USING THERMAL NEUTRON CAPTURE TO CALIBRATE FOR DARK MATTER DETECTION

by

KITTY C. MICKELSON

B.S. University of Colorado Denver, 2020

A thesis submitted in partial fulfillment

of the requirements for the degree of

Master of Integrated Sciences

Integrated Sciences Program

2025

University of Colorado Denver

This thesis for the Master's degree by

Kitty C. Mickelson

has been approved for the

Integrated Sciences Program

by

Anthony Villano, Chair

Amy Roberts

Erin Austin

Date May 17, 2025

Mickelson, Kitty C. (M.I.S.)

Using Thermal Neutron Capture to Calibrate for Dark Matter Detection

Thesis directed by Dr. Anthony Villano

ABSTRACT

The search for dark matter – the possible "missing mass" implied by various measurements of celestial bodies [1, 2] – has historically been restricted by detector capabilities [3], especially by detectors' ability to distinguish energies below 1 keV from one another [4]. As detector technology advances and these energies become easier to distinguish, the ability to calibrate detectors at low energies becomes vital. Unfortunately, most existing calibration methods do not translate well to this low-energy region [5]. Therefore, we developed a calibration process for Si detectors using thermal neutron capture, which results in nuclear recoils at the sub-keV level. The calibration is kept as general as possible to allow easier adaptation and is entirely independent of the thermal neutron source; we fully expect to be able to extend this work to Ge detectors and believe it can be extended to other detector materials as well, albeit with more adjustments. Calibration is broken into the following steps: simulation of expected energy spectra, including both background and calibration sources along with multiple possible detector response models; data acquisition; and statistical and probabilistic analysis of the resulting data compared to the predicted models. The data used in this work did not have enough low-energy peaks for a comparative evaluation of yield models, but I successfully demonstrate the detection of neutron capture-induced nuclear recoils and measurement of their yield values, demonstrating the viability and method of such a comparison.

Approved: Anthony Villano

"The way I see it, the Mountain can’t bring out anything that isn’t already in you."

— Maddy Thorson, *Celeste*

ACKNOWLEDGMENTS

I gratefully acknowledge support we received from the U.S. Department of Energy (DOE) Office of High Energy Physics and from the National Science Foundation (NSF). This work was supported by DOE Grant No. DE-SC0021364 and in part by NSF Grant No. 2111090.

On a more personal note, thanks to everyone who mentored me. Thanks to Anthony Villano for all the time we spent ironing out code and equations, to Amy Roberts for your endless patience and always being able to point me towards a solution, to Erin Austin for teaching me statistics both in and out of class and helping me ensure I have implemented those lessons correctly, and to Martin Huber for helping me navigate all of the administrative tasks and keeping me on track.

I also thank Himangshu Neog at UMN for entertaining many seemingly small – but still crucial – requests regarding the hardware. My eternal gratitude goes to Sten Hansen at FNAL for his wizardry with our DCRC's firmware and always finding the time to help us troubleshoot.

Thank you to all of my peers who worked alongside me, whether in the lab or on coursework. In particular, AJ Biffl, thanks for always being someone I could look up to and reminding me that I had this in me. Ana Gevorgian, you've been wonderful to work with, and both your hard work and stalwart positivity have made this work feel surmountable. And, Layne Hellman – thanks for showing me that I can still do the things I want to do without sacrificing the chance to be myself.

Last, but most certainly not least, I appreciate all of my friends for your endless support and for giving me a reason to keep going.

TABLE OF CONTENTS

LIST OF TABLES

TABLE

LIST OF FIGURES

FIGURE

NOTATION

Before we begin, we should clarify certain points regarding notation. Unless otherwise stated, this work uses natural units; that is, the speed of light c, the elementary charge q_e, and the Planck constant $\hbar$ are treated as having a value of unity. This unit system allows us to treat both masses and momenta using energy units and makes a photon's energy and momentum equivalent to one another; both of these facts are convenient for our calculations. For those who are not familiar with typical physics notation, Table 1 contains some common variables and subscripts that will be used consistently throughout this work. Variables that are not explicitly written in vector notation are either scalars or the scalar components of vectors.

Table 1: Common Physics Variables

Variable	Name	Usage
p	momentum	vector
m	mass	scalar
T	kinetic energy	scalar
γ	gamma particle	subscript

I. INTRODUCTION

SuperCDMS, the Cryogenic Dark Matter Search, uses cold, solid-state detectors to attempt direct detection of dark matter [6–8]. Generally speaking, "dark matter" is a term to describe matter we cannot detect through electromagnetic interactions and serves as a placeholder for mass that appears to be missing because it must exist according to gravitational measurements [1]. When attempting direct detection, we operate under the assumption that dark matter consists, at least in part, of Weakly-Interacting Massive Particles (WIMPs) – low-energy, non-relativistic particles whose exact properties are unknown, but may fit one of several existing models for theorized particles [9–11]. Being able to confirm (or definitively deny) the existence of these WIMPs would be a massive leap forward in understanding our universe, including understanding the motion of celestial bodies in galaxies and the large-scale structure of galaxy clustering [2]. However, to understand the signals from our detectors, we must first calibrate them, and the calibrations presently available for nuclear recoils are at a higher energy than where we expect to find interactions caused by dark matter. This means that small discrepancies in our existing calibrations may lead to large discrepancies at the region of interest [12], and we have no way of knowing whether a seemingly perfect calibration in the high-energy region translates to the low-energy region at all. Therefore, we must perform additional calibrations at lower energies.

We are most interested in characterizing the relationship between the energy deposited in a silicon detector and that detector's response in the low-energy region, that is, the region where detector responses are below 1 keV in ionization energy. This region may be of similar energies to dark matter, but there is some disagreement regarding the detector response in this region [13], making it exceedingly difficult to confirm any potential dark matter detection; either the predominant theoretically-driven yield model is incorrect for low energies, or there are other factors at play that influence the yield that have not historically been accounted for [14]. We aim to calibrate this energy region below where tried and proven methods work; the majority of methods rely on higher-energy reactions that only allow detector responses on the order of 1 keV or larger [12, 15]. While there are some explorations of neutron capture-induced nuclear recoils, many are indirect [5, 16], and those that are not explore the phenomenon in other materials, not for silicon [14, 17]; such studies utilize different methods than presented here, even in the case of direct detection.

I focus heavily on silicon detectors in this study, as they are what was available for calibration. However, I lay out all methodology in a way that is agnostic to the particular detector element, allowing the methodology, albeit not the results, to be applicable to any monatomic solid-state detector. Where possible, I also avoid assumptions specific to monatomic or solid-state systems, although this is more difficult to achieve. This work can also be used, more generally, to discern background noise for dark matter and other rare event searches in this region by generating predicted spectra for the desired detector.

The research presented here is divided into stages representing an experiment run from beginning to end. The first step, comprising Chapter II, is the simulation of the detector signal, including the ionization yield; this yield is a function describing the relationship between the energy deposited as the result of a nuclear recoil and the resulting effect on the electrons in the detector lattice and discussed more thoroughly in section II.1.3. Following simulation is real data acquisition, in Chapter III, including during-run analysis automated through firmware (Section III.2) or software; After acquisition are data processing, such as cuts (Section IV.1) to determine what data is accurate, real, and relevant; and, alongside the processing, analysis, such as the actual calibration (Section IV.2). Together, processing and analysis make up Chapter IV.

1 Neutron Capture for Calibration

We need to calibrate because the CDMS detectors do not output an actual energy value when a particle, WIMP or otherwise, interacts with them. When collecting data, detectors instead report an ADC (Analog-to-Digital Converter) bin value corresponding to an unknown energy; that is, a bin value is reported each time there is an event resulting in a value above a user-specified trigger level. We record a list of these ADC values, which we need to convert to energies in order for the data to have any clear physical meaning; calibration therefore means deriving and understanding the relationship between these reported ADC bin numbers and the sizes of the energy deposits that caused them. There will always be some uncertainty on the calibration, even when using reference points with singular, known quantities, because noise in the detector slightly alters the bin values.

The calibration process has three main components: modeling, data collection, and parameter extraction. The modeling component involves synthesizing physical models for multiple, related processes to create an expected spectrum in terms of energies rather than bin numbers – this synthesis is discussed in greater detail in Chapter II. The spectrum arises because the energy supplied is not the full energy of the incoming thermal neutron; there are many probabilistic factors that come into play that influence how the nuclei within the detector's lattice respond, or recoil, when struck by incoming neutrons. I construct the spectrum through a combination of probabilistic analysis (where practical) and simulation (where pure analysis is impractical and simulation is determined to have minimal impact). The collection of data for the calibration requires custom software and firmware, both of which we write based on the known and synthesized models. Details on the firmware are provided in Section 2 of Chapter III. The parameter extraction component of the calibration process consists of statistically comparing the acquired data with the aforementioned models of detector behavior and with the resulting spectrum; parameter extraction is covered in Chapter IV.

1.1 Neutron Capture

For calibration, we will utilize neutron capture because it generates nuclear recoils with energy values in the region we wish to calibrate. Whatever their source, neutrons in the area or from a directed source pass through the detector (see Fig I.1 for an example with a radioactive source). These neutrons then have a chance to be captured by one of the nuclei in the detector, meaning that the neutron becomes part of the nucleus. When working with thermal neutrons, the initial kinetic energy is nearly zero in comparison to the separation energy, which is the energy required to separate the neutron from the nucleus and the energy gained when binding them [18]. Therefore, due to conservation of momentum, when energy is released in the capture process, it results in the nucleus "recoiling" from an emitted gamma (see Fig I.2). These two motions are effectively back-to-back with one another with equal and opposite momenta and a combined total energy equal to the separation energy. For the case where the neutron captures directly to the ground state, the recoiling nucleus' initial kinetic energy is therefore

$$T \cong \frac{S_n^2}{2M}, \tag{1.1}$$

where M is the mass of the nucleus. This equation is derived in Appendix A and represents the greatest possible initial kinetic energy. The recoiling nucleus deposits kinetic energy into the detector as it slows down; this deposited kinetic energy is what is read out by the detector as a phonon or ionization energy. However, the captured neutron can exist at different excitation levels, and, if it does not capture directly to the ground state, the excitation level can change while or after the nucleus is slowing down. This process is discussed in greater detail in Section II.1; the general result is that the energies deposited exist on a spectrum [19], which I derive and present in Chapter II, and the shape of the spectrum is what allows us to use the nuclear recoils to calibrate the detector.

1.2 Simulation

Various software packages for simulation of neutron captures already exist. The most general such software is GEANT4 [20], which simulates interactions that occur when particles pass through matter, the type and shape of which can be specified by the user. GEANT4 simulation capabilities include neutron capture within a crystal, provided that the nuclear data for the isotope in question is complete and fully programmed for. There are also more specific tools for nuclear interactions; for example, DICEBOX [21] is a Fortran program for simulation

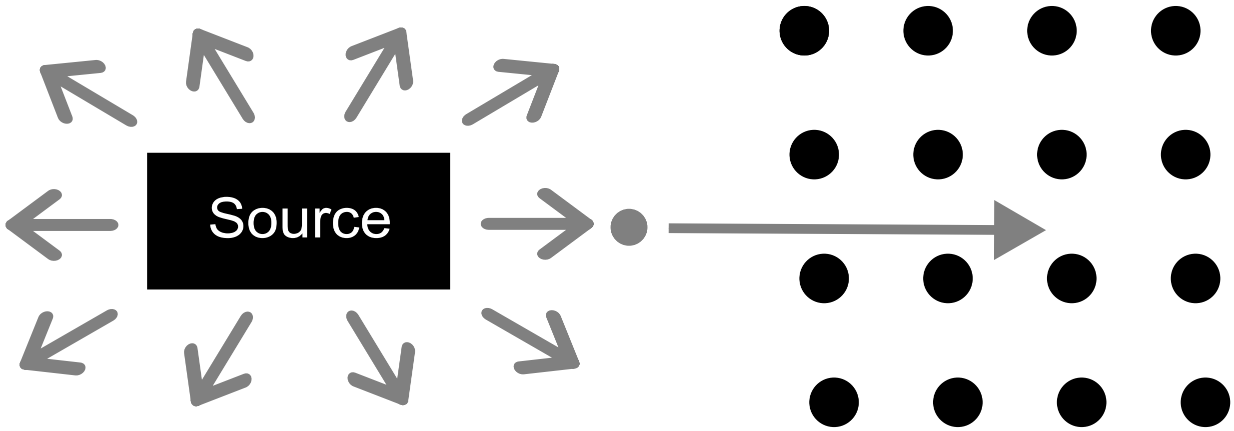

Figure I.1: Radiation Entering Detector

A simplified, two-dimensional demonstration of radiation entering a detector. The detector (right) is made up of nuclei, represented as circles. A source (left) emits particles (grey) in all directions. Many of the emitted particles will not reach the detector because they are emitted in the wrong direction. Those that do enter the detector have a chance to interact with one of the nuclei, but also have a (usually small) chance to pass through the detector without interacting.

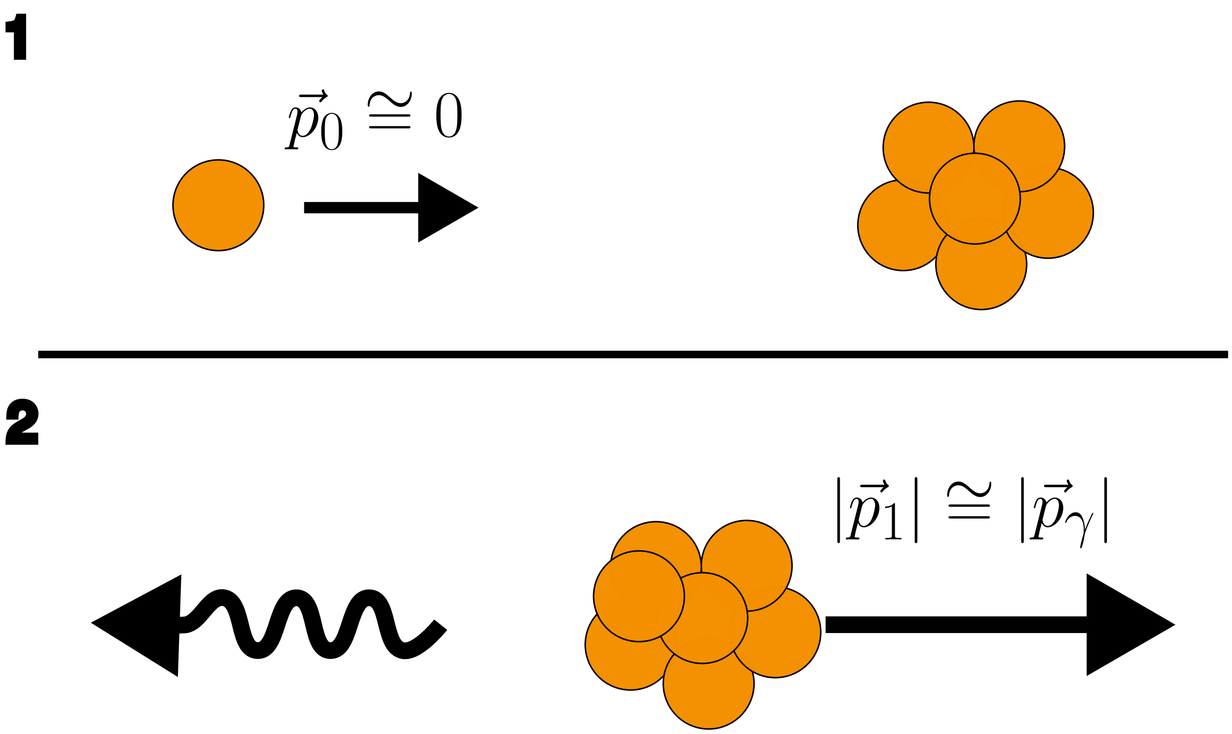

Figure I.2: Back-to-Back Recoil

1: The initial momentum, p_0, of the thermal neutron is very nearly zero. 2: The sum of the resulting momenta p_1 and p_γ must be zero, causing the nucleus to recoil in a direction opposite the emitted gamma.

of neutron capture and subsequent gammas in cases where the energy levels are dense and not well-known. FIFRELIN [22] similarly simulates emissions following de-excitation using level densities; however, it is built with fission in mind, not neutron capture. In contrast to these programs that use level densities, nrCascadeSim [19] is a single-purpose package specifically for simulation of NRs resulting from neutron capture where the energy levels are sufficiently well-known to be treated as fully discrete. Therefore, despite the other aforementioned programs being more established, this work will predominately utilize nrCascadeSim for simulation, in large part to avoid unnecessarily approximating states as continuous.

1.3 Further Applications

Thermal neutrons are of interest to CEνNS experiments as a source of background events. MINER, for instance, measures their near-reactor thermal neutron flux as $5.8\times10^{7}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$ while expecting 5—20 events/kg/day and giving an acceptable background neutron event rate of, at most, 100 events/kg/day [23]. The sea-level terrestrial thermal neutron flux is much lower, at $4\ \mathrm{cm}^{-2}\mathrm{hr}^{-1}$, or $0.0011\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$ [24]; however, flux levels for these experiments would need to be brought as low as $10^{-4}\ \mathrm{cm}^{-2}\mathrm{s}^{-1}$ [25]. Accepting a background much greater than the expected signal creates a clear need for a thorough understanding of said background; without it, events of interest become lost in the sea of background events, and distinguishing them becomes impossible. Therefore, there is a need for background mitigation, such as through vetoing and tagging neutron capture-induced NR data, which requires an understanding of the background spectrum. NR backgrounds are also of interest to dark matter searches, as they have the potential to drown out a possible dark matter signal; while these backgrounds are typically – or, at least, for SuperCDMS experiments – not predominately from neutron capture, neutron capture does still contribute to them [26].

Neutron capture-induced recoils are of interest to a number of experiments. While there have long been indirect measurements of nuclear recoils (NRs) resulting from neutron capture in solid-state detectors [16], we have also recently observed it directly [27], allowing for actual implementation of a capture-based calibration method as laid out in this study. The CRAB collaboration has called for neutron capture measurements in germanium and scheelite, providing simulated data for future calibration [28]. It is also possible to distinguish nuclear recoils from electron recoils in some cases, and this discrimination of different pulse types has been improved in recent years through development of the detector crystals themselves [29]. There has been recent development of new discrimination methods for solids as well, albeit not for silicon [30]. Improved discrimination capabilities mean that the ability to take measurements, including those for calibration, at lower energies improves significantly.

2 Solid-State Detectors

Solid-state detectors convert light, electrical, or heat energy into different types of signals; in this study, I focus on the phonon and ionization signals. Energy deposited by an interaction can produce phonons and ionization, meaning that the detector signal will differ from the exact energy deposited by the initial interaction [31]. The conversion from ionization to phonon signal is proportional to the bias voltage on the detector [32]. With this conversion in mind, efforts are being made to quantify phonons in order to better be able to narrow down possible dark matter detection events [32].

It is important to understand expected ranges of energy deposition within the detector in order to ensure that the calibration method can give precise results – that is, that the resulting can distinguish nuclear recoils of different energies from one another – and to understand detector resolutions and noise so that realistic resolutions can be applied to simulated and analytic models. Detectors have been developed capable of reading very low energies with minimal random noise, so it is vital that we start to develop reliable calibration methods for them. Silicon-based phonon detectors – detectors that read vibrations in a silicon crystal, with the vibrations represented as quasi-particles – can reach detector signal thresholds as low as 70 eV [33]. These detectors can have resolutions down to 3 eV [34] that improve at lower energies [35]. Work is being done to improve the resolution of silicon detectors to less than 1 eV [6]. Other materials can achieve even lower thresholds; for example, $CaWO_4$ (scheelite) can have thresholds as low as 30.1 eV, although the resolutions are a bit larger than for silicon [36]. Lower-energy thresholds allow us to detect smaller-energy recoils, which often means lower-mass particles; smaller-energy resolutions make it easier to distinguish types of interactions.

2.1 Detector Yield Models

Solid-state detector physics still leans heavily on Lindhard theory to describe the relationship between energy deposited into a detector and the ionization signal, and alternative yield models are often modifications of the Lindhard model [37]. However, more recent measurements seem to show that this model is incomplete, at least for silicon [15, 38]. In particular, there are major deviations in the region below 4 keV, indicating that the model is not valid for low energies and requires some modification or a full replacement [39]. Other sources have confirmed that the Lindhard model cannot be extrapolated to low energies; however, a more accurate model has yet to be agreed upon [40]. Experiments that do agree with Lindhard do not go far below the 4 keV mark [41], and there are some experiments that show a deviation from the Lindhard model for germanium as well [12], while those at higher temperatures (77 K) do not [42].

II. SIMULATION

In order to calibrate any data, we must have a theoretical data distribution to compare it to. Therefore, the first step of this research is to generate the predicted distribution. What we need specifically is the distribution of energies read out by the detector following neutron capture. Because our detectors are sensitive to a produced ionization, not the total energy, this means we need a distribution that has been modified according to a yield model; yields are discussed in detail in Subsection 1.3. We also need to adjust for background data, either by including it in the model or subtracting it from the data; backgrounds are discussed in Section 3. Using nrCascadeSim [19], I put the above theory into practice. While I simulated data for silicon to match the detectors we have available for calibration in this study, the simulation, firmware (Section III.2), and modeling (Section 1) for silicon detector responses can be adapted to other monatomic detector lattices, such as germanium, and I therefore leave my work generic wherever practical.

1 Theoretical Basis for Simulated Distribution

Much like electrons in an atom, neutrons in a nucleus can be at one of a finite set of energy levels, but typically will de-excite to a minimum, or ground, state over time. A neutron near a nucleus, prior to capture, can be treated as having an excited state of the neutron separation energy (hereafter, simply "separation energy") S_n. When it captures, it may move to the ground state immediately, or it may capture into a higher, intermediate state. As long as it is not in the ground state, the neutron can continue to de-excite to any lower state; we call such a series of de-excitations a "cascade." Some example cascades are presented in Figure II.1 and Table II.1. Following a capture or de-excitation, the recoiling nucleus interacts with the surrounding lattice, slowing down and depositing its energy into the detector. Yield models, presented in Subsection 1.3, describe the energy absorbed by the lattice during this process, which may be interrupted by further de-excitations. Therefore, because de-excitations can happen before all of the nucleus' kinetic energy has been deposited, each cascade has its own continuous distribution of possible energy deposits, despite the fact that there is a finite set of possible cascades; the momentum gained at the next de-excitation is known for a given cascade, but the momentum before the de-excitation, and therefore the total momentum afterward, is not. That is, a given cascade does not have a single energy that it outputs, but a probability distribution over a range of possible energies. A full distribution of all possible energy deposits is, as a result, the sum of each possible cascade's distribution, weighted by the cascades' probabilities.

The first half of understanding the full distribution of cascade energies is understanding the relative probabilities of each cascade. The thermal cross-sections of isotopes are necessary to determine the probability of a thermal neutron capture occurring and are provided in Mughabghab [44]. Coupled with the cross-sections, we also need to know the relative abundance of each isotope compared to the total abundance of that element; this is provided by Rosman et al. [45]. It is the combination of the thermal cross-sections and the isotope abundances that

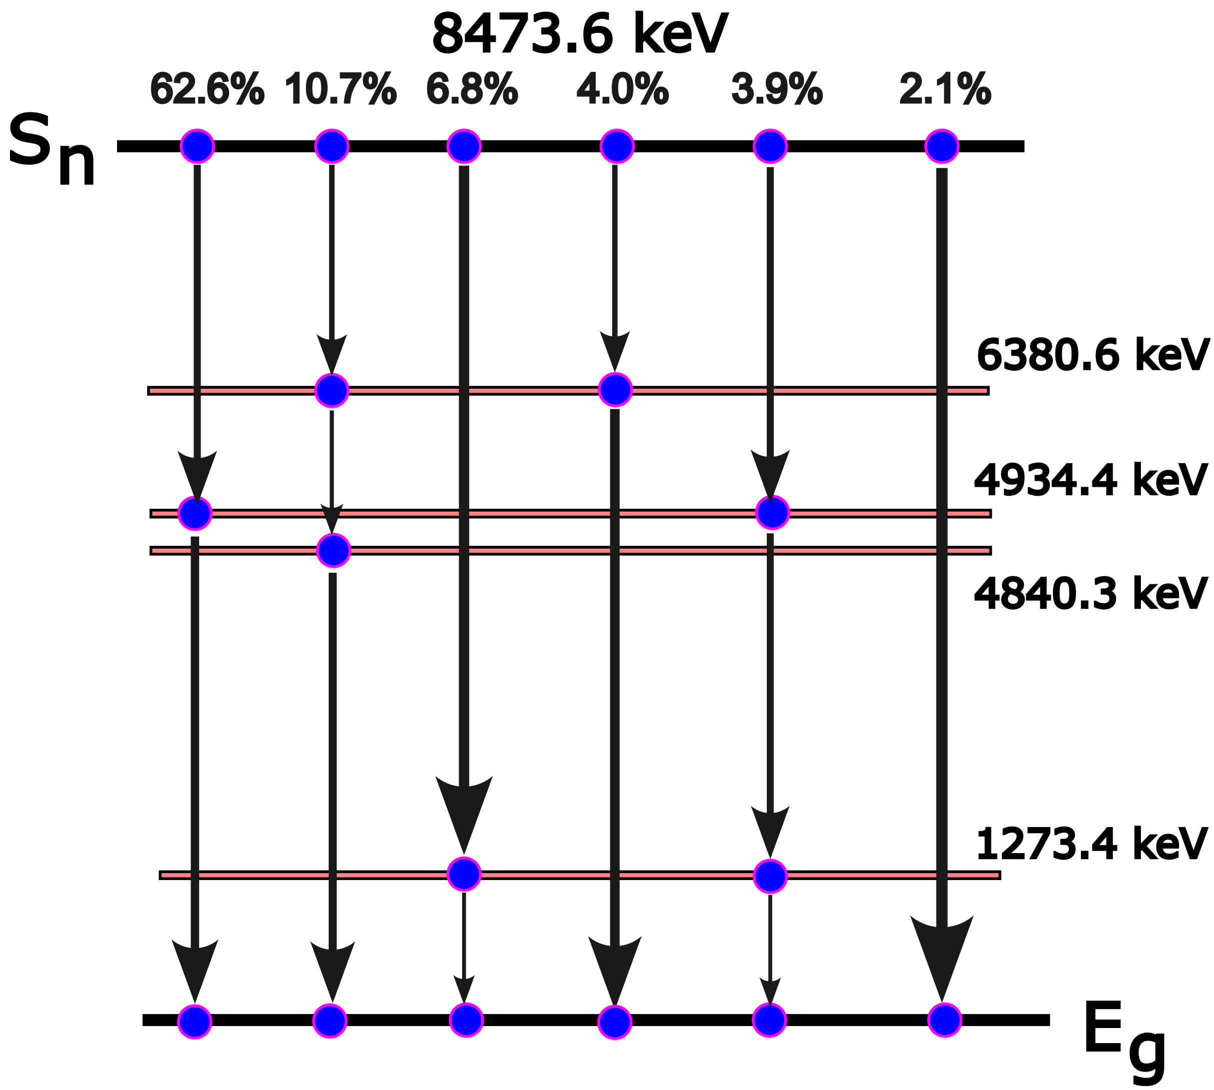

Figure II.1: ^{29}Si Cascades

A diagram demonstrating the six most prevalent cascades for silicon, adapted from Harris et. al. [27] All cascades begin with the capture of thermal neutrons to the separation energy S_n of ^{28}Si and end with the ground state of ^{29}Si.

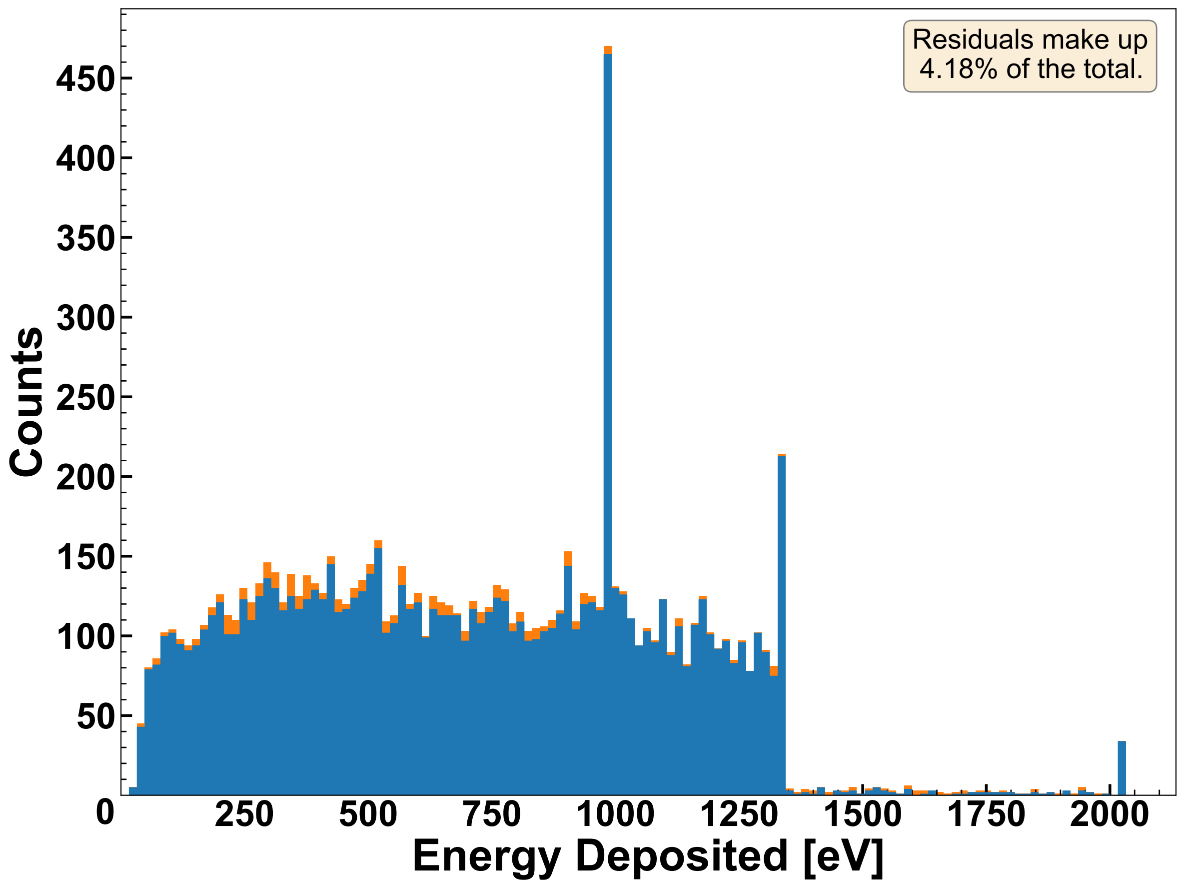

Figure II.2: Residuals

A stacked histogram of energy deposits generated by nrCascadeSim [19] demonstrating the relative influence of the most and least common silicon cascades, adapted from Harris et al. [27]. The lower, blue portion of the histogram represents the 14 cascades listed in Table II.1; the upper, orange portion represents all remaining cascades, which are available in Appendix B.

Table II.1: Si Cascade Probabilities

Isotope	Probability (Percent)	Energy levels (keV)	Half-lives (fs)
^{29}Si	62.6%	4934.39	0.84
^{29}Si	10.7%	6380.58; 4840.34	0.36; 3.5
^{29}Si	6.78%	1273.37	291.0
^{29}Si	3.97%	6380.58	0.36
^{29}Si	3.89%	4934.39; 1273.37	0.84; 291.0
^{29}Si	2.08%	...	...
^{30}Si	1.51%	6744.10	14
^{31}Si	1.41%	3532.9; 752.20	6.9; 530
^{30}Si	1.20%	7507.8; 2235.30	24; 215
^{30}Si	0.401%	8163.20	w(E1)[0.0019]
^{31}Si	0.351%	5281.4; 752.20	w(E1)[0.0069]; 530
^{30}Si	0.341%	...	...
^{31}Si	0.263%	4382.4; 752.20	w(E1)[0.012]; 530
^{31}Si	0.0303%	...	...

The probability of each cascade, adapted from Harris et. al. [27]. This table includes only the most common cascades for silicon, except to ensure the inclusion of straight-to-ground transitions due to their effect on the shape of the overall curve. These fourteen cascades are demonstrated in Fig II.2 to make up over 95% of events. The isotope listed is the isotope after neutron capture, in order to match the energy levels and half-lives listed. A half-life entry in [brackets] preceded by w(E1) specifies that the half-life is unknown and the Weisskopf estimate for an electric dipole transition was used [43]. For a full list of cascades used, including the residuals, see Appendix B.

determines what portion of NR events belong to a given isotope. To model the outcome of a given interaction, we will need to know a variety of information about the nucleus, including excitation energies and lifetimes, which are available from NuDat [46]. Where lifetimes are not available, we use a Weisskopf estimate [43], although these are known by experiment to underestimate the lifetimes [46]. Raman et al. [47] have also verified some of the above information by experimentally finding the relative probability of each gamma emission for neutron capture-induced NRs in silicon.

To combine all of these pieces, we need to understand how to model a cascade where all physical variables have been predetermined. First, a neutron captures into a nucleus at a particular excited state E_1. The post-capture nucleus recoils directly away from the associated gamma, as in Figure I.2, then slows down until it either stops or a de-excitation occurs (assuming E_1 is not the ground state), which is treated as one deposit. To determine how much energy is deposited during a given step of a cascade, we use a constant stopping power α, which represents the time it takes to stop the nucleus using the derivative of the nucleus' kinetic energy T with respect to its total distance traveled within the lattice [48, 49]; that is,

$$\int_0^{x_1} \alpha dx = \int_0^{x_1} \frac{dT}{dx} dx = T(x). \tag{1.1}$$

However, we need to be able to compare the energy lost to the time that has passed so that we can generate a distribution of energies based on the lifetime of the current excited state. Applying the chain rule,

$$T(t) = \int_0^{t_1} \frac{dT}{dx}\frac{dx}{dt}dt, \therefore \tag{1.2}$$

$$\frac{dT}{dt} = \alpha \frac{dx}{dt}. \tag{1.3}$$

Because $T = \frac{m}{2}v^2$ and $\frac{dx}{dt}$ is the velocity v,

$$\frac{dT}{dt} = \alpha\sqrt{\frac{2}{m}T} \tag{1.4}$$

Treating the energy as continuous over the time of interest – discontinuities only occur when a de-excitation occurs, when the movement stops, and before any time has passed – and knowing that it is differentiable everywhere that it is continuous, we can invert both sides of this relationship and maintain differentiability, following the inverse function theorem:

$$\frac{dt}{dT} = \frac{\sqrt{m/2}}{\alpha}T^{-1/2}, \tag{1.5}$$

which, by the inverse and implicit function theorems, means

$$\int dt = \frac{\sqrt{m/2}}{\alpha}\int \frac{dE}{\sqrt{E}}. \tag{1.6}$$

We give the integrals limits corresponding to the start and end times:

$$\int_0^{t_1} dt = \frac{\sqrt{m/2}}{\alpha}\int_0^{t_1} \frac{dE}{\sqrt{E}}, \therefore \tag{1.7}$$

$$t_1 = \frac{\sqrt{2m\left(T_0 - T(t_1)\right)}}{\alpha}. \tag{1.8}$$

Setting $\alpha' = \frac{\sqrt{2m}}{\alpha}$,

$$t_1^2 = \alpha'^2\left(T_1 - T(t_1)\right), \therefore \tag{1.9}$$

$$T(t_1) = T_1 - \left(\frac{t_1}{\alpha'}\right)^2. \tag{1.10}$$

Therefore, the energy deposited by the first step, before the yield model is applied, is

$$D_1 = \left(\frac{t_1}{\alpha'}\right)^2. \tag{1.11}$$

The time spent slowing down and depositing energy is $0 < t < \tau$, where τ is the amount of time it takes for the nucleus to stop (or, equivalently, for the kinetic energy to be zero); that is, the stopping time is

$$\tau = \alpha' \sqrt{T_1}, \tag{1.12}$$

where $T_1 = \frac{p_1^2}{2m}$, m is the mass of the resulting nucleus, and the momentum after the first de-excitation is

$$p_1 = 4m \left(-1 + \sqrt{1 + \frac{2}{m}(S_n - E_1)} \right), \tag{1.13}$$

as derived in Equation A.8.

The distribution of a single-step, or straight-to-ground, cascade is a delta function; all of the energy from capture will be deposited into the detector with a known deposited value. We calculate the readout energy distribution of a multi-step cascade, in a general sense, as follows. The neutron captures into the nucleus, causing the nucleus to have a back-to-back recoil with the emitted gamma, giving it an initial kinetic energy T_1 as described by equation A.12. The first step deposits an energy $D_1(t)$ (1.11) which has an ionization energy of $Y(D_1(t))$. If $D_1 = T_1$, the nucleus stops before de-exciting, and the next step proceeds in the same manner, although T_1, the kinetic energy after the first step, will depend on the energy E_1 of the new excited state, rather than on the separation energy S_n. For all steps, $t \geq 0$ is distributed exponentially:

$$f(t) = \frac{e^{-t/\lambda}}{\lambda} \tag{1.14}$$

where λ is the state's mean life

$$\lambda = \frac{\tau}{\ln(2)} \tag{1.15}$$

for a half-life τ. If the nucleus does *not* stop before de-excitation occurs, however, the calculation of the new recoil's kinetic energy is more complex because of "decay in-flight"; the initial momentum of the nucleus can no longer be treated as negligible, which means that there are a range of possible values of T_2. The only requirements are that energy and momentum are conserved, but we now have to solve for the momentum of the gamma and the nucleus separately. To avoid these extra variables, we can apply a frame transformation. We begin in the lab frame, where the nucleus appears to be moving, before de-excitation occurs. Just before the neutron de-excites, we apply a frame transformation by subtracting the nucleus' momentum from all objects, making the nucleus appear to be stationary; this is the center of mass frame. These two frames are demonstrated in Figure II.3. This introduces only one additional variable in comparison to the first step – an angle β between the pre-de-excitation momentum (which we can declare to be the x axis for this calculation) and the nucleus' new momentum. To represent an even distribution in all directions, we take the distribution of β to be uniform over the surface of a unit sphere:

$$f(\beta) = \frac{\sin\beta}{4\pi} \tag{1.16}$$

The frame transformation is reversed once the post-de-excitation momenta of the gamma and nucleus have been determined by adding the initial momentum back in to all objects' momenta to determine T_2, D_2 is deposited over some time t, and the process repeats until the neutron reaches the ground state at the end of the cascade.

1.1 Two-Step Cascade Distribution

A two-step cascade has a distribution determined as follows. We start from a joint distribution representing the two deposits separately,

$$g(D_1, D_2) = f_{D_1}(D_1) f_{D_2|D_1}(D_1, D_2). \tag{1.17}$$

We know that there are only two probabilistic variables, t_1 and β_2; the first step has no angle component because it lacks an initial momentum, meaning we can define the coordinate system around the initial step's recoil direction, while the second step has no time component because there will be no subsequent steps, meaning it will always deposit all remaining energy. That said, while the first deposit does truly only depend on t_1, the second deposit will depend on both t_1 and β_2 because the energy going into the second step is relevant and depends on t_1. Because both deposits depend on t_1, they are not independent of one another. We begin by finding the distribution of D_1.

The first deposit has two cases: the case where all energy is deposited, and the case of decay in-flight. This gives us a piecewise function

$$f_{D_1}(D_1) = \begin{cases} f^0_{D_1}(D_1) & 0 \leq D_1 < T_1 \\ P(t_1 \geq \tau) & D_1 = T_1 \end{cases} \tag{1.18}$$

which can also be expressed using the Dirac delta function δ:

$$f_{D_1}(D_1) = f^0_{D_1}(D_1) + P(t_1 > \tau)\delta(T_1 - D_1) \tag{1.19}$$

over the range $0 \leq D_1 \leq T_1$. We find $f^0_{D_1}$ by using the equivalence of the probability of a given time (Equation 1.14) and its associated first energy deposit:

$$\int_{t_0}^{t_1} dt \frac{e^{-t/\lambda}}{\lambda} = \int_{t_0}^{t_1} dD_1 f_{D_1}(D_1)^0. \tag{1.20}$$

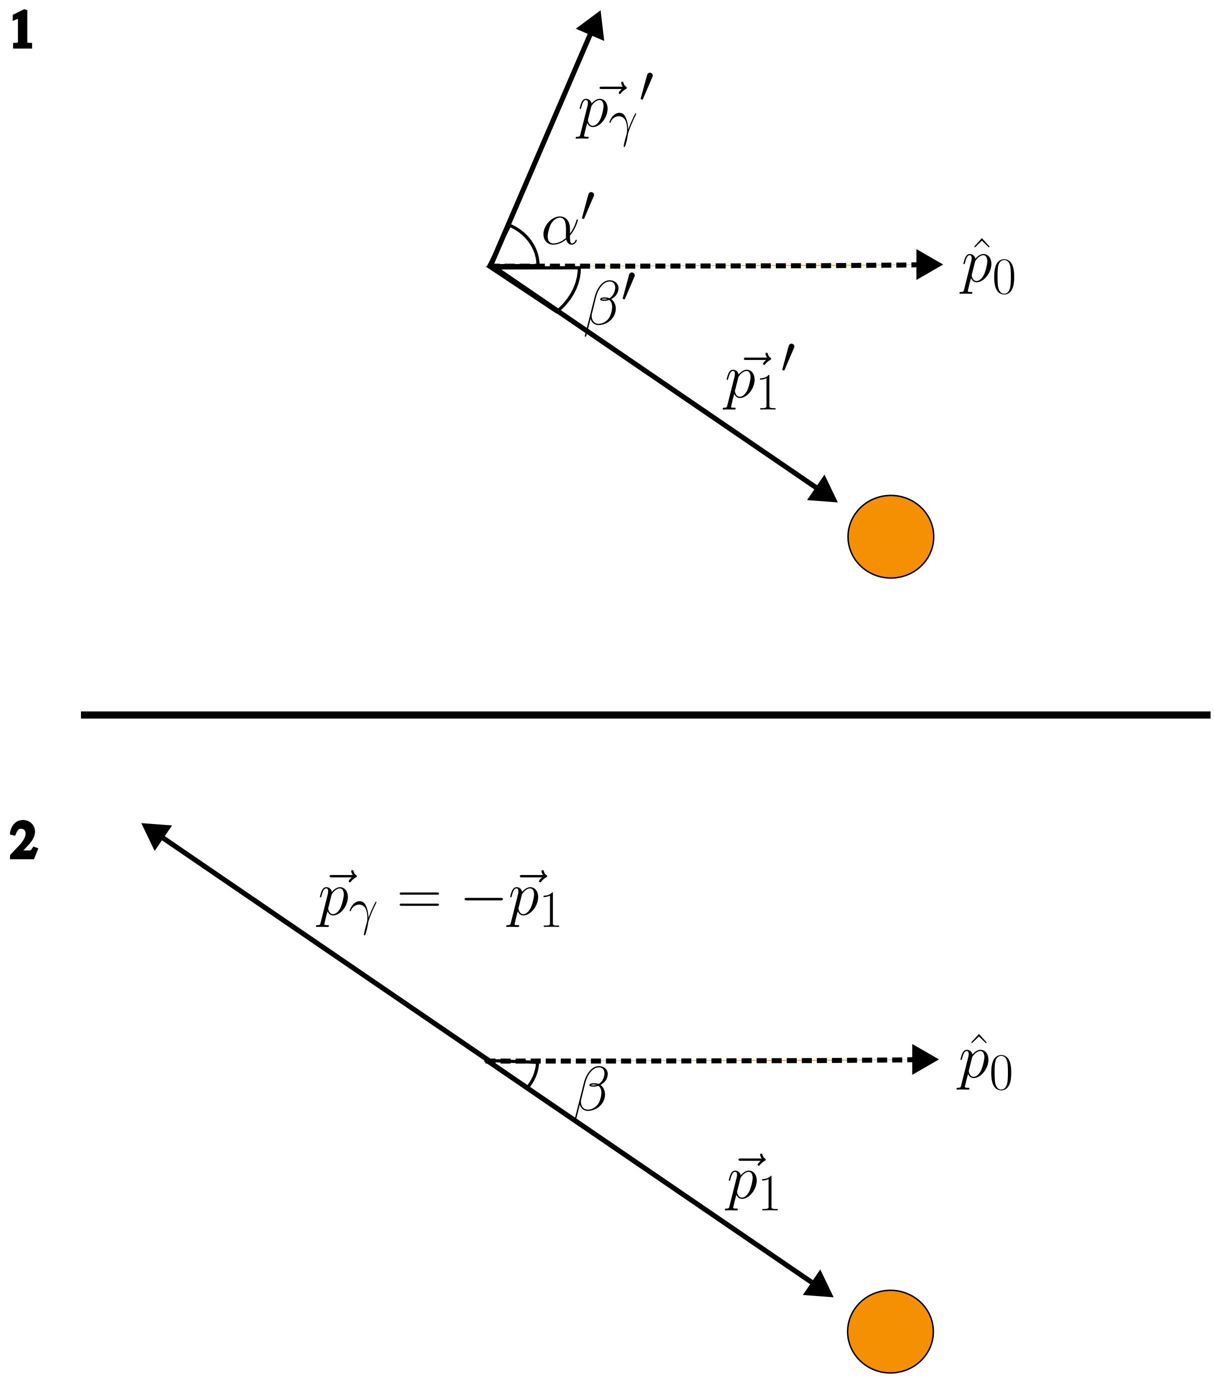

Figure II.3: Recoils with Initial Momentum

Momentum of the nucleus following a de-excitation occurring during a prior recoil. 1: (Lab Frame) Because there is a non-negligible initial momentum associated with the movement of the nucleus, the nucleus' new momentum is at some angle β' with respect to the initial momentum, which we use as the x axis. The resulting gamma must recoil at some other angle α' that satisfies conservation of momentum. 2: (Center of Mass Frame) By applying a momentum transformation to the system such that the nucleus becomes stationary just before the de-excitation, we create another back-to-back interaction, albeit one at an angle β to the transformation vector.

We can rewrite the left side of this equation as

$$\int_a^b dD_1 \left(\frac{dD_1}{dt}\right)^{-1} \frac{e^{-t/\lambda}}{\lambda}, \tag{1.21}$$

which, using equation 1.11, makes

$$\left(\frac{2t}{\alpha'^2}\right)^{-1} \frac{e^{-t/\lambda}}{\lambda} = f^0_{D_1}(D_1). \tag{1.22}$$

In order to convert this into a function of D_1 and remove direct references to t, we invert $D_1(t)$

$$t(D_1) = \alpha' \sqrt{D_1} \tag{1.23}$$

then substitute $t(D_1)$ into equation 1.22 to get

$$f^0_{D_1}(D_1) = \left(\frac{2\sqrt{D_1}}{\alpha'}\right)^{-1} \frac{e^{-2\alpha'\sqrt{D_1}/\lambda}}{\lambda}. \tag{1.24}$$

The probability $P(t_1 > \tau)$ can be easily calculated:

$$P(t_1 > \tau) = \int_\tau^\infty dt f_t(t) \tag{1.25}$$

$$= -e^{-t/\lambda}|_\tau^\infty \tag{1.26}$$

$$= e^{-\tau/\lambda}. \tag{1.27}$$

To find a distribution of D_2, we need to either determine a joint distribution $g(D_1, D_2)$ or convert the conditional distribution $f_{D_2|D_1}(D_2)$. Because there are no further decays possible, all of the new kinetic energy will be deposited, so there must be some minimum possible deposit D_2^{min} and maximum possible deposit D_2^{max} that correspond to specific recoil angles. These two values result from the fact that decay in-flight can cause an increase in kinetic energy by allowing the de-excitation's added momentum to function as a "boost" for an existing momentum, but, due to the randomness of the direction of the added momentum, can subtract from the existing momentum as well. Furthermore, the magnitude of the momentum is constrained but is not predetermined, making it difficult to intuit what these minimum and maximum deposits must be. We will therefore leave them as variables while determining the distribution, solving for them once $f_{D_2}(D_2)$ has been found.

Because D_1 has no dependence on D_2, but D_2 does depend on D_1, we will determine $f_{D_2|D_1}(D_2)$. If $D_1 = T_1$, then D_2 is certain; the new recoil energy is $T_2 = \frac{mp_2^2}{2}$, where

$$p_2 = m\left(-1 + \sqrt{1 + \frac{2}{m}E_1}\right), \tag{1.28}$$

and must be fully deposited. In all other cases, the angle of recoil determines the resulting kinetic and deposited energy. Normally, we would need to track the remaining momentum just before the de-excitation p_1^*; we temporarily subtract this momentum, applying a frame transformation such that the initial momentum is zero, as shown in Figure II.3. This transformation allows us to calculate the new nuclear momentum p_2^* – in the transformed frame, at least – as if there were no initial momentum, just as for p_2 above. Finding the lab-frame momentum is simple matter of reversing the frame transformation:

$$\vec{p}_2 = \vec{p}_2^* + \vec{p}_1^* =< p_2^* \cos\beta, p_2^* \sin\beta > + < p_1^*, 0 >, \tag{1.29}$$

measuring β from the axis of p_1^*. This makes the magnitude

$$\begin{aligned} p_2 &= \sqrt{(p_2^*)^2 \cos^2\beta + (p_2^*)^2 \sin^2\beta + 2p_1 p_2^* \cos\beta + p_1^2} && (1.30) \\ &= \sqrt{(p_2^*)2 + p_1^2 + 2p_1 p_2^* \cos\beta} && (1.31) \end{aligned}$$

and the deposit

$$\begin{aligned} D_2 &= \frac{1}{2m}((p_2^*)^2 + (p_1^*)^2 + 2p_1^* p_2^* \cos\beta) \\ &= \frac{1}{2m}\left(m^2\left(-1 - 2\sqrt{1 + \frac{2}{m}E_1} + 1 + \frac{2}{m}E_1\right)\right. \\ &\quad \left. +2m(T_1 - D_1) + 2\sqrt{2m(T_1 - D_1)}m\left(-1 + \sqrt{1 + \frac{2}{m}E_1}\right)\cos\beta\right) \\ &= \left(E_1 - \sqrt{m^2 + 2mE_1}\right) + (T_1 - D_1) + \sqrt{2m(T_1 - D_1)}\left(\sqrt{1 + 2\frac{E_1}{m}} - 1\right)\cos\beta. \end{aligned} \tag{1.32}$$

Now that we know the dependence of D_2 on other variables, we can find its conditional distribution, starting from the conditional probability:

$$\begin{aligned} P(D_2 \in \mathcal{D}_2 | D_1) &= \int_{\mathcal{D}_2} f_{D_2|D_1}(D_2) dD_2 && (1.33) \\ &= \int_{-1}^{1} f_{cos\beta}(cos\beta) d(cos\beta). && (1.34) \end{aligned}$$

Because β is uniform on the surface of a sphere, $cos\beta$ is uniformly distributed across its allowed values $-1 \leq cos\beta \leq 1$ (see Appendix C). Therefore,

$$\int_{\mathcal{D}_2} dD_2 f_{D_2|D_1}(D_2) = \frac{1}{2}\int_{-1}^{1} d(cos\beta). \tag{1.35}$$

Applying the chain rule,

$$d(cos\beta) = \frac{d(cos\beta)}{dD_2} dD_2, \tag{1.36}$$

which, using the inverse and implicit function theorems, can be rewritten as

$$d(cos\beta) = \left(\frac{dD_2}{d(cos\beta)}\right)^{-1} dD_2. \tag{1.37}$$

The derivative of D_2 with respect to $cos\beta$ is

$$\frac{dD_2}{d(cos\beta)} = \sqrt{2m(T_1 - D_1)}\left(\sqrt{1 + 2\frac{E_1}{m}} - 1\right), \tag{1.38}$$

making

$$f_{D_2|D_1}(D_2) = \begin{cases} 0 & D_2 < D_2^{min} \\ \frac{1}{2}\left(\sqrt{2m(T_1 - D_1)}\left(\sqrt{1 + 2\frac{E_1}{m}} - 1\right)\right)^{-1} & D_2^{min} \leq D_2 \leq D_2^{max} \\ 0 & D_2 > D_2^{max} \end{cases}. \tag{1.39}$$

The limits D_2^{min}, D_2^{max} are

$$D_2^{min} = max\left(0, \left(E_1 - \sqrt{m^2 + 2mE_1}\right) + (T_1 - D_1) - \sqrt{2m(T_1 - D_1)}\left(\sqrt{1 + 2\frac{E_1}{m}} - 1\right)\right) \tag{1.40}$$

$$D_2^{max} = \left(E_1 - \sqrt{m^2 + 2mE_1}\right) + (T_1 - D_1) + \sqrt{2m(T_1 - D_1)}\left(\sqrt{1 + 2\frac{E_1}{m}} - 1\right). \tag{1.41}$$

While it is possible to use $f_{D_2|D_1}(D_2)$ and $f_{D_1}(D_1)$ to find the distribution of the total deposit $f_{D_{tot}}(D_{tot})$, we will find in Section 1.3 that, because of the nature of ionization yields, it is not beneficial to do so at this stage. However, for the sake of demonstration, we will proceed with the remaining calculations with some simplifying approximations. To make integration easier, we make a numerical approximation by taking only the first two terms of the Taylor expansion of $\sqrt{1+x}$ about 0:

$$\begin{aligned} \left(\sqrt{1 + 2\frac{E_1}{m}} - 1\right) &= -1 + \sum_{n=0}^{\infty} \frac{1}{n!}\left(2\frac{E_1}{m}\right)^n \left(\frac{\partial}{\partial x}\right)^n \sqrt{1+x}\bigg|_{x=0} \\ &= -1 + 1 + \frac{1}{2}\left(2\frac{E_1}{m}\right) + O\left(\frac{E_1^2}{m^2}\right). \end{aligned} \tag{1.42}$$

This approximation makes the conditional distribution simply

$$f_{D_2|D_1}(D_2) = \frac{1}{2}\left(\frac{E_1}{m}\sqrt{m(T_1 - D_1)}\right)^{-1}. \tag{1.43}$$

The remainder of the Taylor expansion is then

$$R_2 = \frac{1}{2}\left(2\frac{E_1}{m}\right)^3 \int \tag{1.44}$$

We can then integrate a joint distribution based on this conditional one:

$$f_{D_tot}(D_{tot}) = \int_0^{S_n} f_{D_1}(D_1) f_{D_{tot}-D_1|D_1}(D_1) dD_1 \tag{1.45}$$

An example of $f_{D_{tot}}(D_{tot})$ for silicon, using the above approximation and integrated numerically, is shown in Figure II.4, and an example of the joint distribution $g_{D_1,D_2}(D_1, D_2)$ is shown in Figure II.5.

1.2 Generalizing the Number of Steps

What we are missing, if we wish to induct to find a general case, is the middle step of a three-step cascade, where there are factors from other steps influencing both the starting and ending kinetic energy, as occurs for over 14% of cascades. There is unfortunately, however, a more glaring issue with generating an analytical form for the total deposit. The work above assumes that the relationships between the deposited energy and the time and angle variables are invertible, which is true until we begin to apply yield models. Not all yield models are invertible, preventing a full analytical solution for the total deposited energy of a cascade, regardless of the number of steps. We therefore use the physics outlined in the equations above to simulate the yielded energy totals and the theory presented above to estimate the impact of simulation (see Section 2).

1.3 Yield Models

The readout of a detector does not necessarily have a linear relationship with the energy deposited. Instead, there is an ionization yield – non-electric interactions within the detector cause ionization, which results in a smaller, but proportional, amount of electric energy that the detector reports. Because the time spent at each step is very small, energy transferred to the detector during the slowing-down process, between de-excitations, is a single deposit, rather than a series of deposits, but we treat the energy lost at each step in the cascade as continuous. This greatly complicates the formulation of a general analytic distribution for the total energy of a cascade for a variety of reasons. Changing the ionization yield, in order to test the applicability of different yields for a given detector, means recalculating the distributions of each step. Because we use the inversion theorem in

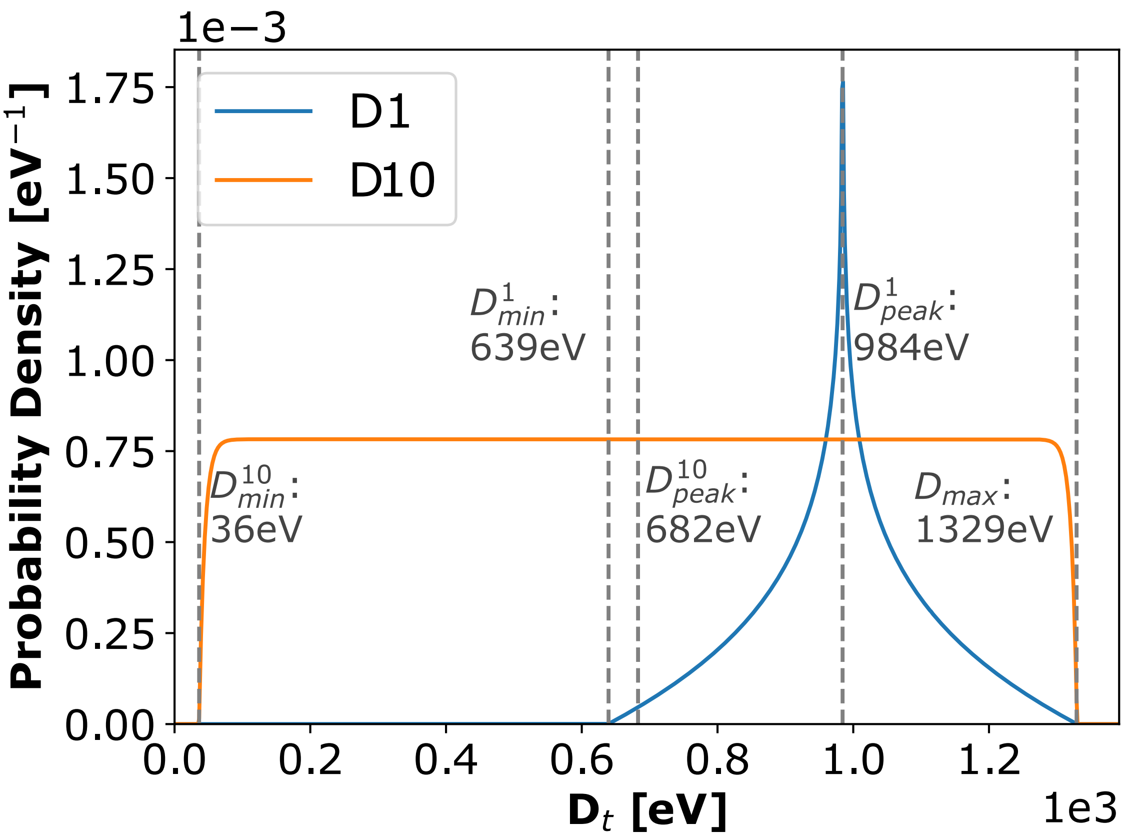

Figure II.4: One-Dimensional PDF of Two-Step Cascade

The PDF for the total energy deposit from a two-step cascade, adapted from Harris et al. [27] D^1 represents deposits where the intermediate step was the first excited state. D^{10} represents deposits where the intermediate step was the tenth excited state, which was chosen because it has a long lifetime, making the probability of fully stopping before de-excitation high.

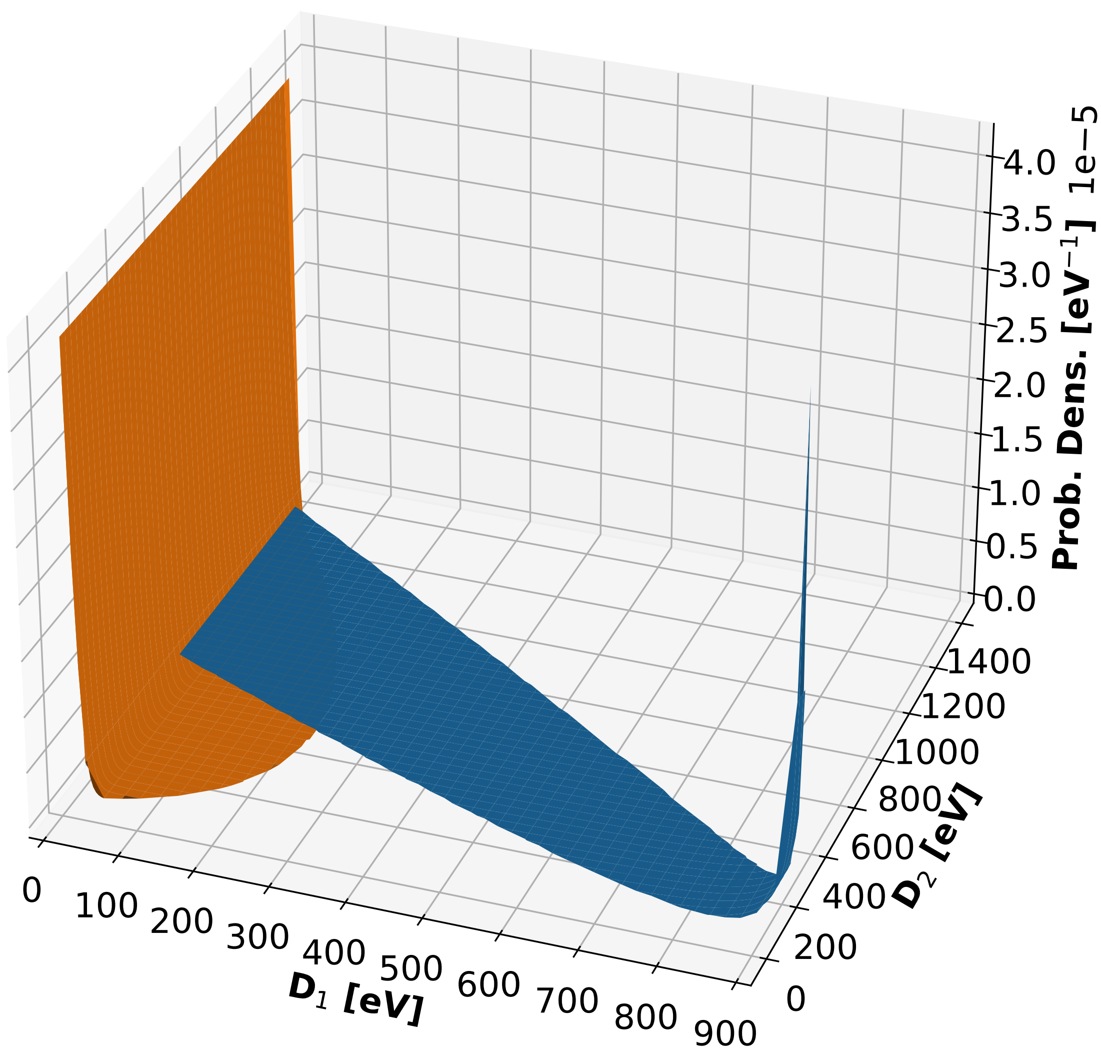

Figure II.5: Two-Dimensional PDF of Two-Step Cascade

The two-dimensional PDF for the two energy deposits from a two-step cascade, adapted from Harris et. al. [27] The blue distribution, which spikes in the high D_1 region, represents deposits where the intermediate step was the first excited state. The orange distribution, which rises near zero, represents deposits where the intermediate step was the tenth excited state, which was chosen because it has a long lifetime, making the probability of fully stopping before de-excitation high.

the above calculations of the analytical distributions, the distributions become more complex to calculate when the yields are applied, and some yield models may not be invertible at all.

As discussed in Section I.2.1, most yield models – of which there are many – are modifications on Lindhard, which describes the ionization yield $\bar{\nu}$ – effectively, the ratio of the yielded energy to the "true" energy deposit E – as

$$\bar{\nu} = \frac{\epsilon}{1 + kg(\epsilon)} \tag{1.46}$$

where

$$\epsilon = \frac{Ea}{2Z^2e^2} \tag{1.47}$$

is a dimensionless, reduced energy; $a = 0.8853a_0Z^{1/3}$ is a multiple of the Bohr radius; Z is the atomic number; e is the fundamental charge; and

$$g(\epsilon) = \alpha\epsilon^{\gamma} + \beta\epsilon^{\omega} + \epsilon \tag{1.48}$$

for unknown parameter k and material-dependent parameters $\alpha, \beta, \gamma, \omega$ [50]. The resulting yielded energy is then

$$Y(E) = E\bar{\nu}. \tag{1.49}$$

The fully expanded expression for $Y(E)$ is not invertible using tools such as Mathematica [51], and I found no way to invert the function when attempting to invert the function myself, but because I have already presented other reasons not to provide a full, general, and analytical solution for $g_{D_i}(D_i)$, I do not provide a formal proof that the function cannot be inverted, as the only consequence of such a proof for this work would be the simulation that has already been adequately justified.

We therefore save all simulation data, such as intermediate step energy values, in a way that makes it possible to adjust yield models and any other post-simulation analysis, as shown in Figure II.6. Keeping this granular data makes the yield model flexible, allowing us to adjust the calibration method as our understanding of the yield models in the low-energy regions improved throughout the study. It also allowed us to run a single simulation for every yield model that we tested calibration for, including applying new models retroactively. The exact models used are discussed in Chapter IV.

2 Stopping Time Approximation

A major assumption is made in this work – an approximation treating the recoiling nucleus as having a constant acceleration, which is calculated from the stopping time. It is therefore important to estimate whether or not this approximation has a major impact on the results of the simulation. I therefore use two-sample Kolmogorov-Smirnov (KS) testing to compare the distributions of multiple simulations with different constant stopping powers.

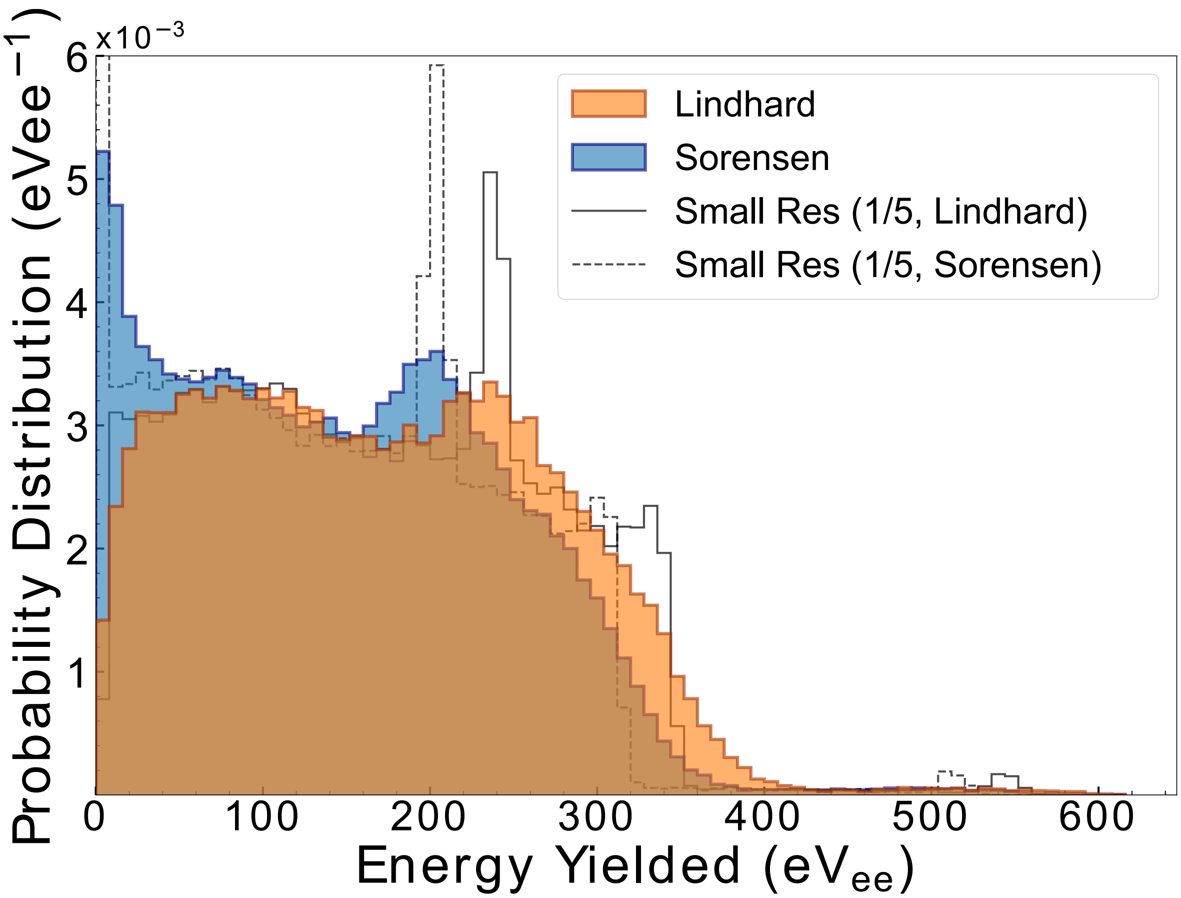

Figure II.6: Sorensen vs. Lindhard

Overlaid histograms demonstrate nrCascadeSim output, which shows a simulated probability distribution of ionization energies. The plot uses two ionization yield models – Lindhard [37] and Sorensen [50] – and two Gaussian noise levels – one equal to the value expected for our detectors, which is a function of energy [52], and one at one fifth of that scale. The underlying data for all four distributions is the same – one million data points generated using the cascades listed in Table II.1 and the seed 1.

I choose to adjust the stopping power because the simulation calculates the constant acceleration from it and uses it for no other calculations, so adjusting either of the two is equivalent. Because the stopping power varies, as shown in Lindhard [37], with energy, this comparison gives some idea of how much the distribution differs from the most accurate values for the relevant regions. A KS test is a comparison of two independent samples that aims to determine if they could feasibly come from the same parent population (the two-sample case), or to see if a given distribution could be the parent of the sample (the one-sample case). The statistic for both tests uses an empirical distribution function $F_n(x)$ – that is, a cumulative distribution function derived from the data – calculated by:

$$F_n(x) = \frac{1}{n}\sum_{i=1}^{n} 1_{(-\infty,x]}(X_i), \tag{2.1}$$

where n is the number of independent and identically distributed (i.i.d.) events, X_i represents the events individually, and $1_{(-\infty,x]}$ is the indicator function, equal to 1 if the argument is within the range $(-\infty, x]$ and 0 otherwise. Because the simulation does not consider previous events during the calculation of cascade energy (events here being entire cascades, not steps within a cascade), events are statistically independent of one another, and because all events have the same chance of coming from any of the cascades' distributions as all other events do, we can treat them as all coming from one identical, composite distribution that represents all cascades. The one-sample test calculates the Kolmogorov-Smirnov statistic D_n as a comparison of the eCDF and a chosen CDF $F(x)$:

$$D_n = \sup_x |F_n(x) - F(x)|, \tag{2.2}$$

where $\sup_x g(x)$ refers to the supremum of a set of all values of $g(x)$ over the domain x, which, in this case, consists of all data points X_i. The two-sample test instead compares the eCDFs $F_{1,n}, F_{2,n}$ of the two datasets:

$$D_{n,m} = \sup_x |F_{1,n}(x) - F_{2,n}(x)|. \tag{2.3}$$

A lower KS statistic is more ideal, and the $p-$value for rejection of the null hypothesis varies depending on the number of samples. In the one-sample test, the null hypothesis is rejected iff

$$\sqrt{n}D_n > K_\alpha \tag{2.4}$$

where

$$P(K < K_\alpha) = 1 - \alpha \tag{2.5}$$

and

$$P(K \leq x) = 1 - 2\sum_{k=1}^{\infty}(-1)^{k-1}e^{-2k^2x^2} \tag{2.6}$$

is the Kolmogorov distribution [53]. In the two-sample test, it is rejected iff

$$D_{n,m} > c(\alpha)\sqrt{\frac{n+m}{nm}} \tag{2.7}$$

for respective sample sizes n and m and where

$$c(\alpha) = \sqrt{-\frac{1}{2}\ln\frac{\alpha}{2}}. \tag{2.8}$$

I evaluated both the statistics and the associated $p-$values using the python package SciPy [54], which provides, by default, a two-tailed exact $p-$value calculated from the KS statistic.

I choose a two-sample KS test for the stopping times because it does not require a well-known parent distribution, which, as discussed in Section 1, we cannot construct. The constancy of the nuclear deceleration is an approximation that greatly simplifies both simulation and formal calculation of the nuclear momentum at different points during its movement. Simplifying the simulation prevents the calculations from becoming unmanageably long; generalizing the acceleration in any way caused the simulation to take upwards of ten minutes per singular event. With even just one thousand events, such a simulation would take a full week to run once, assuming that the machine it is being run on can handle the more intensive calculation at all. With a constant acceleration, however, simulations of a million entries run in minimal time and can be performed on a typical desktop; I was able to run such a simulation in under a minute with only 16 GB of RAM.

I used nrCascadeSim to generate two datasets with one million simulated cascades each, both using the seed 1337 and the full list of cascades provided in Appendix B. Using the same seed and cascade list guarantees the same output every time. The first used the program's default stopping power of $S_2 = 0.1$, while the second used an adjusted stopping power of $S_2 = 0.11$ to simulate faster deceleration. The default value of $S_2 = 0.1$ is taken from Lindhard to be approximately true for the region where the stopping power is flattest to maximize the region where it is approximately true. The only visible difference between the two is an increase in the peak near 1000 eV, which is still the highest peak in both datasets. The KS test does not reject the null hypothesis – that the two samples are drawn from the same distribution – assuming a p-value of 0.05, with a statistic of 0.00156 and a p-value of 0.197. Following the success of this test, I additionally tested the default against $S_2 = 0.2$ and $S_2 = 0.05$; both of these tests did reject the null hypothesis, with $D_n = 0.0110, p = 1.65 * 10^{-50}$ and $D_n = 0.0105, p = 1.34 * 10^{-46}$, respectively, confirming that drastic changes in the stopping power have a major effect on the overall distribution. Running the test only for deposits up to 1000 eV also made the difference

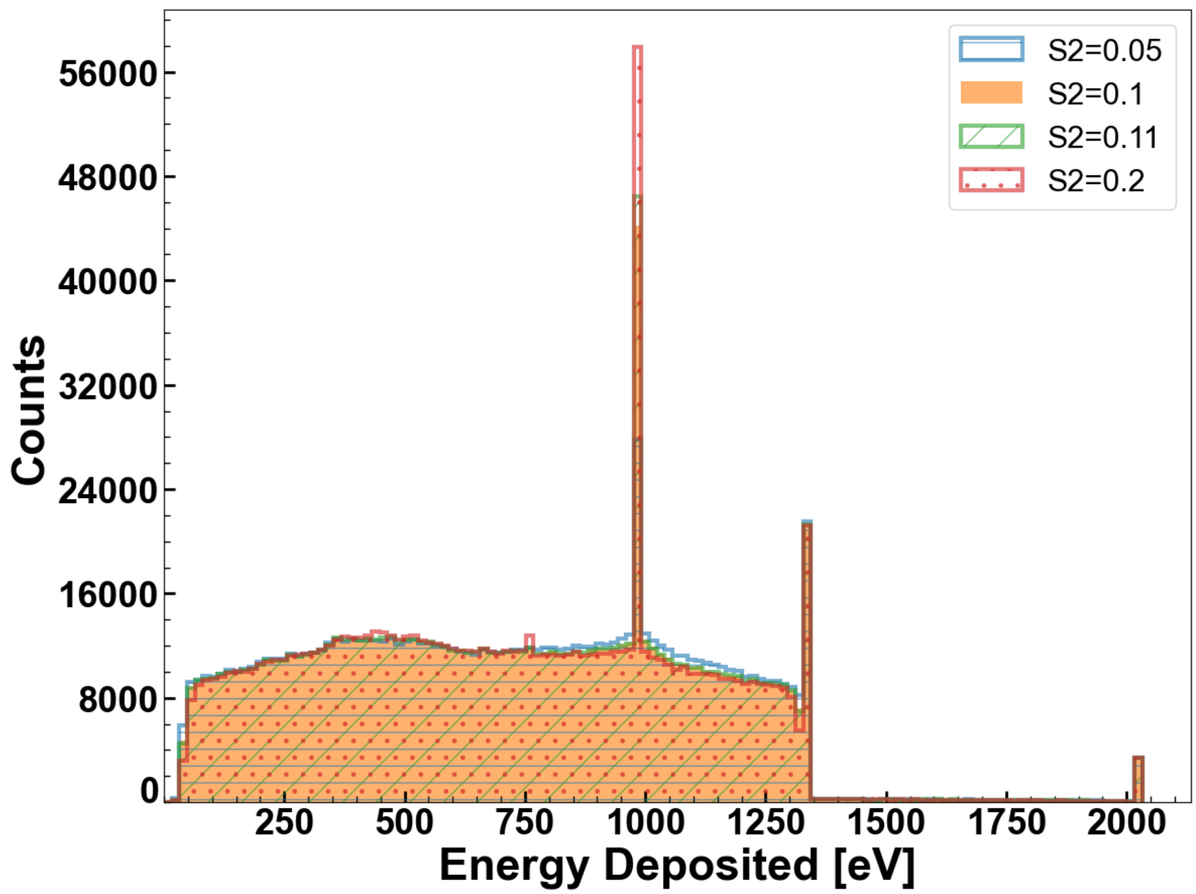

Figure II.7: Comparison of Stopping Powers in Simulation

Figure II.8: Data generated by nrCascadeSim using the same cascade definitions and randomization seed each time, but with different values for the stopping power.

between $S_2 = 0.1$ and $S_2 = 0.11$ statistically significant, with a statistic of 0.00310 and a p-value of 0.00199, indicating that the difference is significant for even small variations in the lower-energy region. The stopping power can change significantly as the energy changes [37], so adjusting the stopping power in the simulation to be calculated from the energy would be a major improvement. However, while the most visible peaks of the distribution have altered heights across stopping power values, they do remain intact and maintain their positions, as shown in Figure II.7.

3 Backgrounds for Calibration Data

While there are many possible sources of NRs, we consider NRs induced by sources other than neutron capture to be background noise and not used for calibration. Fortunately, many of these background sources will be of high enough energy that we can ignore them entirely for the energy region of interest. These sources are also common backgrounds for dark matter, and the neutron distributions discussed in this chapter and in Chapter IV can be DM backgrounds as well.

Neutron-Nucleus Scattering: The interactions we are most interested in are those where an incoming neutron is captured by a nucleus; however, wherever capture is possible, scattering is also possible. In scattering, the incident neutron remains separate from the nucleus when they collide. This results in the neutron's path being diverted and the nucleus potentially recoiling, as shown in Figure II.9. Conservation of energy and momentum still apply; therefore, the maximum energy transferred to the detector this way is proportional to the kinetic energy of the incoming neutron:

$$E_r = \frac{2m_n m_1}{(m_n + m_1)^2} E_0 \tag{3.1}$$

For thermal neutrons, the recoil caused by scattering is therefore below the detector's threshold and therefore won't be picked up. Because the states of the neutron and the nucleus do not change due to elastic scattering, there is no reason for a gamma or other emission to occur [55], and therefore no reason for any scattering to trigger the NaI array or record the event. Inelastic scattering can cause a change in state within the nucleus, although this is less likely as it requires the incoming neutron to have sufficient energy for the state change.

Gamma-Nucleus Scattering: There are a variety of reactions that can occur after incident gamma ray directly interacts with a nucleus in the Si detector. In the case of total absorption, there is some resulting excitation which can cause a decay that may be picked up by the NaI array. It is also possible for the gamma ray to instead scatter [56] and subsequently strike the NaI array, causing a trigger with an associated signal in the Si detector. The lead encasing the experiment (see Chapter III) should have a significant dampening effect on stray gammas.

Gamma-Electron Scattering: While we are most interested in nuclear recoils, the detector also reports electron recoils, and while the responses of each are different, there is no simple way for us to differentiate between the two for a single event. One cause of electron recoils is gamma-electron scattering. Its mechanics are similar to gamma-nucleus scattering; instead of striking the nucleus, the gamma strikes and excites or scatters off of an electron. In practice, however, the impact of the scattering of outer-shell electrons in a solid state detector do not contribute much phonon noise in the region that is of interest to dark matter experiments, so only the electrons more closely bound to the lattice of nuclei have significant impact [57]. The scattering of photons off of charged, particles, including electrons, is known as Compton scattering. Scattering is less likely at lower photon energies [57, 58]. Estimating the backgrounds due to Compton scattering would require first estimating the spectrum of gamma radiation that enters the detector as well as the portion of those that result in a coincident detection.

Coherent Neutron-Electron Scattering: Another potential source of electron recoils is coherent neutron-electron scattering, called "coherent" because no ionization occurs. While neutrons lack any net electric charge, they do contain an internal charge structure due to their quark composition, as well as a spin, allowing them to

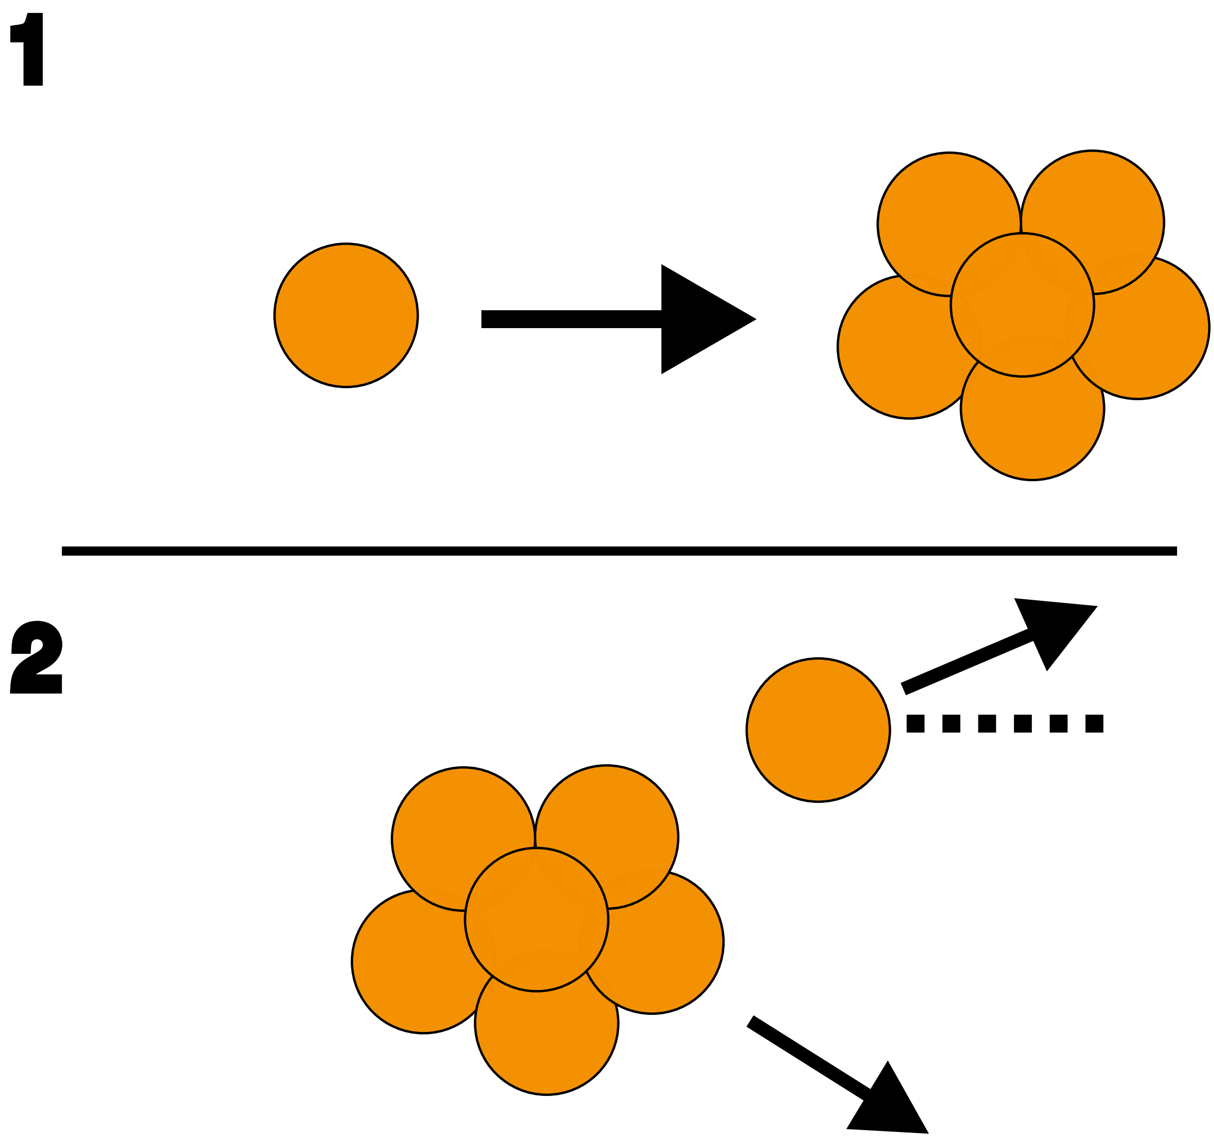

Figure II.9: Scattering

An incident neutron (1) scatters from a nucleus (2), transferring some of its momentum. The dotted line represents the original trajectory of the neutron.

have electromagnetic interactions in rare cases [59]. The neutron-electron scattering length for coherent electric scattering, in which the interaction is due to electric interactions with the charge structure of the atom and the neutron, has been found to be around $-1.32 * 10^{-3}$ fm [59, 60], which, approximating the cross section as a circle with the length as its radius, gives a cross section of $2.19 * 10^{-7} b$. Coherent magnetic scattering is much more likely, which is a result of magnetic structure of the neutron and the electrons, due in large part to spin, with a cross section of 1.69 b per unit solid angle [61, 62].

Neutron Capture in the NaI Array: The NaI array is intended to pick up gamma emissions following nuclear events in the Si detector. However, it is possible for it to detect capture-induced nuclear recoils as well if a neutron emitted by the source does not interact with the Si detector. The measured energy spectrum of such events effectively terminates around 2.25 MeV [63], which is fortunately much lower than the 7.6 MeV minimum that we set for the NaI array to detect. What is more relevant to our region of interest are the gammas following these neutron captures; gammas released by capture within the NaI array can be detected by it, each isotope having its own spectrum of possible associated gammas.

III. EXPERIMENTAL SETUP

Because the goal of this study is calibration, the most basic requirements for our setup are a source to create events of interest, the detector to be calibrated – in this case, a solid-state silicon detector – and a data acquisition (DAQ) system. The main data we require is the total energy of each pulse, or individual event, although the shapes of these pulses – described in more detail in Chapter IV – are also saved, alongside a variety of useful metadata. The data is described in detail in Chapter IV. Other constituent hardware is discussed in greater depth in Section 1.

We want to collect data on neutron capture events while discarding unusable data as early in the process as possible. We discussed backgrounds in Chapter II, which are one example of data we would want to discard or avoid including.3. Both background and data of interest have the same shape and can appear in the same energy regions, so it is difficult to discriminate between the two once they've reached the detector, and we naturally do not want to accidentally eliminate data of interest. The most obvious solution is to use shielding to prevent the inclusion of these background events where possible. While we did use various types of shielding, no combination of shielding can eliminate all backgrounds – or even all events of a single type of background – while allowing all radiation of interest through. We can also eliminate some non-capture events by setting the region of interest, which we do by specifying both a minimum and maximum trigger level; setting the minimum prevents the collection of noise, while setting the maximum prevents the collection of high-energy background event data. The final method for filtering out backgrounds is coincidence detection; we can ensure we are seeing certain types events by only including those which trigger not only the Si detector but also the NaI array used for coincidence detection. Remaining "bad" events are eliminated through cuts (see Chapter IV), and while it was not able to be implemented for the dataset presented here, I propose firmware-based methods for preventing collection of some bad data in Section 2.

1 Hardware

The central piece of equipment is the silicon detector we want to calibrate, which sits inside the cryostat. Keeping the detector at ~30 mK eliminates most thermal noise from within the detector [52]. Outside the cryostat and the detector shielding, within a paraffin drum, sits the plutonium-beryllium (PuBe) source; see Figure III.1. The paraffin forces neutrons emitted by the PuBe source to slow down, thereby thermalizing them, lowering the overall energy they contribute and thereby allowing calibration at lower energies and increasing the chance of capture. There is a stack of lead bricks between the detector and the drum to filter out gammas emitted by the source; this lead also helps to thermalize incoming neutrons. Polyethylene shielding around the cryostat serves a similar purpose; it further thermalizes incoming neutrons as well as blocking some unwanted neutron sources.

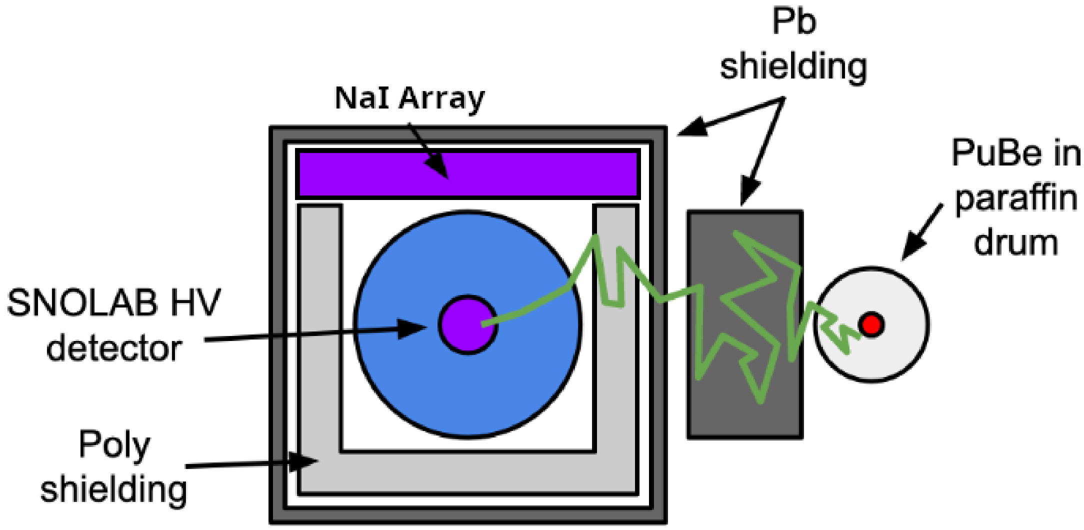

Figure III.1: Detector Shielding

A top view of the source, detectors, and lead and polyethylene shielding between them, adapted from Villano et al. [52]. The green line represents the possible path of a neutron, which slows down each time the direction changes.

Surrounding the cryostat, outside the polyethylene, are sheets of lead, which filter out external gammas that contribute to unwanted background events.

Between the polyethylene shielding and lead sheets sits the sodium-iodide (NaI) array. Data is collected based on a set of conditions determined not by the Si detector, but by the NaI array; Section 2 discusses these conditions in detail. Choosing to trigger on the NaI array allows us to perform coincidence-based detection; we only detect events that have an associated event (ideally, this event is the result of a gamma associated with the initial neutron capture). Non-coincident events become obvious because the Si data will not have an associated pulse for each trigger.

To run the experiment, we use a computer terminal to control two groups of software: the NaI configuration, discussed in Section 2, and the data acquisition (DAQ) system. NaI configuration is run on a local computer and sent through a multi-channel analyzer (MCA), model CAEN DT5560SE [64], which contains the firmware and does a portion of the data processing. The DAQ, which controls only the Si detector, is run from a server on the same network. While the software we used, SuperCDMS DAQ, is not public, it is based on MIDAS, which is [65] Commands to and data from the Si detector must be processed through a DCRC, or detector control and readout card, which is custom SuperCDMS hardware. The DCRC communicates primarily with the DAQ and receives

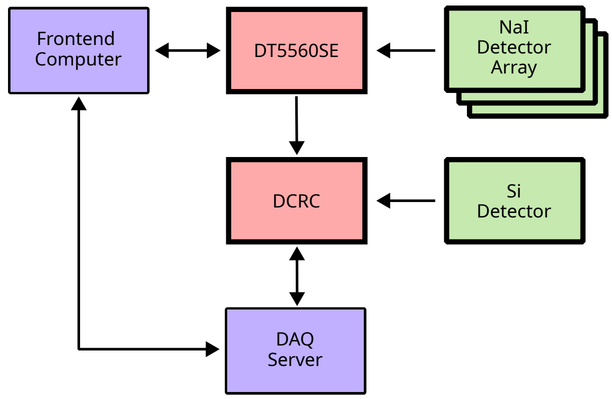

Figure III.2: Network Setup

Network diagram of all computing hardware used for data collection. A front-end computer terminal is used to control the experiment, including controlling the DAQ server via SSH. It connects to the server running the DAQ over the network and has a direct connection to the MCA/DT5560SE via ethernet. The MCA receives data from the NaI array and, under circumstances determined by the front-end computer, sends a trigger to the DCRC. On a trigger, if the DAQ server has told it a run is active, the DCRC sends data to the server to be saved.

triggers from the NaI array. When a trigger occurs, data is saved to the server to be accessed later. Figure III.2 is a diagram of the network of all electronics for this experiment.

While we created boron shielding for the NaI array, this was not used because of issues with collecting further data sets due to a hardware failure of the cryostat. Boron is known to stop some gamma and neutron backgrounds [66] and is cheap to acquire due to the use of Borax as a detergent. The shielding therefore consisted of several long fabric "socks" filled with Borax; these socks have three segments – one for each side – and are depicted in Figure III.3. Similarly, integration between the NaI array firmware and the DAQ is still in progress, and discussed in more detail in Subsection 2.2.1.

2 Multi-Channel Analyzer Firmware

On the MCA, managing all communication with the NaI array and controlled by the local terminal computer, is the firmware. I developed the MCA's firmware and its associated control software. Although I also worked on an MCA + DAQ integration, we were unable to take data that used these changes, but I still briefly discuss them in Subsection 2.1. The primary purpose of the firmware is to output a trigger when certain parameters are met, which was the extent of its functionality at the time of data collection.

The firmware operates by interpreting data from several analog channels. While the MCA supports 32 analog channels, we used only 23 of these channels in data collection, one for each detector. Events appear on one or more channels as "pulses"; an example pulse is shown in Figure III.4. The area within a pulse corresponds to the energy deposited in the detector; because pulses are negative-going, we invert them for the sake of integration, although there may still be some nonzero baseline that we have to subtract out. Beginning with the inversion

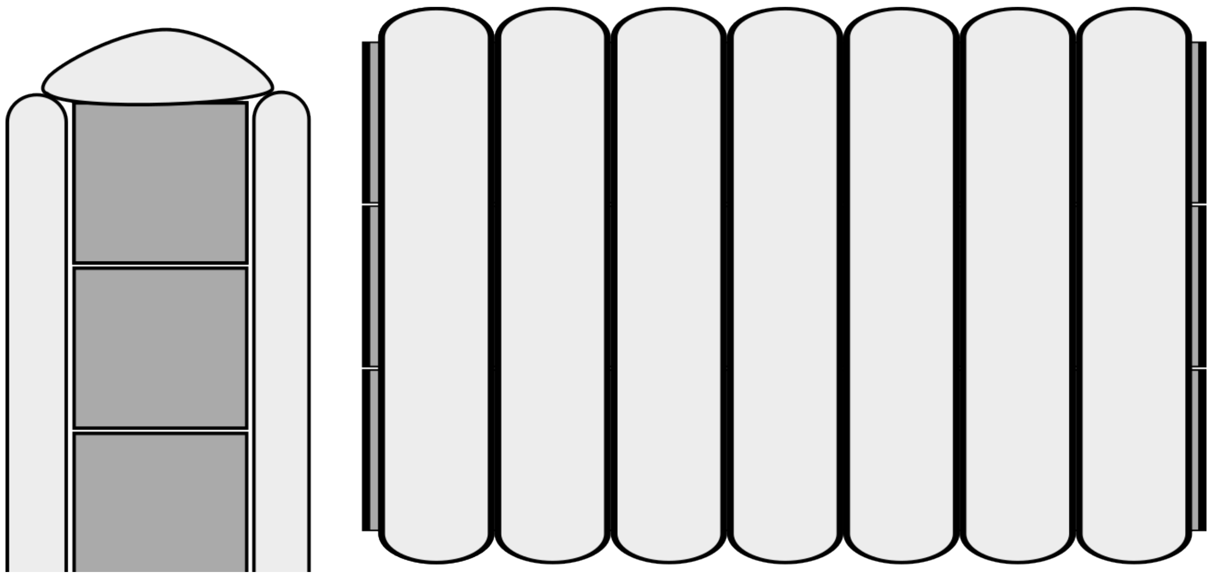

Figure III.3: NaI Shielding

A diagram of the NaI array and boron shielding. Left: view from the side of the array. Right: view from the front and back of the array. The dark boxes indicate individual NaI detectors, while the lighter-colored shapes represent socks filled with borax.

means that higher trigger levels do generally correspond to larger pulses and greater energies. We cannot simply store all data because the files become prohibitively large; a series with approximately 8200 events at around 200 μs is 1.8 GB large but only a bit over 1.5 seconds long, meaning that each terabyte of data could only store about 15 minutes' worth of data.

Each trigger has two components, an upper and lower trigger. Having both creates a window of events that we can trigger on, allowing us to ignore not only noise but also events that are well above the region of interest. The lower and upper triggers are handled separately and on a per-channel basis. If analog data on any channel is at least the lower trigger value, the lower trigger is considered TRUE, and if any channel surpasses the upper trigger value, the upper trigger is considered TRUE. An overall trigger is output if the total lower trigger is TRUE and only if the total upper trigger is FALSE. That is, any channel's upper trigger can veto any channel's lower trigger; it takes only one upper trigger to veto all channels. Gating the upper trigger – extending its boolean signal for a short period of time – and slightly delaying the lower trigger ensures that this process results in the correct vetoing behavior. Figure III.5 diagrams the trigger process for two channels; the OR functions shown can be extrapolated to any number of channels. There is also an inhibition time; for a user-specified number of clock ticks, meant to

Figure III.4: Example Pulse

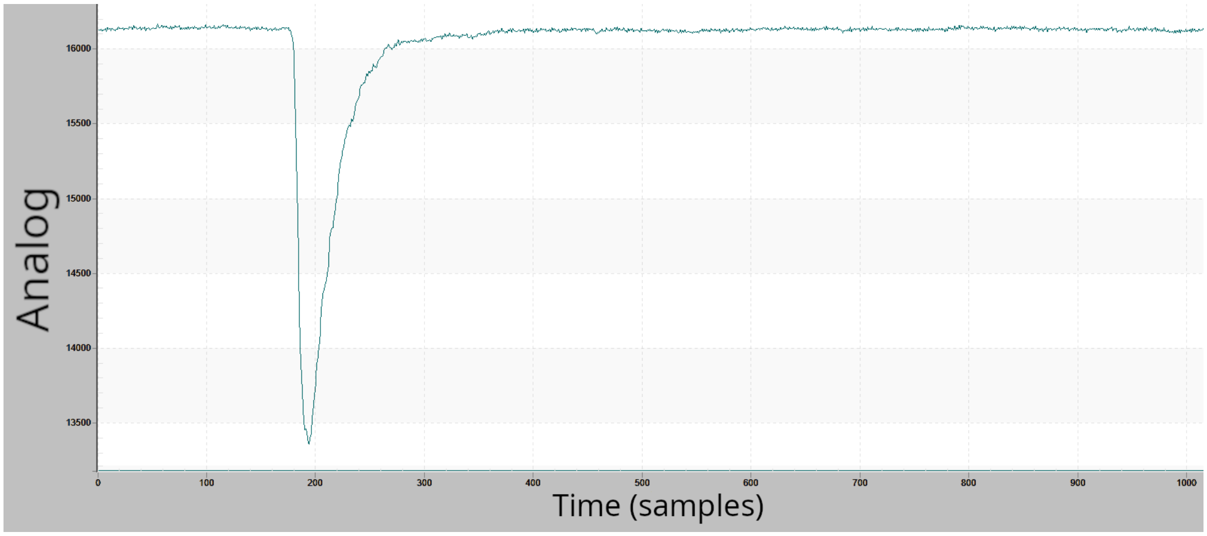

An example pulse. The "ANALOG" values on the y-axis are ADC bin values representing the voltage read from the photomultiplier tubes on the detectors, and the "samples" on the x-axis refer to the MCA's clock cycles. As shown here, pulses are negative-going and must be flipped in order to integrate the area under the curve, which represents the energy.

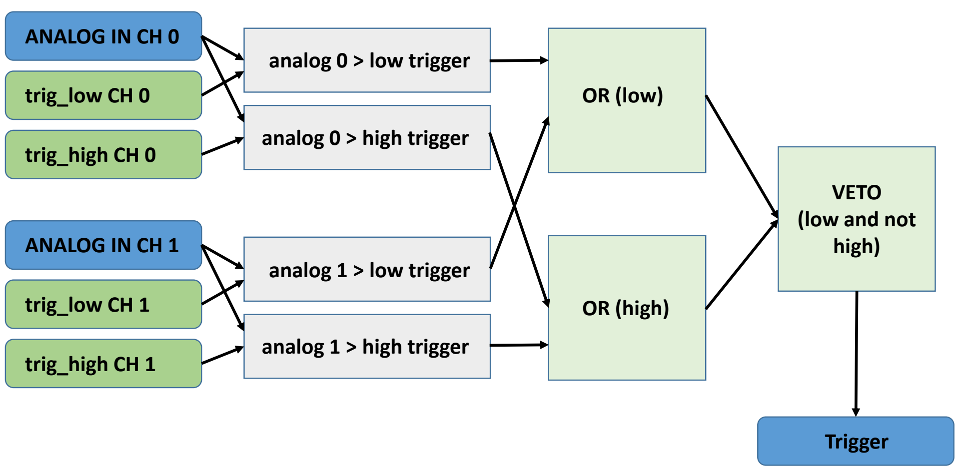

Figure III.5: Simplified Trigger Process

Trigger process for only two analog input channels. Upper and lower thresholds for both channels are configurable. If either analog channel exceeds the upper threshold, the trigger is vetoed; otherwise, if either analog channel exceeds the lower threshold, a trigger occurs.

represent the duration of a pulse, further triggers are ignored, preventing the collection of pileup events – that is, events where the energy baseline is skewed because it falls on top of the tail of a prior event.

Therefore, in order to set the appropriate trigger levels, we must calibrate the NaI array. The pulses, as shown in Figure III.4, provide heights not in units of or including energy, but instead as analog-to-digital converter (ADC) bin values. That is, the MCA has no way of "knowing" what the energy deposited in a detector was and instead takes relative voltage heights at discrete times corresponding to the board clock and stores them as digitized arrays of time and height values. The bins of the histogram are static; that is, they are not adjusted to each individual pulse, so there is consistency in the meaning of the values across events. To appropriately set the thresholds, we need to understand what bin values correspond to what energies, since users set the thresholds in MeV, not in bin numbers. There is an apparent contradiction here; I previously stated that the integral of the pulse over time makes up the energy, which would mean that the bin value can really only correspond to energy per time. However, the shape of pulses is relatively consistent, so we can still calibrate energies to the pulse thresholds. There are two methods we developed for the NaI calibration; we used the method that was simpler to program first, then refined the result using the second.

The simpler NaI calibration method is as follows. We place a radioactive source that will produce clear peaks at well-known energies near the array; specifically, we used ^{22}Na, relying on emission lines at 511 and 1274 keV [46] as well as their sum at 1785 keV. The 1274 keV line is the gamma associated with β decay, while the 511 keV line is from the annihilation of the same β, so they coincide frequently enough that the sum of the two can be used for calibration. The upper threshold is set very high so as to be effectively disabled. The lower threshold is then swept across a wide range of values, starting near the baseline and incrementing by some predetermined value until it reaches the upper threshold. At each threshold, the total event rate is retrieved from the MCA, which measures the rates by averaging the number of events over a limited, rolling time period. The thresholds and associated rates are saved to a file on the local computer as they are retrieved. I then determined the increase in the rate by subtracting the next rate from each value and plotted them to locate the peaks; an example plot is given in Figure III.6. The difference in rates is useful because it tells us what each decrease in the lower threshold would add; that is, it serves as an approximation of what the rate is between the two thresholds. These peaks align with the energies associated with the source's emissions; we can therefore define a linear relationship representing the energy of an event and the ADC bin value of the pulse height that allows us to see it:

$$B = c + sE_{keV} \tag{2.1}$$

gives a bin number B for an energy E_{keV} given a calculated baseline c and scale s. This calibration is done for each channel individually, as they were found to have differing values.

Once the initial calibration was completed, I built the conversion of energy to bin value into the program that sets the thresholds. Doing so allows us to more easily set multiple thresholds systematically for each detector before increasing the complexity of the rate-measuring software used for calibration. Conceptually, the extra step in the updated program is very simple – I made the upper threshold sit close to the lower threshold, sweeping

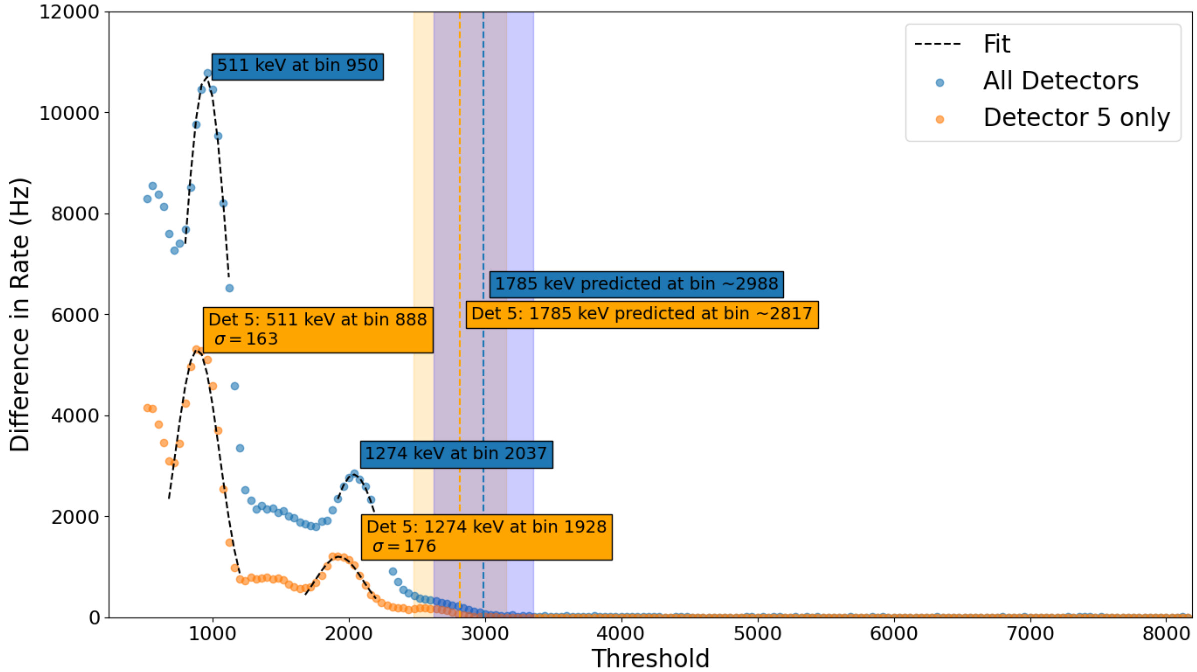

Figure III.6: Example Trigger Rate Measurement

An example plot used for the trigger rate-based calibration for a single detector and the total rate. Vertical lines are placed at peaks of each in the order of their appearance and labeled with the associated energy. The "difference" in rate mentioned is the difference between the current threshold's rate and the next one's. I choose to show the plot for the difference method to demonstrate the ADC bins on the x-axis; later calibrations with the improved window method list provided energies instead of bins, as the thresholds recorded are based on this calibration.

it upward at the same rate as the lower trigger. This creates a small "window" which eliminates the need for subtracting rates from each other. The rest of the method is the same as the first, and the results have been programmed into the threshold control software. The final NaI calibration results are given in Table III.1.

Firmware is constructed using the proprietary software Sci-Compiler [67], which is produced by CAEN and designed for their hardware products. We use this for its compatibility with our MCA, a DT5560 [64]. Due to the complexity of the firmware, a diagram encompassing all of its functions in full detail at once was not readable; therefore, we include the simplified diagrams below alongside the explanation in order to give a full, but still clear, idea of how it works. For access to the full firmware, see [68].

As part of the calibration of the NaI system, we must determine out the baseline of each channel. We did this by taking multiple runs, reading the baseline on the oscilloscope, and adjusting until the minimum energy was near 0. The baseline is subtracted during the integration phase, so triggers are offset by the baseline.

2.1 Further Proposed Firmware Upgrades

Table III.1: NaI Array Calibration Values

Detector	Baseline	Scale
1	288.21	1.44
2	258.69	1.29
3	318.06	1.28
4	233.09	1.29
5	197.68	1.28
6	321.82	1.36
7	261.98	1.15
8	284.21	1.17
9	389.04	1.42
10	329.41	1.29
11	282.30	1.32
12	319.53	1.26
13	389.48	1.22
14	203.09	1.36
15	421.60	1.12
16	272.15	1.34
17	308.12	1.23
18	267.85	1.28
19	293.62	1.22
20	299.07	1.44
21	286.40	1.18
22	243.29	1.06
23	334.94	1.36

A bin number is calculated from an energy following Equation 2.1.

Due to issues with the helium liquefier used by the detector fridge, we were not able to take new data that included features associated with more recent firmware work; the firmware was only able to provide a binary trigger and no additional information. This limitation is a result of the network setup; providing additional data would require some way for the MCA to talk to the DAQ, or to at least tag data identically to the DAQ in a way that would allow the data to be associated later, but altering the DAQ close to the R76 data collection time was too great a risk as resulting issues with the code could halt data acquisition entirely or corrupt the data set beyond readability. I will therefore keep the discussion brief, as these features are useful but do not pertain to the data analyzed in this thesis.

The most impactful change intended for the NaI array that did not make it into the data was the ability to store and correlate energy data from the NaI array to event data from the Si detector. Following a trigger, the firmware integrates the pulse on all channels – also including those that did not cause the trigger – and stores the energies in a custom data packet, but at time of acquisition, there was no way to associate the resulting data with the correct Si data, so it was not stored. One solution is storing, alongside the energies, a timestamp; however, the MCA cannot generate timestamps in a standard time, only based on an internal clock that reaches its maximum and rolls back over to zero frequently. The speed of the clock is such that it will cycle back to zero multiple times during a data run, but because event times are not regular in any way, the timestamp can be treated as a quasi-random number – at least random enough that an event should be identifiable by its timestamp and, if necessary, the stamps of the events near it. Therefore, we needed a way to send the stamps to the DCRC so that they could be saved to the Si data as well. Doing so required altering the DCRC firmware so that it could receive data directly from the MCA, as a connection stopping at one of the computers on the network would be too asynchronous to match events to stamps. Sten Hansen at Fermilab was kind enough to provide us with both the DCRC firmware and the custom VHDL needed for the firmware to communicate with it; the code for both of these is in the firmware repository [68]. The DAQ must also be changed in order to accept and save this data; DAQ changes are still a work in progress.

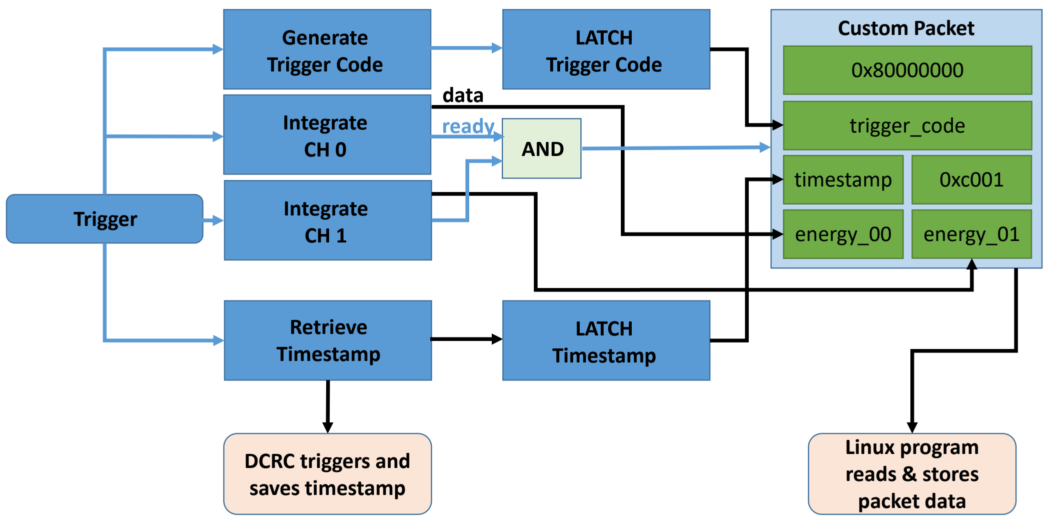

Figure III.7: Basic Firmware Layout

A simplified firmware diagram, focused on the conversion of a trigger event to an output, with only two analog input channels. Following a trigger event, a trigger code is generated, integration of both channels begins, and a timestamp is retrieved, all simultaneously. The timestamp is sent to the DCRC to be recorded; the timestamp and the trigger code are then held in memory until integration is complete. All four pieces of information are then compiled into a single packet which is later read by the computer program.

Storing NaI data would tell us what the associated energy of each gamma component was, helping with the identification of event types. For example, we may simultaneously see captures in two or three detectors that indicate an annihilation event as for the 511 and 1274 keV lines in ^{22}Na, which would indicate (unless we are using a Na source) a background β-decay event. Alternatively, we may be able to discard events if too many detectors are triggered; events in which ten detectors triggered indicate sources much stronger than anything we're interested in and are therefore also backgrounds. Tracking which detectors triggered can give us directional data as well; particles interacting with the Si detector that have non-negligible initial momenta will have resulting gammas in predictable directions, assuming that they come from a known source. This predictability means that we can check the energies of events triggering certain sets of NaI detectors against a list of predicted values in order to determine whether they came from the source of interest. We codify the list of triggering detectors into a single value stored in the packet alongside the rest of the NaI data; this packet and how the data within it interacts with the firmware and the DAQ is shown in Figure III.7. While, as discussed in Section II.1, the initial momentum being negligible means that the resulting recoil is in a random direction, we still study directionality for the sake of future experiments. For example, an experiment may look for dark matter particles passing through the detectors in a particular direction due to the Earth's movement through relatively stationary "clouds" of dark matter, but the probability that the dark matter both passes trough the Si detector and triggers the NaI array is very near zero. By studying patterns associated with directionality now, we can extrapolate directionality from future experiments where coincidence is not a possibility.

In addition to DAQ integration, I have two minor features solely within the firmware that were not available for R76. The first is an improvement to the pileup rejection. While we were able to reject events after another in the case of a pileup using an inhibition value equal to the pulse integration time, we were not able to reject the first event within a pileup, even though its tail is skewed by the rejected event. To some extent, this behavior is intended – once the trigger has been sent to the DCRC, which must happen immediately, it cannot be revoked. However, we could more easily reject pileups in data analysis if there were some flag marking pileup events retroactively, and such a flag could be applied to a data packet before it was sent. This would make pileup rejection a simple process of discarding all events with the pileup flag.

The second feature is automatic baseline detection. The baselines for each channel drift over time in a way that we do not know how to predict, meaning the calibration needs to be redone periodically. However, recalibrating during a long data run is not a practical solution to drift, and baselines can have subtle but meaningful changes over even short spans of time. Therefore, I added an experimental baseline calculation to the firmware itself. Following the end of a pulse, it calculates the baseline and attempts to subtract it from the pulse for the next integration. This behavior is configurable; the threshold-setting software can be configured to provide only the scaled bin values sE_{keV} (see Equation 2.1), but it remains possible to use the calibrated, hard-coded baseline values as well.

IV. ANALYSIS

The dataset used in this thesis, referred to as R76, was collected between March and December 2022 and amounts to approximately 2 TB in total. During this time, the CDMS Si detector was exposed to three different sources: ^{22}Na and ^{241}Am as gamma sources and PuBe as a neutron source; neutrons are thermalized by the shielding, as discussed in Section III.1. These sources were used in conjunction with a variety of shielding configurations, the most "complete" of which uses the shielding shown in Figure III.1 as well as the boron shielding depicted in Figure III.3. The polyethylene was always present, as well as lead surrounding the experiment to minimize exposure, although the lead bricks between the source and detector was not always in place. Most data is done with coincident triggering, but data triggered only on the Si was also taken for specific purposes, discussed later in this chapter. For the placement of each component, see Figures III.3 and III.1. The combinations of source and shielding used are outlined in Table IV.1.

Table IV.1: Sources and Shielding

	^{22}Na	^{241}Am	PuBe
Basic Only	✓		✓
Lead Bricks		✓	✓
Lead Bricks + Borax		✓	✓

A shielding type of "basic" refers to when only polyethylene and external lead were present. All other shielding configurations also include these two components.

Data is stored in the MIDAS [65] format, which contains metadata and a packet with information about multiple events. The packet is made up of a number of "words", or four-byte segments. A header indicates the number of triggers read and a footer indicates the number of triggers written at the end; ideally, these are identical, but there may be cases where a trigger is discarded because there is something wrong with the data, in which case it is read but not written and has no data in the packet. Unrecorded events are the result of data being lost during communication between the DAQ and DCRC; various headers – pre-determined strings of data – allow the DAQ to determine that a piece of data is incomplete, discard it, and successfully pick up from the beginning of the next event. Sandwiched between these are repeating segments that make up the meat of the data, one for each trigger. Each of these segments are made up of three blocks, one for each DCRC, beginning with information about the trigger and the DCRC. Within each of the DCRC blocks are six smaller blocks, one for each channel, although in practice only the first four channels contain meaningful data. The channel blocks begin with an indicator of the channel and the number of data samples associated with the trigger, followed by the data samples; see Figure IV.1. Samples can be reconstructed into a pulse, as shown in Figure IV.2. Also included is various timing information, including the time that the trigger started, when it was retrieved, and how long it took to retrieve.

The data contains signals from a wealth of sources, including both the backgrounds discussed in Section II.3 and the neutron capture events of interest; Section 1 discusses the elimination of unwanted noise and corrupted

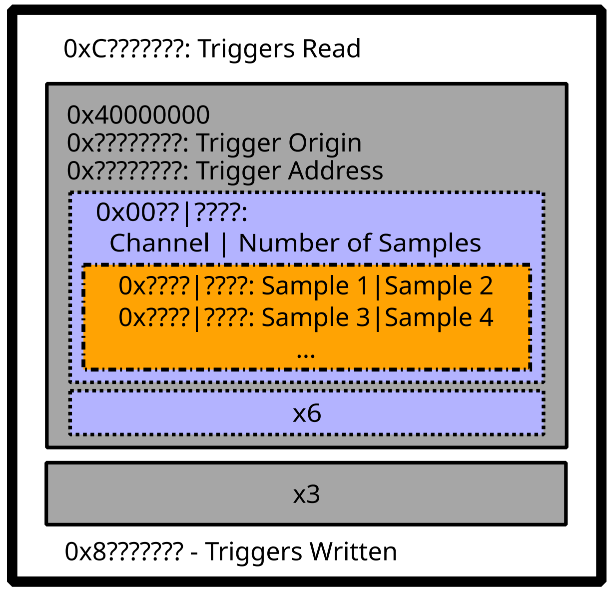

Figure IV.1: DCRC RevC Data Packet Format

The structure of the DCRC RevC data, limited to the most basic information for clarity. A "?" indicates that the hex character represents some specific data that varies, and a "|" indicates not a character but a division of the word into smaller pieces of data, while all other characters are static and used as reference points. The UMN data has a total of three DCRC blocks per packet and a total of six channel blocks per DCRC.

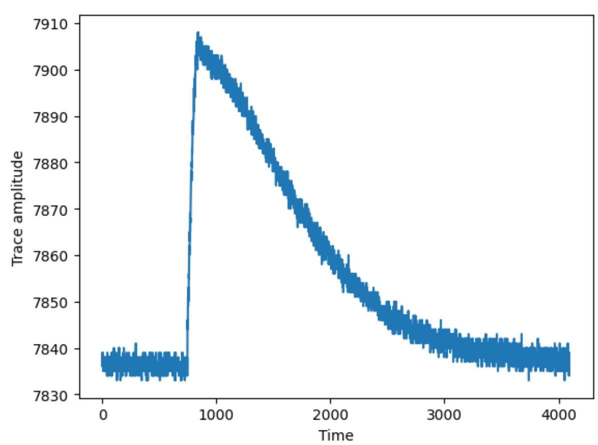

Figure IV.2: Si Detector Pulse

An example of a pulse for a single detector channel. The time on the x-axis refers to clock ticks of the DCRC, which runs at 50 MHz, and the amplitude on the y-axis refers to a digitized ADC bin value.

Table IV.2: Cuts

Name	Event Rejection Criteria
Baseline	The raw pulse baseline's average is above specified values.
Glitch	Excessively fast fall times.
OF χ^2	Poor fit to the optimal filter.
Low-Energy Burst	Events with specific behavior that do not fit any of our models and are correlated to each other.
Pileup	Pileups, based on multiple factors pertaining to the OF.

events through cuts as well as necessary pre-processing. Once the data is collected, pre-processed, and trimmed, we can fit for the best detector calibration, which I do in Section 2.

1 Data Cuts

Before data can be properly analyzed, it must be processed. We can automate some of this processing during data acquisition, such as the conversion of the NaI array's pulses into energies mentioned in Section III.2. This section outlines only processing and analysis that is done on acquired data – that is, this subsection outlines only pre-analysis, not processing that is done automatically or during data collection, such as charge integration that happens within the array firmware.

I apply and adjust various data cuts, beginning with those laid out in Mast [69]. The specific purpose of each individual cut differs, but the general purpose of data cuts is to remove data points that are either from sources we are not interested in analyzing or are of too low quality to be useful for analysis. Unwanted sources can be physical events, such as emission lines from background sources; detector noise; or glitches that occur during data acquisition. While I will remove as much noise as possible from the data, because we want the result to be as accurate and as widely applicable as possible, for the purpose of calibration, we only need to remove enough noise to keep the source's major low-energy peaks visible and prominent, as it is the energy and, to a lesser extent, relative heights of these peaks that we concern ourselves with in calibration. Table IV.2 summarizes all cuts; quality cuts are described in further detail below. The efficiency of a quality cut is defined as the portion of signal events that survive the cut for a given energy; I provide only energy-independent efficiencies here, as they otherwise do not differ significantly from those of R68 [52].

Baseline: The baseline of an event refers to the region where nothing is happening that the pulse effectively sits on top of. For example, the baseline in Figure IV.2 sits somewhere between 7835 and 7840. The baseline of an event is calculated using data well before the trigger time. Events that begin on the tail of a previous event, or pileups, will have an elevated baseline preceding the trigger and are therefore eliminated by limiting the mean. The baseline cut therefore rejects events where the mean of the baseline is too high. We cut all events with an ADC baseline mean of at least 1100 above the adjusted baseline, calculated using 320 μs of data. While Mast [69] also cuts events based on their baseline variance, doing so only works if the data being cut is the same data that was triggered on, as the purpose is to eliminate double-triggers caused by noise in the tail of an event. Because the triggering in R76 is done by the NaI array, it does not make sense to look for double-triggers by analyzing the Si detector data. The baseline cut's efficiency can be measured simply and directly, without energy dependence; since a randomly-triggered event should be near-perfectly flat and cannot be correlated to other events, the efficiency of the baseline cut is the portion of randomly-triggered events that it does not reject; it was found to be 0.820 ± 0.001.

Glitch: Some events have abnormal shapes due to DCRC hardware glitches in which data from an incorrect register is inserted into an event. These events have very fast fall times, to the point of physical impossibility for phonon signals. The fall time is calculated as the time to drop from the maximum pulse height to 20% of the maximum (measured from the baseline), and events with a fall time of less than 35 μs are rejected, following Mast [69].

OF χ^2: We apply to each pulse the optimal filter – this is very similar to fitting a function describing the shape of a pulse [69]. The OF template fits low-energy pulses well, but does not fit high-amplitude events where the pulse shape is distorted due to saturation. We can therefore reject these distorted events by rejecting events where the OF fit is of poor quality, as determined by a χ^2 value calculated from the difference between the OF and the data. The cutoff χ^2 value is $1.25 * 4096$, where 4096 is the number of samples per event, following Mast [69].

Low-Energy Burst: Within the dataset are occasional bursts of low-energy (below 150 eV$_{ee}$) events with rates of 1 kHz or higher. While these events are of interest to other analyses [69], their correlation to each other means they do not match our existing models, so it does not make sense to analyze them for calibration. These bursts do not appear in the background dataset, only in those with a source present; the low-energy burst cut is therefore not used for the background dataset. In the remaining data, a burst is considered to occur whenever two or more events of energy less than 1 keV$_{ee}$ occur within 20 events of one another. Because the energy region the cut is applied to is small, we calculate the efficiency as energy-independent; it was found to be 0.893 ± 0.001 below 1 keV$_{ee}$ and 1 elsewhere.

Pileup: Some events appear to contain multiple pulses, but are not removed by the baseline or OF χ^2 cuts because the OF is able to fit to one of them successfully. These pileups can be detected in multiple ways. Because pulse shapes are fairly consistent, an anomalously large ratio between the pulse's integral and the OF fitted amplitude indicates a region of the pulse has not been accounted for. Therefore, we reject events where this ratio is 3σ or more above the median ratio. Additionally, we can analyze the delay in the optimal filter, that is, the amount of time the start of the pulse has to be shifted for the OF to fit. The OF is given a wide range that begins well before and ends well after the pulse, meaning difficulty fitting within this window indicates that something has gone wrong with the fit. Events that are very close to the beginning (within 1 μs) or end (within 2 μs) of the window are likely trying to fit something outside of the window and can therefore be cut. The third pileup detection approach is to require some consistency between the pulse rise time and the delay. Significant misalignment between the delay and the point at which the pulse first reaches half of its maximum are typically indicative of

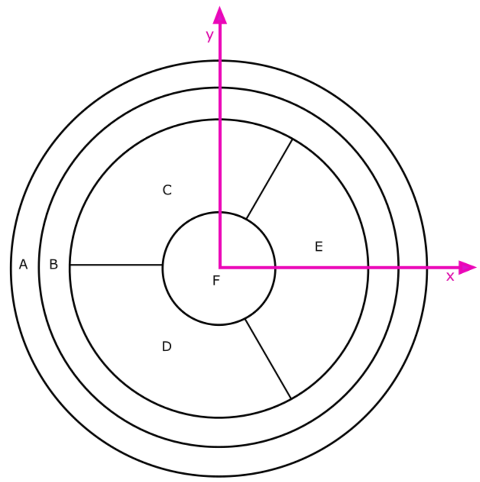

Figure IV.3: Detector Channel Layout

The channels as present on the Si detector. The x and y axes are drawn to be consistent with the orientation of the Phi cut. Courtesy of Layne Hellman.

noise or pileups, so we can reject events where the delay is at least 70 μs before or 10 μs before the halfway rising time. The pileup cut is a combination of Mast's [69] OF cuts, and its overall efficiency was 0.965 ± 0.001.

Phi: The volume of the silicon detector is divided into channels, laid out in Figure IV.3, giving us some information regarding the position of an energy deposition, enhancing coincidence detection. The OF delay can be approximately split into x and y components by analyzing the difference in the delay across channels C, D, and E, giving an angle

$$\phi = \frac{180}{\pi \arctan \frac{x}{y}}. \tag{1.1}$$

Because we trigger coincident events on the NaI array, we did not use this cut for coincidence detection; instead, it is useful for an initial calibration that occurs in the higher-energy region and is based on ^{241}Am events. We took an angle between $5° - 20°$ to indicate a ^{241}Am event.

To work, these cuts require a variety of calculations to be made ahead of time.

1.1 Salting

The basis behind data cuts is sound – not all data we capture is the targeted data, so we can reduce noise by removing the "wrong" data. However, there are two potential issues that can arise from poor cut choices. The first possible problem is that we may accidentally cut out the signal we are looking for. The second, inverse problem is that we may believe we have cut the noise and are now seeing the signal, but part of what we see in the "signal" is remaining background noise. Neither of these are major issues for the calibration, as we are looking for large signals that easily stand out over reduced noise, so even a very rough approach to cuts would be sufficient. When looking for signals such as dark matter candidates, however, even small amounts of noise can easily drown out the signal of interest, meaning that fine-tuning the cuts becomes much more important.

One potential method for verifying the implemented cuts is through the use of salting, which generates real-looking data to allow us to fine-tune the analysis process in a controlled setting. A machine learning model is trained on the coincident-triggered portion of the dataset, and we refer to the synthesized events it produces as salt. If the model and the generation of salt are working correctly, then a comparison of the original dataset and the salt will reveal that they have the same distribution but not identical elements. We can then add to a real data series salt where some known portion of the signal represents dark matter candidates; this portion will be generated by providing constraints on the pulse energy consistent with the region in which we expect to find dark matter signal. Some of the R76 data are randomly-triggered events, or “randoms”, meaning they contain only the baseline with small variation due to noise in the equipment, but no pulse or event signal of any kind; salt events are added to these randoms in order to give them realistic noise, and the randoms are then given a flag indicating that they are salt. From here, the analysis continues normally; the salt flag is completely ignored throughout all analysis tasks and the real data and the salt are analyzed together as if they were one singular data set. Because the DM candidate signal in the salt was known before analysis, we can verify at the end of the process that the cuts used actually select for the DM candidate signal with minimal inclusion of noise. If not, then we can adjust them as needed and repeat the process. Once the cuts are working and verified, the salt can be removed, and whatever remains can be analyzed. This method provides us with a baseline for verifying cuts, as the ability to measure them is currently obscured by our not knowing what the final results should look like.

While some progress was made on the implementation of salt, the salt format does not yet match the format of the data, and there were delays in its generation, so it remains future work. Because the salt does not apply to the current analysis, and because I did not personally create or train the machine learning model, I keep discussion of the current progress brief.

The salting model was trained using a fork of TimeGAN [70, 71], a generative adversarial network (GAN) for generation of datasets with a time component. A GAN, in short, consists of two competing networks; the generative network attempts to generate data that looks “real”, while the discriminative network evaluates whether or not data is from the original or the generated data. The initial training was done on an older dataset, R68 [52], because it was already set up to be shared easily with non-UMN collaborators; this dataset is similar to R76, the primary difference being that it lacks the NaI array, so while the data format, detector, and sources are the same, there is no coincidence-based triggering. The resulting model is able to generate salt, but we are not currently able to salt the dataset – that is, combine the salt with the dataset – which prohibits evaluation of the quality of the generated salt. A variety of methods exist to check the quality of a machine learning model; it is currently proposed that we use the Kullback–Leibler information number, or I-divergence [72], to measure the differences between the salt and the R68 data. The I-divergence, I, is a measure of the distance between a model probability distribution Q and a true probability distribution P,

$$I(P||Q) = \int p(x) \log \frac{p(x)}{q(x)} dx \tag{1.2}$$

for continuous x or

$$I(P||Q) = \sum_{x \in X} P(x) \log \frac{P(x)}{Q(x)} \tag{1.3}$$

for discrete x. It is worth noting that this distance is not symmetric, meaning the order of the arguments matters, as noted by the $||$ symbol. Despite being empirical, we use the data set as P as the "true" PDF because the goal of the model is to replicate the R76 data, not a theoretical distribution.

Once we have verified the salt, we can begin adding physics constraints as well. Machine learning models struggle to accurately model physical systems, particularly when those systems have a considerable level of noise [73–75]. Providing broad constraints based on theoretical models can improve both accuracy and speed of training [73]. We can modify the existing model as well as create a second model designed to only generate dark matter candidate events by restricting the range of permitted energies. Whereas the generative network is normally scored only on how well it matches the data, a second loss function can be added to this one that penalizes generated data that falls outside the provided limitations [74].

Issues with the application of the generated salt boil down to technical issues regarding changes in the CDMS data format over time. The channel mappings used in data collection at UMN, provided in Table IV.3, differ from the one-to-one mappings assumed by the more recent SuperCDMS libraries that the salting program is built on; unfortunately, the recent libraries are not compatible with the legacy equipment used for R68 and R76, but the salting program needs to be compatible with new data, prohibiting the exclusive usage of older libraries, and the two are not similar enough to allow the salting program to be built on either one interchangeably. This difference in channel mappings causes salt to become scrambled when it is applied to the data, shuffling data, salted or not, between channels in the best cases and completely changing the data to nonsensical values in the worst cases.

2 Ionization Measurement

The ultimate goal is to be able to calibrate using thermal neutrons — that is, neutrons with very low energy, around 0.025 eV, corresponding to thermal motion at a temperature of $20°$C [47] -– because we expect dark matter particles to cause similar recoils to those caused by thermal neutron capture, because we know dark matter has either no or minimal interaction with electromagnetic forces, much like a neutron has no net charge. The R76 data is most tailored, however, to events with a NaI coincidence much higher than corresponds to the region of interest. There were plans to take a new dataset as well as to prepare more R76 data for calibration, but these plans were

Table IV.3: Channel Mapping

DCRC Channel	Detector Channel
DCRC1	
Side 2 HV2 pA	Side 2 pC
Side 2 HV2 pB	Side 2 pF
Side 1 HV2 pC	Side 1 pC
Side 1 HV2 pD	Side 1 pF
DCRC3	
Side 1 HV1 pA	Side 1 pE
Side 1 HV1 pB	Side 1 pB
Side 1 HV1 pC	Side 1 pD
Side 2 HV2 pD	Side 2 pE
DCRC4	
Side 2 HV1 pA	Side 2 pB
Side 2 HV1 pB	Side 2 pA
Side 1 HV2 pB	Side 1 pA
Side 2 HV1 pD	Side 2 pD

Table IV.4: Americium Peak Fits

Expected (keV$_{ee}$)	μ (keV)	σ (keV)	Amplitude	Red. χ^2
13.95	153.1 ± 0.1	2.888 ± 0.186	310.0 ± 14.1	9.89
17.74	170.2 ± 1.0	4.060 ± 0.801	371.1 ± 22.8	12.0

waylaid due to equipment failures. Therefore, we proceed with the R76 data despite the limitations. Once cuts (see Section 1) have been applied, we begin with an intermediate calibration above the region of interest to get a reasonable starting point for our parameters. This calibration is fairly simple; we analyze data triggered on the Si detector taken with the ^{241}Am source and complete shielding. The subset of the data used for this pre-calibration is fairly short because it is selected to be a single, contiguous, uncorrupted section of data with a similar voltage bias (-84.3 V) to the PuBe data we will use for the final step and a minimal level of non-^{241}Am events. For this, we can use the 13.95 keV$_{ee}$ and 17.74 keV$_{ee}$ lines shown in Figure IV.4; peaks are assumed to be Gaussian. The fits to the Americium data, with the baseline cut applied, are demonstrated in Figure IV.5 and results are given in Table IV.4.

Once the peaks have been fitted, we construct a fit to the two selected points and to $(0, 0)$, signifying that a non-event (or, equivalently, an event of zero amplitude). Because there are only two (three, counting the origin) points to fit, the function we fit to should have no more than two parameters. The relationship is nearly linear at low energies, but as pulses approach saturation, the reported energies become closer together. We therefore use the empirically-motivated relationship

$$E_{cal} = \frac{cE}{a - E}. \tag{2.1}$$

Figure IV.4: Non-Neutron Calibration

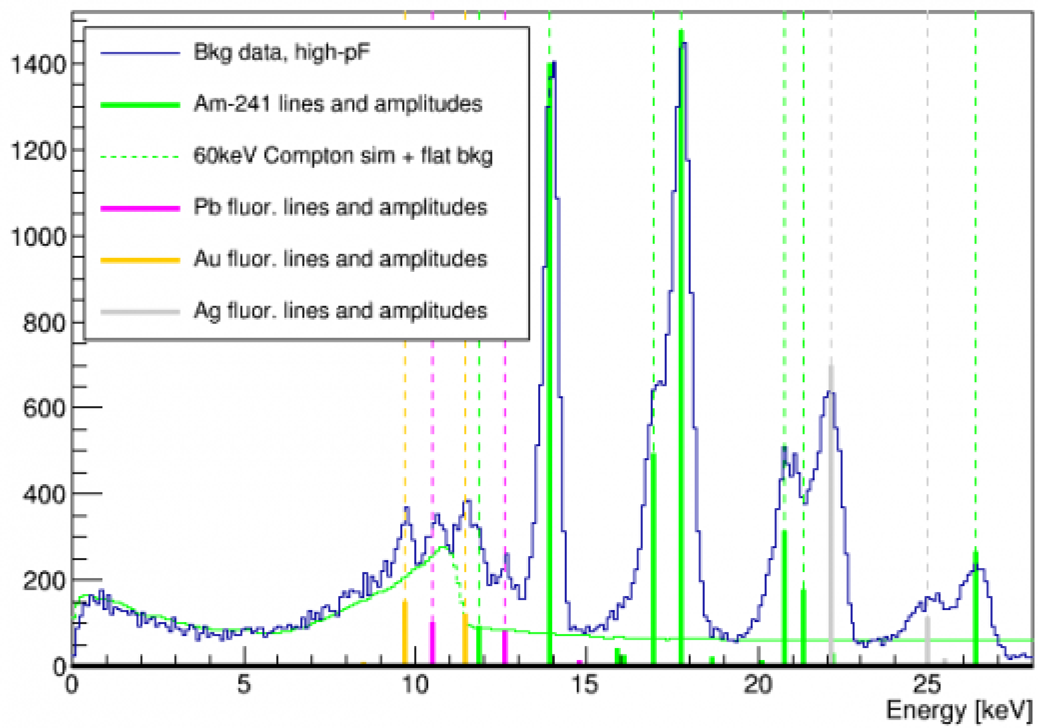

A spectrum of R68 (old) data taken with a ^{241}Am source alongside the relevant calibration lines. Note that the calibration lines have different frequencies relative to one another. Adapted from Villano et al. [52].

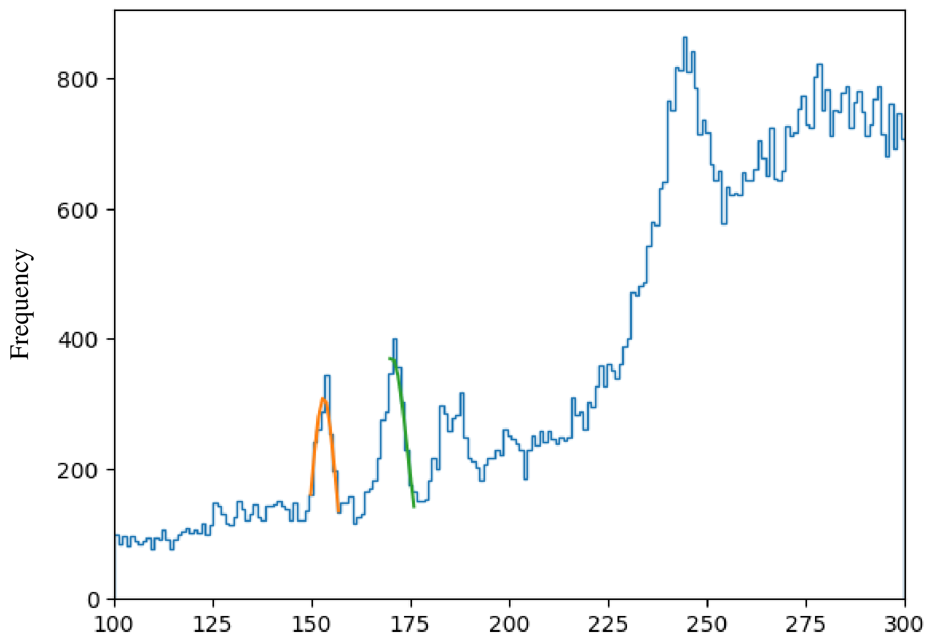

Energy (keV)

Figure IV.5: Americium Calibration

Energy values on the $x-$axis are for relative positioning only, as they are not electron-equivalent. The orange curve (left) and green curve (right) are fits to the 13.95 keV$_{ee}$ and 17.74 keV$_{ee}$ lines, respectively. Due to the proximity of the lines near 18 keV$_{ee}$, we calibrate the right-hand portion of the second peak only to avoid skewing.

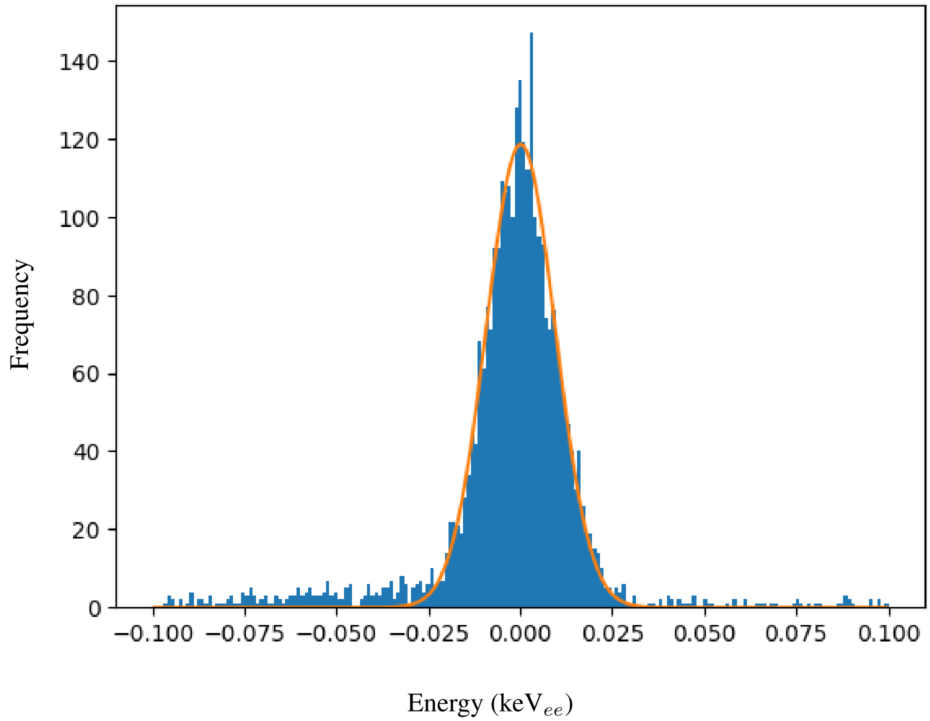

Figure IV.6: Noise Fit

Noise width, generated from ^{241}Am data. The line plotted is a Gaussian fit to the data with a width of 9.34 eV$_{ee}$.

With two parameters and two data points – the function is selected to always cross $(0, 0)$, so it is not added as a third point – there is no room for measuring the goodness of fit. The parameters, using the fits from Table IV.4, are $c = 12.3, a = 288$.

Using the result of this ^{241}Am-based calibration, we can examine the noise width by looking at only random events, applying the baseline cut to prevent the inclusion of randoms taken shortly after an interaction. Knowing the width of the noise gives us a better understanding of the variability within the data as well as how to apply resolution models to simulated or calculated predictions. The noise, shown in Figure IV.6 appears approximately Gaussian, and fitting it as such gives it a width of $\sigma = 9.341 \pm 0.126$ eV$_{ee}$ with a rather high reduced χ^2 of 21. I attempted to fit a Cauchy distribution and to apply both fits to other subsets of the data, but the fit result came out much worse, as seen in Figure IV.7, and the Gaussian more closely matches the center, as seen in Figure IV.7. Therefore, $\sigma = 9.34$ eV$_{ee}$ is my best estimate for the noise width.

The R76 data presently available to me and ready for calibration is at a higher than ideal triggering threshold, but there are still two visible peaks in the sub-keV region of the NaI-coincident data to calibrate. This data was taken with the NaI array set to a window trigger of $7.6 - 11.2$ MeV and uses the pileup and χ^2 cuts. The visible peaks are likely the results of direct-to-ground captures, as single-step cascades have the highest energies and the least spread. They best match those of ^{29}Si and ^{31}Si, which, before yields, would have associated recoil energies

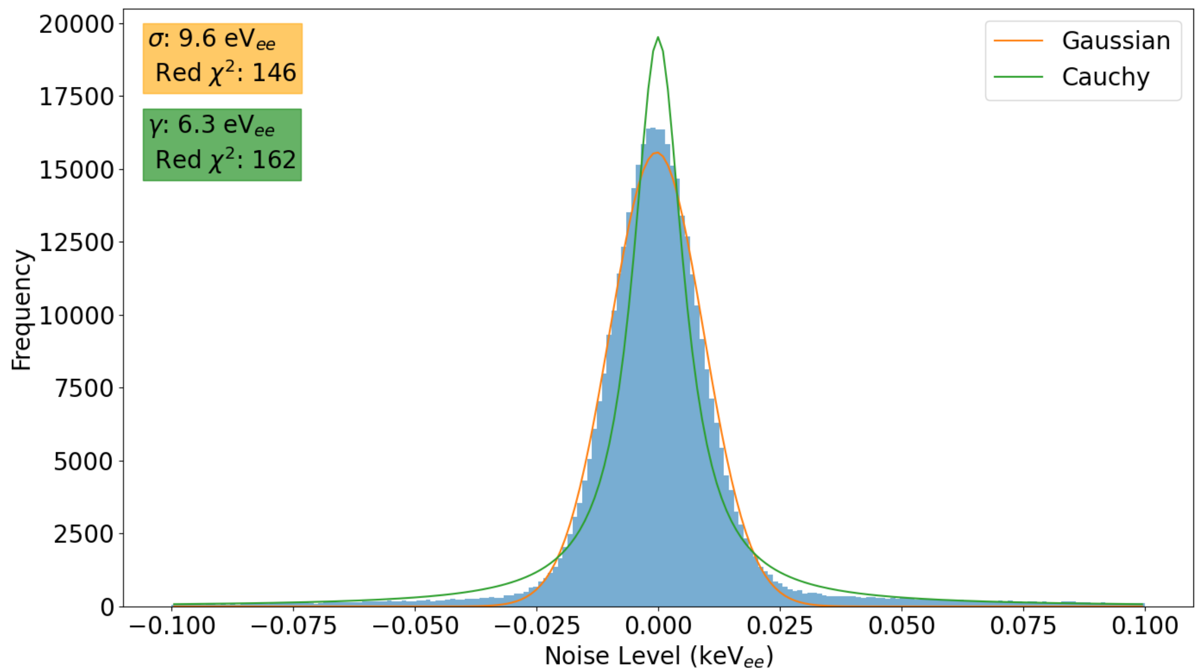

Figure IV.7: Alternate Noise Fit

A second attempt to fit the noise width with a much larger subset of the data, as well as with an additional model with a wider tail.

of 1.33 keV and 2.01 keV, respectively. We performed a log-likelihood fit of the data using Markov Chain Monte Carlo using the emcee python package [76] and the following model for the relative frequency of an energy E:

$$0.00999\left(f_{exp}\frac{e^{-E/\lambda}}{\lambda}+f_g\left(\frac{p_{29}g_{29}(E)}{p_{29}+p_{31}}+\frac{p_{31}g_{31}(E)}{p_{29}+p_{31}}\right)+f_{flat}\right), \tag{2.2}$$

where $g_i(E)$ is the Gaussian function of variable width σ_i and mean $\mu_i = E_i y_i$, E_i are the expected energies given above. The Gaussians represent the two peaks, the exponential the noise spike at the leftmost side of the data, and f_{flat} the baseline. The relative heights of the peaks are defined by p_1 and p_2, but only one needs to be variable, so we fix p_1 at 0.020846, the calculated probability of the single-step cascade for ^{29}Si. Similarly, the leading 0.00999 is a scaling factor that could equivalently be absorbed into f_{exp}, f_g, f_{flat}. We want to fit for the yields y_i and their widths σ_i, and $f_{exp}, \lambda, f_g, p_{31}, f_{flat}$ are all fitted parameters. The resulting fitted curve is shown in Figure IV.8 alongside a corner plot demonstrating the variance and covariances of each parameter; the yields are $y_{29} = 0.1123 \pm 0.0060$ and $y_{31} = 0.1293 \pm 0.0033$. This fit is preliminary; further work can be done to get a fit in this region with more data, potentially improving the fit and even revealing another peak.

With just two points, we cannot truly compare yield models, as they all would be over-parameterized. Using the parameterizations from R68 [52], neither the Sorensen nor Lindhard models appear to fit; the lower of the

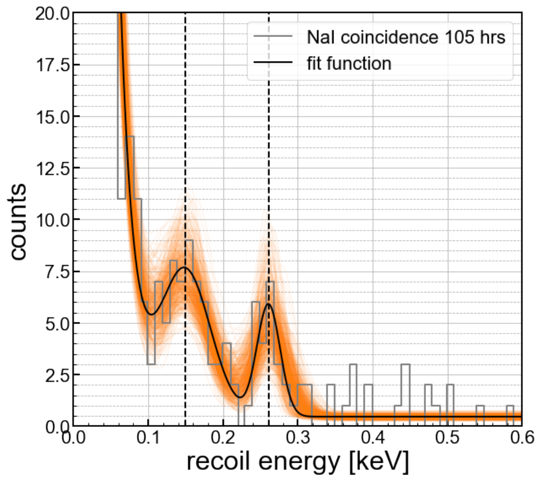

(a) Preliminary. A histogram of NaI coincident data in the low-energy region. The curve is a fitted combination of two Gaussians, representing the peaks, and an exponential decay that accounts for the spike at the lowest energy. The dashed lines highlight locations of the peaks, and the orange lines are from the MCMC walk, presented to highlight a rough uncertainty band.

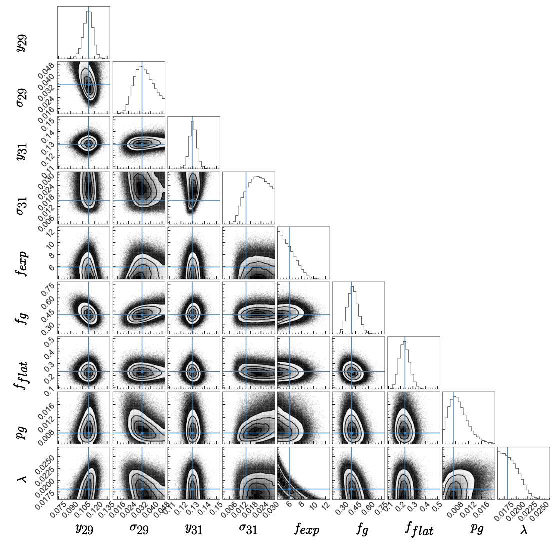

(b) Preliminary. A corner plot of the parameters of the fit in **(a)**. Each line in the two-variable plots represents a σ level.

Figure IV.8: Neutron Capture Lines

two points lies within the Sorensen band and not within the Lindhard band, while the upper of the two points lies within neither. What we do have, however, are strong, low-uncertainty measurements of the ionization yield for nuclear recoils in silicon. Figure IV.9 shows a comparison of our measurement to several other experiments [15, 39, 42, 77, 78]. This shows great promise for both further analysis and future experiments, as it gives us two very good reference points with low uncertainties that will allow us to draw strong comparisons of models once we have more data to do so.

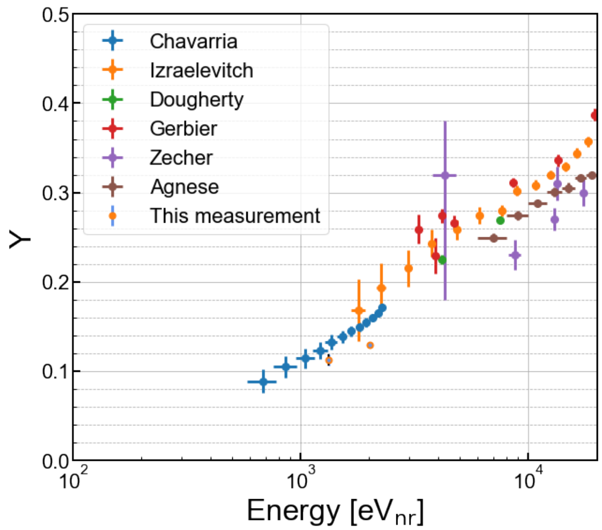

Figure IV.9: Comparison of Experiments

Preliminary. The yield measurements found in R76 alongside those of other experiments.

V. CONCLUSIONS

Understanding detector response in the lower-energy region will also help us to understand yield models for silicon as a whole, which still widely refer to and derive from Lindhard's [37] work despite recent measurements disagreeing with the model [15, 38]. I have laid out the skeleton of the method for analytically predicting the neutron capture spectrum of a solid state detector and simulated the spectrum where full analytical calculation proved impractical or impossible. There is room for improvement in this simulation, but the most prominent peaks remain in place even when statistically significant changes to its approximations are made. The relative positioning compared to the origin of two of the visible peaks in the simulation data, the single-step cascade peaks for ^{29}Si and ^{31}Si, resembles that of the two low-energy peaks in the NaI-coincident data. That is, the ratio of the ^{31}Si peak's energy to the ^{29}Si's in the simulation is similar to the ratio of the upper peak to the lower peak in the coincident data, leaving some "wiggle room" to account for the non-linearity of the yield. From this, we were able to measure the ionization yield of the Si detector at the associated energies. There is room to improve the yield measurements with further analysis, but they are direct measurements of neutron capture events with low uncertainty, a key step in improving our understanding of the low-energy behavior of solid state detectors.

1 Future Work

Throughout this work, I have identified several areas that could be expanded upon in future research. The purpose of this section is to provide a brief summary of these bases for future work, how they connect with each other, and how they could be of interest to other experiments. Chapter II discussed improvements to be made on the cascade modeling and simulation, as well as a variety of possible experiment backgrounds, which could use in-depth modeling. In Chapter III, I discussed the firmware used for the NaI array triggering and laid out remaining firmware improvements as well as DAQ integration. For the analysis in Chapter IV, I discussed the possibility of salting as a method for cut development, which could improve data analysis and calibration.

While I provided purely analytic solutions for the one- and two-step cascades, I did not do so the three-or-more-step cascades. There is no immediate reason these would not be possible to calculate. Such a calculation would require constructing a new formula for the probability distribution of a cascade step that occurs between two others, which would involve integration over multiple variables. Because the allowed angles of recoil following a de-excitation are restricted by conservation of momentum and energy, one would need to determine conditional probabilities of the angles and lifetimes in relation to one another. Once these adjustments have been made, it is likely possible to make a generic form that can be proved by induction. A purely analytic form for the PDF of a given cascade's deposit would be easily adaptable to construct a neutron capture spectrum for any solid state experiment, whether it is looking for captures or treats them as backgrounds. Ideally, we could expand the complete analytic solutions to produce electron-equivalent energy for various models, but first, we would need to prove that the yield models are invertible, which appears – but has not been formally proven – to be untrue.

Analyzing the impact of modifications on the treatment of the stopping power in our simulations indicated that, while the most crucial features of the spectrum remained intact, the spectrum did significantly change. Therefore, it is worth taking time to improve the stopping power used by the simulation. The first and simplest step is to allow user input for the stopping power; while this is not a necessary step, it allows anyone to tinker with the stopping power and see its impact on the data and could, in the future, be adapted to allow users to optionally input their own stopping power models. Inputting different stopping models would allow generating spectra for non-solid state detectors or comparing multiple theories, and the default would remain intact as it is the standard for solid-state detectors. More importantly, the implementation of user-input stopping power values is a simple task that would require tracking down the bug that causes simulation run times to become excessively long whenever the stopping power is modified in any way. Once this bug is fixed, we can more easily implement improved models because testing them would no longer take a prohibitively long time. Improved simulations fill the same role as an expanded analytical model, so they would be of use to any experiment where neutron capture events are present, either as events of interest or as backgrounds.

During R76, I did substantial work on the NaI array firmware, but more remains to be done. Various improvements to the triggering system could allow pileup rejection at the time of data collection, improving the efficiency of pileup cuts later on. Other optimizations to the triggering system, such as automatic baseline calculation, could improve the accuracy of set trigger levels. The biggest improvement to be made to the firmware, however, is direct integration with the DAQ, increasing the data we have on each event by telling us the energy detected by the array as well as providing some useful metadata. Directional data and knowledge of the associated gamma energy deposits would be of significant interest to future experiments, helping identify specific types of events. Section III.2.2.1 outlines the proposed firmware changes in greater detail.

We have a good understanding of backgrounds for R68 [52], which transfers well to the subset of R76 data that was taken without the NaI array for calibration purposes. However, the triggering mechanism for the NaI-coincident data is different entirely. In Section II.3, I laid out some potential background sources, but the precise nature of the backgrounds remains unclear. Constructing them would require a good understanding of what gammas are entering the detector area and what radiation is permeating, in particular, the boron shielding. Both of these subjects warrant further study, and there are already some background-only sections of the R76 data that we could analyze to this end. We can determine what radiation is being blocked by comparing subsets of the data with and without certain types of shielding. Similarly, we can determine, very roughly, the gamma energy spectrum of backgrounds in the NaI array by comparing background data with differing trigger parameters. However, such parameters are quite broad in the current data. The best way to analyze the NaI backgrounds would be to take data with the NaI array directly in a variety of shielding states and trigger levels. Currently, the majority of the NaI energy data is with a source in place, for the sake of calibration. NaI data that does not have a source is either very short or is from early in the firmware development process and would therefore not be representative of the current background levels. Therefore, as an alternative to trying to extract background information from the existing, limited NaI data, we could take a dedicated dataset for a background study after implementing the NaI-DAQ integration. Then, we would have correlated background data for the NaI and Si detectors, allowing us to see how the backgrounds of the NaI affect the backgrounds of the Si. A better understanding of backgrounds improves our ability to select cuts and therefore to accurately analyze this and future data.

There is room for improvement to the measurement of the resolution, which had a low variance but a very high χ^2. The measurement could be improved by either better fitting techniques or by finding a more targeted but still representative subset of the data. It may also be possible to select or build a better model, but I was not able to find one that sits between the over-heaviness of the Cauchy distribution's tails and the over-lightness of the Gaussian's. Improvements on other aspects of the analysis may also help improve this measurement, as it depends on the ^{241}Am calibration. Therefore, further cross-checking of the calibration with more data and improvements to the cuts using the salting techniques laid out in Section IV.1.1 could indirectly improve the noise width measurement. A better understanding of the noise width and distribution allows us to better tailor simulation results to the experiment.

The most crucial improvement to be made is to find more neutron capture events; doing so would allow meaningful comparison of yield models, which would further the field's understanding of solid state detectors and improve our capability of discerning new phenomena. The most obvious path to this is to analyze data with a lower NaI threshold – such data was collected, but is not in a state where it can be analyzed, as more pre-processing must be completed on the raw data first to convert it to an analyzable format. It is also currently not known how experiment-specific factors, such as detector temperature and voltage bias, affect the resulting yield. Studies analyzing the yield under various experimental changes would allow the community to better understand how we can combine data from various experiments.

A. SINGLE-STEP KINEMATICS

Given that the initial energy of a system with no kinetic energy, fields, or interaction between particles is determined exclusively by its masses, the initial energy of our system consisting of one neutron and is approximately $m_n + m_0$, where the subscript 0 represents the initial nucleus and the subscript n represents the lone neutron. Then, for an intermediate state of energy δ, conservation of energy dictates that

$$m_n + m_0 = T + p_\gamma + \delta + m_f, \tag{A.1}$$

where the subscript f represents the resulting nucleus. This can be arranged to

$$m_n + m_0 - m_f = T + p_\gamma + \delta, \tag{A.2}$$

and because $m_n + m_0 - m_f$ is the separation energy S_n,

$$S_n - \delta = T + p_\gamma. \tag{A.3}$$

Using the fact that p_γ is equal to the p_1, the nuclear recoil momentum, and the definition of kinetic energy,

$$S_n - \delta = \frac{p_1^2}{2m_f} + p_1 \tag{A.4}$$

which can be arranged to find p_1:

$$0 = \frac{p_1^2}{2m_f} + p_1 - (S_n - \delta) \therefore \tag{A.5}$$

$$p_1 = \frac{-1 \pm \sqrt{1^+ 4(S_n - \delta)/2m_f}}{2/2m_f} \tag{A.6}$$

$$= -m_f \pm \sqrt{m_f^2 + 2(S_n - \delta)m_f}, \tag{A.7}$$

which, since p_1 is a vector magnitude, must be positive, so

$$p_1 = \sqrt{m_f^2 + 2(S_n - \delta)m_f} - m_f. \tag{A.8}$$

Returning to the kinetic energy,

$$T = \frac{\left(\sqrt{m_f^2 + 2(S_n - \delta)m_f} - m_f\right)^2}{2m_f} \tag{A.9}$$

$$= \frac{m_f^2 + 2(S_n - \delta)m_f + m_f^2 - 2m_f\sqrt{m_f^2 + 2(S_n - \delta)m_f}}{2m_f} \tag{A.10}$$

$$= m_f + (S_n - \delta) - \sqrt{m_f^2 + 2S_n m_f} \tag{A.11}$$

$$= m_f\left(1 + \frac{S_n - \delta}{m_f} - \sqrt{1 + 2\frac{S_n - \delta}{m_f}}\right). \tag{A.12}$$

Because this is a single-step interaction, $\delta = 0$ can be dropped. Now, let $\epsilon = 2\frac{S_n}{m_f}$, which we know to be much less than 1. Then, by Taylor expansion,

$$\sqrt{1+\epsilon} \cong 1 + \frac{\epsilon}{2} - \frac{\epsilon^2}{8} \therefore \tag{A.13}$$

$$T \cong m_f\left(1 + \frac{S_n}{m_f} - \left(1 + \frac{S_n}{m_f} + \frac{S_n^2}{2m_f^2}\right)\right) \tag{A.14}$$

$$\cong \frac{S_n^2}{2m_f}. \tag{A.15}$$

□

B. FULL LIST OF SI CASCADES USED

The table below contains the probability of each cascade, using values adapted from nrCascadeSim [19]. The isotope listed is the isotope after neutron capture, in order to match the the energy levels and half-lives listed. A half-life entry in of the format w(E1) specifies that the half-life is unknown and the Weisskopf estimate for an electric dipole transition was used [43].

Isotope	Probability (%)	Energy levels (keV)	Half-lives (fs)
^{29}Si	62.5846%	4934.39	0.8
^{29}Si	10.7215%	6380.58; 4840.34	0.36; 3.5
^{29}Si	6.7778%	1273.37	291.0
^{29}Si	3.9706%	6380.58	0.36
^{29}Si	3.8974%	4934.39; 1273.37	0.84; 291
^{29}Si	2.0846%	...	...
^{30}Si	1.5106%	6744.1	14
^{31}Si	1.4053%	3532.9; 752.2	6.9; 530
^{30}Si	1.2015%	7507.8; 2235.3	24; 215
^{29}Si	0.8210%	6380.58; 4840.34; 1273.37	0.36; 3.5; 291
^{29}Si	0.6615%	6380.58; 1273.37	0.36; 291
^{30}Si	0.4326%	7507.80	24
^{30}Si	0.4006%	8163.2	w(E1)
^{29}Si	0.3972%	6380.58; 2425.86	0.36; 18.1
^{31}Si	0.3513%	5281.4; 752.2	w(E1); 530
^{30}Si	0.3412%	...	...
^{29}Si	0.2896%	4934.39; 3066.98; 1273.37	0.84; 33; 291
^{31}Si	0.2625%	4382.4; 752.2	w(E1); 530
^{31}Si	0.1376%	4382.40	0.1
^{29}Si	0.1353%	2425.86	18.1
^{29}Si	0.1326%	6380.58; 4934.39	0.36; 0.84
^{29}Si	0.1025%	4934.39; 2425.86	0.84; 18.1
^{30}Si	0.0918%	3498.50	58
^{29}Si	0.0826%	6908.52; 2028.05	0.1; 396
^{29}Si	0.0799%	4934.39; 3066.98; 2028.05	0.84; 33; 396

Isotope	Probability (%)	Energy levels (keV)	Half-lives (fs)
^{30}Si	0.0767%	8898.10	0.1
^{29}Si	0.0744%	7058.00	15.0
^{30}Si	0.0721%	6641.20	21
^{29}Si	0.0698%	6380.58; 2425.86; 1273.37	0.36; 18.1; 291
^{30}Si	0.0692%	9103.70; 6744.1	24; 14000
^{30}Si	0.0672%	3769.50	36
^{29}Si	0.0617%	6908.52; 1273.37	0.1; 291
^{29}Si	0.0545%	2028.05	396.0
^{29}Si	0.0449%	6908.52; 2425.86	0.1; 18.1
^{31}Si	0.0303%	...	...
^{29}Si	0.0303%	4840.34	3.5
^{29}Si	0.0294%	6908.52	0.1
^{29}Si	0.0291%	7996.90	0.1
^{30}Si	0.0278%	9308.10; 2235.3	24; 215000
^{29}Si	0.0238%	2425.86; 1273.37	18.1; 291
^{29}Si	0.0229%	7523.19; 2425.86	14.0; 18.1
^{29}Si	0.0203%	6712.90	0.1
^{30}Si	0.0201%	6914.80	24
^{30}Si	0.0200%	9792.30	0.1
^{30}Si	0.0196%	3787.70; 2235.3	8.3; 215000
^{30}Si	0.0193%	9619.70	0.1
^{29}Si	0.0189%	4934.39; 2028.05	0.84; 396
^{29}Si	0.0180%	4934.39; 2425.86; 1273.37	0.84; 18.1; 291
^{29}Si	0.0159%	6908.52; 3066.98; 1273.37	0.1; 33; 291
^{30}Si	0.0150%	9597.30; 2235.3	0.1; 215000
^{29}Si	0.0135%	3066.98; 1273.37	33.0; 291
^{30}Si	0.0128%	8953.40	0.1
^{30}Si	0.0121%	2235.30	215
^{29}Si	0.0102%	7058.00; 4934.39	15.0; 0.84
^{31}Si	0.0094%	5873.20	0.1
^{29}Si	0.0093%	7058.00; 2425.86	15.0; 18.1

Isotope	Probability (%)	Energy levels (keV)	Half-lives (fs)
^{29}Si	0.0082%	7058.00; 1273.37	15.0; 291
^{29}Si	0.0082%	6380.58; 4934.39; 1273.37	0.36; 0.84; 291
^{29}Si	0.0079%	6908.52; 2425.86; 1273.37	0.1; 18.1; 291
^{29}Si	0.0078%	7523.19	14.0
^{30}Si	0.0057%	10202.3	0.1
^{29}Si	0.0056%	6908.52; 2028.05; 1273.37	0.1; 396; 291
^{29}Si	0.0054%	4934.39; 3066.98; 2028.05; 1273.37	0.84; 33; 396; 291
^{30}Si	0.0049%	5614.00; 2235.3	21; 215000
^{30}Si	0.0049%	8939.00; 2235.3	0.1; 215000
^{29}Si	0.0046%	4840.34; 1273.37	3.5; 291
^{30}Si	0.0046%	8156.10; 2235.3	0.1; 215000
^{29}Si	0.0044%	6908.52; 3066.98; 2028.05	0.1; 33; 396
^{29}Si	0.0040%	7523.19; 2425.86; 1273.37	14.0; 18.1; 291
^{29}Si	0.0037%	3066.98; 2028.05	33.0; 396
^{29}Si	0.0037%	2028.05; 1273.37	396.0; 291
^{30}Si	0.0034%	7667.40	14
^{31}Si	0.0025%	752.200	0.53
^{29}Si	0.0022%	6380.58; 2425.86; 2028.05	0.36; 18.1; 396
^{31}Si	0.0021%	5957.90	0.1
^{30}Si	0.0020%	10275.5; 2235.3	0.1; 215000
^{29}Si	0.0016%	7058.00; 2425.86; 1273.37	15.0; 18.1; 291
^{31}Si	0.0014%	2788.00	14
^{29}Si	0.0013%	4934.39; 2028.05; 1273.37	0.84; 396; 291
^{29}Si	0.0012%	6380.58; 4934.39; 3066.98; 1273.37	0.36; 0.84; 33; 291
^{31}Si	0.0010%	1694.90	0.57
^{29}Si	0.0007%	2425.86; 2028.05	18.1; 396
^{29}Si	0.0007%	7058.00; 4934.39; 1273.37	15.0; 0.84; 291
^{30}Si	0.0007%	8104.80; 2235.3	24; 215000
^{31}Si	0.0007%	2316.90	38
^{29}Si	0.0006%	4934.39; 2425.86; 2028.05	0.84; 18.1; 396
^{29}Si	0.0004%	6380.58; 4934.39; 2425.86	0.36; 0.84; 18.1

Isotope	Probability (%)	Energy levels (keV)	Half-lives (fs)
^{29}Si	0.0003%	6380.58; 4934.39; 3066.98; 2028.05	0.36; 0.84; 33; 396
^{29}Si	0.0003%	6908.52; 3066.98; 2028.05; 1273.37	0.1; 33; 396; 291
^{29}Si	0.0003%	3066.98; 2028.05; 1273.37	33.0; 396; 291
^{29}Si	0.0003%	6908.52; 2425.86; 2028.05	0.1; 18.1; 396
^{29}Si	0.0002%	6380.58; 2425.86; 2028.05; 1273.37	0.36; 18.1; 396; 291
^{29}Si	0.0001%	7523.19; 2425.86; 2028.05	14.0; 18.1; 396
^{29}Si	0.0001%	7058.00; 4934.39; 3066.98; 1273.37	15.0; 0.84; 33; 291
^{29}Si	0.0001%	6380.58; 4934.39; 2028.05	0.36; 0.84; 396
^{29}Si	0.0001%	6380.58; 4934.39; 2425.86; 1273.37	0.36; 0.84; 18.1; 291
^{29}Si	0.0001%	7058.00; 2425.86; 2028.05	15.0; 18.1; 396
^{29}Si	0.0001%	2425.86; 2028.05; 1273.37	18.1; 396; 291
^{29}Si	0.0001%	4934.39; 2425.86; 2028.05; 1273.37	0.84; 18.1; 396; 291
^{29}Si	0.0001%	7058.00; 4934.39; 2425.86	15.0; 0.84; 18.1
^{29}Si	0.0001%	7058.00; 4934.39; 3066.98; 2028.05	15.0; 0.84; 33; 396

C. DISTRIBUTION OF $\cos\beta$

If we assume that the distribution of a solid is angle is uniform over the range $[0, 4\pi]$, we are equivalently assuming that the spatial distribution is uniform over the surface of a sphere with arbitrary radius R:

$$g_{x,y,z}(x,y,z) = \begin{cases} (4\pi R^2)^{-1} & x^2+y^2+z^2 = R^2 \\ 0 & x^2+y^2+z^2 \neq R^2 \end{cases}. \tag{C.1}$$

Switching to polar coordinates,

$$\oiint g_{\beta,\phi|r=R}(\beta,\phi) = \int_{\phi=0}^{2\pi}\int_{\beta=0}^{\pi}(4\pi R^2)^{-1}R^2\sin\beta d\beta d\phi. \tag{C.2}$$

Applying a u-substitution where

$$u = -\cos\beta, \therefore \tag{C.3}$$

$$du = \sin\beta d\beta \tag{C.4}$$

gives

$$\oiint_{r=R} g_{\beta,\phi|r}(\beta,\phi) = (4\pi)^{-1}\int_0^{2\pi}d\phi\int_{-1}^{1}du, \tag{C.5}$$

which indicates a uniform distribution over $u = \cos\beta$ for the range $[-1, 1]$. With two or fewer steps in a cascade, we can choose our $x-y$ plane such that $\phi = 0$; therefore, a uniform distribution over the surface of a sphere for a two-or-fewer-step cascade is equivalent to the distribution

$$f_{\cos\beta}(\cos\beta) = \begin{cases} 0 & \cos\beta < -1 \\ \frac{1}{2} & -1 \leq \cos\beta \leq 1 \\ 0 & 1 < \cos\beta. \end{cases} \tag{C.6}$$

$\square$

References

[1] V. Rubin, W. F. Jr., and N. Thonnard, Astrophysical Journal **225**, L107 (1978).

[2] A. Villano et al., Journal of Open Source Education **6**, 184 (2023).

[3] C. Kouvaris and J. Pradler, Phys. Rev. Lett. **118**, 031803 (2017).

[4] M. J. Dolan, F. Kahlhoefer, and C. McCabe, Phys. Rev. Lett. **121**, 101801 (2018).

[5] A. R. L. Kavner and I. Jovanovic, Phys. Rev. D **110**, 083043 (2024).

[6] I. Alkhatib et al. (SuperCDMS Collaboration), Phys. Rev. Lett. **127**, 061801 (2021).

[7] R. Agnese et al. (SuperCDMS Collaboration), Phys. Rev. D **99**, 062001 (2019).

[8] R. Agnese et al. (SuperCDMS Collaboration), Phys. Rev. Lett. **120**, 061802 (2018).

[9] M. Fabbrichesi, E. Gabrielli, and G. Lanfranchi, *The physics of the dark photon, A primer*, SpringerBriefs in Physics (Springer Cham, Nov. 2020).

[10] P. Sikivie, Phys. Rev. Lett. **51**, 1415 (1983).

[11] L. Roszkowski, E. M. Sessolo, and S. Trojanowski, Reports on Progress in Physics **81**, 066201 (2018).

[12] M. F. Albakry et al., Phys. Rev. D **105**, 122002 (2022).

[13] M. F. Albakry et al. (SuperCDMS Collaboration), Phys. Rev. Lett. **131**, 091801 (2023).

[14] J. I. Collar, A. R. L. Kavner, and C. M. Lewis, Phys. Rev. D **103**, 122003 (2021).

[15] R. Agnese et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **905**, 71 (2018).

[16] K. W. Jones and H. W. Kraner, Phys. Rev. A **11**, 1347 (1975).

[17] C. S. Amarasinghe et al., Phys. Rev. D **106**, 032007 (2022).

[18] R. B. Firestone et al., AIP Conference Proceedings **819**, 138 (2006).

[19] A. Villano, K. Harris, and S. Brown, Journal of Open Source Software **7**, 3993 (2022).

[20] J. Allison et al., IEEE Transactions on Nuclear Science **53**, 270 (2006).

[21] F. Bečvář, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **417**, 434 (1998).

[22] H. Almazán et al., The European Physical Journal A **59**, https://doi.org/10.1140/epja/s10050-023-00977-x (2023).

[23] G. Agnolet et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **853**, 53 (2017).

[24] J. Dirk, M. Nelson, J. Ziegler, A. Thompson, and T. Zabel, IEEE Transactions on Nuclear Science **50**, 2060 (2003).

[25] A. J. Biffl, A. Gevorgian, K. Harris, and A. N. Villano, Phys. Rev. D **107**, 092011 (2023).

[26] R. Agnese et al. (SuperCDMS Collaboration), Phys. Rev. D **95**, 082002 (2017).

[27] K. Harris, A. Gevorgian, A. J. Biffl, and A. N. Villano, Phys. Rev. D **107**, 076026 (2023).

[28] L. Thulliez et al., Journal of Instrumentation **16**, P07032 (2021).

[29] H. Lee et al., Journal of High Energy Physics, 10.1007/JHEP08(2015)093 (2015).

[30] W.-Z. Wei, J. Liu, and D.-M. Mei, Journal of Instrumentation **11**, P07008 (2016).

[31] B. S. Neganov and V. N. Trofimov, Otkryt. Izobret. **146**, 215 (1985).

[32] R. K. Romani et al., Applied Physics Letters **112**, 043501 (2018).

[33] R. Agnese et al. (SuperCDMS Collaboration), Phys. Rev. D **99**, 062001 (2019).

[34] D. W. Amaral et al., Phys. Rev. D **102**, 091101 (2020).

[35] H. Neog et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **1033**, 166707 (2022).

[36] A. H. Abdelhameed et al. (CRESST Collaboration), Phys. Rev. D **100**, 102002 (2019).

[37] J. Lindhard, V. Nielsen, M. Scharff, and P. V. Thomsen, Kgl. Danske Videnskab., Selskab. Mat. Fys. Medd. **33** (1963).

[38] J. I. Collar, A. R. L. Kavner, and C. M. Lewis, Phys. Rev. D **103**, 122003 (2021).

[39] F. Izraelevitch et al., Journal of Instrumentation **12**, P06014 (2017).

[40] A. E. Chavarria et al., Phys. Rev. D **94**, 082007 (2016).

[41] G. Gerbier et al., Phys. Rev. D **42**, 3211 (1990).

[42] B. J. Scholz, A. E. Chavarria, J. I. Collar, P. Privitera, and A. E. Robinson, Phys. Rev. D **94**, 122003 (2016).

[43] V. F. Weisskopf, Phys. Rev. **83**, 1073 (1951).

[44] S. Mughabghab, International Atomic Energy Agency (2003).

[45] K. J. R. Rosman and P. D. P. Taylor, Pure and Applied Chemistry **70**, 217 (1998).

[46] N. N. D. Center, *Nudat database*, version 3.

[47] S. Raman, E. T. Jurney, J. W. Starner, and J. E. Lynn, Phys. Rev. C **46**, 972 (1992).

[48] J. F. Ziegler, M. Ziegler, and J. Biersack, Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms **268**, 19th International Conference on Ion Beam Analysis, 1818 (2010).

[49] J. Ziegler, J. Biersack, and M. Ziegler, *Srim - the stopping and range of ions in matter* (James Ziegler, Jan. 2008).

[50] P. Sorensen, Phys. Rev. D **91**, 083509 (2015).

[51] W. R. Inc., *Mathematica, Version 14.2*, Champaign, IL, 2024.

[52] A. N. Villano et al., Phys. Rev. D **105**, 083014 (2022).

[53] G. Marsaglia, W. W. Tsang, and J. Wang, Journal of Statistical Software **8**, 1 (2003).

[54] P. Virtanen et al., Nature Methods **17**, 261 (2020).

[55] G. Miloshevsky and A. Hassanein, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **749**, 47 (2014).

[56] A. Negret et al., Phys. Rev. C **88**, 034604 (2013).

[57] K. V. Berghaus, R. Essig, Y. Hochberg, Y. Shoji, and M. Sholapurkar, Phys. Rev. D **106**, 023026 (2022).

[58] C.-K. Qiao et al., Journal of Physics G: Nuclear and Particle Physics **47**, 045202 (2020).

[59] L. Koester et al., Phys. Rev. C **51**, 3363 (1995).

[60] L. Koester, W. Waschkowski, and A. Klüver, Physica B+C **137**, 282 (1986).

[61] G. E. Bacon and K. Lonsdale, Reports on Progress in Physics **16**, 1 (1953).

[62] F. Bloch, Physical Review Letters to the Editor, 259 (1936).

[63] A. Baldini et al., Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **570**, 561 (2007).

[64] CAEN and Nuclear Instruments, *DT5560SE*.

[65] *Maximum integrated data acquisition system*, Oct. 2023.

[66] M. Dong et al., Radiation Physics and Chemistry **188**, 109601 (2021).

[67] CAEN and Nuclear Instruments, *Sci-compiler*, version 2024.6.5.1, June 2024.

[68] K. Mickelson, A. N. Villano, L. Hellman, S. Hansen, and A. Sharbaugh, *R76 firmware*, version v0.4, Sept. 2024.

[69] N. G. Mast, “Cryogenic detector development for SuperCDMS”, PhD thesis (UNIVERSITY OF MINNESOTA, 2020).

[70] J. Yoon and D. Jarrett, in Advances in neural information processing systems 32: 33rd conference on neural information processing systems (NeurIPS 2019), Advances in Neural Information Processing Systems 32 (2019).

[71] Z. Kromer, A. Roberts, D. Pradhan, and K. Mickelson, *Timegan-static*, version v0.1.9, 2024.

[72] I. Csiszar, The Annals of Probability **3**, 146 (1975).

[73] G. E. Karniadakis et al., Nature Reviews Physics **3**, 10.1038/s42254-021-00314-5 (2021).

[74] C. Meng, S. Seo, D. Cao, S. Griesemer, and Y. Liu, *When physics meets machine learning: a survey of physics-informed machine learning*, 2022.

[75] Z. Hao et al., *Physics-informed machine learning: a survey on problems, methods and applications*, 2023.

[76] D. Foreman-Mackey, D. W. Hogg, D. Lang, and J. Goodman, PASP **125**, 306 (2013).

[77] B. L. Dougherty, Phys. Rev. A **45**, 2104 (1992).

[78] G. Gerbier et al., Phys. Rev. D **42**, 3211 (1990).

Glossary

baseline is a value describing the no-signal region of data; the base from which the signal raises and lowers. 36

cumulative distribution function refers to a function $F(x)$ that represents the probability of all possibilities up until x; that is, for a probability distribution function $f(x)$ with a domain whose infimum is m, $F(x) = \int_m^x f(u)du$. 24

DCRC stands for Detector Control and Readout Card, which is custom SuperCDMS electronics hardware that allows communication between connected silicon detectors and the collaboration's data acquisition software. 31, 38

energy is a scalar quantity with many possible types, with each type having a different formal definition. Energy is conserved, meaning that, in a closed system, the total energy can never change. This conservation applies only to the total energy of a system, not to any one subtype; energy frequently changes types following and even during interactions. 1

gamma describes a type of radiation consisting of one or more photons, which are particles of electromagnetic radiation, or light. A gamma particle, or photon, is emitted in several types of interactions and spontaneous events. Despite having energy and momentum, it has no mass. 1, 68

infimum of a subset S of set P is the greatest element of P such that all elements of S are greater than or equal to it. In this work we can simplify this to say that the infimum of a set S is the greatest real number that is less than or equal to all elements of S. The infimum differs from the minimum in that it does not have to be an element of S. 68

momentum is a vector quantity describing the motion of an object. It is a typically defined as the mass of the object multiplied by its velocity, representing, in a rough sense, how difficult it is to stop the object. However, momentum can be determined in other ways, and gamma particles still have momentum despite lacking mass. Momentum is conserved, meaning that, in a closed system, the total momentum can never change. 1

phonon describes a packet of energy – essentially heat or sound – that moves through a substance. It is a quasiparticle, exhibiting particle-like behaviors but ultimately being a behavior of the solid's structure. 4, 7

photon See gamma. 68

resolution is a value representing how easily similar values can be distinguished from one another. Resolution can be represented in many ways; in this work, a resolution is a value with the same units as the measurements and represents a width of error. Therefore, larger resolution values mean the data is less clear.. 7

supremum of a subset S of set P is the least element of P such that all elements of S are less than or equal to it. In this work, we can simplify this to say that the supremum of a set S is the least real number such that is greater than or equal to all elements of S. The supremum differs from the maximum in that it does not have to be an element of S. 24

threshold is the lowest value which can be read by a detector. Events that occur below this value but would otherwise be detected cannot be measured because the hardware has no way to meaningfully respond to them.. 7

yield refers to the difference in the deposited and reported energies in a detector. The reported energy is the product of the deposited energy and the yield.. 2