1 was the top detector in the tower, so there was a significant background contribution from the top lid of the copper housing that had a direct line of sight to the top face of the detector. In fact, this was the dominant source of surface background events.

First, the binned models from the surface background simulation were smoothed with an adaptive bandwidth KDE and are shown in Fig. 4.17, Fig. 4.18, and Fig. 4.19.
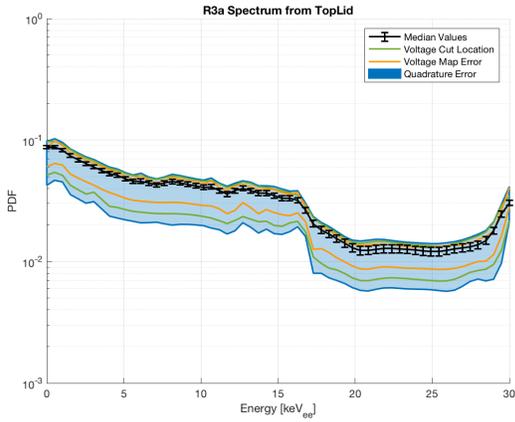
(a) The smoothed surface backgrounds from the copper top lid for R3a.
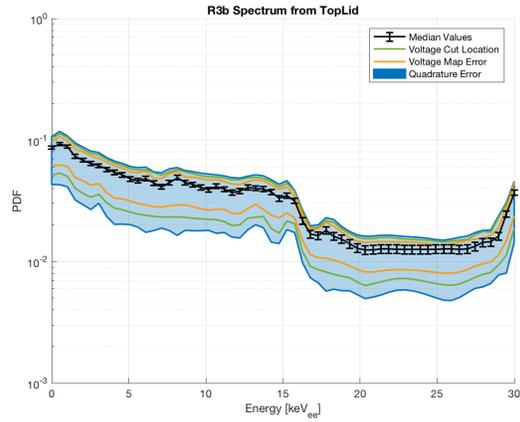
(b) The smoothed surface backgrounds from the copper top lid for R3b.

Figure 4.17: The orange curves represent the systematic error on the voltage map used in the simulation and the green lines represent the systematic error of the voltage cut location. The black median curve has been scaled to a pdf with unit area and is shown with statistical error bars. The green and orange lines have been scaled to maintain the ratio with the black curve, and the blue curves are the quadrature sum of the statistical and systematic errors.
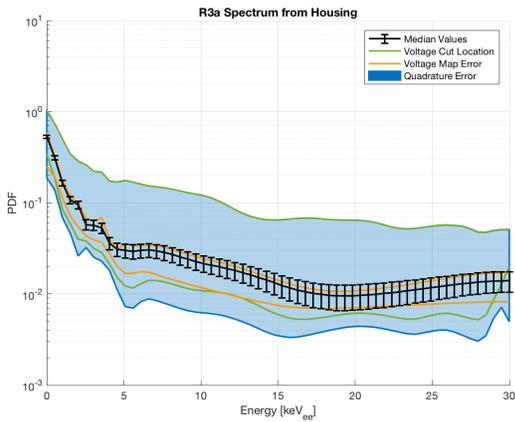
(a) The smoothed surface backgrounds from the copper sidewall housing for R3a.
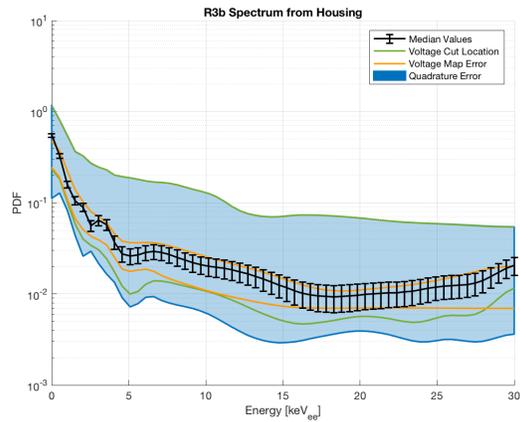
(b) The smoothed surface backgrounds from the copper sidewall housing for R3b.

Figure 4.18: The orange curves represent the systematic error on the voltage map used in the simulation and the green lines represent the systematic error of the voltage cut location. The black median curve has been scaled to a pdf with unit area and is shown with statistical error bars. The green and orange lines have been scaled to maintain the ratio with the black curve, and the blue curves are the quadrature sum of the statistical and systematic errors.
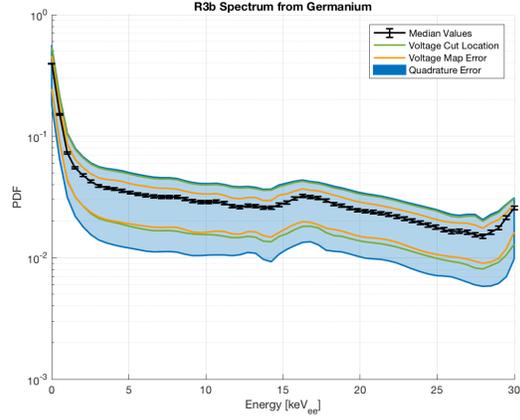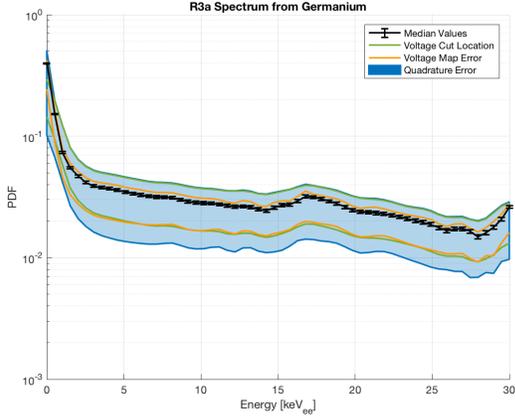
(a) The smoothed surface backgrounds from the germanium crystal for R3a.

(b) The smoothed surface backgrounds from the germanium crystal for R3b.

Figure 4.19: The orange curves represent the systematic error on the voltage map used in the simulation and the green lines represent the systematic error of the voltage cut location. The black median curve has been scaled to a pdf with unit area and is shown with statistical error bars. The green and orange lines have been scaled to maintain the ratio with the black curve, and the blue curves are the quadrature sum of the statistical and systematic errors.

For the $0 - 25$ keV energy range of this analysis, the normalization from the independent measurement of alpha rates gives $N_{TL} = 93.6 \pm 9.8$ events from the top lid (TL), $N_{SH} = 13.2 \pm 0.8$ events from the sidewall housing (SH), and $N_{GZ} = 13.9 \pm 3.5$ events from the germanium iZIP (GZ), with correlation coefficients of $\rho_{TL,SH} = -0.225$, $\rho_{SH,GZ} = -0.3395$, and $\rho_{GZ,TL} = -0.18$. Correcting those numbers for the efficiency of the cuts used in the simulation gives predicted counts of $N_{TL} = 158.4 \pm 16.6$, $N_{SH} = 22.3 \pm 1.4$, and $N_{GZ} = 23.5 \pm 5.9$ events. The uncertainty on these normalizations can be multiplied through the correlation matrix between the sources to give the covariance matrix:

$$
\begin{bmatrix}
274.96 & -5.05 & -17.68 \\
-5.05 & 1.83 & -2.72 \\
-17.68 & -2.72 & 35.07
\end{bmatrix}
\tag{4.11}
$$

4.3.4.2. Morphing the Surface Backgrounds

Following the procedure laid out in Section 4.2.5.2, a single morphing parameter was defined for each of the three surface backgrounds to encode the shape and normalization uncertainty. Figure 4.20, Fig. 4.21, and Fig. 4.22 show the resulting curves from sampling each morphing parameter linearly between $m = -1$ and $m = 1$.
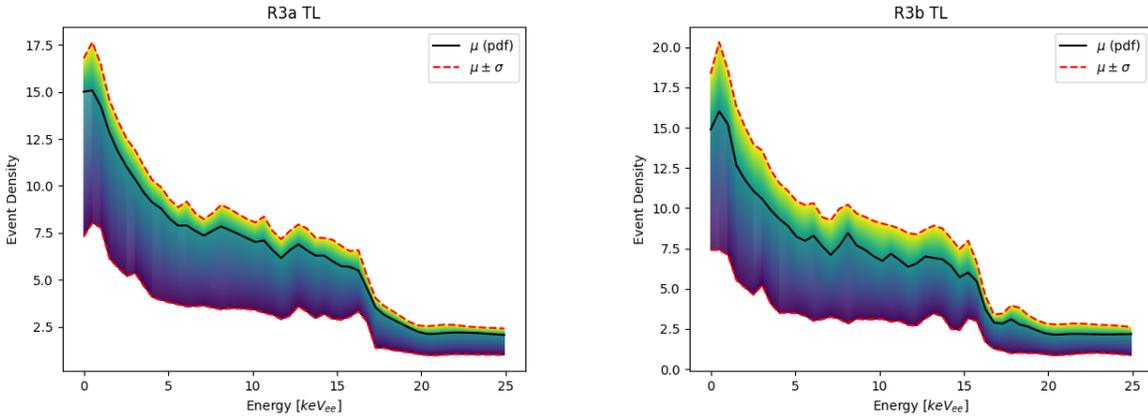
(a) Morphing the event density model of surface events from the copper top lid for R3a.

(b) Morphing the event density model of surface events from the copper top lid for R3b.

Figure 4.20: Linear sampling of the morphing parameter. There are 500 sampled lines on each of these plots with purple curves corresponding to $m = -1$ and yellow curves corresponding to $m = 1$.

The linear sampling demonstrates the behavior of the morphing parameters, but in the likelihood function they were constrained with a Gaussian distribution with mean $\mu = 0$ and width $\sigma = 1$. This forced the likelihood to preferentially fit the median curve or pay a statistical penalty to push the morphing parameter away from the median. In order to relate this constraint to the normalization, a one-dimensional Gaussian was randomly sampled. For each sampling, the model was morphed according to the morphing parameter in Eq. 4.3 and the number of events was calculated using Eq. 4.4. Histograms of the number of events from each morphed model are shown in Fig. 4.23.

Randomly sampling the Gaussian distribution provides an estimate of the uncertainty on the number of events for each surface background source. The covariance matrix due to
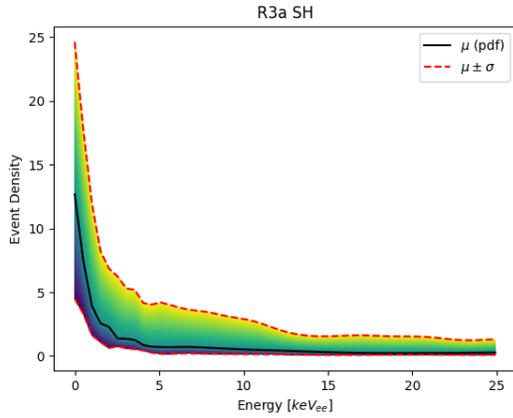
(a) Morphing the event density model of surface events from the sidewall housing for R3a.
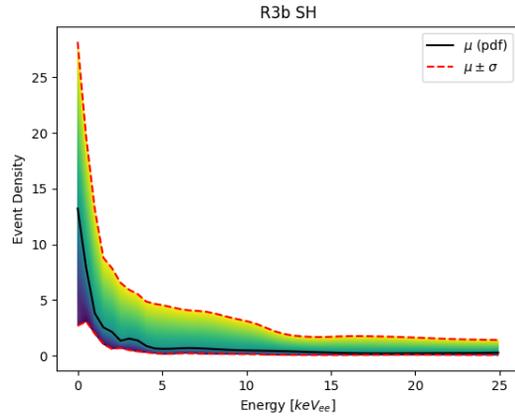
(b) Morphing the event density model of surface events from the sidewall housing for R3b.

Figure 4.21: Linear sampling of the morphing parameter. There are 500 sampled lines on each of these plots with purple curves corresponding to $m = -1$ and yellow curves corresponding to $m = 1$.
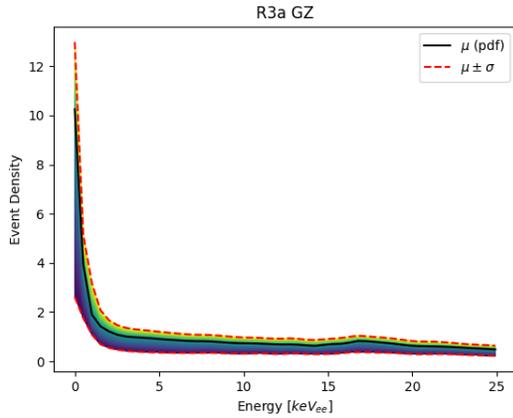
(a) Morphing the event density model of surface events from the germanium crystal for R3a.
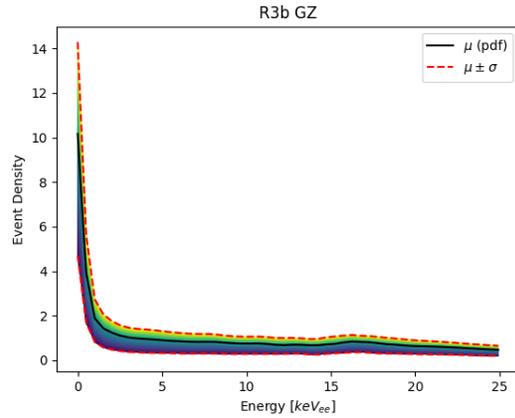
(b) Morphing the event density model of surface events from the germanium crystal for R3b.

Figure 4.22: Linear sampling of the morphing parameter. There are 500 sampled lines on each of these plots with purple curves corresponding to $m = -1$ and yellow curves corresponding to $m = 1$.
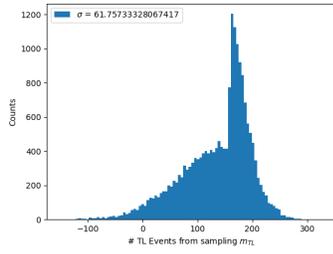
(a) Histograms of the number of events from Gaussian sampling of the top lid morphing parameter. The sampled curves in Fig. 4.20 are much closer together above the median.
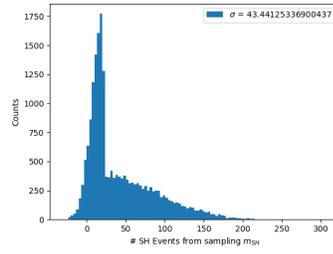
(b) Histograms of the number of events from Gaussian sampling of the sidewall housing morphing parameter. The sampled curves in Fig. 4.21 are much closer together below the median.
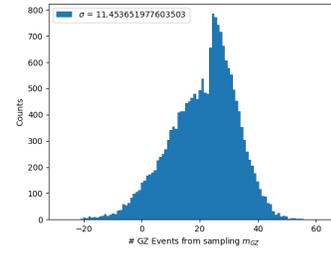
(c) Histograms of the number of events from Gaussian sampling of the germanium morphing parameter. The median curve in Fig. 4.22 is close to the center of the uncertainty band, but the sample curves are slightly closer together above the median.

Figure 4.23: The Gaussian sampling for each of the surface background sources produces an asymmetric distribution because the median curve is not equidistant from the upper and lower uncertainty.

the morphing parameter uncertainty is:

$$
\begin{bmatrix}
3402.79 & 2044.55 & 621.58 \\
2044.55 & 1870.56 & 413.80 \\
621.58 & 413.80 & 116.35
\end{bmatrix}
\tag{4.12}
$$

Adding the matrix in Eq. 4.11 to the matrix in Eq. 4.12 gives the total covariance on the number of surface background events for each source:

$$
\mathbf{V} =
\begin{bmatrix}
3677.75 & 2039.50 & 603.91 \\
2039.50 & 1872.39 & 411.08 \\
603.91 & 411.08 & 151.42
\end{bmatrix} .
\tag{4.13}
$$

The final step was to translate the correlated uncertainty from the number of surface background events to uncertainty on the morphing parameters so that it can be added to the likelihood function as a multivariate Gaussian constraint. This was done by plotting the relationship between morphing parameter and number of events as defined in Eq. 4.4. Figure 4.24 shows this relationship using the data from the linearly sampled morphing parameter in Fig. 4.20, Fig. 4.21, and Fig. 4.22.
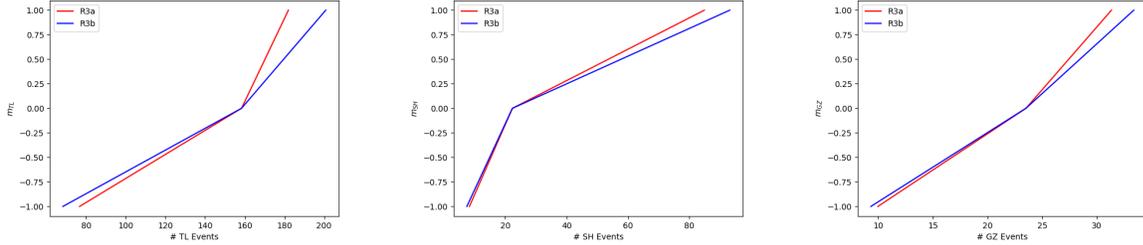
(a) The sharp angle at $m_{TL} = 0$ is related to the asymmetry of the histogram in Fig. 4.23a. This is showing that changing the morphing parameter from 0 to 1 does not result in a large change of number of events.

(b) The sharp angle at $m_{SH} = 0$ is related to the asymmetry of the histogram in Fig. 4.23b. This is showing that changing the morphing parameter from -1 to 0 does not result in a large change of number of events.

(c) The sharp angle at $m_{GZ} = 0$ is related to the asymmetry of the histogram in Fig. 4.23c. This is showing that changing the morphing parameter from 0 to 1 does not result in a large change of number of events.

Figure 4.24: The relationship between morphing parameter and resulting number of events is shown for each surface background source. The red line is for R3a and the blue line is for R3b. The independent variable is plotted on the y-axis for these plots to show to the slopes needed in Eq. 4.7.

The covariance on the morphing parameters can be calculated according to Eq. 4.7 where $m_1 = m_{TL}$, $m_2 = m_{SH}$, $m_3 = m_{GZ}$, $N_1 = N_{TL}$, $N_2 = N_{SH}$, and $N_3 = N_{GZ}$. The resulting matrix shown in Eq. 4.14 is the covariance matrix between the R3 morphing parameters that went into the Gaussian constraint in the likelihood to constrain the surface backgrounds:

$$\mathbf{M} = \begin{bmatrix} 1.08 & 0.86 & 0.92 \\ 0.86 & 1.15 & 0.90 \\ 0.92 & 0.90 & 1.19 \end{bmatrix} . \tag{4.14}$$

A demonstration of the Gaussian constraint is shown in Fig. 4.25, Fig. 4.26, and Fig. 4.27, where the blue lines are drawn by randomly sampling the covariance matrix in Eq. 4.14.
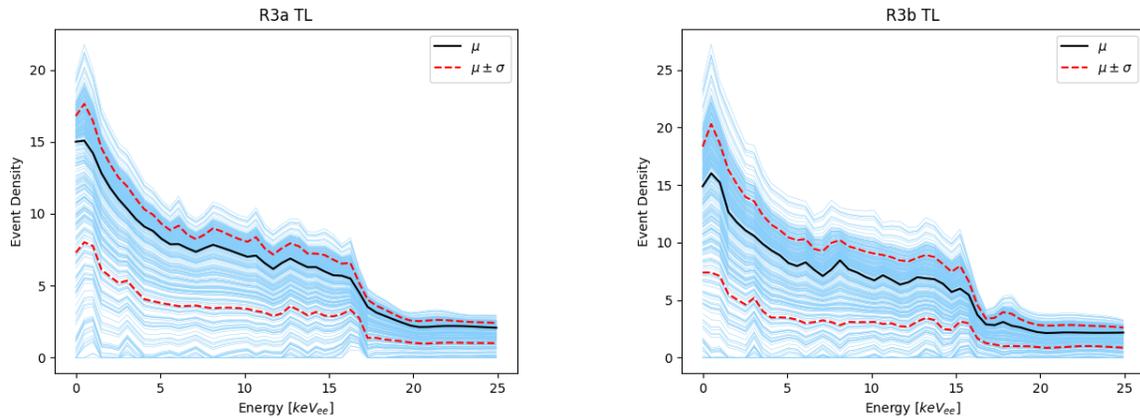
(a) Randomly sampled surface background model for the housing top lid in R3a.

(b) Randomly sampled surface background model for the housing top lid in R3b.

Figure 4.25: Demonstration of Gaussian constraint for morphing the R3 surface background housing top lid models. The black curve is the mean model, and the red curves are the 1σ uncertainty band. The blue curves are created by randomly sampling the covariance matrix in the multivariate Gaussian constraint.

4.3.5. Efficiency Model

The efficiency curves for R3a and R3b are similar to the R2 curve described in Section 4.2.6, but there are a few differences. Below 2 keV, the R3 efficiencies have sizable fluctuations from one of the cuts being defined differently. There is also no "double-hump" feature present because the R3a and R3b efficiencies were not added together. Finally, there was no evidence of high energy saturation in R3, so no linear correction above 10.37 keV was needed. For these reasons, R3a and R3b efficiencies could each be described by 5-parameter piecewise functions, shown in Eq. 4.15 with the free parameters highlighted in orange. The
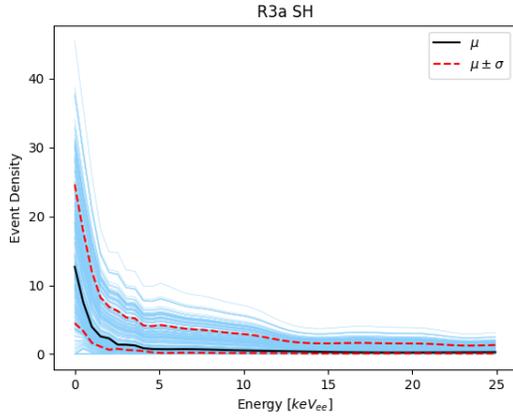
(a) Randomly sampled surface background model for the sidewall housing in R3a.
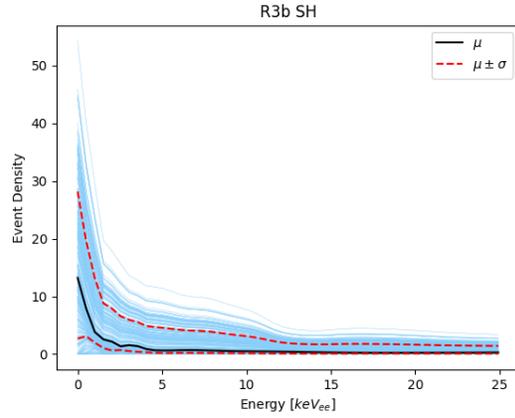
(b) Randomly sampled surface background model for the sidewall housing in R3b.

Figure 4.26: Demonstration of Gaussian constraint for morphing the R3 sidewall housing surface background models. The black curve is the mean model, and the red curves are the 1σ uncertainty band. The blue curves are created by randomly sampling the covariance matrix in the multivariate Gaussian constraint.
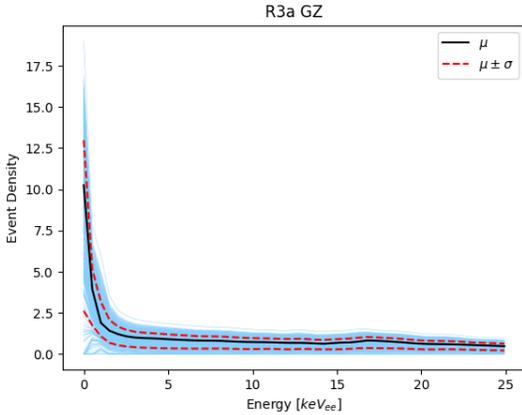
(a) Randomly sampled surface background model for the germanium iZIP in R3a.
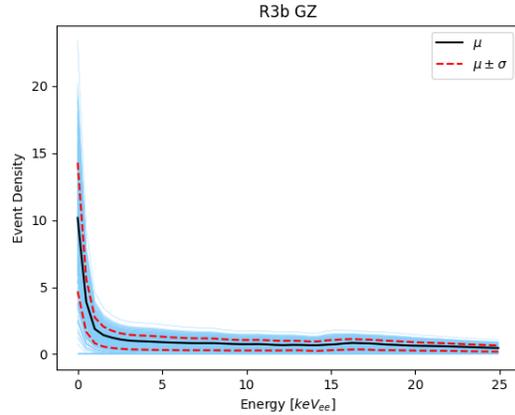
(b) Randomly sampled surface background model for the germanium iZIP in R3b.

Figure 4.27: Demonstration of Gaussian constraint for morphing the R3 germanium iZIP surface background models. The black curve is the mean model, and the red curves are the 1σ uncertainty band. The blue curves are created by randomly sampling the covariance matrix in the multivariate Gaussian constraint.

piecewise function was defined as

$$Eff_{R3}(E) = \begin{cases} f_0 = Erf(E, \mu, \sigma, A_1) & E < E_0 \\ f_1 = m_1(E - E_0) + f_0(E_0) & E_0 \le E < E_1 \\ f_2 = m_2(E - E_1) + f_1(E_1) & E_1 \le E \end{cases} \cdot \qquad (4.15)$$

For the same reasons laid out in Section 4.2.6, it was determined that the morphing procedure was preferable to the piecewise function, so the efficiency of each sub-run was controlled by a single morphing parameter. A demonstration of those parameters are shown in Fig. 4.28 by linearly sampling the morphing parameters between $m = -1$ and $m = 1$.
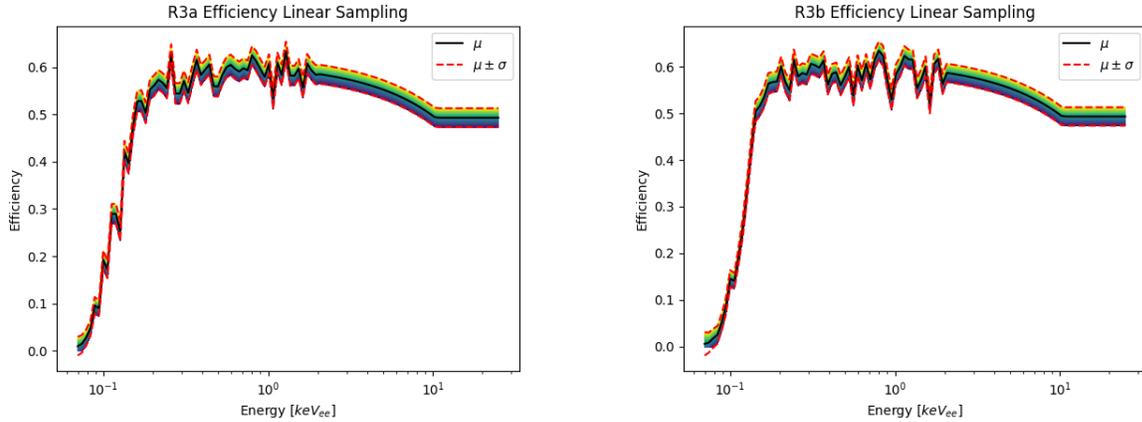
(a) Demonstration of morphing parameter for R3a efficiency.

(b) Demonstration of morphing parameter for R3b efficiency.

Figure 4.28: Linearly sampled morphing parameters. The purple lines correspond to $m = -1$, the yellow lines correspond to $m = 1$, and $m = 0$ returns the median black curve.

Since there was just one morphing parameter for each sub-run, a one-dimensional Gaussian distribution was used to constrain the efficiency models in the likelihood. This creates a penalty for moving further away from the median, yet the parameter was not constrained to $|m| \le 1$. Figure 4.29 shows the distribution from sampling this constraint many times.
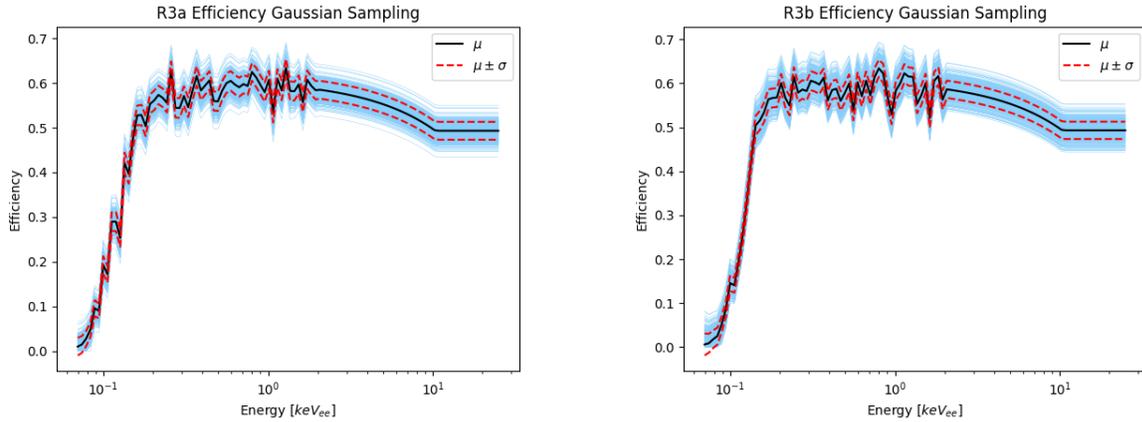
(a) Randomly sampled efficiency model for R3a.

(b) Randomly sampled efficiency model for R3b.

Figure 4.29: Demonstrating the Gaussian constraint for the morphed R3 efficiency models.

The fluctuations in the efficiency change the shape of the models when they are fit to the data, but the changes are small enough that they do not cause any problems in the likelihood convergence.

4.3.6. Resolution Model

The resolution model for R3 was determined using the same method as defined in Section 4.2.7 for R2. In order to fit the resolution function in Eq. 3.1, the baseline noise at 0 ke$V$ and the germanium activation peaks at 10.37 keV, 1.3 keV, and 160 eV were measured. The resulting model fits are shown in Fig. 4.30 and the results are summarized in Table 4.2 [54].

From these results, it is straightforward to write down the correlation matrix and multiply each row and column by the uncertainties of each of the model parameters. The resulting covariance matrices for R3a and R3b are shown in Eq. 4.16 and Eq. 4.17, respectively, and
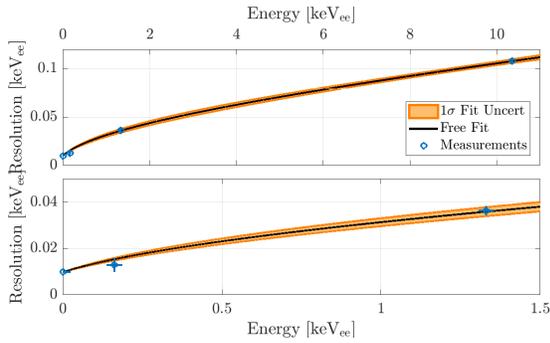
(a) Resolution function fit to measured features in R3a.
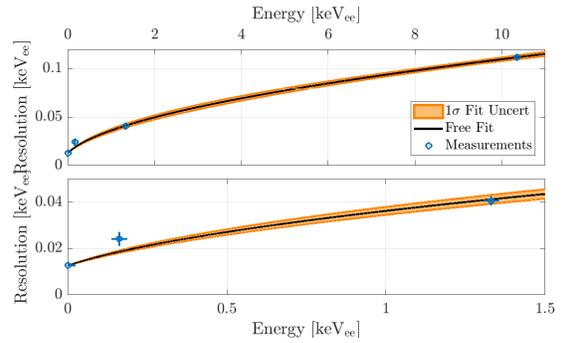
(b) Resolution function fit to measured features in R3b.

Figure 4.30: The blue points are the measured resolutions of the baseline noise and germanium activation peaks. The black curve is the best fit of the resolution model, and the orange band is the uncertainty of that fit. The top plot shows the measured peaks up to 11 keV, while the lower plot is zoomed in below 1.5 keV.

Peak	Energy [keV$_{ee}$]	Resolution [eV$_{ee}$]	Res. Err. [eV$_{ee}$]
R3a Baseline	0	9.87	0.04
R3b Baseline	0	12.67	0.04
M-shell	0.16	13.9	2
L-shell	1.3	36.3	2
K-shell	10.37	108	2

Run	A [eV]	B [eV]	C
R3a	9.87 ± 0.04	0.87 ± 0.12	$(4.94 \pm 1.27) \times 10^{-3}$
R3b	12.7 ± 0.04	0.80 ± 0.12	$(5.49 \pm 1.15) \times 10^{-3}$

Run	ρ_{AB}	ρ_{BC}	ρ_{CA}
R3a	-0.0071	-0.9474	0.0060
R3b	-0.0087	-0.9474	0.0074

Table 4.2: Table showing the measured resolutions of known features in the data (top), resolution model parameters from fitting to the measured widths (middle), and correlations between the model parameters (bottom).

was used to constrain the resolution model in the likelihood function:

$$
\mathbf{R} = \begin{bmatrix}
1.60 \times 10^{-9} & -3.428 \times 10^{-11} & 3.066 \times 10^{-10} \\
-3.428 \times 10^{-11} & 1.475 \times 10^{-8} & -1.467 \times 10^{-7} \\
3.066 \times 10^{-10} & -1.467 \times 10^{-7} & 1.625 \times 10^{-6}
\end{bmatrix} . \tag{4.16}
$$

$$
\mathbf{R} = \begin{bmatrix}
1.60 \times 10^{-9} & -4.233 \times 10^{-11} & 3.393 \times 10^{-10} \\
-4.233 \times 10^{-11} & 1.473 \times 10^{-8} & -1.319 \times 10^{-7} \\
3.393 \times 10^{-10} & -1.319 \times 10^{-7} & 1.316 \times 10^{-6}
\end{bmatrix} . \tag{4.17}
$$

A demonstration of the Gaussian constraint is shown in Fig. 4.31, where the blue lines are drawn by randomly sampling the covariance matrices in Eq. 4.16 and Eq. 4.17.
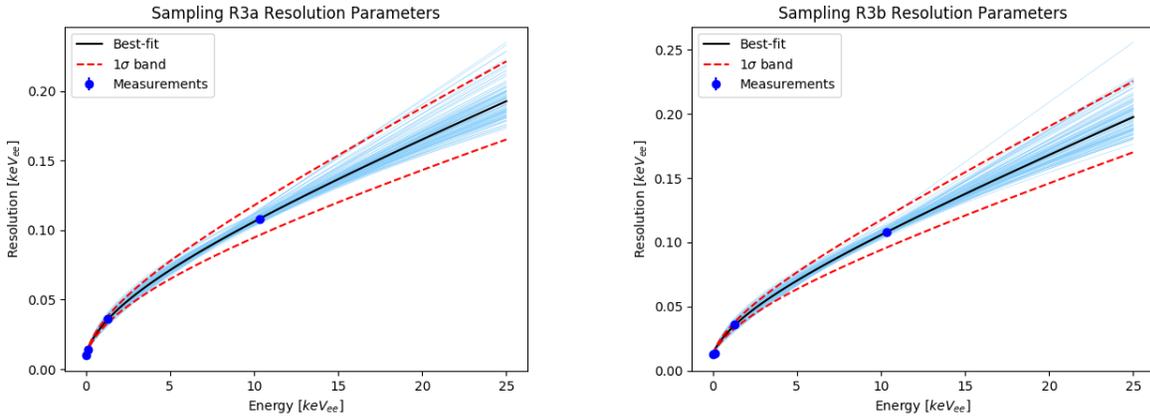
(a) Randomly sampled efficiency model for R3a.

(b) Randomly sampled efficiency model for R3b.

Figure 4.31: Demonstration of Gaussian constraints for the R3 resolution models. The black curve is the mean model, and the red curves are the 1σ uncertainty band. The blue curves are created by randomly sampling the covariance matrix in the multivariate Gaussian constraint. The blue curves account for correlations between the resolution parameters, which is why they do not fully sample the uncertainty band between the red curves.

4.4. Summary

At the time of publication, each of the three CDMSlite WIMP search analyses produced world-leading limits. This is a testament to the success of the SuperCDMS high voltage program, and provides high quality data that is prime for re-analyzing in the low-mass regime.

With all of the likelihood components constructed in this chapter, it is now possible to perform a likelihood analysis from $0-25$ keV using the CDMSlite R2 data, R3 data, or both. The last piece that is needed is the dark matter signal model of interest. In this thesis, there are two signal models discussed in the next chapter that are based on the inelastic scattering of low-mass dark matter.

CHAPTER 5

INELASTIC DARK MATTER

Historically, noble liquid experiments have had difficulty lowering their energy thresholds to be sensitive to dark matter masses below ~ 10 GeV; this range has been almost exclusively probed by solid state detectors. That is why the Large Underground Xenon (LUX) experiment caught the attention of the direct detection community in 2018 when they published limits showing sensitivity down to 0.4 GeV (Fig. 5.1) [115]. They did this by using a new interpretation of their data utilizing inelastic, rather than elastic, dark matter scattering hypotheses.
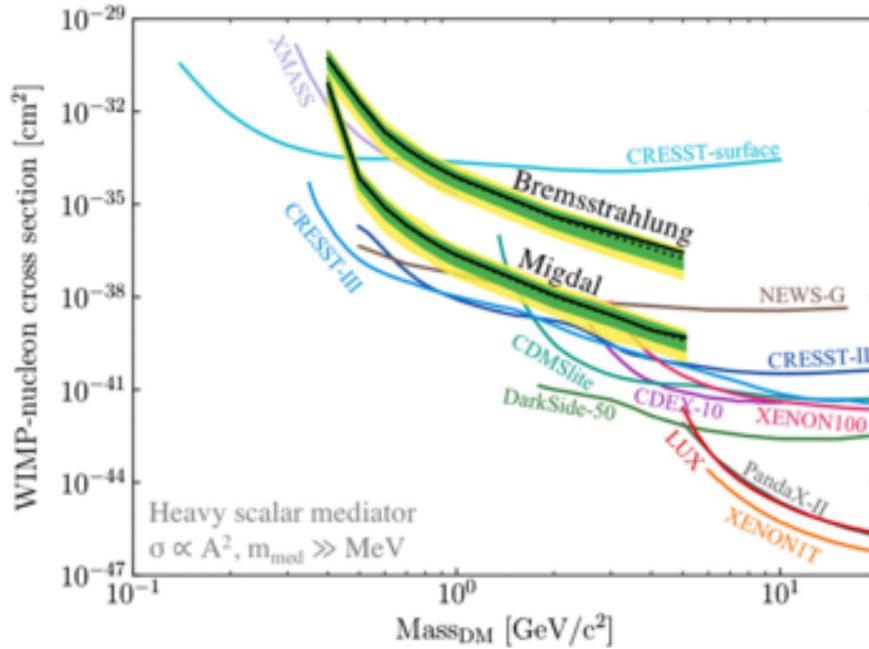
Figure 5.1: Inelastic dark matter upper limits published by LUX in 2018 [115]. The green and yellow uncertainty bands correspond to the 1σ and 2σ uncertainty from background-only trials and the black curves are the median limits. This result demonstrated shocking low-mass sensitivity for a noble liquid experiment.

The success of the LUX analysis shows the power of this approach to gain sensitivity to low-mass dark matter signals that would otherwise not be detectable. Since the CDMSlite experiment has already demonstrated a low energy threshold and sensitivity to 1 GeV dark matter mass, it is logical to believe searching for inelastically scattering dark matter in this data will permit exploration of sub-GeV parameter space. I have led this effort in the SuperCDMS collaboration.

This chapter will discuss the reinterpretation of CDMSlite data for two inelastic scattering scenarios: Bremsstrahlung and Migdal dark matter models. Both of these channels are predicated on a nuclear recoil event, thus the resulting limits can be directly compared to the traditional elastic limits.

5.1. Bremsstrahlung Radiation

Bremsstrahlung radiation is the process of emitting a photon from a charged particle as it is deflected by another charged particle. Since the target atoms in a germanium crystal are not isolated, a similar process could happen for the recoiling nucleus after a dark matter collision. The idea of searching for dark matter employing Bremsstrahlung radiation was proposed by Kouvaris and Pradler [116]. A diagram of photon emission from a nuclear recoil is shown in Fig. 5.2. Conceptually, it is easier to think of this as a two-step process but it should really be treated as a 2→3 body interaction as portrayed in Eq. 5.1,

$$\chi + n \to \chi + n + \gamma. \tag{5.1}$$
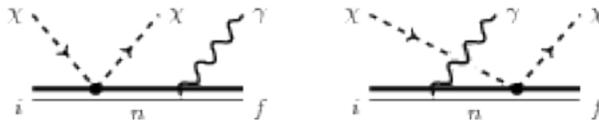
Figure 5.2: Diagrams of Bremsstrahlung radiation from a nuclear recoil [116]. The diagrams are interpreted from left to right, indicated by the initial i and final f states of the nucleus. Both diagrams result in $\chi + n + \gamma$ and are therefore indistinguishable.

Kouvaris and Pradler calculated the differential cross-section of this process as a function of emitted photon energy ω:

$$\frac{d\sigma}{d\omega} = \frac{4\alpha|f(\omega)|^2}{3\pi\omega}\frac{\mu_N^2 v^2 \sigma_0^{SI}}{m_N^2}\sqrt{1 - \frac{2\omega}{\mu_N v^2}}\left(1 - \frac{\omega}{\mu_N v^2}\right) , \qquad (5.2)$$

where α is the fine structure constant, f is the atomic scattering function, μ_N is the DM-nucleus reduced mass, v is the dark matter velocity, and m_N is the mass of the nucleus. Portions of Eq. 5.2 should look familiar from Section 1.4.3.1 where σ_0^{SI} is the spin-independent cross-section shown in Eq. 1.10. The terms on the end of Eq. 5.2 are kinematic terms for the emission of the photon.

The differential rate can be obtained by integrating the differential cross-section over velocity

$$\frac{dR}{d\omega} = \frac{N_a}{A}\frac{\rho_\chi}{m_\chi}\int d\phi \int d(cos\theta) \int v^3 \left(\frac{1}{\pi v_0^2}\right)^{\frac{3}{2}} e^{\left(\frac{v^2 + v_E^2 + 2vv_e cos\theta}{v_0}\right)^2}\frac{d\sigma}{d\omega}dv , \qquad (5.3)$$

where N_a is the number of target atoms, A is the atomic mass, ρ_χ is the local dark matter mass density, m_χ is the dark matter mass, v_0 is the most probable galactic WIMP velocity (220 km/s), and v_E is the mean orbital velocity of Earth. As in Section 1.4.3.1, the velocity distribution is assumed to be Maxwell-Boltzmann.

Figure 5.3 shows the differential rate plotted for target materials that are commonly used in direct detection experiments. For comparison, the traditional elastic nuclear WIMP models are shown as dashed lines. With the exception of Si, the inelastic Bremsstrahlung model extends to higher energy by more than an order of magnitude over the elastic scatter. However, this comes at the expense of a drastically reduced rate by approximately ten orders of magnitude because of the emission of a photon in addition to the nuclear recoil.
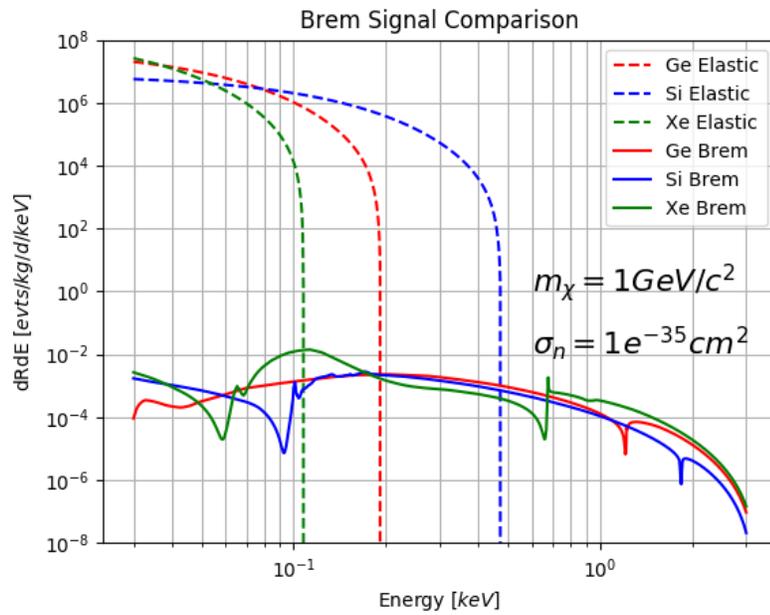
Figure 5.3: Comparison of elastic nuclear recoil (dashed lines) and the Bremsstrahlung model (solid lines) for silicon (blue), germanium (red), and xenon (green). The shapes of the Bremsstrahlung models are mostly driven by the atomic scattering function. It can be seen that the Bremsstrahlung signal extends to higher energies, but at the cost of a significantly decreased rate.

5.1.1. Photoelectric Cross-Section Uncertainty

As seen in Eq. 5.2, the differential cross-section depends on the atomic scattering function f, which is defined as

$$|f|^2 = |f_1 + if_2|^2 = f_1^2 + f_2^2 \,. \tag{5.4}$$

The real part of the atomic scattering function f_1 is given by the dispersion relation

$$f_1 = Z + \frac{1}{\pi r_e hc} \int_0^{\text{inf}} \frac{\epsilon^2 \sigma_e(\epsilon) d\epsilon}{\omega^2 - \epsilon^2} \tag{5.5}$$

and the imaginary part of the atomic scattering function f_2 is

$$f_2 = \frac{\sigma_e}{2 r_e \lambda} \,, \tag{5.6}$$

where Z is the atomic number, r_e is the electron radius, h is Planck's constant, c is the speed of light, ϵ is the dispersion parameter that gets integrated out, and σ_e is the photoelectric cross-section [117].

An exhaustive search of experimental measurements and theoretical calculations was performed to find the photoelectric cross-section at low energies [118]. This search provided a nominal, or "reference," photoelectric cross-section as well as upper and lower bounds. Each of these were used to calculate the atomic scattering factors in Fig. 5.4 and Fig. 5.5. These plots also include data provided by the Henke paper that is most commonly referenced in calculations involving the photoelectric cross-section [119].

Calculating the full atomic scattering function from f_1 and f_2 shows in Fig. 5.6 that there is no consistent trend with energy as far as which curve gives the most aggressive or conservative results. In order to see the full impact of the photoelectric cross-section uncertainty, an envelope was calculated around the atomic scattering functions by taking the minimum or maximum curve for every energy. Figure 5.7 shows the derivation of that
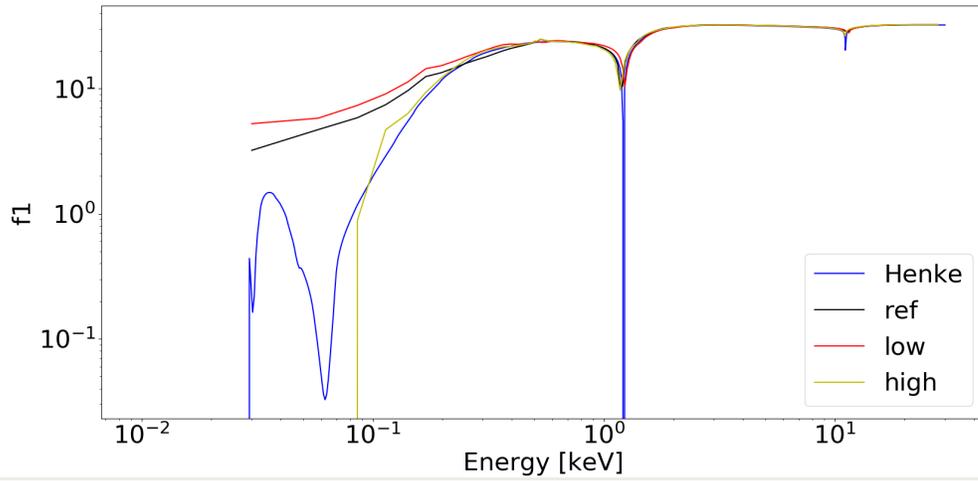
Figure 5.4: The real part of the atomic scattering function for given photoelectric cross-section models. Due to the dispersion relation, the lower bound on photoelectric cross-section leads to the upper bound on f_1.
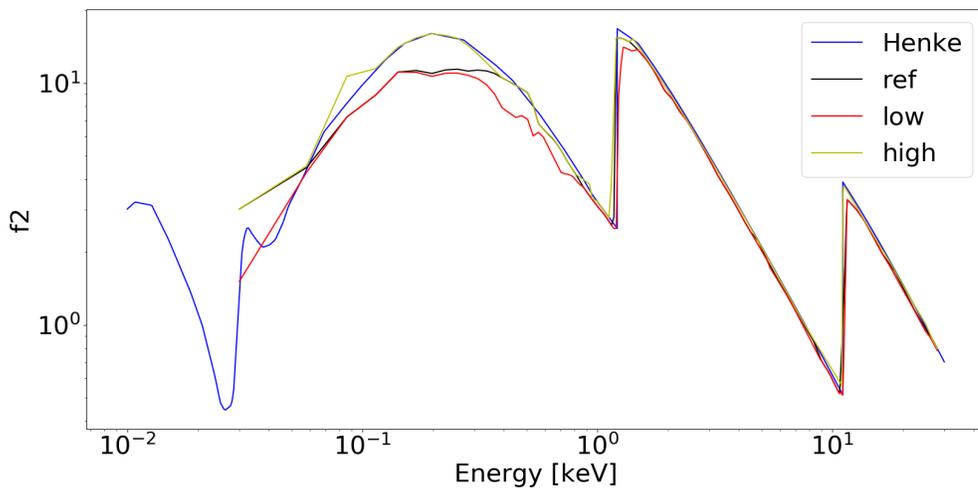
Figure 5.5: The imaginary part of the atomic scattering function for given photoelectric cross-section models. There is no inversion of upper and lower bounds because of the direct relationship between the photoelectric cross-section and f_2.

envelope, labeled as lowf and highf. Calculating the extreme bounds in this way means that the bottom (top) of the envelope will always give the most conservative (aggressive) result, regardless of the analysis or signal model.
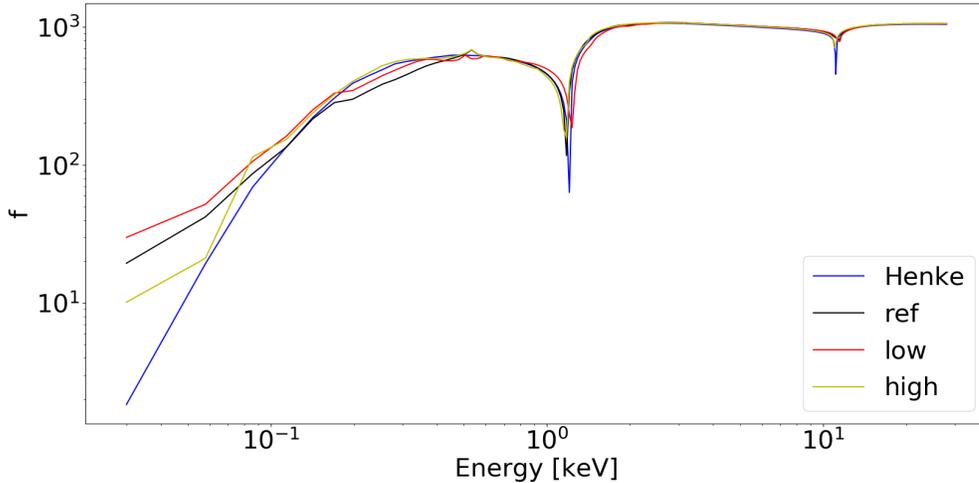
Figure 5.6: The atomic scattering function that goes into the calculation of the Bremsstrahlung signal model.

It was determined that the effect of the photoelectric cross-section uncertainty on the dark matter search results was $< 25\%$ between the most aggressive and conservative atomic scattering functions. On the scale of the limit plots, that is approximately the width of the line used to draw the limits and was therefore considered negligible. This analysis used the conservative bottom of the envelope to calculate the final results.

5.2. The Migdal Effect

The second inelastic dark matter signal considered is based on a theoretically motivated, yet experimentally unverified effect known as the Migdal Approximation [120]. The basis for this phenomenon is that the electron cloud does not immediately follow the nucleus when it is displaced. This is used to search for dark matter collisions that cause the nucleus to recoil and be displaced from the center of the electron cloud, which is equivalent to a simultaneous Lorentz boost for all of the orbital electrons. When the excited electron cloud relaxes, there
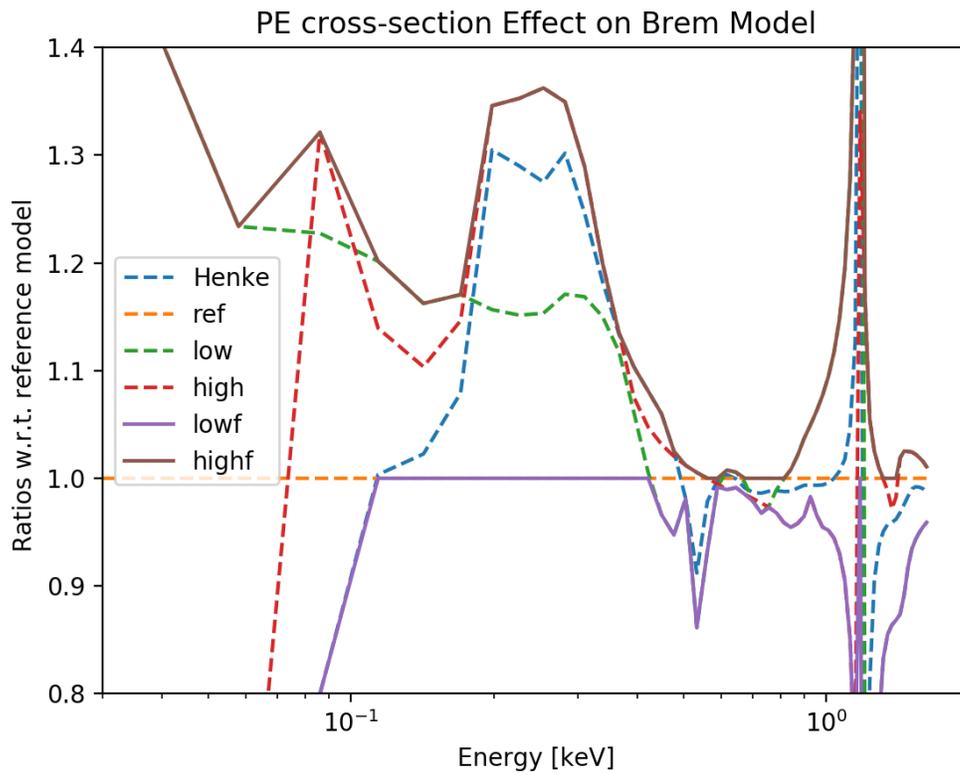
Figure 5.7: The dashed lines are the ratio of each atomic scattering function with respect to the reference model. An envelope was formed around these curves to calculate the most aggressive (highf) and conservative (lowf) final results, as indicated by the solid lines.

is a probability that one or more electrons are ejected. A diagram of this process is shown in Fig. 5.8. The idea to use this effect to search for dark matter was proposed by Ibe, Nakano, Shoji, and Suzuki who calculated the probability that a nuclear recoil would emit an electron [121]. With their paper they included tables of probabilities for a variety of common target materials.
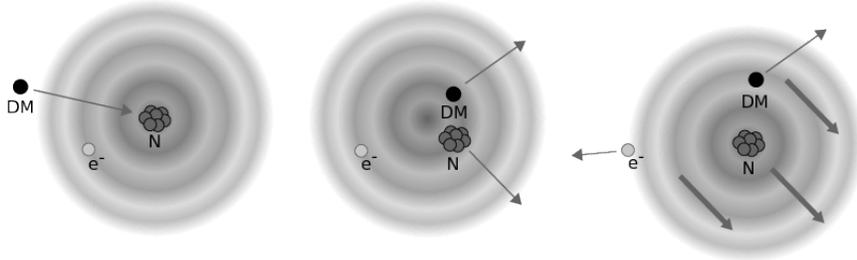
Figure 5.8: Diagram showing a low energy DM-nucleus interaction that pushes the nucleus out of the center of the atom [122]. As the electrostatics cause the electrons to "catch up" to the nucleus, there is a probability that one of them gets ejected from the atom.

Following the formalism of Ibe et al., Dolan, Kahlhoefer, and McCabe laid the groundwork for interpreting data from a dark matter experiment [122]. Through personal correspondence, it was revealed that Dolan et al. assumed isotropic scattering, which makes the integration simpler. The inelastic analysis described in this thesis subscribes to a more complex treatment of angular dependence that excludes non-physical velocity space [104]. Since parameter space is removed, the anisotropic approach leads to a lower expected rate and a more conservative result than the isotropic assumption. Therefore, the isotropic simplification was only used to check our work against others but was not used for the final calculation.

Since the Migdal effect is based on a nuclear recoil, the first step in calculating the signal model is to rewrite Eq. 1.6 as a differential rate over velocity [122],

$$\frac{d^2 R_{nr}}{dE_R dv} = \frac{\rho_\chi \sigma_N}{2\mu_N^2 m_\chi} \frac{f(v)}{v} \,. \tag{5.7}$$

Then the ionization differential rate in Eq. 5.8 can be written by combining the formalism of the Ibe et al. and Dolan et al. papers,

$$\frac{d^3 R_{ion}}{dE_R dE_e dv} = \frac{d^2 R_{nr}}{dE_R dv} \times \sum_{nl} \frac{1}{2\pi} \left(\frac{m_e \sqrt{2E_R/m_N}}{v_{ref} m_e} \right) \frac{dp^c_{v_{ref}}(nl \rightarrow E_e)}{dE_e} , \tag{5.8}$$

where n is the principal quantum number, l is the azimuthal quantum number, v_{ref} is a reference velocity that was used to calculate the transition probabilities (300 km/s), p^c_v is the probability of emitting an electron from state nl with energy E_e that is provided by Ibe et al.

The observed energy is more than just the energy of the emitted electron E_e because the binding energy E_{nl} is released too, so the energy observed in the detector is actually

$$E_{EM} = E_e + E_{nl} . \tag{5.9}$$

In order to obtain the differential rate as a function of the observable energy, the ionization differential rate needs to be integrated over nuclear recoil energy, velocity, and angular space. As with the WIMP and Bremsstrahlung models, a truncated Maxwell-Boltzmann velocity distribution is assumed with a normalization parameter k. The upper limit of the velocity integral is the same as before, the sum of v_{esc} and v_E. However, the lower bound of the velocity integral v_{min} needs to be modified to include the binding energy of the shell that emitted the electron [122],

$$v_{min} = \sqrt{\frac{m_N E_R}{2\mu^2}} + \frac{E_{EM}}{\sqrt{2m_N E_R}} . \tag{5.10}$$

Equation 5.11 shows the differential rate with all the pieces assembled

$$\frac{dR_{ion}}{dE_{EM}} = \frac{\rho_\chi \sigma_N}{2\mu_N^2 m_\chi k}\frac{1}{2\pi} \times \sum_{nl} \frac{dp_{v_{ref}}^c(nl \to E_{EM} - E_{nl})}{dE_{EM} - E_{nl}} \times$$

$$\int_0^{2\pi} d\phi \int_0^{E_{R,max}} dE_R \int_{v_{min}(E_R)}^{v_{max}} dv \int_{-1}^{cos\theta_{max}(v,E_R)} dcos\theta \sqrt{\frac{2E_R}{m_N v_{ref}^2}} v \times e^{\frac{-(v^2+2vv_E cos\theta + v_E^2)}{v_0^2}} .$$

(5.11)

Shown in Fig. 5.9 through Fig. 5.11 are the Migdal signal models showing the contribution of each electron shell.
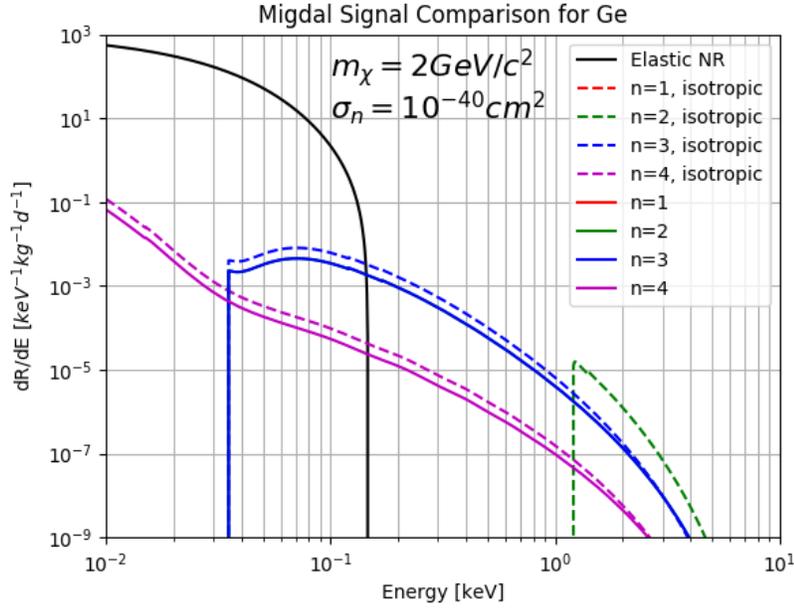
Figure 5.9: Plot of the contributions of each electron shell to the Migdal model for germanium. The dashed lines assume isotropic scattering and the solid lines are anisotropic. The solid black line shows the elastic nuclear recoil spectrum for comparing the rate and energy reach.

As with the Bremsstrahlung model, the Migdal model extends higher in energy than an elastic nuclear recoil from the same incoming particle. The extent of that reach is mostly driven by the inner electron shells that have a higher binding energy that gets released with the electron. However, the Migdal model does not exhibit as much of a decrease in rate as
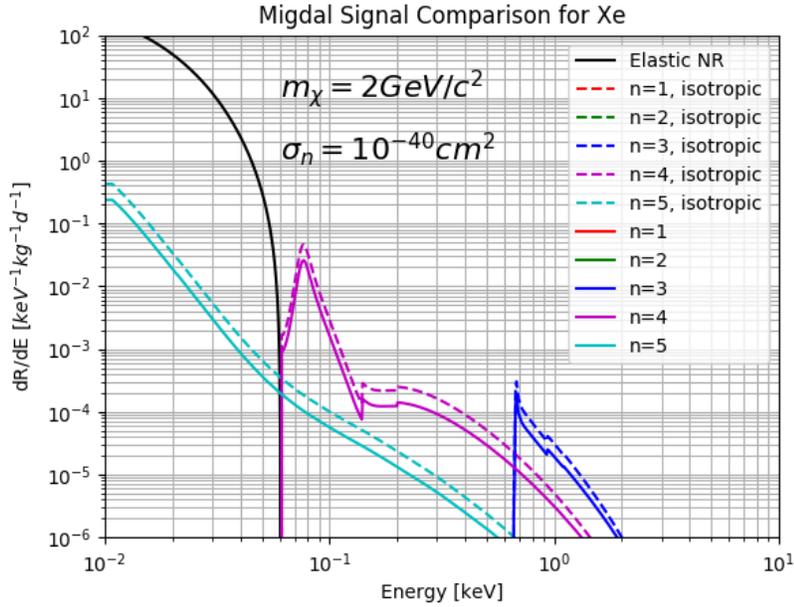
Figure 5.10: Plot of the contributions of each electron shell to the Migdal model for xenon. The dashed lines assume isotropic scattering and the solid lines are anisotropic. The solid black line shows the elastic nuclear recoil spectrum for comparing the rate and energy reach.
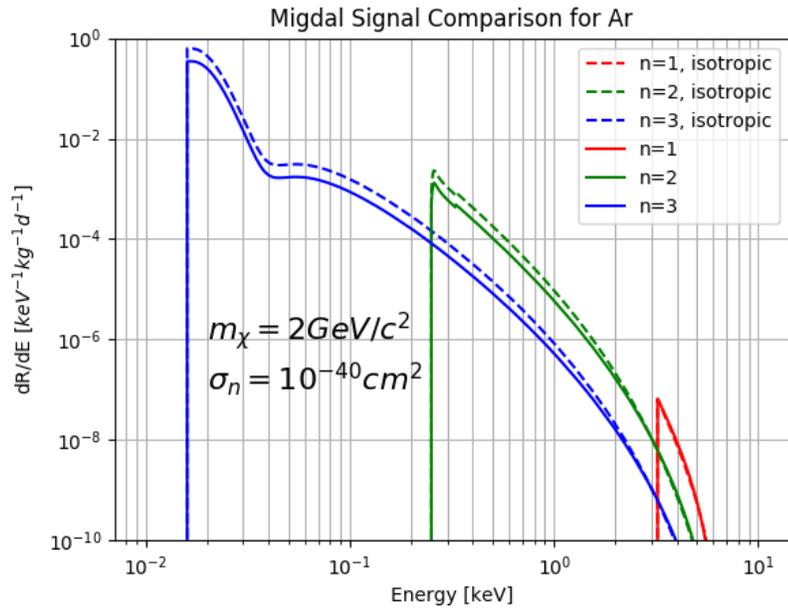
Figure 5.11: Plot of the contributions of each electron shell to the Migdal model for argon. The dashed lines assume isotropic scattering and the solid lines are anisotropic. There is no line showing the elastic nuclear recoil spectrum because it requires making assumptions about the detector response of an argon detector to convert to keV_{ee}.

the Bremsstrahlung model because the electrons already exist in the atom and there is no additional penalty of creating them.

5.2.1. The Isolated Atom Assumption

The biggest uncertainty in the Migdal signal model stems from an assumption made by Ibe et al. while calculating the quantum transition probabilities. In order to do the calculation, they assumed the atom is isolated and ignored potentials from surrounding atoms in the crystal lattice. Ibe et al. specifically address this in a figure caption in their paper, "Since we apply the estimations for the isolated atoms, the ionization spectrum from the valence electrons are not reliable" [121]. In this statement, Ibe et al. is alluding to the fact that the valence electrons in the outermost shell are most likely to feel the effects of surrounding atoms. This analysis adopts the same approach as other published Migdal analyses and excludes the valence shell the signal model is only summed over the remaining inner shells [115, 122, 123].

For the case of the germanium CDMSlite detector used in this analysis, excluding the valence shell has a negligible effect on the signal mode. As shown in Fig. 5.12, the energy cutoff created by excluding the valence shell where the next shell contribution falls sharply is around 35 eV. This is below the energy threshold of the detector, which is 70 eV, so excluding the valence shell does not impact the final results.

Figure 5.13 compiles the limits from Fig. 5.9 through Fig. 5.11 by summing over all non-valence shells.

5.3. Dark Matter Shielding

Given the success of direct detection experiments improving the limits of detector capabilities, DM-nucleon cross-sections on the order of 1×10^{-30} cm^2 are now considered "strong."
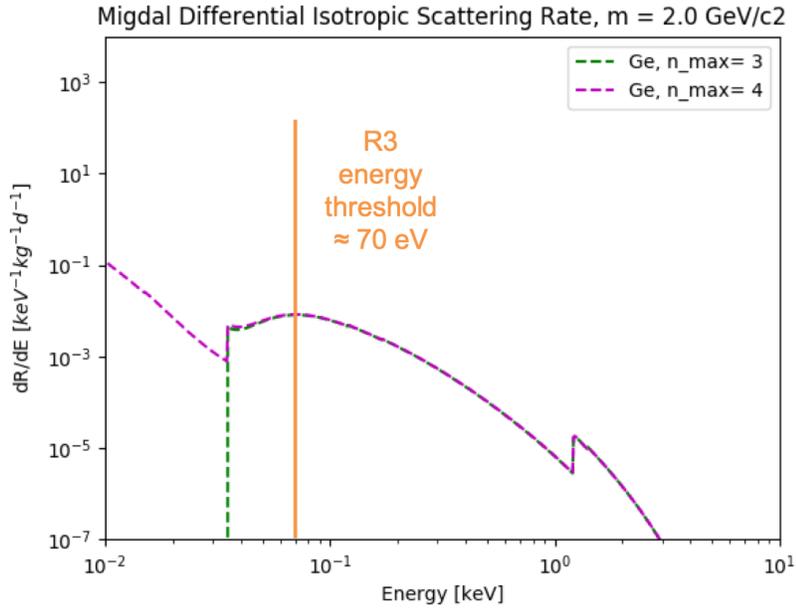
Figure 5.12: Plotting the effect of excluding the valence shell from the Migdal model for a germanium detector. The magenta curve includes the valence shell ($n = 4$ for germanium), and the green line excludes the valence shell from the model. The energy cutoff created by excluding the valence shell is around 35 eV, which is below the 70 eV energy threshold of the detector, so the final results are unaffected by this modification.
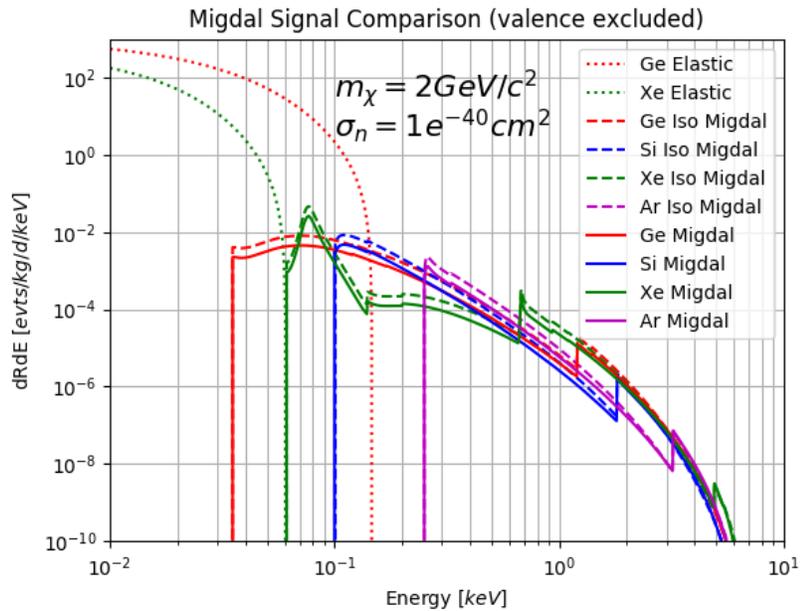
Figure 5.13: Plot comparing the Migdal signal model for common direct detection target materials. The dashed lines assume isotropic scattering, the solid lines are anisotropic, and the dotted lines are the elastic nuclear recoil spectra for germanium and xenon.

For dark matter with stronger cross-sections, there is an increased chance that it will scatter before it reaches the detector. This is especially true for experiments located deep underground to shield from cosmic rays, meaning there is a ceiling to the experimental sensitivity that has largely been neglected.

The concept of a ceiling was popularized in 2018 in a paper by Emken and Kouvaris where they highlight the shortcomings of existing analytic approaches to account for dark matter shielding and proposed a Monte Carlo-based solution, all of which are depicted in Fig. 5.14. The first analytic method calculates the critical cross-section where the experiment no longer has any sensitivity because the overburden completely stops the dark matter. This is the easiest method to calculate, but is not realistic because the dark matter could scatter and lose some energy, but still reach the detector. The second analytic method incorporates a damping factor into the velocity distribution to account for dark matter particles that scatter and lose some energy but are not completely stopped [124]. A modified version of the second analytic method is adopted in this analysis because the Monte Carlo method still included some outstanding issues and was more difficult to modify.
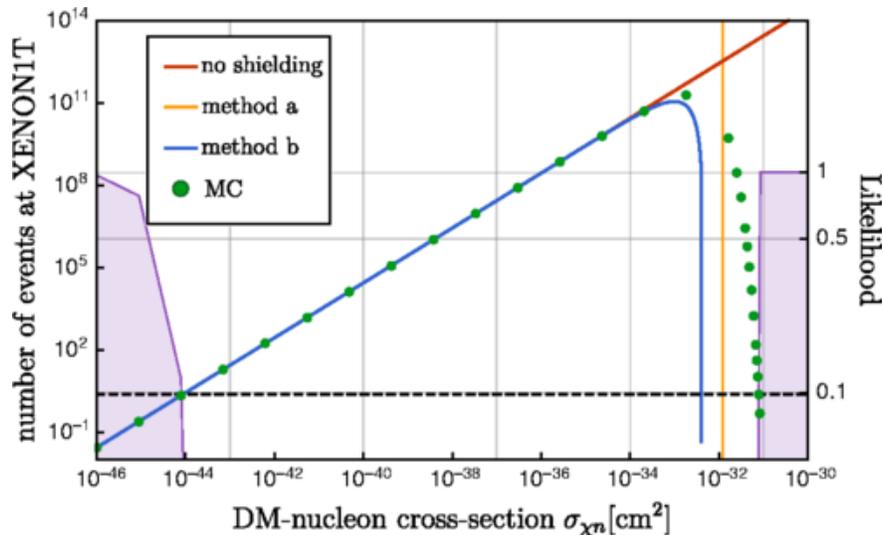
Figure 5.14: Comparison of three different approaches to account for dark matter shielding [124]. The red line has no shielding, the yellow line is the critical cross-section, the blue line uses a damped velocity distribution, and the green points are simulated from a Monte Carlo. All curves assume the dark matter particles are coming directly from above.

122

The first step was to calculate the damping factor D. Equation 5.12 comes from solving Eq. 9 of Emken and Kouvaris and changing some of the variable names for clarity:

$$D = \frac{\sigma_n}{m_\chi \mu_n^2} \left(\sum_i \frac{f_i \mu_i^4 A_i^2}{m_i^2} \right) \rho d \,, \tag{5.12}$$

where σ_n, m_χ, and μ_n are the same variables from previous variations of differential rates. The new term is the summation that accounts for the fractional contribution of each element i composing the overburden where f_i is the mass fraction, μ_i is the reduced mass with the dark matter, A_i is the mass number, m_i is the mass, ρ is the mass density of the medium, and d is the depth.

The second step is to dampen the velocity distribution, according to Eq. 10 of Emken and Kouvaris [124],

$$f_D(v_\chi^{fin}) = \exp[2D] \, f(\exp[D] \, v_\chi^{fin}) \,. \tag{5.13}$$

Incorporating the damped velocity distribution gives the modified differential rate in Eq. 5.14 that was taken from Eq. 8 of Emken and Kouvaris [124],

$$\frac{dR_i}{dE_R} = n_T \frac{\rho_\chi}{m_\chi} \int\limits_{v > v_{min}(E_R)} dv \, v f_D(v) \frac{d\sigma_{\chi i}}{dE_R}(E_R, v) \,. \tag{5.14}$$

Accounting for the layers of the atmosphere, the composition of the rock in the Earth, and the shielding around the experiment results in the damped velocity distributions shown in Fig. 5.15 for different cross-section strengths. Increasing the interaction strength increases the amount of attenuation experienced by the incoming dark matter and pushes the distribution to lower velocities.

One of the shortcomings of the various damping methods that Emken and Kouvaris pointed out was they all assume the dark matter is approaching the experiment from above,
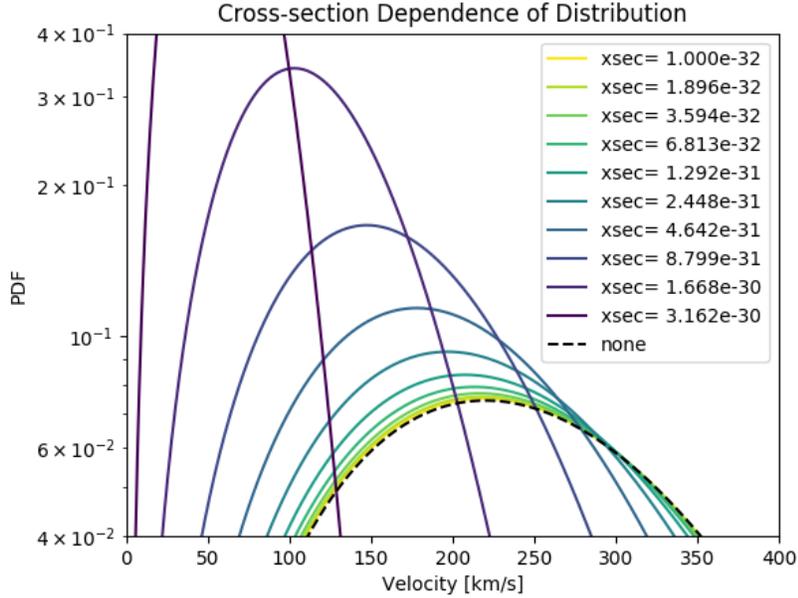
Figure 5.15: Damped velocity distributions for different dark matter cross-sections, with the undamped velocity distribution shown as a black dashed curve. The darker colors indicate stronger interactions, which push the distribution to lower velocities. All of these curves are calculated with a zenith angle of zero.

similar to Fig. 5.15. To first order, this is a reasonable assumption because it is the least amount of Earth shielding and the dark matter is more likely to scatter from any other angle. The assumption of unidirectional scattering simplifies the calculation by identifying the depth of each layer and summing each contribution. However, in the pursuit of a more realistic model, this analysis has included the angular dependence on the damping factor.

Adopting the coordinate system introduced by Kavanagh reveals the following angular dependence where $\theta = 0$ are events reaching the detector from the opposite side of the Earth and $\theta = \pi$ corresponds to an incident angle normal to the surface [125]. As expected, the velocity distribution is attenuated more strongly when the dark matter has to pass through more of the Earth.
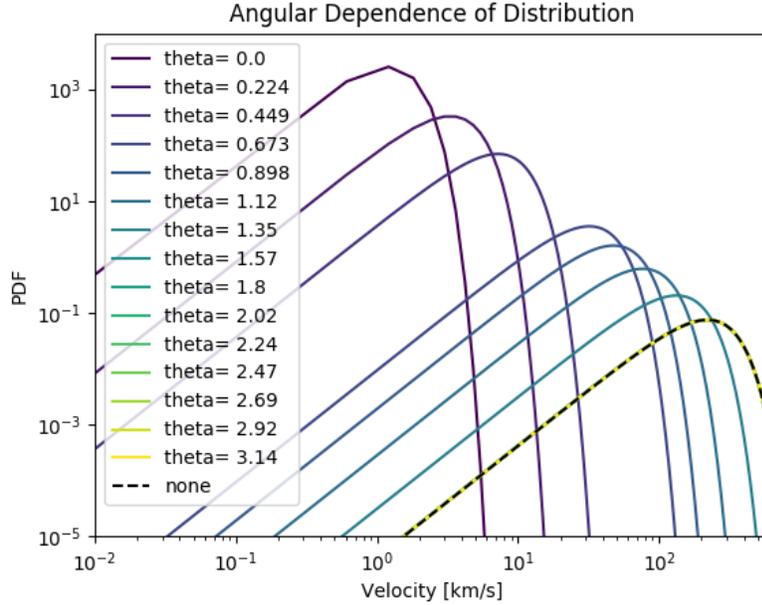
Figure 5.16: Damped velocity distribution for 10^{-34} cm^2 for different incident angles. Darker colors correspond to smaller angles ($\theta = 0$ is from directly below), and therefore experience more shielding than events from directly above ($\theta = \pi$).

Putting the angular dependence together with the spherically layered models of the atmosphere and Earth gives the final damped velocity distributions, shown in Fig. 5.17 for a 1 GeV dark matter particle.

Despite the effort to include realistic effects, there were still a few simplifying assumptions made in the analysis in this thesis. The experimental shielding, Earth, and atmosphere models were all assumed to be spherical. The relationship between the laboratory location and the direction of the Earth's trajectory through the dark matter halo was ignored. Finally, particles were assumed to have straight trajectories so any scattering was considered a loss of energy but not a change in direction.

5.4. Likelihood Function

Both the Bremsstrahlung and Migdal inelastic channels are limited by a reduced rate because of the 3-body final state. The most straightforward way to combat this effect is
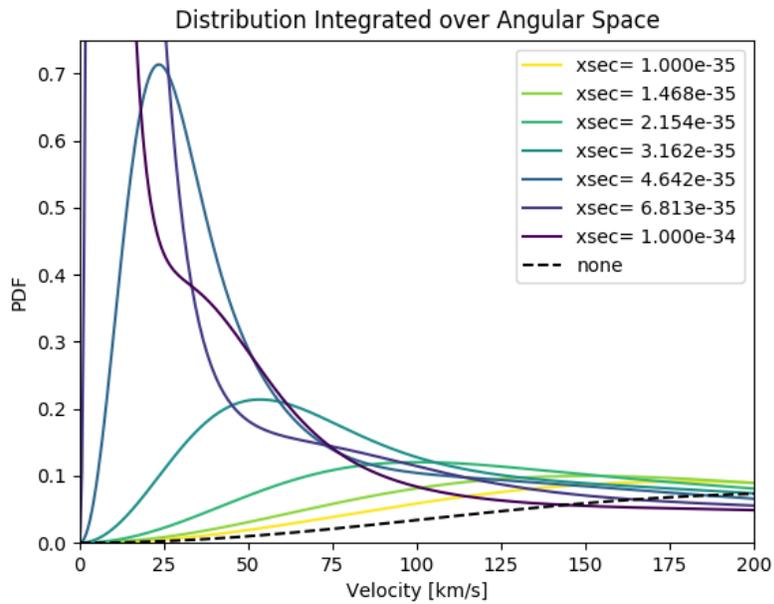
Figure 5.17: Lower half of some damped velocity distributions with angular dependence incorporated. The damping effect turns on very quickly, so smaller cross-sections will closely follow the black dashed undamped model. A bimodal distribution begins to form for stronger cross-sections (darker curves) where events from directly above are not attenuated as much as events originating from the other side of the Earth.

to increase the exposure of the experiment whether that means more mass, longer time operating, or both. Since SuperCDMS is not currently collecting new data, the goal of this analysis was to include both CDMSlite R2 and R3 data. However, as stated in Section 4.2.2, the R2 analysis was optimized for an optimum interval calculation. As such, there are potentially unmodeled backgrounds in the R2 data that need to be addressed before that data can be incorporated into a likelihood framework that requires all backgrounds in the data to be understood and modeled. For example, Fig. 5.18 shows that there is an excess of events in the lowest energy bins that do not fit the summed background model.
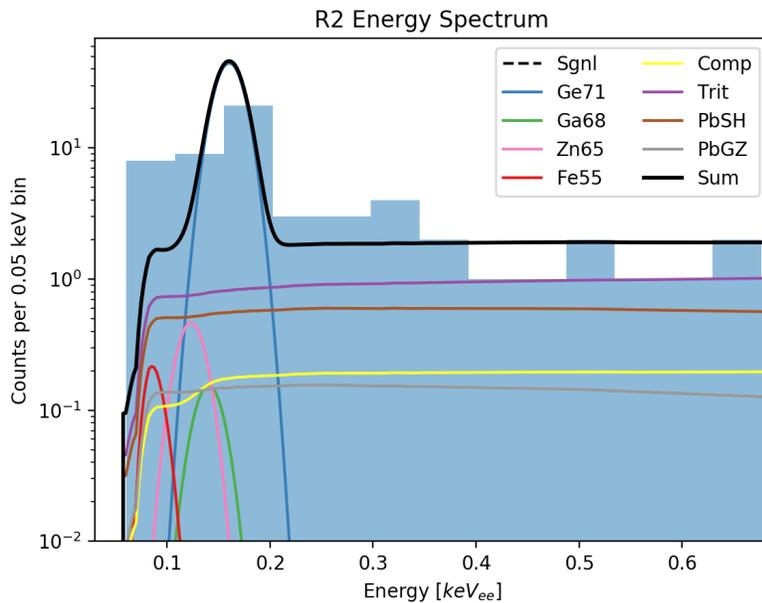
Figure 5.18: Zoomed view of the lowest energy bins of a background-only best-fit of all models summed in the black curve to the blue histogram R2 spectrum. While it is technically possible this excess of events over the summed background model is dark matter, it is incomparably more likely to be an unknown source of background.

The R2 pieces described in Section 4.2 are included in this chapter for completeness. However, only the R3 data described in Section 4.3 was used to calculate these results. The exposure of R2 is 70.1 kg-d and the exposure of R3 is 36.92 kg-d, so combining the datasets and rerunning this analysis is naively expected to give a result that is approximately $\sqrt{3}$ times better than the results presented in this thesis.

Table 5.1 contains a list of the variables used to define the likelihood function.

Variable	Definition
N	number of events in data
i	iterator over N data events
E_i	energy of event i
ν_χ	number of signal events
f_χ	signal PDF
b	iterator over non-surface backgrounds: ^{71}Ge, ^{68}Ga, ^{65}Zn, ^{55}Fe, Compton, ^{3}H
ν_b	number of events in background b
f_b	PDF of background b
sb	iterator over surface backgrounds:
	sidewall housing (SH), germanium iZIP (GZ), housing top lid (TL is R3 only)
ν_{sb}	number of events in surface background sb
ρ_{sb}	event density of surface background sb
k	iterator over systematic nuisance parameters
x_k	value of parameter k
μ_k	expected value of parameter k
σ_k	uncertainty of parameter k
$\boldsymbol{V}$	variance matrix, also known as $(\sigma_k)^2$
d	dimension of Gaussian constraint
s_k	surface background morphing parameter
$\boldsymbol{S}$	covariance matrix constraining morphing of surface backgrounds
r_k	resolution parameter
$\boldsymbol{R}$	covariance matrix constraining resolution model
e	efficiency morphing parameter (one parameter, no iterator)

Table 5.1: Variables used in the likelihood function.

The likelihood function is composed of three different terms. The first term uses a Poisson distribution to constrain the total number of events in the models as compared to the dataset,

$$\mathcal{L}_{poiss} = \left[\nu_\chi + \sum_b \nu_b\right]^N \times \frac{\exp\left[-\left(\nu_\chi + \sum_b \nu_b\right)\right]}{N!} \, . \tag{5.15}$$

There are separate Poisson terms for R2 and R3.

The second term is the core of the likelihood that uses the signal and background PDFs to determine the most likely model for each event in the dataset,

$$\mathcal{L}_{\chi,b} = \prod_{i=1}^{N} \left[f_\chi(E_i) + \sum_b f_b(E_i) \right] .$$

(5.16)

The model PDFs are unique for R2, R3a, and R3b as each sub-run has different efficiencies and resolutions.

The third term, called the "constraint term", controls the nuisance parameters from auxiliary measurements. There were three categories of constraint terms used in this analysis: morphing the surface backgrounds, efficiency, and resolution. All three categories were constrained by either univariate or multivariate Gaussian distributions [126]. Since the inelastic processes considered produce either a photon or an electron, the signal will appear as an electron recoil event in the detector. This means there is no need to introduce additional systematic uncertainty from assuming a Lindhard model to convert to keV_{nr}. The general constraint Gaussian term takes the form

$$\mathcal{L}_{constr.} = \frac{1}{\sqrt{2\pi\sigma_k^2}} \times \exp\left(-\frac{(x_k - \mu_k)^2}{2\sigma_k^2}\right) = \frac{1}{(2\pi)^{d/2}|\boldsymbol{V}|^{1/2}} \times \exp\left(-\frac{1}{2}(x_k - \mu_k)^T \boldsymbol{V}^{-1}(x_k - \mu_k)\right) .$$

(5.17)

There were two terms to control the morphing of the surface backgrounds (R2 and R3), three to control the efficiency models (R2, R3a, and R3b), and three to control the resolution models (R2, R3a, R3b).

A likelihood analysis requires fitting the likelihood function to the data in order to find the combination of parameters that maximizes the value of the function. As most algorithms available find the minimum of a function, the maximum likelihood is found by minimizing the negative of the likelihood. Also, since the likelihood is a product of many probabilities

less than one, it tends to be a very small number and is more convenient to take the natural logarithm of the function. Finally, any constant terms can be dropped to simplify the minimization because only the change in likelihood is relevant. This analysis made use of the MINUIT algorithm [127] via the iminuit [128] Python interface to minimize the negative logarithm of the likelihood.

Combining the first two likelihood terms after taking the logarithm and dropping the constant terms gives

$$
\begin{aligned}
\left[\ln(\mathcal{L}_{poiss}^{R2}) + \ln(\mathcal{L}_{\chi,b,sb}^{R2})\right] &+ \left[\ln(\mathcal{L}_{poiss}^{R3}) + \ln(\mathcal{L}_{\chi,b,sb}^{R3a}) + \ln(\mathcal{L}_{\chi,b,sb}^{R3b})\right] = \\
&-(\nu_\chi^{R2} + \sum_b \nu_b^{R2} + \sum_{sb} \nu_{sb}^{R2}) - (\nu_\chi^{R3} + \sum_b \nu_b^{R3} + \sum_{sb} \nu_{sb}^{R3}) \\
&+ \sum_{i=1}^{N^{R2}} \ln\left[\nu_\chi^{R2} f_\chi^{R2}(E_i) + \sum_b \nu_b^{R2} f_b^{R2}(E_i) + \sum_{sb} \rho_{sb}^{R2}(E_i)\right] \\
&+ \sum_{i=1}^{N^{R3a}} \ln\left[\nu_\chi^{R3} f_\chi^{R3a}(E_i) + \sum_b \nu_b^{R3} f_b^{R3a}(E_i) + \sum_{sb} \rho_{sb}^{R3a}(E_i)\right] \\
&+ \sum_{i=1}^{N^{R3b}} \ln\left[\nu_\chi^{R3} f_\chi^{R3b}(E_i) + \sum_b \nu_b^{R3} f_b^{R3b}(E_i) + \sum_{sb} \rho_{sb}^{R3b}(E_i)\right] .
\end{aligned}
\tag{5.18}
$$

Equation 5.19 shows all of the constraint terms after taking the logarithm,

$$
\begin{aligned}
\ln(\mathcal{L}_{surf}^{R2}) + \ln(\mathcal{L}_{eff}^{R2}) &+ \ln(\mathcal{L}_{res}^{R2}) + \ln(\mathcal{L}_{surf}^{R3}) + \ln(\mathcal{L}_{eff}^{R3a}) + \ln(\mathcal{L}_{eff}^{R3b}) + \ln(\mathcal{L}_{res}^{R3a}) + \ln(\mathcal{L}_{res}^{R3b}) = \\
&-\frac{1}{2}\sum_{k=1}^{2} \left[(s_k^{R2} - \mu_k^{R2})^T \boldsymbol{S}_{R2}^{-1}(s_k^{R2} - \mu_k^{R2})\right] - \frac{(e^{R2} - \mu^{R2})^2}{2(\sigma^{R2})^2} - \frac{1}{2}\sum_{k=1}^{3}\left[(r_k^{R2} - \mu_k^{R2})^T \boldsymbol{R}_{R2}^{-1}(r_k^{R2} - \mu_k^{R2})\right] \\
&-\frac{1}{2}\sum_{k=1}^{3}\left[(s_k^{R3} - \mu_k^{R3})^T \boldsymbol{S}_{R3}^{-1}(s_k^{R3} - \mu_k^{R3})\right] - \frac{(e^{R3a} - \mu^{R3a})^2}{2(\sigma^{R3a})^2} - \frac{(e^{R3b} - \mu^{R3b})^2}{2(\sigma^{R3b})^2} \\
&-\frac{1}{2}\sum_{k=1}^{3}\left[(r_k^{R3a} - \mu_k^{R3a})^T \boldsymbol{R}_{R3a}^{-1}(r_k^{R3a} - \mu_k^{R3a})\right] - \frac{1}{2}\sum_{k=1}^{3}\left[(r_k^{R3b} - \mu_k^{R3b})^T \boldsymbol{R}_{R3b}^{-1}(r_k^{R3b} - \mu_k^{R3b})\right] .
\end{aligned}
\tag{5.19}
$$

The full framework has included R2 and R3 components thus far, but now the R2 pieces are dropped to focus solely on R3. Equation 5.20 is the full likelihood function that was used for calculating the inelastic search results,

$$
\begin{aligned}
\left[\ln(\mathcal{L}_{poiss}^{R3}) + \ln(\mathcal{L}_{\chi,b,sb}^{R3a}) + \ln(\mathcal{L}_{\chi,b,sb}^{R3b})\right] + \ln(\mathcal{L}_{surf}^{R3}) + \ln(\mathcal{L}_{eff}^{R3a}) + \ln(\mathcal{L}_{eff}^{R3b}) + \ln(\mathcal{L}_{res}^{R3a}) + \ln(\mathcal{L}_{res}^{R3b}) = \\
-(\nu_{\chi}^{R3} + \sum_b \nu_b^{R3} + \sum_{sb} \nu_{sb}^{R3}) \\
+ \sum_{i=1}^{N^{R3a}} \ln \left[\nu_{\chi}^{R3} f_{\chi}^{R3a}(E_i) + \sum_b \nu_b^{R3} f_b^{R3a}(E_i) + \sum_{sb} \rho_{sb}^{R3a}(E_i) \right] \\
+ \sum_{i=1}^{N^{R3b}} \ln \left[\nu_{\chi}^{R3} f_{\chi}^{R3b}(E_i) + \sum_b \nu_b^{R3} f_b^{R3b}(E_i) + \sum_{sb} \rho_{sb}^{R3b}(E_i) \right] \\
-\frac{1}{2} \sum_{k=1}^{3} \left[(s_k^{R3} - \mu_k^{R3})^T \boldsymbol{S}_{R3}^{-1} (s_k^{R3} - \mu_k^{R3}) \right] - \frac{(e^{R3a} - \mu^{R3a})^2}{2(\sigma^{R3a})^2} - \frac{(e^{R3b} - \mu^{R3b})^2}{2(\sigma^{R3b})^2} \\
-\frac{1}{2} \sum_{k=1}^{3} \left[(r_k^{R3a} - \mu_k^{R3a})^T \boldsymbol{R}_{R3a}^{-1} (r_k^{R3a} - \mu_k^{R3a}) \right] - \frac{1}{2} \sum_{k=1}^{3} \left[(r_k^{R3b} - \mu_k^{R3b})^T \boldsymbol{R}_{R3b}^{-1} (r_k^{R3b} - \mu_k^{R3b}) \right] .
\end{aligned}
$$
(5.20)

5.5. Limit Calculation

Typically, with a profile likelihood calculation, a parameter of interest (usually σ for dark matter searches) is scanned and a test statistic is calculated for each hypothesis. A standard definition for this test statistic q is shown in Eq. 5.21 that is the negative logarithm of the ratio of the likelihood for each test hypothesis divided by the unconstrained best-fit likelihood,

$$
q(\sigma) = -2\ln \left(\frac{\mathcal{L}(\sigma, \hat{\hat{\theta}})}{\mathcal{L}(\hat{\sigma}, \hat{\theta})} \right) .
$$
(5.21)

In practice, since the minimizer outputs the negative logarithm of the likelihood, a logarithm identity can be used to rewrite this formula as the difference between the two logarithms.

The upper and lower limits of the parameter of interest are obtained by finding where a scan of the test statistic crosses a threshold determined using Wilks' Theorem [129]. Stated

another way, the signal model that corresponds to the extreme parameter values outside of the upper or lower limit are said to be inconsistent with the data at the given confidence level and that disagreement will only get worse further from the best-fit parameter. A cartoon of this process is shown in Fig. 5.19 with a characteristic parabolic shape and an arbitrary confidence level.
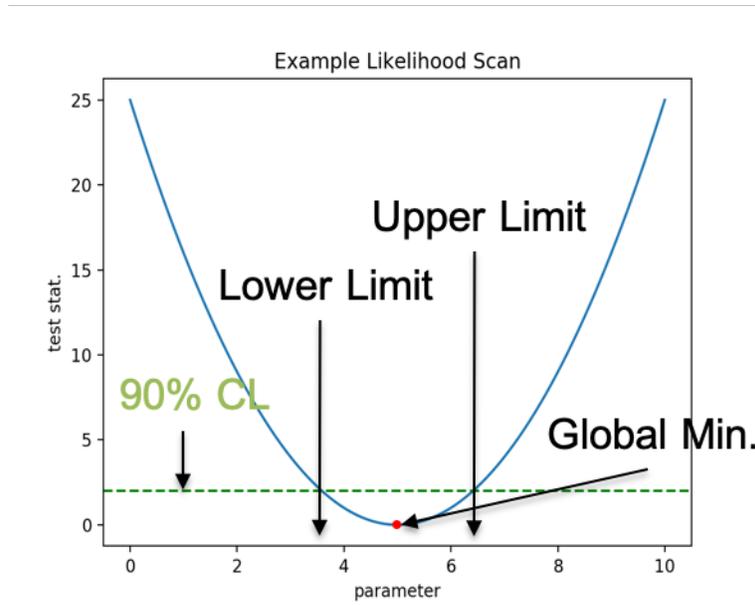
Figure 5.19: Cartoon demonstrating a typical approach of using a likelihood ratio to set limits on a parameter of interest. The red point is the unconditional best-fit that corresponds to the likelihood in the denominator of the test statistic. The hypothesis in the numerator changes and q is plotted as the blue line. The green line is some arbitrary 90% confidence level, so the lower (upper) limit is the parameter value where the blue curve crosses the green line below (above) the best-fit parameter.

This standard approach, however, is not adequate for this analysis due to the velocity damping discussed in Section 5.3. With an unattenuated signal, the number of signal events will increase linearly with increasing cross-section, but the attenuation of the signal causes the number of signal events to decrease. This turnover effect is depicted in Fig. 5.20 where the number of signal events is plotted as a function of cross-section for a given dark matter mass. In other words, two different cross-sections can provide the same number of signal events in the data.
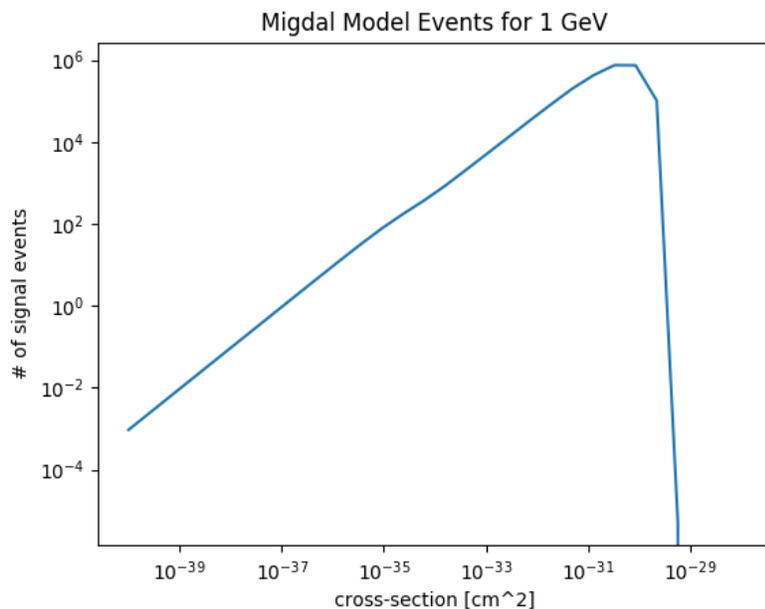
Figure 5.20: Plot of number of signal events found by integrating the model created from cross-section. The number of events increases linearly with increasing cross-section until the attenuation turns on and sharply decreases the number of events. This creates the unusual case where different cross-sections result in the same number of signal events.

The turnover of the signal is also reflected in the likelihood scan, which causes a secondary minimum as shown in Fig. 5.21. Instead of excluding all values above the upper limit, the exclusion region lies between the two minimums where the cross-sections give the strongest signal.

Even more strange is the case where the two minimums have different likelihood values. The left plot of Fig. 5.22 shows an example of this. The region between the minimums, indicated by the orange band, still corresponds to large signals that the experiment should be sensitive to. However, the standard algorithm is not able to find the upper limit if the likelihood scan does not cross the confidence threshold. Plotting the signal model with the cross-section for each minimum provides insight into the different likelihood values, shown in the right plot of Fig. 5.22. Despite integrating to roughly the same number of events, the larger cross-section experiences stronger Earth shielding that attenuates the dark matter signal and pushes more events to low energy. In other words, the cross-section does not just
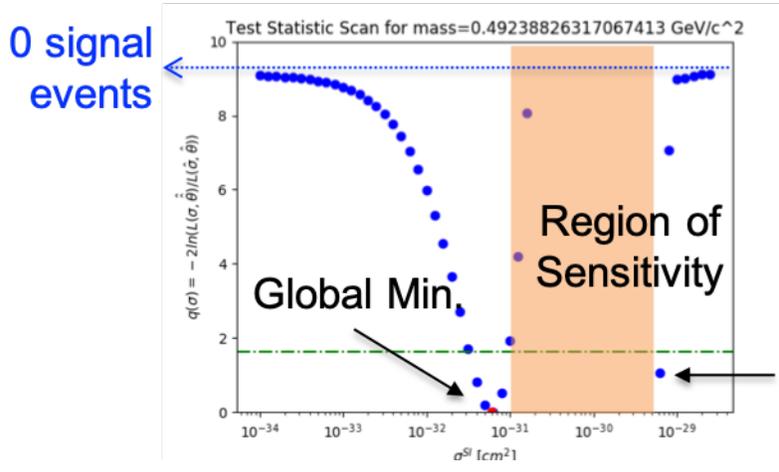
Figure 5.21: Profile likelihood scan that demonstrates the likelihood turning over similar to Fig. 5.20 and leading to a secondary minimum. The orange band roughly highlights the expected region of sensitivity between the minimums where the signal is the strongest. The "bounce" back up at extreme values of cross-section can be ignored. This plot included R2 data and the exclusion of zero signal was the first indicator that the R2 data might have unmodeled backgrounds effecting the likelihood.

control the number of signal events, but also the shape of the signal model. It is apparent that the likelihood found a better fit for the larger cross-section, meaning that the data prefers more low-energy events in the models.

A new method of limit setting had to be devised to handle the attenuated signal. The method decided upon is brute force and computationally intensive, but it is designed to individually treat every dark matter mass and cross-section as a new test hypothesis. For every combination of mass and cross-section in the entire dark matter parameter space, there is a specific signal model shape and expected number of events. The question to ask is whether this signal is consistent with the data. Instead of using cross-section as the parameter of interest in the likelihood scan, the new method scans the number of signal events, which is done by manually setting the normalization of the signal model to different values. By setting the number of signal events, the correlation between shape and normalization is broken and the scan no longer has a double minimum. As with the standard method, the upper limit is determined to be where this scan crosses the predetermined confidence level.
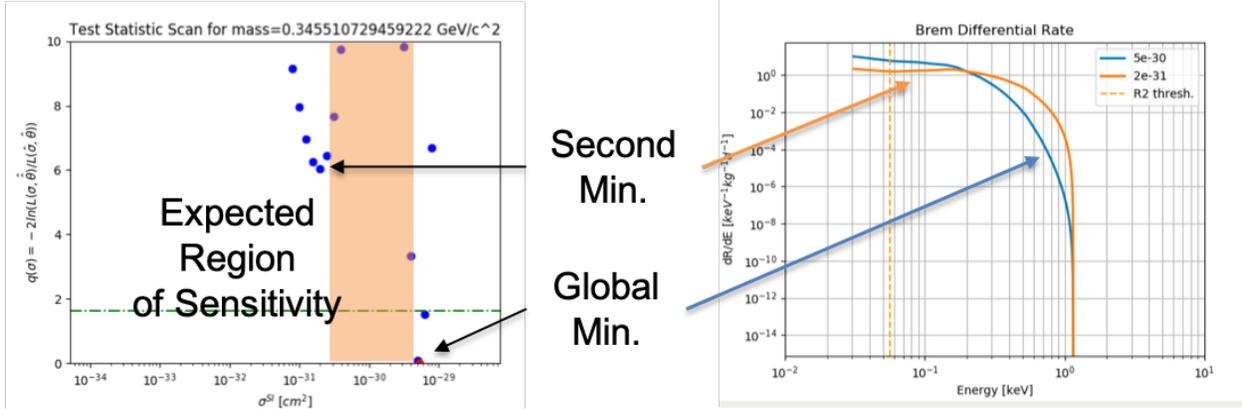
Figure 5.22: The left plot shows a likelihood scan that resulted in two minimums with different likelihoods, and the secondary minimum did not cross back over the confidence threshold. The orange band indicates the expected region of sensitivity, but the traditional algorithm would fail to find this. The plot on the right illuminates the difference in likelihoods. Despite having the same number of signal events, the blue model is a better fit to the data after the attenuation pushed more events to low energy.

5.5.1. Expected Sensitivity

In order to estimate the expected sensitivity in the absence of a dark matter signal, the standard practice is to create numerous background-only Monte Carlo datasets from the background models. Then the full analysis framework is run on the pseudo data to see the effect of the randomly thrown spectra. If there is a significant signal present, then any exclusion results would be worse than the calculated sensitivities.

As mentioned above, the new limit setting method was computationally very intensive. It would have been ideal to run this test with over 100 pseudo datasets, but instead only a handful could be performed in a timely manner. The results of this test are shown in Fig. 5.23.

There are not enough tests to make statistical claims about the probability of seeing the results from the real data, but Fig. 5.23b qualitatively shows that the results in this thesis are consistent with a background-only Monte Carlo pseudo dataset.
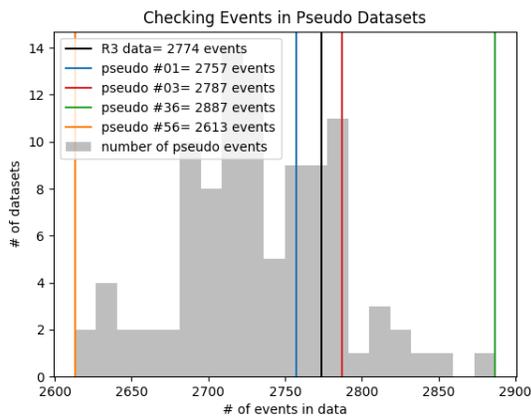
(a) Histogram of the total number of events in each pseudo dataset. The vertical lines indicate the number of events in the real dataset (black) as well as the pseudo datasets used to estimate the sensitivity (colors). The total number of events does not directly correlate to the limit, but datasets were selected from the two extremes and around the real number of events.
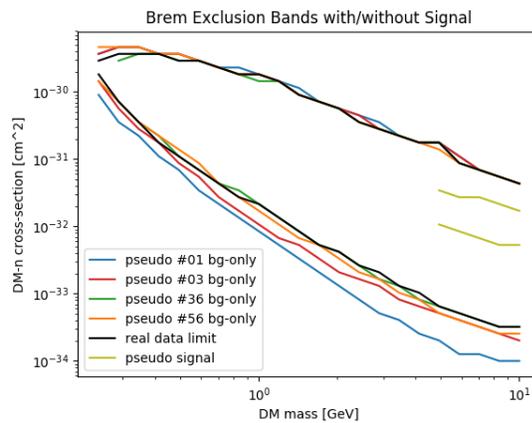
(b) Exclusion bands calculated from the pseudo datasets indicated in Fig. 5.23a. The black curve from the real data is one of the more conservative limits. However, the result was replicated from background-only data with the closest being the green curve. The gold band is the result of injecting a large signal into the Monte Carlo data, clearly resulting in a worse limit.

Figure 5.23: Results of the sensitivity estimate. This qualitatively proves that the final results in this chapter are consistent with background-only exclusion limits.

5.5.2. Exclusion Limit

Since this analysis did not see an excess of events, an exclusion was calculated and only the upper limit was needed. To simplify the algorithm that finds the points of intersection, the test statistic was explicitly set to zero below the best-fit global minimum as defined in

$$
q(\nu_\chi) = \begin{cases} -2\ln\left(\dfrac{\mathcal{L}(\nu_\chi,\hat{\hat{\theta}}|m,\sigma)}{\mathcal{L}(\hat{\nu_\chi},\hat{\theta}|m,\sigma)}\right) & \nu_\chi > N_{fit} \\ \\ 0 & \nu_\chi < N_{fit} \end{cases} .
\tag{5.22}
$$

Then the expected number of signal events is compared to the upper limit calculated. There are two possible outcomes: if the expected number of events is smaller than the upper limit then that mass and cross-section is consistent with the data, and if the expected number of events is larger than the upper limit then that mass and cross-section are excluded. Figure 5.24 and Fig. 5.25 portray the two cases where the signal is consistent with data or excluded, respectively.

The procedure described above is repeated over the entire parameter space and any point that resulted in an excluded model was collected. In practice, every single likelihood fit could be performed independently whether it was an unconditional fit with all parameters in the likelihood free to float, or the signal model was clamped to a particular test hypothesis. This allows the problem to be massively parallelized to take full advantage of ManeFrame II, the high-performance computer cluster at SMU. To calculate the final exclusion bands with high granularity, there were approximately 25,000 (45,000) jobs run in parallel for the Bremsstrahlung (Migdal) signal model.

5.6. Exclusion Band Results

The result of the new limit setting method is a band of points where it can be stated with 90% confidence that, given the CDMSlite R3 data, a dark matter signal was not observed.

PLR Scan for mass=1.7012542798525891, xsec=2.0860644934719524e-3...
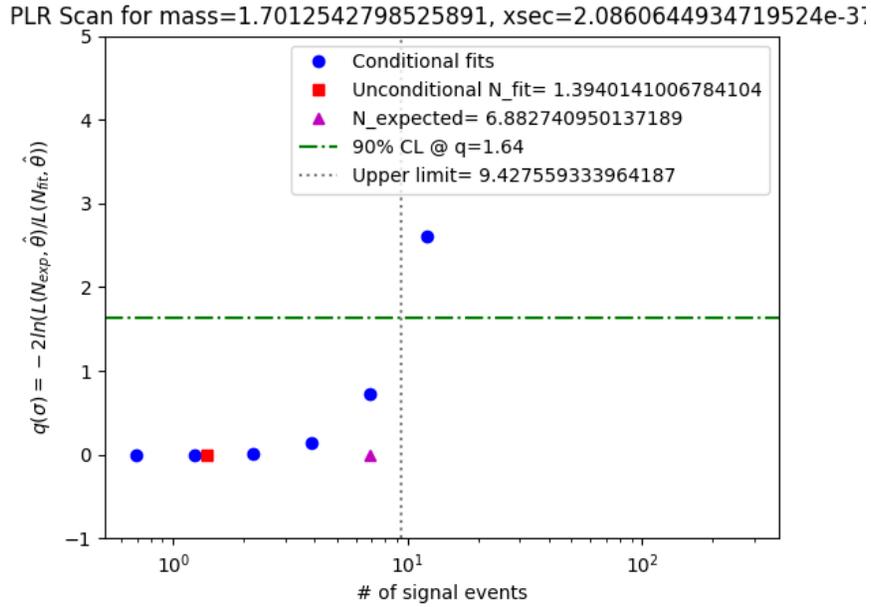
Figure 5.24: Example of a likelihood scan over number of signal events using the new limit setting method. The red square is the best-fit and all of the blue test hypotheses are set to zero to the left of this point. The vertical gray line is the upper limit where the interpolated blue points are determined to cross the green 90% confidence level threshold. The expected number of events for the dark matter mass and cross-section in the title is indicated by the purple triangle. Since the expected number is smaller than the upper limit, this particular signal model cannot be excluded.
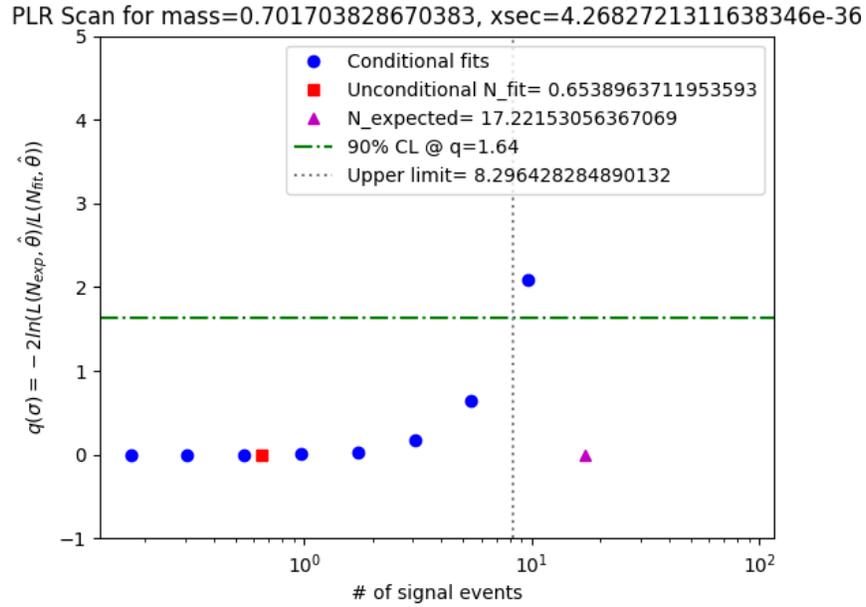
Figure 5.25: Similar plot to Fig. 5.24 where the red square is the best-fit parameter, the blue circles are the test hypotheses, the green line is the 90% confidence level, the gray line is the upper limit on number of events, and the purple triangle is the expected number of events for the signal. In this case, the expected number of signal events is larger than the upper limit, so the signal model for this dark matter mass and cross-section is not consistent with the data. The mass and cross-section in the plot title become a point in the final exclusion band.

The Bremsstrahlung band is shown in Fig. 5.26 and the Migdal band is shown in Fig. 5.27. The bottom edge of each band is the exclusion limit that is historically reported, but the top edge of each band also shows the limit of sensitivity due to the consideration of dark matter shielding.
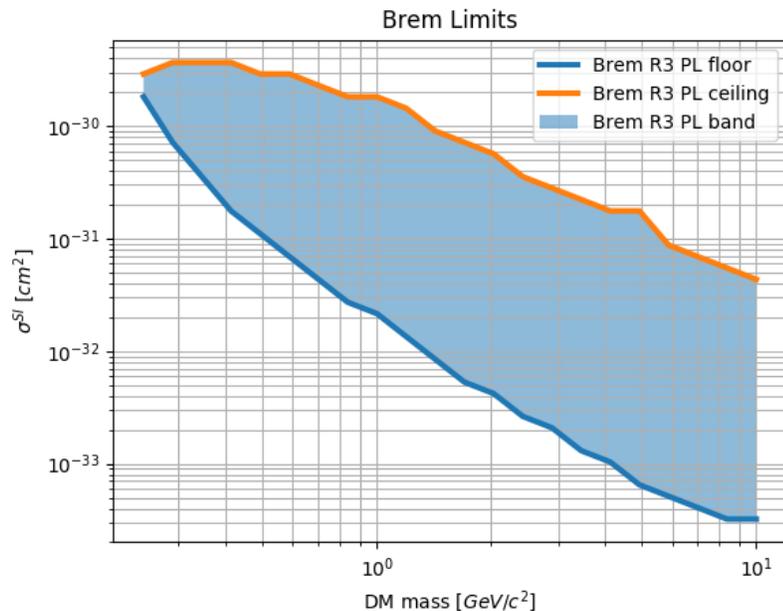
Figure 5.26: Bremsstrahlung exclusion band created with the new limit setting method. The floor and ceiling of the experimental sensitivity are naturally created by the limit setting. Despite the persistence of some small fluctuations around the edges, the parameter space was sampled with a high granularity.

For the most direct comparison, Fig. 5.28 shows all of the inelastic results together. The lower threshold of CDMSlite allows the Bremsstrahlung limit to extend a bit below the 0.4 GeV reach of the LUX Bremsstrahlung analysis, but LUX has 379 times more exposure that helps combat the exposure-limited Bremsstrahlung channel and reach lower cross-sections [115].

For the Migdal channel, the first comparison is to the LUX Migdal result. In this case, the lower threshold of CDMSlite is far more important than the larger exposure of LUX in extending to masses an order of magnitude lower [115]. EDELWEISS has also published a Migdal limit with a small 33.4 g germanium detector operated for 24 hours at the sur-
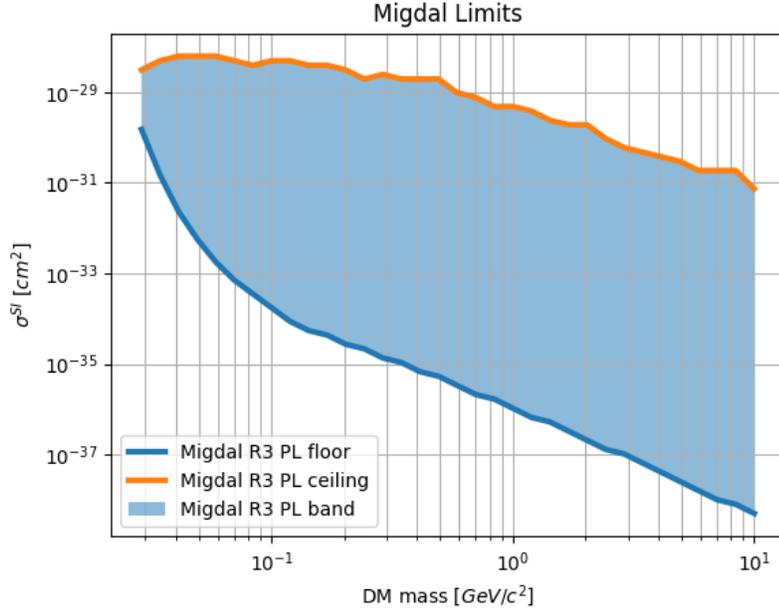
Figure 5.27: Migdal exclusion band created with the new limit setting method. The floor and ceiling of the experimental sensitivity are naturally created by the limit setting. Despite the persistence of some small fluctuations around the edges, the parameter space was sampled with a high granularity.

face [123]. CDMSlite R3 has more than 1000 times more exposure than the EDELWEISS surface data, which allows this result to extend many orders of magnitude lower in cross-section. However, it is interesting to note that the ceiling of the EDELWEISS band is much higher than this result because it is not shielded by the Earth while operating on the surface.

And finally, Fig. 5.29 incorporates the Bremsstrahlung and Migdal results into the current state of low-mass dark matter direct detection limits. Unfortunately, the Bremsstrahlung limit does not exclude any new parameter space because of the drastically reduced rate compared to the elastic nuclear recoil interaction. That is why CRESST was able to exclude lower cross-sections by lowering their nuclear recoil threshold to 30.1 eV, which is enough to be sensitive to sub-GeV masses [130]. The more exciting news is that the Migdal channel results in a world-leading limit below 0.2 GeV.

5.7. Summary
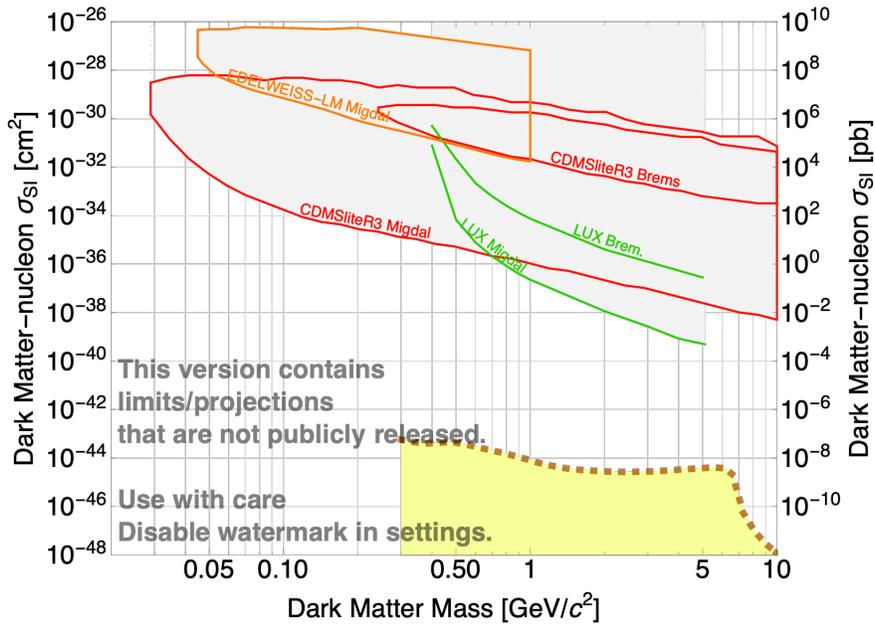
Figure 5.28: Limit plot of the current state of all the inelastic channel results.
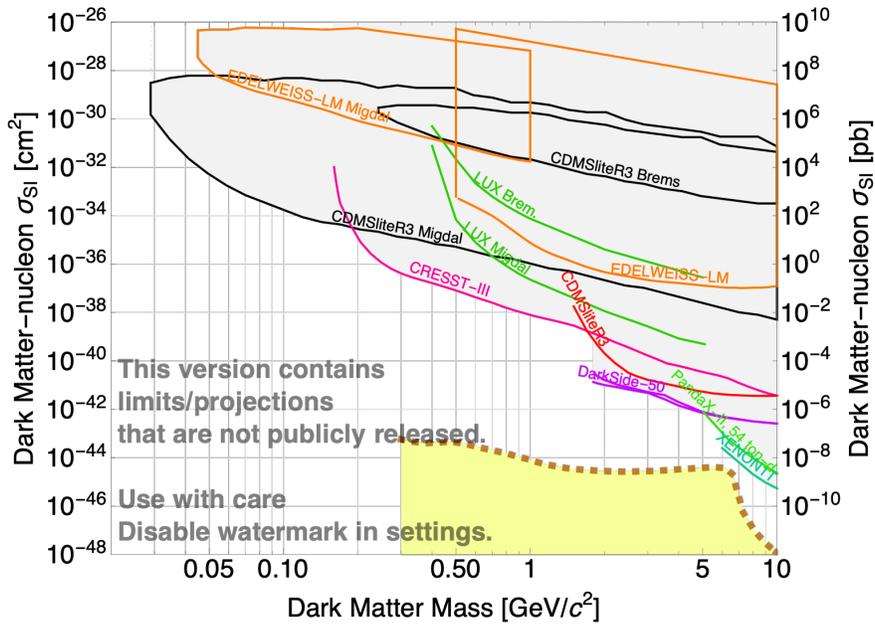
Figure 5.29: The current state of low-mass dark matter direct detection experiments with the new Bremsstrahlung and Migdal channel results. The Bremsstrahlung result does not exclude any new parameter space because of the significantly decreased rate with respect to elastic nuclear recoils. The Migdal channel results in a world leading limit below 0.2 GeV due to the combination of having a 70 eV energy threshold and 36.92 kg-d exposure.

This analysis was a re-interpretation of the published CDMSlite R3 dataset. One clear improvement for a future analysis is to go back on a cut-by-cut basis and optimize the R2 data to remove possible sources of background leakage. Many of the cuts are fairly basic and might be acceptable, but two of the more concerning cuts to investigate are the low-frequency noise cut and the radial fiducialization cut. This will immediately give a factor of 3 increase in exposure that has been shown to have a significant impact on the decreased rates of the inelastic channels.

There were also a number of assumptions and simplifications that went into both signal models and dark matter shielding model. We relied on the work done in the referenced theory papers, but there is a lot of solid state physics and atomic physics that could make more accurate models. For example, we admit our practice of excluding the valence shell is a temporary patch because the remaining shells are still not in isolation. This is a problem that the general dark matter community is aware of, yet no one seems to have figured out how to redo the calculation.

Finally, the Migdal Effect has never been observed experimentally. While everything in the theory seems plausible, it would be much more reassuring if this effect were physically validated. It would lend significantly more credibility to the Migdal Effect as being a reliable channel through which searches for low-mass dark matter can be performed.

CHAPTER 6

CONCLUSION

As expressed in Ch.1, there is an abundance of observational evidence that dark matter exists, yet its constituents have evaded direct detection. There are numerous theories proposing dark matter candidates over a wide range of possible masses and interaction cross-sections. Experiments around the world will continue to push the boundaries of detector technologies to explore as much of that parameter space as possible. Generally speaking, noble liquid detectors are prime for scaling to larger detector mass in order to push to lower cross-sections, and solid state detectors are developing new methods to reduce energy thresholds and probe low-mass dark matter. Besides improvements in hardware design, new analysis techniques are also being developed to improve experimental sensitivity. The Bremsstrahlung and Migdal channels for inelastically scattering dark matter provide the ability to search for low-mass dark matter interactions that would otherwise be undetectable.

This thesis has presented the results of searching for Bremsstrahlung and Migdal signals in CDMSlite R3 data. The Bremsstrahlung search resulted in a leading limit for the Bremsstrahlung channel below 0.4 GeV, but this parameter space has already been excluded by searches for elastically scattering WIMPs. This is because of the significantly decreased rate with respect to traditional elastic nuclear recoils. However, the Migdal search resulted in a world leading limit for dark masses below 0.2 GeV, at a minimum cross-section of 3×10^{-35} cm^2. This result is encouraging for the utility of the Migdal channel in future searches, and there are a couple of obvious paths for improvement.

To begin with, the theoretically-motivated Migdal effect has not been physically observed. Verifying the effect with a dedicated experiment would help to alleviate most of the

skepticism surrounding the Migdal signal. There are technical challenges to designing an experiment capable of measuring the Migdal effect. It would have to use a neutron source to provide a source of nuclear recoils, but the incoming neutrons need to have sub-keV energies. Higher energy collisions will completely dislocate the neutron from the atom, whereas the Migdal effect just needs the neutron to be perturbed from the center of the atom. The experiment also needs to have low background rates with respect to the rate of the neutron source to have a high signal-to-background ratio. One of the challenges is how to direct a low energy neutron source at the detector and not have all of the neutrons thermalize in the required shielding.

The other pressing need relates to the structure of the signal model. The assumption that the target atom is isolated needs to be corrected to account for potentials of the surrounding atoms. It is not clear yet if the best method is to redo the calculation from first principles, or to determine the correction numerically with a simulation. Accounting for the crystal dynamics could also allow the outer shell to be included in the signal model. By excluding the valence electrons from the model, an artificial energy threshold was enforced in the analysis, such that reducing the experimental threshold does not gain any further sensitivity to low masses. Upgrading the calculation to include the valence shell creates the opportunity for increased sensitivity by reducing the experimental energy threshold.

Looking toward the future of SuperCDMS, construction of the next generation experiment at SNOLAB began in mid-2019. When construction concludes, a commissioning period will commence in 2020. There will be interesting results from the data collected in the commissioning period, but the first dark matter search data is expected to begin in 2021.

In the initial deployment at SNOLAB, there will be 10 (2) germanium (silicon) iZIPs and 8 (4) germanium (silicon) HV detectors. These larger detectors will have approximately twice the mass of the Soudan iZIPs. All HV detectors will also be read out simultaneously, whereas the electronics at Soudan could only support a single CDMSlite detector at a time.

Looking at the germanium HV detectors as an example, this deployment scheme means there will be 16 times more exposure than CDMSlite for the same run time. Assuming there are no improvements with the new detectors, the increased exposure alone would naively lead to a sensitivity that is 4 times lower in cross-section. However, SNOLAB has 3 times more overburden than Soudan, so the reduced ceiling of experimental sensitivity will counteract the mass increase.

I believe the SuperCDMS SNOLAB detectors best-suited to search for inelastic dark matter scatters will be the germanium HV detectors. While the Migdal rate is not as suppressed as the Bremsstrahlung channel, it is still a few orders of magnitude below the elastic nuclear recoil rate. Any future inelastic dark matter search will benefit from having more exposure than the CDMSlite detector at Soudan. The Migdal rate is also enhanced for heavier atoms because they have more electrons, and thus a higher probability of emitting one. This is why germanium is favored over a silicon target.

To reference the name of the dark matter summer school I attended at SLAC in 2014, "Shining Light on Dark Matter" is more than a catchy play on words. I have met brilliant physicists from all over the world diligently working to unravel one of the greatest unsolved mysteries in the Universe, and I firmly believe that one day we will successfully conclude this noble endeavor.

XIA ALPHA SCREENING AT SMU

When searching for dark matter, any radioactive decays near the detectors can mimic signals and reduce the sensitivity of the experiment. Two methods of reducing this background are to use materials of construction with inherently low bulk activity, and to develop cleaning methods and manufacturing processes to reduce surface contamination. In either case, it is necessary to have a sensitive counter for screening materials. There is an XIA UltraLo-1800 alpha counter in the Laboratory for Ultrapure Materials and Isotope and Neutron Assessment (LUMINA) at SMU to screen the surface of samples for radon contamination. Refer back to Chapter 2 for details about radon and surface contamination. In particular, the SMU screening program focuses on the long-lived ^{210}Pb that has a half-life of 22.3 years because this is the isotope that will not decay away within the lifetime of most experiments. A reliable way to measure the presence of ^{210}Pb is to detect the subsequent 5.3 MeV alpha decay of ^{210}Po.

A.1. XIA ULTRALO-1800

The XIA UltraLo-1800 is an ionization counter that is composed of a bottom sample tray that is grounded and a pair of positively charged electrodes at the top. The bottom tray is made of stainless steel, but the sample tray in the SMU counter is lined with conductive teflon because it can be easily cleaned or discarded altogether if there is any indication of surface contamination. The teflon is also measured to have an emissivity of 1×10^{-3} counts/(h cm^2), whereas the steel is approximately 3×10^{-3} counts/(h cm^2) [131]. The active volume is filled with boil-off argon gas so an emitted alpha particle will create an ionization track that drifts through the electric field to the electrode. Copper strips separated by resistors line

the sidewalls of the counter and act as field shapers to prevent any fringing and ensure the uniformity of the electric field. An integrating preamplifier is connected to the electrode that measures the induced current of the drifting electrons, so the electrode starts recording the pulse as soon as the ionization track is created, defined as time t_0. The voltage at the electrode is 1100 V and the height of the chamber is 15 cm, so it takes approximately 70 μs for the free electrons to drift to the electrode [131]. When the ionization track drifts across the active volume and first reaches the electrode, defined as time t_S, the electrode begins absorbing the electrons. The risetime is defined as being the difference between t_0 and t_R, where t_R is when the last electron is absorbed by the electrode. A schematic of the device is shown in Fig. 1.1 and a cartoon of the important times associated with the pulse are shown in Fig. 1.2.
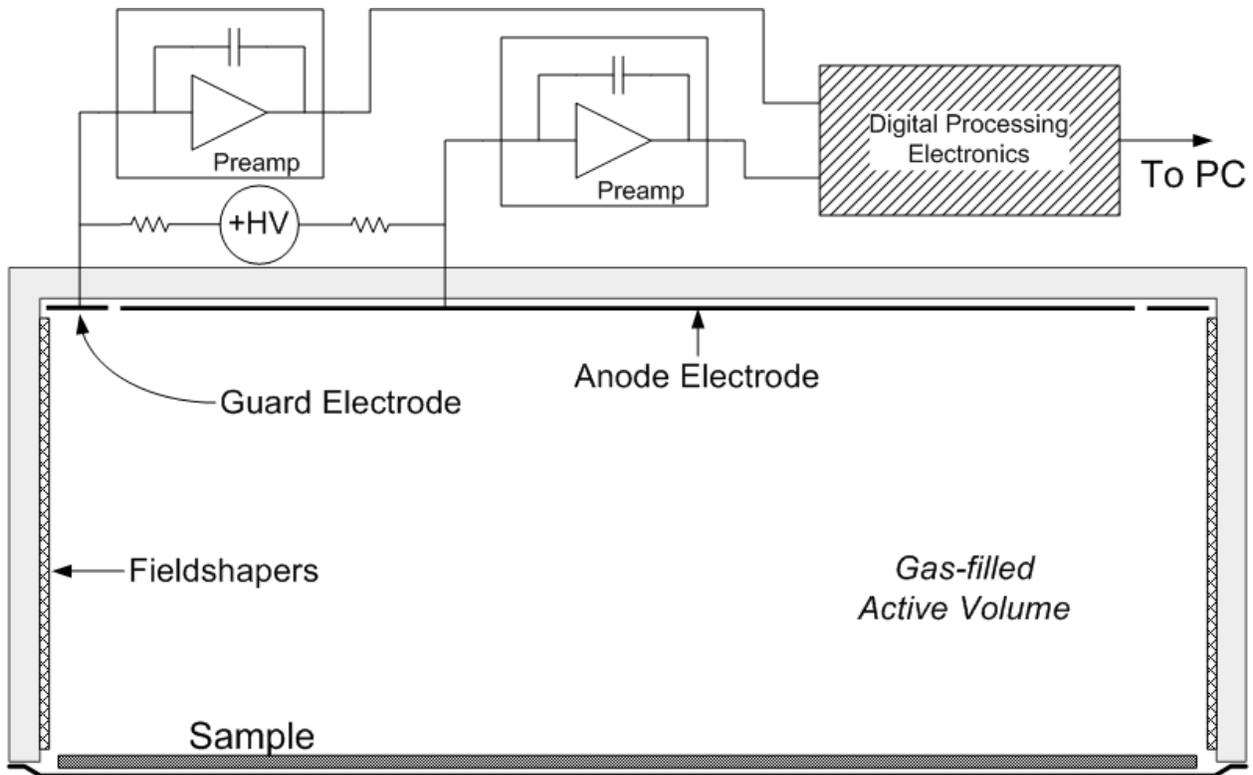
Figure 1.1: Schematic of the XIA UltraLo-1800 alpha counter shows some of the major components of the ionization counter as well as the electronics used to read out the pulses [131].
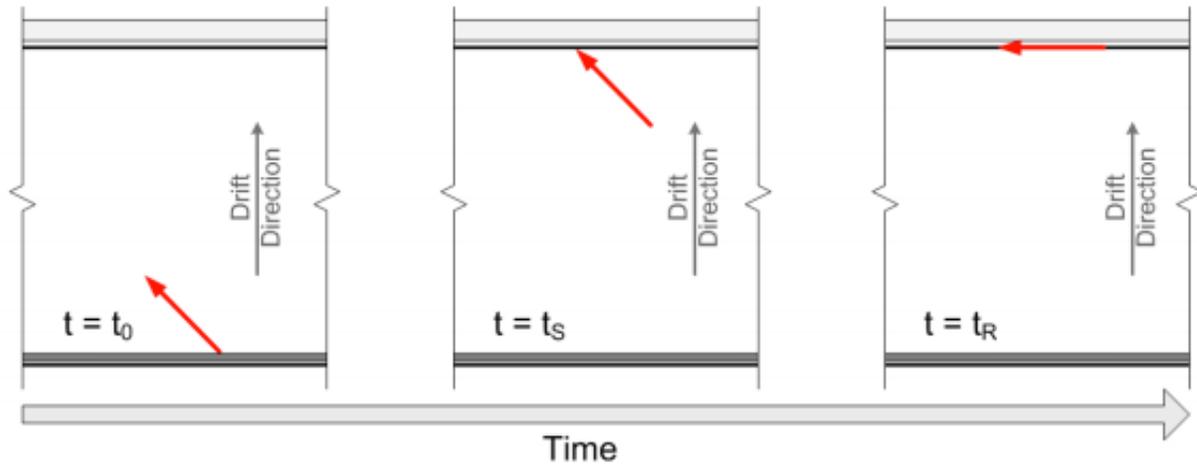
Figure 1.2: Cartoon depicting the definitions of times associated with the pulse: t_0 is when the ionization track is created, t_S is when the track reaches the electrode, and t_R is when the full track has been collected by the electrode [131].

The XIA company provides proprietary software, called CounterMeasure, with each UltraLo-1800 that controls the counter and analyzes the data. Up front, the program is very user-friendly and only a few parameters are needed to start a run such as how long to count for and how long to purge the detector before counting starts. There are also a number of tunable parameters behind the scenes that control more technical aspects of the XIA, such as the energy calibration [131].

There are two counting areas that are selected in CounterMeasure that correspond to the desired anode area. The "full" area is the full square tray with 1800 cm^2, hence the name of the UltraLo-1800. Outside of the full anode electrode is a guard electrode that is used to reject events originating from the sidewalls of the counter, conveniently called "sidewall events." The other area that can be selected is a circle in the center with 707 cm^2, called the "wafer." When this option is selected, the anode outside of the wafer is considered part of the guard electrode and is used to reject sidewall events.

A.1.1. Background Rejection

Besides rejecting sidewall events when the guard electrode measures a significant amount of energy, there are other background rejection techniques performed by CounterMeasure. Events that originate at the electrode are called "ceiling events" and are also easy to reject based on their extremely short time between t_0 and t_S. Figure 1.3 shows the possible origin of alpha decays inside the counter.
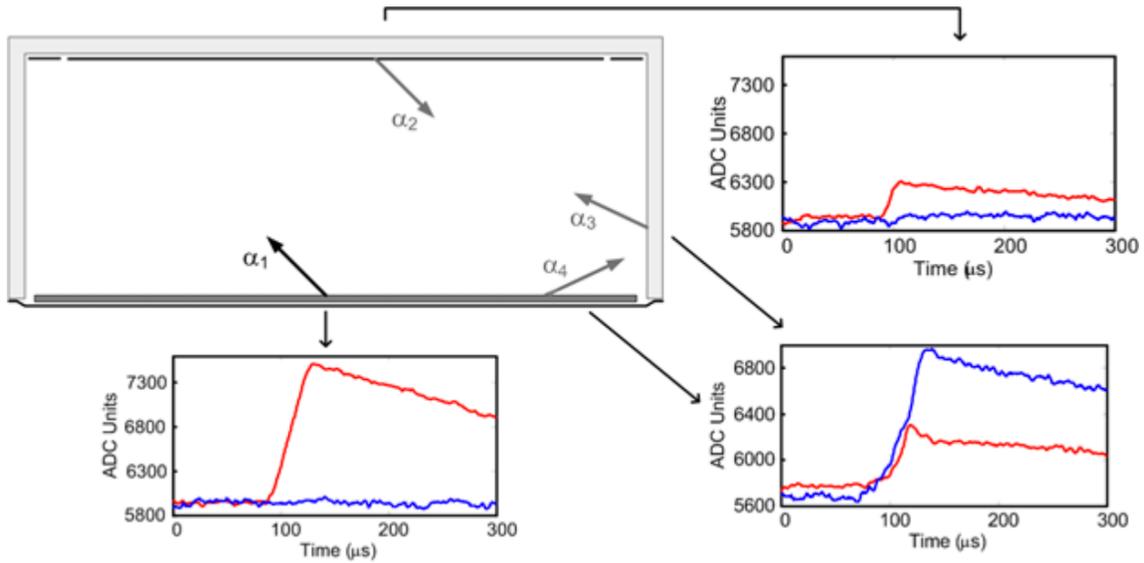
Figure 1.3: Cartoon showing the possible origins of alpha decays and what the resulting pulses look like in the anode (red) and guard electrode (blue) [131]. A "good" pulse from the sample is shown as α_1, while α_2 is a ceiling event, and α_3 and α_4 would be classified as sidewall events because of the significant pulse in the guard electrode.

The pulses shown in Fig. 1.3 all result from alpha decays from different surfaces in the alpha counter, but there are other sources of background that do not originate from a surface. There are a class of events that appear in the XIA that exhibit a significantly more rounded pulse that is not consistent with the short, dense ionization track created by an alpha particle. It is presumed that these events are created by cosmogenic particles, and they are identified using the time difference between t_S and t_R. Figure 1.4 shows an example of what a "round" pulse looks like. According to the XIA User's Manual, the expected rate of these events at

sea level without extra shielding is approximately 2×10^{-4} counts/(h cm^2) and virtually all of them fall below 3.5 MeV, so a simple energy cut can remove them if that energy range is not needed for screening [131].
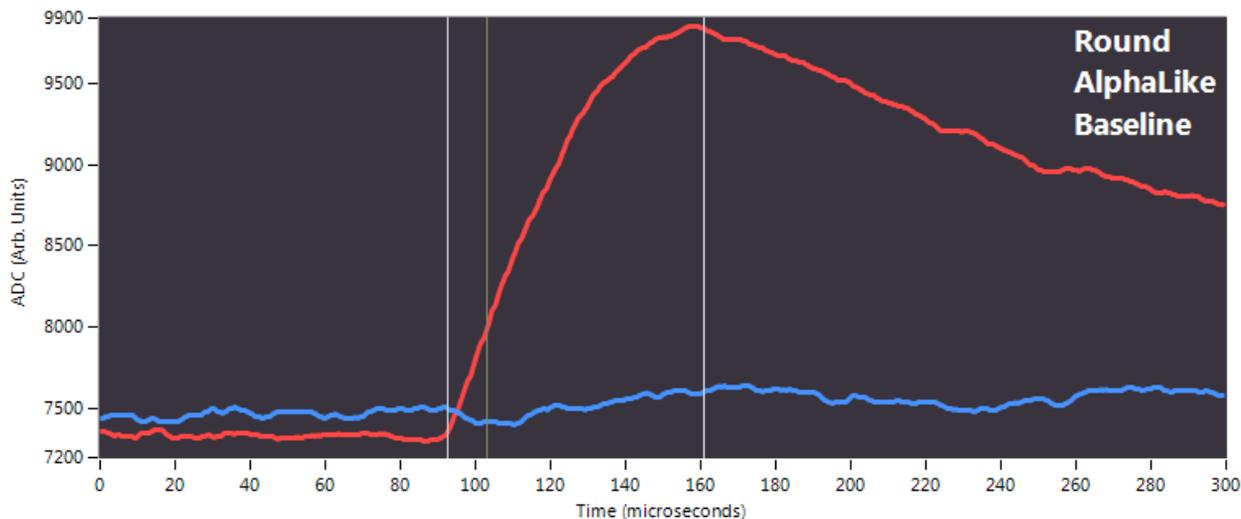
Figure 1.4: Example of a typical round pulse with the anode shown in red and the guard electrode in blue [131]. These events are rejected by CounterMeasure when the time between t_S (second vertical line) and t_R (third vertical line) is sufficiently long.

Another class of events that are not from a surface decay are called "mid-airs" that presumably originate from radioactive decays in the counting gas. The expected rate of these events is 2.5×10^{-4} counts/(h cm^2) and they can have a wide range of energies, so the best rejection method is to use the drift time [131]. As stated above, a typical track from the sample tray will take approximately 70 μs to drift across the active volume, so the assumption is that any event with a shorter risetime must have been from a decay in the gas above the sample tray. The risetime cut is defined when starting a run in CounterMeasure, and is typically set to 60 μs.

Within the last year there have been abnormally high background levels in the XIA that decay away with the half-life of ^{222}Rn. It was originally hypothesized that radon could be contaminating the argon gas if the liquid argon level was low and any radon that froze in the dewar was warmed enough to boil off with the argon. However, a detailed study did

not reveal any correlations between background levels and liquid argon level. Instead, it was discovered that the background level spiked immediately after the argon dewar was refilled and reconnected to the XIA. It is still not clear what introduced the radon, but subsequent runs have shown that venting the argon dewar into the laboratory for $10 - 15$ s before reconnecting the regulator is an effective procedure for preventing contamination of the argon gas.

A.2. Pulse Shape Study

CounterMeasure already uses basic information and pulse shapes to categorize and reject background events, which is what makes the XIA the most sensitive alpha counter that is commercially available. However, direct detection experiments are constantly improving their sensitivity by mitigating backgrounds. As a result, the materials of construction are lower activity, and the counting time of runs must be extended up to a full month in order to gather sufficient statistics. [1] However, counting for longer does not help when the emissivity of the sample is comparable to or lower than the background of the XIA. The target surface background rate for SuperCDMS SNOLAB is 50 nBq/cm^2, which converts to 1×10^{-4} counts/(h cm^2) assuming 50% efficiency of isotropic decays that leave the surface versus those that go into the bulk of the sample. As stated above, the background of the teflon tray liner is approximately 1×10^{-3} counts/(h cm^2), so the background needs to be reduced in order to improve the sensitivity of the alpha counter.

That is why the group at SMU has taken a more aggressive approach to decreasing the background rate of the XIA, and the sample tray is one of the more difficult background sources to identify due to the proximity of the sample. One idea was that the presence of a conducting sample would distort the electric field enough to differentiate between events that originated from the sample and those that originated from the tray. A distorted field

---

[1]The argon dewar only lasts one month before it needs to be refilled, so this is the longest the XIA can continually run for at SMU.

would affect the drifting of the ionization track through the argon gas as well as the Ramo-Shockley field that induces a current in the readout electronics. Both of these effects should be reflected as changes in the rising edge of the raw event pulses. The following study attempts to find a difference in pulse shapes in a dedicated manner.

A.2.1. Source Locations

To test for different pulse shapes, a measurement program was designed to systematically place a ^{230}Th alpha source at various positions on and off a 4 inch square copper sample shown in Fig. 1.5. As a control group, the same positions were repeated without the sample. This approach has two added advantages: it can show whether the pulses have any positional dependence within the counter, and also show how a conductive sample might alter the electric field that the ionization track experiences. The latter can be extracted by comparing pulse shapes at the same position with and without the sample present.
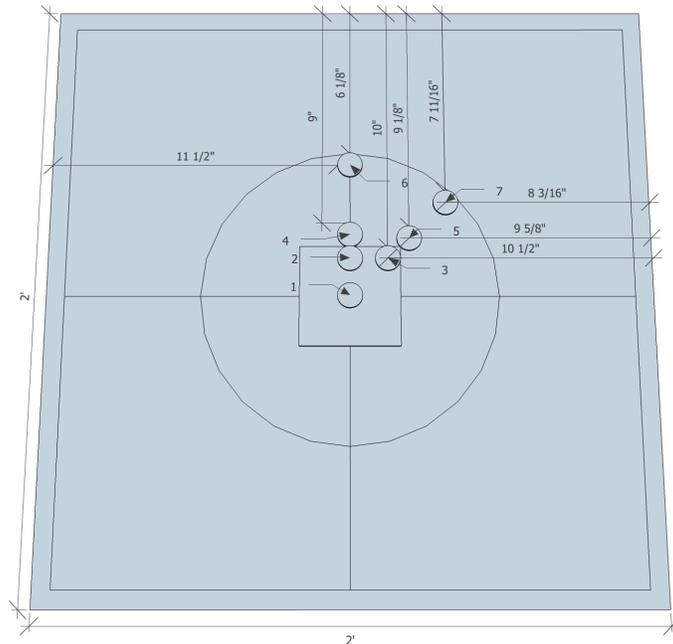
Figure 1.5: XIA tray shown with the square sample and round alpha source at different locations. Notice that positions 1-3 are on the sample and 4-7 are on the tray.

A.2.2. Data Runs

The data taken with the source at each position was saved in the folder `xia_analysis/pulse_shape_study` on the XIA laptop and `/data/xia/pulse_shape_study` on Hertz (previously was Juno). The name of the file is created as: <date>_<time>_purge<time to purge in minutes>_<risetime cut in μs>_<counting area in cm^2>_<identifying information>. The runs are as follows:

- `150706_1307_purge60_40_707_sample_position1`

- `150707_1404_purge60_40_707_sample_position2`

- `150708_1255_purge60_40_707_sample_position3`

- `150709_1226_purge60_40_707_sample_position6_crashed`

- `150710_1121_purge60_40_707_sample_position6`

- `150713_1603_purge60_40_707_sample_position7`

- `150715_1502_purge60_40_707_nosample_position7`

- `150716_1337_purge60_40_707_sample_position5_crashed`

- `150717_0926_purge60_40_707_sample_position5`

- `150720_2038_purge60_40_707_nosample_position5`

- `150722_2116_purge60_40_707_sample_position4`

- `150723_2113_purge60_40_707_nosample_position4`

- `150724_2120_purge60_40_707_nosample_position1`

- `150726_2230_purge60_40_707_nosample_position2`

- 150727_2230_purge60_40_707_nosample_position6

- 150728_1830_purge60_40_707_nosample_position3

A.2.3. Analysis and Results

This study will focus on the rising edge of the pulses here, but parallel studies analyzed the curvatures at the top and beginning of the pulse as well as the negative slope of the pulse after t_R.

In order to focus on the shape instead of the normalization, the baseline noise was subtracted from each pulse. Then a straight line was drawn between the points that corresponded to 20% and 90% of the pulse height [132]. The purpose of choosing these percents was to avoid the curvature at the top and bottom of the pulse. The underlying assumption is that a good pulse originating from the sample should be linear like Fig. 1.6, whereas a curved pulse like Fig. 1.7 is assumed to be from the tray. Three different methods were used to discriminate between pulses based on their shape.

A.2.3.1. Method 1: Rising Slope

The first approach was to linearly fit the pulse from 20% to 90% the pulse height and extract the slope. The results are plotted in Fig. 1.8. However, this will not account for any amount of curvature and would give the same result as a linear rise. In hindsight, a better approach would be to find the slope at the 20% point; this should show a stronger deviation for a nonlinear pulse.

155
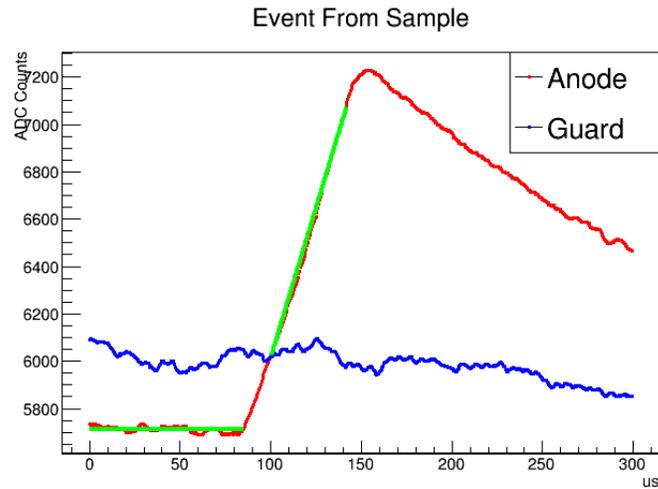
Figure 1.6: An example of linearly fitting a straight pulse from 20% to 90% the pulse height. Also seen here is the baseline fit from 0 μs to about 80 μs. The red line is the signal that comes from the anode electrode and the blue line is from the guard electrode. The lack of any identifiable signal on the guard electrode indicates that this is an event that occurred near the middle of the counter tray; in this case it is from position 1 from the middle of the sample.
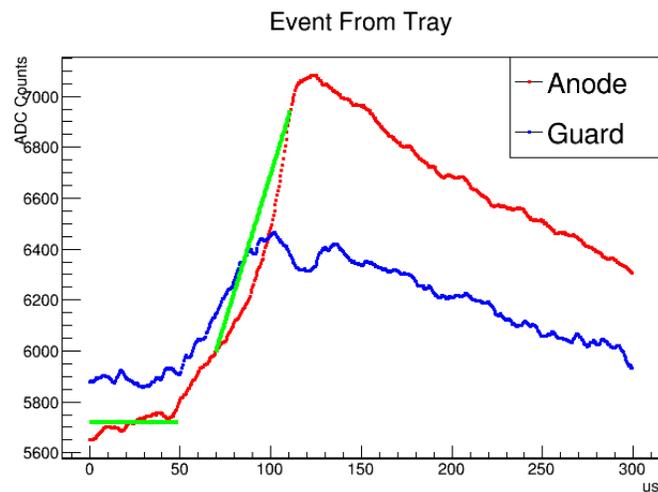
Figure 1.7: An example of linearly fitting a curved pulse from 20% to 90% the pulse height. Also seen here is the baseline fit from 0 to about 50 μs. The red line is the signal that comes from the anode electrode and the blue line is from the guard electrode. Seeing a larger signal on the guard electrode indicates that this is an event that occurred at higher radius on the counter tray; in this case it is from position 7.
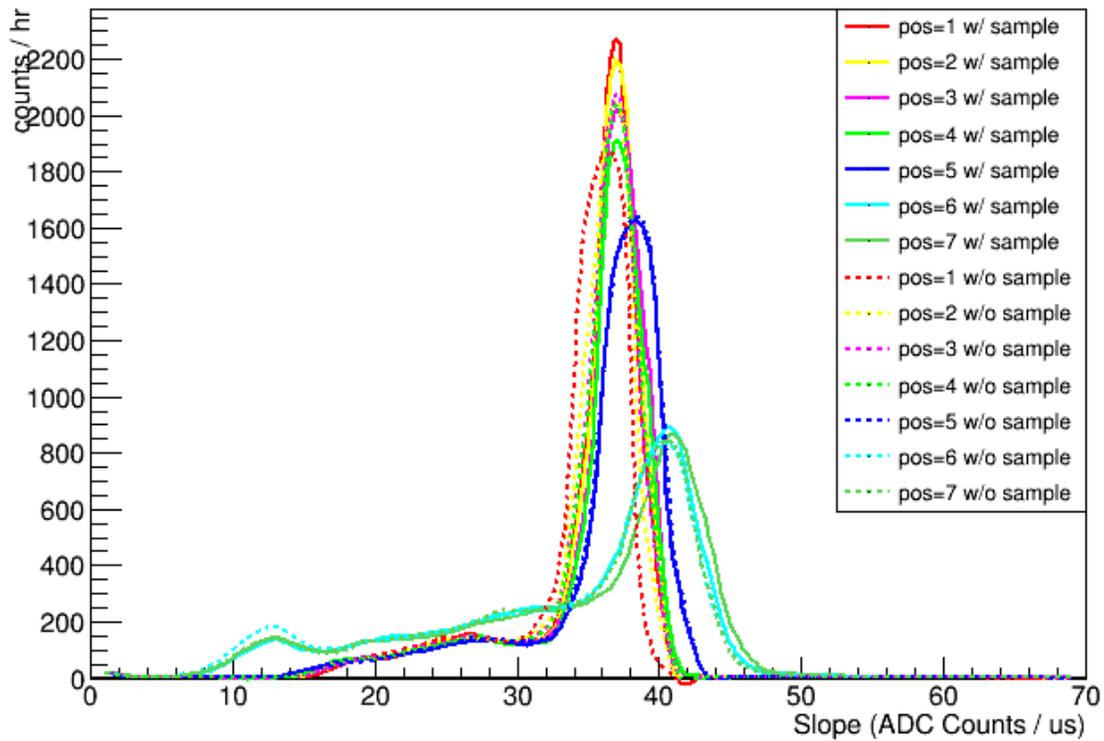
Figure 1.8: This is a smoothed histogram of the rising slope of each pulse. Each color corresponds to a different position. Solid lines are with the sample present and dotted lines are without the sample.

The second approach was to integrate the pulse shape from 20% to 90% the pulse height using a built-in function in ROOT [133]. When taking the integral of a TGraph, ROOT will connect the first point to the last point and find the area enclosed; by definition this is inherently a positive value. This method gives an area that approaches zero for a linear rise and a larger integral for a larger amount of curvature. This is useful for gauging how much curvature is in a particular pulse, but does not provide any information as to whether the pulse is convex or concave. The results from this integration are shown in Fig. 1.9.
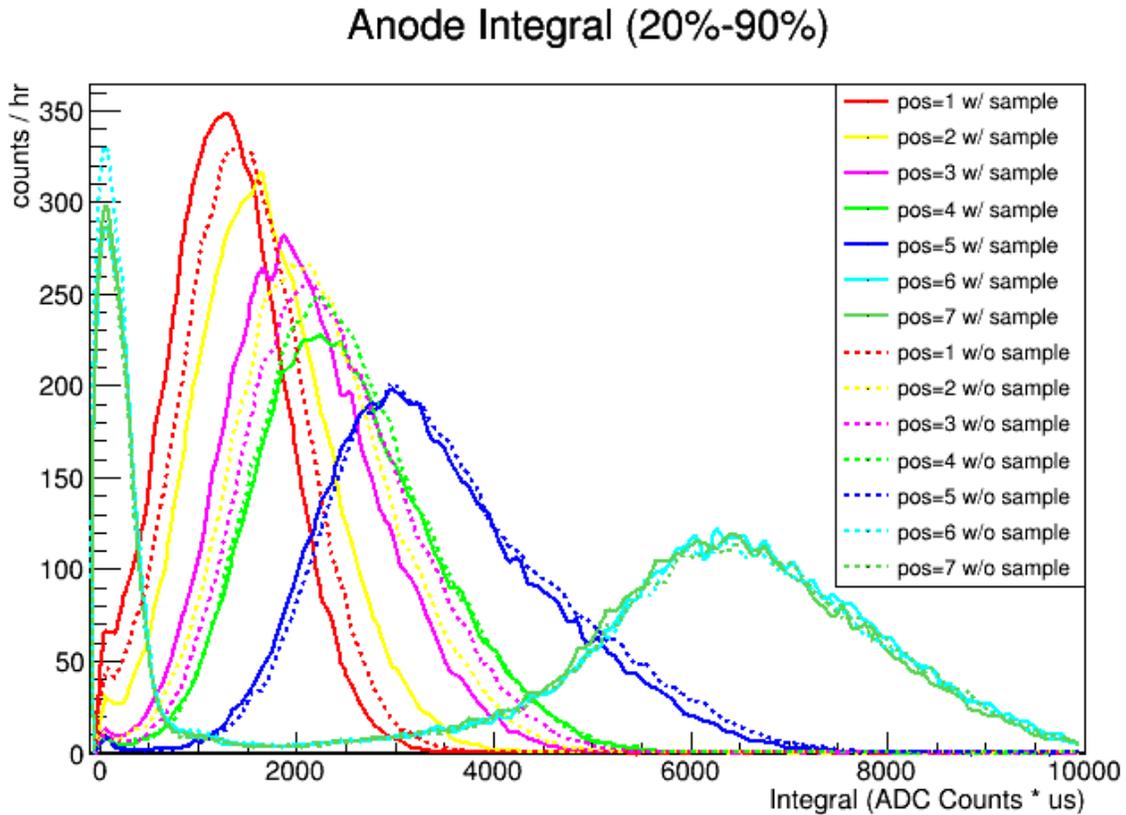
Figure 1.9: This is a smoothed histogram of the integral of the anode using ROOT integration. This always returns a positive integral, which is indicative of the nonlinearity of the pulse but not the concavity.

A.2.3.3. Method 3: Integral Difference

The third approach was to manually calculate the area enclosed between the pulse and the linear fit. This was done by using trapezoidal approximations to integrate under each curve and take the difference. This should sound the same as the method above, except that it gives negative areas if the pulse is convex. Therefore, this method has the most discrimination power of the three. The results are shown in Fig. 1.10.

It appears that placing a cut at 0 would successfully remove most of the events from positions 6 and 7 and some of the events from position 5. This is a promising start to a discriminator because it keeps all of the events from positions 1-3 that originated from the sample. However, more work is needed to refine this approach because the events from position 4 are not from the sample yet still pass this cut. As seen in Fig. 1.5, position 4 is closest to the sample without being on it, so it comes as no surprise that this is the most difficult to differentiate from the sample events.

A.2.4. Conclusion

As current progress stands, it is not clear how to differentiate between a pulse that originated from the sample and one that originated from the tray. I was hoping to see a clear dichotomy either between the positions on the sample (1-3) versus those on the tray (4-7), or for a position with and without the sample present.

As it turns out, there is very little discrimination power from the slopes since all of the distributions peaked between $35 - 40$ ADC/μs. Unsurprisingly, the anode integral and integral difference distributions are similar, except for a relative absolute value between them. It appears that the peaks of these distributions follow more of a continuum across the seven positions instead of separating into two distinct populations. There is not a clear value in any of the calculated quantities to place a cut. However, despite being a simple study that
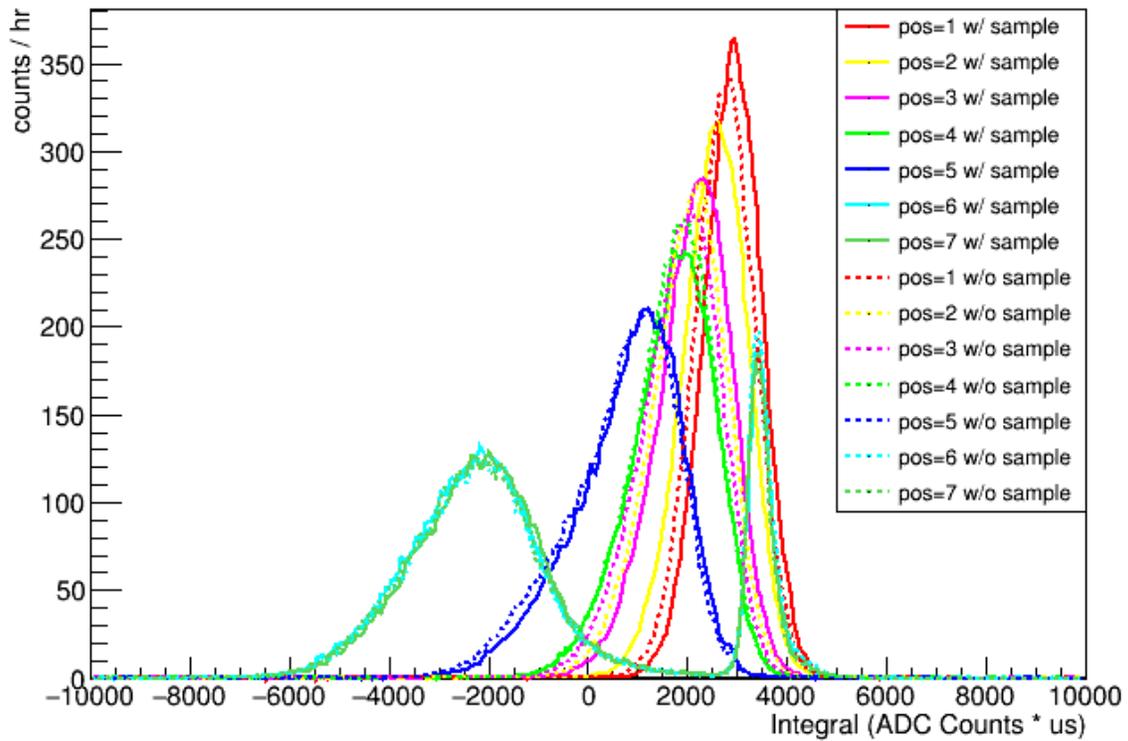
Figure 1.10: This is a smoothed histogram of the integral difference defined as the area under the pulse minus the area under a linear fit. It is similar to the anode integral, but accounts for both positive and negative curvature of the pulse.

could benefit from more powerful analysis tools, there is a small glimmer of hope that a pulse shape discriminator could eventually work, especially with Method 3.

One clear conclusion that can be drawn from this study is that there is not a significant difference between the solid and dotted lines in each of Fig. 1.8, Fig. 1.9, and Fig. 1.10. This means the presence of a conductive sample in the electric field is negligible. This conclusion is verified by simulations of the electric field done in a separate study [134]. A 2-dimensional slice of the electric field from that study can be see in Fig. 1.11.
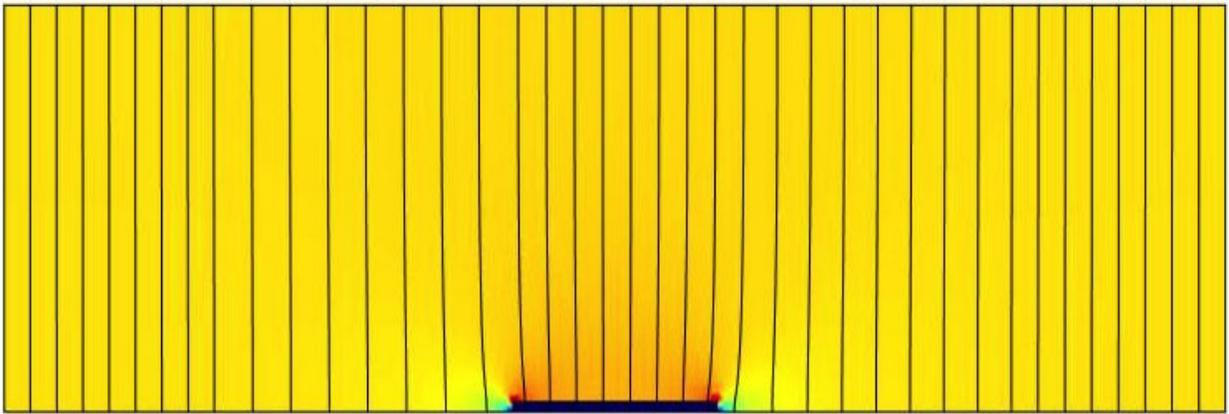
Figure 1.11: Simulated electric field in the XIA with field lines and field density gradient [134]. If you look very closely, you can see the vertical black lines bend slightly inward towards the edge of the sample. This simulation confirms the minimal effect the conductive sample has on the electric field in the alpha counter.

A.2.5. Future Work

One interesting feature in need of more exploration is that positions 6 and 7 curves show two populations in all three calculated quantities. It would require looking at individual pulses to see what might be causing this, but the farthest positions lose some energy to the guard electrode so it is possible an energy cut removes this feature.

There are a few other ideas that could build upon this study.

- As mentioned above, Method 1 should be changed to be the slope at 20% instead of the average slope from 20% to 90%.

- Using another ROOT function, a TProfile should be made for each position. TProfiles can be used to average all pulse shapes to smooth individual fluctuations, which will reveal any systematic differences between the positions.

- Applying quality cuts to the data should clear up the distributions, specifically an energy cut to select events expected from the source.

- Repeat the experiment with a different sized sample. It would be interesting to see if the pulse shapes actually depend on whether they originated from the sample or if they simply depend on the radius of the source location.

Here are also some ideas to improve the sensitivity of the XIA alpha counter:

- Use a tray liner with lower activity than the conductive teflon. We are currently experimenting with ultra-low activity electroformed copper made by SuperCDMS collaborators at PNNL. The biggest drawback is they are thicker than the teflon and the XIA only has 0.25 inch clearance to begin with, so the copper liners will limit the thickness of samples that can be measured.

- Active muon veto panels around the XIA could be used to trigger incoming cosmogenic events so they can be removed using anti-coincidence.

- Although the LUMINA laboratory at SMU is in a basement with three stories above it, moving the XIA to operate underground would greatly reduce the cosmogenic background rate. Unfortunately, this type of space does not exist at SMU (that I know of). Fortunately, a SuperCDMS collaborator at SNOLAB has recently acquired an XIA UltraLo-1800 and plans to eventually move it underground, which will be beneficial for screening samples with better sensitivity than the XIA at SMU.

162

RADON EMANATION AT PNNL

As dark matter direct detection experiments are pushing technological limits to explore smaller dark matter masses and cross-sections, they will require assays of ultra-low radon levels to mitigate sources of background. Commercially available devices do not have the requisite sensitivity, thus the need arises to build a custom system to measure radon emanating from materials. This appendix will describe a new radon emanation system built and commissioned at Pacific Northwest National Laboratory (PNNL).

B.1. Radon Emanation Systems

Radon emanation systems are designed to capture the radon that was emanated, or emitted, from a particular sample and to then pass those radon atoms to a sensitive alpha detector. The basic steps of operation of a typical radon emanation system are as follows:

1. Put the sample in a vacuum chamber, rinse it with low-radon gas, and evacuate to low pressure.

2. Seal the chamber and allow the sample to emanate for ~ 1 week, which should allow the radon to reach equilibrium.

3. Open the chamber, use a carrier gas to transfer the radon atoms to a radon trap, and pump away the rest of the gas from the chamber.

4. Release the radon atoms from the trap and load them into the detector.

5. Connect the detector to a data acquisition system to detect alpha decays from the radon.

B.1.1. The PNNL System

Figure 2.1 shows a schematic of the custom radon emanation system at PNNL. The primary flow path begins with a cylinder of compressed helium that is used as a carrier gas. There are two vacuum chambers of different sizes that can be connected to the system for emanating samples: a "small chamber" (0.3 L) that can be seen in Fig. 2.2, and a "large chamber" (35 L) that is more practical for containing assay samples. Downstream of the emanation chamber is a silica dessicant loop to remove any water that may have outgassed from the sample. Following the silica loop is a cryogenic radon trap made from a coil of narrow tubing that bottlenecks the gas flow and is submerged in a liquid nitrogen dewar. The melting point of radon is 202 K, so the 77 K liquid nitrogen is cold enough to freeze the radon inside the tubing. A gas commonly used for proportional counters, called P-10 (10% methane, 90% argon), is used to load the radon from the trap into the detector. There are two pumps to control the gas flow through the panel: a primary roughing pump for evacuating larger volumes of gas, and a secondary turbomolecular pump for achieving a high vacuum. There is also a Residual Gas Analyzer that can be used to determine if the system is leaking and room air is seeping in.

In order to intentionally insert radon into the system, a 100 Bq calibrated Pylon flow-through source can be connected to the panel via quick-connect fittings. The source consists of ^{226}Ra powder that is contained by a fine mesh. The radium decays to gaseous ^{222}Rn, which can be carried away by gas flowed through the 0.1 L source. Quick connect adapters were added to both sides of the custom gas flow panel, shown in Fig. 2.2. A port on the left side of the system allow radon to be pushed through the panel as if it had emanated from a sample, while a port on the right side of the system provide the ability to load radon directly into the detector. Both of these connections are important for efficiency measurements where the amount of radon detected is compared to the initial activity.
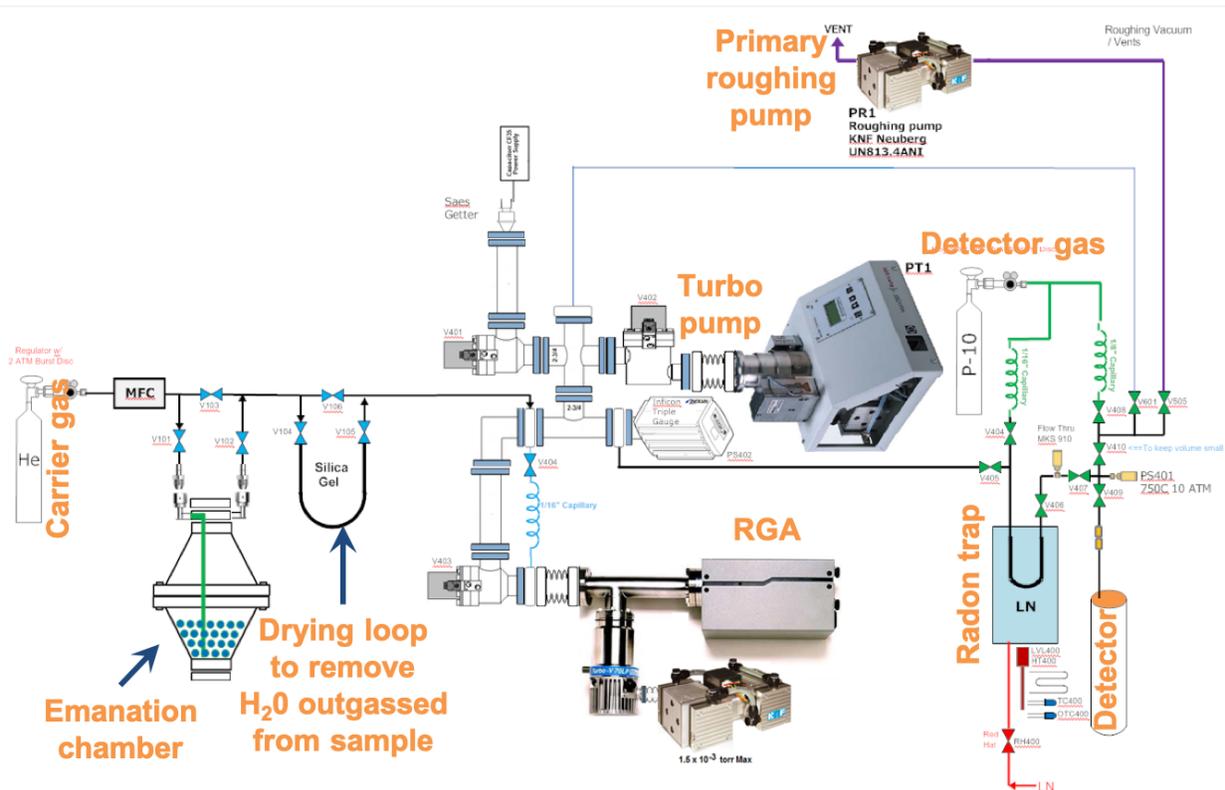
Figure 2.1: Schematic view of the radon emanation system at PNNL. The primary flow path travels from left to right from the He carrier gas, to the emanation chamber, to the silica drying loop, through the cryogenic radon trap, and into the detector.
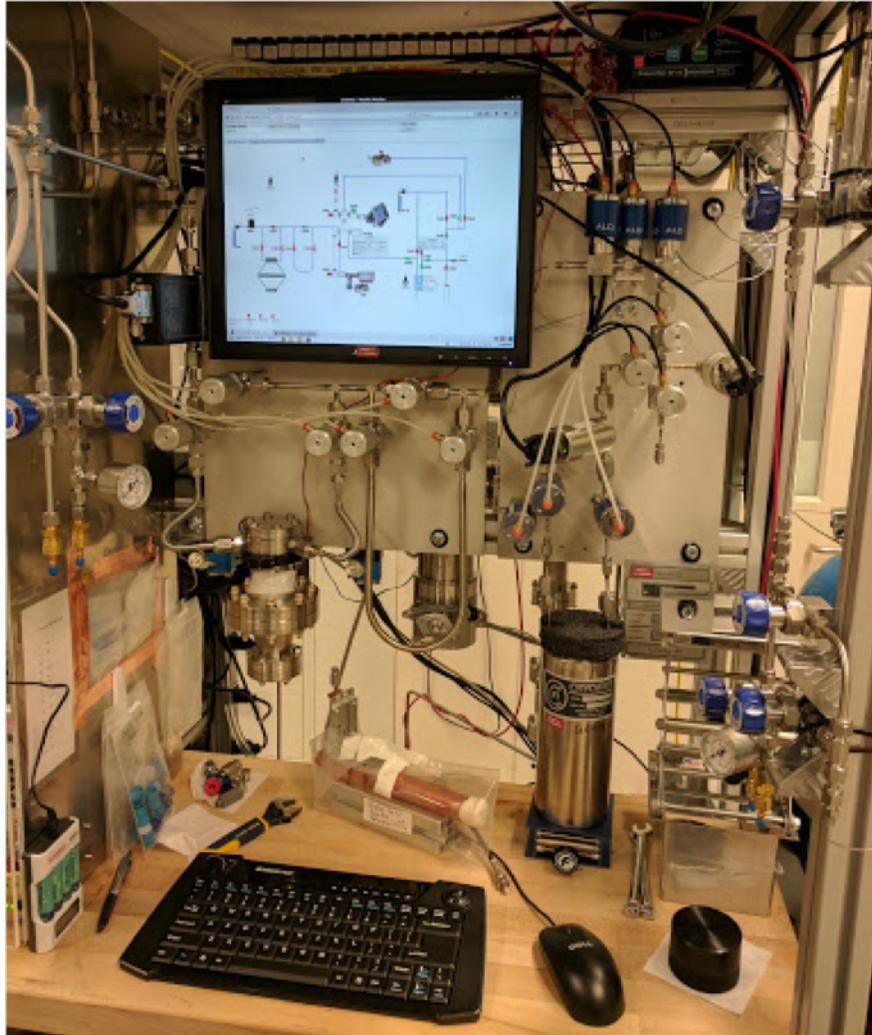
Figure 2.2: Picture of the front panel of the radon emanation system at PNNL. The blue knobs on the left and right side of the panel indicate the ports for connecting the calibrated radon source, which is the black canister in the lower right. The silver device hanging below the monitor on the left side is the small emanation chamber. The tall silver dewar sitting on the table is filled with liquid nitrogen and the radon trap is coiled inside. The prototype ULBPC is the copper tube behind the keyboard. Except for the manual blue knobs, all of the valves on the central portion of the panel are pneumatically controlled by the computer.

B.1.2. ULBPC

The radon emanation system at PNNL is the only such system to utilize an ultra-low-background proportional counter (ULBPC), shown in Fig. 2.3b. The ULBPC was designed by experts at PNNL to have a higher detection efficiency and less background than the silicon diode detectors typically used with radon emanation systems [82]. The ULBPC was traditionally used for detection of low-energy betas and x-rays from argon, but they proved to be adept at detecting high-energy alphas from radon as well.

(a) Pictured is the prototype ULBPC used for most of the measurements shown here. There was a small leak, so the end of the tube was dipped in TorrSeal, which led to the prototype getting its nickname "Marshmallow." It does not have the block of copper shielding that new ULBPCs are made with, so it is essentially just an electroformed copper tube strapped to a piece of aluminum framing.
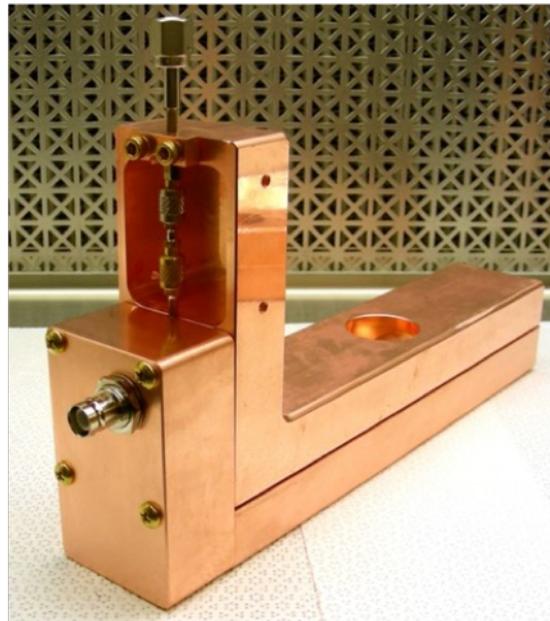
(b) Pictured is a typical ULBPC produced at PNNL [82]. It is quite a bit heavier than the prototype due to the solid block of copper shielding around the inner tube. The rectangular shape allows it to slot right into the lead-shielded counting array in the underground low-background facility.

Figure 2.3: The two ULBPC detector types that were used with the radon emanation system.

There were two ULBPCs used with the PNNL radon emanation system. The first was an early prototype (named "2-2") that had minimal shielding. The lack of shielding meant this

detector had higher background rates and was not sensitive to small amounts of radon, so there were no qualms about contaminating it with high levels of radon to perform calibration and efficiency measurements. The second detector (named "2RO1") was newer, had more shielding, and could be run underground in a low-background environment [135]. All of the data in this appendix was taken with the prototype detector, except for the final spectrum in Fig. 2.6.

B.1.3. Data Acquisition

The data acquisition system consists of the detector, a preamplifier, an in-line attenuator, and a digitizer that connects to a computer. In order to determine the optimal voltage at which to operate the detector, a range of bias voltages was tested from $1665 - 2250$ V. As the voltage increased so did the gain, so the attenuation had to be increased in order to keep the energies within an energy scale that could be detected by the Multi-Channel Analyzer (MCA). The amount of attenuation needed was determined by placing an ^{241}Am gamma source on top of the detector at the beginning of every run. The attenuation was chosen so that the 60 keV x-ray emitted from the source was peaked at channel 500 (out of 7×10^4 channels) of the MCA. This ensured a comparable level of gain between runs, but the increased attenuation resulted in a compressed energy range for higher voltages. Therefore, the energy scale needed to be manually calibrated as a function of voltage using known features in the spectra.

B.1.3.1. Energy Calibration

Calibrating the energy scale can be summarized in a few basic steps:

1. Make a histogram of the MCA channels for a particular voltage, and fit each peak in the spectrum with a Gaussian function to find the mean. Peaks corresponding to the

decay of ^{222}Rn, ^{218}Po, and ^{214}Po were clearly visible when the detector was loaded with radon from the source. An example of this is shown in Fig. 2.4a.

2. Plot the mean channel for each peak versus the energy it is expected to have, as shown in Table 2.1. The ^{241}Am peak was included to constrain the calibration at low energy.

3. Fit a line to the data and extract the slope and intercept, as shown in Fig. 2.4b. This linear function provides the calibration for any given energy, but only for this particular voltage. Other functional forms were tested as well, but the linear function was found to be the most accurate model.

4. Repeat the previous steps for each bias voltage and extract all the slopes and intercepts. As seen in Fig. 2.4c, the slopes and intercepts each appeared to follow quadratic distributions as a function of voltage. By fitting this data, the slope and intercept can be calculated for any voltage and those values provide the energy calibration.

After energy calibration was complete, it can be seen in Fig. 2.4d that the energy resolution degrades for higher bias voltages, so a lower voltage bias was chosen for operation. However, the detector cannot operate below 1665 V because the 60 keV peak from ^{241}Am falls below the detector threshold. All of the remaining data runs shown in this paper are collected with the prototype and new detector operating near 1665 V.

B.2. Energy Region of Interest

After calibrating the energy scale for any detector bias voltage, it is important to measure the intrinsic background rate of the prototype detector. In order to do that, an energy region of interest (ROI) must be defined for the analysis.

169
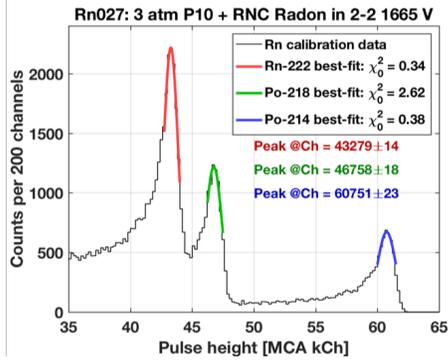
(a) Gaussian distribution fit to each of the radon alpha decay peaks. The plot shown is for the lowest voltage tested (1665 V). The red curve corresponds to the ^{222}Rn peak, the green curve corresponds to the ^{218}Po peak, and the blue curve corresponds to the ^{214}Po peak.
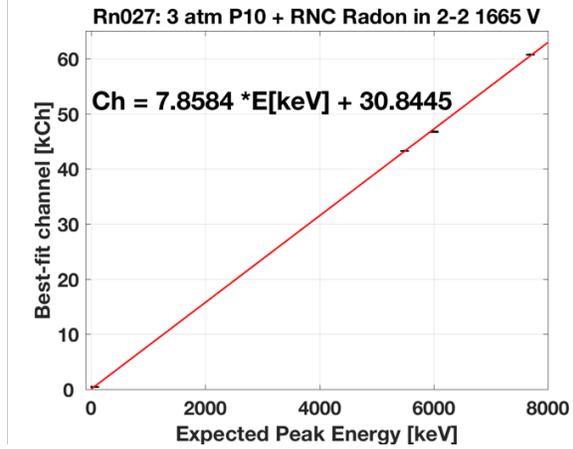
(b) Linear fit correlating measured MCA channel to expected energy for each peak. The plot shown is for the lowest voltage tested (1665 V).
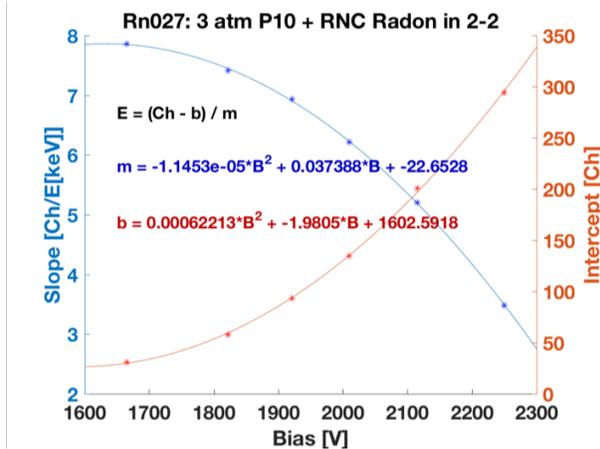
(c) Plot of extracted calibration parameters from the linear fit for each voltage. The blue curve is the slope of the fit in Fig. 2.4b, and the red curve is the intercept, from each bias voltage.
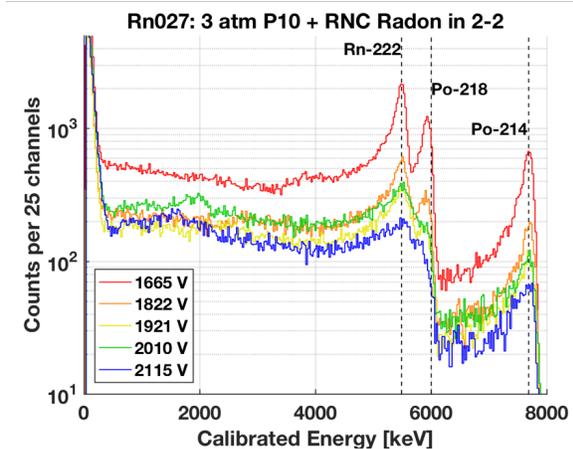
(d) Aligned post-calibration spectra for each bias voltage. The vertical dashed lines indicate the expected energies for each alpha decay.

Figure 2.4: Illustration of each step of the process to calibrate the energy scale.

170

Figure 2.5 shows an energy histogram of the detector background taken with the prototype detector. The first analysis range considered was $1 - 8$ MeV in order to avoid the low-energy noise wall, and the background rate was measured as 10 events/day. One major contribution to this background rate is the falling background shape from running the detector above ground and being exposed to cosmogenics. Another background contribution is the small hint of a ^{210}Po peak that is indicative of long-lived ^{210}Pb contamination. This peak will grow as the detector is sometimes loaded with high levels of radon for efficiency measurements. In fact, over the course of a few months of tests and data taking, the background rate doubled in the region of the ^{210}Po 5.3 MeV alpha.
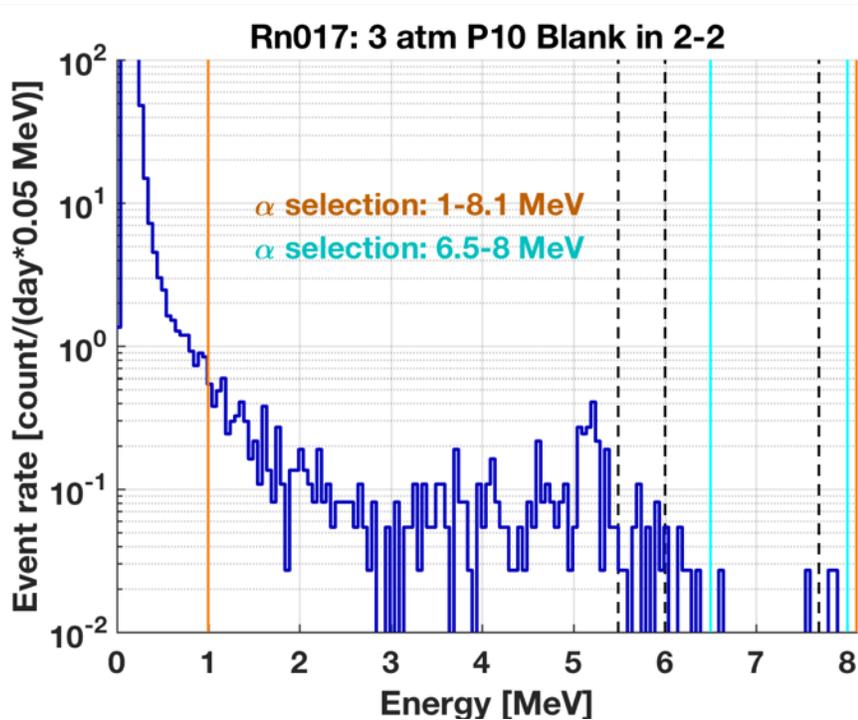
Figure 2.5: Histogram of prototype background spectrum. The vertical orange and cyan bars indicate the ROI bounds, while the vertical black dashed lines correspond to the expected energy of the ^{222}Rn, ^{218}Po, and ^{214}Po alpha decays, respectively.

A background of 10 events/day is much too high for proper commissioning tests, so a cleaner part of the spectrum was considered from $6.5 - 8$ MeV because it encompasses the

^{214}Po peak. This gives a background rate of 0.1 events/day. By avoiding the ^{210}Po peak, this rate also remains constant after months of loading the detector with radon.

B.2.2. Detector Efficiency

Since the analysis ROI is so restricted, it is natural to wonder how this effects the detection efficiency. In order to measure the efficiency, a detector is loaded with a known amount of radon from the calibrated source, the event rate is calculated as a function of time, and then the rate is extrapolated back to the beginning of the run when the radon activity was known. In comparing efficiency of the two energy regions, the efficiency from $6.5 - 8$ MeV is about 10 times lower than $1 - 8$ MeV. However, there is still enough of a signal to work with while commissioning the emanation system.

B.3. Sensitivity

When characterizing a system such as this, it is critical to determine the sensitivity of the system. This requires accurate measurements of the total background rate and the overall efficiency. There could be radon emanating from anywhere in the system, but the most concerning sources are the chamber while emanating a sample, the panel while transferring radon into the trap, and the detector when counting the radon alpha decays. The overall efficiency is also needed to calculate the sensitivity, which is the product of efficiencies associated with each step of the assay procedure described in Appendix B.1.

B.3.1. Small Chamber Background

The background rate of the small chamber was measured by closing it and letting it emanate, but the rate was shockingly 50 times higher than the detector background. It was possible that this was coming from the panel during the transfer from the small chamber, so the panel was isolated by doing a "zero-emanation" run, and background was still 30 times

high. This number cannot be directly subtracted from the original rate to say the chamber is 20 times higher than the detector background because of different counting times, but there was clearly still a source somewhere on the panel. On the back of the panel is a triple filter that is filled with beads of unknown material for filtering the helium as it flows into the panel. After bypassing the helium filter, the silica drying loop, and the small chamber the nominal background rate of 0.1 events/day was recovered. This was important because it demonstrated the ability to identify and remove sources of undesired background.

It is also not obvious that the filter or silica loop are needed for assays, so the system is still fully functional without them. However, the system does not work without an emanation chamber, so the small chamber was reconnected and emanated for 18 days. A relatively small rate of 0.45 events/day was measured. This was not a worrisome rate, but the 3 events in the ROI showed suspicious timing, so there was reason to believe the actual rate might be lower than this.

B.3.2. Large Chamber Efficiency

The large chamber is more practical for assaying samples, so the total background rate and efficiency are needed to be measured to calculate the sensitivity of the system with the large chamber attached. There are multiple stages throughout the emanation procedure that affect the overall efficiency of the system. These include:

1. Transferring the radon to the trap

2. Trapping the radon

3. Loading the radon into the detector

4. Detecting the radon decays

In order to measure the overall efficiency of the system with the large emanation chamber, a known amount of radon was loaded into the chamber. Since the volume of the large chamber is 350 times bigger than that of the radon source, a modified procedure was needed to load radon into the large chamber. The pylon source was closed and allowed to build up radon for a predetermined length of time so that the initial activity is known. Then helium is used to push the radon from the source to the chamber, followed by the normal transfer procedure as if the radon in the chamber had emanated from a sample. This process was repeated twice, resulting in a livetime-weighted average efficiency of $8.08 \pm 0.14\%$.

B.3.3. Large Chamber Background

The final piece needed to estimate the sensitivity is the background of the system with the large chamber connected. The chamber was emanated for 18 days and 3 events appeared in the ROI while counting for 5.5 days. As with the small chamber background run, these events portrayed very suspicious timing behavior. Typically, an exponentially falling rate is expected where most of the events are detected quickly. Instead, with this run there were 0 events for 2.4 days before the 3 events were detected in the last 3 days. It was hypothesized that this behavior was caused by radon getting trapped inside the detector from previous high-radon efficiency runs and diffusing out slowly during subsequent background runs. To test this theory, the previous measurement was repeated after waiting > 1 week. This was enough time for any carryover radon to undergo two half-lives, which is a reduction by a factor of four. This resulted in only 1 event in the ROI with a more believable time signature during 5.7 days of counting. Due to limited statistics, it is difficult to extract much information from a single event, so 100% uncertainty was assumed to calculate an upper bound. Even with this conservative uncertainty, the blank rate of the large chamber was determined to be < 100 μBq or < 50 radon atoms in equilibrium.

This level of sensitivity is comparable with the best radon emanation systems in the United States [136]. It is worth emphasizing that this was achieved with the prototype detector, operating above ground, and using a restricted energy range.

B.4. Outlook

There is reason to believe a superior sensitivity can be achieved by using the new detector in the shallow underground laboratory low-background counting facility [135]. Figure 2.6 shows a calibration and background run taken with this setup. In order to avoid contaminating the new detector, the calibration sample loaded was much smaller than those loaded into the prototype, but it was still enough to see the alpha peaks. The background spectrum is much cleaner than the prototype since the low-energy background from running at the surface is gone and there is no ^{210}Po peak from long-lived contamination.
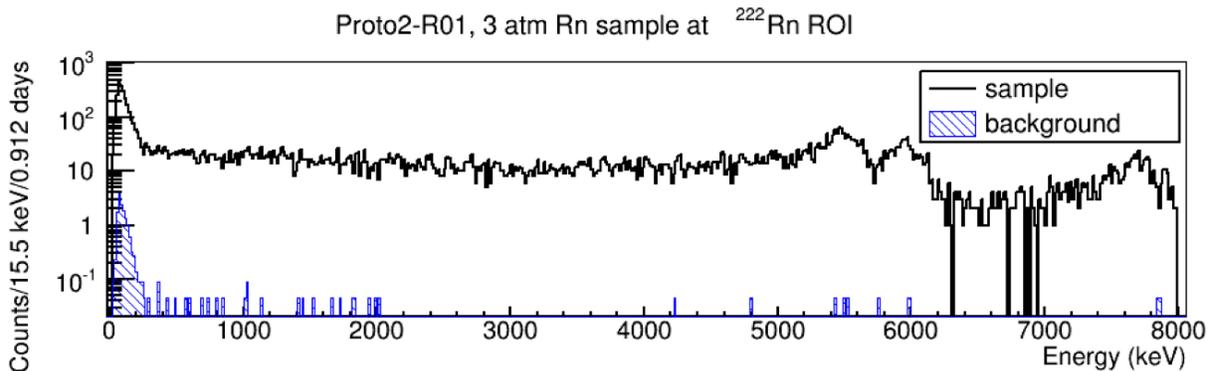
Figure 2.6: Overlaid spectra taken with the new detector in PNNL's shallow underground low-background counting facility. The black calibration data was from loading radon into the new detector, and the blue data is the background of the detector. Notice how much cleaner the blue spectrum is from $1 - 6$ MeV as compared to Fig. 2.5.

The broad energy range from $1 - 8$ MeV results in a measured background rate of 1.07 ± 0.27 events/day. However, the background dwindles around 2 MeV, which suggests that $2 - 8$ MeV is a more optimal energy range. Considering this energy range, the background rate is cut to 0.56 ± 0.21 events/day without a significant loss of efficiency.

175

B.4.1. Future Plans

To begin with, the assay procedure should be optimized. It is possible that the transfer process is not optimal and can be improved by oscillating the flow rate through the chamber to ensure all of the radon atoms are removed. The reproducibility of results can also be improved by automating the transfer procedure.

Once the assay process is improved, more accurate background measurements can be taken to determine the sensitivity of the system using each chamber. A longer will be required to bolster the low statistics of the background runs to calculate the sensitivity.

Thus far only the overall efficiency of the system has been measured, but it might be possible to disentangle the detection efficiency of each alpha peak by using timing information from the decay chain to remove correlations. This will lead to better understanding of the ULBPC performance for detecting high energy alpha decays.

Finally, the PNNL radon emanation system can be used for assaying materials that rare-event search experiments are interested in using. This is particularly relevant to wetted parts of noble liquid experiments, such as LZ, because they are in direct contact with the active volume of the detector [137]. It is also relevant to long-term storage of solid state detectors, such as those used by SuperCDMS, in nitrogen purge boxes because the rubber gaskets used to seal the purge box doors tend to emanate radon. A lot of caution is taken to avoid radon contamination during fabrication, so storing the detector in a sealed box with a radon source would defeat all of those precautions.

SOUDAN IZIP RESOLUTION

Described in Section 3.1.2, the SuperCDMS experiment at Soudan operated 15 germanium detectors that measured ionization and phonon energy deposited by incident particles. In an effort to completely understand the response of each unique detector, it is critical to determine the energy resolution as a function of energy. This is done by measuring the width of known mono-energetic features in the data across a wide range of energy. Then a model is fit to those measurements to obtain the energy resolution as a function of energy for each detector.

Before describing the following study, I must disclose my pet peeve with respect to resolution. When a resolution is better it is said that it increases, but the width of the features actually decreases. Therefore, to avoid confusion, I will refrain from using terms implying size such as increase/decrease and bigger/smaller, and will instead use terms describing performance such as improve/degrade and better/worse. I wish that everybody would adopt this practice.

C.1. Data Used

This study primarily used iZIP data collected between March 2012 and July 2013 (R133), though the type of data depended on which peak was being measured. Each type of data is defined in Section 3.1.3. The low-bg data was used to find the 0 keV, 10.37 keV, and 511 keV peaks. Barium calibration data was used to fit the 356 keV peak. There were not enough statistics in the R133 californium calibration data to clearly identify the 66.7 keV peak, so R134 data (collected between July 2013 and July 2014) was added.

The energy quantities in this study follow the nomenclature of Section 3.1.3.4:

qsummaxOF: max(qi1OF + qo1OF, qi2OF + qo2OF)

qsummaxOF0: max(qi1OF0 + qo1OF0, qi2OF0 + qo2OF0)

ptNF: phonon total pulse energy from summing the pulses in all channels and running the summed pulse through NF

ptOF0: same as above but fixing the start time (see quote below)

From an internal note, "A brief aside on the OF0 quantities vs the OF quantity: The OF0 quantities have a forced start time. This prevents the OF from performing a search for the time that maximizes the OF amplitude which will have a tendency to be biased towards a positive amplitude due to it picking up preferentially on large noise fluctuations. By fixing the start time then we expect the OF0 quantities for random triggers to reflect a Gaussian centered at 0 with a sigma corresponding to the resolution of the particular channel in question." [138, 139]

C.1.1. Cuts

A number of different cuts were used to locate the desired data features. All of the cuts used in this study were developed by the SuperCDMS High-Threshold (HT) analysis group [140, 141]. The cut names are appended with processing tags that are not relevant for this study, such as "_v53," "_133," and "_HT." The only cut applied to all five peaks in this study was the general quality cut, **cGoodEv_v53**.

The following list of cuts were used in various combinations to isolate the data features measured in this study.

cGoodEv__v53: umbrella quality cut to remove bad events, bad series, and poor quality data

cRandom__133: selects randomly triggered events to sample the electronic noise (known as "randoms")

cGoodRandomChi2__v53__HT: uses χ^2 value to remove physical pulses that accidentally show up in random events

cQin1__v53: charge fiducial cut selects the inner charge channel on side 1 of the detector to remove high-radius events

cQin2__v53: charge fiducial cut selects the inner charge channel on side 2 of the detector to remove high-radius events

cPostCf__133__HT: selects the data collected within 48 hours of exposure to the ^{252}Cf neutron source

cQsym__v53: charge symmetry cut defined in Section 3.1.2.2 to select the bulk of the detector

ytNF: not directly a cut, this is the yield quantity defined in Section 3.1.2.1, but a cut was used to select gammas with high yield

Except for the 0 keV peak that will be described below in Appendix C.2.1, a few basic quality cuts were applied to the other four peaks. Due to the shape of the electric field, there is a region at high radius where events do not experience the full NTL gain. These events were removed with the charge fiducial cuts **cQin1__v53** and **cQin2__v53** to prevent low-gain events from creating low-energy tails for each peak and artificially widening the resolution. Random events were also removed with **!cRandom__133**.

C.1.2. Detector Shorts

Each of the crystal detectors is unique and has its own strengths and weaknesses. Without diving into the fabrication history for all 15 detectors, it was important to know which detectors had shorted channels and which did not. Therefore, peaks found to be asymmetric, abnormally small, or mis-calibrated could be attributed to a known channel short. The following lists contain all of the known shorts on charge or phonon channels; some detectors are on both lists.

Charge Shorts:

- **iZIP 2 (T1Z2):** QIS1, QOS1

- **iZIP 3 (T1Z3):** QIS1, QOS1

- **iZIP 10 (T4Z1):** QIS1

- **iZIP 13 (T5Z1):** QIS1, QOS1

- **iZIP 14 (T5Z2):** QOS1

- **iZIP 15 (T5Z3):** QOS1

Phonon Shorts:

- **iZIP 2 (T1Z2):** PAS2, PCS2

- **iZIP 3 (T1Z3):** PAS1

- **iZIP 6 (T2Z3):** PBS1, PDS1

- **iZIP 13 (T5Z1):** PAS1

- **iZIP 15 (T5Z3):** PCS2

C.2. Fitting Peaks

There are too many plots to show the fits for all 15 detectors, so only iZIP 4 (T2Z1) will be shown as an exemplar. This was one of the better-behaved detectors without shorts, and it was the detector used for CDMSlite R3.

Each of the peaks below was expected to be Gaussian or nearly-Gaussian. Some rudimentary background modeling was included to get better fits to the data in the particular energy range of interest. However, the only parameters extracted from the fit were the mean and width of the desired peak, with their respective uncertainties. The width is defined as the full width at half the maximum (FWHM). Each of the following energy histograms has a maximum value at the mean of the Gaussian, so half of that maximum corresponds

to energies above and below the mean; the resolution is the difference between those two energies.

As explained in Section 0.2, the energy scales for charge and phonons are different because of the NTL gain. All of the peak energies are referred to in units of keV_{ee}, which matches the charge energy scale. The phonon energy scale is in keV_t, which is a factor of approximately 2.3 times higher for a bias of 4 V across the crystal.

C.2.1. 0 keV

The 0 keV peak is different from the other peaks because it is essentially a lack of physical pulse, and therefore had its own set of cuts to select empty noise events. The way SuperCDMS samples noise is by randomly issuing a trigger in the electronics without a physical impetus; this class of events are called "randoms". These random events were selected using **cRandom_133**, and **cGoodRandomChi2_v53_HT** was used to remove any events with abnormally high χ^2 valuess.

This peak can be seen in zero-delay quantities from in-run random triggers. These are sometimes referred to as "noise blobs" because they are random noise events that are amorphously centered around zero. Zero-delay quantities (appended with 0) were used to prevent a bias from the OF preferentially finding large noise fluctuations. However, the "max" algorithm selects the side with the larger energy so there is still a slight positive bias and the peaks are not centered exactly at 0 keV. This fit was a single Gaussian without constraints on any of the parameters. Figure 3.1 demonstrates the fit for the total charge and phonon energies on iZIP 4 (T2Z1), while Table 3.1 and Table 3.2 show the parameters extracted from the fits.

The first thing to point out for these fits is that the data is more sharply peaked than the Gaussian function. This is systematically true for all of the detectors, but it is not clear why
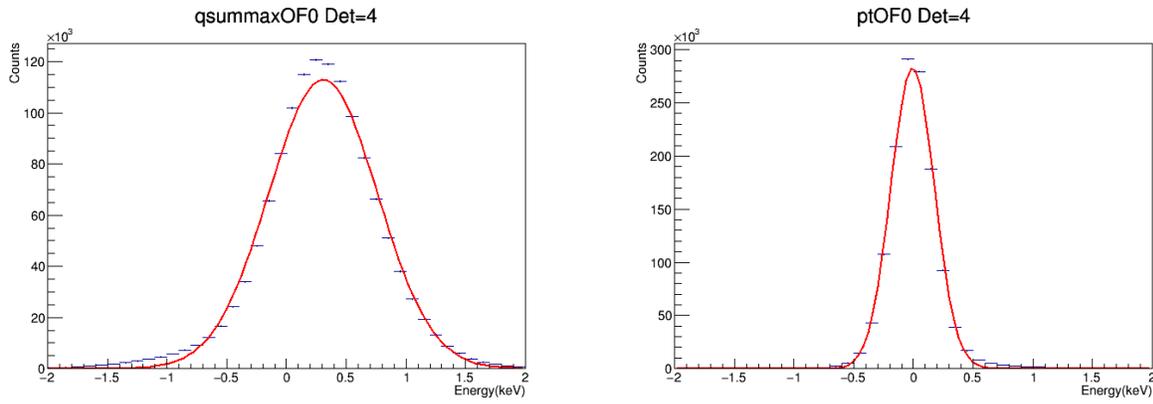
(a) Measuring charge resolution at 0 keV$_{ee}$ for iZIP 4 (T2Z1).

(b) Measuring phonon resolution at 0 keV$_{ee}$ (also 0 keV$_t$) for iZIP 4 (T2Z1).

Figure 3.1: The peaks appear to be more sharp than the Gaussian fit, which is systematically true for all of the detectors. Also, there is usually a tail at low energy, high energy, or both. It is not clear why the data does not follow a strictly Gaussian distribution.

Det	Mean [keV$_{ee}$]	Error	1-sigma [keV$_{ee}$]	Error	χ^2/NDF	% Mean Inaccuracy
1	1.95e-1	3.69e-4	4.08e-1	2.95e-4	372.02	N/A
2	2.18e-1	3.27e-4	3.58e-1	2.58e-4	415.43	
3	2.57e-1	4.27e-4	2.91e-1	3.59e-4	13275.65	
4	3.06e-1	4.01e-4	4.55e-1	3.59e-4	726.84	
5	2.44e-1	3.48e-4	3.89e-1	2.76e-4	519.03	
6	2.71e-1	4.54e-4	4.14e-1	3.45e-4	333.54	
7	2.00e-1	3.20e-4	3.04e-1	2.41e-4	460.04	
8	1.99e-1	3.36e-4	3.45e-1	2.61e-4	1088.53	
9	2.24e-1	3.34e-4	2.99e-1	2.53e-4	473.59	
10	2.20e-1	4.21e-4	3.71e-1	3.58e-4	1234.33	
11	2.37e-1	3.89e-4	3.57e-1	2.94e-4	602.84	
12	2.56e-1	3.35e-4	3.63e-1	2.48e-4	692.93	
13						
14	2.53e-1	4.17e-4	4.31e-1	3.17e-4	605.22	
15	2.38e-1	4.07e-4	4.31e-1	3.25e-4	972.10	

Table 3.1: Parameters extracted from fitting the 0 keV$_{ee}$ charge peak for all detectors.

Det	Mean [keV$_t$]	Error	1-sigma [keV$_t$]	Error	χ^2/NDF	% Mean Inaccuracy
1	-1.66e-2	2.37e-4	2.61e-1	2.12e-4	1179.23	N/A
2	2.50e-2	2.62e-4	2.87e-1	2.27e-4	944.76	
3	2.83e-2	2.40e-4	2.91e-1	2.04e-4	911.32	
4	-7.02e-3	1.60e-4	1.81e-1	1.38e-4	890.17	
5	2.74e-2	1.78e-4	2.03e-1	1.55e-4	828.67	
6	1.04e-2	3.50e-4	3.29e-1	3.15e-4	691.45	
7	-3.31e-3	6.04e-4	5.72e-1	5.63e-4	630.31	
8	-3.38e-2	3.53e-4	3.72e-1	3.29e-4	1456.47	
9	-1.34e-2	3.70e-4	3.32e-1	3.39e-4	1001.61	
10	-2.87e-2	3.34e-4	3.20e-1	3.09e-4	921.04	
11	-3.42e-2	2.97e-4	2.80e-1	2.82e-4	1275.79	
12	-1.83e-2	1.96e-4	2.19e-1	1.77e-4	1862.76	
13	-2.27e-3	7.61e-4	7.02e-1	6.91e-4	297.18	
14	1.91e-2	1.71e-4	1.86e-1	1.47e-4	954.30	
15	1.27e-2	1.88e-4	2.17e-1	1.54e-4	382.17	

Table 3.2: Parameters extracted from fitting the 0 keV$_{ee}$ (also 0 keV$_t$) phonon peak for all detectors.

they do not follow the expected Gaussian distribution. Regardless, these are very stable fits and the Gaussian does accurately represent the rising and falling sides of the peak. The only consequence of the sharpness is that the widths will be slightly wider because the FWHM is calculated at a lower height due to the lower maximum.

Regarding the tables, calculating the percent error for the mean not being located at zero results in errors approaching infinity, so this column was left blank. Also, the χ^2/NDF values are all very large because the fits do not match the peaks of the data where bins that disagree by thousands of events.

Despite not being shown in Fig. 3.1, other detector-specific results are:

- The charge channel on iZIP 3 has very few statistics and a large spike on the lower energy side, which is why its χ^2/NDF is so much worse. However, there is still a low energy tail on the other side of the spike that fits the expected Gaussian distribution so this detector isn't excluded from the results.

- The charge channel on iZIP 13 has no statistics for qsummaxOF0 < 0, so it was excluded altogether.

C.2.2. 10.37 keV

The 10.37 keV peak is also referred to as the germanium K-shell and its origin is described in Section 3.1.3.2. This peak is extremely useful in many analyses and calibrations because it is uniformly distributed throughout the crystal and has a high signal-to-background ratio. The 10.37 keV peak can be found in the data collected immediately after exposure to a ^{252}Cf neutron calibration source, which was selected using **cPostCf_133_HT**.

To accurately fit the peak, a linear background model and an exponential background model were included. A second Gaussian function was also included because there are

cosmogenic x-rays around 8.98 keV and 9.66 keV that can be close enough to create a low-energy shoulder to the K-shell [142]. For this fit, the mean and sigma of the two Gaussians were constrained to ensure they were finding the correct peaks; the linear and exponential background parameters were left free. Figure 3.2 demonstrates the fit for the total charge and phonon energies on iZIP 4 (T2Z1), while Table 3.3 and Table 3.4 show the parameters extracted from the fits.
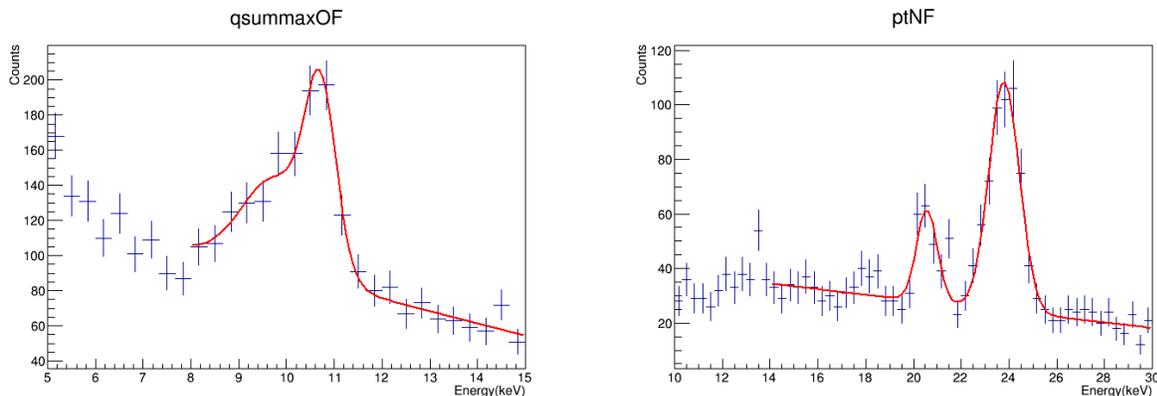
(a) Measuring charge resolution at 10.37 keV$_{ee}$ for iZIP 4 (T2Z1).

(b) Measuring phonon resolution at 10.37 keV$_{ee}$ for iZIP 4 (T2Z1).

Figure 3.2: This detector is a prime example where you can see how the NTL gain in phonon energy actually separates the K-shell from the cosmogenic x-ray peak.

Despite not being shown in Fig. 3.2, other detector-specific results are:

- Due to a phonon short on iZIP 3, the peak appears at $\sim 2/3$ the expected energy.

- Some detectors (such as iZIPs 1, 2, 4, 5, 12, and 14) have a larger phonon χ^2/NDF because there appears to be a third peak that is not accounted for in the fit (e.g. the bump in Fig. 3.2b at 18 keV$_t$). This is likely from the 8.98 keV cosmogenic x-ray getting separated from the 9.66 keV x-ray by the NTL gain. Another peak could have been added to the fit, but the important result is getting an accurate fit for the 10.37 keV peak and this other peak is separated enough that it was not expected to improve the K-shell fit. The other two peaks were not included in the resolution measurements because they did not reliably show up in all detectors, and they are close to 10.37 keV

185

Det	Mean [keV$_{ee}$]	Error	1-sigma [keV$_{ee}$]	Error	χ^2/NDF	% Mean Inaccuracy
1	10.61	0.05	0.43	0.04	1.53	2.46
2	10.53	0.05	0.45	0.04	1.17	1.66
3	10.55	0.04	0.52	0.04	1.55	1.80
4	10.71	0.06	0.33	0.01	0.69	3.42
5	10.50	0.02	0.43	0.02	0.88	1.33
6	10.29	0.21	0.79	0.14	1.17	-0.64
7	10.53	0.08	0.36	0.07	1.42	1.68
8	10.48	0.04	0.35	0.04	1.71	1.16
9	10.54	0.23	0.51	0.14	1.46	1.72
10	10.69	0.24	0.62	0.29	0.81	3.16
11	10.63	0.03	0.57	0.03	2.15	2.64
12	10.59	0.05	0.43	0.04	0.55	2.24
13	10.73	0.13	0.50	0.13	0.73	3.53
14	11.13	0.06	0.33	0.06	1.33	7.41
15	10.67	0.06	0.50	0.07	0.50	3.01

Table 3.3: Parameters extracted from fitting the 10.37 keV$_{ee}$ charge peak for all detectors.

Det	Mean [keV$_t$]	Error	1-sigma [keV$_t$]	Error	χ^2/NDF	% Mean Inaccuracy
1	23.98	0.05	0.67	0.05	1.59	0.64
2	24.73	0.05	0.84	0.05	1.30	3.80
3	17.55	0.06	0.55	0.07	0.93	-26.35
4	23.80	0.05	0.64	0.004	0.93	-0.13
5	23.90	0.02	0.51	0.02	3.14	0.29
6	21.71	0.43	1.78	0.33	0.65	-8.88
7	24.45	0.19	0.53	0.20	0.62	2.63
8	24.11	0.08	0.75	0.07	1.09	1.18
9	24.27	0.14	0.69	0.15	1.03	1.84
10	24.94	0.08	0.53	0.07	1.06	4.69
11	23.85	0.07	0.77	0.07	0.89	0.08
12	24.09	0.04	0.63	0.04	0.95	1.11
13	23.95	0.58	3.57	0.51	1.00	0.52
14	24.19	0.05	0.75	0.04	1.38	1.51
15	24.77	0.08	0.79	0.07	0.92	3.95

Table 3.4: Parameters extracted from fitting the 10.37 keV$_{ee}$ phonon peak for all detectors.

so they do not provide new information when attaining the resolution as a function of energy.

- iZIPs 8 and 12 have a large spike of events at the lowest energy on the plot. This changes the scale on the y-axis to make the peaks appear small, but comparing the number of counts confirms that they contain the same level of statistics as the other well-behaved peaks.

- The low-energy end of the fit function has a strange up-turn for iZIP 12. This is just an artifact of the boundary and can be ignored because the 10.37 keV peak is well-fitted.

- There is no discernible peak in the phonon spectrum for iZIP 13. This detector is consistently one of the poorest.

C.2.3. 66.7 keV

The 66.7 keV peak comes from a prompt gamma ray that is described in Section 3.1.3.2. It is found in the californium neutron calibration data, which is only collected every few months, so there were only four periods of data in R133: April 2012, August 2012, January 2013, and July 2013. This was not enough statistics for a clear peak to emerge, so the data was bolstered by including the R134 californium calibration data, which also had four periods: December 2013, February 2014, May 2014, and July 2014. The charge symmetry cut **cQsym_v53** was also used for additional fiducialization by removing events collected near the faces of the detector. The 66.7 keV gamma events are approximately uniformly distributed throughout the crystal, so selecting the bulk of the crystal provides a higher signal-to-background ratio.

To accurately fit the peak, a linear background model and an exponential background model were included. For this fit, the mean and sigma of the Gaussian were constrained to ensure that it was finding the peak; the linear and exponential background parameters

were left free. The data also had to be re-binned to make the peak more obvious against the background. This was the most biased peak search in this study, but it was useful to have a resolution measurement between the 10.37 keV and 356 keV peaks. Figure 3.3 demonstrates the fit for the total charge and phonon energies on iZIP 4 (T2Z1), while Table 3.5 and Table 3.6 show the parameters extracted from the fits.
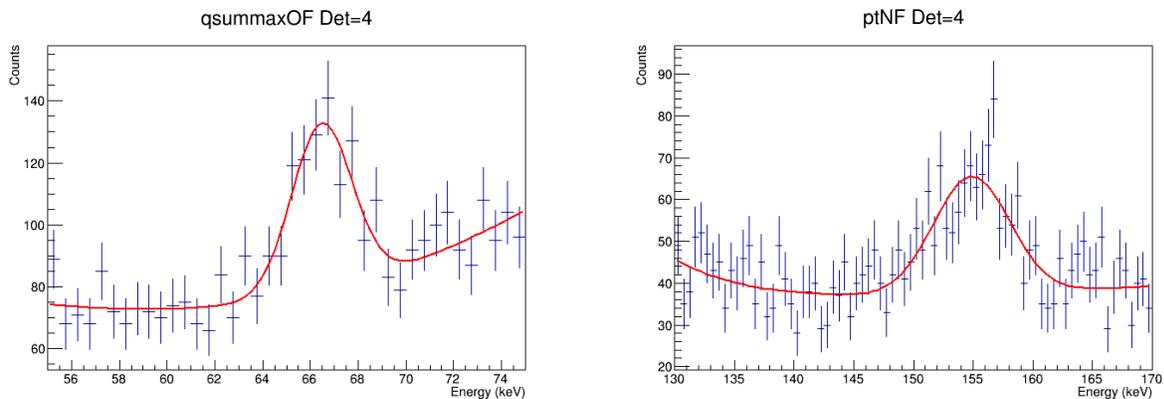
(a) Measuring charge resolution at 66.7 keV$_{ee}$ for iZIP 4 (T2Z1).

(b) Measuring phonon resolution at 66.7 keV$_{ee}$ for iZIP 4 (T2Z1).

Figure 3.3: This peak is the least stable fit as even on the best detector the signal is only tens of counts above background per bin.

In general, shorted channels seemed to have more of an impact on this peak than the others; perhaps this is because it has the lowest signal-to-background ratio.

Despite not being shown in Fig. 3.3, other detector-specific results are:

- Due to a phonon short on iZIP 3, the peak appears at $\sim 2/3$ the expected energy.

C.2.4. 356 keV

The 356 keV peak comes from the decay of ^{133}Ba that is used as a high-rate gamma source. It is deployed regularly, so there is plenty of data to work with as evidenced by the number of events and finer binning in the energy histograms. In order to select the 356 keV

Det	Mean [keV$_{ee}$]	Error	1-sigma [keV$_{ee}$]	Error	χ^2/NDF	% Mean Inaccuracy
1	66.49	0.23	1.22	0.29	1.35	-0.31
2	67.15	0.18	0.97	0.16	1.24	0.67
3	68.02	0.36	2.09	0.45	0.79	1.97
4	66.47	0.16	1.26	0.17	0.74	-0.34
5	66.68	0.19	1.19	0.16	1.05	-0.04
6	64.97	0.32	1.34	0.35	1.42	-2.60
7	67.64	0.15	0.85	0.16	0.82	1.41
8	67.36	0.30	1.27	0.26	1.00	0.99
9	67.13	0.22	1.22	0.22	1.15	0.64
10	68.05	0.80	2.67	0.94	0.63	2.02
11	68.09	0.15	0.36	0.33	1.31	2.09
12	67.27	0.18	1.03	0.20	0.94	0.85
13	68.24	0.65	2.56	0.96	1.19	2.32
14	67.70	0.46	1.32	0.30	0.91	1.49
15	68.58	0.32	0.81	0.26	2.20	2.82

Table 3.5: Parameters extracted from fitting the 66.7 keV$_{ee}$ charge peak for all detectors.

Det	Mean [keV$_t$]	Error	1-sigma [keV$_t$]	Error	χ^2/NDF	% Mean Inaccuracy
1	154.65	0.41	2.72	0.45	1.32	0.81
2	155.04	0.16	0.39	0.16	0.86	1.06
3	115.47	0.37	2.57	0.39	1.05	-24.73
4	154.87	0.36	3.20	0.43	1.05	0.95
5	153.64	0.34	3.12	0.39	0.92	0.15
6	150.00	9.03	4.75	1.85	0.73	-2.22
7	155.62	0.64	3.46	0.57	0.92	1.44
8	154.20	0.74	4.70	0.99	0.90	0.51
9	155.76	0.56	2.96	0.82	0.95	1.53
10	159.02	0.88	3.20	0.74	1.21	3.65
11	157.00	1.06	4.09	1.19	1.25	2.34
12	156.61	0.64	4.62	0.69	0.73	2.08
13	154.30	0.07	1.45	0.72	0.98	0.58
14	154.18	0.98	2.87	0.67	1.73	0.50
15	154.57	0.68	0.81	0.44	0.64	0.75

Table 3.6: Parameters extracted from fitting the 66.7 keV$_{ee}$ phonon peak for all detectors.

gammas and remove the lower energy Compton scatters, a yield cut (**ytNF > 0.95**) was applied.

To accurately fit the peak, a linear background model and an exponential background model were included. Before the addition of the yield cut, there was a low-energy shoulder that sometimes manifested as a separate peak. Therefore, the mean and sigma of the Gaussian were constrained to ensure that the correct peak was being fit. The linear and exponential background parameters were left free. Figure 3.4 demonstrates the fit for the total charge and phonon energies on iZIP 4 (T2Z1), while Table 3.7 and Table 3.8 show the parameters extracted from the fits.
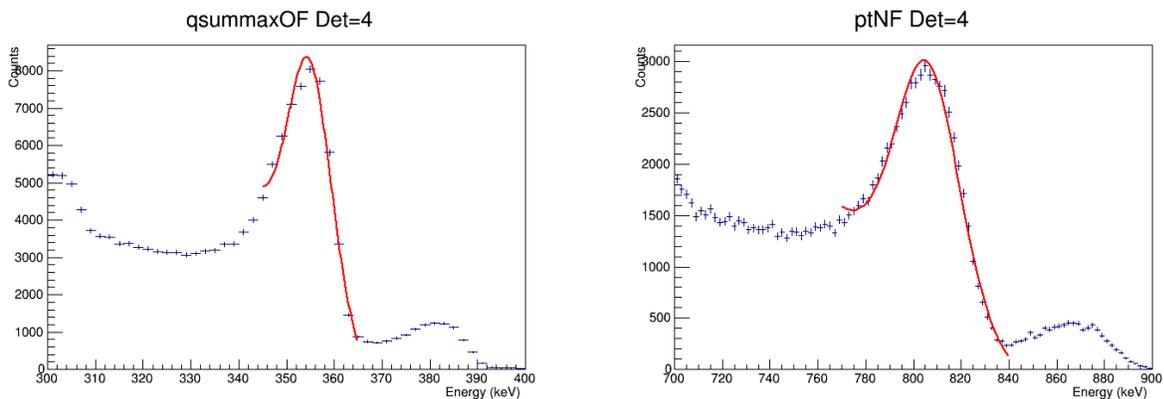
(a) Measuring charge resolution at 356 keV$_{ee}$ for iZIP 4 (T2Z1).

(b) Measuring phonon resolution at 356 keV$_{ee}$ for iZIP 4 (T2Z1).

Figure 3.4: The big 356 keV peak clearly stands out with so much data taken with the gamma calibration source. The 383 keV peak is also visible here, although that is not the case for every detector. The resolution of the 383 keV peak could have been measured for some detectors, but being close in energy to the 356 keV peak makes it a redundant measurement.

Despite not being shown in Fig. 3.4, other detector-specific results are:

- iZIP 12 has at least twice as many events as any other detector (in some cases 10x more). This exaggerates the χ^2/NDF even though the fit looks alright by eye.

- Due to a phonon short on iZIP 3, the peak appears at $\sim 2/3$ the expected energy.

Det	Mean [keV$_{ee}$]	Error	1-sigma [keV$_{ee}$]	Error	χ^2/NDF	% Mean Inaccuracy
1	353.53	0.05	3.59	0.05	260.85	-0.69
2	355.27	0.14	3.47	0.16	1.45	-0.21
3	355.93	0.14	5.38	0.27	3.16	-0.02
4	354.94	0.05	4.44	0.07	65.53	-0.30
5	354.95	0.06	2.90	0.06	63.80	-0.30
6	345.02	0.20	5.21	0.26	4.31	-3.09
7	356.04	0.04	4.35	0.07	23.66	0.01
8	355.28	0.04	4.19	0.07	25.84	-0.20
9	356.26	0.05	4.45	0.12	22.01	0.07
10	360.00	3.40	8.60	0.89	2.68	1.12
11	356.18	0.05	4.16	0.08	12.95	0.05
12	355.53	0.03	4.84	0.05	275.16	-0.13
13	360.00	0.50	9.10	0.34	2.51	1.12
14	359.59	0.004	8.59	0.005	1.86	1.01
15	357.75	3.22e-6	3.63	2.47e-6	3.48	0.49

Table 3.7: Parameters extracted from fitting the 356 keV$_{ee}$ charge peak for all detectors.

Det	Mean [keV$_t$]	Error	1-sigma [keV$_t$]	Error	χ^2/NDF	% Mean Inaccuracy
1	802.60	0.36	18.29	0.56	3.38	-1.98
2	825.74	3.94	21.14	6.09	1.33	0.85
3	583.21	0.24	15.55	0.36	5.90	-28.77
4	806.08	0.14	12.92	0.17	7.24	-1.55
5	814.92	0.35	15.44	0.56	1.85	-0.47
6	749.43	0.49	9.11	0.47	2.74	-8.47
7	799.34	0.33	17.92	0.50	4.43	-2.38
8	799.05	0.24	18.77	0.31	2.42	-2.41
9	809.64	0.39	19.50	0.73	2.20	-1.12
10	823.75	1.31	19.23	2.50	1.37	0.60
11	790.80	0.36	14.55	0.38	2.90	-3.42
12	824.14	0.65	15.90	1.04	2.51	0.65
13	806.52	1.88	34.28	3.29	1.16	-1.50
14	790.00	0.26	14.42	0.46	1.16	-3.52
15	810.90	2.03	17.70	4.62	0.82	-0.96

Table 3.8: Parameters extracted from fitting the 356 keV$_{ee}$ phonon peak for all detectors.

- Despite the yield cut, a few detectors still have a low-energy shoulder in charge that is causing high values of χ^2/NDF. iZIP 1 is the worst case of this.

- Phonon shorts on iZIPs 2, 13, and 15 mean that very little peak is visible. This is why the mean error is an order of magnitude larger than the other detectors.

C.2.5. 511 keV

The 511 keV peak is from the annihilation radiation of a positron emitted from the activated germanium and shows up at the mass of the positron. The data immediately following the activation period was selected using **cPostCf_133_HT**. However, the peak is not as strong as the 10.37 keV, so the charge symmetry cut was needed **cQsym_v53** to select the bulk of the crystal and remove excess backgrounds.

This was a remarkably good fit across all detectors regardless of channel shorts, although there were loose constraints on the Gaussian parameters. As with the other peaks, a linear and an exponential background model were included with free parameters. Figure 3.5 demonstrates the fit for the total charge and phonon energies on iZIP 4 (T2Z1), while Table 3.9 and Table 3.10 show the parameters extracted from the fits.

Despite not being shown in Fig. 3.5, other detector-specific results are:

- This was the only peak where the phonon energy on iZIP 3 was at the correct value instead of 2/3 the expected value.

- There were a few incidents where the rising edge of the peak seemed to be steeper than the Gaussian function. This is especially prevalent for iZIPs 3, 12, 13, and the charge channel on iZIP 6. This is evidenced by the larger χ^2/NDF for these detectors.
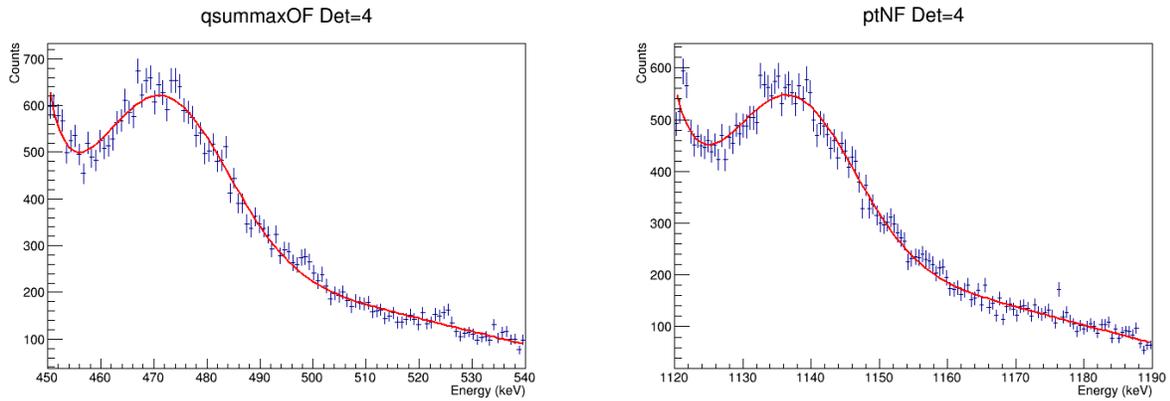
(a) Measuring charge resolution at 511 keV$_{ee}$ for iZIP 4 (T2Z1).

(b) Measuring phonon resolution at 511 keV$_{ee}$ for iZIP 4 (T2Z1).

Figure 3.5: At first glance, both the phonon and charge models appear to be good fits to the data. However, the widths of the peaks do not follow the same trend as the widths at other energies. It is believed that this behavior is caused by saturation in the detectors, so the 511 keV peak was excluded from the resolution model.

Det	Mean [keV$_{ee}$]	Error	1-sigma [keV$_{ee}$]	Error	χ^2/NDF	% Mean Inaccuracy
1	473.25	0.28	8.53	0.40	1.55	-7.39
2	473.72	0.36	12.23	0.45	1.58	-7.30
3	473.72	0.19	10.47	0.31	3.07	-7.30
4	472.20	0.34	12.00	0.51	1.43	-7.59
5	475.46	0.29	10.37	0.35	2.34	-6.95
6	474.29	0.27	10.92	0.42	2.65	-7.18
7	474.64	0.84	7.28	0.99	1.34	-7.12
8	474.47	0.30	10.37	0.44	1.66	-7.15
9	471.02	1.08	11.30	1.02	1.13	-7.82
10	474.30	0.35	9.60	0.49	1.82	-7.18
11	474.95	0.33	11.06	0.51	1.56	-7.05
12	474.36	0.21	11.82	0.34	4.19	-7.17
13	473.42	0.34	10.12	0.49	1.63	-7.35
14	473.66	0.40	12.72	0.59	1.78	-7.31
15	473.84	0.22	10.69	0.37	2.28	-7.27

Table 3.9: Parameters extracted from fitting the 511 keV$_{ee}$ charge peak for all detectors.

Det	Mean [keV$_t$]	Error	1-sigma [keV$_t$]	Error	χ^2/NDF	% Mean Inaccuracy
1	1137.94	0.22	6.90	0.33	1.52	-3.18
2	1138.13	0.18	9.71	0.21	1.17	-3.16
3	1133.61	0.09	6.29	0.16	4.42	-3.55
4	1137.45	0.27	8.99	0.32	1.46	-3.22
5	1137.76	0.24	9.66	0.34	1.78	-3.19
6	1137.71	0.14	8.66	0.21	2.97	-3.20
7	1129.59	7.74	11.47	4.72	1.32	-3.89
8	1138.83	0.22	7.64	0.28	1.47	-3.10
9	1130.82	0.46	13.11	0.43	1.25	-3.78
10	1139.71	0.23	7.51	0.33	1.45	-3.03
11	1139.01	0.23	8.23	0.33	1.42	-3.09
12	1138.85	0.17	9.22	0.25	2.46	-3.10
13	1136.38	0.21	9.99	0.23	1.08	-3.31
14	1138.17	0.16	8.30	0.19	1.32	-3.16
15	1138.21	0.15	8.95	0.28	2.48	-3.16

Table 3.10: Parameters extracted from fitting the 511 keV$_{ee}$ phonon peak for all detectors.

- The ^{210}Pb-implanted plate above iZIP 7 has buried any 511 keV peak and led to the large error on the Gaussian mean and sigma. A similar plate below iZIP 9 also increased the background rate, but the 511 keV peak is still slightly present.

C.3. Summary

For each detector, all of the 1σ widths can be collected from each peak and plotted against the energy of that peak, with their respective uncertainties. This is shown in Fig. 3.6 for iZIP 4 (T2Z1). According to the results, the absolute resolution worsens as energy increases, but the relative resolution improves. Stated in a more graphical way, the trend of the points is increasing but at a slope « 1.
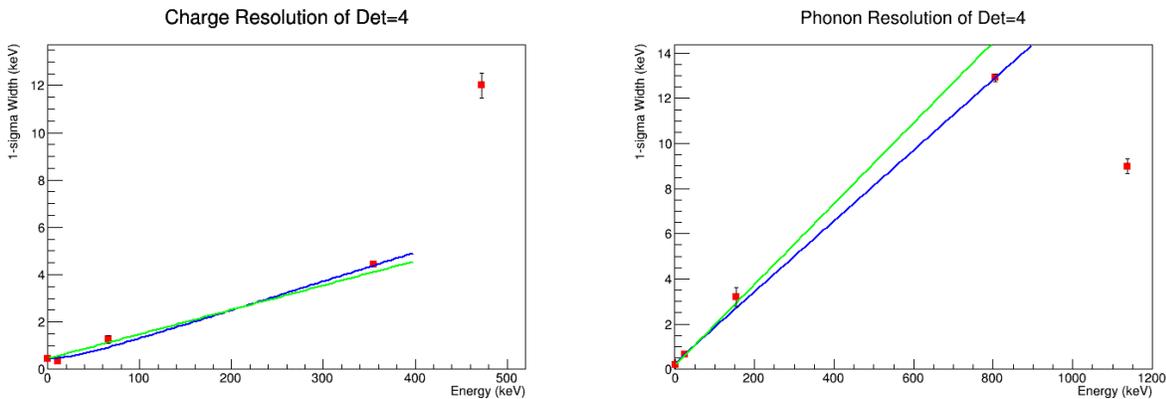
(a) Charge resolution as a functtion of energy for iZIP 4 (T2Z1).

(b) Phonon resolution as a functtion of energy for iZIP 4 (T2Z1).

Figure 3.6: Fitting the resolution model to the measured peak resolutions. As discussed in the text, the 511 keV peak was excluded from the fit due to suspected saturation. The green line is a linear fit because the points appear to be linear by eye; it is more obvious on other detectors that this is not the case. The blue line is the resolution model described in Eq. 3.1.

Based on the results of the other peaks, the width of the 511 keV peak seems to be too large for the charge channel and too small for the phonon channel, by almost a factor of two in both cases. This is suggestive of saturation at high energy. As evidence to support this hypothesis, all of the histogram statistics dropped by roughly three orders of magnitude when the charge and phonon saturation cuts (**!cQsat_133** and **!cPsat_133**, respectively)

were applied to the data. This indicates that those were indeed saturated events, so the measured resolution at this energy has been excluded from the resolution function.

At first, the resolution of the remaining peaks appeared to scale linearly with energy, which is what the green line is showing; some of the other detectors show more clearly that this is not an accurate description of the resolution. However, a more physically motivated model is depicted by the blue line, where the variance is defined as

$$\sigma^2(E) = A + BE + CE^2 \,, \tag{3.1}$$

where σ is the resolution and the parameters A, B, and C are defined as: [143]

- A is noise that is added to every event from the readout electronics, independent of the energy.

- B is an intrinsic variance in the energy due to the Fano factor of the material, directly proportional to the energy.

- C is an extra term to account for the systematic mis-estimation of energy due to feeding a fixed pulse template into the OF, variance scales with the square of the energy.

Table 3.11 and Table 3.12 show the parameters extracted from the resolution fits versus energy.

That these results have been referenced most often by the simulations group to smear perfect, simulated data to better match the energy resolution of data collected with the detectors. The same is true for signal and background models created with analytic functions that need to be smeared to match the performance of the detector. Another application is accounting for the energy uncertainty when setting a threshold to be confident that fluctuations are also contained.

Det	A [keV$^2_{ee}$]	Error	B [keV$_{ee}$]	Error	C	Error
1	0.166004	0.000670044	0.00233716	0.00320014	9.52576e-05	9.52081e-06
2	0.126721	0.00069113	0.0066017	0.00305277	7.59044e-05	1.28096e-05
3	0.0802679	0.00110084	0.0171239	0.00420509	0.00018207	2.66854e-05
4	0.20685	0.000326366	7.21645e-12	3.35021e-05	0.000151331	4.83082e-06
5	0.150299	0.000462304	0.00416849	0.00166944	5.40763e-05	5.51498e-06
6	0.163645	0.00394873	0.0285656	0.014543	0.000142535	4.88897e-05
7	0.0922709	0.000718869	0.000696472	0.00441894	0.00014649	1.11985e-05
8	0.118942	0.000589952	0.000385049	0.00515324	0.000137343	9.32403e-06
9	0.0863211	0.00180246	0.0137235	0.00802272	0.00011678	2.44364e-05
10	0.130918	0.00794812	0.030449	0.0361324	0.000493483	0.0001641
11	0.12363	0.000787667	0.0160938	0.00319961	8.98899e-05	1.04657e-05
12	0.130988	0.000814785	0.00300699	0.00310211	0.000175811	9.59797e-06
13	1.4988e-11	0.755337	0.0183041	0.0127495	0.00058953	6.09428e-05
14	0.185721	0.000273228	0	0.000869786	0.000569211	6.63032e-07
15	0.187189	0.000172272	0.0158603	5.77083e-07	5.71605e-05	1.61276e-09

Table 3.11: Parameters extracted from fitting the charge resolution versus energy for all detectors.

Det	A [keV$_t^2$]	Error	B [keV$_t$]	Error	C	Error
1	0.155393	0.0691325	9.60343e-11	0.00885466	0.000506287	3.0612e-05
2	0.082099	0.000138586	0.0108455	0.00183727	0	2.22402e-06
3	0.0846805	0.000118728	4.71845e-10	0.0108038	0.000702979	3.1949e-05
4	0.0328312	5.00053e-05	0.0100053	0.000280728	0.000244885	6.96688e-06
5	0.0411937	6.97124e-05	0.000549462	0.0010887	0.000361038	2.63019e-05
6	0.10674	0.000583262	0.144297	0.0521945	-4.49364e-05	7.14893e-05
7	0.327138	0.000644738	-0.0137804	0.00904437	0.000520678	3.12901e-05
8	0.0195073	6143.36	0.00802392	6145.57	0.000549911	6148.36
9	0.13223	0.206687	5.55112e-12	0.0310852	0.000572963	4.26256e-05
10	0.102399	0.000197759	0	0.00170764	0.000415507	8.16471e-05
11	0.0788837	0.000225082	0.0141565	0.00469563	0.000321161	1.91979e-05
12	0.0480589	9.00865e-05	0.00535584	0.00251146	0.000389617	5.05697e-05
13	0.492861	0.000972852	0.0238974	0.0313525	0.000810142	0.000199669
14	0.0343254	7.40223e-05	0.0141611	0.00260691	0.000314387	2.2043e-05
15	0.046846	8.63333e-05	0.0191361	0.00429152	2.90118e-05	4.67514e-05

Table 3.12: Parameters extracted from fitting the phonon resolution versus energy for all detectors.

C.4. Future Studies

There are a few interesting ways that this study could be developed further. The first would be to find more features in the data and measure the width of them. Having more data points at more energies would certainly improve the resolution function because there is currently a model with three parameters being fit to four data points.

Another interesting question is whether the resolution changes over time. This entire study would have to be recreated, paying close attention to the dates of the series being analyzed. The limited statistics in some of the fits could make this difficult, especially the 66.7 keV peak since I needed to combine multiple datasets. However, there might be ample statistics using just the 0 keV, 10.37 keV, and 356 keV peaks.

Finally, the simulations group has been working for many years to develop a detector response simulator, called the Detector Monte Carlo (DMC). In order to trust the simulated output, the DMC data needs to be validated against SuperCDMS data. This study would be interesting to repeat using DMC data to see how accurately the resolution compares to these results.

APPENDIX D

TAIL-FITTING ALGORITHM

The ability to accurately calibrate the energy scale of a detector is crucial to understanding and trusting the data. This is usually done with calibration sources that have features in the data that occur at known energies. The energy can be scaled so that these features align at the proper energies, and the scaling can then be applied to other datasets taken with the same detector.

Since there were many layers of shielding around the SuperCDMS Soudan experiment to mitigate external backgrounds, the calibration sources needed to be inserted close to the detectors. However, there were limitations to how close the calibration source could physically get. This was not a problem for the high energy part of the spectrum, because particles such as the 356 keV photon from ^{133}Ba could penetrate the remaining shielding. However, the shielding becomes an issue when attempting to calibrate the low energy part of the spectrum. This is why the electron capture peaks were so helpful because a penetrating ^{252}Cf neutron source was used to activate the germanium and turn it into its own low-energy calibration source.

When considering the energy calibration that will be needed for SuperCDMS SNOLAB, a considerable issue arises. Silicon detectors will be deployed that do not have activation lines in the same way as germanium. This means it will be difficult to calibrate the low energy scale because calibration sources will not be able to penetrate the shielding. This is further complicated by the NTL gain of the HV detectors, described in Section 0.2.1, that will cause high energy sources to saturate the TESs. Hence, an algorithm was devised that could reconstruct the energy from a saturated pulse in order to calibrate using the

higher energy sources such as ^{133}Ba. Of course, this only solves half of the problem, and the calibrations group is still working on alternative methods of calibrating silicon detectors at low energy [98].

D.1. Tail-Fitter Motivation

The problem of saturated events is more prevalent for high-voltage operation such as CDMSlite or HV detectors, so the algorithm to extract energy is only focused on phonon pulses since these modes do not utilize the charge signal.

The two main factors that influence the saturation of a phonon pulse are the operating temperature and the critical temperature T_c of the TES, which was described in Section 3.1.2.1. An event that has enough energy to cause the TES to go normal will result in a flat pulse shape until enough power has dissipated that the TES cools back below T_c and resumes energy collection. Typically, the energy of the event is determined by the amplitude of the pulse, but that is not available in the case of a saturated event where the top of the pulse is not present. However, the pulse from each channel has a characteristic fall-time that is dictated by the thermal response of the TESs. This, in principle, means that the tail of a pulse can be used to extract the energy of that event, even if the pulse exhibits saturation.

D.2. Description of the Algorithm

An algorithm was devised to extract the energy of a pulse, regardless of whether that pulse exhibits saturation. First it reads in a raw pulse, subtracts the pre-pulse baseline noise, and normalizes the pulse to convert it from Analog-to-Digital Converter (ADC) units to physical Amps. Then it fits the tail of a pulse with an exponential function where A is the amplitude, $B = -1/\tau$ is the negative inverse fall-time of the pulse, and C is an offset. The fit function is defined as

$$f_{fit}(t) = Ae^{tB} + C \,. \tag{4.1}$$

Parameter A is initialized to be the value of the pulse at the beginning of the tail. The tail was hard-coded to start at 800 μs, and there is a small adjustment for the delay time of the pulse with respect to the trigger time. As mentioned in the last section, the fall-time of each channel has been determined, so parameter B is fixed. Parameter C is a vertical shift to account for the end of the pulse not returning to zero in the pulse window. This parameter is initialized by fitting the baseline of the last 50 bins of the pulse. All of these initialized parameters are fed to a MINUIT minimizer that minimizes the χ^2 between the exponential function and the pulse.

In order to extract the energy from the fit function, the integral is taken from the start of the tail t_s to the end of the pulse t_e,

$$E = \int_{t_s}^{t_e} (Ae^{tB} + C)dt = (A/B)(e^{t_e B} - e^{t_s B}) \,. \tag{4.2}$$

The offset parameter C accounts for the pulse not returning to the baseline within the readout window, but it is safe to assume the pulse eventually does return to baseline at some infinite time, meaning that $\lim_{t_e \to \infty} e^{t_e B} = 0$. Also, the value of A is set to be the value of the function at t_s, so $e^{t_s B} = 1$ by definition. This simplifies the energy from the algorithm to

$$E = -A/B \,. \tag{4.3}$$

This energy is calculated per channel before some minor corrections are applied. First is the relative channel calibration so that the energy from each channel can be co-added. Then the energy summed over all channels is scaled by a constant that relates the integral of the tail to the amplitude of the pulse. It is important to point out that this energy is not properly calibrated. Appendix D.3 will discuss the calibration, but for now the x-values of the spectra shown here should be considered arbitrary units of energy. To elucidate this point, units will not be used when referring to the extracted tail-fitter energies.

Figure 4.1 demonstrates that the algorithm successfully reproduced the ^{133}Ba spectrum for the detector used in CDMSlite R3 (T2Z1/iZIP 4). However, Fig. 4.2 shows that the ^{133}Ba spectrum does not emerge when the algorithm is applied to the CDMSlite detector used in R2 (T5Z2/iZIP 14).
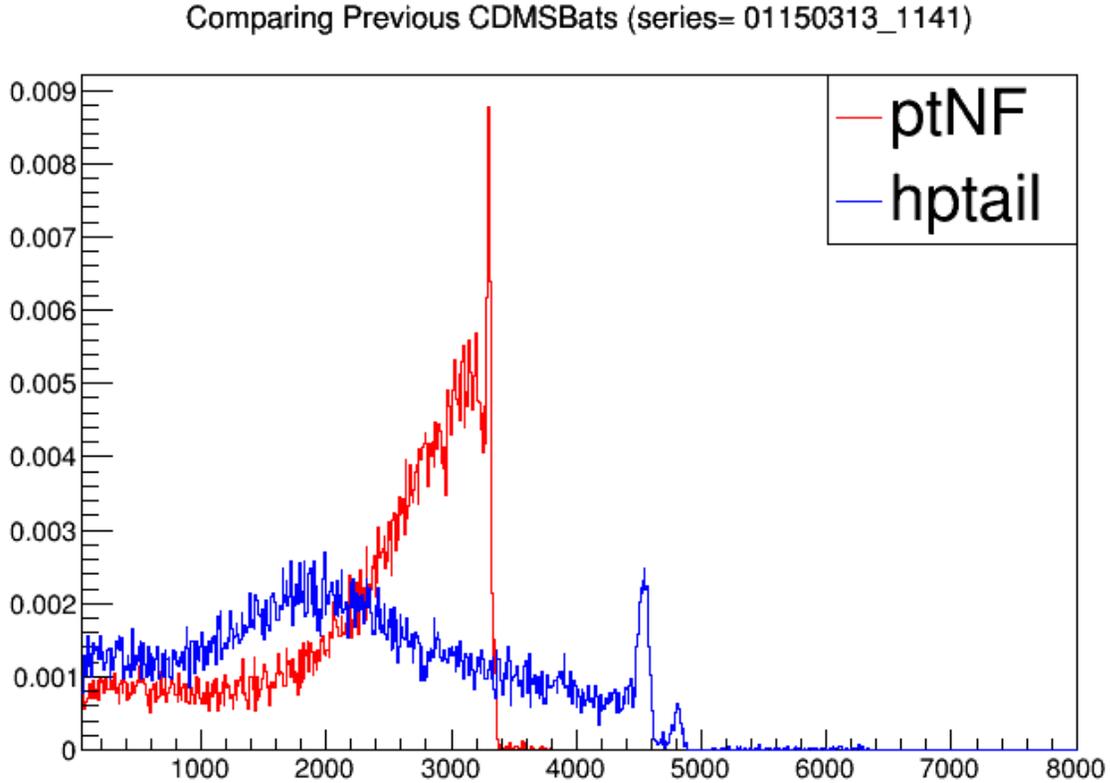
Figure 4.1: Existing status of the tail-fitter in CDMSBats applied to iZIP 4, with both axes in arbitrary units. The red curve is the energy quantity often used from CDMSBats that is based on the amplitude of the template in the OF algorithm described in Section 3.1.3.4. Saturation is making the reconstructed energies pile up. The blue curve is the reconstructed energy from the tail-fitter. There is a clear 356 keV peak visible near 4500 and even the smaller 383 keV peak is visible at 4800.

D.2.1. Replicating the Algorithm

This algorithm was written as part of CDMSBats, which is a complex framework of scripts that processes all SuperCDMS data from the raw data pulses to the final calibrated quantities. This is the correct place for such an algorithm to exist, but makes it difficult to
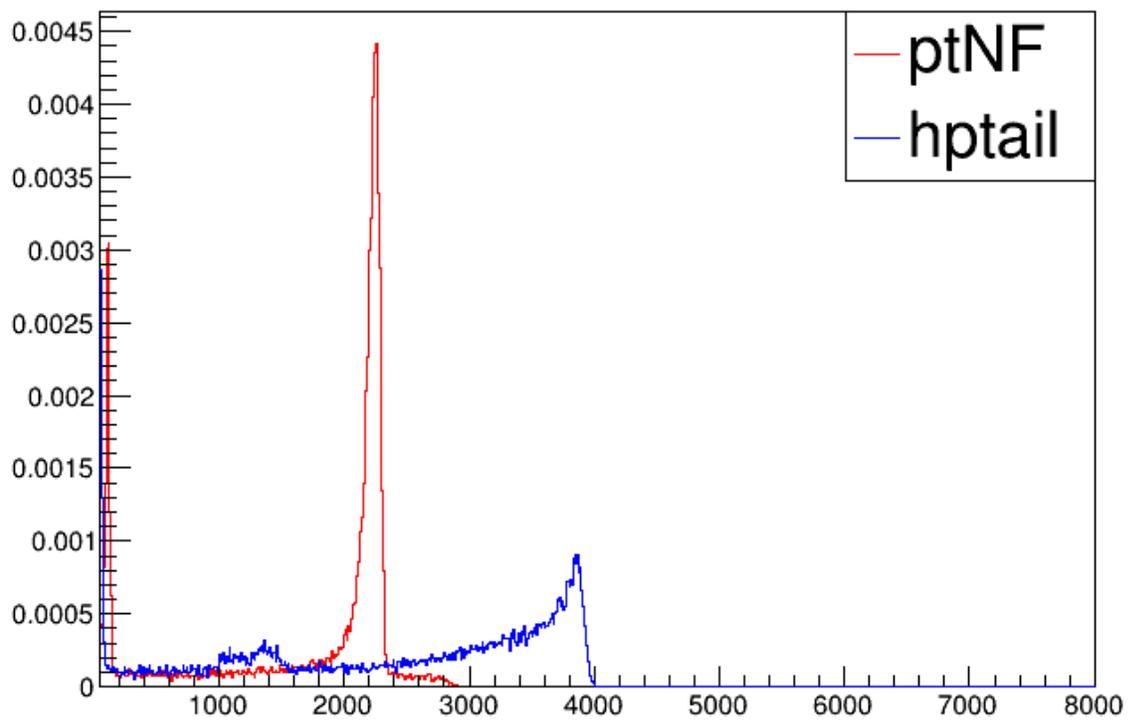
Figure 4.2: Neither curve on this plot displays the expected ^{133}Ba spectrum described in Section 3.1.3.1. Both curves exhibit a large spike of events that is indicative of saturated pulses causing the energy reconstruction to pile them up at the same energy. Detailed investigation of the raw pulses and the behavior of the fitting routine is needed in order to diagnose the energy reconstruction algorithm for this detector (iZIP 14).

test and diagnose. In order to manipulate the algorithm and implement possible fixes, a toy replica of the fitting routine in CDMSBats was written.

The first check of the toy replica was to verify that it was exactly recreating the fit in CDMSBats. To do this, a prime example of an unsaturated pulse and a saturated pulse were chosen to compare the results of the replica to the fit quantities output from CDMSBats. For the unsaturated pulse, Fig. 4.3 shows the initial parameters, Fig. 4.4 shows the results of the replica fit, and Fig. 4.5 shows the output from CDMSBats overlaying the pulse. Table 4.1 shows a quantitative comparison of the values for this particular pulse, event number 20009 in series 01150313_1141 from iZIP 4.
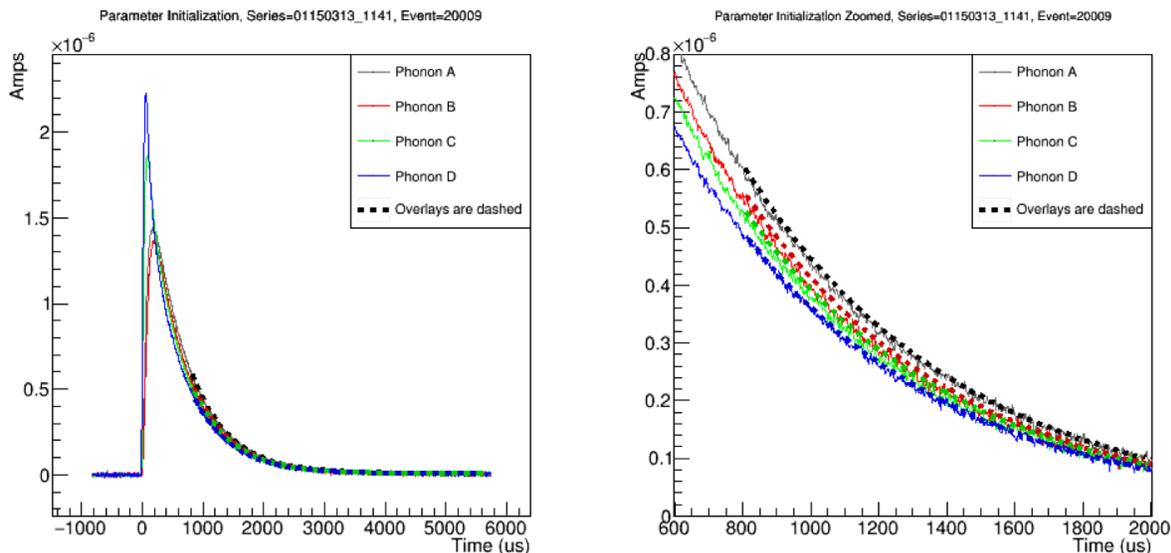
Figure 4.3: The initial parameters of the tail-fitter overlaying an unsaturated raw pulse for each channel of the detector. The dotted lines depict the fit function in Eq. 4.1 overlaying each pulse and they are color-coded to match the channel. The plot on the right is a zoomed version of the left plot.

For the saturated pulse, Fig. 4.6 shows the initial parameters, Fig. 4.7 shows the results of the replica fit, and Fig. 4.8 shows the output from CDMSBats overlaying the pulse. Table 4.2 shows a quantitative comparison of the values for this particular pulse, event number 20001 in series 01150313_1141 from iZIP 4.
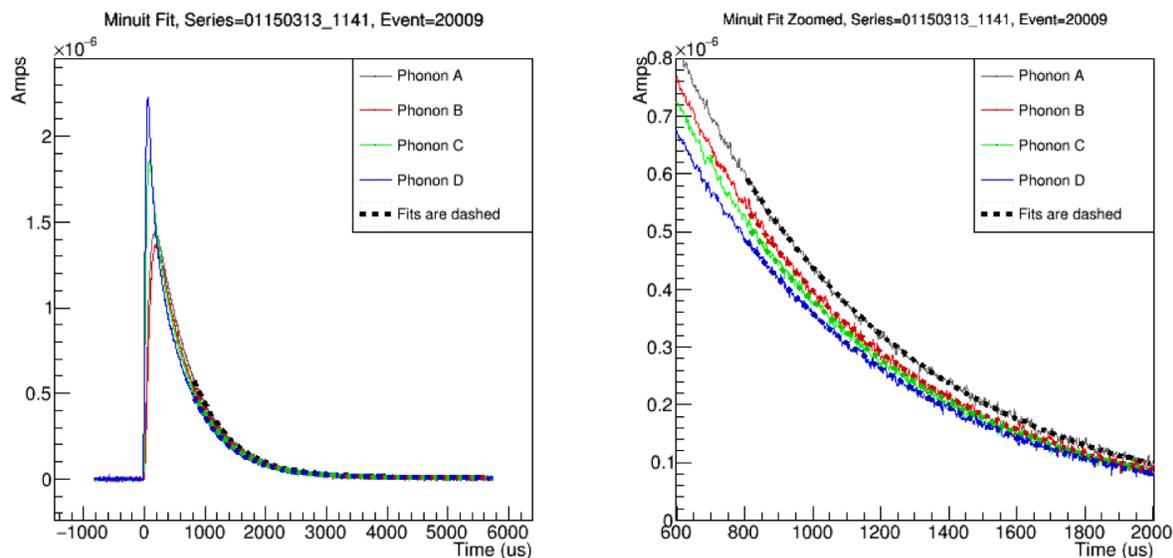
Figure 4.4: Plots demonstrating the behavior of the replica fitter on an unsaturated pulse. The goal was to recreate the fit of CDMSBats, so this figure needed to match Fig. 4.5 and it does.
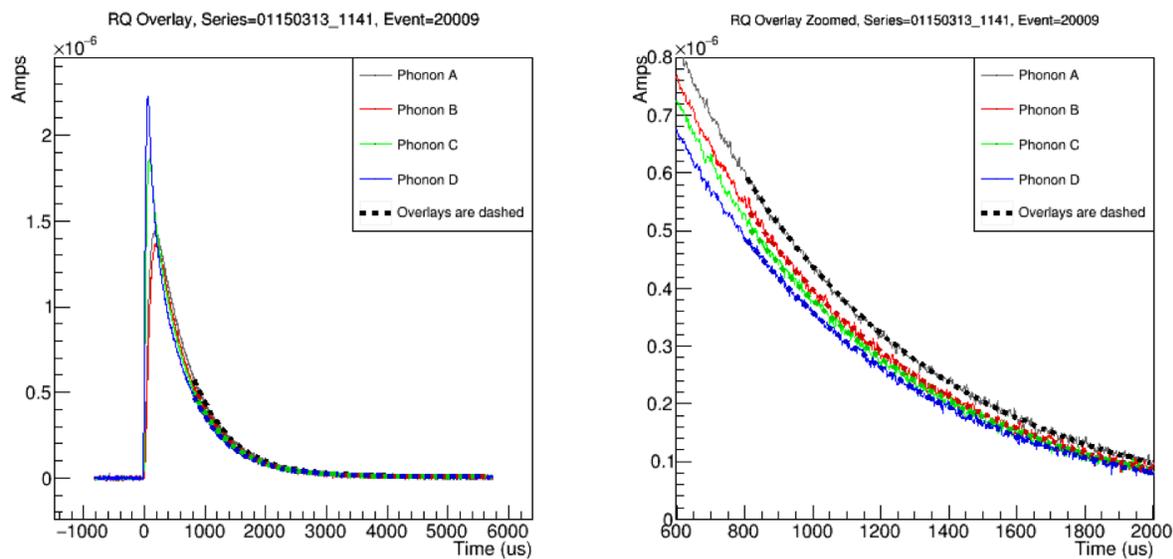
Figure 4.5: The exponential tail-fitter curve drawn using the output parameters from CDMSBats. The curves overlay the raw pulse to show the performance of the previous tail-fitter.

Channel	Parameter	Bats Init	Replica Init	Replica Fit	Bats Fit
	A [nAmps]	602.905	602.905	589.624	589.624
A	B [$1/\mu s$]	-1.56495e-3	-1.56495e-3	-1.56495e-3	-1.56495e-3
	C [nAmps]	6.0791	6.0791	6.91898	6.919
	A [nAmps]	560.387	560.387	539.932	539.932
B	B [$1/\mu s$]	-1.58479e-3	-1.58479e-3	-1.58479e-3	-1.58479e-3
	C [nAmps]	5.05219	5.05219	4.55861	4.55861
	A [nAmps]	529.358	529.358	510.908	510.907
C	B [$1/\mu s$]	-1.55521e-3	-1.55521e-3	-1.55521e-3	-1.55521e-3
	C [nAmps]	4.87061	4.87061	4.78894	4.78933
	A [nAmps]	487.329	487.329	481.097	481.097
D	B [$1/\mu s$]	-1.5748e-3	-1.5748e-3	-1.5748e-3	-1.5748e-3
	C [nAmps]	5.7251	5.7251	6.98992	6.98975

Table 4.1: Quantitative comparison between the initial fit parameters in CDMSBats and the replica tail-fitter for the unsaturated pulse. The last two columns are the resulting parameters from the fits. Parameter B is fixed in both algorithms. The differences are $< 1\%$ and are attributed to rounding error.
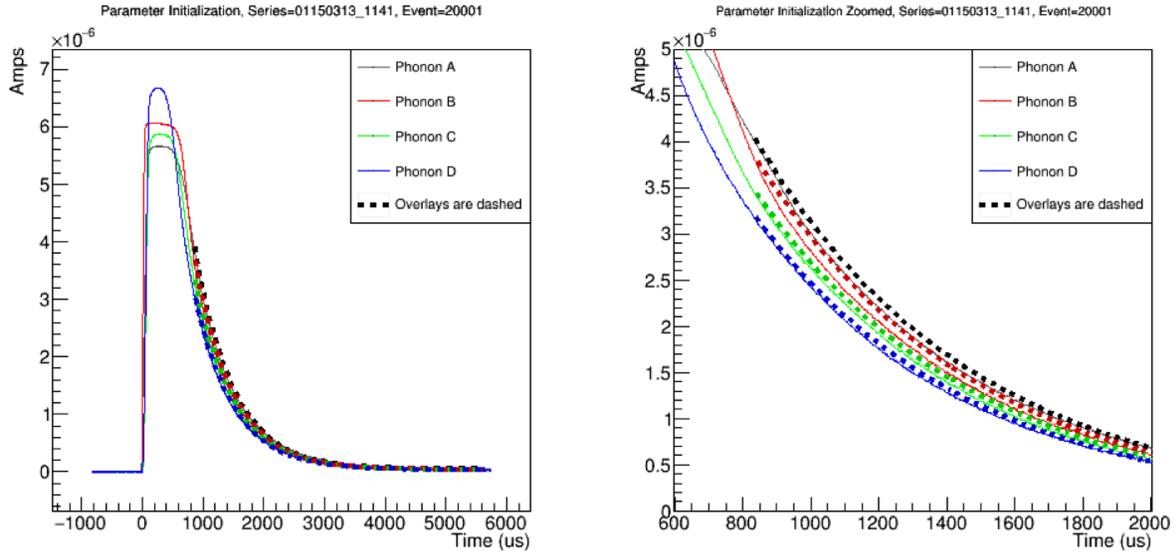
Figure 4.6: Analogous plots to Fig. 4.3 showing the parameter initialization, but for a saturated pulse from the same detector and series. The plot on the right is a zoomed view of the left and the curves are color-matched for each channel.
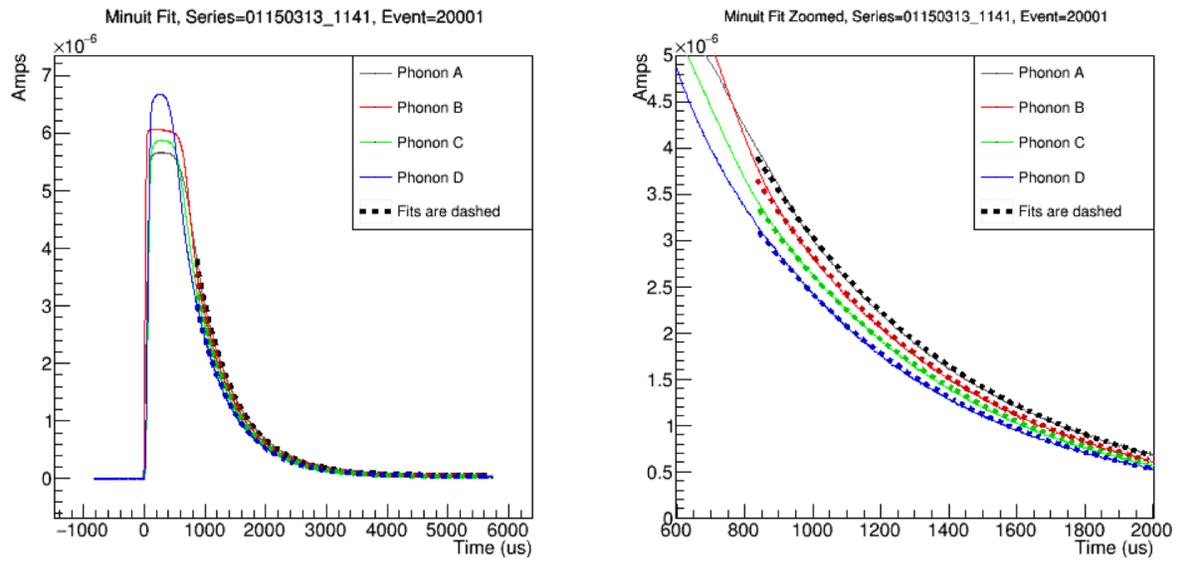
Figure 4.7: Analogous plots to Fig. 4.4 with a saturated event. This figure needed to match Fig. 4.8 to confirm that the replica fitter mimics that of CDMSBats, and it does.
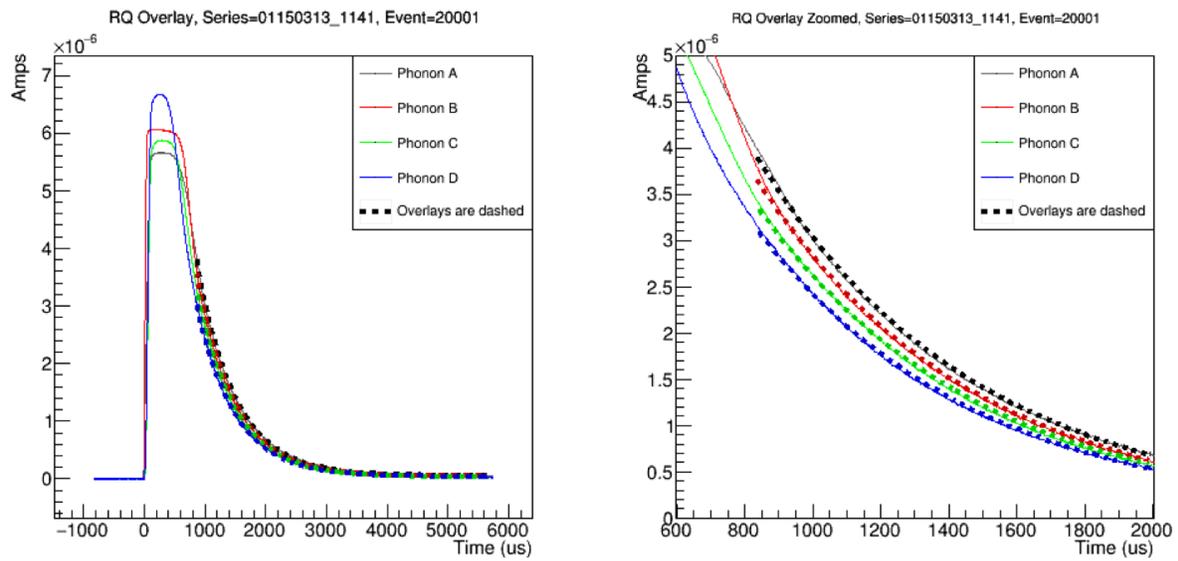
Figure 4.8: The curves are overlaying the raw pulse to show the performance of the CDMSBats tail-fitter for a saturated event, analogous to the unsaturated pulse shown in Fig. 4.5.

Channel	Parameter	Bats Init	Replica Init	Replica Fit	Bats Fit
	A [nAmps]	4027.8	4027.8	3887.63	3887.63
A	B [1/μs]	-1.56495e-3	-1.56495e-3	-1.56495e-3	-1.56495e-3
	C [nAmps]	41.7633	41.7633	53.1672	53.1671
	A [nAmps]	3841.8	3841.8	3653.75	3653.75
B	B [1/μs]	-1.58479e-3	-1.58479e-3	-1.58479e-3	-1.58479e-3
	C [nAmps]	34.7961	34.7961	42.045	42.045
	A [nAmps]	3470.85	3470.85	3376.81	3376.81
C	B [1/μs]	-1.55521e-3	-1.55521e-3	-1.55521e-3	-1.55521e-3
	C [nAmps]	27.2095	27.2095	28.6909	28.6908
	A [nAmps]	3193.35	3193.35	3118.48	3118.48
D	B [1/μs]	-1.5748e-3	-1.5748e-3	-1.5748e-3	-1.5748e-3
	C [nAmps]	29.4693	29.4693	35.6715	35.6714

Table 4.2: Quantitative comparison between the initial fit parameters in CDMSBats and the replica tail-fitter for the saturated pulse. The last two columns are the resulting parameters from the fits. Parameter B is fixed in both algorithms. The differences are $< 1\%$ and are attributed to rounding error.

Satisfied with the accuracy of the replica fitting algorithm, the next step was to use the replica to check the behavior of the algorithm for an event on the other detector to see why the energy reconstruction was not working properly.

D.2.2. Fixing the Algorithm

It was immediately obvious that the initialization of the fitter on iZIP 14 was not close to the pulse and MINUIT was unable to adjust the parameters to properly match the pulse. The example shown here is event number 20033 in series number 01140401_0847 from iZIP 14. As before, Fig. 4.9 shows the initial parameters, Fig. 4.10 shows the results of the replica fit, and Fig. 4.11 shows the output from CDMSBats overlaying the pulse. One positive takeaway is that the replica fitter still matches the output of CDMSBats.
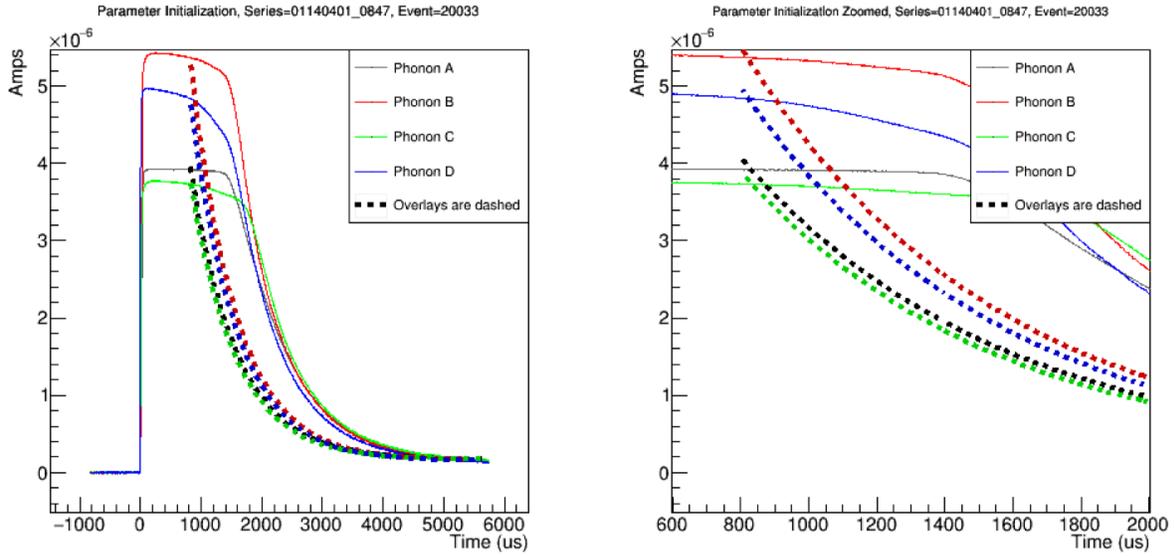
Figure 4.9: Analogous plots to Fig. 4.6 with a saturated event on the other detector. This is the "smoking gun" to the problem because the initial parameters are not close to the tail of the pulse, which increases the probability that the χ^2 fitter will get stuck in a local minimum.
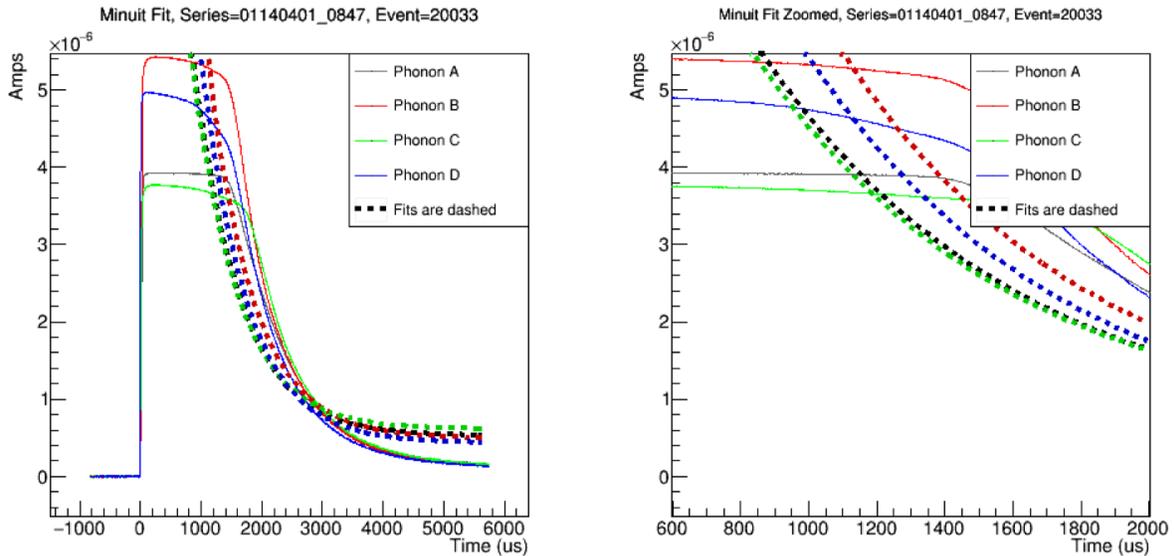
Figure 4.10: Analogous plots to Fig. 4.7 with a saturated event on the other detector. The behavior of the replica fitter needed to match the output of CDMSBats in Fig. 4.11 in order to diagnose the algorithm for this detector.

213
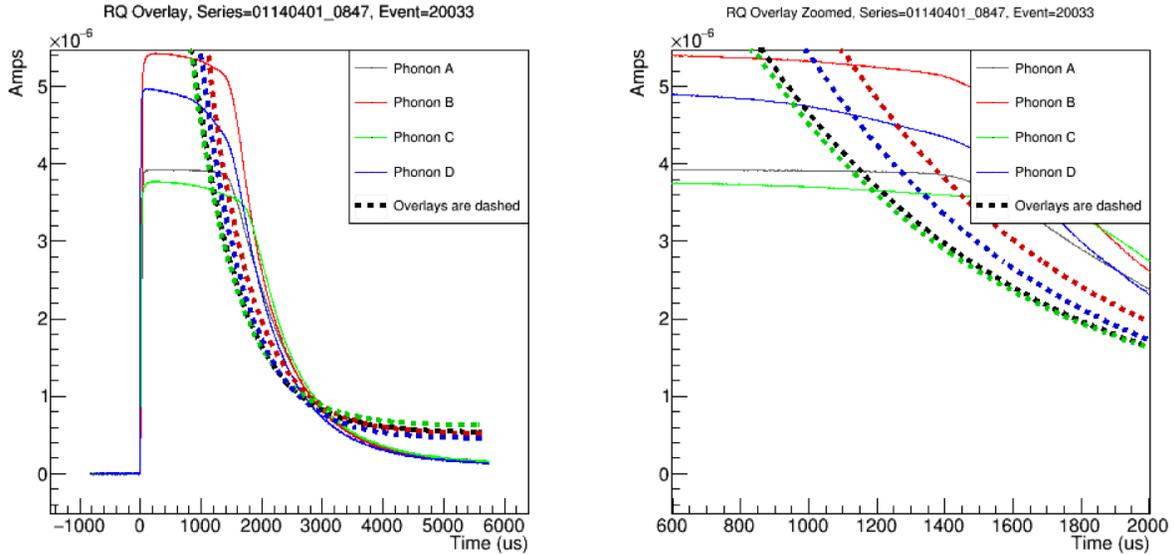
Figure 4.11: Analogous plots to Fig. 4.8 with a saturated event. This clearly shows that the fitter is not able to match the tail of the pulse and get an accurate energy reconstruction.

The biggest difference with the pulse shapes from iZIP 14 was that the duration of the saturated part of the pulse was much longer than saturated pulses from iZIP 4. Also, the saturation "plateau" is at a lower current than the pulse in Fig. 4.6, Fig. 4.7, and Fig. 4.8.

An investigation of differences between these detectors revealed that T_c was systematically higher for iZIP 4 than iZIP 14. Each phonon channel of the detector has a different T_c that is measured after the detector has been fabricated [99]. The values of T_c on iZIP 4 were 76 mK, 80 mK, 79 mK, and 80 mK), while the values of T_c on iZIP 14 were 63 mK, 70 mK, 68 mK, 71 mK. While TES saturation depends on a number of variables, it is likely that this systematic difference in T_c is causing the different detector responses.

The simplest solution to fix the algorithm was to change where the start of the tail was defined. Previously, this was hard-coded to be 800 μs in the configuration file that is input into CDMSBats. Instead, the algorithm was modified to include a fourth fit parameter D that allows the fit function to shift left or right,

$$f_{fit}(t) = Ae^{(t-D)B} + C \, . \tag{4.4}$$

In order to initialize the horizontal shift parameter D, the algorithm calculates the slope at every point of the raw pulse and selects the time where the slope is the most negative. Whether the pulse is saturated or not, this should be the beginning of the exponential tail of the phonon pulse.

Using event number 20018 in series number 01140401_0847 from iZIP 14, Fig. 4.12-Fig. 4.14 shows the improved fit despite the long saturation time. Figure 4.12 shows the initial parameters that are closer to the actual pulse, Fig. 4.13 shows the results of the replica fit, and Fig. 4.14 shows the output from CDMSBats overlaying the pulse to demonstrate how bad the original algorithm was.
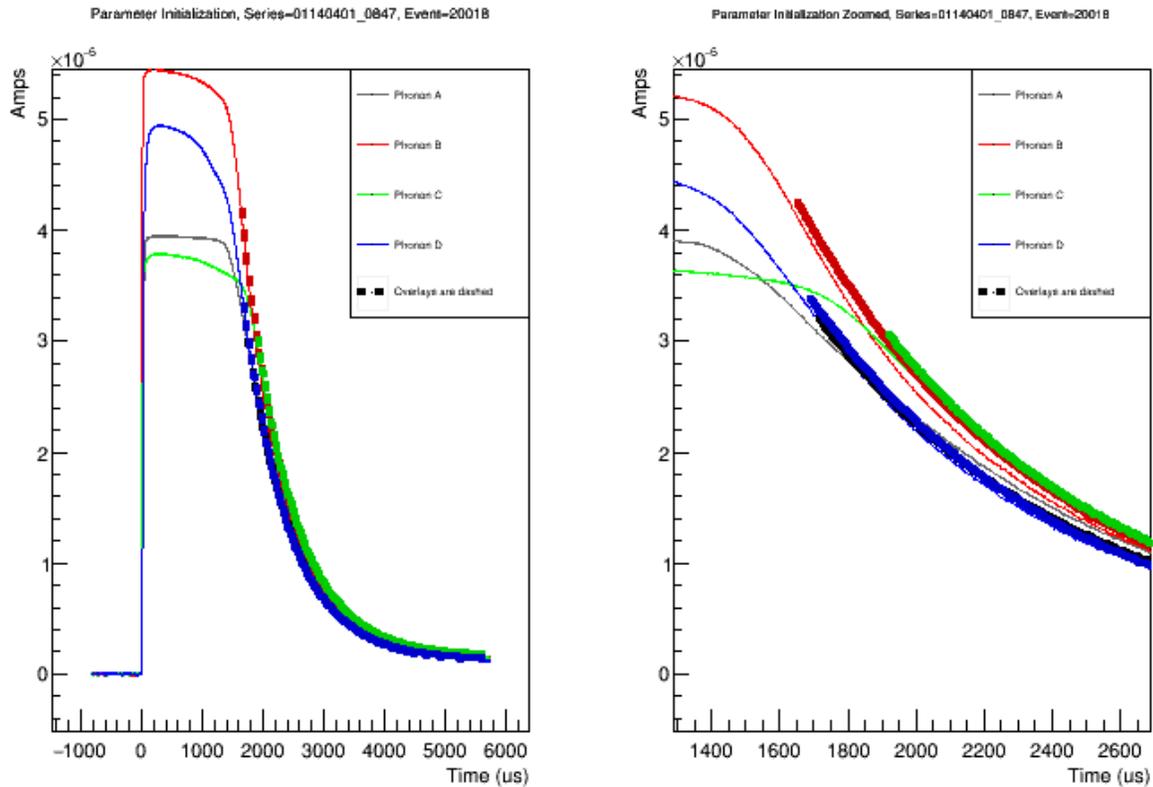
Figure 4.12: Analogous plots to Fig. 4.6 with a saturated event on the other detector. This is the critical component of the modification because the algorithm needs to accurately describe the shape of the pulse.
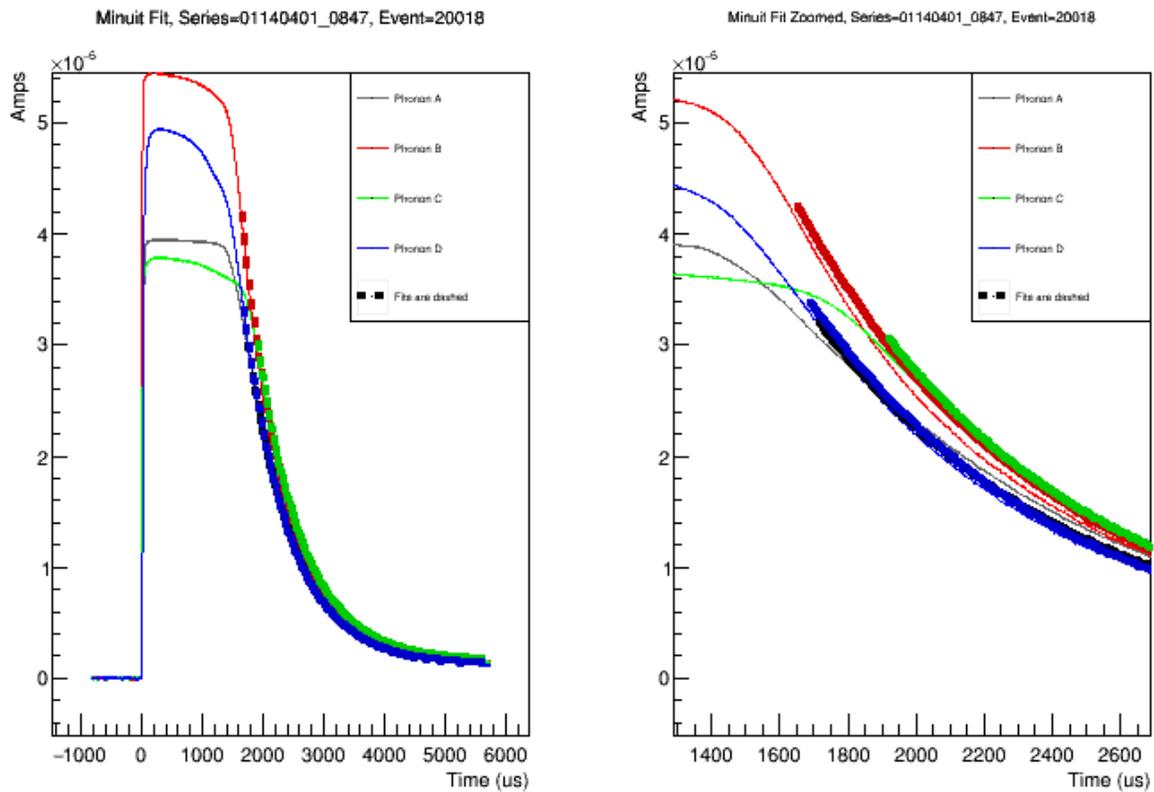
Figure 4.13: Analogous plots to Fig. 4.7 with a saturated event on the other detector. The fitting routine did not have to do very much to fit the pulse because the initialization did such a good job. The difference between this and Fig. 4.12 is imperceptible.
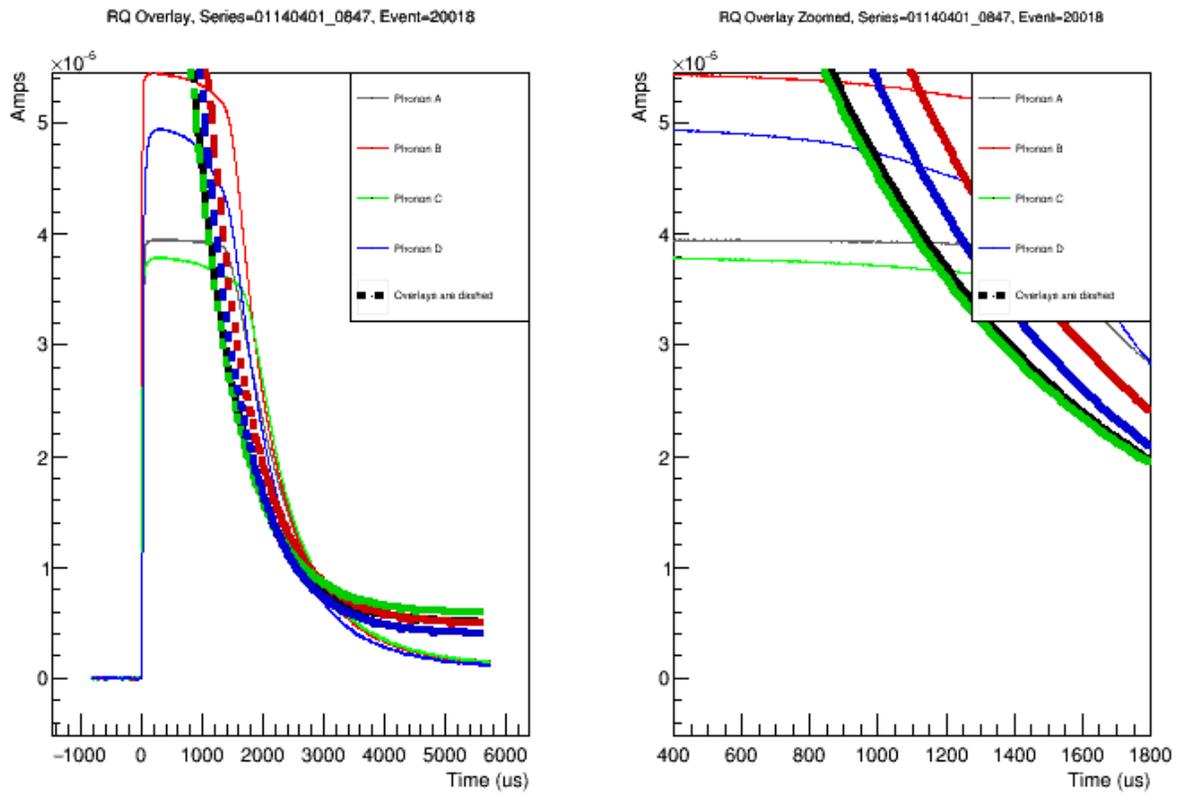
Figure 4.14: Analogous plots to Fig. 4.8 with a saturated event. This plot shows the results of the original fitting algorithm for this event, which is clearly not able to match the tail of the pulse.

Because of the addition of the extra parameter, the integral of the pulse tail that extracts the energy as described in Eq. 4.3 needed to be modified. When the beginning of the pulse tail was hard-coded, the integral started from that time and was the same for all channels and all pulses. However, allowing the exponential to shift horizontally for each channel means that the amplitude A is defined at different times and the relative calibration between channels changes for each event. That is why a specific time was chosen to begin all the integrals. From an implementation standpoint, this means that instead of Eq. 4.3 the integrated energy is

$$E = -f_{fit}(t_s)/B. \tag{4.5}$$

The calculated energy will be calibrated later, so the choice of t_s does not effect the results as long as all fits use the same time. A start time of $t_s = 800$ μs was chosen since that is where the integral in the original algorithm was defined to start.

Figure 4.15 shows the resulting spectrum for the CDMSlite R3 detector (T2Z1/iZIP 4) after modifying the tail-fitting algorithm. This was the detector that the old tail-fitter already worked for in Fig. 4.1, so it was important for the ^{133}Ba spectrum to still be clearly visible for this detector. Figure 4.16 shows the resulting spectrum for the CDMSlite R2 detector (T5Z2/iZIP 14) after modifying the tail-fitting algorithm. The old algorithm did not work for this detector as shown in Fig. 4.2, so it was important for the ^{133}Ba spectrum to emerge for this detector.

All of the energy spectra shown above this point have an energy cut of $ptNF > 100$ keV to remove low energy noise. However, Fig. 4.16 still shows a large spike of events at low ptNF. The tail-fitter is smearing out this peak, but it still creates the large bump below 1000 on the blue histogram. The source of these events is unknown, but increasing the low energy cut to $ptNF > 200$ keV (8.22 keV$_{ee}$) successfully removes them as shown in Fig. 4.17.

D.3. Calibration
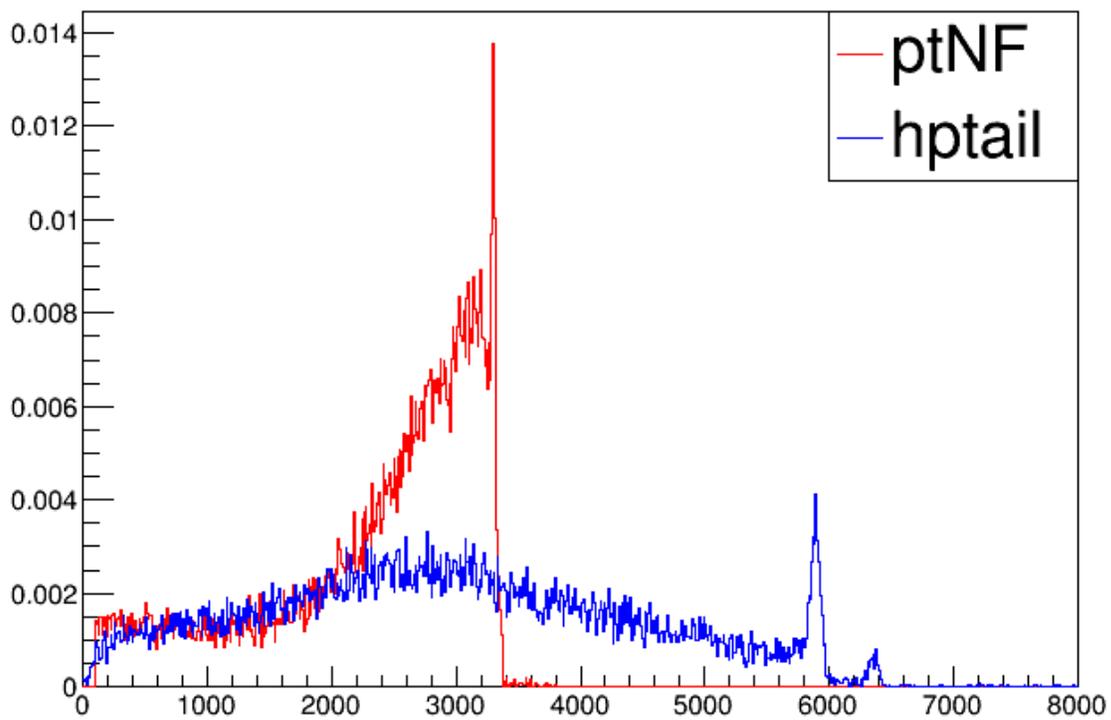
Figure 4.15: The red curve is the energy reconstruction often used from CDMSBats that is based on the amplitude of the template in the OF algorithm described in Section 3.1.3.4. The lower half of the Compton "hump" is visible, but saturation makes the reconstructed energy pile up. The blue curve is the reconstructed energy from the tail-fitter. There is a clear 356 keV peak visible near 4500 and even the smaller 383.8 keV peak visible at 4800.
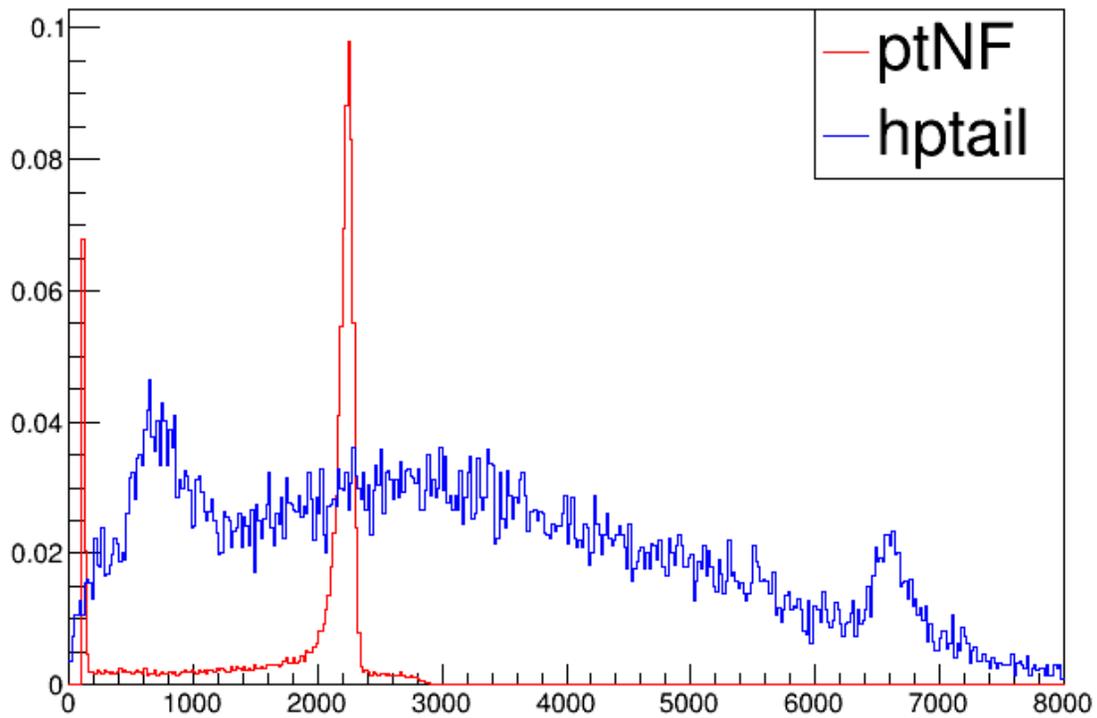
Figure 4.16: The blue spectrum from the tail-fitter was scaled up by a factor of 10 on the vertical axis so that both curves could be seen, and it was scaled down by a factor of 10 on the horizontal axis so that the energies could be plotted in the same range. This is largely a qualitative study, so these scalings have not influenced the results.
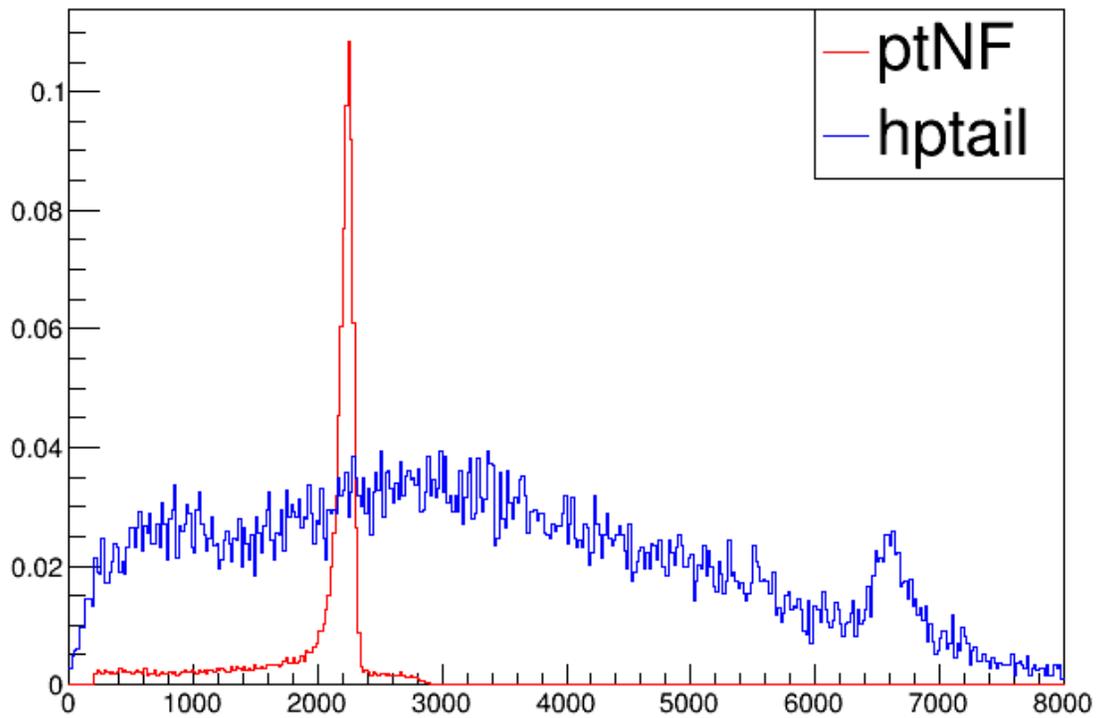
Figure 4.17: Same plot as Fig. 4.16, except the low energy cut was changed from $ptNF > 100$ keV to $ptNF > 200$ keV to remove the large spike in the red distribution. This removed the corresponding bump in the blue distribution below 1000, and leaves the higher energies untouched, which is the more important part of the spectrum for calibration.

Despite the existing calibration between channels and the overall energy scale, the reconstructed energy values in Fig. 4.15, Fig. 4.16, or Fig. 4.17 are not to be trusted. This is why energy units were avoided when referring to the reconstructed spectra. In fact, for these three plots, the reconstructed energy on the x-axis was scaled down by a factor of 10 so that the blue histograms could be more easily compared to the red histograms.

Figure 4.18 shows the reconstructed energy for each channel, and it appears that the relative calibration between channels needs adjustment. If the relative calibrations are off, then the 356 keV peak gets smeared out when co-adding channels; this is why the peak is wider in Fig. 4.16 and Fig. 4.17 than in Fig. 4.15.

The overall calibration depends on the definition of the tail integral, so any modification to that integral means the spectrum features appear at different energies. A new calibration constant is easy to calculate as the current calibration constant times the ratio of where the peak should be over the energy the peak appears at. The calibration constant could also be different for each detector. However, these calibration constants need to be defined in a configuration file that controls CDMSBats and is typically customized by each user that processes the data. It would be very difficult to roll out a broad change to every user that already has a custom configuration file.

For these reasons, the tail-fitter energy calibration should be left to the user, as it will be much easier to calibrate the spectrum after it is produced provided that any known feature is visible, such as the 356 keV peak. Since this is the total phonon energy summed over all four channels, then the NTL gain needs to be included when calculating the expected energy. For example, with a bias of 70 V the 356 keV peak should be at $356 \, \mathrm{keV_{ee}}(70/3+1) = 8662.\bar{6} \, \mathrm{keV_t}$.

D.4. Future Work

Appendix D.3 outlines that it is not worth the time and effort of calibrating this algorithm for the Soudan data that has already been processed and analyzed. However, it would
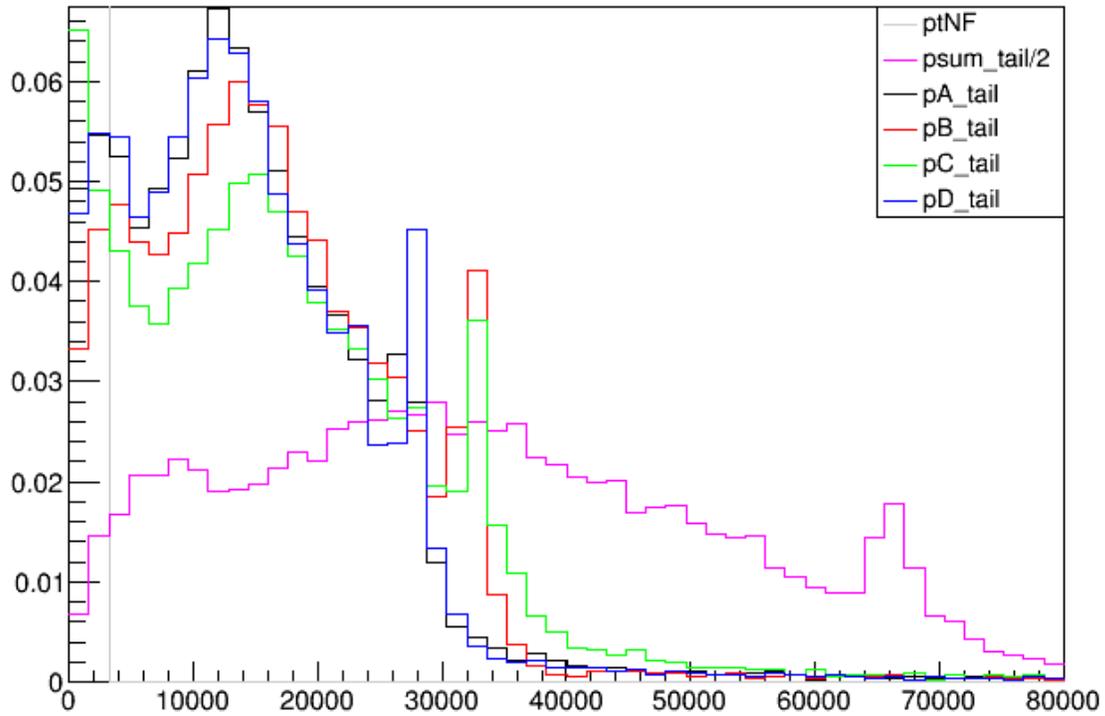
Figure 4.18: This plot shows the histograms of the reconstructed energy for each channel where black, red, green, and blue corresponds to channels A, B, C, and D, respectively. These histograms already have the relative calibration applied, and it can be seen that the 356 keV peak shows up around 2800 for channels A and D, while the peak is at 3200 for channels B and C. The magenta histogram is the sum of these (with $ptNF > 200$ keV cut), and despite the different binning, it can be compared directly with the blue histogram in Fig. 4.17. However, this plot has not been scaled down by a factor of 10 as the others, which explains why the gray ptNF histogram is so low by comparison.

certainly be ideal to have this algorithm reconstructing and outputting accurately calibrated energy quantities for SNOLAB data. I recommend that the calibration is done in conjunction with the development of the data processing, when it is less onerous to make changes to the configuration file.

The tail-fitting algorithm can be tested and developed even before SNOLAB is operational. There are a number of test facilities taking data with SuperCDMS detectors. I have been told they have plenty of saturated events, so this algorithm would be useful for reconstructing otherwise unattainable event energies.

Given the original motivation of this project, a crucial test would be checking the ability of this algorithm to calibrate data collected with a silicon detector when exposed to ^{133}Ba. It might also be possible to model high energy background models such as alpha decays, or to use this algorithm as a crude energy reconstruction when monitoring the quality of data as it is being collected.

Finally, I tried to modify this algorithm to be as smart and unrestricted as possible, but by construction, it was developed to work specifically on the two CDMSlite detectors tested. It would be very interesting to see how it behaves with some of the new technology being developed by SuperCDMS such as the HVeV chip with single electron-hole pair resolution [144].

BIBLIOGRAPHY

[1] K. A. Tomilin, *Natural systems of units. to the centenary anniversary of the planck system,* . 1

[2] P. N. Luke, *Voltage-assisted calorimetric ionization detector,* *Journal of Applied Physics* **64,12** (1988) 6858–6860. 2

[3] B. S. Neganov and V. N. Trofimov, *Colorimetric method measuring ionizing radiation,* *Otkryt. Izobret.* **146** (1985) 215. 2

[4] SuperCDMS collaboration, R. Agnese et al., *New Results from the Search for Low-Mass Weakly Interacting Massive Particles with the CDMS Low Ionization Threshold Experiment,* *Phys. Rev. Lett.* **116** (2016) 071301, [1509.02448]. xv, xx, 3, 72, 73, 89

[5] J. Lindhard, V. Nielsen, M. Scharff and P. Thomsen, *Integral equations governing radiation effects. (notes on atomic collisions, iii),* *Kgl. Danske Videnskab., Selskab. Mat. Fys. Medd.* **Vol: 33: No. 10** (1, 1963) . 4, 93

[6] J. D. Lewin and P. F. Smith, *Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil,* *Astropart. Phys.* **6** (1996) 87–112. 4, 24, 26

[7] CDMS collaboration, R. Agnese et al., *Nuclear-Recoil Energy Scale in CDMS II Silicon Dark-Matter Detectors,* *Nucl. Instrum. Meth.* **A905** (2018) 71–81, [1803.02903]. 4

[8] D. Barker, W. Z. Wei, D. M. Mei and C. Zhang, *Ionization Efficiency Study for Low Energy Nuclear Recoils in Germanium,* *Astropart. Phys.* **48** (2013) 8, [1304.6773]. xv, 4, 5

[9] A. E. Chavarria et al., *Measurement of the ionization produced by sub-keV silicon nuclear recoils in a CCD dark matter detector,* *Phys. Rev.* **D94** (2016) 082007, [1608.00957]. xv, 5

[10] J. Pivarski, *What do we mean by 'cross section' in particle physics?,* 3, 2013. 6

[11] J. Oort, *The force exerted by the stellar system in the direction perpendicular to the galactic plane and some related problems,* *Bulletin of the Astronomical Institutes of the Netherlands* **6** (Aug, 1932) 249. 7

[12] F. Zwicky, *Die Rotverschiebung von extragalaktischen Nebeln,* *Helvetica Physica Acta* **6** (1933) 110–127. 8

[13] V. C. Rubin and W. K. Ford, Jr., *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*, *apj* **159** (feb, 1970) 379. 8

[14] K. G. Begeman, A. H. Broeils and R. H. Sanders, *Extended rotation curves of spiral galaxies : dark haloes and modified dynamics.*, *Monthly Notices of the Royal Astronomical Society* **249** (Apr, 1991) 523. xv, 9

[15] T. Abel, O. Hahn and R. Kaehler, *Tracing the dark matter sheet in phase space*, *Monthly Notices of the Royal Astronomical Society* **427** (11, 2012) 61–76, [http://oup.prod.sis.lan/mnras/article-pdf/427/1/61/18239226/427-1-61.pdf]. 10

[16] R. Kaehler, O. Hahn and T. Abel, *A Novel Approach to Visualizing Dark Matter Simulations*, *IEEE Transactions on Visualization and Computer Graphics* **18** (Dec, 2012) 2078–2087. xv, 10

[17] A. A. Penzias and R. W. Wilson, *A Measurement of Excess Antenna Temperature at 4080 Mc/s.*, *Astrophysical Journal* **142** (jul, 1965) 419–421. 11

[18] R. Alpher and R. Herman, *Evolution of the Universe*, *Nature* **162** (Nov, 1948) 774–775. 11

[19] R. H. Dicke, P. J. E. Peebles, P. G. Roll and D. T. Wilkinson, *Cosmic Black-Body Radiation*, *Astrophysical Journal* **142** (Jul, 1965) 414–419. 11

[20] PLANCK collaboration, P. A. R. Ade et al., *Planck 2013 results. I. Overview of products and scientific results*, *Astron. Astrophys.* **571** (2014) A1, [1303.5062]. xv, 11, 12

[21] PLANCK collaboration, P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, 1502.01589. xv, 11, 12

[22] T. Clifton, P. G. Ferreira, A. Padilla and C. Skordis, *Modified Gravity and Cosmology*, *Phys. Rept.* **513** (2012) 1–189, [1106.2476]. 13

[23] M. Markevitch, A. H. Gonzalez, D. Clowe, A. Vikhlinin, L. David, W. Forman et al., *Direct constraints on the dark matter self-interaction cross-section from the merging galaxy cluster 1E0657-56*, *Astrophys. J.* **606** (2004) 819–824, [astro-ph/0309303]. 13, 14

[24] D. Clowe, M. Bradac, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones et al., *A direct empirical proof of the existence of dark matter*, *Astrophys. J.* **648** (2006) L109–L113, [astro-ph/0608407]. xvi, 13, 15

[25] A. Mahdavi, H. y. Hoekstra, A. y. Babul, D. y. Balam and P. Capak, *A Dark Core in Abell 520*, *Astrophys. J.* **668** (2007) 806–814, [0706.3048]. 14

[26] W. A. Dawson, D. Wittman, M. J. Jee, P. Gee, J. P. Hughes, J. A. Tyson et al., *DISCOVERY OF a DISSOCIATIVE GALAXY CLUSTER MERGER WITH LARGE PHYSICAL SEPARATION*, *The Astrophysical Journal* **747** (feb, 2012) L42. 14

[27] J. R. Primack, *Dark matter and structure formation*, in *Midrasha Mathematicae in Jerusalem: Winter School in Dynamical Systems Jerusalem, Israel, January 12-17, 1997*, 1997. astro-ph/9707285. 16

[28] G. Steigman and M. S. Turner, *Cosmological constraints on the properties of weakly interacting massive particles*, *Nuclear Physics B* **253** (1985) 375 – 386. 16

[29] G. Bertone, D. Hooper and J. Silk, *Particle dark matter: Evidence, candidates and constraints*, *Phys. Rept.* **405** (2005) 279–390, [hep-ph/0404175]. 17

[30] G. Steigman, B. Dasgupta and J. F. Beacom, *Precise Relic WIMP Abundance and its Impact on Searches for Dark Matter Annihilation*, *Phys. Rev.* **D86** (2012) 023506, [1204.3622]. xvi, 17, 18

[31] B. Penning, *The pursuit of dark matter at colliders - an overview*, *Journal of Physics G: Nuclear and Particle Physics* **45** (may, 2018) 063001. xvi, 20

[32] J. M. Gaskins, *A review of indirect searches for particle dark matter*, *Contemp. Phys.* **57** (2016) 496–525, [1604.00014]. 21

[33] T. Marrodán Undagoitia and L. Rauch, *Dark matter direct-detection experiments*, *J. Phys.* **G43** (2016) 013001, [1509.08767]. xvii, 23

[34] M. Schumann, *Dark Matter Search with liquid Noble Gases*, 1206.2169. 24

[35] J. Cooley, *Overview of Non-Liquid Noble Direct Detection Dark Matter Experiments*, *Phys. Dark Univ.* **4** (2014) 92–97, [1410.4960]. 24

[36] L. Necib, M. Lisanti and V. Belokurov, *Dark Matter in Disequilibrium: The Local Velocity Distribution from SDSS-Gaia*, 1807.02519. 25

[37] J. D. Vergados and H. Ejiri, *Can Solar Neutrinos be a Serious Background in Direct Dark Matter Searches?*, *Nucl. Phys.* **B804** (2008) 144–159, [0805.2583]. 27

[38] J. Billard, L. Strigari and E. Figueroa-Feliciano, *Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments*, *Phys. Rev.* **D89** (2014) 023524, [1307.5458]. 27

[39] C. A. J. O'Hare, A. M. Green, J. Billard, E. Figueroa-Feliciano and L. E. Strigari, *Readout strategies for directional dark matter detection beyond the neutrino background*, *Phys. Rev.* **D92** (2015) 063518, [1505.08061]. 27

[40] F. Ruppin, J. Billard, E. Figueroa-Feliciano and L. Strigari, *Complementarity of dark matter detectors in light of the neutrino background*, *Phys. Rev.* **D90** (2014) 083510, [1408.3581]. xvii, 27, 28

[41] D. Akimov, J. B. Albert, P. An, C. Awe, P. S. Barbeau, B. Becker et al., *Observation of coherent elastic neutrino-nucleus scattering*, *Science* **357** (2017) 1123–1126, [https://science.sciencemag.org/content/357/6356/1123.full.pdf]. 27

[42] G. Adhikari, P. Adhikari, E. B. de Souza, N. Carlin, S. Choi, M. Djamal et al., *An experiment to search for dark-matter interactions using sodium iodide detectors*, *Nature* **564** (2018) 83–86. xvii, 27

[43] COUPP COLLABORATION collaboration, E. Behnke, J. Behnke, S. J. Brice, D. Broemmelsiek, J. I. Collar, A. Conner et al., *First dark matter search results from a 4-kg* cf_3I *bubble chamber operated in a deep underground site*, *Phys. Rev. D* **86** (Sep, 2012) 052001. xvii, 27

227

[44] M. Felizardo, T. A. Girard, T. Morlat, A. C. Fernandes, A. R. Ramos and J. G. Marques, *Recent results from the simple dark matter search*, *Journal of Physics: Conference Series* **375** (2012) 012011. xvii, 27

[45] E. Behnke, M. Besnier, P. Bhattacharjee, X. Dai, M. Das, A. Davour et al., *Final results of the PICASSO dark matter search experiment*, *Astropart. Phys.* **90** (apr, 2017) 85–92. xvii, 27

[46] C. Amole, M. Ardid, I. J. Arnquist, D. M. Asner, D. Baxter, E. Behnke et al., *Dark Matter Search Results from the Complete Exposure of the PICO-60 C$_3$F$_8$ Bubble Chamber*, *arXiv e-prints* (Feb, 2019) arXiv:1902.04031, [1902.04031]. xvii, 27, 28

[47] LUX Collaboration collaboration, D. S. Akerib, S. Alsum, H. M. Araújo, X. Bai, A. J. Bailey, J. Balajthy et al., *Results from a search for dark matter in the complete lux exposure*, *Phys. Rev. Lett.* **118** (Jan, 2017) 021303. xvii, 27

[48] D. Akimov, H. Araújo, E. Barnes, V. Belov, A. Bewick, A. Burenkov et al., *Wimp-nucleon cross-section results from the second science run of zeplin-iii*, *Physics Letters B* **709** (2012) 14 – 20. xvii, 27

[49] XENON Collaboration 7 collaboration, E. Aprile, J. Aalbers, F. Agostini, M. Alfonsi, L. Althueser, F. D. Amaro et al., *Dark matter search results from a one ton-year exposure of xenon1t*, *Phys. Rev. Lett.* **121** (Sep, 2018) 111302. xvii, 27, 28

[50] DAMIC Collaboration collaboration, A. Aguilar-Arevalo, D. Amidei, X. Bertou, M. Butner, G. Cancelo, A. Castañeda Vázquez et al., *Search for low-mass wimps in a 0.6 kg day exposure of the damic experiment at snolab*, *Phys. Rev. D* **94** (Oct, 2016) 082006. xvii, 27

[51] TEXONO Collaboration collaboration, H. B. Li, H. Y. Liao, S. T. Lin, S. K. Liu, L. Singh, M. K. Singh et al., *Limits on Spin-Independent Couplings of WIMP Dark Matter with a p -Type Point-Contact Germanium Detector*, *Phys. Rev. Lett.* **110** (jun, 2013) 261301. xvii, 27

[52] The DarkSide Collaboration, *DarkSide-50 532-day Dark Matter Search with Low-Radioactivity Argon*, *ArXiv e-prints* (Feb., 2018) , [1802.07198]. xvii, 27

[53] CRESST collaboration, A. H. Abdelhameed et al., *First results from the CRESST-III low-mass dark matter program*, 1904.00498. xvii, 27, 28

[54] SuperCDMS collaboration, R. Agnese et al., *Search for Low-Mass Dark Matter with CDMSlite Using a Profile Likelihood Fit*, *Phys. Rev.* **D99** (2019) 062001, [1808.09098]. xvii, xxiii, 28, 91, 92, 93, 94, 104

[55] E. Aprile, J. Aalbers, F. Agostini, M. Alfonsi, F. D. Amaro, M. Anthony et al., *Physics reach of the xenon1t dark matter experiment.*, *Journal of Cosmology and Astroparticle Physics* **2016** (2016) 027. xvii, 28

[56] J. Aalbers et al., *DARWIN: towards the ultimate dark matter detector*, *ArXiv e-prints* (June, 2016) , [1606.07001]. xvii, 28

[57] M. Battaglieri et al., *US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report*, in *U.S. Cosmic Visions: New Ideas in Dark Matter College Park, MD, USA, March 23-25, 2017*, 2017. 1707.04591. 29, 31

[58] L. J. Hall, K. Jedamzik, J. March-Russell and S. M. West, *Freeze-In Production of FIMP Dark Matter*, *JHEP* **03** (2010) 080, [0911.1120]. xvii, 29, 30

[59] D. E. Kaplan, M. A. Luty and K. M. Zurek, *Asymmetric Dark Matter*, *Phys. Rev.* **D79** (2009) 115016, [0901.4117]. 29

[60] K. Petraki and R. R. Volkas, *Review of asymmetric dark matter*, *Int. J. Mod. Phys.* **A28** (2013) 1330028, [1305.4939]. 29

[61] K. M. Zurek, *Asymmetric Dark Matter: Theories, Signatures, and Constraints*, *Phys. Rept.* **537** (2014) 91–121, [1308.0338]. 29

[62] H. M. Hodges, *Mirror baryons as the dark matter*, *Phys. Rev. D* **47** (Jan, 1993) 456–459. 31

[63] R. Foot, *Mirror dark matter: Cosmology, galaxy structure and direct detection*, *Int. J. Mod. Phys.* **A29** (2014) 1430013, [1401.3965]. 31

[64] M. Raggi and V. Kozhuharov, *Results and perspectives in dark photon physics*, *Riv. Nuovo Cim.* **38** (2015) 449–505. 31

[65] R. Essig, J. Mardon and T. Volansky, *Direct Detection of Sub-GeV Dark Matter*, *Phys. Rev.* **D85** (2012) 076007, [1108.5383]. 31

[66] SuperCDMS collaboration, R. Agnese et al., *First Dark Matter Constraints from a SuperCDMS Single-Charge Sensitive Detector*, *Phys. Rev. Lett.* **121** (2018) 051301, [1804.10697]. 31

[67] "Radionuclide basics: Radon." https://www.epa.gov/radiation/radionuclide-basics-radon. 34

[68] S. Singh, A. K. Jain and J. K. Tuli, *Nuclear Data Sheets for A = 222*, *Nuclear Data Sheets* **112** (Nov., 2011) 2851–2886. 35

[69] A. K. Jain and B. Singh, *Nuclear Data Sheets for A = 218*, *Nuclear Data Sheets* **107** (Apr., 2006) 1027–1074. 35

[70] S.-C. Wu, *Nuclear Data Sheets for A = 214*, *Nuclear Data Sheets* **110** (Mar., 2009) 681–748. 35

[71] M. Shamsuzzoha Basunia, *Nuclear Data Sheets for A = 210*, *Nuclear Data Sheets* **121** (Sept., 2014) 561–694. 35

[72] F. G. Kondev, *Nuclear Data Sheets for A = 206*, *Nuclear Data Sheets* **109** (June, 2008) 1527–1654. 35

[73] E. Hoppe, A. Seifert, C. E. Aalseth, P. P. Bachelor, A. Day, D. Edwards et al., *Cleaning and passivation of copper surfaces to remove surface radioactivity and prevent oxide formation*, *Nuclear Instruments and Methods in Physics Research Section A-accelerators Spectrometers Detectors and Associated Equipment - NUCL INSTRUM METH PHYS RES A* **579** (08, 2007) 486–489. 35

[74] R. W. Schnee, R. Bunker, G. Ghulam, D. Jardin, M. Kos and A. S. Tenney, *Construction and measurements of a vacuum-swing-adsorption radon-mitigation system*, *AIP Conf. Proc.* **1549** (2013) 116–119, [1404.5811]. 35

[75] J. Street, R. Bunker, C. Dunagan, X. Loose, R. W. Schnee, M. Stark et al., *Construction and Measurements of an Improved Vacuum-Swing-Adsorption Radon-Mitigation System*, *AIP Conf. Proc.* **1672** (2015) 150004, [1506.00929]. 35

[76] W. Jacobi, *Activity and potential alpha-energy of 222-radon and 220-radon daughters in different air atmospheres*, *Health Physics* **22** (5, 1972) 441–450. 37

[77] N. Stevanovic, V. Markovic, V. Urosevic and D. Nikezic, *Determination of parameters of the jacobi room model using the brownian motion model*, *Health physics* **96** (02, 2009) 48–54. 37, 38, 39, 41

[78] R. Schnee, "LZ internal document:The Jacobi Model for Radon Daughter Plateout." 6, 2016. 37, 40, 41, 44

[79] P. Scofield, *Radon Decay Product In-Door Behaviour – Parameter, Measurement Method, and Model Review*, . 39

[80] M. Stein, D. Bauer, R. Bunker, R. Calkins, J. Cooley, B. Loer et al., *Radon daughter plate-out measurements at snolab for polyethylene and copper*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **880** (2018) 92 – 97. 46

[81] B. LaFerriere, T. Maiti, I. Arnquist and E. Hoppe, *A novel assay method for the trace determination of th and u in copper and lead using inductively coupled plasma mass spectrometry*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **775** (03, 2015) . 48

[82] A. Seifert et al., *The design, construction, and initial characterization of an ultra-low-background gas-proportional counting system*, *Journal of Radioanalytical and Nuclear Chemistry* **296** (2012) 915–921. xxxii, 49, 167

[83] CDMS Collaboration collaboration, D. Abrams, D. S. Akerib, M. S. Armel-Funkhouser, L. Baudis, D. A. Bauer, A. Bolozdynya et al., *Exclusion limits on the wimp-nucleon cross section from the cryogenic dark matter search*, *Phys. Rev. D* **66** (Dec, 2002) 122003. 50

[84] CDMS Collaboration collaboration, D. S. Akerib, M. J. Attisha, L. Baudis, D. A. Bauer, A. I. Bolozdynya, P. L. Brink et al., *Low-threshold analysis of cdms shallow-site data*, *Phys. Rev. D* **82** (Dec, 2010) 122004. 50

[85] D. S. Akerib et al., *Installation and commissioning of the CDMSII experiment at Soudan*, *Nucl. Instrum. Meth.* **A520** (2004) 116–119. 50

[86] CDMS collaboration, D. S. Akerib et al., *First results from the cryogenic dark matter search in the Soudan Underground Lab*, *Phys. Rev. Lett.* **93** (2004) 211301, [astro-ph/0405033]. 50, 51

[87] M. C. Pyle, *Optimizing the design and analysis of cryogenic semiconductor dark matter detectors for maximum sensitivity.* PhD thesis, Stanford U., 2012. 10.2172/1127926. 51, 54

[88] MINOS collaboration, J. Evans, *The MINOS Experiment: Results and Prospects*, *Adv. High Energy Phys.* **2013** (2013) 182537, [1307.0721]. 51

[89] R. Bunker, *A low-threshold analysis of data from the cryogenic dark matter search experiment*, . 51

[90] J. Sander, *Results from the Cryogenic Dark Matter Search Using a Chi Squared Analysis*. PhD thesis, UC, Santa Barbara, 2007. 10.2172/1415815. xviii, 52

[91] SuperCDMS collaboration, R. Agnese et al., *Demonstration of surface electron rejection with interleaved germanium detectors for dark matter searches*, *Applied Physics Letters* **103** (2013) –. xix, 57

[92] J. Filippini, "SuperCDMS internal document: Hunting for 66.7 and 511 keV Ge peaks." 11, 2007. 58

[93] M. D. Pepin, *Low-Mass Dark Matter Search Results and Radiogenic Backgrounds for the Cryogenic Dark Matter Search*. PhD thesis, Minnesota U., 2016. 10.2172/1342205. xx, 59, 67, 73, 75

[94] W. A. Page, *Searching for low-mass dark matter with SuperCDMS Soudan detectors*. PhD thesis, U. of British Columbia, 2019. 10.14288/1.0376351. 60, 92, 93

[95] *SNOLAB User's Handbook*, 2 ed., June, 2006. 61

[96] F. Duncan, A. Noble and D. Sinclair, *The Construction and Anticipated Science of SNOLAB*, *Annual Review of Nuclear and Particle Science* **60** (2010) 163–180, [https://doi.org/10.1146/annurev.nucl.012809.104513]. 61

[97] N. R. Council, ed., *Connecting Quarks with the Cosmos: Eleven Science Questions for the New Century*, ch. 7 Realizing the Opportunities. The National Academies Press, Washington, DC, 2003. 10.17226/10079. xix, 62

[98] SuperCDMS collaboration, R. Agnese et al., *Projected Sensitivity of the SuperCDMS SNOLAB experiment*, *Phys. Rev.* **D95** (2017) 082002, [1610.00006]. xix, 62, 63, 64, 65, 66, 204

[99] N. A. Kurinsky, *The Low-Mass Limit: Dark Matter Detectors with eV-Scale Energy Resolution*. PhD thesis, Stanford U., Dept. Phys., 2018. 10.2172/1472104. 64, 214

[100] SuperCDMS collaboration, R. Agnese et al., *Search for Low-Mass Weakly Interacting Massive Particles Using Voltage-Assisted Calorimetric Ionization Detection in the SuperCDMS Experiment*, *Phys. Rev. Lett.* **112** (2014) 041302, [1309.3259]. xx, 67, 68, 69

[101] R. Basu Thakur, *The Cryogenic Dark Matter Search low ionization-threshold experiment*. PhD thesis, Illinois U., Urbana, 2014. 10.2172/1182551. xx, 68, 69

[102] S. Yellin, *Finding an upper limit in the presence of an unknown background*, *Phys. Rev. D* **66** (Aug, 2002) 032005. 72

[103] S. Yellin, *Extending the optimum interval method*, 0709.2701. 72

[104] SuperCDMS collaboration, R. Agnese et al., *Low-mass dark matter search with CDMSlite*, *Phys. Rev.* **D97** (2018) 022002, [1707.01632]. 73, 85, 116

[105] SUPERCDMS collaboration, R. Agnese et al., *Production Rate Measurement of Tritium and Other Cosmogenic Isotopes in Germanium with CDMSlite*, *Astropart. Phys.* **104** (2019) 12, [1806.07043]. 75, 85

[106] E. Fascione, *Measurement of Cosmogenic Production Rates in Ge with CDMSlite*, Master's thesis, Queen's University, Aug, 2018. 75

[107] D. R. Barker, *SuperCDMS Background Models for Low-Mass Dark Matter Searches*. PhD thesis, Minnesota U., 2018. 76, 77, 80, 94

[108] E. Parzen, *On estimation of a probability density function and mode*, *Ann. Math. Statist.* **33** (09, 1962) 1065–1076. 77

[109] M. Rosenblatt, *Remarks on some nonparametric estimates of a density function*, *Ann. Math. Statist.* **27** (09, 1956) 832–837. 77

[110] C. C. Evirgen, *"statistical geometry of the magnetic fields in the multi-phase, turbulent interstellar medium"*, 01, 2016. 10.13140/RG.2.1.4238.4402. xx, 78

[111] D. Scott, *Multivariate Density Estimation: Theory, Practice, and Visualization*. John Wiley and Sons, 1992. 78

[112] H. Shimazaki and S. Shinomoto, *Kernel bandwidth optimization in spike rate estimation*, *Journal of Computational Neuroscience* **29** (Aug, 2010) 171–182. 79

[113] A. Kubik, "SuperCDMS internal document: Surface Alpha Rate Study." 10, 2015. 79

[114] DARKSIDE COLLABORATION collaboration, P. Agnes, I. F. M. Albuquerque, T. Alexander, A. K. Alton, G. R. Araujo, D. M. Asner et al., *Low-Mass Dark Matter Search with the DarkSide-50 Experiment*, *Phys. Rev. Lett.* **121** (Aug, 2018) 081307. xxiii, 94

[115] LUX COLLABORATION collaboration, D. S. Akerib, S. Alsum, H. M. Araújo, X. Bai, J. Balajthy, P. Beltrame et al., *Results of a search for sub-gev dark matter using 2013 lux data*, *Phys. Rev. Lett.* **122** (Apr, 2019) 131301. xxvii, 108, 120, 140

[116] C. Kouvaris and J. Pradler, *Probing sub-GeV Dark Matter with conventional detectors*, *Phys. Rev. Lett.* **118** (2017) 031803, [1607.01789]. xxvii, 109

[117] L. B. N. L. C. for X-ray Optics and A. L. Source, *X-ray data booklet*, October, 2009. 112

[118] M. Wilson, "SuperCDMS internal document: Existing Data Concerning the Photoelectric Absorption Cross Section in Ge and Si." 11, 2017. 112

[119] B. Henke, E. Gullikson and J. Davis, *X-ray interactions: Photoabsorption, scattering, transmission, and reflection at e = 50-30,000 ev, z = 1-92*, *Atomic Data and Nuclear Data Tables* **54** (1993) 181 – 342. 112

[120] A. Migdal, *Ionization of atoms accompanying alpha- and beta- decay, J. Phys.* **4** (1941) 449–453. 114

[121] M. Ibe, W. Nakano, Y. Shoji and K. Suzuki, *Migdal Effect in Dark Matter Direct Detection Experiments*, *JHEP* **03** (2018) 194, [1707.07258]. 116, 120

[122] M. J. Dolan, F. Kahlhoefer and C. McCabe, *Directly detecting sub-GeV dark matter with electrons from nuclear scattering*, 1711.09906. xxvii, 116, 117, 120

[123] EDELWEISS COLLABORATION collaboration, E. Armengaud, C. Augier, A. Benoît, A. Benoit, L. Bergé, J. Billard et al., *Searching for low-mass dark matter particles with a massive ge bolometer operated above ground*, *Phys. Rev. D* **99** (Apr, 2019) 082003. 120, 141

[124] T. Emken and C. Kouvaris, *How blind are underground and surface detectors to strongly interacting dark matter?*, *Phys. Rev. D* **97** (Jun, 2018) 115047. xxviii, 122, 123

[125] B. J. Kavanagh, *Earth scattering of superheavy dark matter: Updated constraints from detectors old and new*, *Phys. Rev.* **D97** (2018) 123013, [1712.04901]. 124

[126] G. Cowan, *Statistical data analysis.* Oxford university press, 1998. 129

[127] F. James and M. Roos, *Minuit – a system for function minimization and analysis of the parameter errors and correlations*, *Computer Physics Communications* **10** (Dec., 1975) 343–367. 130

[128] iminuit team, "iminuit – a python interface to minuit." https://github.com/scikit-hep/iminuit. 130

[129] S. S. Wilks, *The large-sample distribution of the likelihood ratio for testing composite hypotheses*, *Ann. Math. Statist.* **9** (03, 1938) 60–62. 131

[130] CRESST collaboration, A. H. Abdelhameed et al., *First results from the CRESST-III low-mass dark matter program*, 1904.00498. 141

[131] X. LLC, *Ultralo-1800 alpha particle counter: User's manual*, 5, 2013. xxxi, 147, 148, 149, 150, 151

[132] B. Cabrera and S. Coleman, "SuperCDMS internal document: Alpha Counter Study." 7, 2009. 155

[133] R. Brun and F. Rademakers, *Root — an object oriented data analysis framework*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **389** (1997) 81 – 86. 158

[134] A. Posada, "SMU undergraduate thesis: Pulse-shape Discrimination Studies of Th-230 Decay Alpha Particles Using the XIA UltraLo-1800." 4, 2017. xxxii, 161

[135] C. E. Aalseth et al., *A shallow underground laboratory for low-background radiation measurements and materials development*, *Review of Scientific Instruments* **83** (2012) 113503, [https://doi.org/10.1063/1.4761923]. 168, 175

[136] SUPERCDMS AND LZ collaboration, R. Schnee, R. Bunker, J. Street, X. Loose, C. Dunagan, D. Tronstad et al., *The Mines Radon Emanation Detection System*, *AIP Conf. Proc.* **1672** (2015) . 175

[137] B. J. Mount et al., *LUX-ZEPLIN (LZ) Technical Design Report*, 1703.09144. 176

[138] B. Welliver, "SuperCDMS internal document: R133 Charge Noise Sigmas From Low Bg Using Optimal Filter 0 Quantities." 4, 2012. 178

[139] B. Welliver, "SuperCDMS internal document: R133 Phonon Noise Blob From Low Bg Using Optimal Filter 0 Quantities." 4, 2012. 178

[140] SuperCDMS Collaboration collaboration, R. Agnese, T. Aramaki, I. J. Arnquist, W. Baker, D. Balakishiyeva, S. Banik et al., *Results from the super cryogenic dark matter search experiment at soudan*, *Phys. Rev. Lett.* **120** (Feb, 2018) 061802. 178

[141] B. Welliver, *Dedicated Searches for Low and High Mass Wimps with the SuperCDMS Soudan iZIP Detectors*. PhD thesis, U. of Florida, 2016. 10.2172/1350525. 178

[142] T. Saab and L. Baudis, "SuperCDMS internal document: Measuring the Abundance and Decay Rate of the 10 keV Triplet." 5, 2002. 185

[143] R. Bunker, *Email correspondence: Resolution function*, 10, 2015. 199

[144] SuperCDMS collaboration, R. Agnese et al., *First Dark Matter Constraints from a SuperCDMS Single-Charge Sensitive Detector*, *Phys. Rev. Lett.* **121** (2018) 051301, [1804.10697]. 224