

SIMULATION OF THE CHARGE MEASUREMENTS FOR THE SUPERCDMS
SOUDAN EXPERIMENT

A Dissertation

by

JORGE DANIEL MORALES MENDOZA

Submitted to the Office of Graduate and Professional Studies of
Texas A&M University
in partial fulfillment of the requirements for the degree of
DOCTOR OF PHILOSOPHY

Chair of Committee,	David Toback
Committee Members,	Bhaskar Dutta
	Nader Mirabolfathi
	Mahmoud El-Halwagi
Head of Department,	Grigory Rogachev

May 2020

Major Subject: Physics

Copyright 2020 Jorge Daniel Morales Mendoza

ABSTRACT

Astronomical observations lead us to believe that dark matter constitutes the majority of the mass of the Universe, yet it has never been directly observed experimentally. The Super Cryogenic Dark Matter Search Experiment (SuperCDMS) aims to discover the particle nature of dark matter, by developing and using some of the most sensitive detectors ever built. The SuperCDMS detectors are capable of rejecting known particles to search for dark matter interactions, but the discrimination is a complicated analysis task. Previous analyses with SuperCDMS set world leading sensitivity limits to dark matter interactions, but the analysis techniques have not been based on simulations. Now we have developed an efficient simulations infrastructure that will help better understand the detectors, develop a better background model, and optimize the next generation dark matter searches. In this thesis we present the fully-functional simulations infrastructure to produce high quality samples as well as results for the charge system readout simulations, and comparisons to expectations from data. We show that our simulations well-reproduce dominant features of data, bringing SuperCDMS closer to a robust simulation-based analysis method, much needed for the upcoming SuperCDMS dark matter searches.

ACKNOWLEDGMENTS

It goes without saying that this work couldn't have been done without the professional and personal support, hard work, and help from many others. I value this work as representative for the biggest milestone in my long sought ambition to become a scientist. For this reason, I hereby thank deeply all the people with whom I have been fortunate to cross paths, worked closely, and received support. While I wish to name them all, I limit my writing to those mentioned here, and extend gratitude to the rest.

First of all I wish to thank the people who have forged me into the person I am today. Starting with my wife Ana Arnaus, in particular for her unconditional love and support in this and many other projects making me a better person, and for being the center of our beloved family and kids (Jorge and Lucia). I thank my mother Ma. Teresa Mendoza and father Jorge Morales, for planting the seeds of what I am and for teaching me to go further and do better. Also thanks to my brother Rodrigo Morales for his care, advice, and good conversations.

Turning to my professional evolution, I have been lucky to work with David Toback, my advisor, mentor, friend, and example of what 'being a professor' is. I am grateful not only for his support and guidance throughout my professional development (which persisted even though some unconventional circumstances), but particularly for him honestly caring about passing his experience and knowledge for becoming a world-class scientist and professor. I am both proud and honored to have learned this much from him. I then thank the Toback Group members, our A&M research group, especially Jon S. Wilson (former postdoc) for his aid, continuous advice, and always interesting conversation; and Ziqing Hong (former grad and office mate) for his always fun but accurate prophetic advice (as he was living what I was about to live).

I wish to express gratitude to the members of the SuperCDMS Collaboration for the welcoming and warm atmosphere in our quest for developing and using the simulations for

analyses. I will list some people who have had special roles in my development and projects to thank for their particular support and contributions. I start with Ray Bunker for his guidance and interest in finding a good path since the beginning of my work with the collaboration. Dan Bauer has provided support, in particular for the University Research Award (URA) Fellowship at Fermilab to work on SuperCDMS electronics. For the URA project I worked closely with Sten Hansen from whom I learned a lot about readout electronics, FPGAs and other such devices, while having a great time and fun conversations. Scott Oser has also given me good feedback and advice in particular about the readout electronics projects and needs. Bruno Serfass helped me get set up with many aspects of the SuperCDMS software, in particular the Reconstruction package. I also thank Matt Pyle who taught me a lot about the SuperCDMS detectors, operation, readout, and noise for my first projects in SuperCDMS where we actually got to modify the operating conditions of the SuperCDMS Soudan experiment during my first shift at the Soudan mine.

Finally, I turn to the Texas A&M University crew, from professors to teammates and friends, I am thankful for being part of it, and for the always positive spirit carried by. I begin expressing my thanks to the Brazos Cluster team, in particular Guy Almes who always had a genuine interest in supporting my development as a scientist, and Steve Johnson who has also been fully supportive in our needs for our projects. I then turn to thanking my dissertation evaluation committee: Bhaskar Dutta, Nader Mirabolfathi, and Mahmoud El-Halwagi. I do thank in particular Bhaskar Dutta from whom even outside the lecture hall I've had an open door to learn a lot about QFT and cosmology, from major aspects to subtleties. My last words of appreciation in this work go to Sriteja Upadhyayula, who ever since I arrived to campus, has become a good friend, with fun conversations, full support, and always there for a bike ride.

CONTRIBUTORS AND FUNDING SOURCES

Contributors

This work was supervised by a dissertation committee consisting of Professors David Toback [advisor] and Bhaskar Dutta members of the Department of Physics and Astronomy and the Mitchell Institute for Fundamental Physics and Astronomy, Professor Nader Mirabol-fathi of the Department of Physics and Astronomy, and Professor Mahmoud El-Halwagi of the Department of Chemical Engineering.

All work for the dissertation was completed by the student, under the advisement of Prof. David Toback of the Department of Physics and Astronomy, in collaboration with the members of the SuperCDMS Experiment.

Funding Sources

This work was made possible in its majority by the U.S. Department of Energy (DOE), through a research grant for the SuperCDMS Experiment, by support from the Mitchell Institute for Fundamental Physics and Astronomy for work with the Texas A&M Brazos Cluster, by a Teaching Assistantship from the Texas A&M University, and in part by a fellowship from the Consejo Nacional de Ciencia y Tecnología (CONACYT), México.

The contents of this work are solely the responsibility of the author and do not necessarily represent the official views of the U.S. DOE or CONACYT.

TABLE OF CONTENTS

	Page
ABSTRACT	ii
ACKNOWLEDGMENTS	iii
CONTRIBUTORS AND FUNDING SOURCES	v
TABLE OF CONTENTS	vi
LIST OF FIGURES	x
LIST OF TABLES	xxix
1. INTRODUCTION	1
1.1 Astronomical Evidence for the Existence and Particle Nature of Dark Matter .	2
1.2 Cosmological Predictions and the Properties of Dark Matter	7
1.2.1 The Laws of Physics Used in Cosmology	7
1.2.2 A History of the Universe with Dark Matter and Its Use to Make Quantitative Predictions	9
1.2.3 Data and Cosmological Models Make Predictions About Dark Matter: The WIMP Hypothesis	12
1.3 Search Strategies and Current Limits on WIMPs	14
1.3.1 Methods For Observing Dark Matter Interactions and Overview of Direct Detection Experiments	15
1.3.2 Direct Detection WIMP-Nucleon Interaction Rate Model	17
1.3.3 Overview of Current WIMP-Nucleon Cross Section Sensitivity Limits by SuperCDMS and Others	19
1.4 Overview and Outline of this Thesis	21
2. THE SUPERCDMS SOUDAN EXPERIMENT	24
2.1 Overview of Background Sources and How They Interact in the Detectors	25
2.2 Overview of the Experimental Apparatus and Detectors	29
2.2.1 Apparatus Description	30
2.2.2 Detector and Sensor Description	32
2.3 Description of How the Backgrounds Are Mitigated by the Experiment	35
2.3.1 Cosmogenic Sources Backgrounds	36
2.3.2 Contaminants Near the Detectors	37
2.3.3 Energy Mismeasurement	39

3.	DETECTOR RESPONSE TO PARTICLE INTERACTIONS.....	42
3.1	Energy Deposition in the Crystal.....	43
3.1.1	From Recoil Energy to Primary Energy	46
3.1.2	From Primary Interaction Energy to the Number of Phonons and Electron-Hole Pairs Produced	48
3.1.3	Phonon Energy Released by High-Speed Charges Interactions in the Crystal (Secondary Energy) and the Total Energy Expected Values	49
3.2	Motion of Charges and Phonons in the Crystal	51
3.2.1	Charges Pulled by the Electrodes But Subject to the Constraints on the Motion From the Ge Crystal	52
3.2.2	Electron-Hole Scattering and Phonon Release Due to High-Speed Mo- tion of Charges	55
3.2.3	Motion of Phonons in the Crystal	56
3.3	Charge Readout Response to Energy Deposition.....	57
3.3.1	Charge Collection by the Electrodes	57
3.3.2	Charge Waveform Creation	59
3.4	Phonon Readout and Collection by TES Sensors	62
4.	EVENT RECONSTRUCTION: OPTIMAL FILTER.....	64
4.1	Optimal Filter Implementation Overview	65
4.2	Signal Template Construction	68
4.3	Charge Optimal Filter	70
4.4	Charge Energy Observables	73
5.	CALIBRATIONS SOURCES AND DETECTOR RESPONSE, WELL-BEHAVED SELECTION CRITERIA, PREVIOUS WIMP SEARCH METHODOLOGIES AND RESULTS WITH AN EYE TO THE FUTURE WITH SIMULATIONS	77
5.1	The Calibration Sources: ^{133}Ba and ^{252}Cf	78
5.2	Calibration and Noise Samples and Detector Response.....	81
5.2.1	The ^{133}Ba Sample and Its Use in Electron Recoil Response and Energy Scale Calibration	82
5.2.2	The ^{252}Cf Sample and Its Use in Nuclear Recoil Response	83
5.2.3	The InRunRandoms (No Interactions Sample) and Its Use In Measur- ing Detector Readout Noise.....	84
5.3	Well-Behaved Event Selection Criteria: Low Threshold (LT) Fiducial Require- ments	85
5.4	Outline of WIMP Searches with SuperCDMS Soudan	86
5.5	Previous SuperCDMS Analyses.....	88
5.6	Analyses Improvements with Simulations	90
6.	DETECTOR NOISE SAMPLE, READOUT MISMEASUREMENT EVENTS AND THEIR REJECTION AND CORRECTION	92

6.1	Overview of Readout Mismeasurement Sources	93
6.2	Characterization of the InRunRandoms Sample: Noise Events Contaminated by Pileups and Blips	97
6.2.1	InRunRandom Sample Creation and Energy Distribution	97
6.2.2	Pileup Event Removal	99
6.2.3	Single-Detector Blips and TowerBlips Description	101
6.3	Finding and Measuring Blips and TowerBlips	103
6.3.1	The BlipFinder Tool	103
6.3.2	The TowerBlipFinder Tool	107
6.4	Correcting for TowerBlips	110
7.	THE SUPERCDMS SIMULATIONS	115
7.1	Overview of Simulations and the Simulations Infrastructure	117
7.2	Simulation of the Sources and Recoils in the Detector: SourceSim	120
7.3	Simulation of Detector Physics and Readout: DetectorSim	122
7.3.1	CrystalSim: From Recoil Energy to Creation and Propagation of Charges and Phonons	123
7.3.2	Simulation of the Charge Readout: FETSim	126
7.3.3	Simulation of the Phonon Readout: TESSim	127
7.3.4	Simulation of Data Acquisition and Noise: DAQSim	127
7.4	Event Reconstruction with CDMSBats	128
7.5	Simulated Samples	129
7.5.1	Pure Noise Simulation	129
7.5.2	Barium Calibration Simulation	131
7.5.3	Enhanced WIMP Simulation	135
8.	COMPARISON OF NOISESIM TO DATA AND BASELINE ENERGY RESOLU- TION	137
8.1	NoiseSim Expectations and Energy Resolution Measurement	137
8.2	Estimation of BlipCleaned PSDs for Use in the Simulations	138
8.3	Comparison of Data and Simulations, Final Resolution Results	140
9.	RESULTS FOR REAL ^{133}Ba DATA AND SIMULATIONS	144
9.1	Expectations, Goals, Methods, and Overview of the Results	145
9.1.1	Goals for the Simulations Comparisons and Measurements	145
9.1.2	Preview of the Results	147
9.2	Data Comparison and Simulation Predicted Accuracy and Resolution Mea- surement Methods	148
9.2.1	Overview of the Methods	149
9.2.2	Variables Used for the Comparisons	149
9.3	Using the Single-Scatter Sample for Finding the Simulation Detector Region of Well-Measured Events (SimFiducial)	150
9.4	Study of LT Fiducial Variables in ^{133}Ba Data and Simulations	154

9.4.1	Full Energy Range Comparison of Inner Energy Estimates	155
9.4.2	Events with No Energy Deposition in One Side.....	156
9.4.3	Events with One Side Well-Measured and the No-Outer Energy Requirement	157
9.4.4	Results After the No-Outer Energy Requirement and Before the Symmetric Constraint	159
9.4.5	What the Simulation Tells Us About Poorly Measured Events that Pass the LT Fiducial Requirements	162
9.5	Comparison of Simulated and Real ^{133}Ba Data.....	165
9.5.1	Comparison of Simulation and Data for Events in the 356 keV Reconstructed Peak	165
9.5.2	Full Energy Range Comparison of Simulation and Data	170
10.	CHARGE SIMULATION RESULTS FOR WELL-MEASURED EVENTS, INTRINSIC DETECTOR RESOLUTION, CAUSES OF MISMEASUREMENT, AND NEXT STEPS FOR DEVELOPMENT	172
10.1	Overview of the Predicted Accuracy and Resolution Measurement Methods ...	173
10.2	Charge Collection Explains SimFiducial Volume and Linearity in Measurements	175
10.3	Predicted Accuracy and Resolution of Electron Recoil Events	177
10.4	Predicted Accuracy and Resolution of Nuclear Recoil Events	180
10.5	Using the Simulation to Study Causes of Energy Mismeasurement.....	181
10.5.1	Lowest Energies Limitations: Noise and Digitization	182
10.5.2	LT Fiducial Comparison	184
10.5.3	A Study of the Qimean Measurement Drop as a Function of Energy: Z-Position Dependence	185
10.6	Next Steps for the Simulations Development	187
10.7	Summary	189
11.	CONCLUSIONS	191
11.1	Overview of Optimizing a Dark Matter Search with the Simulations and Clear Examples to Follow Up From This Work	191
11.2	Final Conclusions	193
	REFERENCES	196
	APPENDIX A EXCESS OF ZIP4 BLIPS AT $t=0$	204
	APPENDIX B TOWERBLIP RATES	205
	APPENDIX C TOWERBLIP COINCIDENCE RATES	207
	APPENDIX D ZIP5+6 BLIPFINDER PREDICTED ENERGY FITS	211

LIST OF FIGURES

FIGURE	Page
1.1 Image of the Coma Cluster, where the position and brightness of the light shows the location and density of stars, and gives a measure of the atomic mass in the galaxy cluster reprinted from Ref. [13]. It provided the first evidence for the existence of dark matter as the speeds of the galaxies exceeded the escape velocity expected from the observed amount of mass. This suggested a large, non-visible source of matter/mass.	3
1.2 Left: Image of a galaxy with luminous matter and the calculated dark matter halo engulfing the galaxy (in blue) where the amount of light shows the mass distribution from the stars in the galaxy, reprinted from Ref. [14]. Right: Recent observations of the speeds of the stars around the M33 Galaxy as function of distance from the galaxy center, showing data (points) along with a combined rotation model (solid line) from mass contributions from a number of different sources, reprinted from Ref. [4]. The dark matter halo contribution is drawn with a dashed-dotted line, the contributions from observable stellar sources are in short dashed line, and the surrounding gas in long dashed line. The dark matter halo model well describes the observed rotation data, which suggests that there is a non-visible substance with mass that is distributed as a halo in the galaxy.	4
1.3 Left: Diagram of gravitational lensing, showing the view of two galaxies at different distances from the Earth, reprinted from Ref. [18]. The ‘far’ galaxy emits light that passes through, or around, the ‘near’ galaxy and which is gravitationally lensed, bending the path of the light towards us. Right: An example gravitational lensing observed in nature reprinted from Ref. [15], where the blue horseshoe shape is gravitationally lensed light from a distant galaxy which shows an Einstein Ring created by the near galaxy in the path bending the light towards us as observers. In the center, the orange bright source is a massive near galaxy (called LRG). There is greater gravitational bending than expected from mass of light emitting sources, evidence supporting that there is a non-visible substance that interacts gravitationally.	5

1.4	A set of images of the ‘Bullet Cluster’ which is a colliding cluster of galaxies that provide evidence that dark matter is a particle [5]. Top-left: Calculated contours of mass contributions as measured using gravitational lensing techniques, reprinted from Ref. [5]. Top-Right: Overlay of the mass contours with X-ray picture showing the amount of light from X-rays indicating the amount (location and density) of interacting atomic mass, reprinted from Ref. [5]. Bottom: Three overlaid images; mass distribution (from gravitational lensing) in blue, mass distribution of X-rays showing the area of atomic interactions (in pink), and the optical image shows galaxies in orange and white, reprinted from Ref. [20]. This data supports the hypothesis that the dark matter is separated from the atomic matter of the system, due to it being a weakly interactive massive particle, in agreement with observations when fitting with a particle-based dark matter halo model within each galaxy cluster.	6
1.5	A diagram showing the known particles of the Standard Model of Particle Physics, reprinted from Ref. [29]. It has 3 types of particles: Quarks (in purple), leptons (in green), and force carriers (in orange and yellow). All the matter particles are grouped into three generations of matter, Generation I is the most stable version, involving the most common particles: electrons, up and down quarks (which make protons and neutrons), and electron-neutrinos. Force carriers are photons (mediating electromagnetic force), gluons (mediating strong-force interactions), and Z and W bosons (mediating weak interactions). None of these particles have the properties of a dark matter particle.	9
1.6	Data from Cosmic Microwave Background (CMB) measurements and their interpretation in the Λ CDM model with best-fit values of the cosmological parameters, figure reprinted from Ref. [31]. Center bottom is a picture of the CMB temperature variations heat-map, where blue and red represent colder and hotter regions respectively, by 1 part in 10^4 relative to the average temperature of the Universe of 2.7260 ± 0.0013 Kelvin [32]. Center top is the analyzed data of temperature variation as function of angular size of patches in the CMB along with the best-fit from the Λ CDM model. On the left is the best-fit parametrization of the Λ CDM model that indicates that dark matter is about 22% of the mass-energy of the Universe and thus about $\sim 85\%$ of the mass in the Universe [7].	13

1.7	On the y-axis, the number density of dark matter particles as a function of temperature (proxy for time) of the early Universe on the x-axis, for a 100 GeV dark matter particle, figure reprinted from Ref. [8]. The solid line represents the calculations for a cross section that gives the mean measured relic density (agreeing with the observations from the CMB), and the yellow, green, blue, represent $10^1, 10^2, 10^3$ deviations from such cross section. This suggests that the dark matter annihilation cross section is at the weak-scale, and that the dark matter may be weakly interacting, thus the name Weakly Interacting Massive Particle, or WIMP.	14
1.8	A Feynman diagram of dark matter particles interacting with SM particles with an unknown interaction/coupling. Depending on the direction of time, we can have different types of interactions (and detection methods). Annihilation occurs when time flows left to right, and dark matter annihilation creates Standard Model particles which can be observed by astronomical methods [40]. Combination occurs when time runs in the opposite direction, when Standard Model particles collide to produce dark matter which could be observed in experiments like the LHC [33]. Direct detection, the case considered in this work, is the interaction of a dark matter particle with a Standard Model particle, where both particles recoil (time flowing bottom to top).	16
1.9	Picture of a SuperCDMS detector (left) and a diagram of a WIMP interaction with a nucleus in the germanium crystal (right).	17
1.10	The number of expected WIMP events interactions as a function of recoil energy interacting in a SuperCDMS Ge detector, for multiple WIMP masses, assuming a dark matter halo in the Milky Way [43] and a wimp-nucleon cross-section of 10^{-42} cm ² for simplicity. Figure reprinted from Ref. [46]. The quickly dropping rate as a function of recoil energy suggests that optimizing for low energy interactions not only improves low WIMP-mass searches, but also increases the sensitivity for higher WIMP masses, as the number of interactions drops quickly with recoil energy.	18
1.11	This figure shows the regions of dark matter-to-nucleon cross section limits vs mass which have been excluded by multiple direct detection experiments [23] Figure generated with the SuperCDMS limit plotter tool [50]. The area above and to the right of the solid lines has already been excluded. The dashed lines and pink space represents the expected sensitivity of near-future searches. The dotted brown and yellow space represents the solar neutrino background, which will impose a difficult background in the future.	20

2.1	A diagram of how particles interact with the electrons (left) and the nucleus (right) in the crystal. Because the energy is deposited based on the amount of energy given to the electron or nucleus, we will refer to these types of interactions as electron recoils and nuclear recoils respectively. Electron recoils are scatters with electrons in the crystal, knocking the electrons out of the atoms and liberating electron-hole pairs. Nuclear recoils are interactions with nuclei in the crystal, inducing mechanical vibrations propagating as phonons. WIMP interactions are expected to be nuclear recoils, while for our purposes, electron recoils are considered background interactions.	27
2.2	Top: Diagram of energy distributions in the charge and phonon signal depending on particle interaction type. As electrons and holes propagate through the crystal and pick up energy in the electric field, that energy is deposited into phonon energy, denoted as $E_{Ph,NTL}$, where NTL stands for the name of the process, Neganov-Trofimov-Luke gain. Note that the total amount of collected energy for an electron recoil event is larger for a nuclear recoil event, thus charge energy for photons and electrons dominates. For neutron and WIMP interactions most of the energy is released into the phonon system, but some fraction is deposited into the liberation of e/h pairs (which causes additional phonon NTL energy). Bottom: Pair of plots that show the amount of energy recorded in the ionization/charge processes as function of phonon recoil energy (left), and the ionization yield as a function of the amount of energy recorded from phonon recoil energy (right). In both plots electron recoils (^{133}Ba calibration in blue) and nuclear recoils (^{252}Cf calibration in green) are clearly separated in the ionization yield quantity, but the challenge is to separate at low energies where they overlap.	28
2.3	A set of diagrams showing the SuperCDMS Soudan detectors, the apparatus surrounding them, and its location which is deep below the surface of the Earth in the Soudan Underground Mine. The experiment is underground to block from particles incoming from space, and the apparatus is designed to reduce the rate at which the remaining SM particles enter the detector. The top middle figure (reprinted from Ref. [52]) shows the apparatus shielding of the detectors by both lead and dense plastic material to prevent from other sources of electron and nuclear recoils (both from contaminants in the environment and from remnant cosmic particles). The right figure (reprinted from Ref. [53]) shows the shielding of each tower of detectors. The apparatus consists of 15 detectors: 5 towers, with 3 detectors each, shown in the bottom middle figure. This work is focused in comparing simulations with detector 4 (labelled 1104 in the figure), the germanium, middle detector in tower T2, as it is known to be the most sensitive detector	30

2.4	Left: Zoomed view of the electrodes and phonon sensor arrays, running interleaved along the surface of the detector. Right: Cross section diagram of a portion of the electric field in the crystal created by the arrays, reprinted from Ref. [53], showing the location of the electrodes on the surface of the crystal as well as electric field lines, which in turn determine the location of the collection of electrons and holes. In red and blue rectangles are the interleaved bias voltage values.	33
2.5	Two diagrammatic views of one side of the charge readout system. On the left, the charge readout is composed of electrode lines at ± 2 V, interleaved with the phonon sensor lines at 0 V, figure reprinted from Ref. [53]. On the right, the detector channel map where the inner electrodes are combined into a single channel, denoted B in the figure, and the outer electrodes are combined and denoted as A. The top and bottom systems provide an independent measurement of charge energy. The four separate measurements can be used help determine if the primary charge energy is fully reconstructed, as well as provide information about the location of the primary interaction.	34
2.6	The lines containing the TES collection elements (left) and the combined detector phonon readout channel map (right), on both surfaces of the detector. Figure on the left reprinted from Ref. [53] and figure on the right from Ref. [56]. Each channel is an array of collection elements which act as variable resistors to produce the phonon signal.	35
2.7	This figure (reprinted from Ref. [57]) shows the multistage decays of Uranium which, if within the shielding around the detectors, can decay and cause difficult-to-remove background events in our experiment. The final states (^{210}Bi , ^{210}Pb , and ^{206}Pb) can produce α particles (^4He nuclei) and electrons, which can produce nuclear and electron recoils in the crystal. The hexagon shows the nuclear states, and the notation above the arrow shows the particle created in the transition, α and β emission.	39
3.1	A diagram of an incoming particle interacting with the detector, creating both primary phonon and electron/hole interactions which then propagate throughout the detector. Electron-hole pairs are created from the effects of the initial interaction in the crystal, either by electron recoils or from by-products of a nuclear recoil. Primary phonons are created from a nuclear recoil interaction which quickly convert into secondary photons. From the interactions of electrons and holes in the crystal additional energy is added to the system, as they are accelerated by the bias voltage, however all of this energy is released into the phonon energy.	44

3.2	The Lindhard yield describes the proportion of energy released into charge and phonons (vibrations) of a crystal lattice, from a single scatter with a nucleus in the crystal. The figure shows the Lindhard yield ($Y_{Lindhard}$) as function of the nuclear recoil energy [60]. The fraction is typically below 1/3. For electron recoils all the primary energy is converted into liberated electron-hole pairs, and none released into primary phonons, so the Lindhard Yield doesn't apply and can equivalently be set to a value of 1.	48
3.3	The total expected energy for the charge and phonon readout systems, as a function of recoil energy, for electron and nuclear recoils. Left is for electron recoils (like photons and electrons), where all recoil charge is from the recoil interaction and all phonon energy is from NTL gain. Right is for nuclear recoil events, like neutrons or WIMPs. The fraction of the energy in the phonon system is significantly larger because of 2 added effects: the Lindhard Model scale factor (well below 1/3) which puts a larger fraction of the energy into phonons, and the phonon energy from NTL phonons, phonons created from the moving e/h in the accelerating field, which is correspondingly smaller. The added effects account for phonon energy to be ~ 4 times the charge energy. ..	51
3.4	Two figures describing the lattice structure and propagation properties of germanium crystal. Left: The Brillouin zone for the Ge crystal which is a diamond cubic lattice. Notice that while we have indicated many of the labels on the left hand side, we have only indicated the Γ and L directions on the right hand side. For electrons there are eight L points, which dictate the propagation direction. Image reprinted from Ref. [65]. Right: Germanium band structure. The lowest conduction valley is at L (blue arrow), which makes electrons propagate in the [111] plane. The maximum valence band value is at Γ (red arrow), and is roughly parabolic around the maximum, hence holes propagate ballistically. Image reprinted from Ref. [64]. The electric field in the detector is aligned in the Z direction, so for interactions away from the edges of the detector (where the field changes rapidly) electrons propagate in 4 narrow cones (following L-direction), towards the positive biased detector side (side1, top), while holes propagate freely in the potential pulling them to the negatively biased side (side2, bottom).	54

3.5	A set of figures showing the charge readout of the detectors and its response in the presence and absence of a particle interaction. Top left shows the charge signal amplification schematics used in the FET readout system, where the detector serves as a Capacitor (C_d) in the circuit (figure reprinted from Ref. [73]). Each of the four channels in the detector functions as a circuit responding to the capacitance, as seen by the electrode's orientation. The typical output real-interaction event looks like the top-right figure, while the amplification circuit noise is shown in the bottom-left plot. The bottom right figure shows the noise profile or Power Spectrum Density (PSD) which shows the intrinsic power as function of frequency (described in more detail in the next chapter), of the charge circuit of one channel in the absence of real interactions, taken from 3 different samples, showing that noise levels remain constant as baseline output from the filter-circuit. It is clear that such noise is not white noise, from the construction of the noise profile it is known that it is also not cross-correlated in between frequencies.	61
3.6	A diagram showing a TES readout sensor, where the red indicates a particle that interacts in the crystal which creates a high energy phonon (yellow) which in turn is absorbed in the Al fin, breaking up a pair of bounded electrons (green) which can then be collected as part of the readout electronics. If the phonon energy is above the Cooper pair gap energy, the pair is broken, and the excited particles relax in the decreasing energy gap of the Al / W interface, releasing energy that heats the electron system of the W in the TES, which is carefully maintained in the superconducting transitioning state. Image reprinted from Ref. [76]. The transitioning state is hyper-sensitive to temperature changes, strongly responding with increased resistance as the W heats-up, and producing an observable signal.	63
4.1	A diagram of inputs and results of the optimal filter calculation. The algorithm finds a minimization of the χ^2 as a function of amplitude and time-delay.	66
4.2	Example side1 templates: inner and outer channels, and 2 cross-talk templates. Notice that the templates are normalized to the channel-templates (not cross-talk), it is expected for the channel-signal to be much stronger than it's cross-talk. An optimal filter is performed for each of these and the minimized coupled system determines the amplitude of the event.	70
4.3	An example of a simulated event along with its best fit, showing the resulting amplitude and time delay values for this event. The fitting template, normalized by the best-fit amplitude and delayed by the respective time is in red, while the waveform (with noise) in green. Noise is de-weighted by the fit. Notice that this is a simulated event where we know the energy that the waveform should represent, and the difference between the observed and expected is 0.9 eV.	73

4.4	This figure shows two different ways of determining the amplitude in a sample of events with no interaction (done with noise data, which is slightly contaminated by mismeasurement effects described in Chapter 6 causing the left-hand-side excess). In the top, we allow the optimal filter to look at all times to pick a minimum χ^2 , and the result is clearly biased away from zero due to fluctuations (we allow the amplitude to be negative). The bottom plot shows the same data, but where the time-delay is fixed to be the same value for all events, the 0-delay time ($t = 1024$ time-bins). In this case the average energy is expected to be zero, and be well-described by a single Gaussian centered at zero with an RMS determined by the resolution of the noise (charge readout circuit in the absence of a particle interaction).	76
5.1	A simplified diagram of a calibration source, in this case ^{133}Ba , near the detector and how it produces photons which enter the detector along with the particles that propagate within it in response. This source is used as a calibration input for the detector. We note that the primary photons may or may not reach the detector, and can interact in the crystal or just pass by. Some of them are blocked before reaching the detector, others scatter before being fully absorbed (and their subsequent scatters can reach the detector), others can reach it without scattering, and some can pass through without depositing energy.	79
5.2	This figure shows a energy level diagrams of ^{133}Ba and a schematic of the ^{252}Cf decay processes. Sources which are placed near the detectors for calibration and detector response studies. Top: Schematic representation of the energy levels and transition energies that produce high energy photons in ^{133}Ba , reprinted from Ref. [78]. A barium atom decays into Cs, and the electrons release photons with discrete energies, due to the energy levels they transition to. The energies are represented by the horizontal lines, and the energy values (in keV units) are in the left. The down-arrows show possible transitions, and the difference between the energy levels gives the energy of the released photon. These energies are characteristic to the sample so they can be used in our detectors as signature peaks in the data. Bottom: Diagram of ^{252}Cf decays producing photons and neutrons that reach the detector. Of particular interest are neutrons used as a source for nuclear recoils for a variety of energies. The ^{252}Cf source is used for calibration runs of the experiment by placing it near the detector to study nuclear recoil response.	80

5.3	Two figures showing the expectations and results for the SuperCDMS Low Threshold (LT) analysis, reprinted from Ref. [47]. On the left is the background model. On the right is the observed signal events. After the quality and basic event selection criteria, the background model was used for constructing an arbitrary parameter to differentiate signal from background. The final result determined that 11 events were signal-looking, but upon close inspection they were related to noise and other types of possible fluctuations.	90
6.1	Pileup examples in noise-data events in the four different charge readout channels. The set of events shown were in noise-collection mode (random readout time) but when the ^{133}Ba was placed near the detector and a large E_0 energy is recorded.	95
6.2	Example Blip events which show small signals in the waveform as a function of time. These figures show the readout from each of the four channels from a noise sample. Individual events are shown in blue, green, and yellow. Given the low amplitude of the Blip-feature it is easier to see the average Blip shape (black). The orange band shows $\mu \pm \sigma$	96
6.3	Example of a Glitch readout malfunction waveform (left), affecting both charge and phonon readout, and a low frequency noise example affecting a phonon event (right).	96
6.4	This plot shows the 0-delay qimean energy measurement from the full set of events in the IRR sample taken when the ^{133}Ba source was placed near the detector. While the bulk of the distribution is from readout noise (main peak at zero energy) it is contaminated by Blips and Pileups. The high energy events are due to pileups, while the non-Gaussian left hand side tail is from Blips.	98
6.5	The distribution of the maximum value (over baseline) in ADC counts, called Pulse max, for QIS1 for iZip4. The simulated noise events (using the methods of Chapter 7) are shown on the left and IRR data is shown on the right. Since the data appears to be well described by a combination of noise and pileup events we set the threshold for removing pileup events at 150 ADC counts (dashed orange) which is the lowest threshold that does not remove simulated events.	100
6.6	The plot shows the E_0 distribution of the subsample of IRR events after pileup removal. The data is purely Gaussian on the right hand side, but still has a left tail from Blip contributions. We also note that the mean is not zero, which as will be shown later is also due to Blips.	100

6.7	This figure shows the 0-delay energy measurement of each of the 4 charge readout channels for all the events in the IRR sample with $q_{\text{imean}_0} < -1.2$ keV. The measurements in all channels are correlated. In all distributions the sample is negatively biased, suggesting that the Blips are real, correlated, and not a random fluctuation.	102
6.8	The measured 0-delay energy in all iZip5 channels when there is evidence of a Blip in iZip4. This set of events, selected using the LHS sample, provides evidence that the LHS sample is correlated in iZip5-side1 and anti-correlated in iZip5-side2 channels. In all distributions the sample is biased away from zero, suggesting that these Blips are observed in both detectors.	103
6.9	The BlipFinder results (BlipFinder time and energy) applied on noise simulations (NoiseSim - red crosses), the full IRR data sample (in blue), and on the sample of LHS events (in yellow). There is no preference in BlipFinder time for simulations as well as for data (except for the first time-bin which is ignored because it is far from t_0). On the other hand the distribution is very different for the LHS sample and is correlated with Blips found at t_0 . The BlipFinder energy of the IRR data is greater than that of NoiseSim, suggesting Blips are common.	106
6.10	The 0-delay energy measurement as a function of BlipFinder time. Only Blips occurring before time-bin 375 and after 1450 do not affect the measurement. The most affecting Blips are those near the 0-delay time (time-bin 1024), causing the negative (LHS) tail, while Blips occurring before or after the 0-delay cause the positive offset in the Gaussian distribution shown in Figure 6.6.	107
6.11	This figure shows the time difference, in time-bins, between the biggest blips in iZip4 and the combined iZip5+6 BlipFinder result (100 bins around iZip4 BlipFinder time). It is further used as a weighting function for the iZip5+6 BlipFinder search (100 bins around iZip4 0-delay). The fit function is a Gaussian with $\mu = -4.9 \cdot 0.8 \mu s$ and $\sigma = 7.5 \cdot 0.8 \mu s$	109
6.12	This figure shows the time and energy predicted in iZip4 from results for the TowerBlipFinder which looks for evidence in iZip5+6side2. The results are for both the full IRR sample after pileup removal (blue) and the LHS subsample (yellow). We note that the LHS events (known to be iZip4 Blips) are correlated in time and energy with the iZip5+6 BlipFinder results. The overflow bin (time bin 1074) in the left hand plot results are not from the algorithm finding a large negative excursion, but rather zero-values from the de-weighting effects.	110

6.13	Top-Left: iZip5+6 Correction Data. X-axis is iZip5+6 BlipFinder time, Y-axis is iZip5+6 BlipFinder amplitude, and Z-axis is observed iZip4 qimean 0-delay energy. Top-Right: Fitted shape for the iZip5+6 Correction 3D fit based on the data from the left with the same axes. Bottom: iZip5+6 prediction fit example, from a sample of events with iZip5+6 BlipFinder Amplitude amplitude of $-230 \pm 50 \mu V$. Data in green, profile in red and fit in orange. For fitting planes of the full data see Appendix D. These figures show that the 3D fit is a good representation of the data and can be used to predict the effect of iZip5+6 Blips on iZip4 0-delay estimates.	112
6.14	These plots show the amount of energy predicted in iZip4 based on the TowerBlipFinder tool that only looks at iZip5 and 6. Left: Recalculated qimean 0-delay iZip5+6 predicted energy for the sample. Notice all but 2 LHS events (original qimean 0-delay tail) are biased to lowest values of the iZip5+6 prediction. Right: iZip4 qimean 0-delay (not corrected) distribution of sub-samples separated by iZip5+6 Prediction above of below -0.3 keV. Notice events with high-negative correction (below -0.3 keV) make up the LHS tail of the iZip4 distribution and is highly negatively biased.	113
6.15	This figure shows two versions of the qimean measurement before (left) and after (right) the TowerBlip correction energies. After correction the sample is purely Gaussian (few events on LHS), mean offset is gone, and RMS is lower, showing that the correction works as expected.	113
6.16	The iZip5+6 prediction fit to near peak region ($\mu - 1\sigma, \mu + 2\sigma$). Observe that the correction method introduces the prediction's resolution to the corrected energy distribution, added in quadrature with the pure-Gaussian noise resolution.	114
7.1	A diagram showing the various stages of the simulation (on the left) which mock up interactions in the real detector (on the right). The first stage (SourceSim) is about simulating the sources and the primary interactions with the detector (and any shielding material in the experiment). The next stage (DetectorSim) simulates the response to the interactions, from electron, hole, and phonon propagation, to sensor readout and noise. Next the same format as data is produced to analyse with the same Reconstruction software as used in real data. Finally, the events are run through Reconstruction to produce Processed events.	116
7.2	Detailed view of the simulations workflow (courtesy of Richard Lawrence), showing all the subcomponents of the various pieces of the simulation.	118

7.3	Parallel diagrams of the propagation, interactions, and readout comparing the real experiment (top) and simulations (bottom). A radiation source creates events (can also be WIMP signal events, or cosmogenic showering, etc.) that pass through the experimental apparatus and reach the detector crystal. After that the incident particles deposit energy. All instances of scattering are simulated up to interactions in the crystal; multiple interactions can occur within a single event and are referred to as multi-scatter events. The charge and phonon sensors produce a signal which is recorded by the readout electronics, and stored in a file for Reconstruction. The simulations replicate the process by propagating the electron-hole pairs through the system and simulating the response of the charge and phonon sensors.	122
7.4	A figure showing the trajectories of electrons (left) and holes (right) for an example event which scattered in the middle of the detector. As expected from detector physics, charge trajectories well-represent the expected physics: electrons propagate in the restricted [111] plane (L-valleys) of the crystal, while holes propagate ballistically. All of these electron-holes were absorbed at the top/bottom electrodes.	126
7.5	Example real data noise event (left), and NoiseSim simulated event (right) randomly sampled from the input PSD. The simulated event is indistinguishable from typical real data noise.	130
7.6	Figures showing the energy distribution of photons emitted from a ^{133}Ba source at various stages in the simulation. Barium emits photons with energies following delta functions at the excitation-transition levels (top left), as given in Figure 5.2. As some of the photons are absorbed in the copper shielding, not all of the photons reach the detectors, the original/primary energy of photons reaching the detectors is in the top right figure. The distribution of energy deposited (bottom) in the detector results from the interactions of electrons and photons either directly from the source or after scattering in other material, hence a Compton spectrum in addition to the smeared signature Ba peaks. The expected charge energy is equal to the true recoil energy (all energy is converted to liberated e/h).	133
7.7	This figure shows the simulated number of scatters in the detector for a sample of 356 keV photons from the ^{133}Ba source. We note that for these high energies all events are multiple-scatters, which increases the possibilities of mismeasurement due to position dependent electron and hole collection.	134
7.8	The recoil (left) and expected (right) charge energy distributions for our two WIMP samples. The expected charge energy is calculated per Lindhard Theory [60] from the recoil energy. The WIMP sample is created with a 50 GeV WIMP mass, and the other sample is created by sampling a uniform random spectrum of 40 to 200 keV recoil energy (equivalent to ~ 10 to ~ 80 keV expected charge energy).	135

8.1	This figure shows the theoretical resolution of the 0-delay qi1 energy measurement as a function of BlipFinder energy sub-sample selection. The horizontal dashed lines show the value from the original PSD (red) and the observed from the corrected IRR fit (orange). The BlipFinder energy sub-sample selection value that intersects the orange dashed line (vertical dashed green) is the value for the final sub-sample selection to create the BlipCleaned PSD. ...	140
8.2	This figure shows two different versions of the PSDs. The first is from the BlipCleaned sub-sample that reproduces the observed RMS of the BlipCorrected sample. The second is the original IRR PSD (as calculated by the event reconstruction method). Note the difference in the sub-10kHz range.	140
8.3	The measured 0-delay energy distribution of a fully simulated and processed NoiseSim sample of events from qi1, where we fix the time delay of the fit to be 0 as consistent with noise estimates described in Chapter 5. Note that it is a single Gaussian with a measured mean of zero, that there are no tails, and that the expected RMS is 317 eV while the observed RMS is 312 ± 4 eV.	141
8.4	Overlaid noise simulations (blue) and Blip-Corrected IRR sample (red error bars) 0-delay energy measurements for qi1, qi2, qo1, qi2, qimean, and qsummax. The fit region for mean and RMS measurements is shown in grey, with the results shown in Table 8.1. Simulated samples are in agreement within error, except for qsummax LHS tail, which is produced by correlations of single-detector Blips.	142
9.1	Heat-plot histogram of the number of scatters as a function of recoil energy for the ^{133}Ba simulation. Above 230 keV all events are multiple scatters. Most of the events have less than 8 scatters, and below 100 keV less than 3.	151
9.2	The figures show the mismeasurement (ΔE) and relative mismeasurement ($\Delta E/E_Q$) of the single-scatter nuclear recoil (Enhanced WIMP) sample, with energy above 3 keV (to avoid the lowest energy mismeasurement which is described later). On the left plots we show the full-sample where we can see many badly mismeasured events, and on the right plots we zoom-in to the well-measured region. In the top we note that the peak is not Gaussian due to energy dependence, so we turn to the relative mismeasurement as the resolution is linearly dependent on the energy (in the bottom). The well-measured region is a single Gaussian peak, centered at zero. We will select a clean sample of badly measured events by selecting events, below $\Delta E/E_Q < -0.015$, 7σ below the mean, to look at the mismeasurement position dependence.	152

- 9.3 This figure shows the true position of a clean sample of mismeasured events ($\Delta E/E_Q < -0.015$ with $E > 3$ keV) with a single nuclear scatter. On the left we show the Z-dependence limits, selected from events with low radius. On the right we show the R vs. Z position of the mismeasured events which are contained in the truncated cone volume (dashed orange), we denote the inner-volume as the SimFiducial region with the limits defined in Table 9.1. .. 154
- 9.4 Comparison of qi1 and qi2 energy measurements of the data (left) and simulation (right) full samples. The first thing to note is that the bulk of the distributions fall along the line where the two have equal values, but the distribution is much wider in data (worse resolution). While the yellow dashed lines indicate the LT fiducial requirement (which appears to be reasonable for real data, but would be much too loose for simulations), both samples show events where the energies are very badly measured and the values from the two sides are very different from each other, in some cases where one side reads out zero energy while the other side has a large energy. The colored events in the simulated sample indicate which events are shown for study in Figures 9.5 and 9.6. 156
- 9.5 The true energy deposition physical location in the detector (Z vs. R) for badly mismeasured events, left is for 0-energy measurement in qi1 (no electron collection) and right is for 0-energy measurement in qi2 (no hole collection). The events are selected as being either the orange or yellow from Figure 9.4 where only between 50 and 250 keV of energy was deposited on the other side to ensure a large energy scatter. The points represent only scatters with more than 100 keV deposited from such events. This shows that the mismeasurement is related with position dependence in particular due to the interleaved potential as scatters interacting close to it get absorbed by the near-end surface while mismeasured by the opposite side. Furthermore, events with large radius for inner-side2 produced a 0-energy observation because all the holes either got trapped in the detector side-wall or entirely got absorbed by the side2-outer electrode (because holes propagate freely), and we note that for 0-energy in qi1 there are no scatters with large radius (left plot) because electrons propagate in the L-valleys (with a wider angle) causing a non-zero deposition in qi1. 157

9.6	<p>Top: The true position of events that are poorly measured in one side while well measured in the other one. The events are selected by requiring one side to be within 2 keV from the 356 keV peak, while the other side to be greater than 50 keV and lower than 250 keV. The points show the scatters with more than 100 keV energy from such events. This figures show that at least one of the scatters in the event deposits a substantial amount of energy outside SimFiducial. Bottom: Outer channel energy measurements of side1 and side2, for 3 samples of 356 keV events: well-measured (356 ± 2 keV on both sides), poorly measured in qi1 (while well-measured in qi2), and poorly measured in qi2 (while well-measured in qi1). The peak at zero is when things are well measured, showing the utility of requiring no energy in the outer channel, and the large number of events with bad mismeasurement on side1 when we have large amounts of energy in the outer channel shows that this requirement removes badly mismeasured events. We note that while this works well for side1, it does not for side2. The reason for this difference is that electrons propagate in the L-valleys with a wider normal angle, thus reaching higher radii and depositing energy in the outer channels.</p>	159
9.7	<p>Plots of the qi2 energy measurement as function of the qi1 energy measurement (data on the left and simulations on the right) after the requirement that there be no energy measured in the outer charge channels. The orange dashed line shows the region selected for defining the LT Fiducial (symmetric measurement) criteria. While both simulations and data show that both sides are linearly correlated it is clear that data has much worse resolution on both sides, which makes it hard to trust if the mismeasurement effect really exists in data. Indeed, the LT Fiducial constraint (the symmetry selection) removes these outliers in our sample (and data). The outliers in the simulation sample, in particular events with a side1 energy that appears to be fully measured (356 keV) but much less energy in side2 are due to some (or all) scatters depositing energy outside SimFiducial. We select 4 events to inspect closely in Figure 9.8.</p>	160
9.8	<p>The true energy deposition physical location in the detector (Z vs. R) for all scatters of the 2 most poorly measured events on each side. The sample is selected from the 356 ± 2 keV events in one side and the lowest measurement above 50 keV in the other) after the requirement of no energy in the outer channels, the 2 events with mismeasurement in maroon (qi1) and red (qi2) in Figure 9.7. In all cases there is a scatter with a position that is close to the interleaved potential of the opposite side, preventing charges from reaching their collection side.</p>	161

9.9	The ΔE measurement which compares the true and observed qimean from fully simulated ^{133}Ba events where we only consider input photons with a true energy $E_{true} = 356.013 \pm 0.100$ keV. The events in green, selected as having passed the LT Fiducial requirements, show that LT Fiducial removes most outliers, none below $\Delta E = -14$ keV, but still has a few poorly measured events. The right hand side shows a zoomed-in version of the plot on the left so we can see the most poorly measured events which pass the LT Fiducial requirement. The four events with $\Delta E < -10$ keV are shown in Figure 9.10 where we will see that the mismeasurement is due to at least one of the scatters being close to the edges of the detector.	163
9.10	This set of plots shows the Z vs. R deposition positions in the detector for the four simulated events from our 356 keV photon with the worse ΔE measurement, but pass the LT Fiducial requirements. All scatters for each event are shown showing that all (or all of the high energy ones) are located near the edges of the detector. Compare with the simulation-based fiducial volume (dashed orange), defined in Section 9.3.	164
9.11	Histogram of difference between observed and expected energy for the simulated ^{133}Ba events that pass the LT Fiducial requirements, above 3 keV, single scatter events, compared with multiple scatter events (multiple scatter events histogram normalized to the single scatter events). Multiple scatters are affected by position dependence because while the events can pass the LT Fiducial requirement some of their scatters can still be near the edges accounting for energy loss.	164
9.12	The qi1 and qi2 energy distribution of the sample with LT Fiducial selection where we show only the data around the 356 keV peak, from simulation and data ^{133}Ba . We show results for qi1 (left plots), and qi2 (right plots), for the ^{133}Ba calibration data results (top), simulations with noise (middle) and without noise (bottom). The fits are done with a skewed Gaussian function on $\mu - 2\sigma$ to $\mu + 4\sigma$. The samples are well calibrated since the fit maximum value location (x_{max}) is at 356.013 keV within uncertainty, regardless of noise addition in the simulations, telling us that noise resolution is not a dominant issue at high energies. We note that the resolution on data is much worse than it is in our simulations, in particular qi1.	167
9.13	The best-estimate energy observable qimean 356 keV peak measurement with events passing the LT Fiducial requirements for data (left) and simulations (right). As previously noticed, the RMS in data is much worse than in our simulations. Of particular note is that while the resolution in qimean improved in simulations, it worsened with respect to qi2 in data due to the large RMS of qi1, an indication that qi1 ought to be treated differently in data, although for large energy searches qsummax is used.	170

9.14	The energy distribution of events from the Ba calibration sample from both data and our simulation after the LT Fiducial requirement. We note that the observed energy estimate (qimean) well represents the expected charge energy, and is comparable to data, although the resolution of the Ba peaks is much better than in data.	171
10.1	This figure shows the energy, as measured by charge collection in two different ways. On the left is the SimFiducial $\Delta E_{Coll}/E_Q$ (side1), compared with the full sample. Note that there are no events that pass the SimFiducial requirements that go beyond -0.02 . On the right is the position dependence of the sample of events with poor charge collection on either side ($\Delta E_{coll}/E_Q < -0.02, E > 3$ keV), which shows that the cause of the mis-measurement outside the SimFiducial is poor charge collection.	176
10.2	Collected charge energy by the electrodes, divided by the expected (maximum) charge energy (we call this charge collection efficiency) in qi1 and qi2, for the full Enhanced WIMP sample for the full set of events as well as the subset which pass the SimFiducial requirements. The top plots show the distributions for the full sample and the SimFiducial subsample. The bottom plots show skewed Gaussian fits for the SimFiducial events. The plots show a zoomed-in version for just the subset of events that pass the SimFiducial requirements showing that they remove any case of incomplete collection (below 96%). Furthermore the fits are consistent with the absorption condition of 2% difference between the electrode and the final position (described in Chapter 3.3) because the peak values are exactly at 98.5%, and have small resolution (0.3% for qi1 and 0.6% for qi2). It is expected that the left-hand-side tail is caused by charge trapping effects.	177
10.3	The accuracy of the electron recoil energy measurement as determined from our simulated sample of ^{133}Ba events that pass the LT Fiducial requirements. The observed minus expected mean value as function of energy is shown for qi1 (top left), qi2 (top right), and qimean (bottom), with noise (green) and without noise (blue). Each of the results are basically well behaved, linear in energy, and centered at zero. The expectation of a single constant calibration works as expected for the two separate sides. A caveat is that qimean drops at a rate of 0.15% (very small effect with respect to variations and other mismeasurement reasons), which is a function of true Z-position and true energy, described in Section 10.5.3.	179

10.4	The resolution of the electron recoil energy measurement as determined from our simulated sample of ^{133}Ba events that pass the LT Fiducial requirements. The RMS as function of energy is shown for qi1 (top left), qi2 (top right), and qimean (bottom), with noise (green) and without noise (blue), and the simulated noise RMS (dashed orange). Noise dominates below 120 keV, but intrinsic detector resolution rises linearly with energy, at the rate of $0.003 \text{ keV}\cdot E$ (for qimean).	180
10.5	Two figures showing the detector response (accuracy and resolution) to single scatter nuclear recoil events from the Enhanced WIMP sample. Left: Single scatter low energy measurements are well behaved, linear and centered at zero, except for lowest energies ($< 2.5 \text{ keV}$). Right: RMS behaves as expected, linear response without noise (blue), but dominated by noise (green). Noise causes large variations when energy goes to zero because of pure noise measurements, as described in Section 4.4.	181
10.6	Lowest energies show large mismeasurement below $E < 2.5 \text{ keV}$, both in sample with noise (left) and sample without noise (right). The sample without noise also shows a mismeasurement, an artifact of the pulse digitization.	183
10.7	This figure shows the charge system readout waveform, in ADC counts as a function of time, for a simulated low-energy event ($E_Q = 1.31 \text{ keV}$) with no noise to show the problems of the digitization effect on our energy measurement. Superimposed is the best fit template shape. We note the the readout waveform values are so low that the shape from the ADC has been truncated.	183
10.8	The $\Delta E/E_Q$ distribution for the Enhanced WIMP sample for the subsamples of events which pass the LT Fiducial and the SimFiducial requirements. We use samples with $E > 3 \text{ keV}$ to ignore lowest energies limitations. We see that while we already know LT Fiducial does a good job at removing mismeasured events there is a large population of mismeasured events (even with $\Delta E/E_Q$ near -1). SimFiducial does not have this long tail, so making LT Fiducial closer to SimFiducial would represent a major improvement in data analysis.	184
10.9	Two figures showing the location of energy deposits for the enhanced WIMP sample, but have charge energy mismeasurement. Left: Shows true position of all of the events which pass LT Fiducial requirements. We note that many are outside SimFiducial volume (dashed orange). Right: The true position of events which pass the LT Fiducial requirement but which are poorly measured (selected with $\Delta E/E_Q < -0.015$), their true position is outside SimFiducial.	185

10.10	This figure shows the amount of energy mismeasurement on both sides of the detector as a function of Z for a large energy sample that pass all the SimFiducial requirements. For both sides there is a linear dependence as a function of Z which partially, but not completely, offsets the full qimean measurement. This explains the qimean measurement as a function of energy shown in Figures 10.3 and 10.5. The qimean energy measurement drops at a rate of 0.15%.	186
A.1	Waveforms of Zip4 Blips occurring within the first 2 time-bins show that these Blips wrap around the readout, a possible indication that Blips are a DAQ electronics issue.	204
C.1	BlipFinder (ignoring first/last 248 bins), time on the left and energy on the right, for Zips 4 (top), 5 (middle), and 6 (bottom). In none of these independent searches there is a preference for Blips at any time, not even when looking at Zip4 Big Blips. We define BigBlips as events with BlipFinder total energy less than -3.5keV (Zip4), -2.5keV (Zip5), and -1.5keV (Zip6).	208
C.2	The difference between the BlipFinder time (performed within 100 time-bins of the BigBlip time), and the time of the BigBlip, for Zip4 BigBlips (top), Zip5 BigBlips (middle), Zip6 BigBlips (bottom), the left column shows cases when a second detector is coincident, and the right column events where the second detector is not coincident, and the distributions are for the third detector. The timing thresholds to determine if they are coincident are in dashed orange. The samples are sub-divided from the results of the second detector (indicated in the top-left label in each plot), to look at the distribution in the third one (indicated in the x-axis label), we estimate the background by taking the average value outside the coincident time (in dashed red), and finally we calculate the background-subtracted number of events. The rates are displayed in the legends of the plots.	209
C.3	Zip5 (top) and Zip6 (bottom) BlipFinder energy (searching within Zip4 Biggest Blip time ± 50 bins), for Big Blips (left) and Small Blips (right), showing that the energy is also correlated with coincidence of Blips on another Zip. This is true for both detectors. We note that even fro Small Blips there is correlation, the only difference being a higher population of non-correlated Blips (Single-Detector Blips).	210
D.1	Zip4 qi1 0-delay energy as function of Zip5+6 BlipFinder time. The data profile is shown in red crosses (mean and with uncertainty on the mean), the 3D fit is shown in orange.	212
D.2	Zip4 qi2 0-delay energy as function of Zip5+6 BlipFinder time. The data profile is shown in red crosses (mean and with uncertainty on the mean), the 3D fit is shown in orange.	213

LIST OF TABLES

TABLE	Page
2.1 Summary of the background sources, what type of interaction they have with the detector (recoil type), and method for dealing with them. Nuclear recoils from Pb sources are the most affecting background since they mimic WIMP events.	26
4.1 Summary of calibrated energy observable quantities for the charge readout system, including 0-delay measurements for noise. Types notation: I is for independent measurement, R is for calculated, combined quantity. Side1 is the top, which collects electrons, and side2 collects holes, X (in the last row) can be replaced with any other estimator.	75
7.1 Summary of simulation variables additional to real-data reconstruction. These variables provide insight to true information to help understand mismeasurement causes.	120
8.1 For each of the six different 0-delay energy measurements, taken from Figure 8.4, we give the theoretical and observed resolution of simulated and real noise data. In all cases the resolutions are in agreement within the 4 eV measurement uncertainty.	143
9.1 The set of criteria used to remove events as being poorly measured, the Sim-Fiducial region. We use the the shape of a truncated cone. The height of the cone is given by Z_{top} and Z_{bottom} to avoid the interleaved potential trapped events. The top radius is given by R_{top} , and bottom radius by R_{bottom} , and from their intersection with the Z-limits we calculate the side limit line which is a function of R dependent on Z	153

9.2	The best fit mean measurements for the ^{133}Ba 356 keV peak locations, from simulations and data as determined for events passing the LT Fiducial requirements. These data are taken from Figure 9.5.1) and the full sample (not shown), comparing with and without noise, and comparing with real data, we also compare the full sample with the LT Fiducial selection criteria. In both data and simulations, both channels are well calibrated, within error the peak location is consistent with the true energy of the incident photon of 356.013 keV. In simulations the answer has little change when removing the LT Fiducial requirements, or adding noise. With respect to LT Fiducial with noise, it changes by $2Err_\mu$ for qi1 full sample, by $10Err_\mu$ for qi2 full sample. In data the full sample peak is also in the correct place, even though the error in qi1 increases to 24 keV due to the large resolution (which will be compared in the next table). In simulations qimean is always below the expected value by 200 to 300 eV (due to position and energy dependence as will be shown in Section 10.5.3), while this effect is not seen in data it is possible it is obscured by the large resolution.	169
9.3	The best fit RMS measurements for the ^{133}Ba 356 keV peak LT Fiducial samples comparing the simulation without and with noise, as well as with real data. While the noise adds little variation to qi1 and none to qi2, the RMS is much better in the simulation than it is in data (by a factor of $\times 16$ for qi1 and $\times 6$ for qi2).	169
B.1	Results of background subtracted fraction calculations. Note that iZip6 only considers side2 channels. For details on the background subtraction see Appendix C.....	206

1. INTRODUCTION

Sir Isaac Newton first published his theory of gravitation in 1687 to explain the movement of the planets around the sun by describing mass as providing a force that attracts other objects in proportion to their mass and squared-inversely proportional to their separation [1]. This was refined by Einstein's theory of General Relativity in 1916, which describes the same attraction, but using the notion of mass curving space-time, and objects moving through that curved space-time [2]. This proved very successful until 1933, when Fritz Zwicky observed that the velocities of galaxies bound within the Coma Cluster exceeded those predicted by the gravitation from visible objects [3]. While many tried new versions of gravity, a better explanation is that this is the first evidence that there is a large amount of matter 'out there' which can't be seen directly. It is now known as *dark matter*. Even though we don't really know what this substance is, other observations have provided additional evidence that dark matter is the most likely explanation for this effect [4], that it is a particle [5], it plays an important role on the evolution of our Universe, and makes up for $\sim 85\%$ of the matter density of its constituents [6, 7, 8].

The work presented in this dissertation is focused on simulations to be used in dark matter searches using the Super Cryogenic Dark Matter Search (SuperCDMS) experiment. It is part of a broader effort to understand the particle nature of dark matter, in particular to have a first observation of dark matter interacting in a detector here on Earth. Since this particle is expected to be highly elusive (with a very small interaction probability), there is good reason to believe that simulations ought to be a key component, as they have been for the recent discoveries of the top quark [9] and the Higgs Boson [10] particles.

We will need a fair amount of understanding before describing our work, so we begin with this chapter which provides an overview of the motivation for the existence and the predictions of the properties of dark matter, as they provide a justification for choosing the direct detection methods used in our experiments. Starting with Section 1.1, we describe the

astronomical evidence supporting the existence of dark matter as well as suggesting that it behaves like a particle. Next, Section 1.2 describes the predictions of dark matter properties from the combination of General Relativity and quantum mechanics into cosmological models of the evolution of Universe and the particles which constitute it. In Section 1.3 we describe how current experiment efforts are attempting to observe dark matter, as well as the method followed in this thesis, direct detection with our experiment. Finally, in Section 1.4 we give an overview of this thesis work. In the next chapter we will describe the experimental design and how it looks for direct detection of dark matter interactions while rejecting non-dark matter interactions from multiple sources.

1.1 Astronomical Evidence for the Existence and Particle Nature of Dark Matter

In this section we present the astronomical evidence for the existence and particle nature of dark matter. First we will show how the motion of galaxies around their cluster center and stars around their galactic center suggest the existence of a large amount of matter that engulfs the galaxies but cannot be seen, giving it the name dark matter. Then we will show a second observation which suggests that a large amount of dark matter is present in galaxies from the bending of light as it travels through and around them. Finally we will show that dark matter is likely to be a particle, given the observed distributions of matter and dark matter after the collision of a galaxy cluster.

Since the amount of light emitted from a star gives a good estimate of how much mass it has, this can be used to estimate how objects near it are expected to interact with it gravitationally [11]. While this was expected to correctly describe how galaxies are bound within a galaxy cluster, the observations of the Coma Cluster, see Figure 1.1, showed that the gravitation from visible objects is not enough to explain the high speeds of the galaxies within it [3]. Since modified versions of gravity could not explain the data [5], the data was interpreted via the hypothesis that there is a large amount of matter that exists that can't be seen directly, what we now call dark matter. This observation lay mostly dormant until in

1970 when Vera Rubin and Kent Ford made observations of the speeds of stars that orbited the outer part of the Andromeda galaxy in a way that could also be explained by a large amount of mass that could not be seen [12]. Using data like that shown in Figure 1.2 of the Galaxy M33, a better explanation of the data is that $\sim 80\%$ of the mass in the galaxy is in a set of particles that interact neither with light nor the atomic matter and is distributed in a large gas-like volume that we refer to as a halo that surrounds the stars [4].

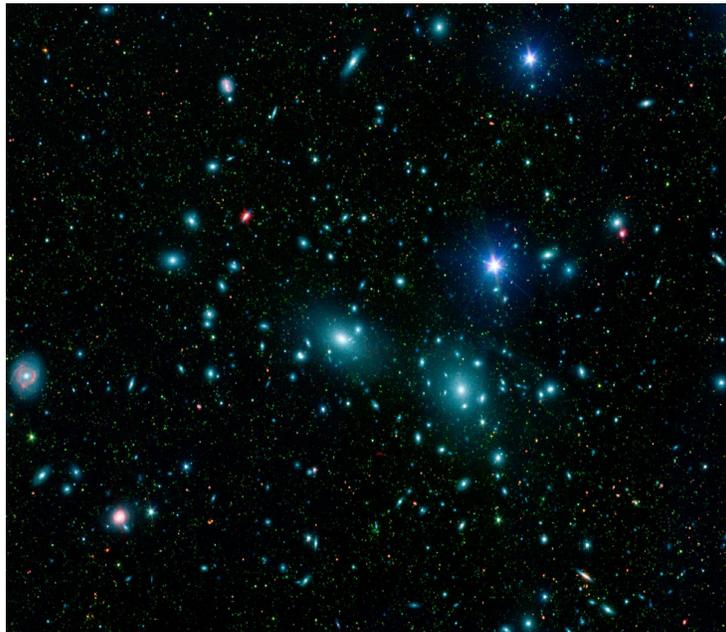


Figure 1.1: Image of the Coma Cluster, where the position and brightness of the light shows the location and density of stars, and gives a measure of the atomic mass in the galaxy cluster reprinted from Ref. [13]. It provided the first evidence for the existence of dark matter as the speeds of the galaxies exceeded the escape velocity expected from the observed amount of mass. This suggested a large, non-visible source of matter/mass.

Additional evidence for dark matter comes from the bending of the path of light as it passes through and around the large mass contained within a galaxy, a General Relativity effect known as gravitational lensing [15]. This lensing is most readily observed when a near galaxy (very massive object) bends space-time and is in the path of light coming from a light-emitting source (like another galaxy) to the observer, as shown in Figure 1.3. In

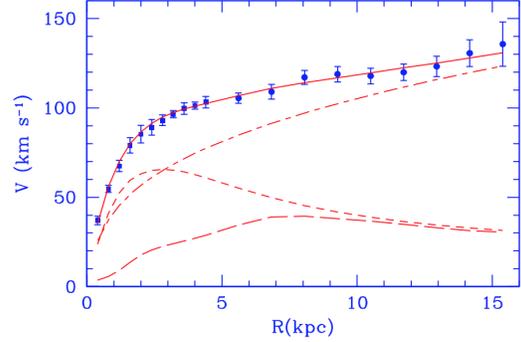
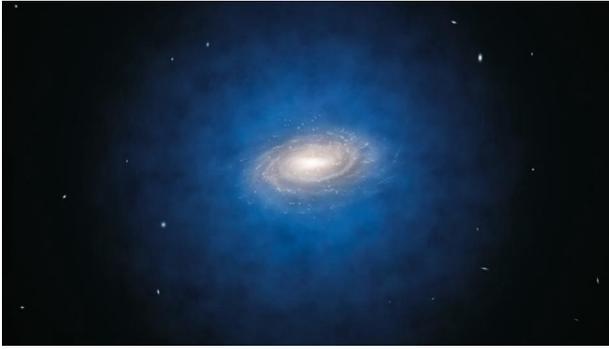


Figure 1.2: Left: Image of a galaxy with luminous matter and the calculated dark matter halo engulfing the galaxy (in blue) where the amount of light shows the mass distribution from the stars in the galaxy, reprinted from Ref. [14]. Right: Recent observations of the speeds of the stars around the M33 Galaxy as function of distance from the galaxy center, showing data (points) along with a combined rotation model (solid line) from mass contributions from a number of different sources, reprinted from Ref. [4]. The dark matter halo contribution is drawn with a dashed-dotted line, the contributions from observable stellar sources are in short dashed line, and the surrounding gas in long dashed line. The dark matter halo model well describes the observed rotation data, which suggests that there is a non-visible substance with mass that is distributed as a halo in the galaxy.

the idealized case of one object directly behind the other in the view of an observer, first proposed by Einstein [16], the bending of space-time makes light go around the object, forming a circular or repetitive pattern around the near galaxy to the observer known as an Einstein Ring [15]. In practice, distributions of mass sources typically give the appearance of a ‘distortion’ from a lens. Measurements of the size of these distortions, as shown in the right hand side of Figure 1.3, indicate that there is more than just stellar (and interstellar gas) mass, suggesting that a significant amount of non-light emitting matter is present. Gravitational lensing is now widely observed in nature and, purely non-light interacting (dark) objects forming gravitational lenses have been observed [17].

The collisions of clusters of galaxies provide evidence that dark matter behaves like a particle [5] by looking at the separation between the bulk of the mass and atomic matter distribution after the collision as shown in Figure 1.4. The explanation of the data is that the dark matter portion of the cluster is mostly non-interacting and it will continue its travel at

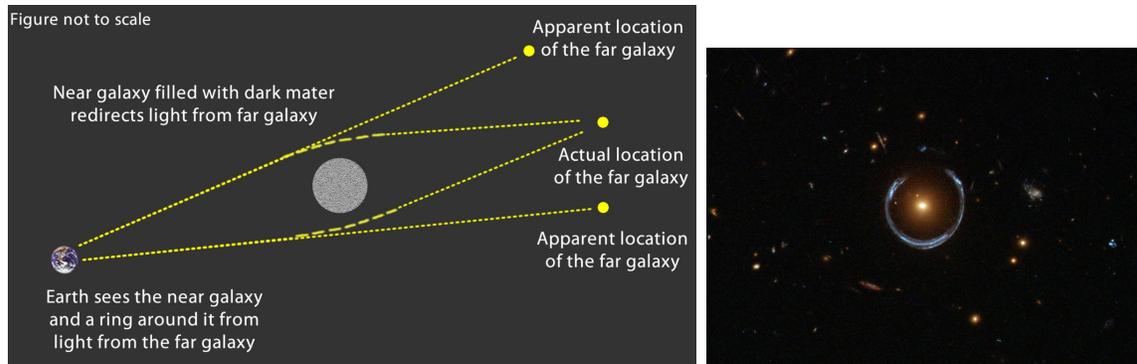


Figure 1.3: Left: Diagram of gravitational lensing, showing the view of two galaxies at different distances from the Earth, reprinted from Ref. [18]. The ‘far’ galaxy emits light that passes through, or around, the ‘near’ galaxy and which is gravitationally lensed, bending the path of the light towards us. Right: An example gravitational lensing observed in nature reprinted from Ref. [15], where the blue horseshoe shape is gravitationally lensed light from a distant galaxy which shows an Einstein Ring created by the near galaxy in the path bending the light towards us as observers. In the center, the orange bright source is a massive near galaxy (called LRG). There is greater gravitational bending than expected from mass of light emitting sources, evidence supporting that there is a non-visible substance that interacts gravitationally.

the same rate, while the atomic matter in the two clusters interacts electromagnetically and emits light as it slows down [5]. The data from examples like the El Gordo cluster [19] and the Bullet Cluster [5] are consistent with the hypothesis that dark matter behaves in a particle-like matter way, with a very low probability of interacting (or scattering) [5]. There are no known modified versions of gravity which can adequately explain these observations [5].

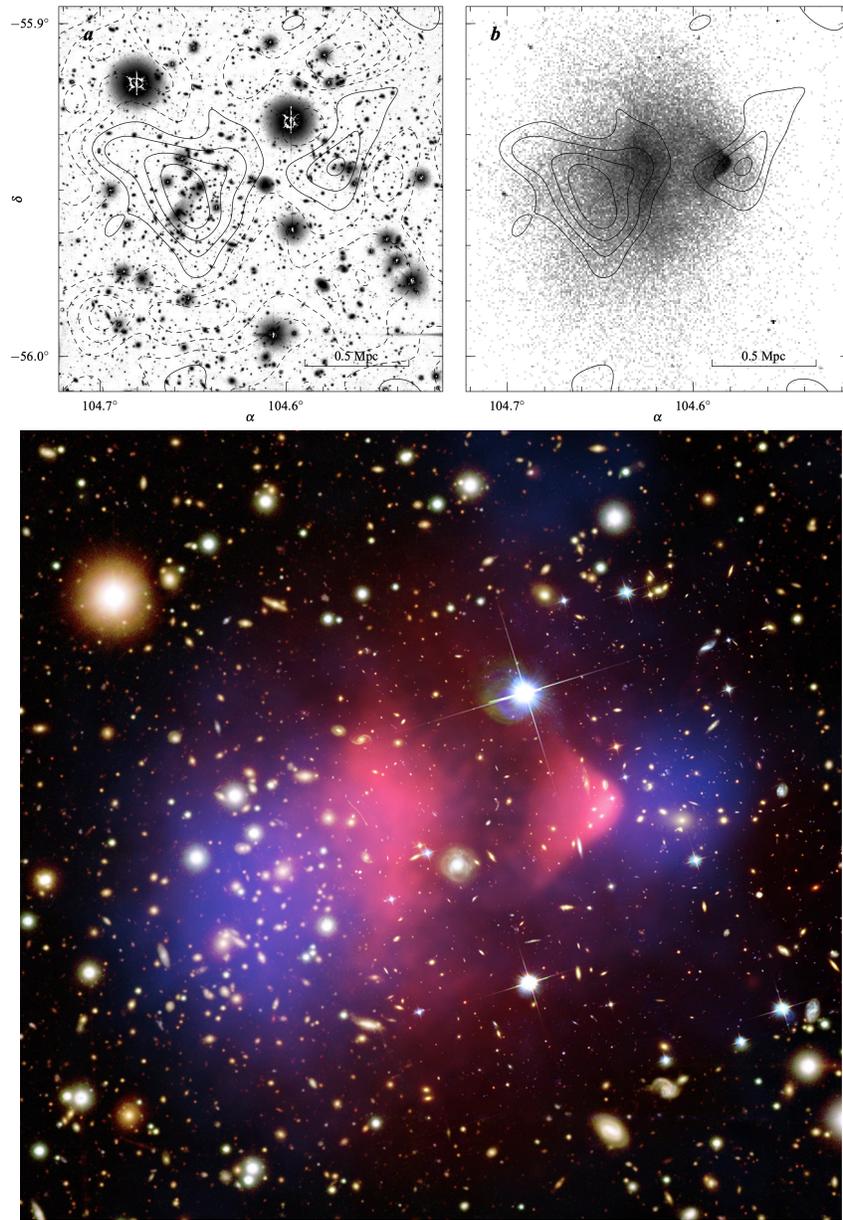


Figure 1.4: A set of images of the ‘Bullet Cluster’ which is a colliding cluster of galaxies that provide evidence that dark matter is a particle [5]. Top-left: Calculated contours of mass contributions as measured using gravitational lensing techniques, reprinted from Ref. [5]. Top-Right: Overlay of the mass contours with X-ray picture showing the amount of light from X-rays indicating the amount (location and density) of interacting atomic mass, reprinted from Ref. [5]. Bottom: Three overlaid images; mass distribution (from gravitational lensing) in blue, mass distribution of X-rays showing the area of atomic interactions (in pink), and the optical image shows galaxies in orange and white, reprinted from Ref. [20]. This data supports the hypothesis that the dark matter is separated from the atomic matter of the system, due to it being a weakly interactive massive particle, in agreement with observations when fitting with a particle-based dark matter halo model within each galaxy cluster.

1.2 Cosmological Predictions and the Properties of Dark Matter

Since the existence of dark matter affects the Universe itself as well as how all the other particles in nature have interacted within it from the Big Bang until today, we can use cosmological models to learn more about the properties of dark matter [21]. In particular we can use the laws of General Relativity (GR) and the Standard Model of Particle Physics (SM) to build up a model incorporating dark matter interactions. GR describes space-time itself and how the temperature or energy of the particles within the Universe drops as it expands, and the SM describes how particles interact within space at each of these energies [22, 23]. Taking into account dark energy [21], and the current assumption that it is described as a cosmological constant, Λ [21, 24], and the best guess of dark matter as a single type of massive particle that is non-relativistic in the early Universe (cold dark matter or CDM for short) [25], we can interpret our results in what is known as the Λ CDM model, which is the most widely accepted explanation of the large-scale structure and evolution of the Universe and is consistent with modern observations [6, 21].

In Section 1.2.1 we provide an overview of GR and SM considerations that will go into the present best-understanding of the physics behind our cosmological models. Next, in Section 1.2.2 we describe the Λ CDM model and point out the moments in its history that allow for predictions of dark matter properties and particle interactions. We continue with Section 1.2.3 and the description of the observations that justify the cosmological model and its implications for building a dark matter particle hypothesis, the *Weakly Interacting Massive Particle* (WIMP).

1.2.1 The Laws of Physics Used in Cosmology

Explaining the evolution of the Universe requires building up a theory connecting the large-scale (GR) with the small-scale (SM). We next describe each in turn, with an eye towards experimental measurements of further astronomical data that provides additional clues, and eventually a hypothesis for dark matter properties, and how we might observe

dark matter in an experiment.

General Relativity relates matter with space-time with the Einstein Equations [21]. They incorporate purely geometrical implications with energy and matter density, describing the connection between mass and space-time distortion, as well as their evolution over time. Applying these ideas to the Universe as a whole is a difficult undertaking, but is simplified with a number of assumptions that appear to be consistent with all modern observations [21]. The most important are that the Universe is roughly the same everywhere, equally distributed in every direction (isotropic and homogeneous), and that our observation point is not particular in any way [21, 24, 25]. Upon imposing this condition to the energy-momentum tensor, the Einstein Equations have an exact solution: the Friedmann-Robertson-Walker (FRW) metric [21]. Since space is expanding, and stretching the wavelengths of all the particles within it, according to quantum mechanics the particle energies are consistently dropping; since the energies of the particles are a measure of temperature, temperature can be seen as a proxy for time [21]. From measuring the current expansion rate of the Universe (using the speeds and distances of distant galaxies) [21, 26], the temperature (energies of the particles) and the matter density can be extrapolated at any point in time [8, 21, 24].

The Standard Model of Particle Physics, with the particle content shown in Figure 1.5, is the quantum field theory that describes the known particles and their interactions [22] within space. It has been incredibly successful at predicting the existence of many particles, and has produced high-precision agreement with many experiments about their interactions and properties [22, 27]. While the simplest guess is that dark matter is a particle, none of the known particles have the properties of dark matter [23, 28]. Since everything else in the Universe is related to a quantum field manifesting through particles [23], and evidence suggests that the dark matter is fairly massive (described below) the next simplest guess is that there are new particles and/or interactions in nature with dark matter being a single type of massive particle with small couplings to the SM particles. Nonetheless, using the framework of quantum field theories, cosmology may point the way for describing it, [21], in

particular with the interactions in the early Universe and how dark matter may have stopped being produced by other SM particles.

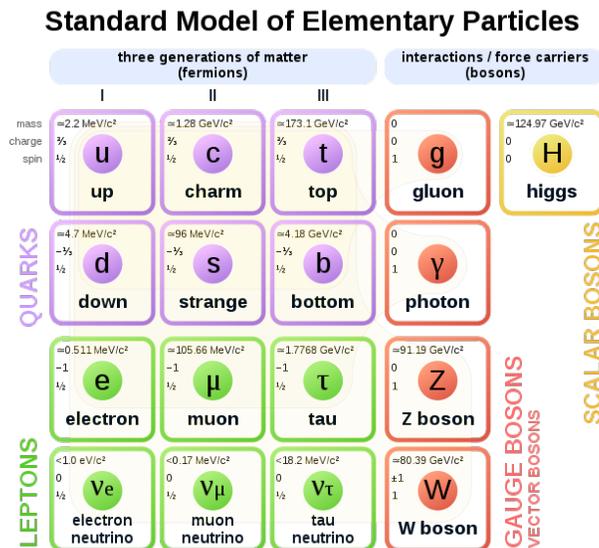


Figure 1.5: A diagram showing the known particles of the Standard Model of Particle Physics, reprinted from Ref. [29]. It has 3 types of particles: Quarks (in purple), leptons (in green), and force carriers (in orange and yellow). All the matter particles are grouped into three generations of matter, Generation I is the most stable version, involving the most common particles: electrons, up and down quarks (which make protons and neutrons), and electron-neutrinos. Force carriers are photons (mediating electromagnetic force), gluons (mediating strong-force interactions), and Z and W bosons (mediating weak interactions). None of these particles have the properties of a dark matter particle.

1.2.2 A History of the Universe with Dark Matter and Its Use to Make Quantitative Predictions

Combining our understanding of GR and the SM, and applying to an expanding Universe has led to the Λ CDM and the Big Bang model of the origin and evolution of the Universe [21] which makes predictions that can be compared to observations made today. As we will see, if dark matter were a particle it would be affected along with the rest of the known particles, which allows for us to consider it in experiments. As the Universe expands the energies of

the particles (i.e. temperature) drop, which in turn affects the types of interactions between particles and which also affects their relative abundances over time [21]. This understanding will allow us to measure the amount of dark matter in the Universe as well as constrain our understanding of its mass and interactions.

We next list four important points in our history, and then discuss why each is relevant to our story. They are:

1. When quantum fluctuations governed the early Universe
2. When the Universe stopped expanding exponentially
3. When the average energy of the particles in the Universe could no longer produce pairs of dark matter particles
4. When the average energy of the photons is no longer high enough to break apart atoms

We will next present a brief history of time so it becomes clear how each is relevant to the measurements that tell us about the properties of dark matter.

Starting from the earliest moments after the Big Bang, the Universe was small and hot and dominated by quantum interactions with quantum fluctuations causing different energies in different locations in space [21]. A short time thereafter, the Universe went through a period of exponential expansion (due to the vacuum energy of that epoch and the absence of any other type of field or matter) in a period known as inflation [21]. In addition to creating large spacial separation between these different temperature regions, which we refer to as temperature anisotropy, the rapid acceleration in the expansion produced gravitational waves which, in combination with the quantum fluctuations, produced the first variations in gravitational effects in space-time, which we call sink-holes [21].

This leads us to the second moment relevant to our story, which is when inflation ended and the quantum fluctuations that created the gravitational sink-holes for small variations in density and temperature were locked in place. After inflation, the Universe continued

expanding but at a slower (not exponential) rate, disconnecting the locations of different temperature as they were no longer in causal contact. These initial sink-holes will drive the oscillations of particles as they are created and travel through space at a rate governed by their energy, mass and interaction rates. These waves are called Baryon Acoustic Oscillations (BAO) [6] and determine the final locations of the hot and cold spots, the anisotropy, and the angular distances between them. As we will see, the temperature map of the Universe, along with the BAO predictions, allows for a measure of the density of dark matter and other particle masses at later times.

Eventually the Universe cools and the number of dark matter particles in the Universe stabilizes. This occurs for two reasons. The first is that dark matter particles are so massive that they can no longer be created in high energy interactions of other (SM) particles. The second is that the Universe got big enough that dark matter particles stopped finding each other to annihilate. This is the time of dark matter freezeout [21, 25]. Since dark matter is no longer produced and it is unlikely to annihilate, the abundance of the dark matter stabilizes and is known as the dark matter relic density. The value of the left-over dark matter relic density at a specific point in time can be measured today and compared with predictions from the Λ CDM models to determine intrinsic dark matter properties.

The last critical time in our history is when the temperature dropped below the energies needed to break apart atoms, which allows for measurements of temperature (and its variations) in different directions of space today, which in turn allows for measurements of dark matter properties of the epoch. This time is known as Recombination [24, 30]. The location of the last scatter of a photon by a charged particle (i.e. before it was able to stay in a neutral atom) is determined by the local temperature, and since the Universe has expanded by the same amount, looking at the temperature distributions of the photons in the Universe today allows for a measurement of the temperature in every direction back then, which in turn unveils the anisotropy of the Universe (and dark matter density) at the time of Recombination. These remnant photons from this point are known as the Cosmic

Microwave Background (CMB).

With the history in place, we can see how measurements of the temperature anisotropy imprint from Recombination, the CMB, allows for measurements of the dark matter at Recombination as it oscillated through space with the other particles in the early Universe. [21, 25]. We will compare the predictions and observations in the next section.

1.2.3 Data and Cosmological Models Make Predictions About Dark Matter: The WIMP Hypothesis

The Λ CDM model, along with the BAO, allows us to interpret the measured anisotropy in the CMB in terms of the dark matter properties. The current measurements from WMAP, SPT, Planck, and other experiments [6, 7, 26], are shown in Figure 1.6. There are a number of results worth noting. First off, the CMB observations tell us that the temperature today is mostly the same in all directions, confirming that in the earliest moments after the Bang the Universe had a roughly constant energy density everywhere, and that each portion of the Universe came into local thermal equilibrium [21, 30]. However, the fine-detail measurement of the temperature anisotropies, which show fluctuations by 1 part in 10^4 relative to the average temperature, confirm the fluctuations in the early Universe [6, 7, 26]. Measurements interpreted in conjunction with the Λ CDM model and the BAO suggest that at the time of Recombination 84.27 ± 0.96 % of the matter of the Universe was dark matter [7] and that it is cold (massive and non-relativistic) [6, 8, 21]. With these hints in hand, and evidence that it is cold (massive), we are ready to look in more detail to determine more about the dark matter annihilation cross section from the dark matter relic density.

Assuming that there is a single massive (cold), stable dark matter particle that has existed since well before Recombination, we can estimate the dark matter density as a function of time (temperature). Picking an example mass of 100 GeV (in agreement with constraints from various experiments and theories [6, 23, 33, 34, 35]), the results are shown in Figure 1.7, where the shaded regions give the variation as a function of annihilation cross section from the observed dark matter relic density. At early times/higher temperatures the density

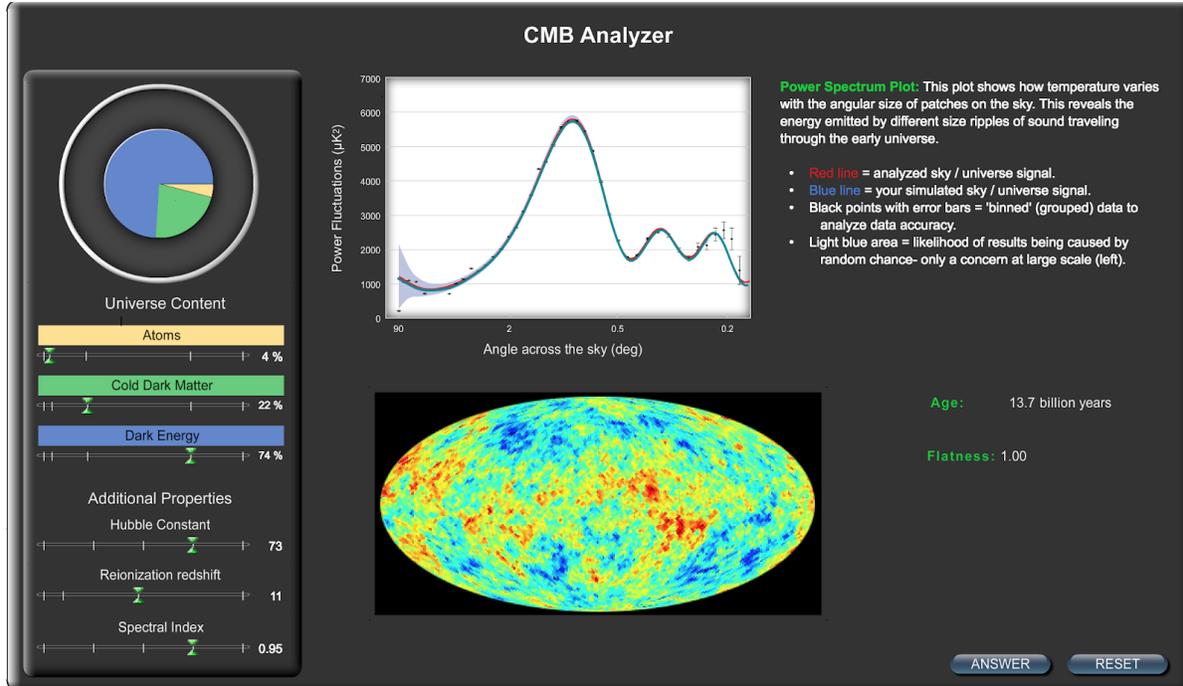
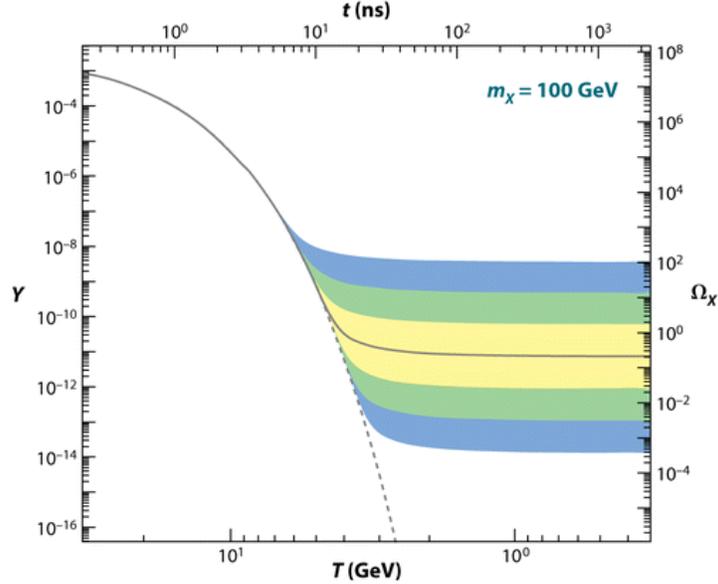


Figure 1.6: Data from Cosmic Microwave Background (CMB) measurements and their interpretation in the Λ CDM model with best-fit values of the cosmological parameters, figure reprinted from Ref. [31]. Center bottom is a picture of the CMB temperature variations heat-map, where blue and red represent colder and hotter regions respectively, by 1 part in 10^4 relative to the average temperature of the Universe of 2.7260 ± 0.0013 Kelvin [32]. Center top is the analyzed data of temperature variation as function of angular size of patches in the CMB along with the best-fit from the Λ CDM model. On the left is the best-fit parametrization of the Λ CDM model that indicates that dark matter is about 22% of the mass-energy of the Universe and thus about $\sim 85\%$ of the mass in the Universe [7].

drops exponentially as the temperature drops, and levels off when the other particles in the Universe are no longer able to create it. It flattens out to its final value once the Universe is big enough that dark matter particles no longer find each other to annihilate, and its final value is determined by the dark matter self-annihilation cross section. Remarkably, the annihilation cross section is very close to that observed from other pairs of heavy, weakly interacting particles [23, 35]. This ‘coincidence’ suggests that we should be looking for Weakly Interacting Massive Particles (WIMPs) and is known as the WIMP miracle [36].



 Feng JL. 2010.
Annu. Rev. Astron. Astrophys. 48:495–545

Figure 1.7: On the y-axis, the number density of dark matter particles as a function of temperature (proxy for time) of the early Universe on the x-axis, for a 100 GeV dark matter particle, figure reprinted from Ref. [8]. The solid line represents the calculations for a cross section that gives the mean measured relic density (agreeing with the observations from the CMB), and the yellow, green, blue, represent $10^1, 10^2, 10^3$ deviations from such cross section. This suggests that the dark matter annihilation cross section is at the weak-scale, and that the dark matter may be weakly interacting, thus the name Weakly Interacting Massive Particle, or WIMP.

1.3 Search Strategies and Current Limits on WIMPs

With evidence that the dark matter is massive and weakly interacting, we next describe how it guides our searches for interactions with SM particles that we can observe in experiments. While multiple theories have been developed that ‘predict’ the existence of stable, massive particles with weak-scale cross section interaction rates and explain their relation to the SM, none are clearly favored by the community, so we will use general considerations from each to point the way forward in terms of search strategies. We say a few words about some of the most commonly discussed models for completeness. The first is Supersymmetry (SUSY) which predicts a symmetry between Fermions and Bosons in nature [28] as an ex-

tension to the Standard Model, and provides a natural candidate for WIMPs: the lightest of the supersymmetric particles [23, 34]. A second alternative is the Singlet Scalar Model, which introduces a simple extension to the Standard Model via a single additional field [37] and is viable for a wide range of WIMP masses. Other theories are typically referred to as Dark Sector, and include axions and dark photons as viable candidates [38]. Of course, to distinguish which, if any, of these theories are correct, the dark matter particle must be directly observed first, and its properties measured. For our purposes, we don't follow the predictions of any one theory, but use the general features to guide our search strategy.

We begin with the three main search strategies in Section 1.3.1. Next, in Section 1.3.2 we focus on the one in this thesis, direct detection, and provide the interaction rate model for collisions in our detector on Earth, which resides in the dark matter halo of the Milky Way. We conclude with a summary of the state of the art limits on direct detection searches in Section 1.3.3.

1.3.1 Methods For Observing Dark Matter Interactions and Overview of Direct Detection Experiments

A WIMP interaction with SM particles, in our case a quark inside a nucleon, can be simplified into a single Feynman diagram [39] with an unspecified interaction, as shown in Figure 1.8. There are 3 alternatives which can potentially allow for WIMP observations:

- Indirect Detection: In astronomical observations of high energy photons that are produced from the annihilation of WIMPs near the center of the galaxy where the dark matter density is expected to be largest [40]
- Collider Production: In controlled observations of high energy collisions of SM particles in large colliders which can produce dark matter particles and then study them with detectors [23, 35]
- Direct Detection: In the observation of the scattering of a WIMP particle from the halo of the Milky Way with a Standard Model particle in a sensitive detector here on

Earth (this is the subject of this thesis)

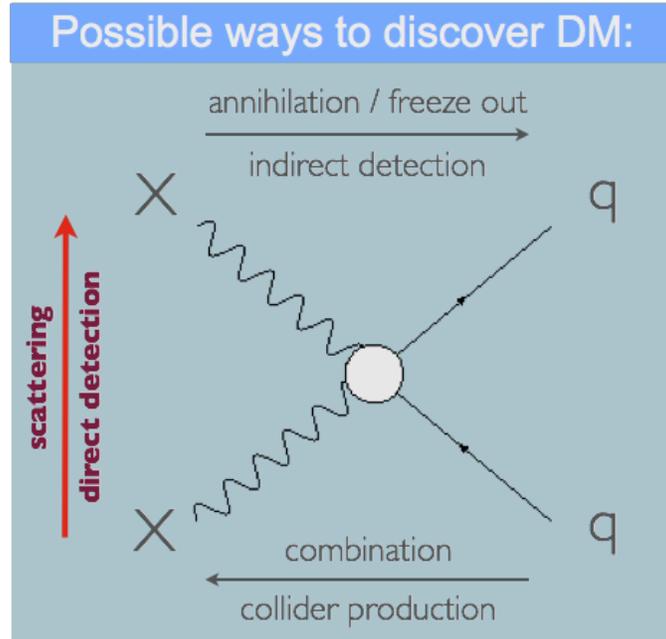


Figure 1.8: A Feynman diagram of dark matter particles interacting with SM particles with an unknown interaction/coupling. Depending on the direction of time, we can have different types of interactions (and detection methods). Annihilation occurs when time flows left to right, and dark matter annihilation creates Standard Model particles which can be observed by astronomical methods [40]. Combination occurs when time runs in the opposite direction, when Standard Model particles collide to produce dark matter which could be observed in experiments like the LHC [33]. Direct detection, the case considered in this work, is the interaction of a dark matter particle with a Standard Model particle, where both particles recoil (time flowing bottom to top).

This work is focused on direct detection methodology with the SuperCDMS Experiment. As the Earth is moving through the dark matter halo of the Milky Way [41], SuperCDMS is looking for an interaction between a WIMP and a heavy nucleus in a sensitive detector [42], as shown in Figure 1.9. More specifically, we are looking for the WIMP-recoil off a germanium nucleus in the crystal lattice of our detector.

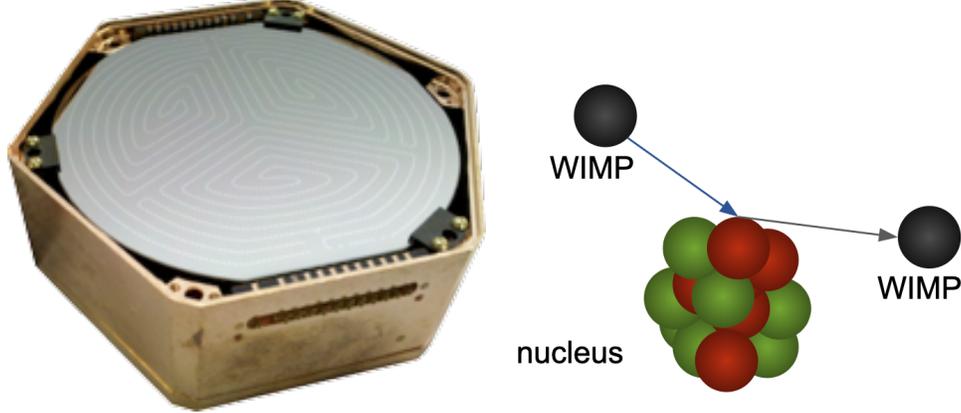


Figure 1.9: Picture of a SuperCDMS detector (left) and a diagram of a WIMP interaction with a nucleus in the germanium crystal (right).

1.3.2 Direct Detection WIMP-Nucleon Interaction Rate Model

We next derive the predictions for the expected rate of WIMP-nucleon interactions as a function of nuclear recoil energy, and then use it, with the particulars of the SuperCDMS germanium detectors and the expectations from our known location within the dark matter halo of the Milky Way, to make quantitative predictions. First we will give an overview summary of the calculation. Next we will describe the pieces needed to make the quantitative estimate and show the numerical values for several WIMP masses. Finally, we comment on the implications for our searches.

The expected WIMP-nucleon interaction rate, R , as a function of nucleon recoil energy, E_r , can be predicted as a function of the dark matter mass M_D , the target (Ge detector nuclei) mass M_T , the dark matter density ρ_D , the dark matter-nucleon cross section σ , the kinematic factor r (from momentum conservation $r = 4M_D M_T / (M_D + M_T)^2$), and the dark matter velocity distribution F . Following the calculation in Ref. [43], the WIMP-nucleon interaction rate as a function of recoil energy [43] is given by:

$$\frac{dR}{dE_r} = \frac{\sigma}{r M_D^2 M_T} \cdot F \cdot \rho_D \quad (1.1)$$

Quantitative predictions for our experiment come from inserting details of the dark matter halo estimations from models of the Milky Way [44] and our detector mass into the equation. In particular the velocity distribution of dark matter can be predicted from the observation of stellar distributions and speeds [43, 45]. Since it is expected that the dark matter velocity distribution is given by a function of the Maxwell-Boltzman distribution we take $F(v_0, v_{esc}, v_E)$ [43], where v_0 the most likely velocity, v_{esc} the galactic escape velocity, and v_E is the Earth’s velocity [45]. Astronomical measurements indicate that $v_0 \sim 300$ km/s, $v_{esc} \sim 600$ km/s, and $v_E \sim 244$ km/s [43, 45], and that $\rho_D \sim 0.3$ GeV/cm³ near our location at the Sun. Considering our detector, by setting the mass of a Germanium nucleus $M_D = 67.6$ GeV, and assuming a WIMP-nucleon cross section of 10^{-42} we find the values shown in Figure 1.10. Note that the speed of the Earth relative to the halo changes with the direction of Earth’s orbit, causing an annual modulation which can also potentially be measured [43].

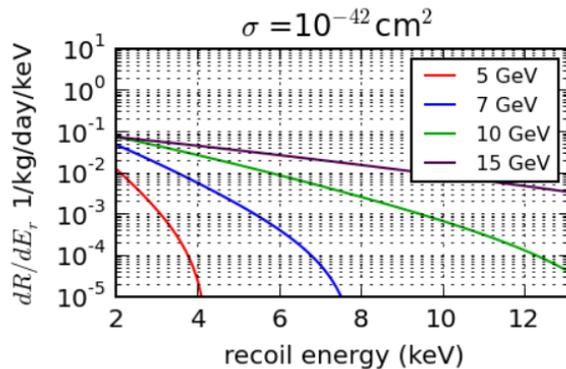


Figure 1.10: The number of expected WIMP events interactions as a function of recoil energy interacting in a SuperCDMS Ge detector, for multiple WIMP masses, assuming a dark matter halo in the Milky Way [43] and a wimp-nucleon cross-section of 10^{-42} cm^2 for simplicity. Figure reprinted from Ref. [46]. The quickly dropping rate as a function of recoil energy suggests that optimizing for low energy interactions not only improves low WIMP-mass searches, but also increases the sensitivity for higher WIMP masses, as the number of interactions drops quickly with recoil energy.

Perhaps the most important thing to notice is that most interactions transfer very little energy (give very little recoil energy to the SM particle), and the rate drops quickly as a function of recoil energy [43, 46, 47]. This is an indication that experiments need to aim for high sensitivity to low energy interactions. SuperCDMS focuses in low mass WIMP searches, by focusing on having sensitivity to interactions which deposit as little as few keV of energy.

1.3.3 Overview of Current WIMP-Nucleon Cross Section Sensitivity Limits by SuperCDMS and Others

Multiple experiments have searched for WIMPs with ever-better sensitivities to tiny WIMP-nucleon interaction cross sections [23]. Direct detection searches are focusing in the mass region between 500 keV and 1 TeV WIMP mass [23, 42]. The challenge is to observe WIMP-like interactions (call these signal events), while rejecting any other interaction that can mimic dark matter (call these background events). As shown in Figure 1.3.3 the SuperCDMS Soudan experiment, the main topic in this thesis and described in Chapter 2.1, has led the effort in low mass searches (1.5 to 10 GeV) and will be leading future searches with SuperCDMS SNOLAB (0.4 to 5 GeV) [48]. Its direct competitor, the Lux-Zeplin (LZ) experiment has leading sensitivity in the high mass search (4 to 1000 GeV) search as their predecessors (LUX and ZEPLIN) have lead before [49]. Others, including DarkSide, XENON, and PandaX complement the search, particularly in the high mass range [23]. Currently no experiment has provided compelling evidence of the existence of dark matter, but the search space keeps being pushed to smaller WIMP-nucleon interaction cross section over all masses [23]. It is worth noting that the next-generation experiments (SuperCDMS SNOLAB [48] and LZ [49]) are approaching interaction sensitivity levels that will allow for the observation of interactions from neutrinos radiated from the Sun (solar neutrinos). This will allow for exciting measurements, but will also become a challenging background to discriminate against.

While state of the art efforts are improving the technology, and analysis techniques allow the detectors to be measure smaller energy deposits, this also means that they accept

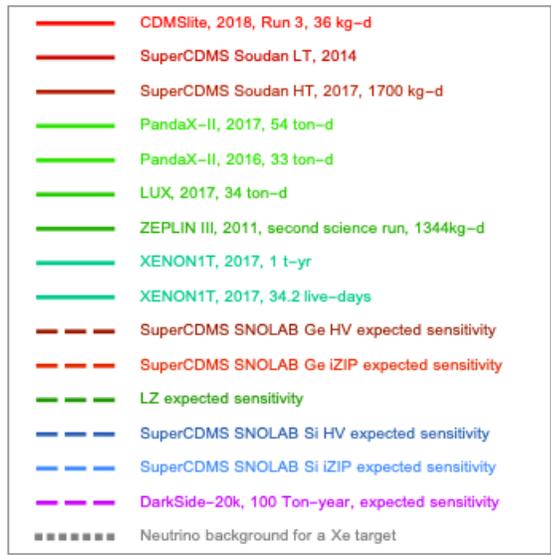
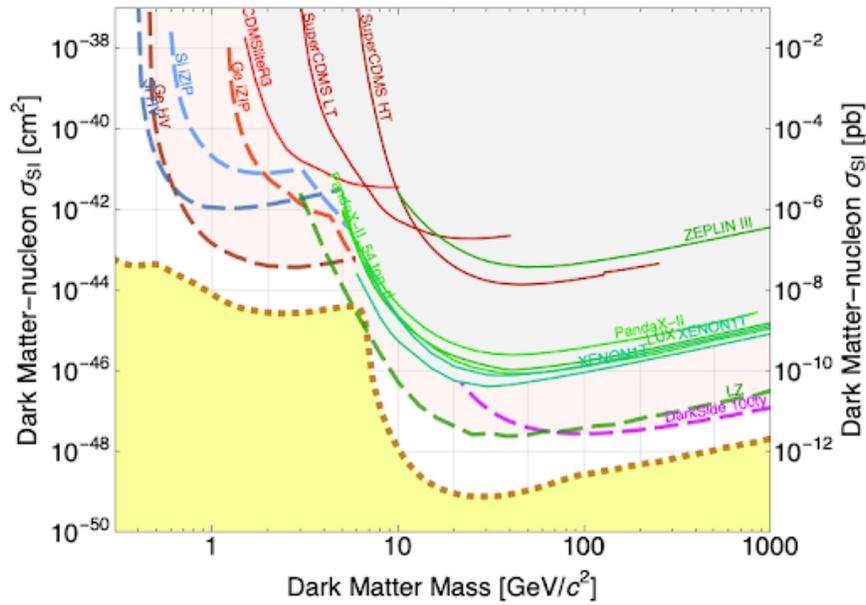


Figure 1.11: This figure shows the regions of dark matter-to-nucleon cross section limits vs mass which have been excluded by multiple direct detection experiments [23] Figure generated with the SuperCDMS limit plotter tool [50]. The area above and to the right of the solid lines has already been excluded. The dashed lines and pink space represents the expected sensitivity of near-future searches. The dotted brown and yellow space represents the solar neutrino background, which will impose a difficult background in the future.

more sources of ‘background’. In this context next-generation experiments require a better understanding of the background and signal response in their detectors and new tools, such as simulations, to do so.

1.4 Overview and Outline of this Thesis

We now have the motivation for the search for dark matter, a hypothesis to observe it in a direct detection experiment, and an experiment which has completed highly sensitive searches. With this in mind, this thesis is about the simulations of dark matter and other SM interactions in the SuperCDMS Experiment with an eye towards using them for optimizing a WIMP search to improve the sensitivity moving forward. Our goal is to provide the tools to enable simulation-based analyses to optimize dark matter searches for current (SuperCDMS Soudan) and upcoming (SuperCDMS SNOLAB) experiments expected to start taking data in 2020. As the SuperCDMS Experiment is transitioning to better technology, and since WIMPs have not been observed, the next-generation analyses and detectors will have to discriminate ever more challenging backgrounds and will require even more robust analysis techniques. As the experiments become more sensitive to lower recoil energies, the backgrounds become harder to mitigate and need to be discriminated against in analysis, which will be particularly complicated as we become sensitive to neutrinos and other sources of low energy interactions. As simulations have played a key role in the recent top quark [9] and Higgs boson [10] discoveries, the time has come to bring modern simulation techniques to SuperCDMS.

To make it easier to see where we are going, we will summarize what we have done. The primary task of this work was to create a full simulation package and show that it works as expected. While we have made a full simulation infrastructure, and have fully simulated events, the entire endeavor is enormous and will require many more years with contributions from lots of individuals. While we have completed many of the most important first steps, we have not completed the studies of how well all of the parts work, nor finished the multi-year process of studying and tuning them, nor using them in an optimized search. This work

focuses on what we have done: the simulation of the noise, charge, and phonon systems, and studies of the noise and charge measurement, comparing with data and expectations about the detector's response in a sample of calibrations where the inputs are well understood. We will show what we have learned, and point the way forward for the future.

Chapter 2 contains a brief description of the SuperCDMS Soudan Experiment, including a description of the backgrounds and the apparatus design to mitigate against them. We follow with Chapter 3 where we describe the physics of the detectors which allow us to measure the energies of any interactions as well as both distinguish between dark matter and background events. Chapter 4 describes the most relevant data analysis techniques to turn detector readout into energy observables, the optimal filter, which is used to extract the best-estimate energy measurement of an event which is contaminated by noise from the readout electronics. Chapter 5 is about going from the energy estimates to understanding and studying the response, and the common analyses techniques and methods in dark matter searches with SuperCDMS to separate between signal and background events. Chapter 6 describes the known causes of mismeasurement due to readout malfunction and the methods used to correct them along with the treatment to clean them from our noise data prior to the simulations comparisons. Chapter 7 describes the simulations infrastructure and the individual pieces of the simulation which take us from a description of the physics to fully reconstructed samples of particle interaction events in the detector. The results of comparisons to data are described in Chapters 8 and 9. Chapter 8 is about the validation of the noise simulations, and the measurement of the contributions of noise to the overall energy resolution. Chapter 9 is about the quantitative and qualitative comparisons of the simulations to data, and showing what the simulation tells about well-behaved selection criteria. Chapter 10 is about the simulation-based predictions, by measuring detector accuracy and resolution and looking in more detail at the causes of energy mismeasurement with an eye towards the next steps for development. Finally, Chapter 11 provides an overview about optimizing future analyses with our simulations, provides some clear examples of studies to

be followed up on, and concludes this work.

2. THE SUPERCDMS SOUDAN EXPERIMENT

With the last chapter providing evidence that dark matter exists and an outline of our plan to discover it via direct detection, we next describe the Super Cryogenic Dark Matter Search Soudan experiment (SuperCDMS Soudan) used in the search. The apparatus consists of a set of germanium and silicon detectors (shown in Figure 1.9) that are good at separating between low energy WIMP interactions with a single nucleus and SM particle interactions with an electron in the detector, and surrounded by an infrastructure that is designed to reduce the number of SM particles that even enter the detector [42, 47, 51]. The goal is to obtain a clean sample of WIMP-like events (call this signal). This means having a good balance between being able to efficiently accept signal events, while discriminating against any interactions of SM particles from contaminants in the environment or poorly measured background events that mimic signal events (call both of these background). This is particularly difficult at low energies where the signal-to-noise ratio for interactions is the poorest. We will describe the backgrounds and how the experimental apparatus mitigates against them, and in the next chapters we will describe how the particle interactions produce the readout from the effects on the crystal and the dedicated sensors with an eye towards how we will simulate or reject them.

In this chapter we provide an overview of how the SuperCDMS Soudan experiment is designed, with an emphasis on how signal and background sources interact with the detectors so we can accept signal and reject backgrounds. Since the design includes a number of ways for keeping backgrounds from entering the detector, we first give an overview of the background sources in Section 2.1, with an eye towards how they will be reduced. Next, in Section 2.2 we describe the apparatus design and the detector components. And finally, in Section 2.3 we expand on the considerations that mitigate the backgrounds either by blocking, shielding or lowering the rates, or at the analysis level.

2.1 Overview of Background Sources and How They Interact in the Detectors

The SuperCDMS detectors are highly sensitive to many types of particle interactions and while the goal is to observe WIMP interactions, other sources of particle interactions are a problem. The list of background sources is given in Table 2.1 and each item will be discussed in more detail soon. As described in later sections, the detectors and the apparatus that surrounds them are designed to reduce how often these background particles enter the devices themselves, and distinguish between the ones that do enter from signal events. The background sources can come from near the detectors, within the detectors themselves, as well as from other cosmic sources (WIMPs are not the only cosmic particles hitting the Earth), and while most of the backgrounds are blocked (some to a point of becoming negligible), others need to be identified and rejected in the WIMP search analysis.

In general, particle interactions are expected to scatter with either an electron or a nucleus in the detector as shown in Figure 2.1. We will refer to these two types as electron recoils and nuclear recoils respectively. The interaction type depends on the incoming particle: electrons and photons would interact with the atomic electrons that surround the nucleus, but neutrons and WIMPs would interact with the nucleus itself [8, 42, 45]. While the total recoil energy is given by the kinematics of the interaction, each type of interaction releases two types of energy into the crystal of the detector: charge energy (from ionization of electrons in the crystal), and phonon energy (from vibrations of nuclei in the lattice), as shown in Figure 2.2 [42]. While the mechanism of energy release for each type, and how much goes into each, will be described in Chapter 3.1 we note for now that the reconstructed energy measurement of charge and phonon signals differ significantly between the two and allow us to discriminate between an electron or nuclear recoil type event. As shown in Figure 2.2, both types of recoils have measured energies that rise as a function of energy, while the constant of proportionality is different. Thus, the ionization yield, defined as the ratio of the measured charge and phonon recoil energy, creates a clearly defined method for identifying background-like (electron) from signal-like (nuclear) interactions.

Cosmogenic Backgrounds		
Background Source	Recoil Type	Mitigation Method
Electrons and Photons	Electron recoil	Mostly shielded otherwise discriminated, also referred to as Compton background (because of the scattering process)
Muons	Electron recoil	Shielded or vetoed by the apparatus to negligible rates
Neutrons	Nuclear recoil	Shielding reduces these to negligible rates
Neutrinos	Nuclear recoil	Interactions rates and energy deposits are so small that they can be neglected for Soudan

Contaminants In or Near the Detectors		
Background Source	Recoil Type	Mitigation Method
Ge Activation	Electron recoil	Discriminated against in analysis
Pb Implantation	Electron and Nuclear recoil	Discriminated against in analysis

Mismeasured Events		
Background Source	Recoil Type	Mitigation Method
Cross-talk	None	Considered in event reconstruction
Pileup Events	None	Discriminated against in analysis
Blips	None	Discriminated against in analysis
Glitches and Low Frequency Noise	None	Discriminated against in analysis

Table 2.1: Summary of the background sources, what type of interaction they have with the detector (recoil type), and method for dealing with them. Nuclear recoils from Pb sources are the most affecting background since they mimic WIMP events.

The most important analysis challenge is for the lowest energy interactions where signal rates are expected to be highest but where mismeasurement occurs, resolution worsens, and it becomes difficult to separate between the two types of interaction responses. We will get back to how these quantities are reconstructed in Chapter 4, and to the samples used to understand the response in Chapter 5. Ultimately it is our goal to have simulations help in this process. For now we give an overview of the experimental apparatus and then expand on how it reduces the number of particles that enter the detectors to reduce the overall number

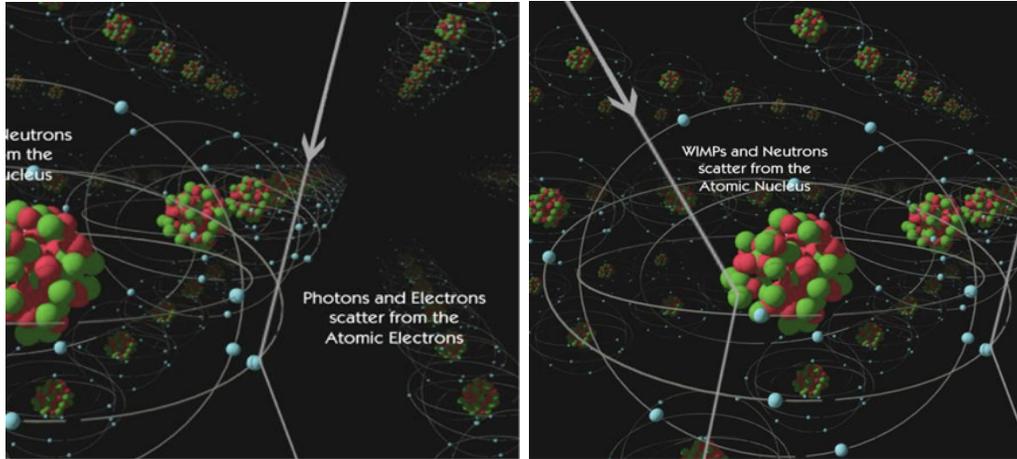


Figure 2.1: A diagram of how particles interact with the electrons (left) and the nucleus (right) in the crystal. Because the energy is deposited based on the amount of energy given to the electron or nucleus, we will refer to these types of interactions as electron recoils and nuclear recoils respectively. Electron recoils are scatters with electrons in the crystal, knocking the electrons out of the atoms and liberating electron-hole pairs. Nuclear recoils are interactions with nuclei in the crystal, inducing mechanical vibrations propagating as phonons. WIMP interactions are expected to be nuclear recoils, while for our purposes, electron recoils are considered background interactions.

of background events.

Particle	Primary Interaction	Ionization Energy E_Q	Primary Phonon Energy $E_{Ph, prim}$	Charge-Released Phonon Energy $E_{Ph, NTL}$
Photon	Electron Recoil			
Electron	Electron Recoil			
Neutron	Nuclear Recoil			
WIMP	Nuclear Recoil			

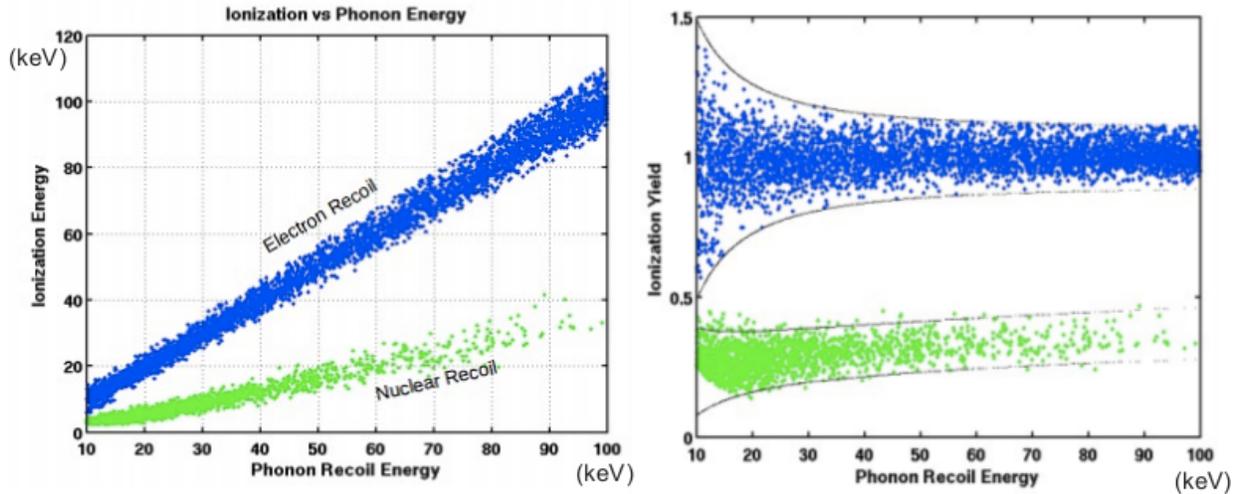


Figure 2.2: Top: Diagram of energy distributions in the charge and phonon signal depending on particle interaction type. As electrons and holes propagate through the crystal and pick up energy in the electric field, that energy is deposited into phonon energy, denoted as $E_{Ph, NTL}$, where NTL stands for the name of the process, Neganov-Trofimov-Luke gain. Note that the total amount of collected energy for an electron recoil event is larger for a nuclear recoil event, thus charge energy for photons and electrons dominates. For neutron and WIMP interactions most of the energy is released into the phonon system, but some fraction is deposited into the liberation of e/h pairs (which causes additional phonon NTL energy). Bottom: Pair of plots that show the amount of energy recorded in the ionization/charge processes as function of phonon recoil energy (left), and the ionization yield as a function of the amount of energy recorded from phonon recoil energy (right). In both plots electron recoils (^{133}Ba calibration in blue) and nuclear recoils (^{252}Cf calibration in green) are clearly separated in the ionization yield quantity, but the challenge is to separate at low energies where they overlap.

2.2 Overview of the Experimental Apparatus and Detectors

The experiment, shown in Figure 2.3 consists of fifteen detectors, arranged in five towers of three detectors each, and placed within an experimental infrastructure deep underground that is designed to reduce the rate at which SM backgrounds enter the detector, ideally to negligible levels. For this reason we use two different types of shielding: the Earth's crust to reduce the rates of sources of high energy particles from space (cosmogenic sources), and local shielding that directly surrounds the detector to reduce the backgrounds which are produced near the detectors. The events that do interact with the detector leave phonon and charge energy deposits that are read-out by dedicated sensors.

In this section we will describe the surrounding apparatus that supports the detectors while reducing the SM interactions, and how it is equipped to place radioactive sources to study the detectors response to known particle interactions (Section 2.2.1). Next, we will describe the detectors and the sensors that measure the charge and phonon energy, with an eye towards analyses and background discrimination (Section 2.2.2).

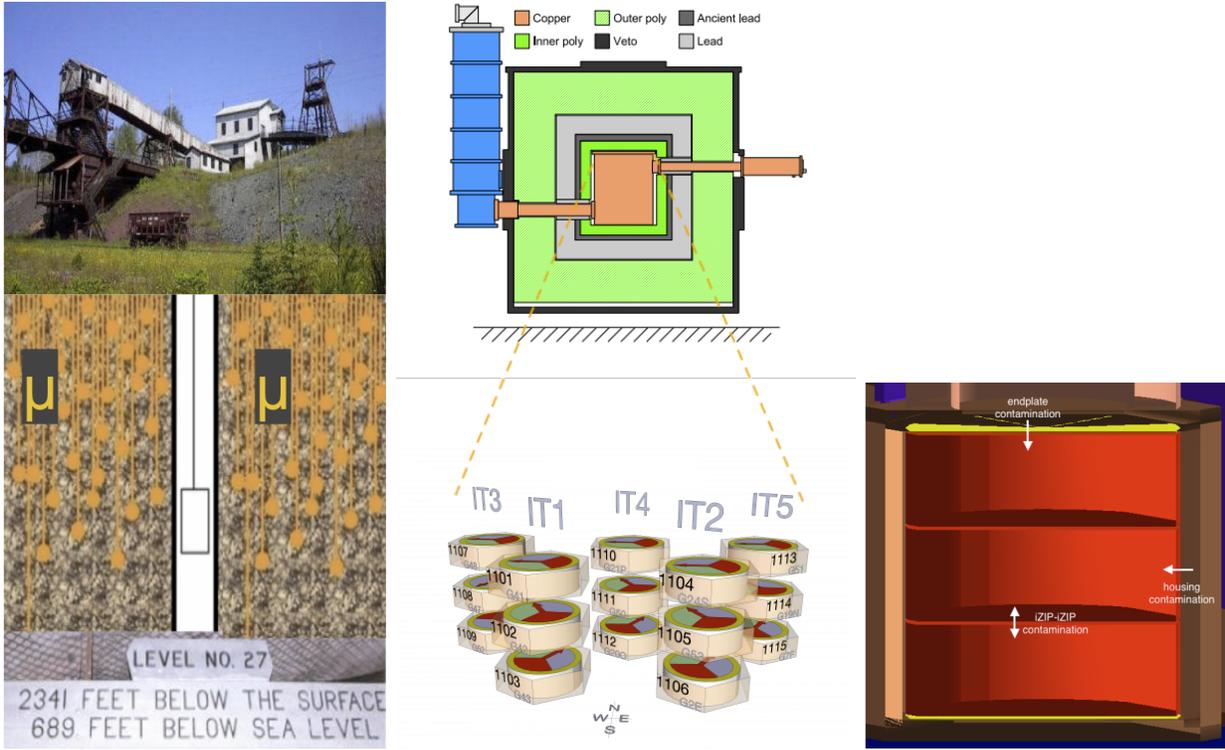


Figure 2.3: A set of diagrams showing the SuperCDMS Soudan detectors, the apparatus surrounding them, and its location which is deep below the surface of the Earth in the Soudan Underground Mine. The experiment is underground to block from particles incoming from space, and the apparatus is designed to reduce the rate at which the remaining SM particles enter the detector. The top middle figure (reprinted from Ref. [52]) shows the apparatus shielding of the detectors by both lead and dense plastic material to prevent from other sources of electron and nuclear recoils (both from contaminants in the environment and from remnant cosmic particles). The right figure (reprinted from Ref. [53]) shows the shielding of each tower of detectors. The apparatus consists of 15 detectors: 5 towers, with 3 detectors each, shown in the bottom middle figure. This work is focused in comparing simulations with detector 4 (labelled 1104 in the figure), the germanium, middle detector in tower T2, as it is known to be the most sensitive detector

2.2.1 Apparatus Description

The detectors and the apparatus that surrounds them are located underground at 2341 ft below the surface (689 feet below sea level), at the Soudan Underground Mine in the Vermilion Range of Minnesota. We will describe the apparatus and the physical components, with an eye towards their purpose to mitigate the backgrounds. We begin by explaining how the

experiment is underground to block from space-bound backgrounds. We then describe how the experiment is equipped to place radioactive sources to observe its response to well-known interactions and set the measurement expectations. We follow by describing the detectors structure, prior to describing the sensors in the next section.

The experiment is located underground to use the Earth's crust to block particles coming from space (mostly muons), known as Cosmic Rays, that could directly fake, or produce other particles such as neutrons which could fake a WIMP interaction [54]. While most of this background is mitigated by placing the detector underground, some of the particles from cosmogenic sources do reach the detector, so muons that are not blocked are identified using a Veto system, which tags them and rejects the recorded event [42].

The apparatus is designed for ready placement of radioactive sources that produce photons, electrons and neutrons of known energies near the detectors which assist in understanding the detector response to each. In particular we use ^{133}Ba and ^{252}Cf [47, 53] as sources of such particle interactions, as they have specific energies from nuclear decay chains. We will refer to each as a calibration source and they will be described in more detail in Chapter 5. Data is taken when these sources are near the detector which can be used to make samples of well-understood events called calibration samples. Since many of the expected signatures from the emitted particles have known energies the detector response and measurements for both electron and nuclear recoil interactions can be studied and used. For example they are used to help determine the good-quality event selection criteria as well as the observed ratio between charge and phonon energies for electron and nuclear recoils to discriminate between signal and background, as seen in Figure 2.2. With these capabilities in mind, we will later describe how the backgrounds are mitigated by the experiment.

Each detector is enclosed in a hexagonal prism copper housing and the detectors are surrounded by multiple layers of lead and dense plastic (poly) (see top right of Figure 2.3) to prevent other types of radioactive sources that produce high energy particles in decays from entering into the detectors [42]. The detector itself is built from a single cylindrical

crystal of either germanium or silicon with the size of about a hockey-puck and weighing about 0.6 kg.

2.2.2 Detector and Sensor Description

Each detector has a charge and phonon readout system that collects the charge and phonon signals from energy deposits from a particle interaction and turns them into a measurement [42], as shown in Figure 1.9. We will first overview the charge and phonon sensors location and detector configuration. Next we will detail the charge readout components, followed by the phonon's. We then describe the charge and phonon readout in turn. This will leave us well set to describe the background mitigation in the next section.

The detectors have both charge and phonon sensors, at the top and bottom surfaces, operating at superconducting temperatures, where the signals are collected, as seen in Figure 2.4. Charge electrodes (red rectangles in the figure) at the top (+2V) and at the bottom (-2V) are the potential source of the electric field in the crystal, and the inter-spaced phonon circuit runs with ground lines (blue rectangles in the figure), creating an additional effect in the electric field near the top/bottom which is designed to help identify interactions which occur near the top and bottom and are known to be poorly measured (described in more detail in Chapter 3.3). The electric field is uniform in the bulk of the detector, while near the surfaces it changes periodically creating an inter-leaved field (also shown in the same figure), thus the response has a vertical-position dependence (labelled the Z-direction). For these reasons the SuperCDMS Soudan detectors are called Interleaved, Z-dependent Ionization and Phonon Detectors or iZips for short.

The charge readout on each side is performed by the electrode lines, as shown in Figure 2.5. The voltage bias (top to bottom) attracts the electron/holes uniformly in the bulk (electrons to the top and holes to the bottom), but when near the surfaces they get trapped in the interleaved potential. The collection process will be described in detail in Chapter 3.3. The charge collection elements are commonly referred to as FET sensors, even though inaccurate (because the sensing device is simply the electrode), it is because the amplifier circuit

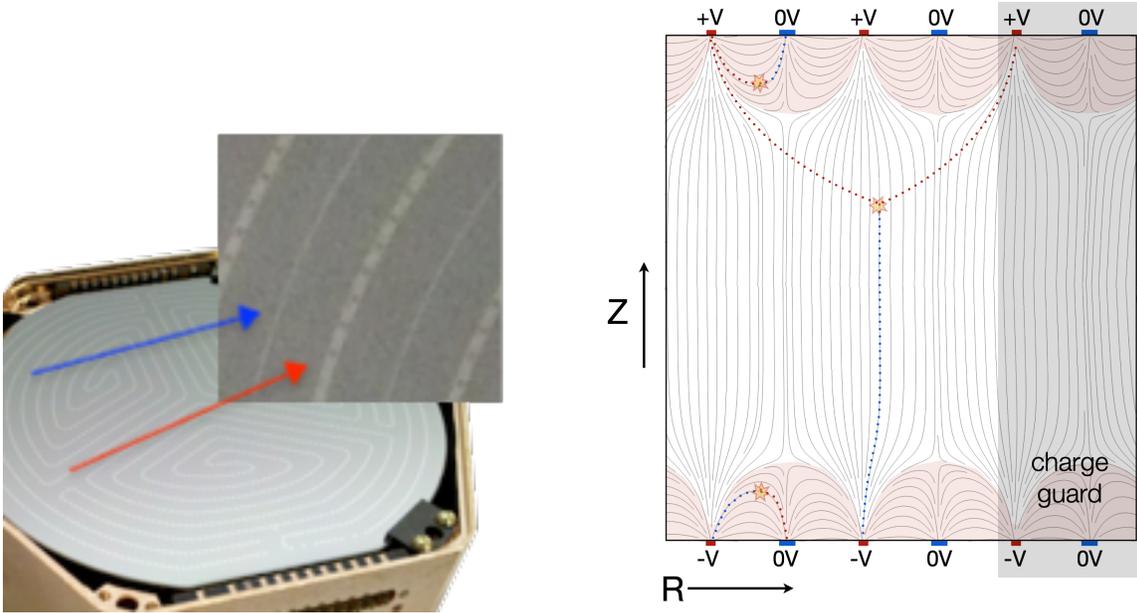


Figure 2.4: Left: Zoomed view of the electrodes and phonon sensor arrays, running interleaved along the surface of the detector. Right: Cross section diagram of a portion of the electric field in the crystal created by the arrays, reprinted from Ref. [53], showing the location of the electrodes on the surface of the crystal as well as electric field lines, which in turn determine the location of the collection of electrons and holes. In red and blue rectangles are the interleaved bias voltage values.

uses a junction Field-Effect Transistor [42, 55]. The individual charge collection response is combined into two channels on each side, denoted inner and outer [42], given their radial partition, as seen in Figure 2.5. In principle, the combination of the response on each side provides an independent measurement of charge energy while in practice the two can be very different if the interaction occurs near the top or bottom of the detector, allowing us to readily remove the near-surface Z position interactions. The inner and outer portions share the deposited energy, and provide a measurement of the radial location of the interaction which is useful for removing events that occur near the radial edges, which are also known to not have their energy well-measured.

The phonon energy is read and recorded using a particular type of technology referred to as a set of Transitioning Edge Sensors (TES), which are composed of small sensitive cells, arrayed in series to construct the sensing channels [42] at the top and bottom of the detectors

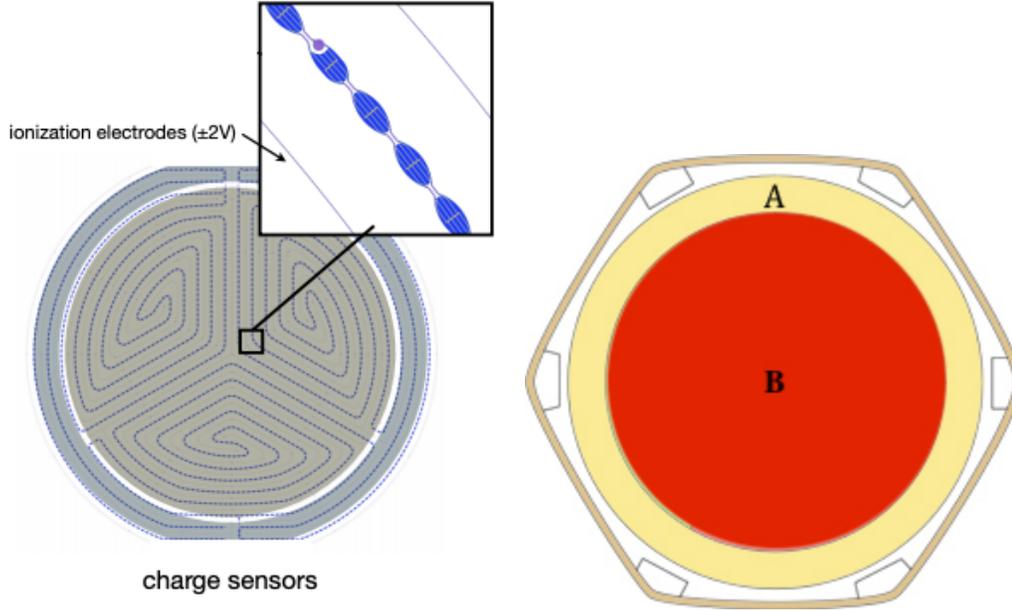


Figure 2.5: Two diagrammatic views of one side of the charge readout system. On the left, the charge readout is composed of electrode lines at ± 2 V, interleaved with the phonon sensor lines at 0 V, figure reprinted from Ref. [53]. On the right, the detector channel map where the inner electrodes are combined into a single channel, denoted B in the figure, and the outer electrodes are combined and denoted as A. The top and bottom systems provide an independent measurement of charge energy. The four separate measurements can be used help determine if the primary charge energy is fully reconstructed, as well as provide information about the location of the primary interaction.

which ‘collect’ the phonons. The location of the TES sensors are as shown in Figure 2.6, and the results from each are summed from the individual TES cells and collected into one of 4 channels on each side, where each side provides an independent measurement of the phonon energy [42]. As shown in the figure, one channel is an outer ring and the other three are in equipartitioned fractions of the inner circle. Each TES acts as a resistor which changes its resistivity as phonons hit the unit [42], and the channel is the added effect of many of them in series. The higher resistance allows charge to flow, and the total collected charge is a measure of the total deposited phonon energy. A broader description of the individual TES cell functioning will be described later in Section 3.4.

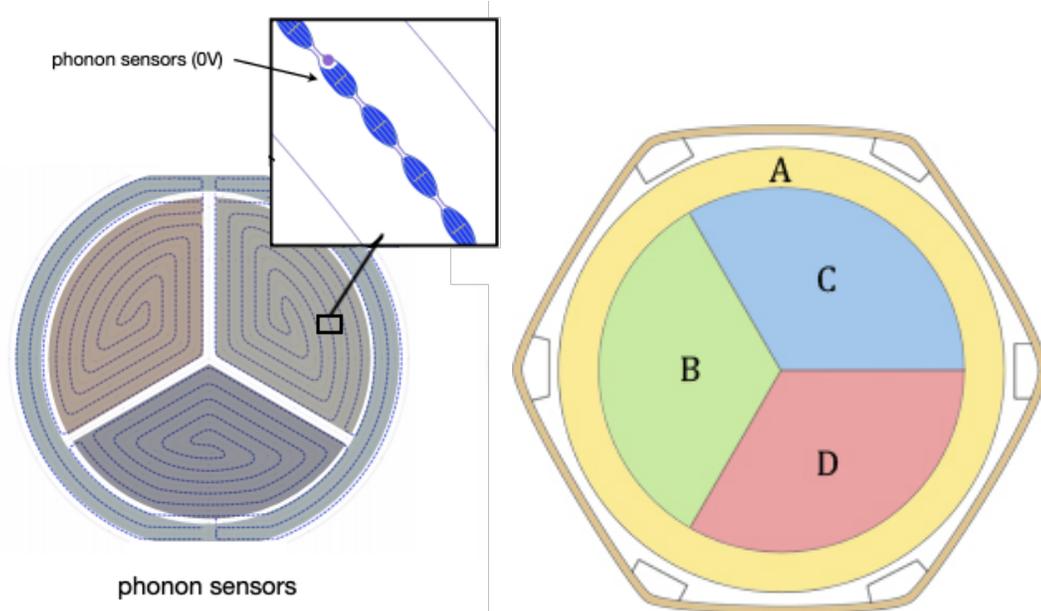


Figure 2.6: The lines containing the TES collection elements (left) and the combined detector phonon readout channel map (right), on both surfaces of the detector. Figure on the left reprinted from Ref. [53] and figure on the right from Ref. [56]. Each channel is an array of collection elements which act as variable resistors to produce the phonon signal.

2.3 Description of How the Backgrounds Are Mitigated by the Experiment

We now describe how the experiments reduces or treats each of the backgrounds. We follow the list from Table 2.1 and expand on how the background interacts with the detector as well as the mitigation method. Each will be described in turn.

In the upcoming sections we start with a description of the cosmogenic sources (those from cosmic particles), beginning with electrons and photons produced by their interactions, followed by muons, neutrons and neutrinos. Next, we will describe the backgrounds from contaminants in the detectors (or the apparatus), which are from germanium activation (in the crystal of the detector), and lead decay from the implantation isotopes from the radon decay chain that surround the detectors. Finally we will describe event mismeasurement, which includes contamination of the signal from other channels within the detectors (cross-talk), multiple events occurring within the same readout time period (pileups), and data acquisition (DAQ) malfunction cases (Blips, Glitches, and Low Frequency Noise).

2.3.1 Cosmogenic Sources Backgrounds

We begin with cosmogenic sources. Most of these sources come from by-products of the high energy particles emitted from the Sun and interacting with the atmosphere, the Earth's crust, and the experimental material. We will describe them in order of dominance on the rate at which events are observed in the detectors. We begin with photons and electrons, by-products that produce electron recoils. We follow with muons, which are removed due to their strong electromagnetic interaction that can be identified before reaching the detectors. We next describe neutrons, from scatters produced by muons in the surrounding material. We finish with neutrinos, a background that is currently negligible but as future experiments become more sensitive to such low energy deposits will become an important background.

After the shielding from the Earth's crust has done its work, the dominant background to the experiment is from photons and electrons which are by-products of interactions from cosmic origin [47]. The lead shielding material is expected to block most cosmic inbound events, but some still get through. Photons and electrons create electron recoils typically via Compton scattering [47, 51, 53], either directly interacting with the detector, or as a by-product of another interaction prior to reaching the detector. These are typically rejected as part of the electron recoil rejection methods in analysis.

While the second most common background is from muons, a particular type of high energy cosmic particle [23], after the methods described above, the contribution of this background to the typical search is negligible. Due to its location deep underground most of the muons do not reach the detector [54] and the ones that do will interact electromagnetically outside the detector in the Veto system, so they are easy to observe before they enter. After rejecting events that occur in coincidence with a hit in the veto system, the background rates from muons is negligible, even further, any muons that do enter the detector without being identified by the veto system would be rejected in the analysis as they have an electron recoil signature.

The third cosmic background is from cosmic muon-induced neutrons which have a nuclear

recoil in the detector, but are also expected to be negligible due to the shielding. Between being deep underground and the detector shielding, the rate at which cosmogenic-produced neutrons enter the detector is reduced to a point where this background can also be safely ignored [42, 54, 47].

Neutrinos from the Sun typically pass through Earth's crust, as well as the lead and poly shields, but given their low interaction cross section rate with the nuclei and electrons in our detectors they almost never deposit enough energy to produce a measurable event so this background is also neglected [42, 47]. The next-to-next generation detectors are expected to be sensitive to such low energy interactions such that, if dark matter is not observed before then, this background is likely to dominate. On the other hand it may provide a fruitful area for future physics study.

2.3.2 Contaminants Near the Detectors

The next background category is from contaminants in the environment surrounding the apparatus, in the apparatus material, or in the detector itself. Only two sources are relevant in this case. The first one is germanium activation, an electron recoil background which is controllable because it occurs after exposure of our ^{252}Cf . The second is the most damaging and complicated to reject, signal-like nuclear recoils from uranium decays from contaminants in the apparatus material. We will describe each in turn.

Germanium activation is the first and occurs after ^{252}Cf calibration source exposure which excite the nuclei (with an 11 day half life and a known energy spectrum), releasing photons and electrons, thus producing electron recoils [47, 51] within the detector itself. Since the interaction type is electromagnetic they are rejected in analysis, unless they are mismeasured. On the other hand a side-benefit is that these known backgrounds can be used for studies because of their known energies [47, 51].

The most damaging backgrounds are contaminants from the radon/uranium decay chain of lead atoms (see Figure 2.7) which implant either near or in the apparatus. The right hand side of Figure 2.3 shows the location of where implantation is most difficult to reject

against. The end-product can cause either an electron or a (more problematic) nuclear recoil from an α emission [47, 51, 53]. This background tends to be more difficult to reduce since the material inside the shielding layers can be contaminated, so these events do reach the detector, cause energy deposition, and have to be properly identified during an analysis. While the most effective way to mitigate this background is to reduce the contamination levels during the fabrication stage of the apparatus, this ends up as a dominant background which needs to be treated at the analysis level (particularly the α particles which have the WIMP-signature nuclear recoil) [47, 51]. Simulations will be a key component for discriminating these backgrounds since contaminants creating nuclear recoils are signal-like depositions and have to be identified with more sophisticated analysis techniques [47] or the rate at which they produce events that pass all the final selection criteria precisely estimated.

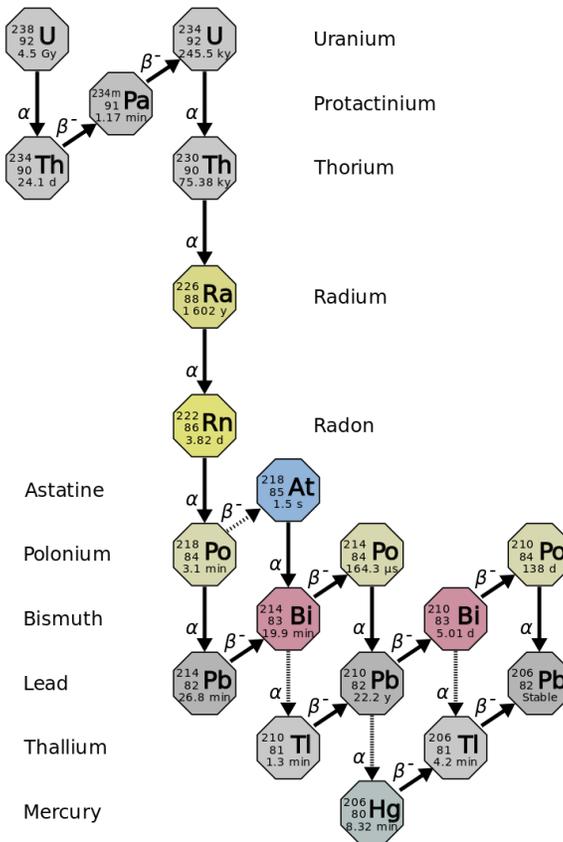


Figure 2.7: This figure (reprinted from Ref. [57]) shows the multistage decays of Uranium which, if within the shielding around the detectors, can decay and cause difficult-to-remove background events in our experiment. The final states (^{210}Bi , ^{210}Pb , and ^{206}Pb) can produce α particles (^4He nuclei) and electrons, which can produce nuclear and electron recoils in the crystal. The hexagon shows the nuclear states, and the notation above the arrow shows the particle created in the transition, α and β emission.

2.3.3 Energy Mismeasurement

Having described all the backgrounds due to real particle sources, the last background type to consider is energy mismeasurement of a SM particle which has entered the detector. We will describe the causes of this background in more detail in Chapter 6, but in this section we will briefly describe them to complete our summary of all the backgrounds listed in Table 2.1. These backgrounds are relevant because if the amount of observed charge energy is lower than expected, an electron recoil event can be misidentified as a nuclear

recoil event, as shown in Figure 2.2. This is most problematic at the lowest energies, where noise already dominates and affects our ability to make an accurate energy measurement. While there are a number of these types of problems (and we will describe where each gets its name in a moment), most can be handled using analyses techniques that we will describe. We begin with a description of cross-talk, then we will describe pileups, we will describe Blips next, and finally we will describe Glitches and Low Frequency noise.

The first of these backgrounds is undesired electronics communication in-between channel measurements, in a process known as cross-talk (described in more detail in Chapter 6). Since this background is well understood and readily measured it is handled by the event reconstruction algorithm described in Chapter 4.

The second background is when for example, two interactions occur so close in time that they both show up in the same readout time window. These are known as pileup events. This is a problematic since most analysis techniques assume a single interaction in the time window. As we will see in Chapter 6, these are readily handled by rejecting events where there is evidence of two interactions separated in time.

The third background is Blips (and TowerBlips) which are a particular type of electronics noise occurrence where a pulse shape that has both a negative and positive contributions in time can randomly enter the readout time window. Depending on when it arrives, it can either raise or lower the measured energy. While they are not a problem for large energy interactions, they are potentially problematic if they occur near the time of a low energy interaction, and they are an important background in our ability to measure the resolution of the detector in the absence of a signal. They typically occur in more than one detector of the same tower and are handled by a TowerBlip correction method. These will be described in detail in Chapter 6.4.

Finally, Glitches and Low Frequency Noise are readout problems that get their name from readout electronics malfunctioning: Glitches show up as a highly discontinuous feature in the readout, and Low Frequency noise as wave-like features with a wavelength similar to half

of the readout length (and are sometimes coupled with mechanical vibrations). Since these signals have a readily identifiable and particular pulse shape, the reconstruction algorithms allow for discrimination with goodness-of-fit methods. They will be described later but not in much detail in this work, since they are well understood and do not show up in the samples we will study.

With our overview of the backgrounds and the methods to mitigate each complete, we are ready to turn to the next chapter. In it we describe how the experiment is expected to respond to particles that interact in the detector, as we prepare for understanding the response and use them for the discriminating technique for a WIMP-search analysis in SuperCDMS.

3. DETECTOR RESPONSE TO PARTICLE INTERACTIONS

Now that we have described the apparatus, the expected signal and the backgrounds, we are ready to describe the detector response, from the interaction in the crystal up until the waveform produced by the sensors. For particles that do interact with the detectors, the discrimination power comes from the ratio of the two different energy measurements, in particular because the two types of interactions, electron or nuclear recoils, produce different responses, either by ionizing electrons (charge system) or creating vibrations in the crystal lattice (phonon system). We will describe how the detectors respond to these energy depositions, and how the response is turned into a readout measurement. In the next chapter we describe how we turn the waveform readouts into best-energy measurements, with an eye towards performing a dark matter search in Chapter 5 and eventually their simulation in Chapter 7.

In this chapter we describe the full set of steps that turn a particle interaction into the detector electronics readout. There are many steps involved which we list below:

- The energy deposition mechanisms in the crystal
- The primary phonon and electron-hole pair creation
- How the charges and phonons propagate in the crystal
- How the phonons from the primary interaction, primary phonons, down-convert to secondary interactions
- How the charge propagation produces additional secondary phonons which contribute to the total energy
- How both charges and phonons are readout by their respective sensor systems

We will first describe the overall energy response to electron and nuclear recoils in Section 3.1, how an electron or nuclear recoil in the crystal produces the primary energy into the

phonon and charge parts of the detector at the interaction point (Section 3.1.1), then how we estimate the number of carriers (phonons and electron-hole pairs) produced by such primary energy (Section 3.1.2), and finally how secondary phonons (produced from electron/hole motion) contribute to the secondary energy (Section 3.1.3). Next, in Section 3.2 we describe details on how both propagate through the crystal including the motion of charges (Section 3.2.1), the emission of phonons from the motion of charges (Section 3.2.2), and the motion of phonons (Section 3.2.3), all of which are needed for understanding the readout as well as used in the simulations implementation. Finally we describe the non-trivial way that the two different readout sensors collect the energy and produce waveforms, for charge in Section 3.3 and phonon in Section 3.4. Since the topic of this thesis is on charge, we provide more detail on the charge collection first (Section 3.3.1) and then the waveform construction by the readout electronics (Section 3.3.2).

3.1 Energy Deposition in the Crystal

In every particle interaction in the detector, as shown in Figure 2.1, energy is deposited into the crystal. In this section we overview how the energy of the incident particle is converted into the charge and phonon signatures. We first describe the primary and secondary processes that contribute to the energy. This is followed by a summary of the propagation of charges and phonons and the conversion of the energy the charges pick up in the electric field into phonon energy. Finally we provide the quantitative estimates of the total energy contributions to both phonon and charge energy. This will leave us well set for the detailed description of each of the contributions in the upcoming sections.

Energy is deposited into the crystal charge and phonon systems in two processes: from the incident particle interaction and as charges are accelerated in the crystal lattice. We call the energy deposited into the crystal from the incident particle the primary or recoil energy, and it is partially deposited into crystal lattice vibrations (phonons) and partially into ionizing electron-hole pairs, with a fraction determined by whether it is an electron or nuclear recoil. A secondary source of energy in the phonon system comes from the bias voltage across the

detector as charges drift at high speeds in the crystal. As shown in Figure 3.1, after the initial interaction electrons and holes are continuously accelerated by the electric field and interact with the lattice structure, in a process called Neganov-Trofimov-Luke (NTL) Gain [58, 59], to produce what we call NTL phonons.

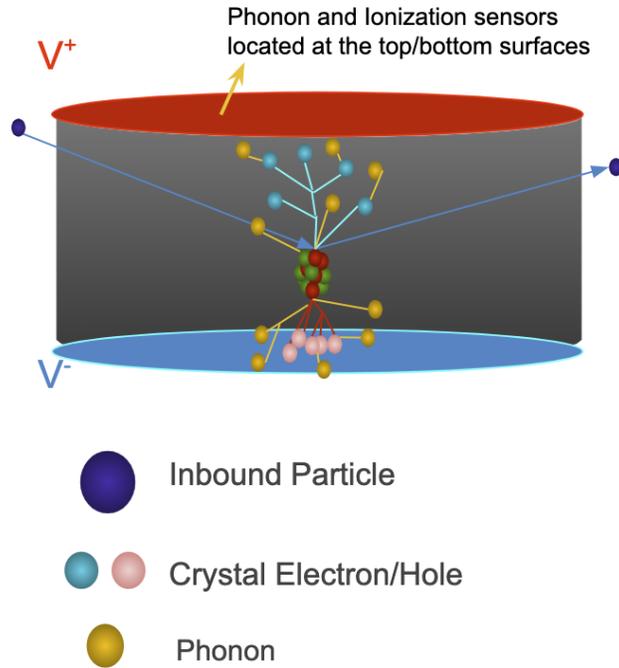


Figure 3.1: A diagram of an incoming particle interacting with the detector, creating both primary phonon and electron/hole interactions which then propagate throughout the detector. Electron-hole pairs are created from the effects of the initial interaction in the crystal, either by electron recoils or from by-products of a nuclear recoil. Primary phonons are created from a nuclear recoil interaction which quickly convert into secondary photons. From the interactions of electrons and holes in the crystal additional energy is added to the system, as they are accelerated by the bias voltage, however all of this energy is released into the phonon energy.

Before going into detail, we summarize the propagation and creation of energy in the two systems. Charge is only released at the interaction point, but phonons are produced via two mechanisms, at the interaction point and during charge propagation. We will refer to phonons as primary when they come from the interaction point, or as secondary if they were

released from charge or phonon propagation interactions with the lattice. For each type of response, we describe them in terms of their energy propagation:

- Phonon Energy
 - Energy is deposited directly into the crystal lattice creating primary phonons
 - The energy propagates/vibrates through the crystal as secondary phonons
 - The motion of charges, through NTL gain, creates more secondary phonons due to the interaction with the crystal lattice as charges are pulled by the bias voltage
 - Phonons are collected on both sides, each providing an independent measure of the phonon energy deposited in the crystal

- Charge Energy
 - Particle interactions either ionize the electrons in the crystal directly or as part of a nuclear recoil, this produces electron-hole pairs
 - Charges are accelerated in the electric field supplied by the potential voltage difference, and scatter within the lattice either changing directions or depositing their energy into the crystal vibration as secondary phonons
 - We collect the charge on both sides, each of which provides an independent measure of the charge energy deposited in the crystal

This is made quantitative by describing each of these energy mechanisms by their contributions. We start by noting that the total recoil energy, E_r , is the sum of the primary energies by writing

$$E_r = E_{Q,primary} + E_{Ph,primary} \quad (3.1)$$

where $E_{Q,primary}$ is the primary charge energy and $E_{Ph,primary}$ is the primary phonon energy. Taking into account the secondary phonon energy, $E_{Ph,NTL}$, due to the charges accelerated

in the electric field, the expected readout energy is given by:

$$E_Q = E_{Q,primary} \quad (3.2)$$

$$E_{Ph} = E_{Ph,primary} + E_{Ph,NTL} \quad (3.3)$$

We describe each of the three contributions in turn.

3.1.1 From Recoil Energy to Primary Energy

At the interaction point the recoil energy, E_r , transferred to the crystal as energy is converted into either vibrations in the lattice and/or liberated electron-hole pairs. We describe how this total primary energy is partitioned into the charge and phonon primary energies by a well-determined ratio, or yield. We follow with the description for how to determine such yield value depending on the interaction type. We next describe the value for nuclear recoils, which follows from condensed matter physics calculations (Lindhard theory [60, 61]), which leaves us well set for describing the number of charge and phonon carriers produced by the interaction in the next section.

The fraction of the recoil energy that goes into $E_{Q,primary}$ and $E_{Ph,primary}$ energy is given by Lindhard theory [60], which depends on whether it is an electron or nuclear recoil interaction. The fraction of the recoil energy that is deposited into charge energy is known as the expected yield, Y_{Exp} , and its complement $(1 - Y_{Exp})$ describes the phonon primary energy. Rewriting Eqs. 3.2 and 3.3, the total energy in both systems becomes:

$$E_Q = E_r \cdot Y_{Exp} \quad (3.4)$$

$$E_{Ph} = E_r \cdot (1 - Y_{Exp}) + E_{Ph,NTL} \quad (3.5)$$

The value of Y_{Exp} is determined by how the energy is transferred to electron ionization or lattice vibration, in what is called a stopping mechanism described by condensed matter

physics [60, 61]. Since electron recoils do not produce primary phonons, all the interaction energy is converted into releasing electron-holes so $Y_{Exp} = 1$. Nuclear recoils on the other hand go through a complex process of down converting their energy into both released charge, and vibration (phonon) energy with the yield value referred to as $Y_{Lindhard}$. Summarizing we take:

$$Y_{Exp} = \begin{cases} Y_{Lindhard} & , \text{ nuclear recoil} \\ 1 & , \text{ electron recoil} \end{cases} \quad (3.6)$$

Since the process which determines the value of $Y_{Lindhard}$ is complex, and described by condensed matter physics, we summarize the results here [60]. The value only depends on the crystal material properties (atomic number Z and mass A), the interaction energy E_r , and the details of the stopping mechanism [60, 61]. To describe this quantitatively, we note that there are three terms which are ϵ (which describes the amount of energy deposited per proton, because it is charged, called reduced energy), k (which describes the amount of energy loss in reduced energy units, per released electron-hole pair, called electronic stopping constant), and $g(\epsilon)$ (the term that describes the number of collisions with electrons in the crystal) [61]. Taking into account the atomic number and the mass the reduced energy value is given by $\epsilon = 0.0115E_r Z^{-7/3}$ [60, 61] (where the constant also takes into account the atomic mass A). The electronic stopping constant is also dependent on Z and A , given by: $k = 0.133Z^{2/3}A^{-1/2}$ [60]. Finally $g(\epsilon)$ is a numerical expression that is energy dependent and given by: $g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon$ [60]. Taken together for a specific atom we get:

$$Y_{Lindhard}(E_r) = \frac{k \cdot g(\epsilon)}{1 + k \cdot g(\epsilon)}, \quad (3.7)$$

where, after fixing with the Ge Z and A , k equals 0.157. Thus, for germanium the value of $Y_{Lindhard}$ is a recoil energy dependent function that is quickly increasing for nuclear recoil energies below ~ 8 keV and typically below 1/3, as shown in Figure 3.2.

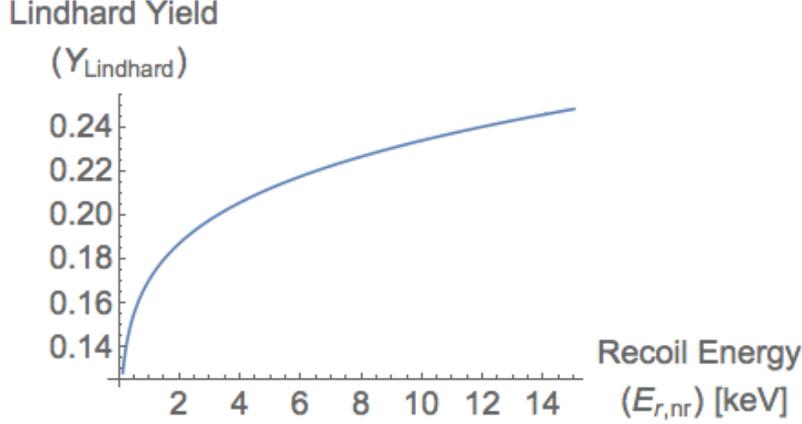


Figure 3.2: The Lindhard yield describes the proportion of energy released into charge and phonons (vibrations) of a crystal lattice, from a single scatter with a nucleus in the crystal. The figure shows the Lindhard yield ($Y_{Lindhard}$) as function of the nuclear recoil energy [60]. The fraction is typically below 1/3. For electron recoils all the primary energy is converted into liberated electron-hole pairs, and none released into primary phonons, so the Lindhard Yield doesn't apply and can equivalently be set to a value of 1.

3.1.2 From Primary Interaction Energy to the Number of Phonons and Electron-Hole Pairs Produced

While only the number of charges produced is needed for the charge energy measurement (the topic of this thesis), we also describe the number of primary phonons produced as both are used in the simulations (presented in Chapter 7.3). These are the initial population of electron-holes and phonons to be propagated in the crystal. We begin with phonons and continue next with electron-hole pairs.

The number of primary phonons is determined from $E_{Ph,primary}$ and the maximum allowed phonon energy (frequency) in Ge, denoted E_{Debye} as is determined from the Debye model, as described in Ref. [62]. The value depends on the material density and temperature. For Ge at mK temperatures it has a value of 8.1 eV ($\sim 2 \cdot 10^{12}$ Hz) [62]. Since, to a good degree of approximation all but one of the phonons has the maximum allowed energy we

find $N_{Ph,primary}$ to be:

$$N_{Ph,primary} = Quotient\left(\frac{E_{Ph,primary}}{E_{Debye}}\right) + 1 , \quad (3.8)$$

where the quotient notation is used because the division is rounded down to the near-lower integer, as a number of produced phonons with E_{Debye} is quantized, with an additional primary phonon produced with the residual energy.

While the number of primary electron-holes is similarly quantized, the additional energy is lost to the crystal and can be neglected. The total total number of e/h pairs is given by the ionization energy $\epsilon_\gamma = 2.96eV$ [62]. Quantitatively we find, for a germanium crystal, the number of electron-hole pairs $N_{e/h}$ is determined from the the charge energy E_Q and the electron binding energy ϵ_γ [62],

$$N_{e/h} = Quotient\left(\frac{E_Q}{\epsilon_\gamma}\right) , \quad (3.9)$$

3.1.3 Phonon Energy Released by High-Speed Charges Interactions in the Crystal (Secondary Energy) and the Total Energy Expected Values

Next we describe the estimate of the NTL phonon energy $E_{Ph,NTL}$ created as electron/holes drift at high speeds interacting with the crystal lattice while being accelerated by the electric field. This will completely describe the mechanisms that produce the observable energy, and gives us the basis for the detailed description of the charge and phonon emission and propagation processes described in the next section.

Since all the energy of the accelerated electrons is converted to NTL phonon energy, we can estimate the total secondary phonon energy from the number of charges and the potential accelerating them. Quantitatively we can estimate the $E_{Ph,NTL}$ (from Eq. 3.9) as:

$$E_{Ph,NTL} = N_{e/h} \cdot e\Delta V_{bias} = \frac{E_Q}{\epsilon_\gamma} \cdot e\Delta V \quad (3.10)$$

where e is the electron charge and ΔV_{bias} is the bias potential in the detector (4V in normal operation).

Summarizing, we find that the total amount of expected phonon and charge energy, from Eqs. 3.4 and 3.5 is:

$$E_Q = E_r \cdot Y_{Exp} \quad (3.11)$$

$$\begin{aligned} E_{Ph} &= E_r (1 - Y^{Exp}) + E_Q \frac{e\Delta V}{\epsilon_\gamma} \\ &= E_r \left(1 + Y^{Exp} \cdot \left(\frac{e\Delta V}{\epsilon_\gamma} - 1 \right) \right) \end{aligned} \quad (3.12)$$

Taking into account that the material properties and the voltage in the crystal are fixed, the equations depend only on the particle interaction type (nuclear or electron recoil) and the recoil energy, with the results shown for Ge in Figure 3.3. Of particular importance is that the ratio of phonon energy to charge energy is measurably different between electron and nuclear recoils, as was described in Figure 2.2. Looked at from another point of view, from the two measurements, we can determine the Yield on an event-by-event basis to determine whether it is likely to be from an electron or nuclear recoil.

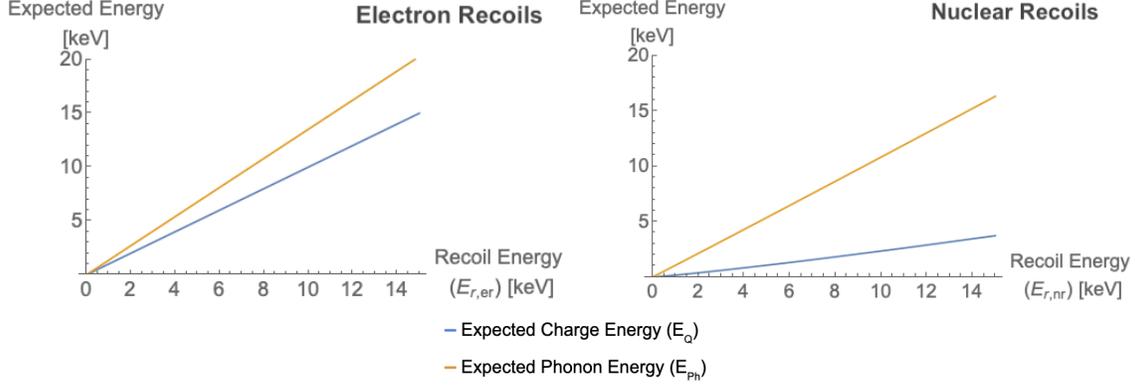


Figure 3.3: The total expected energy for the charge and phonon readout systems, as a function of recoil energy, for electron and nuclear recoils. Left is for electron recoils (like photons and electrons), where all recoil charge is from the recoil interaction and all phonon energy is from NTL gain. Right is for nuclear recoil events, like neutrons or WIMPs. The fraction of the energy in the phonon system is significantly larger because of 2 added effects: the Lindhard Model scale factor (well below $1/3$) which puts a larger fraction of the energy into phonons, and the phonon energy from NTL phonons, phonons created from the moving e/h in the accelerating field, which is correspondingly smaller. The added effects account for phonon energy to be ~ 4 times the charge energy.

3.2 Motion of Charges and Phonons in the Crystal

Now that the number and energy of the created electrons, holes and phonon has been described, we can proceed with the details of the propagation of each in the crystal lattice. This includes how they move through the lattice structure, how they scatter (including how secondary phonons are produced) and how they are eventually lost or absorbed by the sensors. This will be important both for understanding the readout described in next section as well as when we get to the simulations description in Chapter 7. First in Section 3.2.1 we will describe the motion of charges taking into account the constraints of the Ge crystal, as they are pulled by the electric field, as well as the absorption or trapping scenarios to end their motion. Next, we describe the phonon emission by the motion of the high-speed electron-holes in Section 3.2.2. Finally, Section 3.2.3 describes the motion of phonons in the Ge crystal from either source. Beware, that as simplistic as we tried to make these descriptions, the following sections are dense in crystal physics theory.

3.2.1 Charges Pulled by the Electrodes But Subject to the Constraints on the Motion From the Ge Crystal

The motion of charges is constrained by the lattice properties of the material (germanium in our SuperCDMS detectors) and described by solid state physics. The propagation of the electrons and holes in the crystal are different, and cause measurement effects which must be understood and simulated accordingly. The reason for this difference is that electrons and holes propagate along the path of least-resistance, which is determined by the effects of the potential in the lattice. In this section we first describe the momentum variables and crystal physics concepts to understand the motion of each in a crystal lattice. Next we will use these concepts to describe the motion of electrons and holes in the germanium lattice structure in the context of the electric field in which the detectors operate. We will focus on germanium (not silicon) since it is the subject of this thesis. We quickly describe the scattering process which causes phonon emission and give more detail in the next section. Finally we will describe how the charges end their motion, either at the electrodes or in a ‘trapping mechanism’, either by impurities in the crystal or by edge effects near the vertical sides of the detector.

To understand the motion of charges in a crystal lattice we need to review some crystal physics concepts, since the path of least resistance of charges is more readily described by transforming the geometrical position of the lattice into an equivalent representation in momentum variables (given the repetitive properties of the crystal), which is known as the reciprocal space [63]. The transformed smallest representation of the crystal primitive cell into reciprocal space is called the Brillouin zone [63]. The Brillouin zone of a crystal depends on the geometrical array of the lattice, and has several high-symmetry points, common labels particular to the germanium lattice are Γ , L, K, U, W, X [64, 52], as shown in left of Figure 3.4. These symmetry points relate to a system for representing the basis of the reciprocal lattice, which help identify the orientation of such points in the form of vectors, labelled Miller-index planes [63]. By solving the energy system of the crystal lattice

in the context of Miller-indices, obtaining the potential energy as function of Miller-index or Brillouin symmetry point helps understand the direction of motion of charges as they would propagate in the least-resistance (lowest potential) point (or valley) [63].

Given the band structure of Ge, the motion of electrons propagating through the crystal follows that of the [111] Miller-index plane (described below), while holes propagate freely (also called ballistically) [64]. It is easiest to see the geometrical implication of the propagation constraints with the Brillouin zone of the Ge crystal, as shown in Figure 3.4 (left). The [111] plane corresponds to the L point (pointing from the origin to the cube vertex), while the [000] corresponds to Γ (origin). The L points imply propagation in 8 lines (one to each vertex), while the Γ point implies free particle propagation. Put in the context of our detectors, with a voltage bias of +2V at the top and -2V at the bottom electrodes, the propagation is further constrained. In the presence of the Z-oriented electric field, electrons would only propagate and accelerate in the direction of the 4 top vertices of the cube, while holes are simply pulled directly to the bottom plane electrodes. The overall trajectory of charges is well described by these planes.

The trajectories are also affected by scattering a process with the lattice which can both change their direction as well as release phonons. This is the mechanism of phonon emission, which will be described in the next section, and is the reason for NTL gain. This process is the same for both electrons and holes, with the only difference being the constraints of the crystal momentum described above.

The motion of the electrons and holes ends either by being absorbed by the electrode, or by being lost in a trapping mechanism. There are two trapping mechanisms which can cause them to not be collected by the readout system, and cause an undermeasurement of the charge energy. The first is for charges to be trapped by impurities in the crystal and get reabsorbed by the lattice in a process called charge trapping [66]. The second is due to edge effects near the sidewall of the crystal, where charges also get stuck or (if the interaction happens too close to the edge) get ejected outside of the lattice in what is referred to as

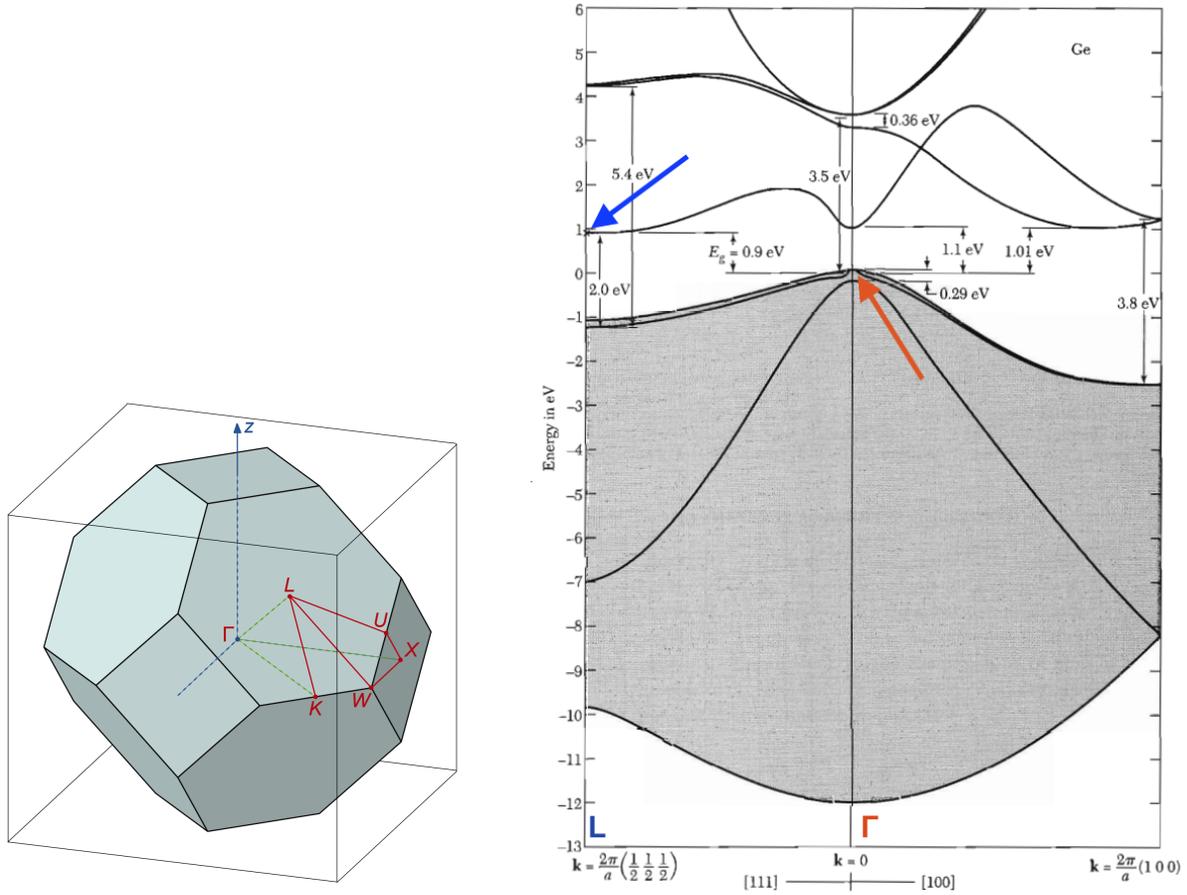


Figure 3.4: Two figures describing the lattice structure and propagation properties of germanium crystal. Left: The Brillouin zone for the Ge crystal which is a diamond cubic lattice. Notice that while we have indicated many of the labels on the left hand side, we have only indicated the Γ and L directions on the right hand side. For electrons there are eight L points, which dictate the propagation direction. Image reprinted from Ref. [65]. Right: Germanium band structure. The lowest conduction valley is at L (blue arrow), which makes electrons propagate in the $[111]$ plane. The maximum valence band value is at Γ (red arrow), and is roughly parabolic around the maximum, hence holes propagate ballistically. Image reprinted from Ref. [64]. The electric field in the detector is aligned in the Z direction, so for interactions away from the edges of the detector (where the field changes rapidly) electrons propagate in 4 narrow cones (following L -direction), towards the positive biased detector side (side1, top), while holes propagate freely in the potential pulling them to the negatively biased side (side2, bottom).

Auger Electron Ejection [67].

3.2.2 Electron-Hole Scattering and Phonon Release Due to High-Speed Motion of Charges

As charge carriers accelerate in the electric field they can accumulate enough energy to create phonons which we will refer to as NTL phonons. From a descriptive point of view we can think of them as emitting phonons if they move above the speed of sound in the crystal [62]. Any such phonons contribute to the total phonon energy as they are a type of secondary phonon and the process of phonon release from a charge is called phonon scattering [42, 62, 68]. Emission of phonons from electrons and holes is the same, except for the direction of motion constraints given by the crystal implications of each (as given in the previous section). We will first describe the properties of the emitted phonons as they make the phonons propagate differently. Next we will describe the phonon emission from holes, followed by the phonon emission from electrons, which will leave us well set to describe the phonon propagation in the next section.

All phonons can be described by their energy (frequency) and mode (direction of oscillations with respect to the motion of the phonon). Each mode is also limited in its motion by the crystal physics, giving directional preference and velocity restrictions [62, 63, 69], in what is called phase velocity. In germanium the resulting phonon can have one of 3 modes: longitudinal, fast-transverse, and slow-transverse (the fast and slow connotation is because one has a higher phase velocity than the other) [62, 69], and the phonon emission direction depends on its phase velocity which is different for each mode. The phase velocities for all such modes can be calculated from the polarization eigenstate solutions describing the system, as described in Refs. [62, 69], for example, the longitudinal group velocity is the speed of sound in the crystal [62]. We next describe the phonon emission by holes and electrons separately.

Phonon emission by hole propagation is described by simple drift equations from the electric field, but restricted by the speed limitation imposed by the phonon scattering process (speed of sound) [62]. The scattering rate, the probability of creating a phonon (of

a particular mode), can be estimated by using quantum perturbation approximations for the transition probability of one eigenstate (hole) into eigenstates of a system (phonons), a derivation from Fermi's Golden Rule [62]. It is possible to solve a system of equations relating the phonon emission angle, the emitted phonon momentum, and hole out-bound momentum, as derived in Ref. [62, 68].

The emission of phonons from electron propagation is the same as holes, except for the restriction of the charges motion in the Ge L-valleys. The system is readily solvable by performing a transformation into a space that is aligned with the Miller-index, known as the Herring-Vogt transformation [62, 70]. Upon performing the transformation, the electron system is solved in the same way as holes [62]. The scattering rate is also governed by the speed of sound, which remains the same (since it is given by the phonon system) [62].

3.2.3 Motion of Phonons in the Crystal

The motion of both the primary and secondary phonons through the crystal is governed by their phase velocity, which depends on their mode and whether there are any impurities in the crystal. The primary phonons will down-convert quickly to lower energy phonons (for our purposes, secondary phonons) for reasons that will be clear soon, but the propagation processes are the same regardless of the phonon type, just dependent on the phonon energy. We describe these processes next.

Phonons will travel freely, up to one of three scenarios: a scatter in the lattice, a decay process (into lower frequency phonons), or absorption by the dedicated sensors or by impurities in the lattice. Either of the first two can change the mode of the phonons and the scattering rate rises as a function of the phonon energy. At the scattering or decay point, the output phonons will be directed according to the conserved momentum from the process and the phase velocity constraints for the outbound phonons modes. Longitudinal, high frequency phonons can decay into lower energy phonons in a process called anharmonic decay [69], where the longitudinal phonon can decay into any combination of modes, with different probabilities. The mode populations fluctuate around the measured proportions

of the population of modes in germanium, slow-transverse (55%), fast-transverse (35%) and longitudinal (10%) [62, 69]. Both, longitudinal and transverse phonons can scatter off a mass defect in the crystal (an impurity from a different Ge isotope [69]) and change their mode, a process named isotope scattering where the outgoing modes are estimated probabilistically according to the expected populations. These reasons are why primary phonons quickly down-convert into secondary phonons [69, 62]. Secondary phonons will continue to propagate until absorbed by the dedicated sensors (as it is also a thermal effect) or ultimately down-converted to such low energies that the sensors can no longer absorb them (described in more detail in Section 3.4).

3.3 Charge Readout Response to Energy Deposition

Having described the phonon and charge energy and propagation, we are ready to explain how both are collected by the dedicated sensors to produce the recorded signals in the form of detector waveform readout. We describe the charge readout creation in this section, and the phonon readout in the next. Section 3.3.1 describes the charge collection as described by Ramo theory [71]. Next, Section 3.3.2 describes the readout by the charge system and the constructions of the output waveform. The readout output will be processed by the event reconstruction software, which is described in the next chapter.

3.3.1 Charge Collection by the Electrodes

As mentioned in Chapter 2.1, the charge readout is performed by the electrode lines of the FET system, running at the top and bottom surfaces, as shown in Figure 2.5. Since the total effective charge readout is dependent on when and where the charge stops its motion, we first give a quantitative summary of the Shockley-Ramo theory which describes the absorption, then make our predictions quantitative for the top and bottom sides of the detector. We follow with remarks on charge collection position dependence, as they will be useful for our simulation understanding and comparisons. With these descriptions we will move on to the readout description in the next section.

The observable signal read out from the electrodes depends on the value of the electric field at the location of where the electron or hole is at the time when it is sensed. Since this happens almost instantaneously, the charge collection description is readily described by the Shockley-Ramo Theorem [71, 72], where the electric field is transformed into a unitary field that represents the effect of distance in the sensed charge, which is called the Ramo Field and denoted ϕ_{Ramo} . For every i -th charge, q_i , the Ramo Field value $\phi_{i,Ramo}$ is taken at the final location point at the moment in which it is sensed. Note that the sign of the charge is used here to obtain the net charge collection. The ϕ_{Ramo} parametrization is convenient when the field has a known value at every point in space, and is a simple scaling normalization which takes $\phi_{Ramo} \rightarrow 1$ at the electrode while grounded at the rest of the electrodes. The total effective charge absorbed is given by:

$$Q_{Ramo} = \sum_i q_i \phi_{i,Ramo} \quad (3.13)$$

where given charge trapping and ejected electrons, we approximate things such that the final location of an electron or hole can end up in either one of the following end-states:

- Reach an electrode: get near unity Ramo field value, good collection, $\phi \approx 1$
- Be trapped in the crystal: partial Ramo field value, incomplete collection, $\phi \approx 0$
- Escape the crystal: zero Ramo field, no charge collection, $\phi = 0$
- Reach the opposite-side interleaved potential line (grounded): zero Ramo field but also opposite contribution in opposite-side, causing incomplete collection in both sides, $\phi \approx -1$

Knowing the number of produced electron-hole pairs from Eq. 3.9, we can estimate the Ramo collected energy, E_{Coll} , by the top/bottom electrodes using:

$$E_{Coll}^{Top/Bottom} = \frac{Q_{Ramo}^{Top/Bottom}}{\pm e} \epsilon_e/h, \quad (3.14)$$

where $Q^{Top/Bottom}$ is the Ramo collected charge at the top/bottom electrodes, divided by the electron charge, and $\epsilon_{e/h}$ is the electron-hole creation energy. If electrons reach the appropriate side (top) and holes reach their appropriate collection side (bottom), then the full energy is collected for each side giving us two separate measurements of the charge energy. With this equation in mind we can now see why the sign of the q_i does matter in the total Q ; if both electrons and holes are collected in the same side, their net collection cancels. Similarly, if electrons or holes do not reach the electrode (i.e. get trapped in the crystal), their Ramo field value is not unity, so the energy is undermeasured.

It is expected that the charge collection is dependent on the position of a particle scatter, which will later be relevant to our simulations studies. Since there are trapping effects at the sidewall (large values of the radius) we expect poor collection ($\phi \approx 0$) near the edges. Furthermore, since electrons propagate in the L-valleys we expect the mismeasurement region to be different at the top than the bottom, and different than hole's mismeasurement region as well. It is also expected that near the top or bottom surfaces, due to the interleaved potential that events near the surfaces are poorly measured. As we will show in later chapters, the simulation results reflect these effects, and the analysis event selection criteria (used in previous publications) are based on them.

3.3.2 Charge Waveform Creation

With the charge collection description we now turn to how the total effective charge is converted to a readout signal using an amplifying circuit. We begin with a brief description of the amplifying circuit functionality. We next provide a description of the signal outputs. We then provide details on the noise and the noise estimation, as it will be useful for the reconstruction algorithms described in the upcoming chapter. We finish with remarks about readout imperfections which tie to the backgrounds defined in Chapter 2, and will connect to the readout mismeasurement in Chapter 6.

The charge collection readout waveform is created when charges reach the the electrodes which are part of an amplifying circuit, as shown in Figure 3.5. The combination of compo-

nents designed to be very efficient for driving low-power, low-noise signals [55], can operate at low temperature, and is located in the cold electronics of the towers, near the detectors. As charge propagates in the crystal and is absorbed at the electrodes, the crystal mimics a capacitor, C_d , in the charge signal amplification circuit. Since the drift velocity of electron-holes is 20 to 30 km/s, much faster than a digitization bin [42, 62, 68], all of the charge collection occurs essentially instantaneously. When the event’s charge is absorbed a change in current is released to the circuit which is up-scaled and filtered by the junction gate field-effect transistor (FET) amplifier [42] (this is why the charge readout is referred to as a FET signal).

Since the analog amplitude of the event is driven by the amplifying choices in the circuit, and to a good degree of approximation the collection happens instantaneously, the signal output quickly rises to a value determined by the total charge collected and has a fall time of approximately 200 μs . The final result from the charge system (single-channel) is a digitized waveform of 1,638.4 μs in 2048 bins (0.8 μs per bin) and is represented by a 14-bit Analog-to-Digital Converter (ADC), ranging from 0 to 16,834 ADC counts [42], with the total charge collected (and thus the energy) readily described by either the peak value or the integrated area under the curve. See top right of Figure 3.5 for an example of the circuit response to an incident particle, and bottom left for the readout when there is no input, known as a ‘noise event’.

As expected for an RC circuit, the noise is correlated in the time domain, but is readily described independently in the frequency domain. As shown in Figure 3.5, the power as a function of frequency, denoted Power Spectrum Density (PSD) [74], has a number of features worth noting. The first is that it is not white-noise (the noise has a frequency dependence, it is correlated frequency-to-frequency and not cross-correlated between frequencies) and second is that the PSD changes only over time scales that are much longer compared to the readout time (as it is the baseline output from the circuit) [74, 75]. These properties will be used extensively in the reconstruction algorithm described in Section 4.1.

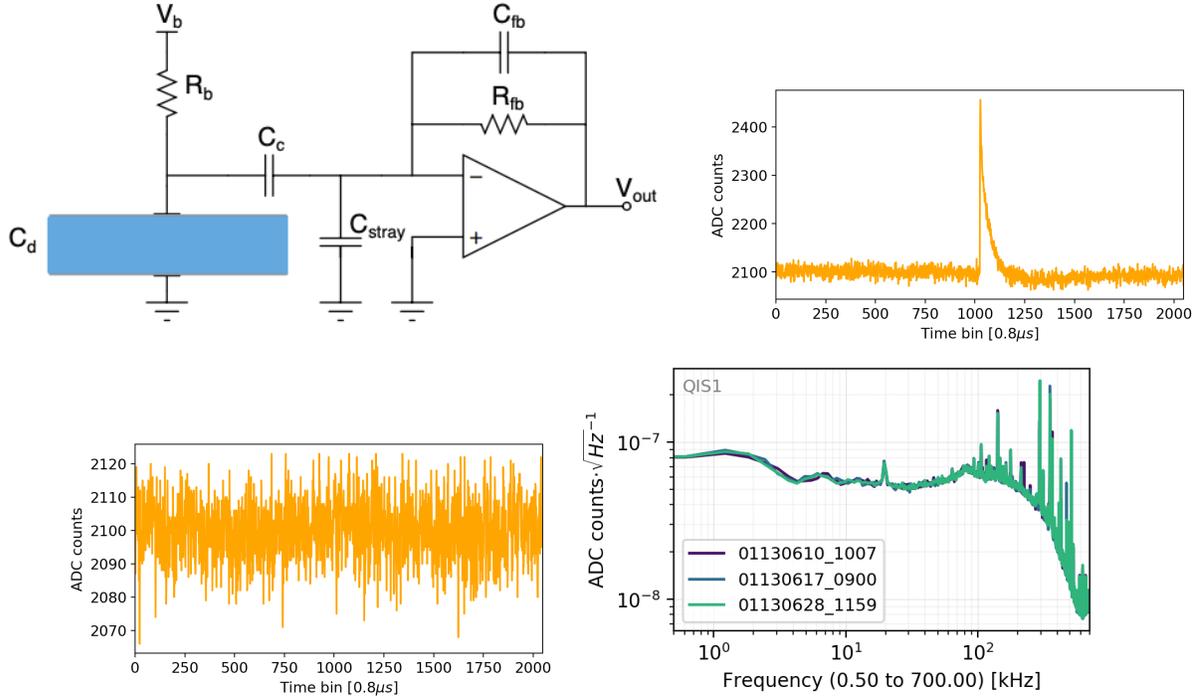


Figure 3.5: A set of figures showing the charge readout of the detectors and its response in the presence and absence of a particle interaction. Top left shows the charge signal amplification schematics used in the FET readout system, where the detector serves as a Capacitor (C_d) in the circuit (figure reprinted from Ref. [73]). Each of the four channels in the detector functions as a circuit responding to the capacitance, as seen by the electrode’s orientation. The typical output real-interaction event looks like the top-right figure, while the amplification circuit noise is shown in the bottom-left plot. The bottom right figure shows the noise profile or Power Spectrum Density (PSD) which shows the intrinsic power as function of frequency (described in more detail in the next chapter), of the charge circuit of one channel in the absence of real interactions, taken from 3 different samples, showing that noise levels remain constant as baseline output from the filter-circuit. It is clear that such noise is not white noise, from the construction of the noise profile it is known that it is also not cross-correlated in between frequencies.

Before concluding, we note that there are imperfections in the readout system as mentioned in Chapter 2, Table 2.1. These include cross-talk between the readout channels, as well as Low Frequency noise, Glitches and Blips. Each will be described in more detail in Chapter 6, and cleaned up to produce a pure noise sample in Chapter 8, as they they are sources for mismeasurement our intrinsic noise measurement studies with noise samples.

3.4 Phonon Readout and Collection by TES Sensors

Having described the charge readout now we turn to the phonon collection and readout. We will give an overview of the operation of the TES which follows from their arrangement for readout described in Chapter 2.2 in Figure 2.6. While the phonon system is fully simulated, as described in Section 7.3.3, we will not give much detail since it is not the topic of this thesis. More information about its operation and amplification can be found in Ref. [53] and in the SuperCDMS Soudan Technical Design Report [42].

As shown in Figure 3.6, a TES cell consists of thin aluminum (Al) layer, often called a fin, placed (deposited) over the crystal, and covered (masked) with a thin layer of tungsten (W) [42]. The TESs measure the phonon energy by virtue of sitting right at the critical temperature for being superconducting (at ~ 40 mK), such that when energy is deposited by phonons in them, they heat up so the resistance changes, which in turn allows a current to flow which is proportional to the amount of energy deposited. Any phonon entering the Al system will break a Cooper pair if above the binding energy of 0.34 meV [58]. Once Cooper pairs in Al have been broken, electrons are freed and get absorbed by the W layer, heating up the W system and producing an observable highly responsive change in resistance which is measured out by its readout circuit [42]. From the phonon emission distribution in Ge it is calculated that in normal operation 91.5% of the electron-emitted phonons and 99.7% of hole-emitted phonons are above the Al Cooper binding energy, and thus would be measurable [58].

With our description of the waveforms created by particle interactions complete we are now ready in the next chapter to describe how the readout response is treated to obtain the best-estimate energy measurements. This will be used as a central tool as we prepare for understanding the response and use them for the discriminating technique for a WIMP-search analysis in SuperCDMS.

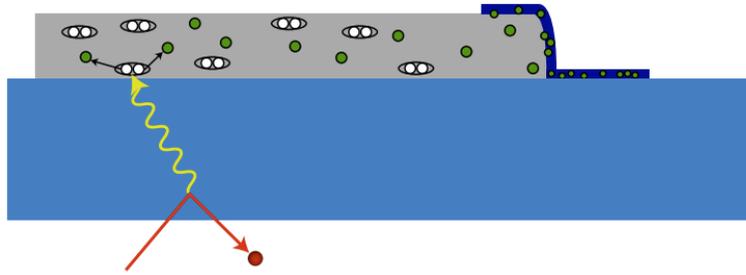


Figure 3.6: A diagram showing a TES readout sensor, where the red indicates a particle that interacts in the crystal which creates a high energy phonon (yellow) which in turn is absorbed in the Al fin, breaking up a pair of bounded electrons (green) which can then be collected as part of the readout electronics. If the phonon energy is above the Cooper pair gap energy, the pair is broken, and the excited particles relax in the decreasing energy gap of the Al / W interface, releasing energy that heats the electron system of the W in the TES, which is carefully maintained in the superconducting transitioning state. Image reprinted from Ref. [76]. The transitioning state is hyper-sensitive to temperature changes, strongly responding with increased resistance as the W heats-up, and producing an observable signal.

4. EVENT RECONSTRUCTION: OPTIMAL FILTER

With a description of the SuperCDMS detectors and their creation of digitized waveforms response complete, we now describe how we turn the output from the electronics into best-estimate energy measurements. Since the size of the signal from the smallest energy deposits are comparable to that from noise, the noise is correlated in the frequency domain (with many frequencies having poor signal to noise), and there are challenges from cross-talk between channels we use a special algorithm known as an optimal filter [74]. By measuring each of the four channels independently, but using the combined information from them, we can get a useful set of measurements for our searches.

Since the algorithm searches the waveform for the presence of an interaction and tries to determine when and how energetic it was, we give a quick overview of how it works. One can think of it as being a simple χ^2 minimization routine that determines the best-fit amplitude and time. In the case where there is no interaction, we can fix the time, and obtain the corresponding χ^2 and amplitude to get a measurement of the electronics noise and baseline resolution of the detectors. We describe the optimal filter in this chapter along with its implementation in the event reconstruction software package and in the next chapter we will describe how we use the best-estimate energy observables to validate the detector response, discriminate the backgrounds, and perform WIMP searches in SuperCDMS.

In this chapter we present the methods to determine our charge energy observables from the detectors. In Section 4.1 we give an overview of the optimal filter implementation and the assumptions. Next, in Section 4.2, we describe the signal and cross-talk templates used as inputs and how they are created from data. We then describe the particular implementation for the readout for the charge system in CDMSBats in Section 4.3. Finally we describe the energy observables derived from the algorithm, as used in analyses, and for pure-noise estimates in Section 4.4.

4.1 Optimal Filter Implementation Overview

We begin with an overview of the optimal filter algorithm, describe how it works and the assumptions we make in its use. First we provide an overview of the method, inputs and results. We then describe the inputs translated in the frequency domain, as needed by the algorithm. We follow with presenting the χ^2 estimate to be minimized. Next we describe the noise estimate PSD calculation and its limitations due to the discrete digitization of the signal. Finally, we show the χ^2 minimization result, and give some remarks on the results from the algorithm.

Because of the readout method, as described in Chapter 3.3 and shown in Figure 3.5, the noise is correlated in time, so a simple time-domain χ^2 is not an ideal solution for an accurate measurement [74, 75]. For this reason we have chosen to turn our waveforms into energy measurements using an algorithm that works in frequency space, known as the optimal filter, we next describe why and provide an overview. A diagrammatic representation for a single channel reconstruction is shown in Figure 4.1. The optimal filter takes 3 inputs: the observation (a waveform as a function of time), a signal template (the shape we expect from real energy depositions) and a noise estimate (in the form of the PSD described in Chapter 3.3). Since both the energy of the interaction and arrival time of the particle interaction relative to the start of the readout time is important for proper measurement, this is done by calculating a χ^2 with both the amplitude, A , and time-delay, t , of the signal as free parameters. Since the dominant uncertainty in the waveform measurements is due to fluctuations from noise, the expected variations due to noise are taken as the sole uncertainty. With these pieces a χ^2 minimization can readily be done. The optimal filter produces 3 outputs: the best estimate amplitude of the expected signal, the time-delay of the signal with respect to the template, and the minimum value of the goodness of fit parameter, χ^2_{min} (note that we will often refer to this quantity as just χ^2).

Before moving to the χ^2 minimization equation, which is done in frequency space for the reasons mentioned above, we note that our methodology must take into account that

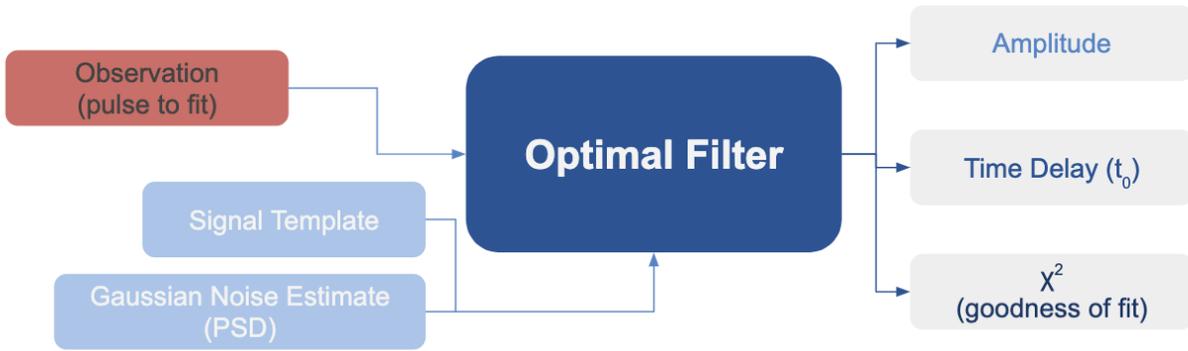


Figure 4.1: A diagram of inputs and results of the optimal filter calculation. The algorithm finds a minimization of the χ^2 as a function of amplitude and time-delay.

the digitized waveforms come integrated into time-bins. For this reason we first must turn our time-domain expected signal shape (our template) $s(t)$, and our observation $v(t)$ into frequency-domain. In frequency domain, the observation $\tilde{v}(f)$ is described by the signal $\tilde{s}(f)$ times the amplitude A , as given by:

$$\tilde{v}(f) = A \cdot \tilde{s}(f) , \quad (4.1)$$

where we have used ($\tilde{}$) above the variable to indicate they are FFT versions, a notation that will be used throughout this Chapter. Instead of a Fourier Transformation (FT) for infinite-continuous calculations, we use a Discrete Fourier Transform (DFT) for the actual implementation since our samples are not continuous (binned) [74] or infinite length; computationally, the implementation is achieved with a Fast Fourier Transform (FFT) [74, 75].

For our minimization we assume that the expected variation between the signal template shape and the data is given by the variations from the noise (which is frequency dependent). As described in Section 3.3, the noise distribution is well described by the expected power as a function of frequency PSD (its calculation method is described below) and we denote it as $J(f)$ (note we do not use the $\tilde{}$ notation on J for convenience). The χ^2 , as in Ref. [75],

can be written as:

$$\chi^2(A, t) = \int_{-\infty}^{\infty} df \frac{|\tilde{v}(f) - Ae^{-i2\pi ft} \tilde{s}(f)|^2}{J(f)}, \quad (4.2)$$

where the time-shift t in frequency domain is given by $e^{-i2\pi ft}$, and we integrate over all frequencies.

Before describing the solution for the minimum χ^2 , we describe how $J(f)$ is computed. We will begin with the idealized description in continuous space and move to our discrete-sampling requirements. In principle, the values of $J(f)$ as a function of frequency must be constructed from an infinitely long sample of noise (denoted $n(t)$ in the time domain and $\tilde{n}(f)$ in the frequency domain). The function is constructed from the mean power of a signal (readily understood as power in a circuit, but extrapolated to any fluctuating signal), ignoring constants it is a term depending on the integral (in frequency or time space) of the squared value of the ‘current’. Assuming no correlation between frequencies, the expected power as function of frequency of the sample is given by [74]:

$$J(f) = \lim_{T \rightarrow \infty} \langle |\tilde{n}(f)|^2 \rangle, \quad (4.3)$$

where T is the length of the continuous sample (which is infinitely long). This shows that the PSD is the measure of the expected variance as function of frequency (following the definition of variance when the mean value is at zero), also denoted auto-correlation, since $\tilde{n}(f)$ is multiplied by itself, with the same ‘frequency argument f ’, as opposed to another frequency [74, 75]. In our case the waveforms are measured in discrete time bins of $dt = 0.8\mu s$ and of finite length (with $N = 2048$ bins), so the frequency resolution is $610.3 Hz$ [42]. We thus convert the expression into a discrete space, by denoting the frequency of the k^{th} bin as f_k we can calculate $n(f_k)$, and simplifying the notation we represent as \tilde{n}_k with N frequency-bins. The binned PSD J_k describes the typical distribution of power across the frequency

components of the discrete FFT decomposition, and is given by:

$$J_k = \langle \tilde{n}_k^2 \rangle \quad (4.4)$$

With the binned frequency distribution we can simply move to the minimized, best-fit value of the amplitude denoted \hat{A} . As given in Ref. [75], and assuming the best-fit time, \hat{t} , in our calculation we take:

$$\hat{A} = \frac{\sum_{k=0}^N e^{i\omega\hat{t}} \frac{\tilde{s}_k^* \tilde{v}_k}{J_k}}{\sum_{k=0}^N \frac{|\tilde{s}_k|^2}{J_k}}, \quad (4.5)$$

where the sum limit factor N refers to the number of frequency bins, and is symmetrically mirrored at the $N/2$ bin (called Nyquist frequency) because of the discrete FFT [74], w is the angular frequency (given by the digitization rate dt and N : $w = \frac{1}{dt \cdot N}$) which is the discrete version of f in the exponential of the χ^2 definition, while t is the time delay, an integer-multiple of dt , and \hat{t} is the best-fit time delay which is determined in two-dimensional minimization, or fixed (as in the case of our noise studies).

We make two notes about the result. The first is that the denominator is a normalization constant to the equation, also called the signal-to-noise ratio $\sum \frac{|\tilde{s}|^2}{J}$, which gives an estimate of the expected resolution of the measurement. The second is that the numerator shows why this is a filter that de-weights dominating noise components. While the $e^{i\omega t}$ is just the best-fit time (with the FFT operator), the second part is a type of filter in that it de-weights noise-dominating frequency components, as \tilde{s}_k^*/J_k which is divided by the noise power, multiplies the observation \tilde{v}_k . This has the advantage, that we can (by hand) actually remove certain frequency bins from the fit if they are expected to be problematic for systematic reasons.

4.2 Signal Template Construction

Now that we have seen how the optimal filter works, we can move on to describing the various signal templates which serve as an input and how they are created with an eye to our final implementation taking into account the cross-talk and all four channel inputs from

the charge system. While we will not describe the details of the methods for the phonon channels, the big picture is the same [53]. We first describe the overall template construction method, following with the description of the cross-talk versions. This will leave us well set for the implementation in the charge energy observations described in the next section.

Since it is expected that all pulses from the same RC circuit (see Figure 3.5) are the same shape, the template in SuperCDMS can be created by calculating the average pulse shape from of a sample of well-behaved events taken from ^{133}Ba calibration samples. All template creation processes follow a similar, iterative method, which we summarize quickly. What makes it iterative is that we use a first guess shape and the optimal filter to measure the amplitude and then normalize, and then we take the template as the average pulse shape from the sample of events. Finally, the resulting average is normalized by its maximum value.

While the main templates for the top and bottom channels are done as above, the cross-talk templates used in the final optimal filter implementation are created using the residuals from a sample of events after the template shapes have been subtracted off. They are calculated after obtaining both channel templates by taking the same sample of events, normalizing by the opposite channels's amplitude and finding the difference (referred to as the residual). We take as the cross-talk template the average of the residuals, with the results shown in Figure 4.2. Since the variation in the cross-talk templates is small, and the templates themselves are small, we assume there is no separate noise contribution to them in all the calculations below.

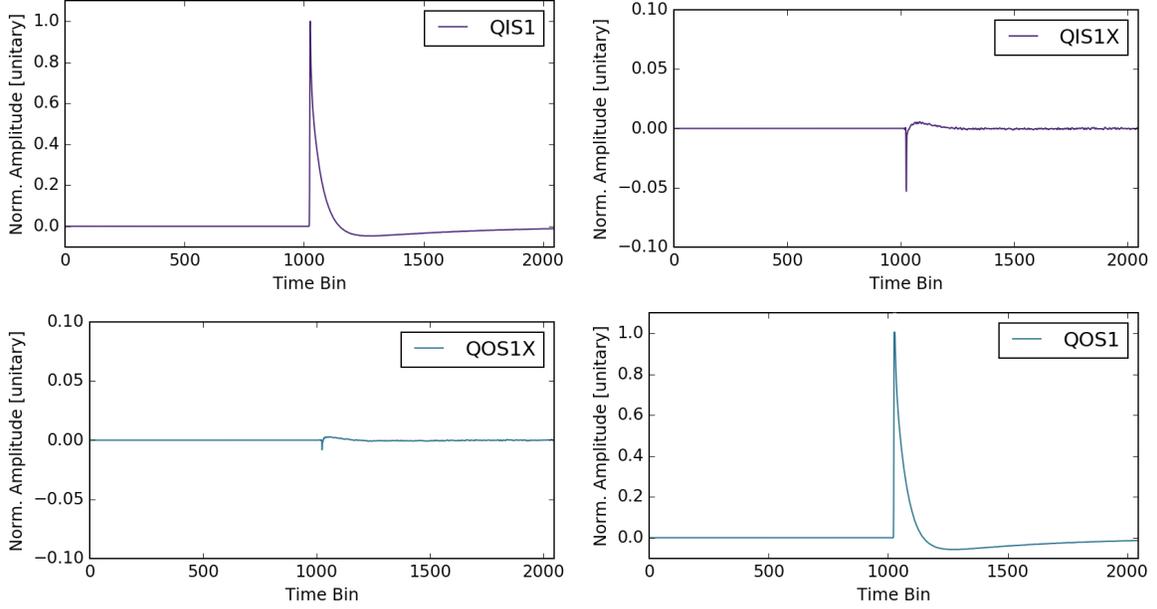


Figure 4.2: Example side1 templates: inner and outer channels, and 2 cross-talk templates. Notice that the templates are normalized to the channel-templates (not cross-talk), it is expected for the channel-signal to be much stronger than it's cross-talk. An optimal filter is performed for each of these and the minimized coupled system determines the amplitude of the event.

4.3 Charge Optimal Filter

With the basics of the optimal filter described, the shapes of the signal and cross-talk templates determined, and the PSDs of all four channels measured, we can describe the details of the specific implementation of the optimal filter for the charge readout measurements with the assumption that a signal should show up at the same time in all four charge readout channels. We will overview the fitting procedure, we then describe the linear system of equations for the observation, to follow with the χ^2 and the description of the minimization solution pieces, and then finishing with the constraint imposed coupling both sides' systems and an example fit of an event. In the next section we will describe how these optimal filter results get turned into our charge energy observables.

The fitting is done as a coupled system where we simultaneously fit both channels (inner and outer) from each side, and constrain the t from the χ_{min}^2 from each to be within a small

time from each other. Here we summarize the overall methodology, for a fuller description see Ref. [73]. Each side will provide an independent measurement of the energy (constrained by a small difference in the time-delay), and later we can use the difference as a measure of well behaved events and a combination to be a best estimator (as will be shown later in Chapter 5.3).

With 2 channels per side, and 2 cross-talk templates, a linear system of 2 equations per side is given by:

$$\begin{bmatrix} \tilde{v}_{\sigma I} \\ \tilde{v}_{\sigma O} \end{bmatrix} = \begin{bmatrix} \tilde{S}_{\sigma I} & \tilde{S}_{\sigma IX} \\ \tilde{S}_{\sigma OX} & \tilde{S}_{\sigma O} \end{bmatrix} \begin{bmatrix} A_{\sigma I} \\ A_{\sigma O} \end{bmatrix} \quad (4.6)$$

where a similar notation as that in Eq. 4.1 is used, but we have added the σ index to denote side 1 or side 2, and the indices I, O, X to denote inner-channel, outer-channel, and cross-talk (IX refers to cross-talk from outer as seen in inner, and vice-versa).

The χ_σ^2 for this system is constructed in the same way as the simple optimal filter, but considering the two channel's coupled equations:

$$\chi_\sigma^2(A_{\sigma I}, A_{\sigma O}) = \sum \left(\frac{|\tilde{v}_{\sigma I} - e^{i\omega t_\sigma}(A_{\sigma I}\tilde{S}_{\sigma I} + A_{\sigma O}\tilde{S}_{\sigma IX})|^2}{J_{\sigma I}} + \frac{|\tilde{v}_{\sigma O} - e^{i\omega t_\sigma}(A_{\sigma I}\tilde{S}_{\sigma OX} + A_{\sigma O}\tilde{S}_{\sigma O})|^2}{J_{\sigma O}} \right), \quad (4.7)$$

where the PSDs for the channels are denoted as $J_{1I}, J_{1O}, J_{2I}, J_{2O}$, and the solution to the best-fit amplitude $A_{\sigma I}, A_{\sigma O}$, which is also used to determine \hat{t}_σ , is thus given by (for derivation details see Ref. [73]):

$$\begin{bmatrix} \hat{A}_{\sigma I} \\ \hat{A}_{\sigma O} \end{bmatrix} = M_\sigma^{-1} \begin{bmatrix} \sum \text{Re} \left(e^{i\omega \hat{t}_\sigma} \frac{\tilde{S}_{\sigma I}^*}{J_{\sigma I}} \tilde{v}_{\sigma I} + \frac{\tilde{S}_{\sigma OX}^*}{J_{\sigma O}} \tilde{v}_{\sigma O} \right) \\ \sum \text{Re} \left(e^{i\omega \hat{t}_\sigma} \frac{\tilde{S}_{\sigma IX}^*}{J_{\sigma I}} \tilde{v}_{\sigma I} + \frac{\tilde{S}_{\sigma O}^*}{J_{\sigma O}} \tilde{v}_{\sigma O} \right) \end{bmatrix}, \quad (4.8)$$

with M_σ defined as:

$$M_\sigma = \begin{bmatrix} \sum \frac{|\tilde{S}_{\sigma I}|^2}{J_{\sigma I}} + \frac{|\tilde{S}_{\sigma OX}|^2}{J_{\sigma O}} & \sum \text{Re} \left(\frac{\tilde{S}_{\sigma I}^* \tilde{S}_{\sigma IX}}{J_{\sigma I}} + \frac{\tilde{S}_{\sigma O}^* \tilde{S}_{\sigma OX}}{J_{\sigma O}} \right) \\ \sum \text{Re} \left(\frac{\tilde{S}_{\sigma I}^* \tilde{S}_{\sigma IX}}{J_{\sigma I}} + \frac{\tilde{S}_{\sigma O}^* \tilde{S}_{\sigma OX}}{J_{\sigma O}} \right) & \sum \frac{|\tilde{S}_{\sigma O}|^2}{J_{\sigma O}} + \frac{|\tilde{S}_{\sigma IX}|^2}{J_{\sigma I}} \end{bmatrix}, \quad (4.9)$$

where Re denotes the real part, and the quantity equivalent to the signal-to-noise ratio is M_σ , which divides the optimal filter terms acting together in the right hand side of Equation 4.8, and it is also the expected resolution of the measurement.

Since, in the above, both sides have their own time optimization, we can get a single best-fit by adding the time constraint separately. Both sets of equations are solved to find the χ_{side1}^2 and χ_{side2}^2 as a function of t_{side1} and t_{side2} , and then a combined (by simple addition) χ^2 is determined by constraining t_{side1} and t_{side2} to be within 10 time-bins of each other as well as within $-35\mu s$ to $+10\mu s$ from the phonon (total channel sum) simple optimal filter best-fit time [42, 53].

An example of an event, along with its best-fit optimal filter results superimposed, is shown in Figure 4.3. The figure shows the amplitude \hat{A} and the time-delay $t = 0.8\mu s$ for the minimum χ^2 , as well as the template normalized by the best-fit amplitude and delayed by the corresponding time. The event has noise but the measured answer is off by barely 0.001% from the expected value (which we know since it is a simulated event).

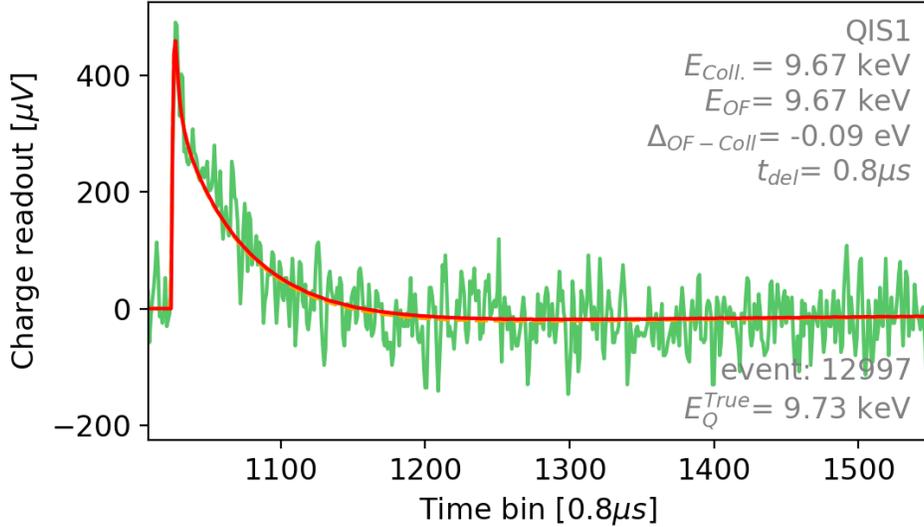


Figure 4.3: An example of a simulated event along with its best fit, showing the resulting amplitude and time delay values for this event. The fitting template, normalized by the best-fit amplitude and delayed by the respective time is in red, while the waveform (with noise) in green. Noise is de-weighted by the fit. Notice that this is a simulated event where we know the energy that the waveform should represent, and the difference between the observed and expected is 0.9 eV.

4.4 Charge Energy Observables

With the optimization now described we turn to the details of the measurement of the energy when there is an interaction, and how we measure the resolution of the detector when there isn't one. Ultimately at the analysis level, the charge and phonon observables are the essence of what is analyzed. Since this work is only focused on charge simulations, only the charge energy observables will be described. We first describe how the best-fit results are converted into our best-estimate energy measurements, and follow with presenting them. We finalize by presenting how to use the algorithm to measure the resolution in the absence of a physical interaction as it will be of use in our noise studies in Chapters 6 and 8.

Since the total expected energy is expected to be linearly related to the total observed charge, see Eqs. 3.11 and 3.12, the best-fit amplitude is turned into the best energy estimator by multiplying the measured amplitude by a constant scale-factor calibration. While the

calibration procedure will be described in Chapter 5.2.1, we note for now that it is determined by measuring interactions of samples with well known features (like the photon line with an energy of 356 keV from ^{133}Ba sources described in more detail in Chapter 5.2.1). We get one optimal filter fit per channel with amplitude (in ADC units) and each is then scaled with a calibration constant. After calibration the four energy variables are summarized in Table 4.1 and are referred to as qi1, qo1, qi2 and qo2 where q refers to the charge, the letter refers to the inner or outer channel, and the number refers to the side (with 1 denoting the top which reads out charge from the electrons).

With the independent quantities we can use their separate information to tell us about where the interaction occurred, whether the event is well measured, and we can use the combined values to reduce the uncertainty. This will be discussed in more detail in Section 9. We will refer to the two combined values as qimean (average from both side's inner channels), and qsummax (maximum of either side's sum of inner and outer channels). As noted in previous SuperCDMS analyses [47, 77], qimean is most useful for low energies (1 to 50 keV) because both sides measurements are expected to be similar (symmetric) for well measured events, and qsummax is best for high energy measurements (50 to ~ 400 keV) because the measurement from one side is expected to be better when the interaction is closer to it and we are less worried about whether there is an overall mismeasurement at high energies.

Since we will want to measure the resolution of the detector in the absence of an interaction, the resolution of 0-energy deposited, we point out some details which must be taken into consideration. In the cases described above, the optimal filter is looking at all times for the location of the interaction signal maximal value of A , so in principle there is only one correct time, but when there is no true time (true interaction) the optimal filter considers all times and essentially picks the biggest excursion it can find anywhere in the data (and we allow negative excursions). Said differently, it always picks the largest variation, which is biased compared to zero; the more times considered the larger the expected bias due to fluctuations. As shown in Figure 4.4, this shows up as a bi-modal distribution of best-fit amplitude values.

Quantity	Type	Description
qi1	I	Inner-side1 energy measurement
qo1	I	Outer-side1 energy measurement
qi2	I	Inner-side2 energy measurement
qo2	I	Outer-side2 energy measurement
qimean	R	Average of both side's inner measurements
qsummax	R	Maximum of either side's sum of inner and outer measurements
qX-0delay	I,R	All of the above with fixed (0-delay) time to measure noise

Table 4.1: Summary of calibrated energy observable quantities for the charge readout system, including 0-delay measurements for noise. Types notation: I is for independent measurement, R is for calculated, combined quantity. Side1 is the top, which collects electrons, and side2 collects holes, X (in the last row) can be replaced with any other estimator.

This is not the quantity of interest in the context of 0-energy measurements. For that reason, when looking at events with no expected signal interaction, the optimal filter time-delay is constrained to the 0-delay single value (we will refer to this time as t_0 , and pick the time when the data acquisition system typically places the pulse $t = 1024$). This measurement is denoted E_0 , it provides a true estimate of noise energy fluctuations that would affect a real measurement and is a single Gaussian distribution, centered at zero. As expected, see Figure 4.4, the distribution of energies in a sample of events with no interactions is a single Gaussian with an RMS that is determined by the resolution due to noise alone.

Now that we understand how the event reconstruction produces the energy observables, we are ready to show how the detector response is understood and validated to determine the background discrimination techniques for use in dark matter search analyses.

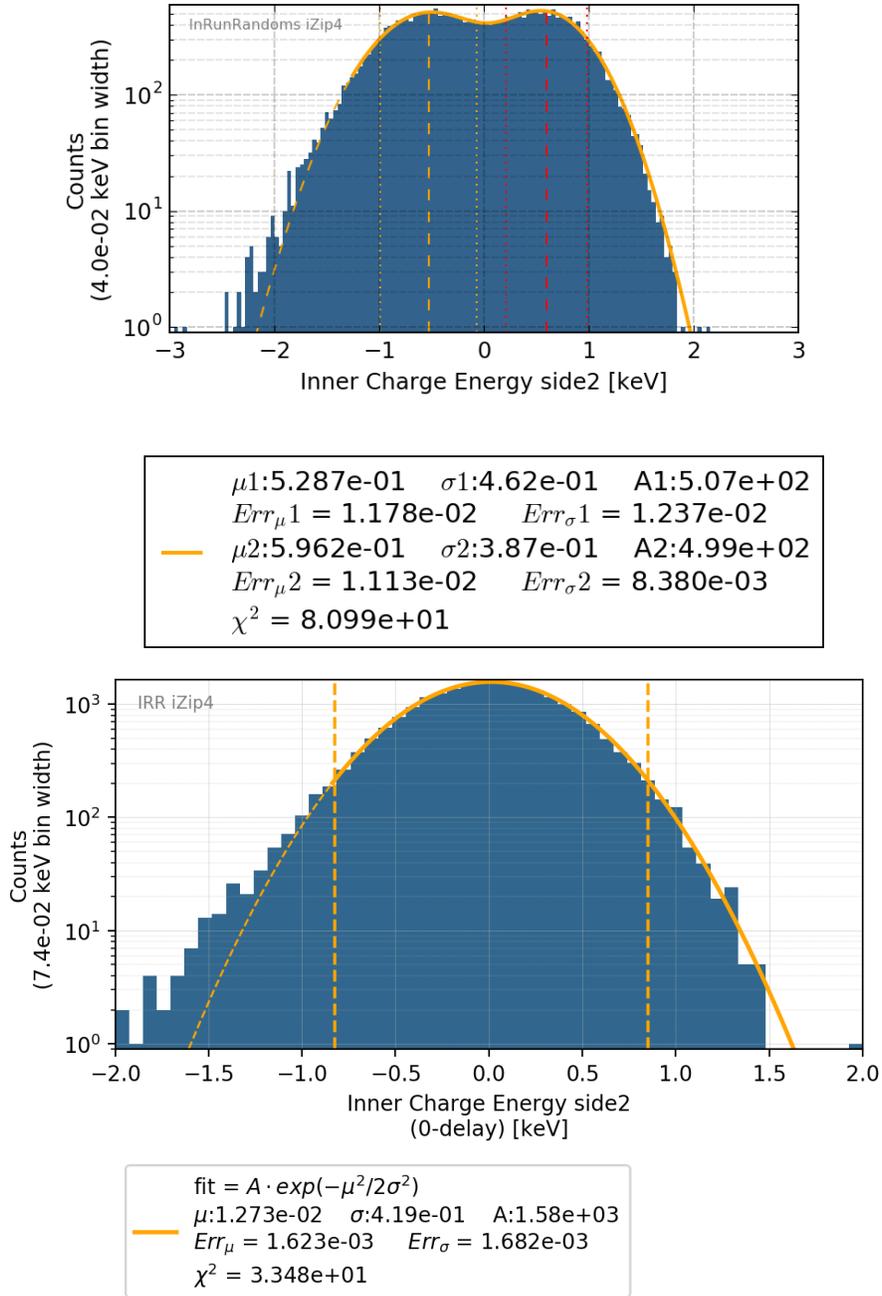


Figure 4.4: This figure shows two different ways of determining the amplitude in a sample of events with no interaction (done with noise data, which is slightly contaminated by mismeasurement effects described in Chapter 6 causing the left-hand-side excess). In the top, we allow the optimal filter to look at all times to pick a minimum χ^2 , and the result is clearly biased away from zero due to fluctuations (we allow the amplitude to be negative). The bottom plot shows the same data, but where the time-delay is fixed to be the same value for all events, the 0-delay time ($t = 1024$ time-bins). In this case the average energy is expected to be zero, and be well-described by a single Gaussian centered at zero with an RMS determined by the resolution of the noise (charge readout circuit in the absence of a particle interaction).

5. CALIBRATIONS SOURCES AND DETECTOR RESPONSE, WELL-BEHAVED SELECTION CRITERIA, PREVIOUS WIMP SEARCH METHODOLOGIES AND RESULTS WITH AN EYE TO THE FUTURE WITH SIMULATIONS

With our description of the readout and response of the detectors complete we are ready to overview the pieces we will need for analysis, previous results, and outline our next steps. There is a lot to cover including calibration sources and how calibration data is taken and used to study the detector response as well as about how we use samples of events where no physical interaction occurs (pure noise samples) to study the baseline effects in the electronics readout. With our well understood samples we will be able to set the scale of our readout, from unscaled ADC units into energy units we trust. This allows us to understand the overall response of our detector to the electron and nuclear recoil calibrations to determine criteria and requirements for the quality and variations in our detector response. With all these pieces in mind we can describe how WIMP searches are performed in SuperCDMS, what some of the latest results have shown, and how our simulations can have a major impact on future dark matter searches. This will leave us well set to describe our noise data in the upcoming chapter and our simulations in Chapter 7.

In this chapter we describe the calibration samples and how they are used to understand the detector response and define event selection criteria for use in analyses, with an eye towards previous WIMP searches in SuperCDMS and how are simulations will help optimize the searches when they are ready. We begin with Section 5.1 describing the ^{133}Ba photon and ^{252}Cf neutron sources themselves and the key features of their decay chain which make them useful for our purposes. Then, in Section 5.2 we describe the samples needed to understand our detector and simulations, including ^{133}Ba , ^{252}Cf , and data taking at random times so there is little chance of an interaction for noise studies. Since we want to only consider well-measured events in our studies, in Section 5.3 we describe event selection criteria used in the previous analyses to attempt to identify the well-measured events, which will be used

extensively in our simulations results and comparisons in later chapters. Next, in Sections 5.4 and 5.5 we give a summary of the common analysis methodology to search for WIMP interactions in the detector, and the results and summary of the previous analyses. Finally in Section 5.6 we describe how we envision the possible improvements with simulations, with an eye towards our simulation results.

5.1 The Calibration Sources: ^{133}Ba and ^{252}Cf

Since we have completed the description of our detector, readout and energy measurement techniques, we now describe how we use our radioactive calibration sources. In this section we describe why and how we use the calibrations, followed by the description of each of the two primary sources we use: ^{133}Ba and ^{252}Cf , for electron and nuclear recoil response respectively.

The goal for placing the radioactive sources is both to calibrate the energy scale of the detectors and to help us differentiate between background and signal-like events, by understanding the expected patterns from both. As described in Section 2.1, the apparatus that surrounds the SuperCDMS detectors is well set up to put calibration sources near the detectors for understanding the response to electrons, photons and neutrons, as shown diagrammatically in Figure 5.1. Unfortunately, there are many complicated effects involved that need to be taken into account if we want to understand how the emitted particles from the source produce the observations in our detectors. For example, the particles don't just deposit full energy in our detectors, but rather travel from the source, through the detector apparatus (where they can scatter and lose energy or be stopped) and then interact in the detector one or more times before either escaping or being absorbed the observed energy will look quite different than the emitted energy spectra. Untangling these differences are part of understanding the detector response. To understand the electron recoil response we place a ^{133}Ba radioactive source of photons (and electrons) with a well understood emission spectra. And to understand nuclear recoil response we place the ^{252}Cf source of neutrons. We next describe the emission processes of each in turn.

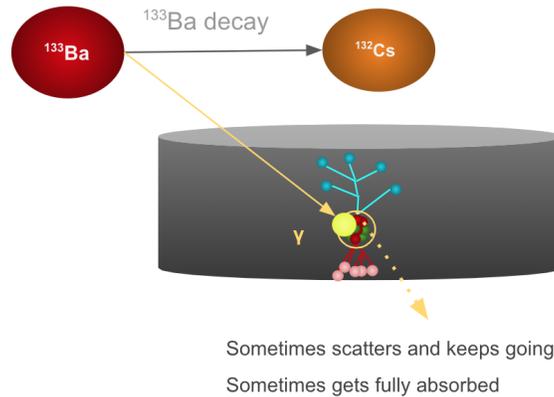


Figure 5.1: A simplified diagram of a calibration source, in this case ^{133}Ba , near the detector and how it produces photons which enter the detector along with the particles that propagate within it in response. This source is used as a calibration input for the detector. We note that the primary photons may or may not reach the detector, and can interact in the crystal or just pass by. Some of them are blocked before reaching the detector, others scatter before being fully absorbed (and their subsequent scatters can reach the detector), others can reach it without scattering, and some can pass through without depositing energy.

The ^{133}Ba isotope is a source of photons and electrons with very particular nuclear emission lines from electron capture decay [78], as shown in Figure 5.2 which is useful because these photons of known energy leave a particular signature (from the emission lines) in the detectors and for that reason they are useful as some of the features can be identified in the reconstruction. This sample is useful for understanding the electromagnetic background response and thus guide discrimination decisions (like the electron recoil band in Figure 2.2), and for setting the calibrations of the energy observables. The emitted photons eventually deposit energy in the detector via electron recoil in a process known as Compton Scattering [47, 51]. Figure 5.2 on the left describes the barium isotope energy states, and the down-arrows describe the changes in energy (the difference is the emitted photon energy). A plot of the number of photons as a function of energy is thus given by a series of delta functions at the emission energies and shown later (in Chapter 7.5.2).

The ^{252}Cf calibration source is used in similar ways and allows for studying the detector

spontaneous fission is a cascading process, the produced neutrons are abundantly emitted at all energies of interest for SuperCDMS (up to 400 keV).

5.2 Calibration and Noise Samples and Detector Response

Data samples are constructed in three different running modes:

1. Detector is readout when there is evidence of an interaction to search for WIMPs (no calibration source nearby)
2. Detector is readout when there is evidence of an interaction to study calibration samples (a single calibration source is placed nearby)
3. Detector is readout at a random time to study what noise looks like (either with or without a calibration source nearby)

These samples will help us calibrate the energy response of charge and phonon readout systems, as well as understand the detector response, other background sources which occur in time and the contributions to the energy measurement resolution due to noise. Each of the calibration samples, and readout noise is used for a different purpose given the particle interactions they produce in the detector. These samples are taken periodically to also understand variations of the detector response in time or other external variables.

In this section we describe the ^{133}Ba calibration sample first, along with the calibration method to scale the observations to energy units. This sample is of interest for our simulations comparisons and simulation calibration described in Chapters 9, and 10. We will follow with the ^{252}Cf calibration sample, which we show for completeness. We finish with the sample of events with no interactions as we will use it to understand electronics readout effects as well as provide a sample of events which can be used in our noise measurement studies. We note that we describe the contributions to the noise sample in the next chapter, as well as how we will remove contamination of the non-pure-noise sources from the sample, for use in our noise studies and comparisons to simulation in Chapter 8.

5.2.1 The ^{133}Ba Sample and Its Use in Electron Recoil Response and Energy Scale Calibration

The ^{133}Ba calibration sample produces photons and electrons with features that are highly useful for the detector response understanding and basic for the energy estimates. We will briefly describe the uses of these samples (many of which will be described later) and then describe how it is used for calibrating the energy scale next.

The ^{133}Ba calibration sample is used in a number of ways in our work. The first is that it is used to determine the constants that take us from readout units (ADC counts) to energy units (keV) for the charge system. It is also similarly useful for the phonon system calibration. The second is to define the electron recoil discrimination methods, as seen on Figure 2.2. Finally, since this sample has clearly defined features (photon emission peaks) it is used for understanding the quality of the events, and determine well-measured criteria [47, 51, 77]. In our work, we simulate the full process (described in Chapter 7.5) so we can better understand the details of what we observe in the data to disentangle the various effects.

Since the best-fit detector readout amplitude, in ADC counts, will show the peaks the energy calibration scale is readily set by a single calibration constant to convert from ADC units to energy units. The ^{133}Ba emission spectra has a strong high energy peak (at 356.013 keV) that is easily measurable in the energy observables (even without any selection criteria imposed in data) [78]. Low energy lines are blocked by the shielding and housing of the detector, and medium-energy ones are obscured by the Compton-scattering, but the 356.013 keV peak is both above the bulk of the Compton spectrum and with enough energy to (sometimes) reach the detector without scattering along the trajectory [47]. This feature in the ^{133}Ba sample is measured in the energy reconstruction, and thus allows to calibrate by positioning the peak in the expected value, and measure experimental resolution (by measuring the RMS of the peak). The peak location is estimated by fitting a skewed Gaussian function to the observed data, after a well-behaved event selection criteria (as described in Section 5.3), and finding the location of the maximum value of such fit-function. We will follow this procedure

in our simulated samples as well.

5.2.2 The ^{252}Cf Sample and Its Use in Nuclear Recoil Response

The ^{252}Cf calibration sample is used for estimating the nuclear recoil response (from the Figure 2.2) to understand the response to signal-like events. We want to achieve two things with the ^{252}Cf calibration samples: define the selection criteria for separating between electron and nuclear recoils, and understand nuclear recoil response to discriminate them (or other non-WIMP but also non-electron recoil backgrounds). In addition, after exposure to this source the Ge crystal atoms activate and become a source of well-defined electron recoils at lower energies. We describe each in turn.

The first goal is to determine the selection criteria around the expected charge-phonon yield since this is the response to signal-like events, and this is achieved by defining the nuclear recoil band as previously shown in Figure 2.2. With this (and the electron recoil band from ^{133}Ba) we can readily discriminate between electromagnetically interacting particles and our WIMP-like events.

Second we wish to be able to identify neutrons from background sources in our dark matter search data and we use the ^{252}Cf source to create samples of such events to recognize how to remove them. This is complicated because in many ways the neutrons (and other nuclear recoils) are very similar to a WIMP interaction. The neutron-nuclear recoils from this source both cover and vastly surpass all energy ranges of interest for SuperCDMS (0 to ~ 400 keV) [79] so it is useful for estimating the response of any WIMP-like interaction.

Finally, the detectors exposure to the ^{252}Cf emission produces an activation effect in Ge, which is also useful for studying the response to electron recoils in a different energy range (10 keV). The exposure to this source turns some of the atoms in the Ge lattice into ^{71}Ge , which produces photons (x-rays) similar to the ^{133}Ba source, with clearly defined peaks useful in the same ways as the barium spectrum.

5.2.3 The InRunRandoms (No Interactions Sample) and Its Use In Measuring Detector Readout Noise

Data samples taken when no sources are nearby, and selected randomly in time, known as InRunRandoms (IRR), or simply as a noise sample, are useful in many ways. These include noise studies, understanding the baseline detector resolution, making PSDs for the reconstruction expected variations, as well as understanding electronics readout problems. We first describe what noise data is and what it is useful for. Finally, we describe how noise samples are created and the types of subsamples produced with them. With noise samples we complete the descriptions of the types of samples taken for understanding the detector response and move on to describing how event selection criteria are generated.

Samples of events with no interactions are recorded at random times before, after, or inter-spaced during any other type of data-taking mode, and is useful for quality control and baseline resolution estimates. Since the waveform readout the detector in the absence of a particle interactions is expected to be dominated by electronic noise, creating a sample of ‘noise events’ is useful. We will refer to this as noise data, and a sample of these events is known as a noise-only sample. Noise data is useful for measuring the estimated variation in the detector response, understanding it from the absence or physical interactions and extrapolating to the actual measurements gives us the baseline detector resolution. This task is particularly important at lowest energies (our region of most interest) where noise dominates the observed resolution. Finally, they are used to generate noise profiles (PSDs) for event reconstruction and to understand the baseline noise in the detectors, which is mostly from the readout electronics.

Noise samples are created by periodically reading the detectors’ response in the absence of physical interactions, as shown in Figure 3.5 bottom plots. For that, an event is recorded randomly in time (as opposed to being recorded after a signal threshold). Data is commonly divided in data-taking periods labelled runs. Noise samples are taken in 3 different ways: before the run (named BeginningOfRunRandoms), after the run (EndOfRunRandoms), and

randomly during the run (InRunRandoms). Typically the BeginningOfRunRandoms are used for data quality monitoring and for the event reconstruction algorithm (creating a noise profile that serves as an estimate of the variations in the signal). The InRunRandoms (IRR) are often used for standalone studies to understand changes in the response of the detector and in the environment conditions and we will use a sample of these events, in conjunction with simulated versions of them, extensively in Chapter 7. While we will not consider EndOfRunRandoms in our work, we note that they are useful for comparing with BeginningOfRunRandoms to search for evidence of the data taking conditions having changed during the course of the run.

5.3 Well-Behaved Event Selection Criteria: Low Threshold (LT) Fiducial Requirements

With an understanding of how the detector is expected to respond to interactions, we are now set to describe how we to determine criteria for selecting events which are likely to be well measured, or are easily corrected. We will use this criteria throughout our charge measurement studies to compare well-measured events in simulations and data and to look for mismeasurement causes. We describe the only event selection criteria used in this work first, which is related with the well-measured part of the detector, as we list its requirements. We next describe how these are tuned from the calibration samples. We note at the outset that these criteria are based only on observables from the detector; the limitations of these criteria will be studied with simulations and described in Chapters 9 and 10.

The most common, and the only event selection criteria used in this work, is the set of requirements to select events used in the Low Threshold (LT) WIMP search Ref. [47]. Since these criteria are designed to only select the well-measured events, presumably because they are in the well-measured, or fiducial, portion of the detector we will refer to them as the LT Fiducial requirements. Later in this work (in Chapter 9) we will compare with simulations, to show that the LT Fiducial cut does a good job at rejecting most mismeasured events, but can be improved. The subject of this work, however, is not to optimize the selection

but to show that simulations give insight to improvements in these analysis decisions (even though our simulations still need tuning and full validations). The criteria are composed of 2 requirements that are expected to relate with fiducial volume, which are:

- No energy (above noise) recorded in outer channels: keeping only events consistent with zero energy in outer channels (lower than the maximum expected from noise variation, 1 keV). It is expected that events where the interaction occurs near the side-wall of the detector are mismeasured and deposit more energy in the outer channels [47, 51, 53]; the radial partition of the channels (Figure 2.5) was designed for this purpose
- Symmetric measurement of inner channels: since we expect the amount of energy measured on both sides of the detector to be equal, significantly different values between them is an indication that at least one is mismeasured (i.e. the event didn't occur in the well-measured portion of the detector) [47, 53]

Collectively these are known as the LT Fiducial requirements.

The parameters for both LT Fiducial requirements are determined from ^{133}Ba calibration samples, by observing the response of energy measurements and using algorithms to determine the requirements [53]. While we will show more on this in Chapter 9, we simply note that the no-outer measurement requirement is tuned from fitting the charge response near-zero energy and rejecting events above 2σ from noise fluctuations (and then simplified to 1 keV for iZip4). The symmetric cut is also created from ^{133}Ba data, by defining a 2σ band from the symmetric mean of side1 and side2 inner channels. We will show the values of the requirements in Chapter 9 when we compare with simulations.

5.4 Outline of WIMP Searches with SuperCDMS Soudan

With the tools from previous searches in hand we are ready to give an overview of the previously completed WIMP searches as performed by SuperCDMS. The previous WIMP searches used the above tools to understand the detector response, define the selection criteria for selecting well measured events and for discriminating between electron and nuclear recoils.

With these in mind we begin by giving an outline of the WIMP searches. We next describe the three major selection criteria levels:

1. Operation decisions, data quality, and apparatus malfunctioning
2. Identification of particle interactions with detector response
3. Background motivated event selection and background estimations

This will set up for giving an overview of previous SuperCDMS analyses with an eye towards the improvements we hope to achieve with the simulations in the next section.

To overview the search process we begin by noting that the goal in a generic WIMP search analysis is to simultaneously maximize the number of signal events while minimizing the number of background events. The event selection is based on the knowledge acquired from the calibration samples where the response is understood, and from detector physics intuition and experimental conditions observations. The basic procedure starts with selecting a sample of events where a significant amount of energy is deposited and well measured, and then selecting a subsample of WIMP-like events. As shown in Chapter 2.1, by better understanding what backgrounds and signal look like, measurement techniques can be developed to better separate the two types, like the charge yield as function of phonon energy from Figure 2.2. A background and signal rate estimation model is constructed to make final event selection criteria, typically with sophisticated machine learning algorithms. Previous analyses have done that with real data only, from calibration sources, and sometimes aided with stand-alone simulations [47, 77]. The long term goal of this work is to base the analysis in simulations insight and understanding which is the topic in next chapter. We next describe the three levels of selection criteria that have been common in previous analyses.

In a common analysis framework the first level of selection criteria is related with operations decisions, apparatus malfunction, data quality monitoring acknowledgement, etc. Since this work is not really about a fully optimized analysis, but rather about inspecting the behavior of the simulations those will not be described. More detail can be found in

Refs. [46, 47, 51, 53, 77]. These cuts are usually referred to as ‘good quality’ event selection cuts, and the goal is that if an event passes all these requirements then we have the expectation that the detectors and apparatus are functioning as expected.

The second layer of cuts is motivated by the actual response of the detector from different calibration samples. Statistical tests, energy response, the nuclear and electron recoil band are defined (as shown in Figure 2.2), and criteria that is thought to resemble a fiducial volume selection, like the LT Fiducial described above. These cuts can be referred to as ‘well-behaved’ and ‘model motivated’ cuts.

The third layer is a background-motivated event selection and background estimations. A signal and background model that describes the expected number of events of each of the sources in Table 2.1, as well as the expected response from each source is needed for this, as it is based on a sophisticated method for doing a final event selection. An example which is based on using the expectations from background types and signal types that separates between the two on an event-by-event bases is the Boosted Decision Tree [47], where for every event the BDT determines a single value based on its likelihood of being signal or background. Events with a value less than zero are background-like, while events with a value above zero are signal like. Combining the expectations from the BDT for each background produces the expectations for the final sample of events, these expectations can be used for the final criteria optimization or used in the sensitivity limits calculations, as was done with an Optimal Interval Method on the observed events [47, 51].

5.5 Previous SuperCDMS Analyses

With all the applicable techniques and methods in place we now describe the previous SuperCDMS analyses results. First we give an overview of the former results, along with the methods followed in the context of using simulations, and how we see room for improvement with them. We next describe the background model proposed by the Low Threshold (LT) analysis and the resulting WIMP-like events which were not conclusive evidence for dark matter interactions due to unconsidered background sources, again an eye towards the

relevance of robust simulation-based methods.

The SuperCDMS Soudan Experiment has conducted several important searches for dark matter over the last decade [47, 51, 77]. As previously shown in Figure 1.11, dark matter has not been discovered but the limits keep getting pushed to lower cross-sections. All analyses background discrimination techniques are based on selection requirements to reject the background from calibration-based observable responses. None of these analyses have been simulation-based, but simulations have been used to make some decisions in the event selection in the development of a Boosted Decision Tree (BDT) discriminator quantity for separating between signal and background [47]. We simply note that on the outset the event selection was not completely understood, leaving room for improvement, in particular with the use of simulations.

In the Low Threshold analysis (LT) a background model was used to create a BDT score for event and create the expectations from all the sources of background that were expected. The expected final distribution, along with the expectations for signal events and the observed data are shown in Figure 5.3. A number of observations are in order. The first is that from the expectations of the signal and background sources described in Chapter 2 we can see that the bulk of the background distribution is well estimated. The second is that there are a number of events that don't appear to be consistent with the background model, but are consistent with a signal model. While we will not go into detail, a study of each event showed that it was not that they are signal-like, but rather that they are from a type of background that was not included in the analysis (but now described in Chapter 2).

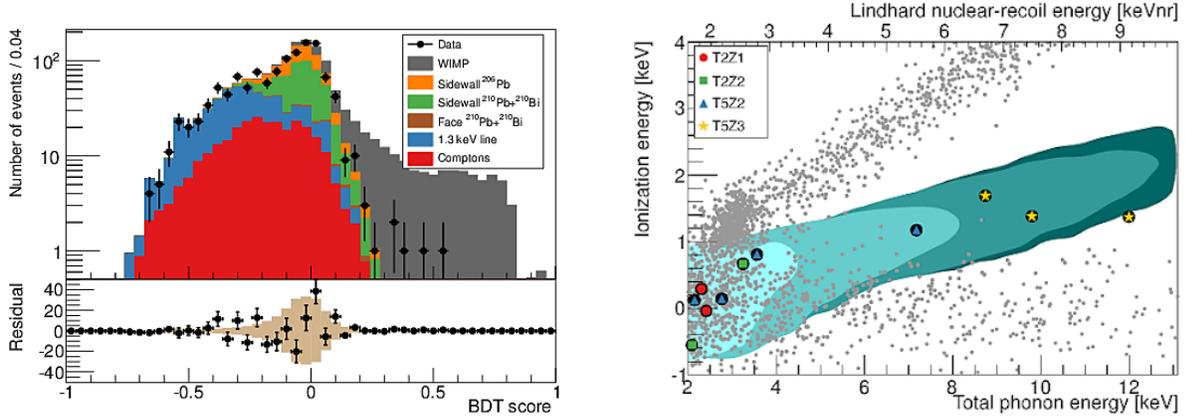


Figure 5.3: Two figures showing the expectations and results for the SuperCDMS Low Threshold (LT) analysis, reprinted from Ref. [47]. On the left is the background model. On the right is the observed signal events. After the quality and basic event selection criteria, the background model was used for constructing an arbitrary parameter to differentiate signal from background. The final result determined that 11 events were signal-looking, but upon close inspection they were related to noise and other types of possible fluctuations.

5.6 Analyses Improvements with Simulations

Looking forward, the goal of this work is for the simulations to be used as a tool to improve the selection criteria and then for an optimal discrimination method. To conclude this section we first give an overview of the improvements for pushing the dark matter searches further and how simulations may lead the way in the future. Finally, we describe the two major improvements that simulations will help achieve as we ready for our simulations infrastructure description in the next chapter.

The improvements required for pushing the dark matter searches further are for a better modeling of the backgrounds and better methods for signal versus background discrimination. Simulations may well lead the way. Of course, discriminating for such faint signals is not straight forward as mismeasurement backgrounds come into play. Multiple effects need to be accounted for to only select a good sample of events before identifying them. As we have shown, in the past the selection criteria has been mostly guided through detector research, development knowledge, and good intuition from observations in the data, but the aim for

the future is to have trustworthy simulations that help optimize the selection criteria and create a well-understood background model.

With fully trusted and validated simulations in hand, two major improvements can be done: first a trustworthy background model for the sophisticated final selection algorithms and background estimation methods, and second an optimized event selection criteria after enhanced understanding of the detectors response. In Chapter 10 we describe in much more detail the current status and goals for a simulation-based optimization to our future WIMP searches.

With our calibration techniques, detector study samples, and search strategies described we look forward. The goal of this work is to provide a common ground to trust simulations of the SuperCDMS Experiment. We will describe the simulations and the infrastructure in Chapter 7, but first we will expand on causes of mismeasurement due to electronics sources which hurt the previous analyses, while we clean up the detector noise sample for use in our comparisons with our noise simulation.

6. DETECTOR NOISE SAMPLE, READOUT MISMEASUREMENT EVENTS AND THEIR REJECTION AND CORRECTION

Now that we know how the detector is supposed to respond in the absence of an interaction, we are ready to talk about the noise sample, how the readout is contaminated by various effects and how we will deal with them. As previously described in Chapter 2.1, readout mismeasurement is a relevant source of background events because it can make SM interactions fake our WIMP signal (as happened in the previous analyses events shown in Figure 5.3). While it is not always possible to know when this happens, we use our noise sample to identify patterns of mismeasurement, as well as suggest rejection and correction techniques. We will focus on the readout mismeasurements that show up in our data (pileups and Blips) and how we will handle them as we create a pure noise sample for use in comparisons. With this sample we will be well-set to make the proper input to the simulations and our simulation comparisons in upcoming chapters.

In this chapter we describe the detector noise sample, readout mismeasurement events and their rejection and correction. We begin in Section 6.1 giving an overview of the known cases of readout failures, how they affect our measurements, and how they are dealt with, with an emphasis on pileups and Blips (previously described in Chapter 2.3.3). Then, in Section 6.2 we characterize our InRunRandoms (IRR) sample, and show how it can be well-described as basically being a Gaussian noise event sample contaminated by pileups (which we can remove) and Blips. Next in Section 6.3 we will first introduce our BlipFinder tool and describe its results on iZip4 Blips (biggest Blips in the entire events waveform), and then use it to identify and measure events which show evidence of Blips in multiple detectors in the same tower known as TowerBlips, specifically from iZip5 and iZip6. Finally, in Section 6.4 we describe the method for correcting the energy measurement in iZip4 when there is evidence of a TowerBlip in iZip5 and iZip6 and use it to produce our corrected IRR sample for use in comparisons with the simulation in Chapter 8.

6.1 Overview of Readout Mismeasurement Sources

Even though only pileups and Blips show up in the IRR sample we will use, we will give an overview of all sources of readout mismeasurement for completeness. The full list of readout mismeasurement sources was given in Table 2.1. In this section we start with an overview of each and describe each in turn. We note before beginning that while we have already described how to deal with cross-talk, and pileups treatment is straight forward, our understanding of Blips is mostly new so we will focus on them. On the other hand, since Glitches and Low Frequency (LF) noise don't occur in our sample we will not describe them in detail. Our emphasis will be on characterization and correction of TowerBlips for the creation of our corrected IRR sample.

While we will always have noise in every event, there are 4 other important sources of readout mismeasurement. They are::

- Cross-talk: the observed effect that when energy is measured in one channel it causes a readout response in it's neighbor
- Pileups: when a particle interaction in the detector causes an additional pulse during the recorded waveform
- Blips: negative spikes with a rise-time of approximately $50\mu\text{s}$ in length followed by a tail of $300\mu\text{s}$ which, for a short period of time, goes positive before asymptotically going to zero. They correlate in all channels of a detector (Single-Detector Blips), and most often also correlate between all iZips of the same tower (TowerBlips)
- Glitches and Low Frequency noise (LF): when short spikes or specific patterns occur in the observed waveform. Glitches are short spikes in the readout less than $40\mu\text{s}$ in length, and LF noise are events with a wavelength near the full-span of the readout

The first mismeasurement cause, cross-talk, occurs in every event and is dealt with by the algorithms of the event reconstruction, as described in Chapter 4. We have seen that the

event reconstruction implementation takes the signal and cross-talk components into account to measure the best-estimate energy of the event. Moving forward we will no longer consider this source of mismeasurement as its effects are negligible after taking it into account.

The second case, pileup contamination, is due to a real particle interaction occurring and depositing energy in the detector at a random time during the waveform readout. The two primary cases are when two interactions occur during the same event, or when an interaction occurs when we are trying to collect an IRR event to study noise. In both cases the effect is that there can be a bias towards an over-measurement of the energy in the real region of interest. They are a highly common mismeasurement cause during high-rate calibration data taking, but rare during WIMP-search data taking. Figure 6.1 shows IRR events with pileup contamination. In WIMP searches, the closer the pileups occur to the WIMP interaction, they become more damaging to the measurement. For noise studies, the closer it occurs to the 0-delay time, the more bias it will have on our 0-delay energy measurement, E_0 . These events are discarded at the analysis level, with event quality cuts. In Section 6.2.2 we will show how they affect our IRR data sample and how they are removed.

The third source of mismeasurement, Blips, appear to be an electronics issue as they are correlated between the channels in a detector and sometimes between detectors in a tower. Their true origin is unknown as they were only uncovered during the studies presented in this dissertation. They can be categorized as Single-Detector Blips and TowerBlips depending on the number of iZips they affect. While Blips can occur randomly in time they have a common shape, as shown in Figure 6.2, where the average pulse shape is in black and the 1σ deviation in orange (we will describe the creation of these distributions in Section 6.2.3). They are a problem when they occur near the time of an energy deposition, and are a dominant feature in our noise studies. One of the things that makes the correction possible in our data is that they appear to always be correlated between channels in a detector, and often correlated between detectors in a single tower. For this reasons we will refer to the later case as TowerBlips. This background will be studied further in Section 6.2.3.

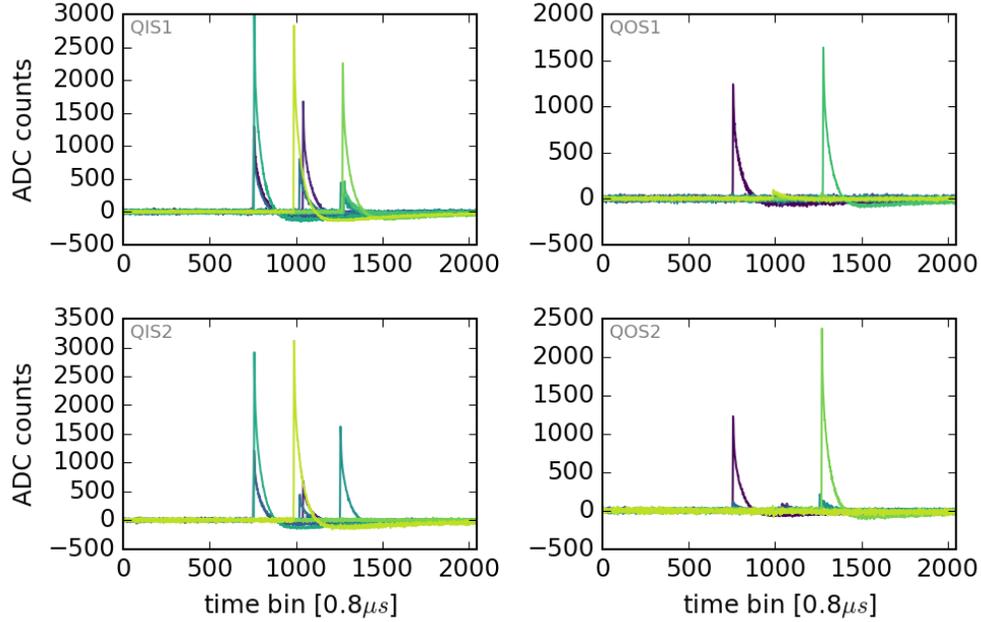


Figure 6.1: Pileup examples in noise-data events in the four different charge readout channels. The set of events shown were in noise-collection mode (random readout time) but when the ^{133}Ba was placed near the detector and a large E_0 energy is recorded.

The fourth case consists of Glitches and LF noise. While they do not appear to affect our studies we summarize them here for completeness (for more information see Ref. [53]). As with any effect that produces ADC counts above baseline they can damage the energy measurement when they occur near the interaction time. They need to be removed at the analysis level, since there is no known way to prevent them from happening. As shown in the left of Figure 6.3, Glitch events are an exotic pulse shape in the charge and phonon channels, characterized by a rapid rise and fall compared to a typical interaction pulse shape. They occur in multiple detectors at the same time, and have been shown to be associated with spikes in the bias voltage of the phonon circuit [53]. A special fitting algorithm is devised for Glitches and a simple goodness-of-fit (χ^2) constraint efficiently removes them. Low frequency noise is another similar source, shown in the right of the same figure. Mechanical vibrations in the experiment are known to produce this type of pattern in the phonon readout output but these are not observed in the charge readout. The same approach as for Glitches is

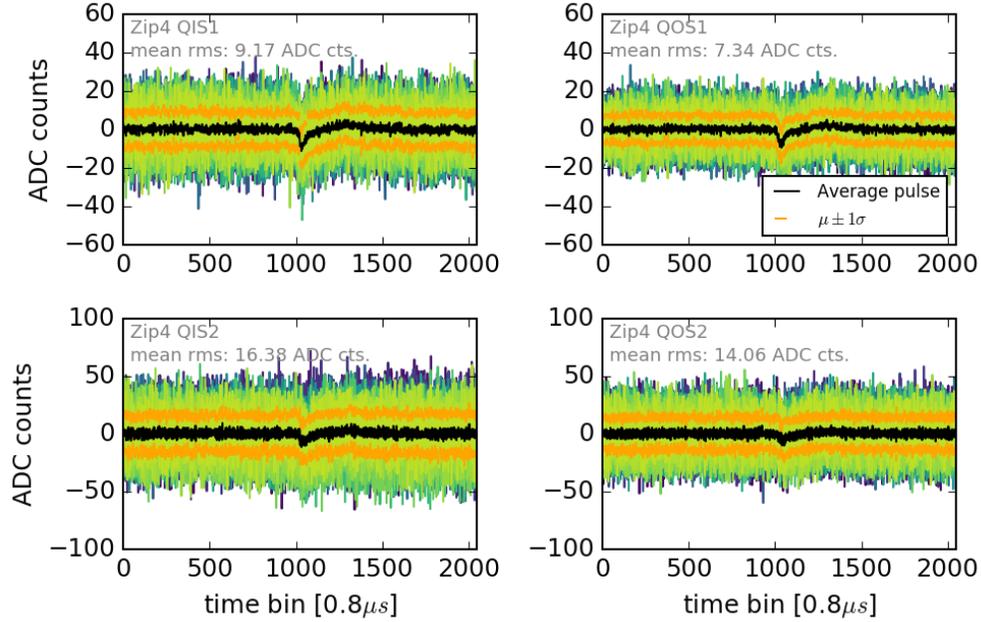


Figure 6.2: Example Blip events which show small signals in the waveform as a function of time. These figures show the readout from each of the four channels from a noise sample. Individual events are shown in blue, green, and yellow. Given the low amplitude of the Blip-feature it is easier to see the average Blip shape (black). The orange band shows $\mu \pm 1\sigma$.

followed, and while not being perfect most of this background is removed in all analyses. For our studies, we consider backgrounds of these types to be negligible.

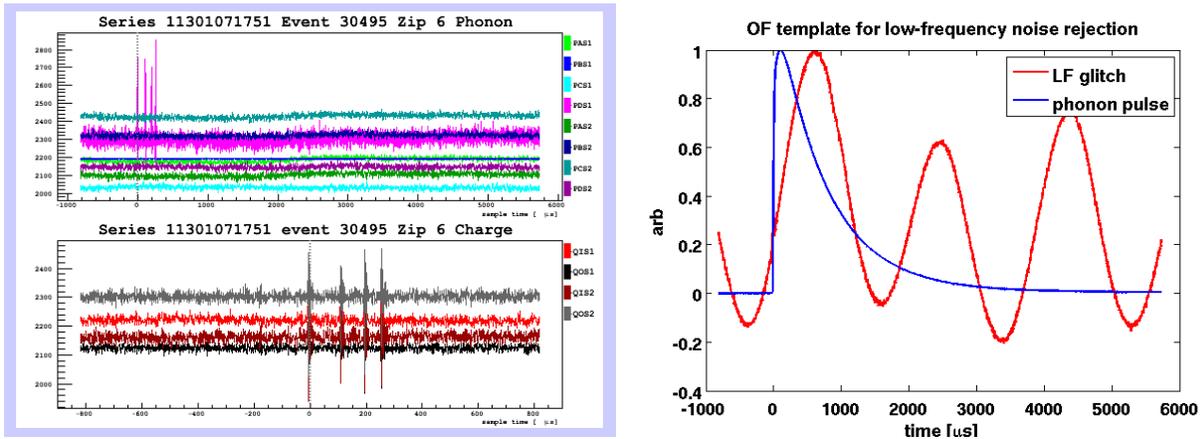


Figure 6.3: Example of a Glitch readout malfunction waveform (left), affecting both charge and phonon readout, and a low frequency noise example affecting a phonon event (right).

6.2 Characterization of the InRunRandoms Sample: Noise Events Contaminated by Pileups and Blips

Now that we have described the dominant sources of contamination to the detector readout and how they can cause energy mismeasurement, we are ready to describe their contributions to our IRR sample. Events are collected using the method described in Section 5.2.3, and a number of methods are used to clean the sample and correct the events leaving only Gaussian noise for our comparisons with simulations. In Section 6.2.1 we will first describe the IRR sample straight out of the box to characterize its non Gaussian features. We will proceed with detailing each, starting with pileups and their removal in Section 6.2.2, and with the description of Single-Detector Blips and TowerBlips in Section 6.2.3 with an eye towards how we will correct for them in upcoming sections.

6.2.1 InRunRandom Sample Creation and Energy Distribution

Our IRR sample can be thought of as a sample of noise events contaminated by pileups and Blips which affect the energy estimates. We begin with a brief overview of our data and how we will use it later, we will then show the IRR energy observation straight out of the box, to show how the measurements are affected by the pileups and Blips, as we ready for removing pileups in the upcoming section.

Since we will be using the sample of IRR events and using it to validate our noise simulations and measure the baseline energy resolution of the detector, we need a more detailed description of the contributions to the sample. Before beginning we note that our IRR sample was intentionally taken during a ^{133}Ba calibration period because we expect it to be more susceptible to showing large readout issues since the barium calibrations have the highest interaction rates. For our studies we will consider events from tower 2, which includes iZip4, iZip5, and iZip6, and present data from iZip4 unless otherwise specified. After cleanup, our cleaned IRR sample will be used in Chapter 8 for comparisons in our simulation, but for now we will show how to remove its readout mismeasurements and contamination.

The 0-delay qimean energy measurement for the full IRR sample is shown in Figure 6.4. The dominant feature in the figure is the peak near zero which is from pure noise events. As we will see this portion is mostly Gaussian and centered close to, but not at zero for reasons that will be described soon. The next most dominant feature is the right hand side tail, which is due to pileups. We remove them with a simple raw-pulse maximum value (over baseline) threshold cut (in Section 6.2.2), similar to that which was done in previous analyses [53, 47]. The tail on the left hand side (LHS) is from Blips. While we can remove them it turns out it will be more effective to correct for them with estimates from the correlated measurement from detectors of the same tower (described in Section 6.2.3). This, as we will see, is also the cause of the mean shift of the main distribution away from zero. We describe each in turn.

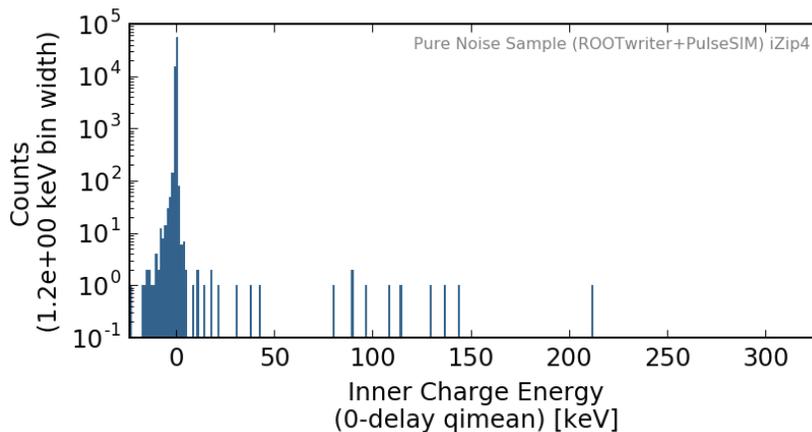


Figure 6.4: This plot shows the 0-delay qimean energy measurement from the full set of events in the IRR sample taken when the ^{133}Ba source was placed near the detector. While the bulk of the distribution is from readout noise (main peak at zero energy) it is contaminated by Blips and Pileups. The high energy events are due to pileups, while the non-Gaussian left hand side tail is from Blips.

6.2.2 Pileup Event Removal

Since pileups are from real particle interactions that occur during the readout time, they are readily removed using standard techniques [53, 47]. As shown in Figure 6.1, pileups are always positive. We begin with a description of our modified version to remove them, and then we will show the IRR sample after their removal.

Since we expect the noise to only have small energy fluctuations around zero, and a pileup event should have a large energy deposition somewhere in the event, our procedure to remove pileups is a straightforward set of steps to remove any event where any of the energy bins in time is above a threshold picked to well separate noise from pileups. The first step is to remove the baseline from the event (like those in Figure 6.1), which is done by computing the average number of ADC counts in the first 500 bins and subtracting it from all bins in the waveform. Then the value of every time bin is checked, and the maximum value is recorded, we call this Pulse max. Since a threshold must be picked, we preemptively use our noise simulation to create a sample of noise events and determine Pulse max from it. Figure 6.5 shows the distributions for the Pulse max values from events in the simulation and IRR samples. By comparing the two by eye we can see that the data is well described by a combination of noise events and pileup events. For simplicity, we have picked a threshold of 150 ADC counts as it keeps 100% of the simulated events. To remove any possible inter-detector contamination similar methods are used to identify thresholds for all channels for all detectors, including for the phonon readout channels. Any event is removed if any channel measurement exceeds the threshold.

The results, after removing all pileup events, are shown in Figure 6.6. We see that all the high energy events have been removed and the 0-delay energy distribution is nearly Gaussian. On the right hand side (RHS), there is no tail and the Gaussian function is a good fit to the data. There is a non-Gaussian LHS tail however, which will be shown that is due to Blips. Taking a closer look at the mean offset of the distribution, we note that it is not zero (many Err_μ above zero). As we will see this effect is also due to Blips. Now that the contribution

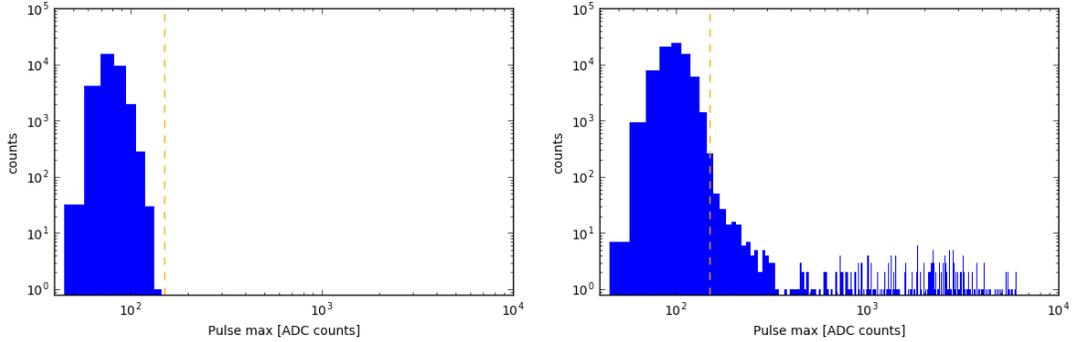


Figure 6.5: The distribution of the maximum value (over baseline) in ADC counts, called Pulse max, for QIS1 for iZip4. The simulated noise events (using the methods of Chapter 7) are shown on the left and IRR data is shown on the right. Since the data appears to be well described by a combination of noise and pileup events we set the threshold for removing pileup events at 150 ADC counts (dashed orange) which is the lowest threshold that does not remove simulated events.

from the pileup background has been removed, we can now proceed with the description of the second dominant background in the IRR sample: Blips.

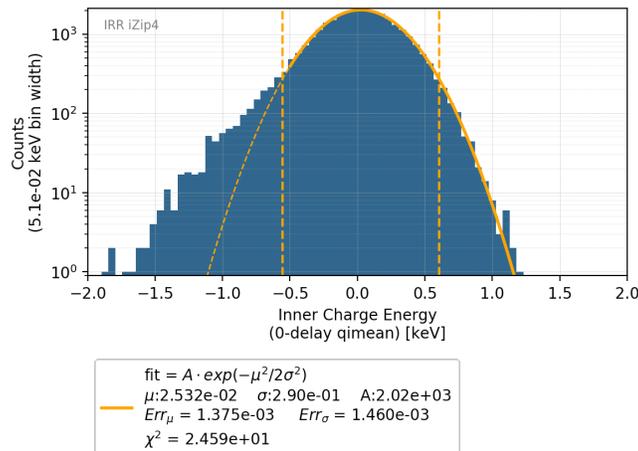


Figure 6.6: The plot shows the E_0 distribution of the subsample of IRR events after pileup removal. The data is purely Gaussian on the right hand side, but still has a left tail from Blip contributions. We also note that the mean is not zero, which as will be shown later is also due to Blips.

6.2.3 Single-Detector Blips and TowerBlips Description

The remaining events in our IRR sample after pileup removal are pure noise and Blips. Since Blips are a more complicated background, and they affect our energy measurements in a number of ways, we will describe them in this section, and in Section 6.3 we describe the BlipFinder tool we have created to look for them in the readout data, the different ways Blips affect our energy measurement, and how we will use the BlipFinder tool to understand the rates of correlation between Blips in a tower, or TowerBlips. Next, in Section 6.4, we describe the modified version of the BlipFinder for TowerBlips, TowerBlipFinder, and how we use it for correcting the energy measurement from them, based on the BlipFinder tool. After this, will have a clean enough sample of IRR events so that they can be used to estimate and simulate the baseline resolution of the detector and allow for full measurement of the PSDs for use in the simulations which is described in Chapters 7 and 8.

Before getting into our BlipFinder methods, tools and corrections, in this Blip overview section we begin by describing the creation of the sample of Blip waveforms that were shown in Figure 6.2. We follow by showing that they are correlated in all channels of the same detector and that most often they are also correlated in a different detector of the same tower (TowerBlips). With this vantage point at the outset, we will have shown motivation for finding them and correcting for them.

While the shape of Blips was shown in Figure 6.2 we now comment on how plot was made. The sample of events was drawn from the E_0 distribution, in particular all events with $q_{\text{imean}_0} < -1.2$ keV from our IRR sample (all of the LHS events from Figure 6.6). This is our clean sample of Blip-affected E_0 events and we will refer to it as the LHS sample. We will use the LHS sample next to study Blips, in particular to study the correlations in the measurements of all channels.

Using our LHS sample we observe energy in all four channels of the same detector. Figure 6.7 shows that the energy measurement for all four charge channels for all the events in our LHS sample. The energy measurement in all channels is negatively biased (actually

correlated), showing that the Blip exists in all channels and with the same polarity; we see no such correlation with the phonon channels. This observation will be used in the creation of our BlipFinder tool.

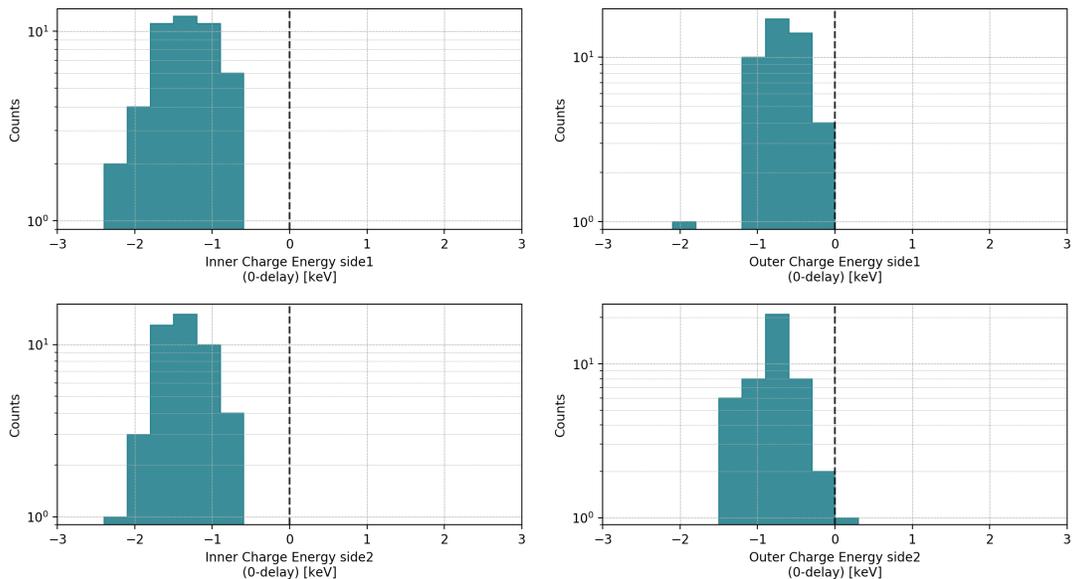


Figure 6.7: This figure shows the 0-delay energy measurement of each of the 4 charge readout channels for all the events in the IRR sample with $q_{\text{imean}_0} < -1.2$ keV. The measurements in all channels are correlated. In all distributions the sample is negatively biased, suggesting that the Blips are real, correlated, and not a random fluctuation.

When Blips occur we measure negative energy in either a single detector in the tower (Single-detector blips) or in more than one detector in a tower (TowerBlips). Figure 6.8 shows the energy measurement in iZip5 for same sample shown in Figure 6.7 (of events which were selected as having significant energy in iZip4). The average value is clearly biased (with some channels correlated and some channels anti-correlated), and shows that the correlations between detectors occurs frequently. Similar, but not identical, results are observed for iZip6. They are our first evidence for TowerBlips and this effect will be used to create our TowerBlipFinder.

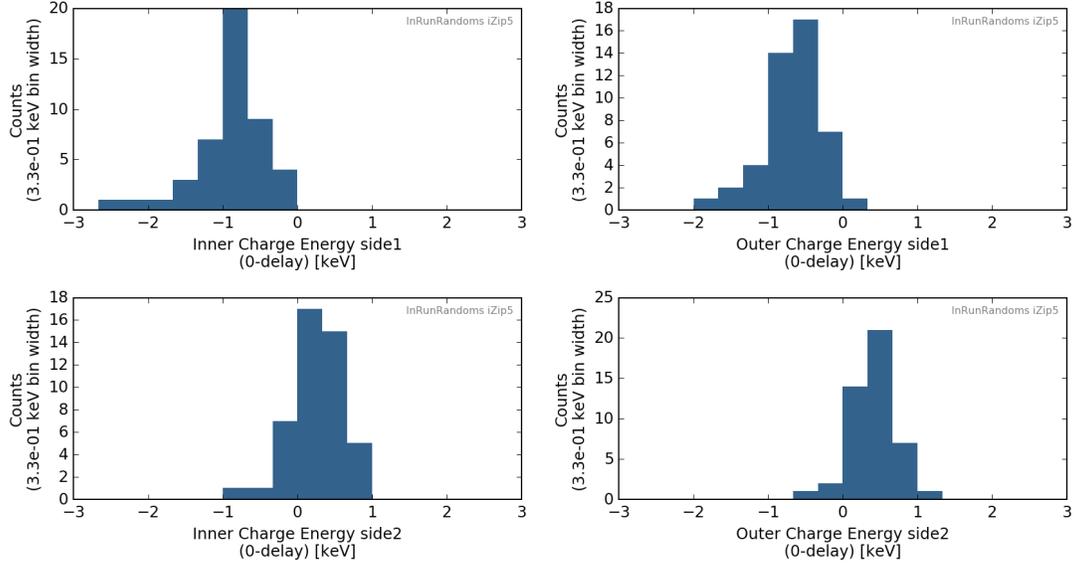


Figure 6.8: The measured 0-delay energy in all iZip5 channels when there is evidence of a Blip in iZip4. This set of events, selected using the LHS sample, provides evidence that the LHS sample is correlated in iZip5-side1 and anti-correlated in iZip5-side2 channels. In all distributions the sample is biased away from zero, suggesting that these Blips are observed in both detectors.

6.3 Finding and Measuring Blips and TowerBlips

Since, as we will show, it is very inefficient to remove Blips from the sample, our goal is to use the correlations in other iZips to correct our energy measurement for the presence of Blips (described in Section 6.4). Before doing that, we describe the tool to identify and measure them, known as BlipFinder. In this section we first describe the BlipFinder tool in Section 6.3.1. Next, in Section 6.3.2, we describe how the BlipFinder is used to look for TowerBlips with an eye towards using the iZip5 and iZip6 information to correct for Blips in iZip4 in the next section. We note at the outset that most Blips are in fact TowerBlips; for more information on the rates of occurrence see Appendix B.

6.3.1 The BlipFinder Tool

Now we present our BlipFinder tool which identifies and measures Blips in our readout waveforms. We will first describe the intended use cases. Next we describe the assumptions

prior to an overview of its functionality. We will follow with the description of the BlipFinder algorithm. And finally, we will show the results for looking for the biggest Blip in our detector (iZip4) in preparation for corrections in the upcoming sections.

We use the BlipFinder for three purposes: identifying the biggest Blip first, identifying if a Blip occurs near the biggest Blip in another iZip, and identifying the Blip that affects most the t_0 energy measurement. Depending on the intended use we can select the detectors to consider (single detector, or multiple detectors). We can also consider searching in different time windows (near t_0 , full window, or near the BlipFinder time of another detector), to look for the biggest Blip in a pulse, or the most offending Blip near a specific location in a narrow window (for example look for an offending Blip in one detector given the biggest Blip in another one). The former two are for characterizing Blip rates and correlation, and the latter for correcting for TowerBlips as will be described in the next section.

The BlipFinder is designed to identify Blips in a detector, or group of detectors, and works using a modified version of the optimal filter, with four assumptions about Blips (following properties identified in the previous section):

- They are correlated across all channels in a detector
- They are negatively biased
- They can occur at any point in the readout time
- Sometimes they occur in all iZips of the same tower

With these ideas in mind we overview the algorithm, and then give the steps. The BlipFinder uses the optimal filter methods described in Chapter 4 to find the the maximum negative excursion in the combined set of charge channels (as opposed to looking for the minimum χ^2). It gives the Blip time (location in the event), energy, and the optimal filter χ^2 value. The signal templates are the same used for event reconstruction (Chapter 4.2). Since the Blip search is enhanced by the addition of each channel's optimal-filtered waveform,

because correlated Blips add constructively, while Gaussian noise cancels, instead of doing all channels separately, in this case we add them up. Since we know that Blips correlate in time and shape in at least all channels of a single detector, we add-up the optimal-filtered term (optimal filter value as function of time-delay) of all channels of the same detector.

The steps of the BlipFinder Tool algorithm, applied in a given iZip, denoted A , are listed below:

1. Obtain optimal-filtered energy as function the location of the time-delay of the optimal filter template for each channel
2. Add all correlated optimal-filtered channels
3. Define a search window
4. Look for biggest-negative energy value, and obtain BlipFinder time t_A and energy E_A

To show how well the BlipFinder works we show some quick results to understand how the biggest blip in the events affect our 0-delay IRR energy measurement. We use the BlipFinder in the iZip4 4-channel readout, looking in the full time-window. Even though we have not described the pure noise simulation we will show the result of the simulated sample as evidence that there is no bias in the detection technique (since the distribution is known to be a pure Gaussian sample). As shown in Figure 6.9, Blips occur anywhere in time, for the simulation there is no preference in time, agreeing with our expectations. As for the energy measurement of our simulations, as expected we get a negative bias due to ‘finding the maximum negative excursion’, the distribution is peaked at -2 keV and extends to -4 keV.

We next compare our simulation-based expectations with the full IRR sample and the LHS sample. We notice that there is no preference in the time of the biggest Blip, except for the first bin. While the cause of this is not yet understood, we ignore it for now as they can be safely ignored for our resolution studies since they are occur far away from t_0 , where

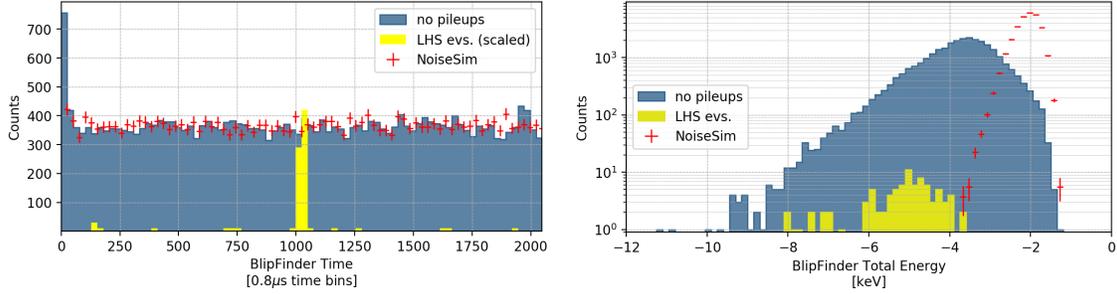


Figure 6.9: The BlipFinder results (BlipFinder time and energy) applied on noise simulations (NoiseSim - red crosses), the full IRR data sample (in blue), and on the sample of LHS events (in yellow). There is no preference in BlipFinder time for simulations as well as for data (except for the first time-bin which is ignored because it is far from t_0). On the other hand the distribution is very different for the LHS sample and is correlated with Blips found at t_0 . The BlipFinder energy of the IRR data is greater than that of NoiseSim, suggesting Blips are common.

they could affect our 0-delay energy measurement (we give more detail of these events in Appendix A). We note that in almost all cases the LHS events also correlate with the biggest Blip occurring near t_0 . We also note that a number do not. A study of these events shows that they are occurrences where there are multiple Blips in the same readout window. In this case there is one at t_0 but the biggest one does not occur near t_0 . As for the energy measurement results, we note that all of the LHS events are beyond our noise simulation expectations, and that most of the full IRR sample is, an indication that Blips are a common occurrence and that they affect more than just our LHS events.

Finally, the smoking gun of our BlipFinder is shown in Figure 6.10, which shows the 0-delay energy measurement as function of BlipFinder time. This figure shows that only Blips near t_0 (time-bin 1024) have a large impact on the measurement, while those far away from the 0-delay time do not have an impact. A closer inspection shows that while Blips closest to the 0-delay time are the cause of the LHS tail, those before or after it (but not too far) are the cause of the positive offset in the Gaussian distribution in Figure 6.6.

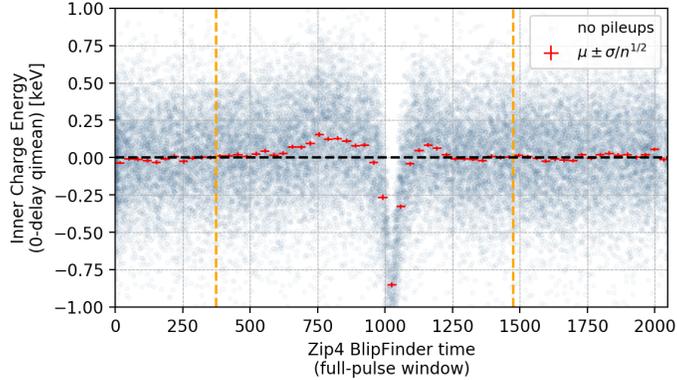


Figure 6.10: The 0-delay energy measurement as a function of BlipFinder time. Only Blips occurring before time-bin 375 and after 1450 do not affect the measurement. The most affecting Blips are those near the 0-delay time (time-bin 1024), causing the negative (LHS) tail, while Blips occurring before or after the 0-delay cause the positive offset in the Gaussian distribution shown in Figure 6.6.

6.3.2 The TowerBlipFinder Tool

With the BlipFinder in hand and working as expected, we describe the TowerBlipFinder tool. In this section we begin with the goals and how the TowerBlipFinder will be used, with an emphasis on using it for correcting the energy measurement in iZip4. We next describe our assumptions for our tool, based on observations made with the BlipFinder. We follow with the steps of the algorithm. Next we present the de-weighted amplitude search method, which is optimized to find the most offending TowerBlip, near the t_0 region. Finally we end with the TowerBlipFinder results, as preparation for our Blip correction method in the next section.

Blips occur in most events and the goal of the TowerBlipFinder is to correct the energy measurement in a detector using the correlations between detectors to make a prediction. For the case of estimating how much energy is expected for the t_0 energy measurement in iZip4 we use the TowerBlipFinder to search for Blips in iZip5 and iZip6 that would have the biggest effect on that measurement. In Appendix B we show that TowerBlips are the most common type of Blips, so since we are only interested in the energy observations of iZip4,

we can use iZip5 and iZip6-side2 information to make a prediction of the effect in the iZip4 energy measurement.

Before describing our TowerBlipFinder tool, we list our assumptions. Some of these will be based based on observations from the blip rates as described in Appendix B and some we will state and then show they are true with data. Our assumptions are:

- We are only interested in Blips that affect the iZip4 E_0 measurement (at 0-delay time) for this reason we will only consider times of 100 bins around the t_0 time
- Most often (over 65% of the time) iZip4 Blips are correlated with iZip5 and 6, showing up as TowerBlips
- Blips in iZip4 are correlated with iZip5-side1 channels and anti-correlated with iZip5-side2 and iZip6-side2 channels
- Since they are correlated we can add the optimal-filtered amplitudes as a function of time, to enhance correlated features while canceling any potentially uncorrelated noise
- We measure the distribution of the time difference between our correlated Blips as seen in iZip4 and iZip5 or iZip6, which is a Gaussian function with a mean of -4.9 time-bins, and RMS of 7.8 time-bins, and since we only care about Blips near t_0 , we will de-weight our amplitudes to search with this function (as will be shown after describing the algorithm)

With our goals for the TowerBlipFinder and the assumptions in hand we now describe the tool algorithm itself. The steps for identifying and measuring TowerBlips in iZip4 are:

1. Add optimal-filtered channels of both iZip5 and iZip6-side2 (flipping polarities of anti-correlated channels)
2. De-weight the amplitude predicted for iZip4 by the Gaussian Δt_0 window ($\mu = -4.9$ time-bins, RMS = 7.8 time-bins)

- Obtain the TowerBlipFinder time $t_{iZip5+6side2}$ and energy $E_{iZip5+6side2}$ of the maximum negative excursion (biggest blip in the de-weighted amplitude window) which can be subtracted from the observed 0-delay energy in iZip4

After correlation in-between channels, the next crucial part of our TowerBlipFinder is to determine de-weighted amplitude function which we now describe. Figure 6.11 shows the time difference, in time-bins, between the biggest Blip in iZip4 and that found in iZip5+6. From this figure we know that a 100 time-bin window is a good range to look for correlated Blips across multiple detectors, so we force the BlipFinder to only look in a 100 bin window centered at t_0 . Just in case there is more than one Blip within this window and since we are only interested in Blips that would affect the t_0 energy in iZip4 the TowerBlipFinder is further optimized for finding the most-offending Blip (as opposed to the biggest one regardless of the distance from t_0). To do that, the BlipFinder searches for Blips in the added iZip5 and iZip6-side2 optimal-filtered amplitude after being de-weighted as function of distance from t_0 . The de-weighting function is determined from the best-fit values of the Gaussian to the distribution in Figure 6.11. We find $\mu = -4.9 \cdot 0.8 \mu s$ and $\sigma = 7.5 \cdot 0.8 \mu s$.

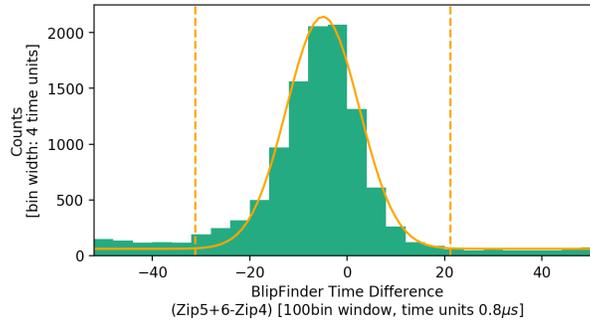


Figure 6.11: This figure shows the time difference, in time-bins, between the biggest blips in iZip4 and the combined iZip5+6 BlipFinder result (100 bins around iZip4 BlipFinder time). It is further used as a weighting function for the iZip5+6 BlipFinder search (100 bins around iZip4 0-delay). The fit function is a Gaussian with $\mu = -4.9 \cdot 0.8 \mu s$ and $\sigma = 7.5 \cdot 0.8 \mu s$.

Using the de-weighting function with our BlipFinder, the resulting iZip5+6 BlipFinder

time and total energy are shown in Figure 6.12 for the full IRR sample and the LHS subsample. For both samples most of the events are correlated (near time-bin 1019), and the energy distribution shows that LHS tail events are also large-negative iZip5+6 BlipFinder Energy events. We note that there is an overflow bin in the BlipFinder time result (time bin 1074), this is not from the algorithm finding a large negative excursion, but rather from the amplitudes having zero-values after the de-weighting effects. These will not affect the correction because (as will be shown) the correction is also zero-valued for edge events.

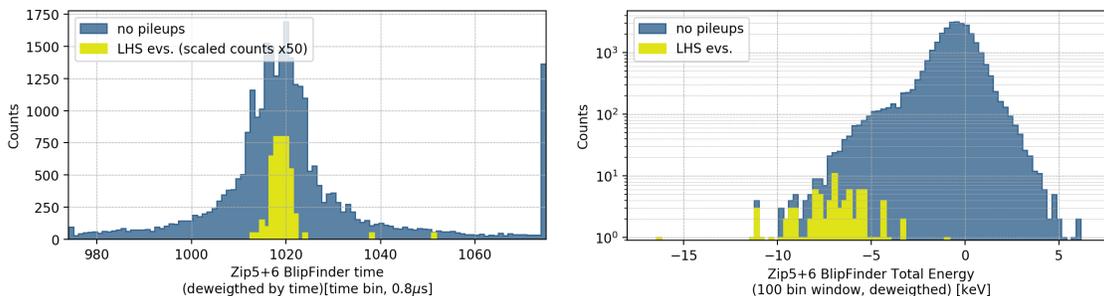


Figure 6.12: This figure shows the time and energy predicted in iZip4 from results for the TowerBlipFinder which looks for evidence in iZip5+6side2. The results are for both the full IRR sample after pileup removal (blue) and the LHS subsample (yellow). We note that the LHS events (known to be iZip4 Blips) are correlated in time and energy with the iZip5+6 BlipFinder results. The overflow bin (time bin 1074) in the left hand plot results are not from the algorithm finding a large negative excursion, but rather zero-values from the de-weighting effects.

We will use the TowerBlipFinder in the next section to make a prediction for the contribution to the t_0 energy measurement in iZip4 based on the information from iZip5 and iZip6-side2.

6.4 Correcting for TowerBlips

In this section we show how we correct the energy measurement in a detector when there is evidence of a TowerBlip from the other detectors. This will be useful for when we clean our IRR sample for comparisons to our noise simulation, and measuring the contribution

to the detector resolution from noise in Chapter 8. We will use the information from the rest of the iZips in the tower to make a prediction of its effect in our energy estimate, to correct it by subtracting the prediction. In this section we begin by showing the observed 0-delay amplitude functional dependence on BlipFinder time and energy, which is the basis of our correction. We proceed by presenting the resulting energy prediction and separation of the original data distribution from our prediction. Finally we show the Blip-corrected IRR data, which is predominantly purely Gaussian and we give some remarks on the resolution estimate and resolution effects from the correction.

The iZip5+6 BlipFinder results (time and energy) are highly correlated with the observed iZip4 0-delay energy measurement. We make a 3D fit of the iZip4 energy measurement (for each channel) as function of iZip5+6 BlipFinder time and energy. As shown in Figure 6.13, the fit (described in Appendix D) well-represents the profile of the data. We will use the fitting function for each channel to obtain a predicted value of the iZip4 effect from iZip5+6 information.

Upon computing the iZip5+6 Predicted energy, from the 3D fit of each channel, we have predictions for all 4 channels, and can recalculate the qimean and qsummax energy estimators. Figure 6.14 shows the energy prediction of the 0-delay iZip4 qimean distribution. We show the resulting prediction for both the full IRR and LHS subsamples, as well as the corrected uncorrected 0-delay energy measurement for the IRR sample sub-divided into samples with a large and a small predicted value. The LHS events indeed have large predicted values, and the large predicted value sample does correspond to a negatively biased distribution in iZip4 0-delay. This predicted energy could be used to remove Blips but would not be very efficient and will not remove the offset in the rest of the distribution.

We can use the iZip5+6 predicted energy to correct the iZip4 0-delay observation. We will call this the corrected energy estimate. We do this for all 4 channels and recalculate

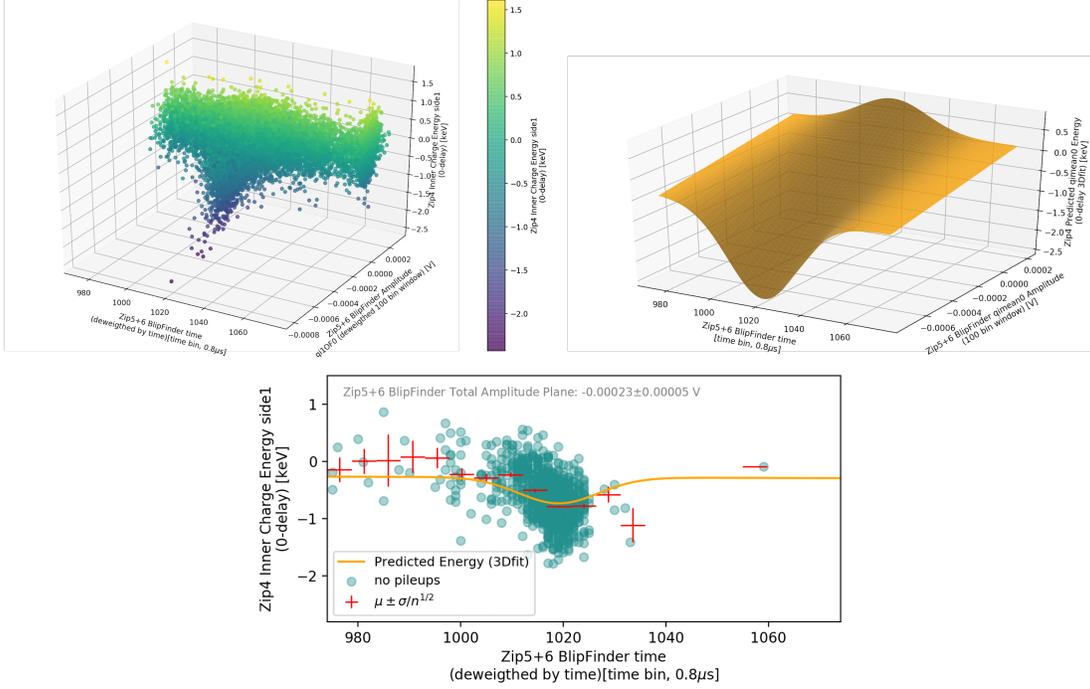


Figure 6.13: Top-Left: iZip5+6 Correction Data. X-axis is iZip5+6 BlipFinder time, Y-axis is iZip5+6 BlipFinder amplitude, and Z-axis is observed iZip4 qimean 0-delay energy. Top-Right: Fitted shape for the iZip5+6 Correction 3D fit based on the data from the left with the same axes. Bottom: iZip5+6 prediction fit example, from a sample of events with iZip5+6 BlipFinder Amplitude amplitude of $-230 \pm 50 \mu\text{V}$. Data in green, profile in red and fit in orange. For fitting planes of the full data see Appendix D. These figures show that the 3D fit is a good representation of the data and can be used to predict the effect of iZip5+6 Blips on iZip4 0-delay estimates.

qimean and qsummax. We take:

$$E_0^{\text{corrected}} = E_0^{\text{measured}} - E_0^{\text{TowerBlip}} \quad (6.1)$$

where the correction is simply a subtraction of the prediction from the observation. Figure 6.15 shows the uncorrected and corrected distributions with Gaussian fits (and we will discuss remarks on the prediction resolution next). The corrected distribution is Gaussian (only few LHS events remain, but with much lower energy), the mean offset is removed, and the RMS is lower.

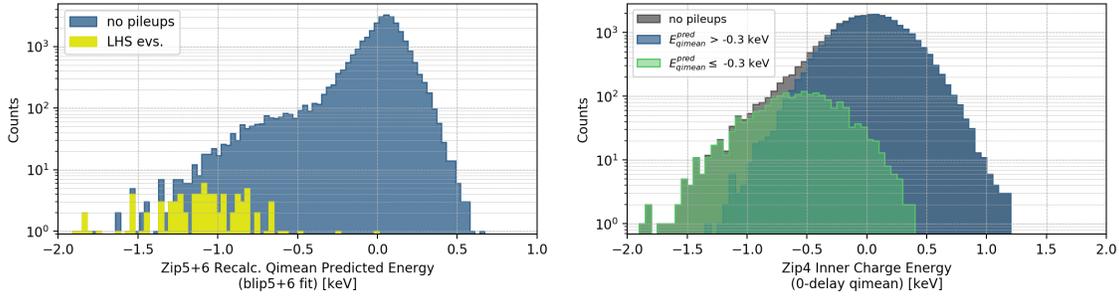


Figure 6.14: These plots show the amount of energy predicted in iZip4 based on the TowerBlipFinder tool that only looks at iZip5 and 6. Left: Recalculated qimean 0-delay iZip5+6 predicted energy for the sample. Notice all but 2 LHS events (original qimean 0-delay tail) are biased to lowest values of the iZip5+6 prediction. Right: iZip4 qimean 0-delay (not corrected) distribution of sub-samples separated by iZip5+6 Prediction above of below -0.3 keV. Notice events with high-negative correction (below -0.3 keV) make up the LHS tail of the iZip4 distribution and is highly negatively biased.

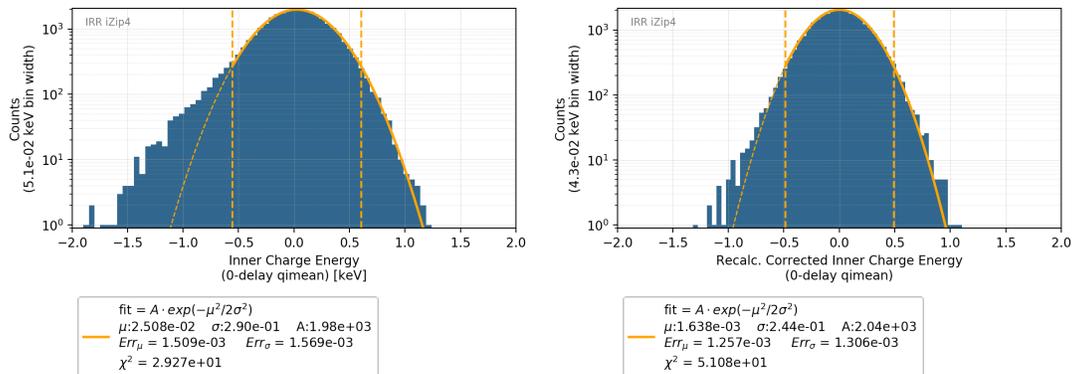


Figure 6.15: This figure shows two versions of the qimean measurement before (left) and after (right) the TowerBlip correction energies. After correction the sample is purely Gaussian (few events on LHS), mean offset is gone, and RMS is lower, showing that the correction works as expected.

Finally, it is worth noting the effects of the correction in the 0-delay observation. While the corrected samples have lower RMS, it is appropriate to consider that the correction introduces a resolution effect from the prediction. The prediction effect can be estimated by making a Gaussian fit near the prediction peak and measuring the resolution (RMS), see Figure 6.16. We can assume that the real pure-noise resolution adds up with the prediction

resolution, and we observe the quadrature sum of both in the corrected distribution. For example, for q_{imean_0} measurements, the observed corrected sample resolution is $244 \text{ eV} \pm 1 \text{ eV}$ (shown in Figure 6.15), the prediction resolution is $90 \text{ eV} \pm 3 \text{ eV}$ (from Figure 6.16), so the true pure-noise resolution of our IRR sample after subtracting in quadrature is $227 \text{ eV} \pm 1 \text{ eV}$.

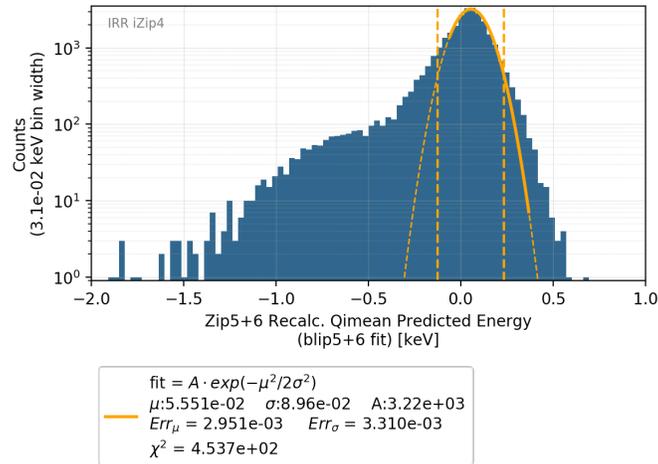


Figure 6.16: The iZip5+6 prediction fit to near peak region ($\mu - 1\sigma, \mu + 2\sigma$). Observe that the correction method introduces the prediction’s resolution to the corrected energy distribution, added in quadrature with the pure-Gaussian noise resolution.

Now that we have described the sources of readout mismeasurement as well as how to deal with them and have an IRR sample that is clean from Blips and pileups effects we are ready to move forward. With a pure-Gaussian noise sample we will be able to compare to our simulations in Chapter 8. But first, in the next chapter we describe our simulations and the simulations infrastructure.

7. THE SUPERCDMS SIMULATIONS

With a full description of the detector, data, and analysis techniques in place we are now ready to describe the SuperCDMS simulations and the infrastructure designed to produce high-quality samples of events interacting in the detectors. We note at the outset that the simulations infrastructure is fully functional and is capable of producing large quantities of both high quality signal and background events for use in analysis. An important option we have added to the simulation is the ability to simulate events with and without noise, as this allows us to separate out which response features in samples of simulated events are due to noise and which are due to detector response effects. After we describe the tools and the data-sets we have made we will have laid the ground work for the next three chapters where we will show results, measurements, and comparisons with data. While the overall goal is to use the simulations for understanding the detector response to optimize an analysis, we begin by showing that they respond as expected in known cases. In Chapter 8 we will compare our simulated noise sample to the cleaned IRR sample described in Chapter 6. Next in Chapter 9 we will compare our simulated ^{133}Ba sample to the data taken and described in Chapter 5, following with Chapter 10 where we make studies from the simulation predictions from truth information and use those studies to help us with understanding of our WIMP-like events sample. We note at the the outset that while all the pieces qualitatively work as expected, only tuning of the noise simulations is complete and additional work will need to be done on the detector response simulations, as described in the second part of Chapter 10.

Before outlining the chapter we quickly overview the infrastructure so it is easier to see where we are going. The simulations infrastructure is composed of 3 basic pieces, as seen in Figure 7.1, which allows us to create simulated events in the same basic path as was done in real data. In many ways, it mirrors the descriptions of Chapters 2 and 5. The first stage simulates the creation of particles that enter the experimental apparatus and interact to create an electron or nuclear recoil in the detector crystal, this is SourceSim. The

second stage is DetectorSim which simulates e/h and phonon creation and propagation in the crystal, then simulates the charge and phonon sensor response, adds noise (optionally), and stores the results in the same format as real data. Finally Event Reconstruction, which is done by running the same algorithms, as described in Chapter 4, as used in real data reconstruction, to turn the raw signals into energy and other measurement values [53, 80].

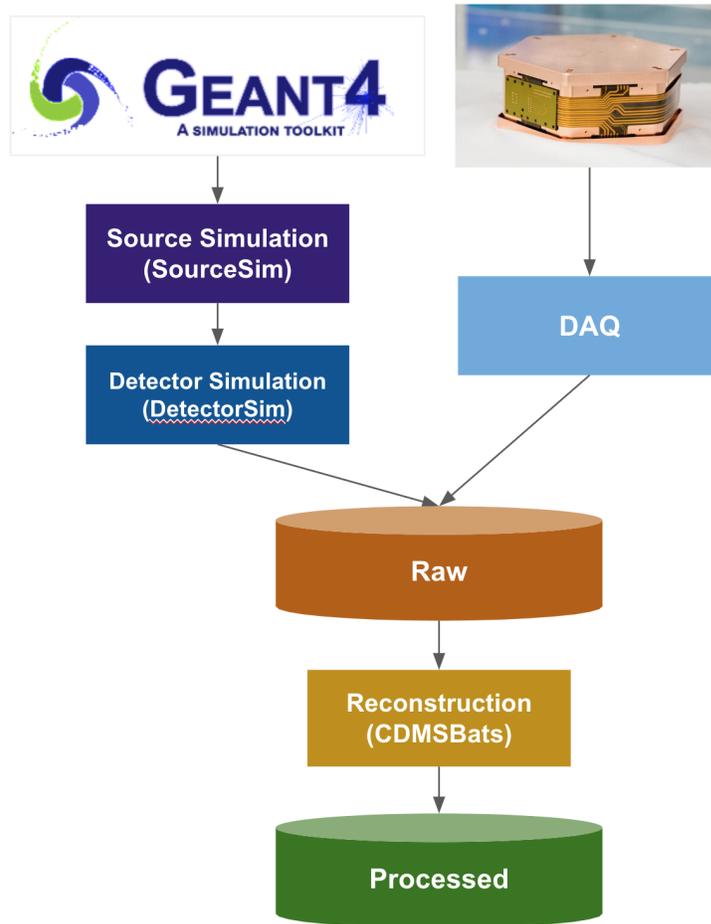


Figure 7.1: A diagram showing the various stages of the simulation (on the left) which mock up interactions in the real detector (on the right). The first stage (SourceSim) is about simulating the sources and the primary interactions with the detector (and any shielding material in the experiment). The next stage (DetectorSim) simulates the response to the interactions, from electron, hole, and phonon propagation, to sensor readout and noise. Next the same format as data is produced to analyse with the same Reconstruction software as used in real data. Finally, the events are run through Reconstruction to produce Processed events.

In this chapter we will describe our fully functional set of simulation software packages, and how they are used to create the samples we use in this work. First, in Section 7.1 we give an overview of the simulation pieces and the infrastructure that integrates them together. In Section 7.2 we describe SourceSim, which simulates the sources, the propagation of the emitted particles through the detector apparatus including the primary interactions in the detector with the determination of the recoil energy and the recoil type. Next, Section 7.3 describes DetectorSim which simulates the detector response, from the Lindhard Model to the distribution of the primary energy into charge and phonons, their propagation in the crystal, the sensor response, and the simulation of the data acquisition to add noise and formatting of the output. We then give a summary of the event reconstruction in Section 7.4. Finally, in Section 7.5 we give an overview of the detector noise simulation sample as well as the primary energy deposition in the ^{133}Ba and WIMP simulated samples used for this thesis work, which will be presented as a precursor to use them for their analysis in the upcoming chapters.

7.1 Overview of Simulations and the Simulations Infrastructure

We begin by describing the parts of the SuperCDMS simulations infrastructure. A more detailed version of the simulations part of Figure 7.1 is shown in Figure 7.2, which contains all the individual pieces that are integrated to produce simulated samples of events. We will briefly describe how they work together to produce the samples, starting with SourceSim, then DetectorSim and finally Reconstruction along with the simulation-based truth quantities for analysis, and in the upcoming sections we will describe each part in detail.

The first stage is SourceSim, composed of a number of subcomponents known as SuperSim [48], and WimpSim which will be described in more detail in Section 7.2. The core component does the simulation of the particle-level interactions, from the source to the crystal, called SuperSim [48]. It is based on a standard particle physics package called Geant4 [81] which simulates particle interactions and propagation within materials. For the simulations of WIMP interactions (since non-SM interactions are not part of Geant4), we use a custom

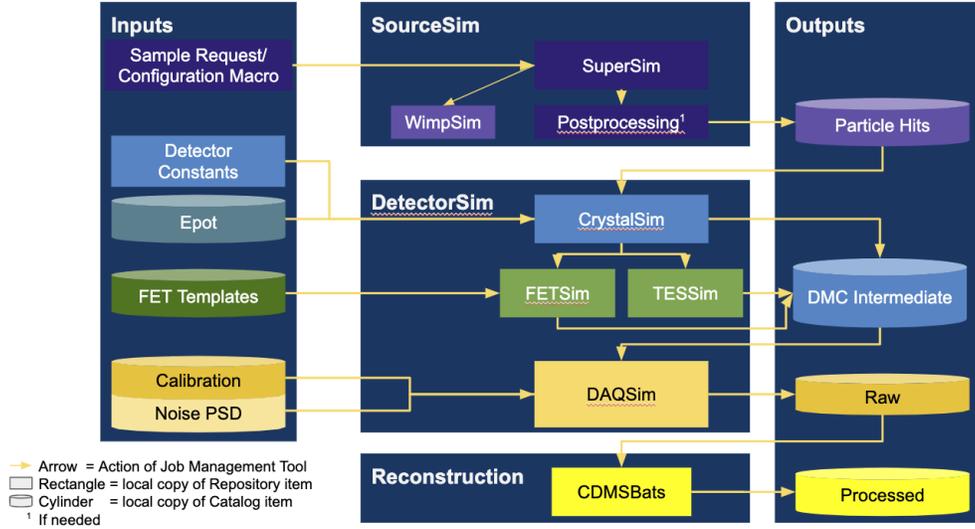


Figure 7.2: Detailed view of the simulations workflow (courtesy of Richard Lawrence), showing all the subcomponents of the various pieces of the simulation.

package known as WimpSim which simulates signal-like events following the astrophysical implications defined for Eq. 1.1 (note that this functionality has been integrated into SuperSim, after the analysis presented in this thesis). The final output, after any post-processing to select-out subsets of events (as shown in the Postprocessing box of the figure for completeness), contains all the information needed by DetectorSim including the particle interactions in the detector crystal (whether electron or nuclear recoils), the amount of energy deposited, and the location where the interaction occurred. We note that a single particle can interact multiple times within a detector, called multiple scatters. The information for all scatters are put in an output file called ParticleHits.

Next, the detector readout response simulation is performed by DetectorSim, which simulates the charge and phonon creation, as well as their propagation through the crystal and all collection by the sensors. This will be described in more detail in Section 7.3. It has 4 working pieces (see center of Figure 7.2) which take input from 4 separate places (ParticleHits, Detector Constants, the electric potential map of the detectors, and the FET templates), and stores its output in the same format required for Reconstruction, known as Raw, as

shown in Figure 7.2. The first stage, known as CrystalSim, takes all the particle interaction recoil energies and interaction types from ParticleHits, and using the processes described in Chapter 3.1, determines the amount of primary energy deposited in e/h and phonon creation. The simulation of the propagation requires Detector Constants (material parameters, propagation constraints, etc.) and a description of the electric field at every point in space. The latter is simulated using a mesh-triangulated file with the numerical solution of the electric field inside the detector at every point in space (called the Epot file) to simulate the acceleration of charges throughout the crystal. The Epot file is created before running the simulation using a commercial program for finite element analysis, called COMSOL® [82]. After propagation to the electrodes, the readout of the FET and TES systems is modeled using FETSim (charge) and TESSim (phonon), which output a series of waveforms that mimic real data. The combined output of CrystalSim, FETSim, and TESSim is referred to as DMC Intermediate, as shown in the figure. The last stage in DetectorSim is the calibration from simulation units to ADC counts, data formatting, and noise simulation which occurs in a package known as DAQSim. As shown in the figure, DAQSim package takes inputs from DMC Intermediate, a calibrations constants file (for scaling the waveforms into expected ADC amplitudes as in data), and noise PSDs for the noise simulations. With these inputs it scales the waveforms accordingly, simulates and adds noise (optionally). The noise addition is performed by NoiseSim, which has two purposes: first it simulates the expected baseline-noise sample necessary for Reconstruction, and second it (optionally) adds sampled random noise to the simulated (and already scaled) TESSim and FETSim events. The noise addition is optional and turning it on and off allows the simulation to be used in a way that can separate between noise and detector effects. Finally, DAQSim stores the waveforms in the same format as real SuperCDMS Soudan output, shown as Raw in Figure 7.2.

Finally CDMSBats, the SuperCDMS event reconstruction software, including the algorithms described in Chapter 4, analyzes the simulation events using the same tools as done in real data (see bottom of Figure 7.2). CDMSBats is based on C++ [83] and ROOT [84]

(common framework for data analysis in physics). Since in the simulations we have access and use the true information from the interactions and throughout the response of events and scatters, we have added useful truth quantities to the final reconstructed files. Table 7.1 summarizes the additional simulation-based variables.

Quantity	Description
SimEventNumber	Simulation event number (to map with DetectorSim output)
SimAvgX (Y, R, and Z)	Energy weighted average scatter's location
E_r	True recoil energy (summed over all scatters per event)
E_Q	True/Expected charge energy from Lindhard Model (for nuclear recoils) or equal to E_r (for electron recoils)
E_{Ph}	True/Expected phonon energy from Lindhard Model and NTL Phonons
Q_{max}	Total charge amount released in the recoil interaction
Q_{Ramo} (i1,i2,o1,o2)	Collected charge at any of the 4 electrodes
E_{Coll} (qi1,qi2,qo1,qo2)	Collected charge energy at any of the 4 electrodes

Table 7.1: Summary of simulation variables additional to real-data reconstruction. These variables provide insight to true information to help understand mismeasurement causes.

Having described all the simulation components, we are now ready to describe in detail how each does their work.

7.2 Simulation of the Sources and Recoils in the Detector: SourceSim

The first stage and primary workhorse within SourceSim is SuperSim [48] which is developed in a publicly available toolkit package for the simulation of particles passing through matter, called Geant4 [81]. This package is essentially about declaring the geometry and materials of the entire experimental apparatus and letting Geant4 do all the heavy lifting (using cross section calculations to simulate interactions on an particle-by-particle basis).

Geant4 allows for the incorporation of non-standard model particle theories such as WIMPs (as people can contribute and release their own packages, like SuperSim) in WIMPSim.

With the apparatus defined, similar description tools in Geant4 can be used to simulate a source which creates particles that can then be propagated through the experiment in multiple steps to create events. As shown in Figure 7.3, the simulated source of events (radiation source or signal events) produces the particles that travel through the material in the apparatus. The source can be a radioactive material, or a virtual flux of particles; any source, any process, for any particle is allowed. All the steps the generated particle takes as it travels through, and interacts with, the experiment's materials are simulated until either being absorbed in the crystal, annihilating, or going off into infinity. For our purposes, we are interested in scatters in the experiment, and in this context a scatter is any interaction (including absorption) that produces an energy deposition in the crystal, either via electron or nuclear recoil. We note that particles that interact with the detector can scatter within the crystal any number of times before finally being absorbed or exiting the crystal. Whenever there is more than one deposition (either from the same or multiple input particles) these are known as multiple scatter events. The final output (ParticleHits) contains the information about all the scatters and interactions that occur within the detector crystal (which can be for more than one detector, up to all 15 detectors in the SuperCDMS Soudan configuration).

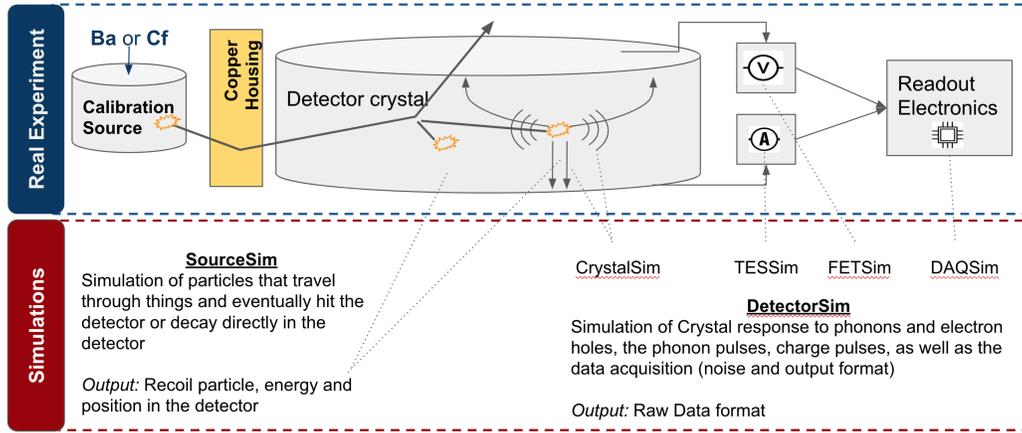


Figure 7.3: Parallel diagrams of the propagation, interactions, and readout comparing the real experiment (top) and simulations (bottom). A radiation source creates events (can also be WIMP signal events, or cosmogenic showering, etc.) that pass through the experimental apparatus and reach the detector crystal. After that the incident particles deposit energy. All instances of scattering are simulated up to interactions in the crystal; multiple interactions can occur within a single event and are referred to as multi-scatter events. The charge and phonon sensors produce a signal which is recorded by the readout electronics, and stored in a file for Reconstruction. The simulations replicate the process by propagating the electron-hole pairs through the system and simulating the response of the charge and phonon sensors.

7.3 Simulation of Detector Physics and Readout: DetectorSim

The outputs of SourceSim contained in ParticleHits are used as the input to DetectorSim which consists of 4 modules that simulate all the steps to turn a scatter (or scatters) in the crystal into a simulated detector response readout for use in reconstruction. As seen in Figure 7.3, the simulation begins with the full set of scatter information from the ParticleHits, and outputs a collection of waveforms for event reconstruction.

First we will describe CrystalSim, in Section 7.3.1, starting with the implementation of the Lindhard Model to distribute the scatters energy into primary charge and phonon creation, and the charge and phonon propagation (and secondary phonon emission) implementation. The sensor readout response is described in Sections 7.3.2 and 7.3.3 for charge (FETSim) and phonons (TESSim) respectively. Finally, Section 7.3.4 describes the calibration, digitization, formatting, and noise (NoiseSim) implemented in the DAQ simulation (DAQSim).

7.3.1 CrystalSim: From Recoil Energy to Creation and Propagation of Charges and Phonons

CrystalSim simulates the response to the electron and nuclear recoils, as well as the charge and phonon motion in the crystal for use in TESSim and FETSim according to the models described in Chapter 3. This section describes how the charge and phonon creation and propagation mechanisms are implemented in CrystalSim, first by describing how the amount of energy given to the creation of primary electron-hole pairs and phonons is determined, and next by describing their propagation in the crystal up to their end-state in the simulation, according to the crystal physics.

The first step is to take the information from ParticleHits (the position, energy and recoil type) and use the Lindhard calculation (Eq. 3.7), to determine the primary charge and phonon energies, and create the primary e/h and phonon particles. In the simulation these initial phonon and charge carriers are created at the interaction point, each with the energies specified in Eqs. 3.8 and 3.9. Since the recoil is with the entire crystal, rather than the single particle, we approximate the creation direction of motion of all phonons and charges with a directionally-random distribution, which is valid approach due to crystal momentum physics constraints [62].

The second step is to propagate the charges and phonons, and simulate the creation of secondary of phonons as described in Chapter 3.2.2, for use as input to TESSim and FETSim. The implementation is reduced to two processes: propagation of phonons, and propagation of charges with secondary phonon creation (from NTL gain as described in Section 3.2.2) [62, 68]. These processes can be solved at every time-step of the computation taking into account the crystal propagation constraints, for electrons and holes (from Chapter 3.1.3). We next describe some details of the phonon and charge propagation in turn.

For the phonon propagation, as described in Section 3.2.3, our simulation calculates rates, directions, and modes of the scattering processes (anharmonic decay and isotope scattering). The rates and mode proportions, and mode conversions are calculated from the

results from Ref. [69] which are frequency dependent [62, 69] and the modes are randomly chosen from a population of modes in agreement with the populations described in Section 3.2.2. The phonon group and phase velocities (which are only directionally dependent), are pre-computed with the computation step-size for all phonon modes [62]. From the rates and expected proportions, the simulation calculates the time-step until the next interaction point for every phonon, where each phonon undergoes one of the 2 scattering options: either just changes direction and mode (isotope scatter), or down-converts into 2 lower energy phonons also changing mode (anharmonic decay, only applicable for longitudinal phonons). The isotope scatter process is simple: first the phonon mode is selected from random fluctuations around the populations described in Section 3.2.2, next the propagation direction is randomly chosen. The anharmonic decay is more complicated (given the multiple end-state mode population fractions), but a similar approach is implemented; first the number of outgoing phonons of each phonon-mode is selected from random fluctuations around the expected populations, then for each of the phonons (with a defined mode), a random scattering direction is simulated, preserving momentum and energy given by the scattering solutions from Ref. [62, 69]. Finally, for both processes, the output states are checked against the dispersion limits, and random values for the direction are regenerated for cases that exceed them [62]. Phonon propagation stops when phonons get absorbed at the TES cells, or they down-convert into phonons with energy below the Al band gap energy ($340 \mu\text{V}$ or 0.084 THz) [68, 62] in which case the energy is lost for measurement purposes.

Charge and hole propagation with phonon emission is simulated step-by-step by solving the equations of motion and constraints described in Section 3.2. Since there are multiple interactions, for each charge interaction we probabilistically determine whether phonon emission occurs, and the outgoing charges and photons are forced to conserve 4-momentum. A fine-enough time step in the charge propagation is calculated from the empirically obtained maximum charge momentum as described in Ref. [62]. The equations of motion take into account crystal propagation constraints as well as acceleration by the electric field (esti-

mated by using the triangulated Epot mesh) and the phonon production ‘scattering’ process as described in Section 3.2.2. In the case of scattering, the angular solution from Fermi’s Golden Rule [62] is randomized. For holes, the equations are straightforwardly solved from the momentum evolution in the electric field and limited by the phonon scattering, while for electrons, we use the Herring-Voigt transformation [70], as described in Section 3.2.2, at each time-step, simplifying the system and making it solvable with the same equations as for holes. An example is shown in Figure 7.4 and shows the electrons following the trajectories along the L-valleys and the holes moving ballistically.

The charges motion stops when they are ‘absorbed’ by the electrode or stopped by any ‘trapping mechanism’. A charge is considered absorbed (and removed from the simulation) when the charge’s position Ramo-field value difference from the electrode’s value is less than 0.02. Most of charges reach the respective electrode because the field is mostly uniform, but depending on the scatter location some can get trapped in the interleaved potential, or in one of the two trapping mechanisms, as described in Chapter 3.3. While detector edge absorption occurs when the charge’s transverse momentum is higher than a given value at the edge of the detector (either at the top/bottom or in the sidewall), the present simulation is set to always absorb charges at the sidewall (whenever their momentum value is higher than zero, which is always true), while charges are always reflected at the top/bottom surfaces (impossible threshold momentum value of 10^{20} eV) other than at the electrodes. Since random trapping is a probabilistic effect (motivated by trapping due to crystal impurities) [62], we roll the dice every $100 \mu\text{s}$ to see if the charge is randomly trapped.

Now that we have described the charge and phonon propagation implementation, we are set to describe how the sensor readout, by FETSim and TESSim, is simulated.

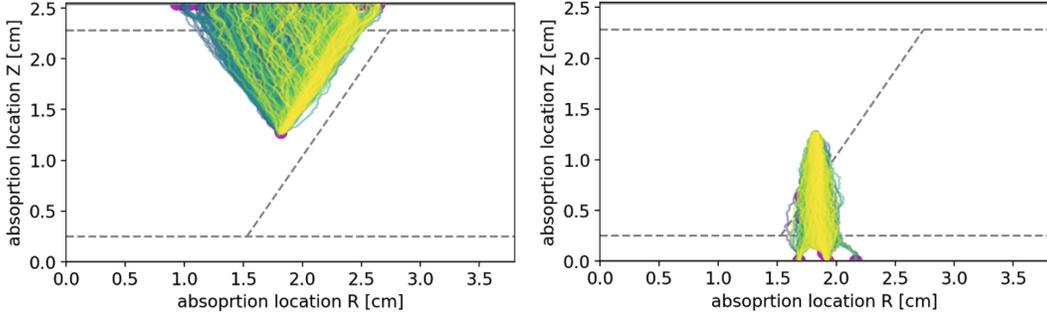


Figure 7.4: A figure showing the trajectories of electrons (left) and holes (right) for an example event which scattered in the middle of the detector. As expected from detector physics, charge trajectories well-represent the expected physics: electrons propagate in the restricted [111] plane (L-valleys) of the crystal, while holes propagate ballistically. All of these electron-holes were absorbed at the top/bottom electrodes.

7.3.2 Simulation of the Charge Readout: FETSim

The FET Simulation (FETSim) package calculates the charge collected by the electrodes and produces the waveform simulations. In this case the Shockley-Ramo theorem [71, 72] is used to estimate the total amount of charge that is collected by each electrode, simply converting the Ramo field value for every absorbed or trapped charge into the final collected charge for each electrode, and summed by channel. The value is then used to scale the template waveforms, described in Chapter 4.4 to simulate both the true readout and cross-talk. As described in Chapter 3.3, the real-detector charge collection occurs so quickly (during one time-bin of $0.8 \mu s$), that the simulation is basically that of the charge circuit which is measured in data as described in the template creation in Chapter 4.2, so it is appropriate to simulate the charge signal by scaling the templates (see Figure 4.2) and adding the cross-talk, as opposed to simulating the dynamics of the pulse rise-time. For each channel, the charge collection value is used to scale the template and cross-talk templates amplitudes to build the final response. The results are stored in DMC Intermediate.

7.3.3 Simulation of the Phonon Readout: TESSim

The TES simulation, implemented with TESSim, takes the simulated phonons that reach the TES cells, collects them and outputs the digitized waveform using a simple Ordinary Differential Equation (ODE) solution mechanism. The simulation starts with the collective response of the TES cells to the total number of phonons (and their energies) that reach the device and takes into account the thermo-electrical effects as they evolve over time. Each phonon increases the temperature, hence resistivity, of the TES cell which produces an observable effect in the simulated TES circuit, from which a waveform response is recorded. We note that saturation, when TES cells heat-up and continue receiving phonons before falling back to the transitioning superconducting stage, is taken into account. The result is that the pulses rise-time slows down and the peak becomes wider. Each phonon channel is simulated independently, and their response is recorded. The results are stored in DMC Intermediate. Since this thesis is about charge readout, this simulation will not be described in great detail. For more details of this implementation see Ref. [85].

7.3.4 Simulation of Data Acquisition and Noise: DAQSim

The DAQSim portion of DetectorSim takes the DMC Intermediate file output, optionally adds noise, and prepares the files for event reconstruction. It does 4 essential tasks: converts the TESSim and FETSim waveform outputs into detector readout units (with appropriate calibration constants), adds noise with NoiseSim, digitizes the events in the same way as data (particularly important for the lowest energy deposits), and formats the events for use by the standard event Reconstruction tools. The noise addition portion of DAQSim, known as NoiseSim, is done separately for the FET and TES systems. We will describe each task in more detail in turn.

Since the outputs of both FETSim and TESSim are in arbitrary units, the waveforms need to be converted to the same scale as real data. This is done in a single step. Since for each charge and phonon channel, Reconstruction uses a single constant to take ADC units

to energy units, determined as described in Section 5.2.1 using ^{133}Ba data, the simulation does the inverse, but takes into account the measured conversion scale from simulation units to ADC units. In practice, we simply run a full ^{133}Ba sample through the full simulation and reconstruction procedure and determine the constant directly. We also note that creating an event with zero energy from TESSim and FETSim can be used as the input to NoiseSim to create pure-noise events.

Next DAQSim creates noise for the event reconstruction by simply creating a fake noise waveform and adding it to FETSim or TESSim output. The noise as a function of time in an event is created from the PSD measured in real data (like the PSDs in Figure 3.5), by randomly sampling the amplitude and phase contributions for each frequency with the expected Gaussian variation as in the PSD definition in Eq 4.4.

Since data is digitized by the ADC components in the data acquisition as described in Chapter 3.3, DAQSim next digitizes the waveforms using two simple steps. The first is to add the sensor and noise waveforms. Next, since the simulation results are floating point variables DAQSim simulates the digitization done in real data by rounding down the waveform values at each time-bin to the nearest integer in absolute value.

Finally DAQSim stores events in the same binary-file format as real SuperCDMS Soudan data [86]. This is known as Raw Data, as shown in Figure 7.3. This is the last piece of the infrastructure that is particular to simulations only, as the next step follows the same process as real data.

7.4 Event Reconstruction with CDMSBats

The last stage in the production of simulated events is Reconstruction, using the standard reconstruction software, CDMSBats [53], which takes the in the Raw Data files and outputs the best energy estimators and other calculated quantities of interest for use in analysis (labelled as Processed in Figure 7.2). CDMSBats reads the Simulated Raw Data events as if it was real data, and produces the same output files as used in analysis. It executes all the standard algorithms, like the optimal filter, described in Chapter 4, but it does not do

corrections based on environmental variables (like temperature, bias voltage in the crystal, etc.) as they are assumed to be nominal values only. All the charge and phonon energy best-estimates are obtained in this stage.

7.5 Simulated Samples

With all the pieces to create fully simulated and reconstructed events described, we next list the set of samples we have created for this work. They are:

- Pure Noise Simulation: A sample of events with no particle interaction (zero TES and FET output). Used for NoiseSim and the baseline noise resolution measurement
- ^{133}Ba Calibration Simulation: A sample of events simulated with ^{133}Ba as an input source and run through the full simulations infrastructure. Used for high energy electron recoil studies and for comparison with ^{133}Ba calibration data
- WIMP Simulations: Two sets of events simulated with WIMPSim as an input source and run through the full simulations infrastructure. Used to study nuclear recoil events and detector response for WIMP-like low energy events

We next describe more on each in turn.

7.5.1 Pure Noise Simulation

Pure noise samples are created by simulating events with no interactions. The pure noise simulation sample will allow us to confirm that noise is simulated correctly when compared with data and we will compare the measured detector resolution in the absence of particle interactions against data. We first describe how pure noise samples are created, and then we describe the two types of noise samples we create that way (one as a prerequisite for Reconstruction, the other for our noise studies).

Pure noise samples are created by adding noise-simulated events to zero-energy (empty) TESSim and FETSim events, specifically using the equivalent of a zero-energy Raw Data file as the output of DetectorSim. Since there is no energy deposition in pure-noise there is

nothing to simulate in SourceSim nor DetectorSim, only the electronics noise as in Figure 3.5, which is handled by DAQSim with the primary work being done by NoiseSim. In some sense this sample is a longer version of the pure-noise sample required for reconstructions, the BeginningOfRunRandoms as described in Chapter 5.2.3. While it is known that real detector noise samples are contaminated, as shown in Chapter 6, we will not attempt to simulate non-Gaussian components nor contamination, we will compare to the corrected IRR sample instead.

We create two pure noise samples for use in our studies. The first is a sample of events to study for their energy measurement. The second is a BeginningOfRunRandoms file to be used as the input for use in the Reconstruction for PSD estimates. The pure-noise sample is created from the BlipCleaned PSDs from Zip4, taken during a ^{133}Ba calibration run (from June 2013). The noise-only data sample contains 31,432 events and the three BeginningOfRunRandoms files (one for each of the three samples) contain 500 events each, for a total of 1,500 events. The simulated events are indistinguishable from a typical real data noise event, as seen in Figure 7.5.

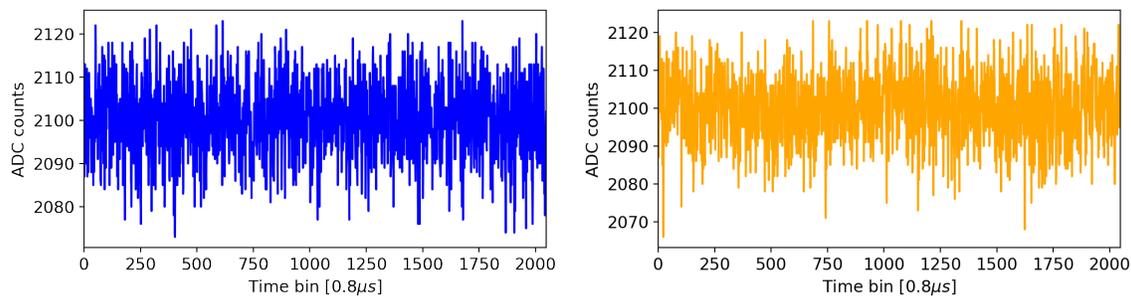


Figure 7.5: Example real data noise event (left), and NoiseSim simulated event (right) randomly sampled from the input PSD. The simulated event is indistinguishable from typical real data noise.

7.5.2 Barium Calibration Simulation

A large sample of ^{133}Ba events is simulated for use in our studies. It will allow for calibration of the simulation energy scale, studying electron recoil response in the full energy range (0 to 400 keV), and studying both well-measured and poorly measured events to understand the detector responses better, in particular with the 356 keV peak which is measurable in both simulations and real data. We first describe how the samples are created and then we make a quick check on the SourceSim stages about the primary particles emitted energies.

The simulation begins with the decay of the radioactive barium isotope (as previously seen in Figure 5.1) in SuperSim and letting the infrastructure handle each stage of the production. As the photons travel through the apparatus material the detectors will observe a combination of photons directly coming from the source as well as electrons and photons from by-product scatters in the apparatus and detectors. The BeginningOfRunRandoms file described above is pre-pended to the sample. Our ^{133}Ba simulation contains 80,000 events of fully simulated ^{133}Ba decays propagated through the apparatus and interacting with iZip4. Two separate versions of the sample are made with the same TESSim and FETSim events, with and without noise.

As a check of the SourceSim stages, before DetectorSim, we show that the simulations are working as expected. As described in Chapter 5.2.1, the ^{133}Ba source produces photons and electrons with known fractions and energies. Figure 7.6 shows the distributions of photons and electrons out of the source, in the top left, which is well-representing the emission distribution in discrete lines [78]; noting that this is before scattering they produce the expected delta functions. After the emitted photons have passed through the apparatus only some of them reach the detector, as seen in the top right of same figure (the others are absorbed by the shielding). Finally, the recoil energy deposited in the crystal (bottom of same figure) is the sum of a number of effects. These include events where the incident photon or electron is fully absorbed by the detector (the peaks in the distribution), where there is a

single scatter (the Compton edge which peaks at approximately $356 \text{ keV}/2 = 178 \text{ keV}$), and cascade-scattering effects from photons and electrons that travelled through the experimental material (in a Compton scattering process) [47, 51]. We note that events with the largest energy deposition in the detector in a ^{133}Ba sample are typically multi-scatter events, implying that the deposited energy was released into more than one primary electron recoil. This is important as it means that not all the recoil energy is deposited in a single location in the crystal, which is a potential cause of mismeasurement since a larger number of scatters can occur in a region of the detector where there is incomplete absorption. Figure 7.7 shows the number of scatters in the crystal for the subsample of events with an incident 356 keV photon, we note that at high energies all events are multiple scatters. Finally, in ^{133}Ba , since all interactions are electron recoils, the recoil energy in the crystal (deposited energy) is the same as the expected charge energy.

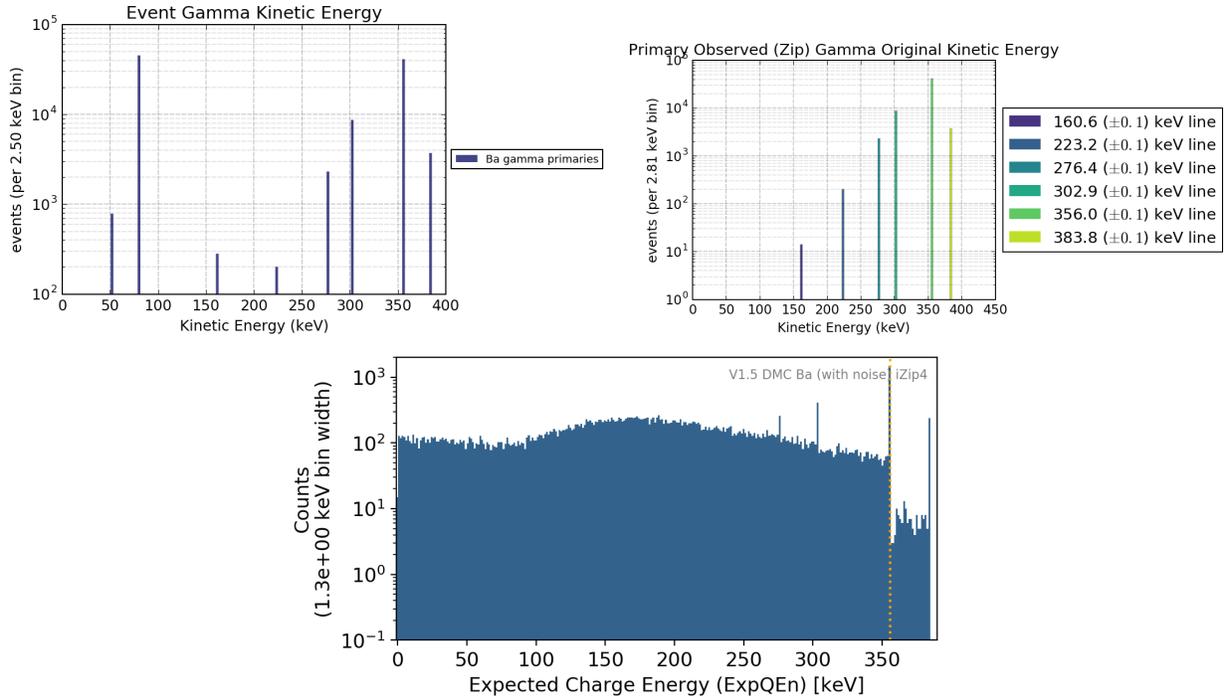


Figure 7.6: Figures showing the energy distribution of photons emitted from a ^{133}Ba source at various stages in the simulation. Barium emits photons with energies following delta functions at the excitation-transition levels (top left), as given in Figure 5.2. As some of the photons are absorbed in the copper shielding, not all of the photons reach the detectors, the original/primary energy of photons reaching the detectors is in the top right figure. The distribution of energy deposited (bottom) in the detector results from the interactions of electrons and photons either directly from the source or after scattering in other material, hence a Compton spectrum in addition to the smeared signature Ba peaks. The expected charge energy is equal to the true recoil energy (all energy is converted to liberated e/h).

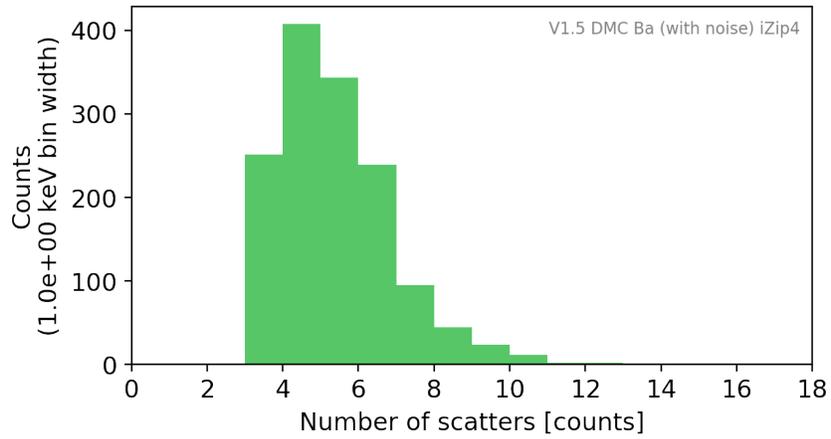


Figure 7.7: This figure shows the simulated number of scatters in the detector for a sample of 356 keV photons from the ^{133}Ba source. We note that for these high energies all events are multiple-scatters, which increases the possibilities of mismeasurement due to position dependent electron and hole collection.

7.5.3 Enhanced WIMP Simulation

The WIMP sample is our signal nuclear recoil event simulation. With this sample we hope that that we can learn new ways for optimizing our search by understanding how the simulations respond to nuclear recoils. We simulate two different samples of WIMP interactions from SourceSim, using the same methods as used for ^{133}Ba but with WIMPSim as the input. In both cases these are single scatter, nuclear recoil events and we assume that the position of the interactions is randomly distributed in the crystal. The energy distributions, however are different. For the first sample, which we will refer to as the WIMP sample, we assume a WIMP mass of 50 GeV and a recoil energy distribution given from Eq. 1.1. This sample consists of 80,496 events for iZip4. The second sample consists of a total of 90,498 events randomly distributed with energies between 40-200 keV. We will refer to this sample as the Uniform Spectrum sample. We will refer to the combination of the WIMP sample and the uniform spectrum sample, as the Enhanced WIMP Sample. Two versions are simulated from the same DMC Intermediate Raw Data files, with and without noise. In Figure 7.8 we show the recoil energy spectrum and the expected charge energy (after the Lindhard yield [60], as given by Eq. 3.11).

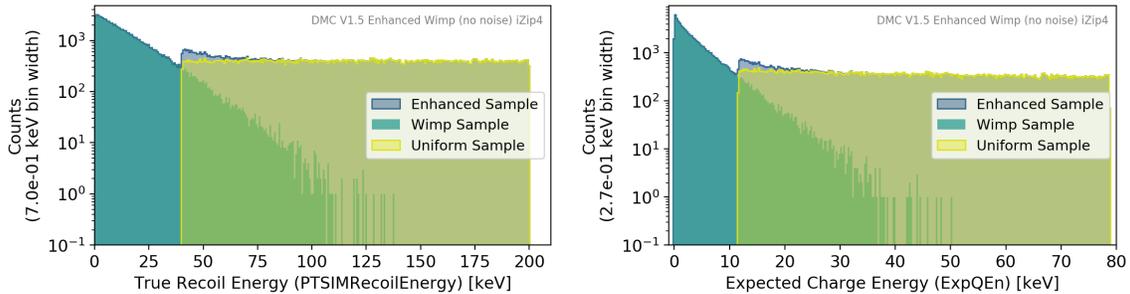


Figure 7.8: The recoil (left) and expected (right) charge energy distributions for our two WIMP samples. The expected charge energy is calculated per Lindhard Theory [60] from the recoil energy. The WIMP sample is created with a 50 GeV WIMP mass, and the other sample is created by sampling a uniform random spectrum of 40 to 200 keV recoil energy (equivalent to ~ 10 to ~ 80 keV expected charge energy).

Now that we have described the process for generating simulated samples, and the samples themselves, we are ready to move on to the comparison with data and to what we have learned from them. Chapter 8 is about comparing the pure-noise simulation to data and about measuring the detectors baseline resolution. Chapter 9 is about the charge simulation results from events with real interactions, in particular comparing the charge readout results with the ^{133}Ba calibration to understand the full spectrum of electron recoils, and Chapter 10 is about the simulation predictions, about the accuracy and resolution as calculated from truth information and what the simulation tells us about mismeasurement causes including our WIMP signal simulation to look closely at low energy nuclear recoil events response.

8. COMPARISON OF NOISESIM TO DATA AND BASELINE ENERGY RESOLUTION

With the samples of simulated noise events and InRunRandoms (IRR) corrected and cleaned we are now ready to validate NoiseSim and to get a good measurement of the baseline energy resolution due to noise. In particular, we will use the expectations from Chapter 5 to predict the optimal filter resolution from a measured PSD, and compare to the IRR sample after readout mismeasurement cleanup, as described in Chapter 6. We show that the noise simulations agree within error to the Blip-corrected data although there is still a small non-Gaussian effect from Single-Detector Blips that we cannot remove. With our noise simulations in hand we will use them in the next chapter on charge simulations for particle interactions, where we can separate the noise contributions and effects among the rest of the properties and causes of mismeasurement from other effects to the total resolution.

In this chapter we will compare our noise simulations with data, to show that they are in agreement, and measure the baseline energy resolution. We will use the Blip corrected IRR sample (described in Chapter 6.4) and use the BlipFinder tool (Chapter 6.3) to generate a BlipCorrected PSD for input to the simulations. We begin with Section 8.1 describing the expectations from the noise simulation, and the method for measuring the baseline energy resolution. Next, in Section 8.2 we show how the BlipFinder (from Chapter 6.3) is used to create a BlipCleaned PSD for use as input to the simulations. We then show in Section 8.3 that NoiseSim well reproduces the same characteristics of the optimal filter observables on IRR data (centered mean, same RMS, and same Gaussian shape within error) taking into account the contamination from Single-Detector Blips.

8.1 NoiseSim Expectations and Energy Resolution Measurement

With the expectations from Chapter 4 for the 0-delay energy observations we can compare data and simulations to each other as well as to predictions. Prior to showing our results,

we briefly describe the methods we will follow for validating our noise samples and measure the resolution. Next we will describe why we need to create a noise-profile (PSD) clean from Blip effects, for use as input to NoiseSim.

For our pure-noise samples, in all 0-delay versions of the variables in Table 4.1 we expect the E_0 distribution to be a Gaussian with a mean of zero and the RMS to be a measure of the resolution. This can be compared to the theoretical resolution (Eq. 4.9), as well as between simulations and real data. We will start from our Blip corrected IRR sample, but we note that since there is still some contamination of the sample due to Single-Detector Blips, which affect the low side tail of the energy measurement, we will do Gaussian fits in the region from $\mu - 2\sigma$ to $\mu + 4\sigma$, to avoid the region where the tail effects could skew the fit. Finally to compare our simulated results, we will perform the same fit to our noise simulation and we will present both distributions overlaid with the expected uncertainty of each bin of our IRR sample to show that noise is within the expected variation. For our comparisons we will present both simulation and IRR distributions overlaid for all channels and the qimean and qsummax estimators, and we will report the observed resolution of both along with the expected.

As described in Chapter 7, while a simulated version of noise requires an input PSD, we note that the originally created PSDs were contaminated by Blips. Thus before comparing our BlipCorrected IRR sample we need a matching noise simulation since the original PSDs are contaminated by Blips. In the next section we describe how to make a modified version of the input PSDs to the simulation. With these BlipCleaned PSDs we will be well set for the measurement of the resolution and comparison of the BlipCorrected IRR sample with the pure-noise simulation 0-delay energy estimates.

8.2 Estimation of BlipCleaned PSDs for Use in the Simulations

With our understanding of Blips and TowerBlips from Chapter 6, we now explain why we need a new PSD calculation for use in our studies. Since Blips contaminate the IRR sample, and the original input PSDs (from the event reconstruction) do not remove them,

the original PSDs do not reflect the Gaussian noise, but rather the contributions of both pure-noise and Blips. A version of the PSD from a BlipCleaned sample is needed for a simulation that reproduces the Gaussian (corrected) 0-delay measurements of the corrected IRR sample (which is our best pure-noise data sample).

In this section we will show how we use the BlipFinder tool to identify the BlipFinder energy threshold that selects events that reproduce the observed resolution from the Blip-corrected IRR sample, to create an input PSD for simulation. First we will show how we create subsamples from our IRR (not corrected) data to calculate the theoretical resolution of each. Next we show how we pick a sample that well-represents the observed resolution of the BlipCorrected IRR sample.

A set of BlipCleaned PSDs can be selected from the no-pileup IRR sample assuming that we can select a sub-sample that well-represents the expected optimal filter resolution of pure-noise effects. We begin by sub-dividing the IRR data with the BlipFinder tool utilized in Zip4, full search window. The data is sub-divided by Blip energy ranges (from -5.9 to -1.5 keV, with steps of 0.2 keV). For each sub-sample a PSD is calculated, and for each PSD we can estimate the theoretical resolution of the optimal filter (Eq. 4.9). Figure 8.1 shows the theoretical resolution as function of BlipFinder energy bins.

Since we have measured the BlipCorrected RMS, Figure 6.15, using Figure 8.1 we know which sample of events best reproduces that RMS. We select the events within ± 100 eV from the intersect value (-2.18 keV for qi1), and take that PSD as our BlipCleaned PSD. The BlipCleaned PSD is shown in Figure 8.2, along with the original IRR PSD. The low frequency contributions are much lower, as expected for a long peak structure like Blips. This BlipCleaned PSD is used for the input to the simulation to be compared with the IRR data (Blip corrected).

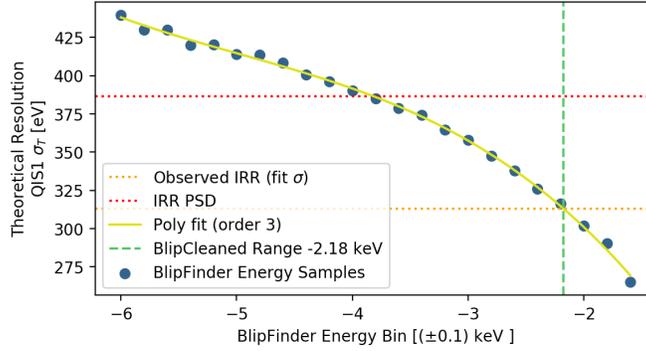


Figure 8.1: This figure shows the theoretical resolution of the 0-delay q11 energy measurement as a function of BlipFinder energy sub-sample selection. The horizontal dashed lines show the value from the original PSD (red) and the observed from the corrected IRR fit (orange). The BlipFinder energy sub-sample selection value that intersects the orange dashed line (vertical dashed green) is the value for the final sub-sample selection to create the BlipCleaned PSD.

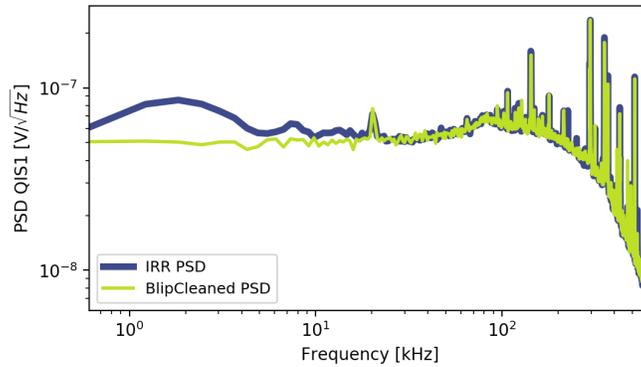


Figure 8.2: This figure shows two different versions of the PSDs. The first is from the BlipCleaned sub-sample that reproduces the observed RMS of the BlipCorrected sample. The second is the original IRR PSD (as calculated by the event reconstruction method). Note the difference in the sub-10kHz range.

8.3 Comparison of Data and Simulations, Final Resolution Results

With the matching input to the noise simulations and data in hand we are finally ready to qualitatively and quantitatively compare the BlipCleanedPSD simulated sample with the BlipCorrected IRR data. We first show an example of the Gaussian fit to our noise simulation and how the mean is consistent with zero and the RMS is a measure of the

resolution. Next we show the overlaid simulations and data to show that they are consistent within uncertainties with the expectations and with each other. Finally, we comment on the small discrepancy between the qsummax noise and simulations left-hand-side tail, which is due to effects correlated amongst all channels (Single-Detector Blips).

As an example of the fit applied to a simulated sample and taking account the Blip-Cleaned PSD and the expected resolution from Eq. 4.9, we show Figure 8.3 (inner charge channel, side1, electron readout side), and we note that we expect a mean of zero and an RMS of 317 eV, and observe a mean of 0 ± 2 eV and an RMS of 312 ± 4 eV is observed, giving us some basic confidence in our simulation tools and validation methodology that are working as expected before comparing to data.

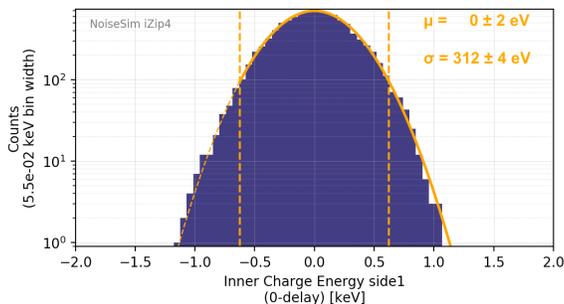


Figure 8.3: The measured 0-delay energy distribution of a fully simulated and processed NoiseSim sample of events from qi1, where we fix the time delay of the fit to be 0 as consistent with noise estimates described in Chapter 5. Note that it is a single Gaussian with a measured mean of zero, that there are no tails, and that the expected RMS is 317 eV while the observed RMS is 312 ± 4 eV.

Figure 8.4 shows the data from the BlipCleaned IRR sample along with the simulated data, where the number of simulated events is normalized to be equal to the number of observed events in the data. The six different 0-delay energy distributions (qis1, qos1, qis2, qos2, qimean and qsummax, as described in Table 4.1), along with their Gaussian fits are given in Table 8.4. A number of qualitative and quantitative remarks are in order. The first is that the simulations clearly look like the data and most of the distributions do not

have non-Gaussian tails; we will comment on the ones that do in a in a moment. For both the real data and the simulations all the means are zero within uncertainties as expected. A comparison between the theoretical resolution (predicted from Eqn 4.9 from each of the BlipCleanedPSDs), the measured from the RMS from NoiseSim (the expected resolution) and from the real data (observed resolution) are consistent within uncertainties. Qimean is also in agreement, but qsummax has a smaller LHS tail in the simulation, due to correlation of Single-Detector Blips (which are not corrected), concluding our noise simulation studies.

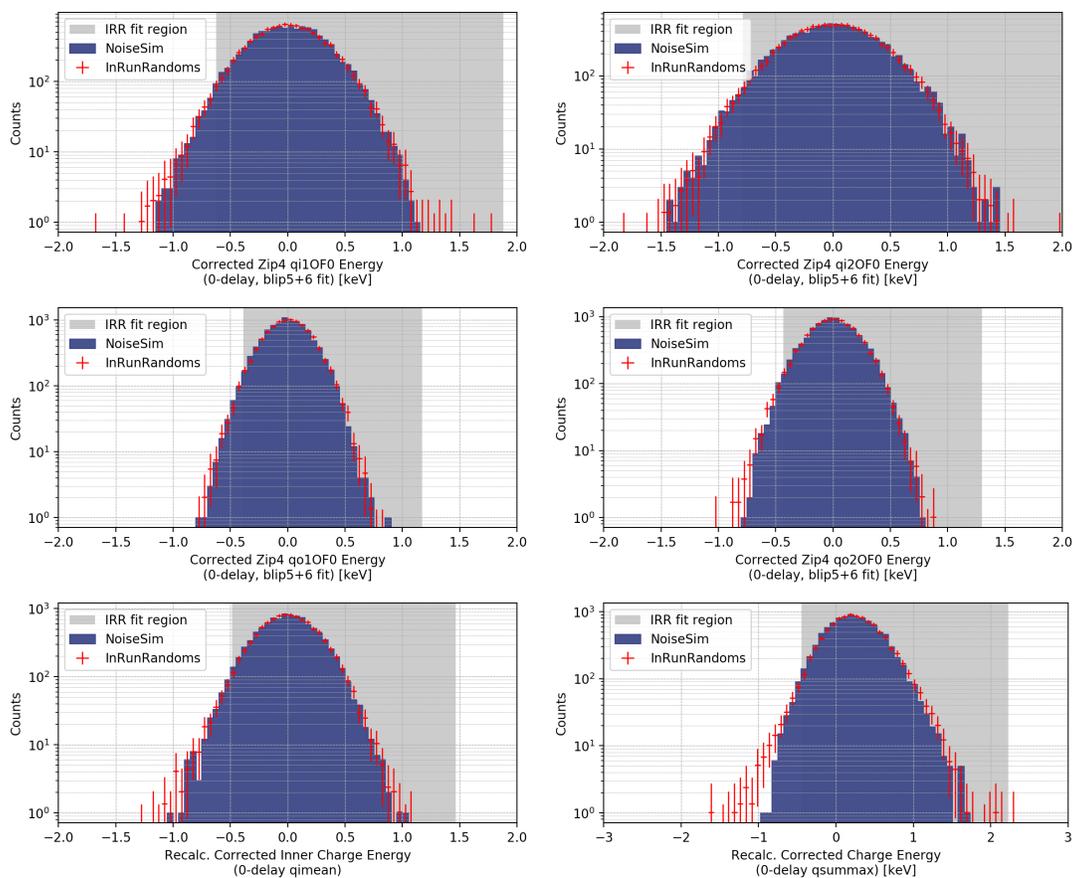


Figure 8.4: Overlaid noise simulations (blue) and Blip-Corrected IRR sample (red error bars) 0-delay energy measurements for qi1, qi2, qo1, qo2, qimean, and qsummax. The fit region for mean and RMS measurements is shown in grey, with the results shown in Table 8.1. Simulated samples are in agreement within error, except for qsummax LHS tail, which is produced by correlations of single-detector Blips.

0-Delay Energy Measurement Resolution						
Sample	qi1 [eV]	qo1 [eV]	qi2 [eV]	qo2 [eV]	qimean [eV]	qsummax [eV]
Theoretical (Blip-CleanedPSD)	317	196	398	217	NA	NA
Observed NoiseSim (BlipCleanedPSD) <i>Err</i> ~ 4 eV	312	193	393	222	250	344
Observed IRR (Corrected) <i>Err</i> ~ 4 eV	313	195	394	217	244	334

Table 8.1: For each of the six different 0-delay energy measurements, taken from Figure 8.4, we give the theoretical and observed resolution of simulated and real noise data. In all cases the resolutions are in agreement within the 4 eV measurement uncertainty.

Before moving to the next chapter some notes on the non-Gaussian tails in the qsummax data are in order. We note that since we do not have a way for correcting the IRR sample contamination from Single-Detector Blips, and since they are correlated in both sides, the contamination will affect the data, but not the simulation. This shows up as a bigger LHS tail in qsummax IRR than there is in the simulation (where there is no correlation).

Now that we have shown that NoiseSim works as expected, and well-reproduces (qualitatively and quantitatively) real-data Gaussian noise, we can use it in the simulations. With our samples with and without noise in hand we can move to the next chapter to study the accuracy, resolution and causes of mismeasurement in our simulations of ^{133}Ba calibrations and WIMP-like interactions.

9. RESULTS FOR REAL ^{133}Ba DATA AND SIMULATIONS

With fully simulated events in hand, and data to compare to, we start our studies of the charge simulation response to particle interactions. There are a number of things we want to know:

- How well does the simulation reproduce the expectations of Chapter 2 and 3?
- How well does the simulation reproduce the data? (The focus of this chapter)
- What does the simulation tell us about sources of mismeasurement? (Both in this chapter and the next)

While this thesis will not fully answer all of these questions, we will show the progress so far. In particular we have learned a lot from event-by-event comparisons of fully simulated ^{133}Ba events to input truth values as well as from comparisons of samples of simulation data to calibration data. Since there are an infinite number of things we could compare, we will preview what we have learned as a way of motivating the steps we took to get there. Some of the results will be shown in this chapter and some in the next, for pedagogical reasons. As we will see, the simulations show that, as expected, events depositing the bulk of their energy in the fiducial part of the detector are fully measured, and can be badly undermeasured when they don't. Qualitative comparisons of the energy measurements in the individual charge channels data show areas of agreement, but quantitatively they show that the energy resolution in simulation is much better than observed in data. With a first set of comparison results in hand we are well set for a fuller set of what the simulations predict about the detector response and next steps discussed in the next chapter.

In this chapter we present the results of our studies of what the simulations predict for our detector model and compare to data with an emphasis on the events that are expected to be fully measured. Since there are many complicated and subtle tasks, before presenting the full

results we give an overview, in Section 9.1, of the expectations, goals, big picture methods, and a preview of the results of our studies. In Section 9.2 we follow with a description of the methodologies for our studies. Next, in Section 9.3 we use the simulations to predict what physical region of the detector produces well-measured events. Next, in Section 9.4 we study and compare the simulations and data variables used for the LT Fiducial criteria, with an eye towards what the simulation tells us about them. Finally in Section 9.5 we compare the results from real and simulated ^{133}Ba data.

9.1 Expectations, Goals, Methods, and Overview of the Results

In this section we describe our expectations and goals in Section 9.1.1 and give a preview of our results in Section 9.1.2. With our path well motivated we will be well set to describe the methods we used to accomplish them in Section 9.2.

9.1.1 Goals for the Simulations Comparisons and Measurements

The ideal for simulations is for them to be so trustworthy that they can be used in optimized searches for WIMPs. This means they must match the data in cases where the outcomes are known as a way of showing that we understand that we intrinsic accuracy and resolution of the detector, as well as the causes of mismeasurement. Since, as we will see, this is a complicated task and we are not there today, there will be future versions of our simulations software which will require many iterations as well as tuning of the simulation parameters. To motivate and understand the results in this work, we start off describing our expectations for what we can and cannot do in our simulations. With those in mind, we next describe our more limited set of goals. Then we overview the methods we use to compare and understand our samples. Finally we note the next clear tasks and studies that we have not done yet, prior to anticipating our results in Section 9.1.2.

Before presenting the comparisons to data it is important to describe our expectations for what we can and cannot do with our simulations and data so it is clear where we are going and why. This is especially true when comparing to samples of events where we don't

understand the full set of effects in the data. In the last chapter, we understood what we were measuring and what the expectations were, but in this chapter we know far less about our real data. In this work we do not show that the samples are fully understood and well-representing data in all possible ways, but rather that we are initiating the iterative process of understanding them and comparing with data. The full validations of these complex simulations will require many iterations, and here we present a major leap in clarity for where they stand and where we are going with them.

The goals in this study are limited to understanding what the simulations predict for the detector response and doing first comparisons to data and identifying the dominant issues that need to be worked on. This includes studying how well the simulation reproduces the expectations for the detector response expectations from Chapter 2 and 3 (results shown in Sections 9.3 and 9.4), how well the simulations reproduce our data (Section 9.5), and what the simulation suggests for causes of mismeasurement (Chapter 10.5). With this information in hand, we will describe the next steps to be worked on in the simulations development (Chapter 10.6).

Before showing our results and comparisons we note the clear tasks that we haven't done yet. The first is that while we compare the event reconstruction results for electron recoils from ^{133}Ba data with the simulations, we will not show them for nuclear recoils from ^{252}Cf as that study is not completed. For nuclear recoils, in Chapter 10.4, we compare the expected simulation-based truth quantities with the observed values. Similarly, we have not completed the phonon studies so nothing from them will be presented. Finally, the process of fully iterating the improvements and tuning of the simulation parameters have not been completed, so our results will not fully match data observations and we will simply point out places where they differ. Comments on these issues will be discussed in Chapter 10.6. With these caveats in mind we turn to what we did: focusing on the simulation-based fiducial volume, the understanding of the data event selection criteria, and the data comparison in this chapter; and following with simulation-based expectations, new understanding, and

future steps in the next chapter. We begin with a preview of the results as a way of motivating the methods used.

9.1.2 Preview of the Results

Since we have not completed all the studies, and we don't understand some of the results yet, it is useful to preview what we have learned as a way of helping make the description of the methods we will use (next section) and the set of results we will cover (Sections 9.3-9.5, and 10.2-10.5) more understandable in their context. Our results suggest that the simulations of the charge energy readout measurements behave as expected qualitatively (reproduce the expectations for the detector response) and well-reproduce the overall features of the ^{133}Ba sample. On the other hand there are quantitative differences, for example the resolution of our simulations is much better than it is in data for reasons we don't yet understand. While this will require more study and work, we provide a brief summary of our results and describe how they tie to our goals.

Our limited goal is to understand how well the simulation reproduces the expectations for the charge energy measurement (Chapters 2 and 3), how well it reproduces data, and what it suggests for causes of mismeasurement, so the following is an overview of the results that our data and simulations tell us to answer these questions:

- The data and simulations samples are well calibrated by the 356 keV method (in this chapter, Section 9.5.1)
- In simulations we observe that the energy response is linear in charge collection, which is consistent with the expectations of Chapter 5, and with the use of single calibration constants (next chapter, Section 10.3)
- With simulations the mismeasurement is mostly associated with position dependence, and we can determine what the detector model predicts that the well-behaved region should be. We call this the SimFiducial region (this chapter, Section 9.3 and 10.2)

- The real ^{133}Ba data shows events which are badly mismeasured but are removed by the LT Fiducial cuts. With simulations we can compare the LT Fiducial cuts with the SimFiducial expectations (both chapters, Sections 9.4, and 10.5.2)
- Qualitatively our simulation well-reproduces the ^{133}Ba data in that the energy measurements of side1 (electrons) is different than side2 due to differences between electron and hole propagation effects. On the other hand we see quantitative differences between simulations and data in that the resolution of the 356 keV peak is much better in simulations than data (this chapter, Section 9.5.1)
- Simulation events show that when we have multiple-scatter events, if a large amount of energy is in the non-SimFiducial region we get poor measurements (this chapter, Section 9.4.5)
- Simulations tell us that LT Fiducial does a good job of removing events that are in the poorly measured region, by 1) removing the ones that produce a lot of energy in the outer part of the detector (i.e., scatters fall outside SimFiducial, with high radius), and 2) where the two sides do not measure the same thing (typically because one of them is mismeasured because a large amount of energy is not deposited in SimFiducial). While the LT Fiducial requirement does well in most cases, it has the downside that it lets in events when both sides are mismeasured (both chapters, Sections 9.4, and 10.5.2)

9.2 Data Comparison and Simulation Predicted Accuracy and Resolution Measurement Methods

While the list of goals and answers here was presented in an order that built up the understanding, we will order our discussion by focusing on the methods we used to determine these conclusions. We begin with a discussion of the methods in Section 9.2.1, and follow with a description of the simulation-based variables used for the comparisons in Section 9.2.2. We will start the presentation of the results in Section 9.3.

9.2.1 Overview of the Methods

With our goals, expectations, and preview of our results in hand we next describe the methods we use to help us understand the detector response and the mismeasurement causes. Before our simulation existed, the typical starting place was to use the one piece of known information in the data: the 356 keV line which is readily observed in ^{133}Ba data. We will take a different tact. Specifically, we will start with Section 9.3 where we use our simulation-based mismeasurement quantity to make a clean selection of badly measured events and look at their position dependence to define our simulation-based fiducial region, SimFiducial. We next follow by looking at the LT Fiducial variables used to identify well-measured events for the full set of ^{133}Ba events for both real data and simulations in Section 9.5. With these we can begin to see the differences as well as understand more about what the variables do and don't do. Next, in Section 9.5.1 we compare our simulations with data by observing the dominant 356 keV photon energy peak from ^{133}Ba calibrations and measuring its mean value and resolution, and showing that our samples are well calibrated. Finally in Section 9.5.2 we do a qualitative comparison of the full energy range in data and simulations.

We next describe the simulation mismeasurement estimators (from truth information), what the expectations are for energy distributions of well-measured events, and how we use them to select well-measured and poorly-measured samples.

9.2.2 Variables Used for the Comparisons

To begin our studies of what events are well measured (as predicted by simulations) we first need a simulation prediction of mismeasurement, from the truth-based quantities, like we did in the last chapter. The expectation for this value should be far from zero for poorly-measured events, close to zero for well-measured events, and an ensemble of them should allow for a measure of their resolution. Quantitatively we define mismeasurement as ΔE :

$$\Delta E = E_{\text{observed}} - E_Q , \tag{9.1}$$

where $E_{observed}$ is the measured energy (either by qi1, qi2 or qimean) from the fully processed reconstruction of the simulated events, and E_Q is the expected charge energy (simulation truth) as given in Eq. 3.11 from the recoil energy. Any event with a large value of ΔE is considered mismeasured (i.e. indicating that a real particle in a real detector would have been mismeasured), so samples of such events allow for studies of their causes.

Since not all events have the same resolution, and it is not always possible to have pure samples of well-measured events we will use variations of these variables and measurement fitting techniques. For example, as we will see, the resolution gets worse linearly as a function of energy so we will also use the variable $\Delta E/E_Q$ to compare ensembles of events with different expectations of resolution. Similarly we will sometimes do asymmetric fits, for example, in the case of multiple-scatter samples, where a significant part of the total energy deposition is mismeasured such that a left-hand-side tail in the distribution (observation of undermeasurement) is prominent.

With these goals and methods in mind we can compare our simulation results, both to truth and to data for samples where we have some understanding about the particle interactions that produced them. Before doing this comparison, we start with the definition of our SimFiducial region in the next section, since it is of use for understanding mismeasurement.

9.3 Using the Single-Scatter Sample for Finding the Simulation Detector Region of Well-Measured Events (SimFiducial)

As Section 3.3 predicted, and as we will show, the dominant cause of mismeasurement is the location(s) within the detector where energy is deposited. We begin our study by seeing what the full simulation results tell us, including the simulation-based true fiducial volume. We denote this volume SimFiducial and it delimits the region of poorly measured events. We first describe why we selected the Enhanced WIMP sample for determining the SimFiducial volume. Next we show that our sample is well calibrated relative to the simulation-based truth quantity, and from it we select a clean sample of mismeasured events.

Before starting, we note why we begin this study with our Enhanced WIMP sample. As

shown in Figure 9.1 for the ^{133}Ba simulation, most events, especially at high energy, will scatter multiple times in our detector. As we will see, multiple scatter events can deposit a significant amount of energy in very different places in the detector such that some of the energy within the event is well measured and some is not. Specifically, when a scatter with enough energy is near the edges (as will be clear in later sections), energy is lost as charges get trapped by the edges or the interleaved potential. We also note that, as we will later show, electron and nuclear recoils appear to behave in a similar fashion so there are no additional considerations that we will have to take into account. For these reasons, for our SimFiducial region determination studies we use our WIMP samples as all events are single scatters.

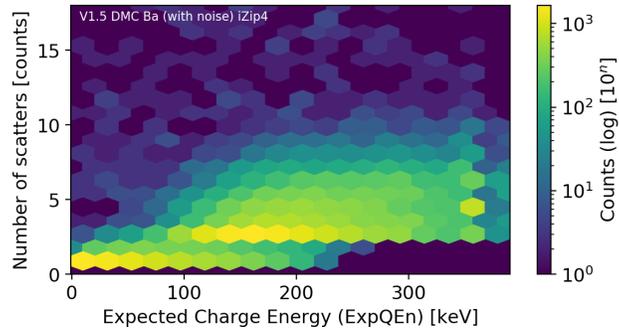


Figure 9.1: Heat-plot histogram of the number of scatters as a function of recoil energy for the ^{133}Ba simulation. Above 230 keV all events are multiple scatters. Most of the events have less than 8 scatters, and below 100 keV less than 3.

We begin showing the simulation mismeasurement distribution from our Enhanced WIMP sample to show that it is well calibrated. Two notes are in order. The first is that, as we will show, the resolution is energy dependent so we will also use $\Delta E/E_Q$ and the second is that since there is an important measurement effect for energies below 3 keV (described in Section 10.5.1), we will remove those events from our sample for now. The results are shown in Figure 9.2. The top two plots shows the ΔE measurement (the right hand side is a blow

up of the left hand side). We see that while many of the events are well measured (peak at zero), many are very poorly measured (tails on both sides of the distribution). These are due to two effects: mismeasured events (far left hand side) and that the resolution of the measurement is energy dependent (both sides). The bottom plots show the same events, but where we have used $\Delta E/E_Q$ which takes into account the resolution. In this variable the right hand side tail is gone, and we are left with the sample of poorly-measured events. We make a fit on the well-behaved events with an asymmetric Gaussian fit just around the peak, which immediately suggests that the sample is well calibrated (mean at zero within errors, shown in the right of the same figure) and with an RMS of 0.21%.

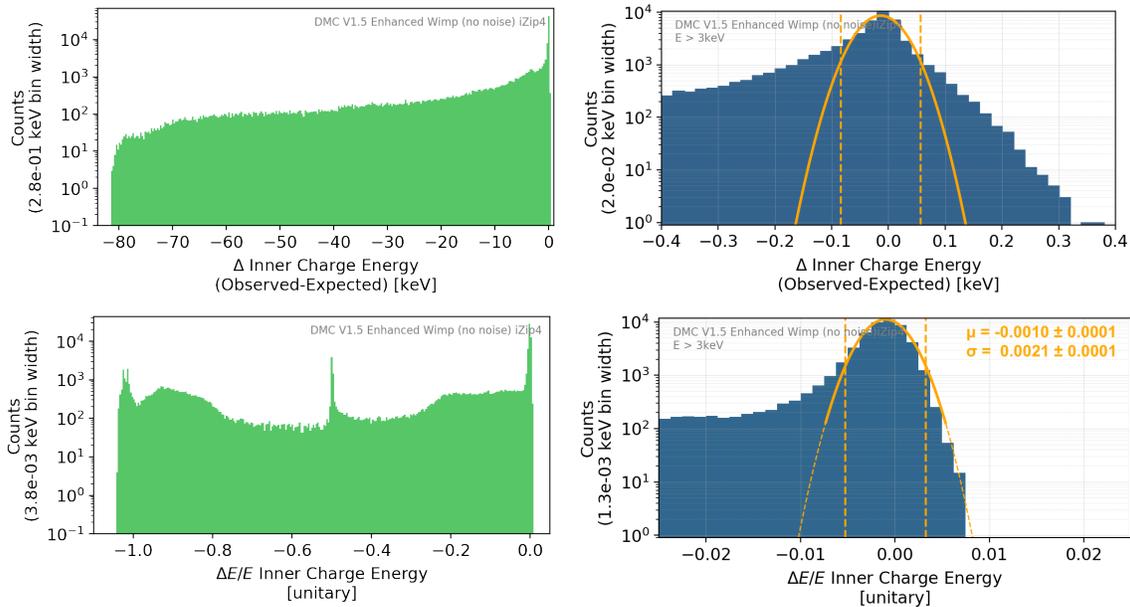


Figure 9.2: The figures show the mismeasurement (ΔE) and relative mismeasurement ($\Delta E/E_Q$) of the single-scatter nuclear recoil (Enhanced WIMP) sample, with energy above 3 keV (to avoid the lowest energy mismeasurement which is described later). On the left plots we show the full-sample where we can see many badly mismeasured events, and on the right plots we zoom-in to the well-measured region. In the top we note that the peak is not Gaussian due to energy dependence, so we turn to the relative mismeasurement as the resolution is linearly dependent on the energy (in the bottom). The well-measured region is a single Gaussian peak, centered at zero. We will select a clean sample of badly measured events by selecting events, below $\Delta E/E_Q < -0.015$, 7σ below the mean, to look at the mismeasurement position dependence.

To define SimFiducial we plot the interaction location of the poorly measured events. These events are selected with $E/E_Q < -0.015$ as they are over 7σ below the mean. The results are shown in Figure 9.3, and the yellow dashed lines indicate where we have chosen to define SimFiducial. As expected the events that interact near the top and bottom interfaces are trapped by design and not fully measured, and the ones deposited near the outer-parts of the detector have energy loss on the sides. The asymmetric shape is due to differences in propagation between electrons and holes. We first define the Z limits by selecting a sub-sample events away from the sides, as shown in the left plot. For the events with large R we use a single line. The values are given by the limits on Table 9.1.

Removed Region	Description
$Z > 0.0228 \text{ m}$	Remove events near the top
$Z < 0.0025 \text{ m}$	Remove events near the bottom
$R > 0.596 Z + 0.014 \text{ m}$	Linearly determined from $R_{top} = 0.0274 \text{ m}$ and $R_{bottom} = 0.0153 \text{ m}$ and their respective Z -intersections

Table 9.1: The set of criteria used to remove events as being poorly measured, the SimFiducial region. We use the the shape of a truncated cone. The height of the cone is given by Z_{top} and Z_{bottom} to avoid the interleaved potential trapped events. The top radius is given by R_{top} , and bottom radius by R_{bottom} , and from their intersection with the Z -limits we calculate the side limit line which is a function of R dependent on Z .

With a demonstration that mismeasurement effects are position dependent we move to a comparison of data and simulations starting with the LT Fiducial criteria.

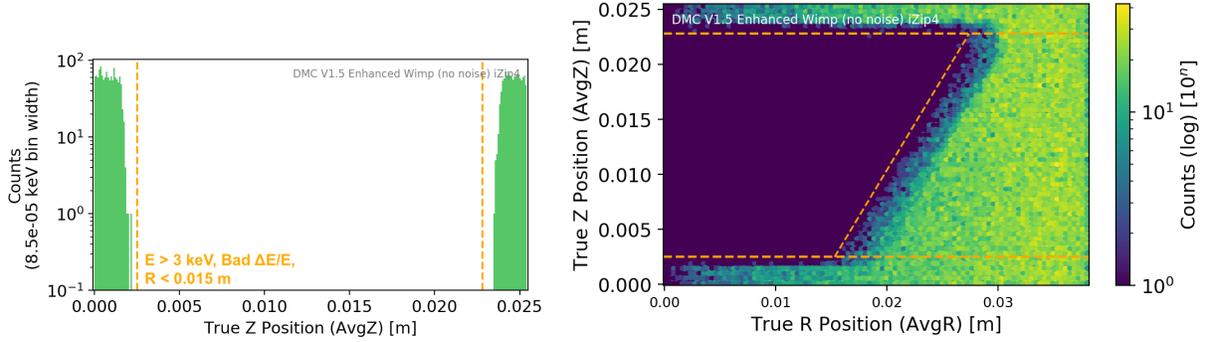


Figure 9.3: This figure shows the true position of a clean sample of mismeasured events ($\Delta E/E_Q < -0.015$ with $E > 3$ keV) with a single nuclear scatter. On the left we show the Z-dependence limits, selected from events with low radius. On the right we show the R vs. Z position of the mismeasured events which are contained in the truncated cone volume (dashed orange), we denote the inner-volume as the SimFiducial region with the limits defined in Table 9.1.

9.4 Study of LT Fiducial Variables in ^{133}Ba Data and Simulations

Our comparisons between ^{133}Ba data and simulations begin with the distributions of the energy measurements from the top and bottom, qi1 and qi2 respectively. These are important as they are used by the LT Fiducial criteria and what the simulations tell us about the mismeasurement and effects of such criteria. We will show that the LT Fiducial criteria remove many of the badly measured events but still, due to position dependence of the scatters in the events, long tails and mismeasured events are still expected. Since we do not have the true position of the event in real data we start by comparing the measured values of the qi1 and qi2 measurement variables and LT Fiducial construction criteria in data and simulated events. While ideally they would agree, our goal is understand where they disagree and why. As a reminder, the LT Fiducial definition requires both no energy in the outer channels, and that the qi1 and qi2 measurements be the same within tolerance. After this section we will be ready to use the LT Fiducial requirement in our comparisons between data and simulations.

In this section we start with a comparison of energy measurements of the top and bottom

charge readout sensors, in Section 9.4.1, where well-measured events are expected to show the same energy deposited on both top and bottom, both collecting energy according to Eq. 3.14. While both sides are calibrated to produce the expected mean energy measurement, as we will see, the simulation results show that both sides produce a very similar response, while data appears quite different for reasons we will describe only briefly (as they have not been fully understood). In Section 9.4.2 we describe, from the simulations, the set of events with no energy deposition in one side, to see that these are events with most of their scatters outside our SimFiducial volume. We follow with Section 9.4.3, where one side provides a good measurement (of the 356 keV peak), while the other is poorly measured and how the no-outer energy constraint helps remove most of these. We then look at events that pass the no-outer LT requirement in Section 9.4.4 to look at the remaining events and the symmetry constraint. Next in Section 9.4.5 we discuss what the simulation tells us about poorly measured events that pass both of the LT Fiducial requirements to show what the simulations suggest that the reason for mismeasurement is mainly due to the location of the scatters being near the edges. We will present more on this in Chapter 10.5.2. Section 9.5 will compare the full data samples.

9.4.1 Full Energy Range Comparison of Inner Energy Estimates

We begin our LT Fiducial study with Figure 9.4 which shows qi_{12} vs. qi_1 for all events in the calibration samples (before the LT fiducial requirements) for both data and the simulation. A number of features are readily observed. The first is the large band of events at 45 degrees, where the events are well measured. On the other hand both plots show large numbers of events where the two measurements are very different (i.e., at least one is badly mismeasured). From the figure on the RHS, the simulation, we can study the causes of the most straightforward mismeasurement cases, highlighted with the colored dots. We select from simulations 4 mismeasurement regions to look at position dependence, 2 are from events with no energy measurement on the opposite side (orange and yellow) and 2 of them from 356 keV peak events where one side provides a good measurement but the other does

not (maroon and red). We deal with each in turn.

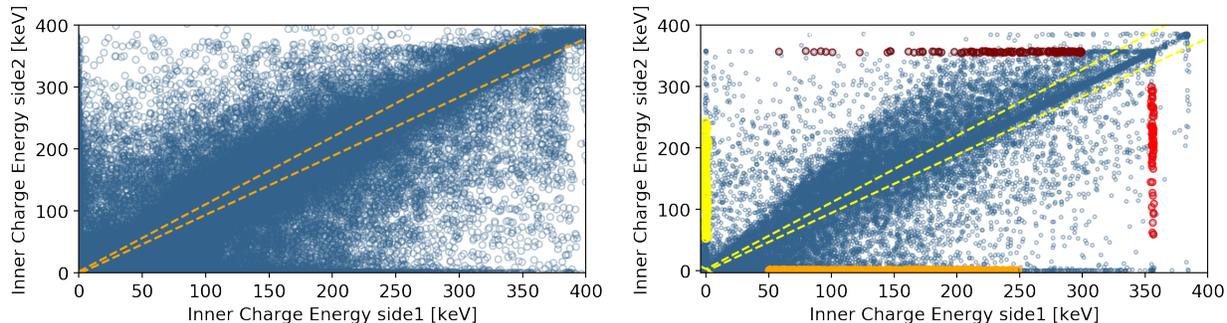


Figure 9.4: Comparison of qi1 and qi2 energy measurements of the data (left) and simulation (right) full samples. The first thing to note is that the bulk of the distributions fall along the line where the two have equal values, but the distribution is much wider in data (worse resolution). While the yellow dashed lines indicate the LT fiducial requirement (which appears to be reasonable for real data, but would be much too loose for simulations), both samples show events where the energies are very badly measured and the values from the two sides are very different from each other, in some cases where one side reads out zero energy while the other side has a large energy. The colored events in the simulated sample indicate which events are shown for study in Figures 9.5 and 9.6.

9.4.2 Events with No Energy Deposition in One Side

We next select the 0-observed energy mismeasurement region events from the simulation and look at the position of the scatters that deposited more than 100 keV of energy on the other side. The results are shown in Figure 9.5. The first thing we notice is that the mismeasurement is due to position dependence of the scatters. In particular the mismeasurement occurs when the events deposit energy close to the top or bottom surfaces. When they are this close to the surfaces, most of the scatters are absorbed at the interleaved potential causing complete mismeasurement in the opposite side (explaining the bottom and right selected regions), also for inner-side2 there is a 0-energy measurement for events with large radius because holes (which travel ballistically) either get stuck at the side-wall or get entirely absorbed at the outer-side2 electrode. We note that the left shows no events with high radius

because electrons propagate in the L-valleys, which accounts for a wider cone that results in some deposition in qi1. On the other hand, for holes which propagate ballistically it is possible to have events near the edges (like on the left of the figure) that deposit all of their energy in the outer channels and have no deposition in qi2.

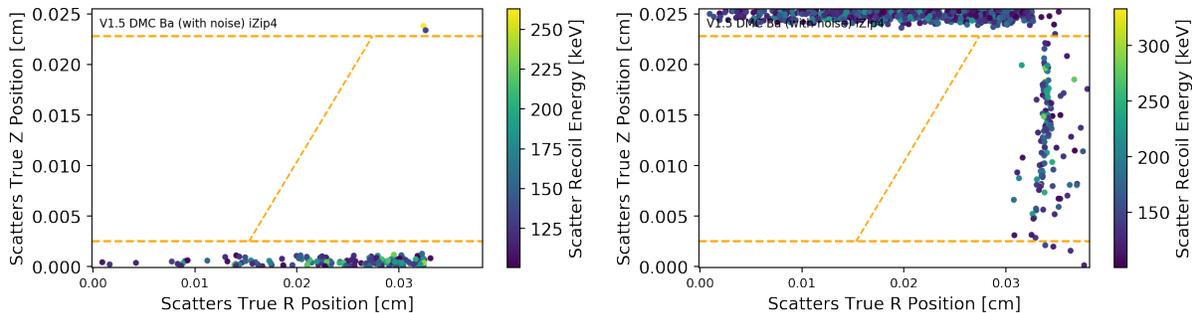


Figure 9.5: The true energy deposition physical location in the detector (Z vs. R) for badly mismeasured events, left is for 0-energy measurement in qi1 (no electron collection) and right is for 0-energy measurement in qi2 (no hole collection). The events are selected as being either the orange or yellow from Figure 9.4 where only between 50 and 250 keV of energy was deposited on the other side to ensure a large energy scatter. The points represent only scatters with more than 100 keV deposited from such events. This shows that the mismeasurement is related with position dependence in particular due to the interleaved potential as scatters interacting close to it get absorbed by the near-end surface while mismeasured by the opposite side. Furthermore, events with large radius for inner-side2 produced a 0-energy observation because all the holes either got trapped in the detector side-wall or entirely got absorbed by the side2-outer electrode (because holes propagate freely), and we note that for 0-energy in qi1 there are no scatters with large radius (left plot) because electrons propagate in the L-valleys (with a wider angle) causing a non-zero deposition in qi1.

9.4.3 Events with One Side Well-Measured and the No-Outer Energy Requirement

Next we look events where one side had a full energy measurement, but the other had a bad energy measurement. These are marked as red and maroon in Figure 9.4 and are selected as being well-measured in one side (within 2 keV from 356 keV) while poorly measured in the other (between 50 and 250 keV in the other side). While we start with the truth-location

for these events, see top of Figure 9.6, the results show that the bulk of them are near the edges or near the outer portion of the detector, which helps us understand the use of the outer charge channel portion of the LT requirement. It also suggests we can also look at the observables we have in data, in particular the measure of the amount of energy in the outer charge channels, shown in the bottom of same figure. We quickly note that many of the events mismeasured in qi1 (top region, maroon) have a high energy deposit in qo1, whereas it does not happen for mismeasurement in qi2. This difference confirms expectations from charge propagation, since it is expected that electrons deposit more energy in outer channels because they propagate in the L-valleys. The no-outer deposit threshold for the LT Fiducial selection is at 1 keV, so it will remove most qi1 mismeasurements but not so for qi2. We note that the propagation differences between holes and electrons produce similar effects as in Figure 9.5, but because this time we look at cases of ‘some’ deposition and since electrons travel in the wider-cone (L-valleys) mismeasurement in qi1 occurs not only near the bottom surface but also when near the sidewall. Since some of the effects are thus explained by side-wall mismeasurement and are effectively removed by the LT no-outer constraint we will impose this criteria and look at the remaining events next.

When the events are partially collected by the interleaved potential, which happens close to the surfaces but the effect smears into the detector, energy is partially measured at the opposite side (explaining the top and right selected regions). In addition, due to the asymmetry of the charges propagation, this effect is different on each side, of particular note is the side-wall events in the qi1 mismeasurement region plot (bottom left in the figure) also deposit large amounts of energy in the outer channel making the no-outer requirement effective for removing them. This is not the case for side2 because holes travel ballistically and do not deposit large amounts of energy in qo2.

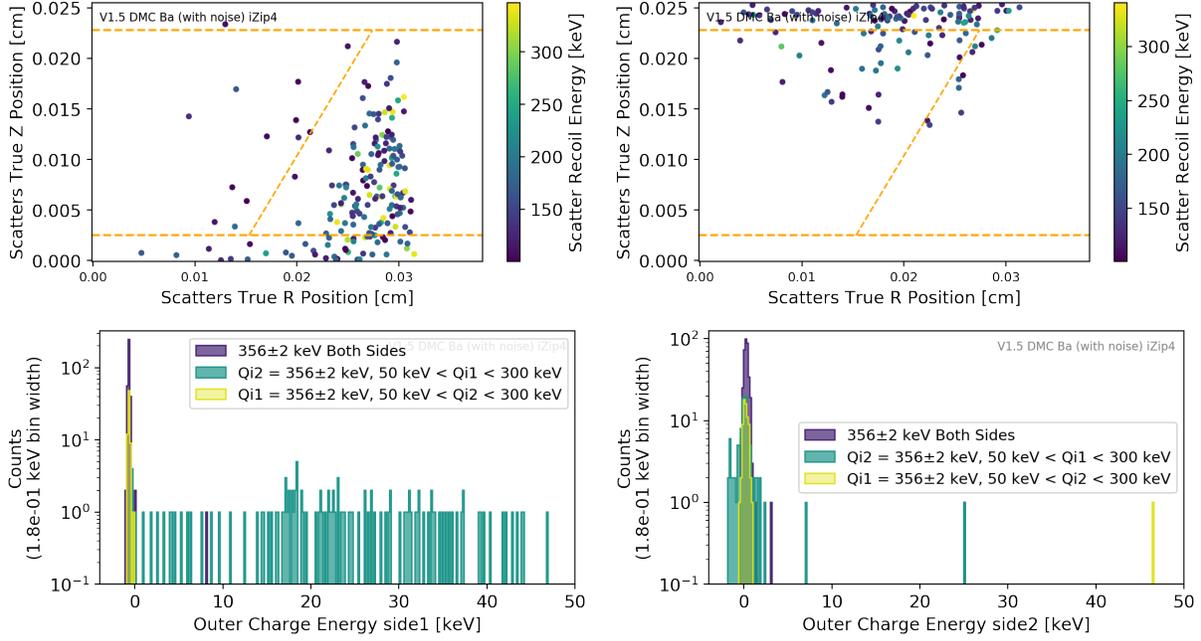


Figure 9.6: Top: The true position of events that are poorly measured in one side while well measured in the other one. The events are selected by requiring one side to be within 2 keV from the 356 keV peak, while the other side to be greater than 50 keV and lower than 250 keV. The points show the scatters with more than 100 keV energy from such events. This figures show that at least one of the scatters in the event deposits a substantial amount of energy outside SimFiducial. Bottom: Outer channel energy measurements of side1 and side2, for 3 samples of 356 keV events: well-measured (356 ± 2 keV on both sides), poorly measured in qi1 (while well-measured in qi2), and poorly measured in qi2 (while well-measured in qi1). The peak at zero is when things are well measured, showing the utility of requiring no energy in the outer channel, and the large number of events with bad mismeasurement on side1 when we have large amounts of energy in the outer channel shows that this requirement removes badly mismeasured events. We note that while this works well for side1, it does not for side2. The reason for this difference is that electrons propagate in the L-valleys with a wider normal angle, thus reaching higher radii and depositing energy in the outer channels.

9.4.4 Results After the No-Outer Energy Requirement and Before the Symmetric Constraint

We next compare the qi1 and qi2 measurement for events that pass the requirement that there be no energy measured in either outer channel. Figure 9.7 shows that in both samples linear correlation region is much clearer and that the data still has much wider variations in both sides. Due to the poor resolution in data it is hard to tell if there is a qi1 versus qi2

mismeasurement in the same way as there is for simulations. As we expected from Figure 9.6, the no-outer deposit constraint removed most mismeasurements of qi1 for the 356 keV peak events, while several are still present for qi2. We note that it is hard to tell if this effect is present in data due to the much wider resolution. We also note the presence of a secondary band at 45 degrees but with a 100 keV offset in side1, we will not study these events further because they are beyond the symmetric requirement but we note them for future studies.

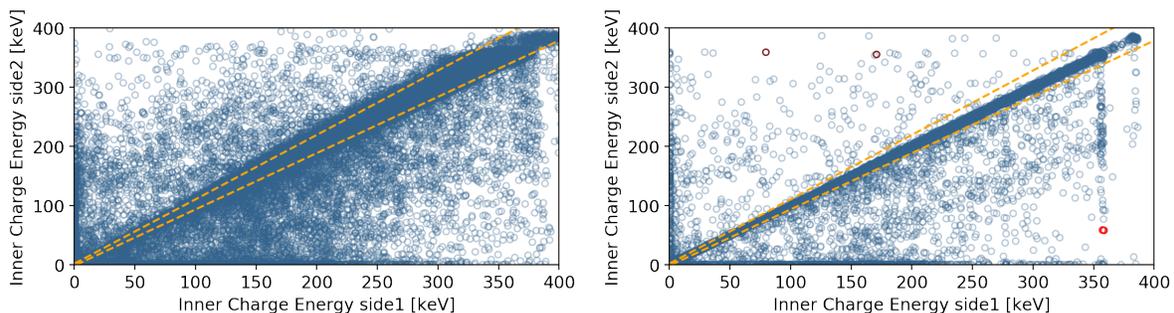


Figure 9.7: Plots of the qi2 energy measurement as function of the qi1 energy measurement (data on the left and simulations on the right) after the requirement that there be no energy measured in the outer charge channels. The orange dashed line shows the region selected for defining the LT Fiducial (symmetric measurement) criteria. While both simulations and data show that both sides are linearly correlated it is clear that data has much worse resolution on both sides, which makes it hard to trust if the mismeasurement effect really exists in data. Indeed, the LT Fiducial constraint (the symmetry selection) removes these outliers in our sample (and data). The outliers in the simulation sample, in particular events with a side1 energy that appears to be fully measured (356 keV) but much less energy in side2 are due to some (or all) scatters depositing energy outside SimFiducial. We select 4 events to inspect closely in Figure 9.8.

Most of the mismeasured events would be removed by the symmetry region (in dashed orange), but it is instructive to understand more about a few that are not. For that, we select the 2 worst qi1 and qi2 events well-measured in one side 356 ± 2 keV, and poorly measured in the other (but with a measurement above 50 keV), these 4 events are marked in maroon and red in Figure 9.7. Figure 9.8 shows these events, the 2 worst from qi1 in the top, and the 2 worst from qi2 in the bottom, and confirms that the scatters energy deposit

is outside SimFiducial. We can easily extrapolate and say that the rest of the mismeasured events is due to some scatters being outside SimFiducial while some inside (contributing to a better measurement), until no scatters are outside SimFiducial and we get well-measured events.

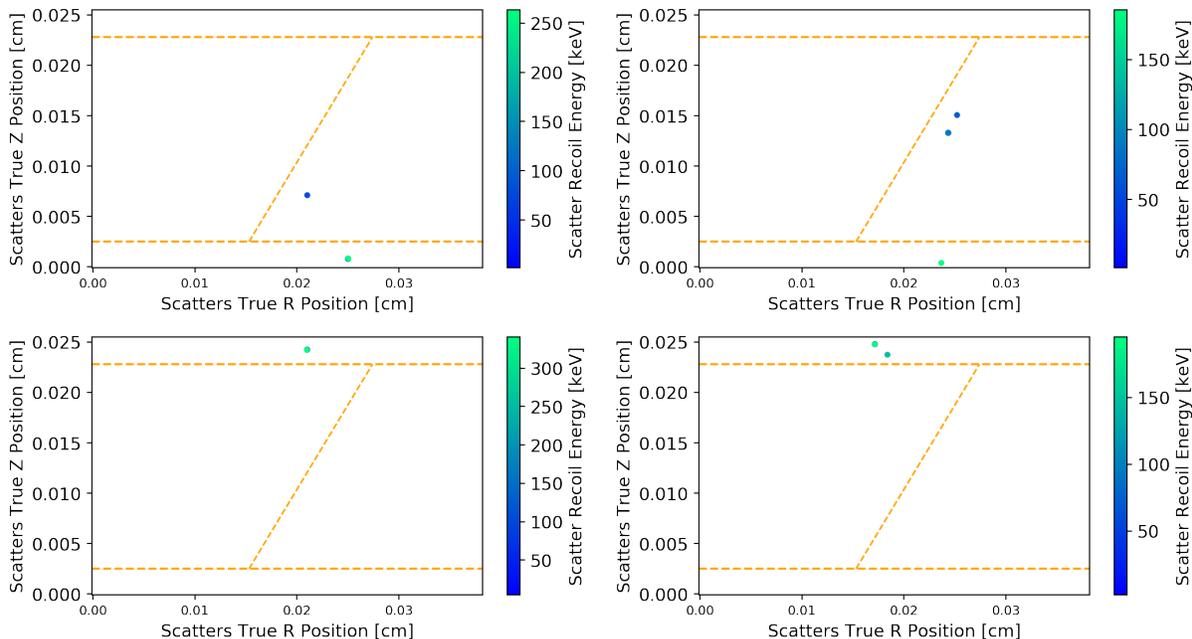


Figure 9.8: The true energy deposition physical location in the detector (Z vs. R) for all scatters of the 2 most poorly measured events on each side. The sample is selected from the 356 ± 2 keV events in one side and the lowest measurement above 50 keV in the other) after the requirement of no energy in the outer channels, the 2 events with mismeasurement in maroon (qi1) and red (qi2) in Figure 9.7. In all cases there is a scatter with a position that is close to the interleaved potential of the opposite side, preventing charges from reaching their collection side.

We have seen from the 356 keV peak, with the aid of our simulations that the dominant cause of mismeasurement is indeed position dependence; scatters depositing energy outside SimFiducial. We will provide quantitative details when comparing the calibration results with the 356 keV peak observations but before doing that, with our LT Fiducial event selection in hand, we will use the simulation to describe how well the LT Fiducial performs

in removing mismeasured events and where it fails next.

9.4.5 What the Simulation Tells Us About Poorly Measured Events that Pass the LT Fiducial Requirements

While the simulation has a resolution that is better than the data, we can use the events which pass the LT Fiducial requirements, but are still mismeasured for study. We have already seen the explanation for mismeasurement of events rejected by LT Fiducial, but we would still want to know how well it does for the events that do pass the criteria. We will see that it is the same effect from position of the scatters in multi-scatter events. First we compare the mismeasurement of a sample of 356 keV events with and without the selection requirements. Next we show the single and multiple scatters mismeasurement differences.

We compare the mismeasurement of all events with a true 356 keV of total recoil energy deposited with and without the LT Fiducial criteria, to show that LT Fiducial does a good job at removing most mismeasurement. We begin by selecting the subsample of events with the photon energy of 356 keV (selected from true recoil energy: $E_{true} = 356.013 \pm 0.100$ keV). The ΔE distribution, where we plot qimean, is shown both with and without the LT Fiducial requirements in Figure 9.9. As expected, LT Fiducial does a good job at removing most of the mismeasured events, but does not remove all of them. Selecting the 4 worst events from the LT Fiducial distribution, and showing the true locations and energies of the scatters in the detector in Figure 9.10, we can see that all of them have most of their scatters near the detector edges but pass the symmetry requirement since they are border-line events. This is interesting because it tells us about events where both sides are undermeasured in a way that makes them both mismeasured by the same amount.

Since the presence of multiple scatters affects the measurement of the energy we compare the resolution for events with multiple scatters and when there is only a single scatter. Figure 9.11 shows the ΔE measurement for the full ^{133}Ba simulation sample of events which pass the LT fiducial requirements and have an energy of 3 keV or more deposited in the crystal. We note that the resolution appears to be roughly the same, the RMS of single

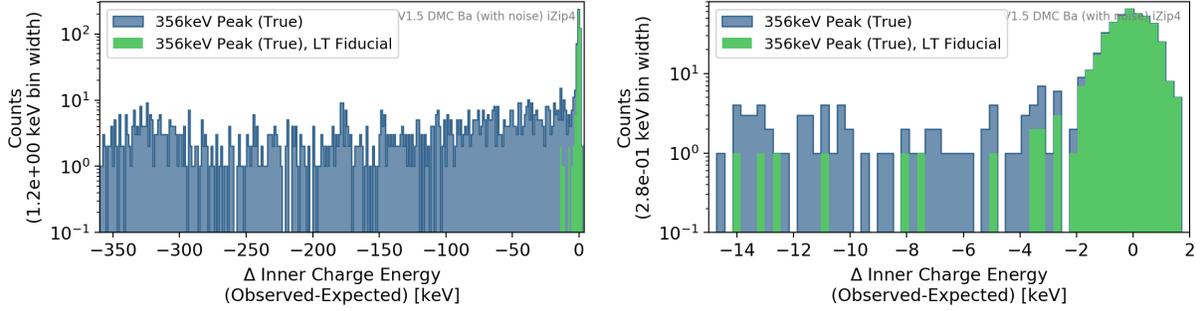


Figure 9.9: The ΔE measurement which compares the true and observed qimean from fully simulated ^{133}Ba events where we only consider input photons with a true energy $E_{true} = 356.013 \pm 0.100$ keV. The events in green, selected as having passed the LT Fiducial requirements, show that LT Fiducial removes most outliers, none below $\Delta E = -14$ keV, but still has a few poorly measured events. The right hand side shows a zoomed-in version of the plot on the left so we can see the most poorly measured events which pass the LT Fiducial requirement. The four events with $\Delta E < -10$ keV are shown in Figure 9.10 where we will see that the mismeasurement is due to at least one of the scatters being close to the edges of the detector.

scatters is smaller than for multiple scatters, and the multiple scatter events have a larger fraction of mismeasurement. While this is a cause of mismeasurement it is not an important effect for WIMP searches as WIMP are expected to be single scatter events only. On the other hand, as we will see in upcoming sections, this effect matters for the calibrations and overall detector response.

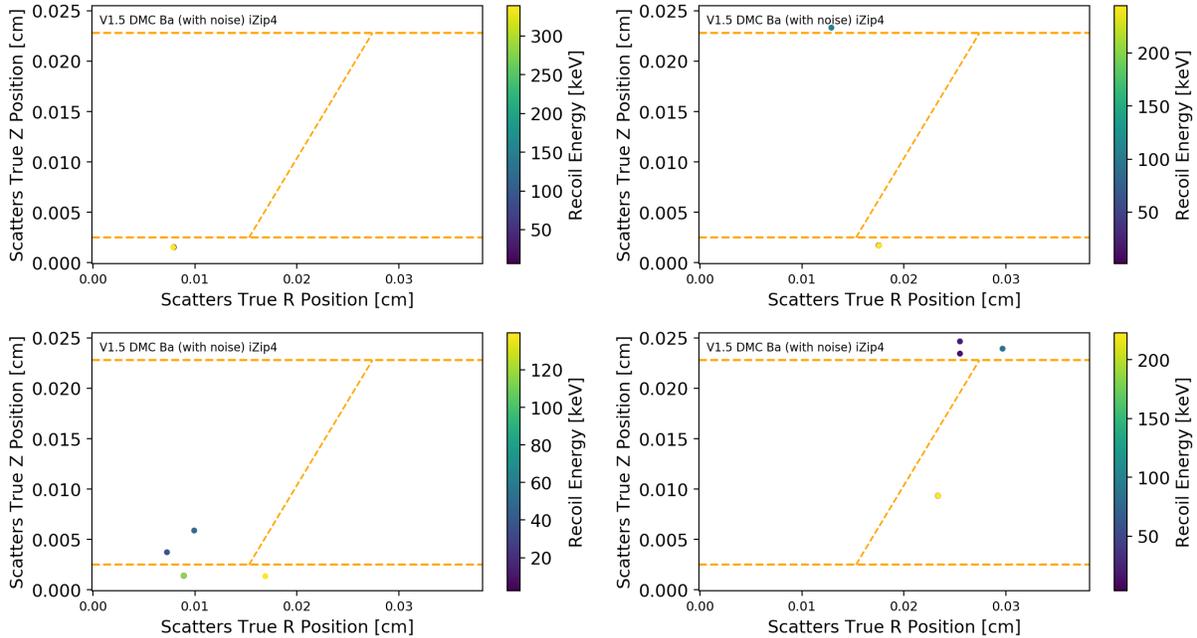


Figure 9.10: This set of plots shows the Z vs. R deposition positions in the detector for the four simulated events from our 356 keV photon with the worse ΔE measurement, but pass the LT Fiducial requirements. All scatters for each event are shown showing that all (or all of the high energy ones) are located near the edges of the detector. Compare with the simulation-based fiducial volume (dashed orange), defined in Section 9.3.

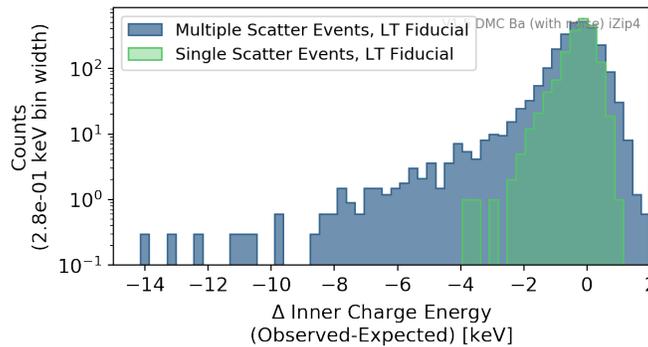


Figure 9.11: Histogram of difference between observed and expected energy for the simulated ^{133}Ba events that pass the LT Fiducial requirements, above 3 keV, single scatter events, compared with multiple scatter events (multiple scatter events histogram normalized to the single scatter events). Multiple scatters are affected by position dependence because while the events can pass the LT Fiducial requirement some of their scatters can still be near the edges accounting for energy loss.

9.5 Comparison of Simulated and Real ^{133}Ba Data

We now study our ^{133}Ba sample and compare the distributions of q_{i1} , q_{i2} and $q_{i\text{mean}}$ to simulations after requiring the LT Fiducial selection criteria. As we will see, this is a complicated sample both due to multiple scatter effects as well as due to the fact that simulations don't reproduce the expected resolution. We will show that both samples are well calibrated and we can compare the resolution for the two samples after the LT Fiducial requirements.

We will begin with Section 9.5.1 where we show a comparison of simulation and data for events in our 356 keV peak, we will note that there is a discrepancy with data (the resolution in the simulations is much better by a factor of $\times 6$ for side1 to $\times 16$ for side2). Finally, in Section 9.5.2 we do a qualitative comparison of the full energy range for both real data and simulations with the ^{133}Ba calibration data.

9.5.1 Comparison of Simulation and Data for Events in the 356 keV Reconstructed Peak

Having understood some of the effects and limitations of LT Fiducial, we next compare the results of the simulated and real data versions of the ^{133}Ba sample near the 356 keV peak for LT Fiducial events. The results are shown in Figure 9.5.1, and summarized in Tables 9.2 and 9.3. First we will show that both samples are well calibrated, and then we comment on the differences between the two. We will compare quantitatively with data (measuring the resolution) to point out that upgrades to the simulation or better understanding is still needed to understand the difference of the resolution of both, and we will understand the effects of noise and the selection criteria in preparation for our accuracy and resolution results in the later sections.

Figure 9.12 shows a zoomed-in region of the full ^{133}Ba sample after LT requirements for both side1 (left) and side2 (right) for three different samples: simulations without noise, simulations with noise, and real data. The first thing we notice is that both sides inner

channels are well calibrated within error, showing the fit-peak location (maximum value of the fitted Skewed-Gaussian distribution) at 356.013 keV (within uncertainty), both for data as well as simulations and regardless of noise addition or selection criteria. This is useful as the simulations samples calibration was performed on the sample with noise. The means and their uncertainties are given in Table 9.2. In addition, we report the values where we have removed the LT fiducial requirement. A comparison shows that the noise effect is negligible for these (higher) energies, and that the LT Fiducial requirement has a small effect in simulations because peak events are only there if they were well measured to begin with (recall the peak is originally a delta function in truth value), while in data there is a significant difference in the LT Fiducial selection for the data q11 measurement (mostly because it reduces the RMS as will be shown next).

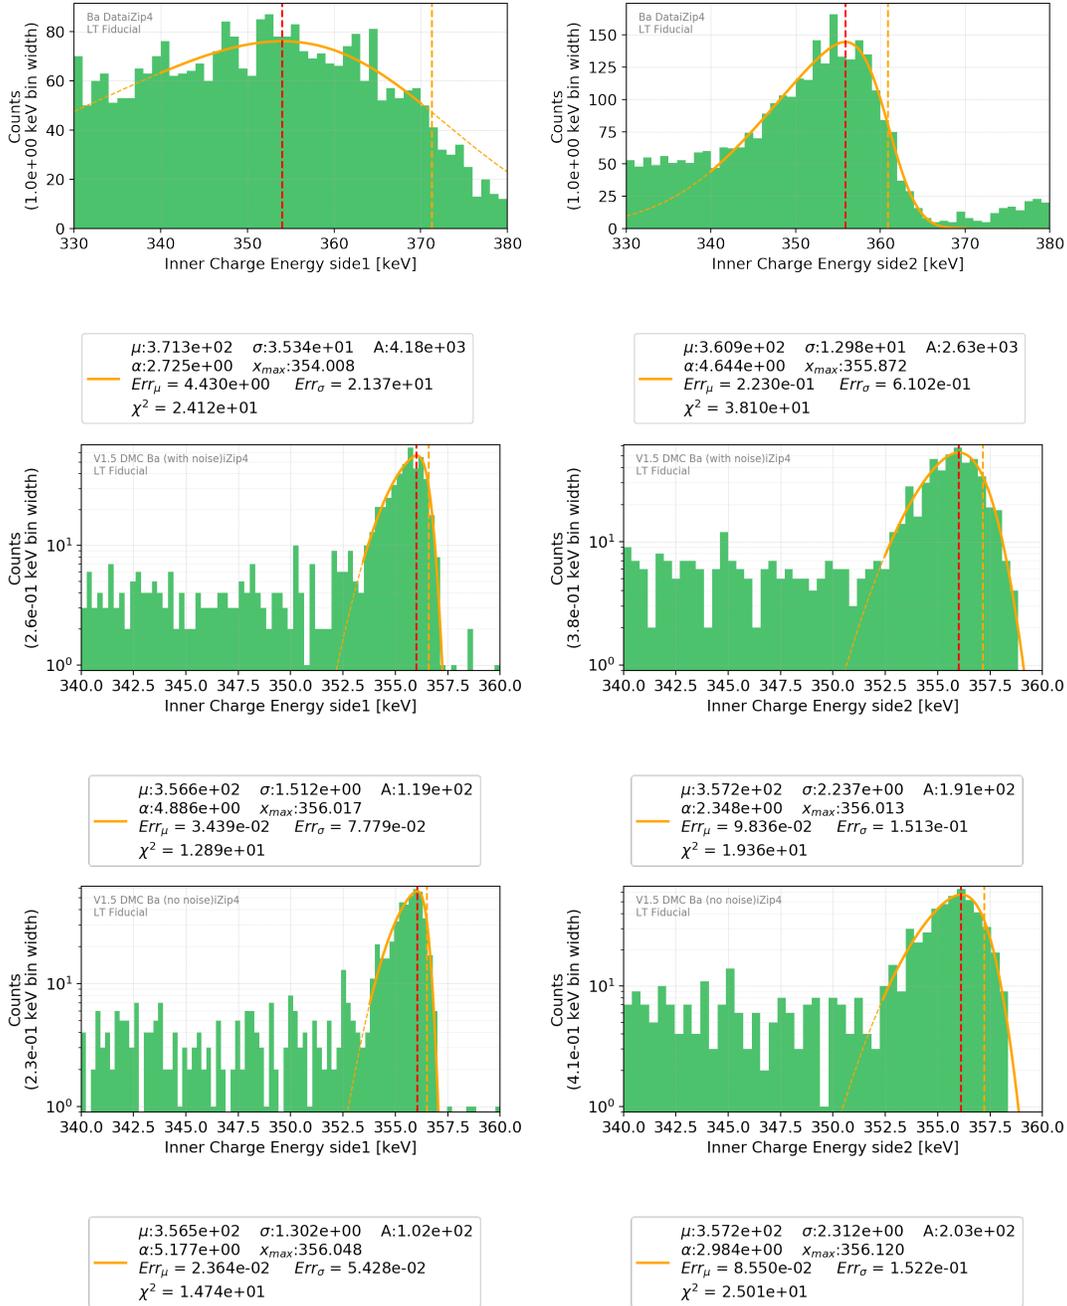


Figure 9.12: The qi1 and qi2 energy distribution of the sample with LT Fiducial selection where we show only the data around the 356 keV peak, from simulation and data ^{133}Ba . We show results for qi1 (left plots), and qi2 (right plots), for the ^{133}Ba calibration data results (top), simulations with noise (middle) and without noise (bottom). The fits are done with a skewed Gaussian function on $\mu - 2\sigma$ to $\mu + 4\sigma$. The samples are well calibrated since the fit maximum value location (x_{max}) is at 356.013 keV within uncertainty, regardless of noise addition in the simulations, telling us that noise resolution is not a dominant issue at high energies. We note that the resolution on data is much worse than it is in our simulations, in particular qi1.

The most striking difference in these plots is that in the data the RMS is much larger than in the simulations, and `side1` is very large compared to `side2`, by a factor of $\times 16$ for `qi1` and $\times 6$ for `qi2`. The full set of values are given in Table 9.3. The RMS at these high energies does not change much with or without noise, indicating that the intrinsic detector resolution dominates at high energies (more details in Chapter 10.3). While we also note a large discrepancy in the observed resolution in data, we will not study it further because it is likely due to tuning effects and possibly also related with corrections in real data. It also suggests that since the resolutions are so different, a weighted version of `qimean` may be appropriate. On the other hand, for high energy studies usually the `qsummax` variable is used for the reasons mentioned in Chapter 4.4, and in more detail at [53, 77].

Mean value of the energy distribution near 356 keV									
		Simulation With Noise				Simulation Without Noise			
Charge	En- ergy Measure- ment Variable	LT	Fiducial	Full	Sample	LT	Fiducial	Full	Sample
		$(x_{max} \pm Err_{\mu})$		$(x_{max} \pm Err_{\mu})$		$(x_{max} \pm Err_{\mu})$		$(x_{max} \pm Err_{\mu})$	
		[keV]		[keV]		[keV]		[keV]	
qi1		356.02 ± .03		355.94 ± .05		356.05 ± .02		356.08 ± .03	
qi2		356.01 ± .10		356.47 ± .06		356.12 ± .09		356.66 ± .05	
qimean		355.70 ± .08		355.71 ± .05		355.76 ± .07		355.75 ± .05	
Real Data									
Charge	En- ergy Measure- ment Variable	LT	Fiducial	Full	Sample				
		$(x_{max} \pm Err_{\mu})$		$(x_{max} \pm Err_{\mu})$					
		[keV]		[keV]					
qi1		354.01 ± 4.43		351.60 ± 24.16					
qi2		355.87 ± .22		356.11 ± .24					
qimean		355.87 ± .46		354.31 ± .48					

Table 9.2: The best fit mean measurements for the ^{133}Ba 356 keV peak locations, from simulations and data as determined for events passing the LT Fiducial requirements. These data are taken from Figure 9.5.1) and the full sample (not shown), comparing with and without noise, and comparing with real data, we also compare the full sample with the LT Fiducial selection criteria. In both data and simulations, both channels are well calibrated, within error the peak location is consistent with the true energy of the incident photon of 356.013 keV. In simulations the answer has little change when removing the LT Fiducial requirements, or adding noise. With respect to LT Fiducial with noise, it changes by $2Err_{\mu}$ for qi1 full sample, by $10Err_{\mu}$ for qi2 full sample. In data the full sample peak is also in the correct place, even though the error in qi1 increases to 24 keV due to the large resolution (which will be compared in the next table). In simulations qimean is always below the expected value by 200 to 300 eV (due to position and energy dependence as will be shown in Section 10.5.3), while this effect is not seen in data it is possible it is obscured by the large resolution.

Resolution of the energy distribution near 356 keV				
LT Fiducial $\sigma \pm Err_{\sigma}$ [keV]				
Charge	En- ergy Measure- ment Variable	Simulation Without Noise	Simulation With Noise	Real Data
qi1		1.30 ± .05	1.51 ± .08	35.34 ± 21.37
qi2		2.31 ± .15	2.24 ± .15	12.98 ± .61
qimean		1.09 ± .10	1.15 ± .10	20.76 ± 1.84

Table 9.3: The best fit RMS measurements for the ^{133}Ba 356 keV peak LT Fiducial samples comparing the simulation without and with noise, as well as with real data. While the noise adds little variation to qi1 and none to qi2, the RMS is much better in the simulation than it is in data (by a factor of $\times 16$ for qi1 and $\times 6$ for qi2).

While additional studies are required to understand the resolution discrepancy with data we end this section with a final remark about the best-energy estimate (qimean) measurements of the 356 keV peak, shown in Figure 9.13. Indeed, as previously shown the resolution in data is much worse than in our simulations, but we note a striking difference about the improvement in resolution with respect to the independent measurements. In our simulations the resolution of qimean is better than both qi1 and qi2, while in data qimean is not better than the qi2 resolution, an indication that work needs to be done to incorporate considerations of the qi1 measurement without hurting the best-energy estimate.

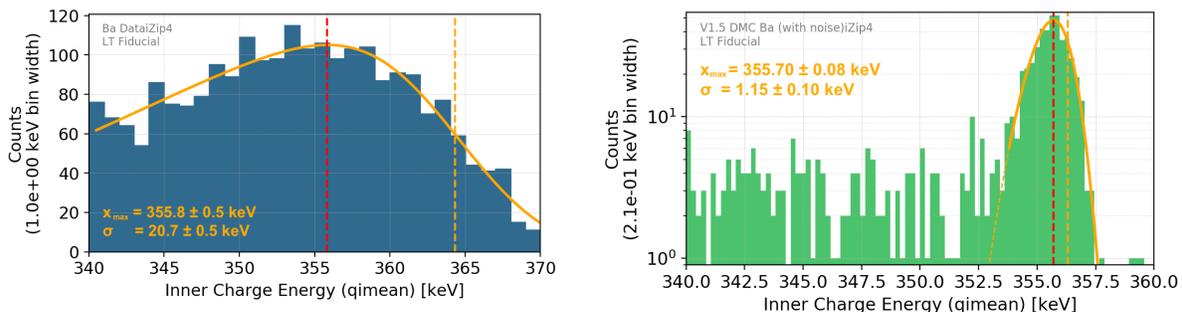


Figure 9.13: The best-estimate energy observable qimean 356 keV peak measurement with events passing the LT Fiducial requirements for data (left) and simulations (right). As previously noticed, the RMS in data is much worse than in our simulations. Of particular note is that while the resolution in qimean improved in simulations, it worsened with respect to qi2 in data due to the large RMS of qi1, an indication that qi1 ought to be treated differently in data, although for large energy searches qsummax is used.

9.5.2 Full Energy Range Comparison of Simulation and Data

Before moving to the accuracy and resolution measurements by comparing the true and reconstructed simulation (in the next chapter) we do a qualitative comparison of the full energy range of our sample. We show the full energy measurement of our ^{133}Ba samples in Figure 9.14, for both simulations and data for the sample of events that pass the LT requirements. Qualitatively we see that the simulation is in agreement with data, the Compton

spectrum is in place, and the peaks appear in the right place. The only effect that is appreciably different is the resolution, which has been quantified already in the previous section.

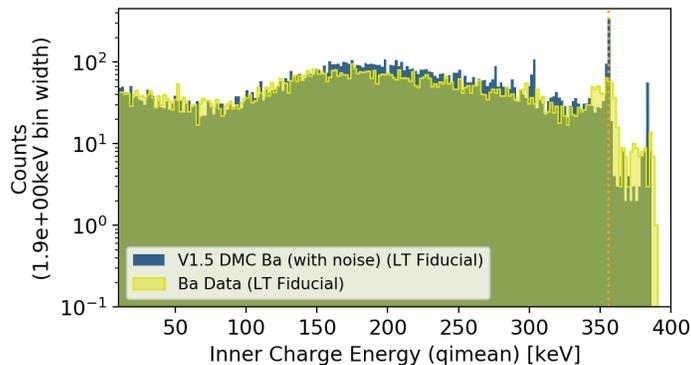


Figure 9.14: The energy distribution of events from the Ba calibration sample from both data and our simulation after the LT Fiducial requirement. We note that the observed energy estimate (q_{imean}) well represents the expected charge energy, and is comparable to data, although the resolution of the Ba peaks is much better than in data.

We now have shown that the sample is well calibrated, that the LT Fiducial cut removes most mismeasured events, and that qualitatively the simulation well-reproduces the data, even though that quantitatively it still needs tuning or we have missing effects in the simulations. We are ready to move on to the predicted accuracy and resolution studies for both the electron and nuclear recoil samples in the next chapter with an eye towards future development for improvements in dark matter searches.

10. CHARGE SIMULATION RESULTS FOR WELL-MEASURED EVENTS, INTRINSIC DETECTOR RESOLUTION, CAUSES OF MISMEASUREMENT, AND NEXT STEPS FOR DEVELOPMENT

Now that we have performed a set of comparisons of the simulations against data, we are set to compare our fully simulated events to input truth. This will constitute our final results from this thesis and focuses on what the simulations tell us about when and how well the detectors measure the deposited energy. We show that the charge response is basically linear, that the charge collection causes the SimFiducial volume, the resolution rises as a function of energy, and that there is already insight to be gained about the limitations and causes of mismeasurement. With these in hand we describe the next steps to work on with an eye to the short term and longer term future as we move towards the conclusions in the next chapter.

In this chapter we complete our studies by comparing the results of the fully simulated events to truth values and look forward to the next steps for development. We begin with Section 10.1 providing an overview of the methods for estimating the predicted accuracy and resolution measurements. Next, in Section 10.2 we show that SimFiducial is due to the charge collection (simulation-based quantity based on Ramo field values) and show that when events deposit their energy in this region the charge energy response is linear. Breaking our study into fiducial and non-fiducial parts we present results for when events are well measured in Section 10.3, where we focus on comparisons of the simulated events to the input energies that created them by measuring the accuracy and resolution of well-behaved events for electron recoil events in the full energy range. We follow with the same study but for nuclear recoil events using the Enhanced WIMP sample in Section 10.4. Then, in Section 10.5 we turn to our poorly measured events to describe what the simulation tells us about three causes of mismeasurement: due to digitization effects at the lowest energies in Section 10.5.1, from events that pass LT Fiducial but are not in the SimFiducial region

as described in Section 10.5.2, and from events within the LT Fiducial region but have systematic mismeasurement from position and recoil energy deposition dependent effects in Section 10.5.3. In Section 10.6 we describe a number of things to be done next for the simulations development to get them ready for use in a dark matter search. Finally, in Section 10.7 we give a summary of what we learned as we move towards our conclusions in the next chapter.

10.1 Overview of the Predicted Accuracy and Resolution Measurement Methods

In order to have a proper understanding of what the simulation tells us we need quantitative measurement methods, and the simulations are well equipped to give us simulation-based truth quantities for that. The goals and expectations have already been shown in Chapter 9.1, our resolution and accuracy methods already outlined in Chapter 9.2, so in this chapter we focus on the understanding from simulations based variables. In addition we will introduce the simulation collected charge energy (at the electrodes) and quantities derived from it. First we will use the simulation collected charge energy to explain SimFiducial and the linearity in the measurements in Section 10.2. Next we use the methods described in Chapter 9.2 to measure the predicted accuracy and resolution as a function of energy, for full energy with multiple-scatter electron recoils in Section 10.3 first, and single-scatter nuclear recoils in Section 10.4 next. Finally we will do studies on highly mismeasured events in Section 10.5, at the lowest energies first, following with comparisons with the LT Fiducial criteria, and finishing with the measurement drop of q_{imean} as function of Z-position and energy. Since most of our methods have already been described, we simply provide a summary of the simulation-based variables first, and expand on our simulated charge collection quantities and how they are useful.

To summarize, we have the following six simulation-based variables that we will use in our studies:

- E_Q : as used in the previous chapter, the total energy available into charge liberation.

This is a truth value given in Eq. 3.11

- Q_{max} : The total charge released in the recoil interaction. This truth value is given in Eq. 3.9, by multiplying $N_{e/h}$ by q (electron charge)
- Q_{Ramo} : This is the total charge absorbed by the sensors after the full crystal propagation simulation. It is given by Eq. 3.13
- E_{Coll} : Set to be $C \cdot Q_{Ramo}$, where the constant C is given by Eq. 3.14. It can be thought of as the simulation truth sensor absorption value and can be compared to after readout simulation and reconstruction
- Collection Efficiency: Set to be Q_{Ramo}/Q_{max} , or equivalently E_{Coll}/E_Q (as is most often presented in this work, because at reconstruction we get energy quantities). It can also be thought of as a simulation truth value to be compared to after readout simulation and reconstruction
- Collection Relative Mismeasurement: Set to be $\Delta E_{Coll}/E_Q$, where $\Delta E_{Coll} = E_{Coll} - E_Q$, as a mismeasurement estimator comparing the collected with the total released charge

In addition to our previously described methods, we now introduce the use of the charge collection quantities, useful for comparisons of energy expectations before the simulation produces the waveforms and the events reconstructed. In other words we can distinguish between propagation and absorption effects from waveform and reconstruction inefficiencies. We will use the charge collection efficiency as described above, E_{Coll}/E_Q , and the charge collection relative mismeasurement $\Delta E_{Coll}/E_Q$, as measures of how much charge propagated and reached the electrodes to contribute to a good measurement. By making similar fits (skewed Gaussian) and looking at largely mismeasured events we will confirm two things with these quantities: that the position dependence is due to charge collection (as opposed to other readout or reconstruction issues), and that the energy response is indeed linear as a function of Q_{Ramo} in the SimFiducial volume.

10.2 Charge Collection Explains SimFiducial Volume and Linearity in Measurements

In this section we will use the charge collection quantities to explain the cause of position dependence of the well-measured region and show the linearity of the detector response. We will use the single-scatter Enhanced WIMP sample for this study since the fiducial volume is clearly defined when no multiple-scatter effects are in play, as has been previously shown. We will first show that the well-measured region is explained by the charge collection, with $\Delta E_{Coll}/E_Q$. Second, with E_{Coll}/E_Q , we will show that the energy response is linear for events in the well-measured (SimFiducial) region.

Figure 10.1 shows $\Delta E_{Coll}/E_Q$ for the full sample of events and for SimFiducial events, and their position dependence. The peak at $\Delta E/E_Q = 0$ in the left hand side plot shows us that the SimFiducial events are indeed well measured (or in this case well collected). In the right hand side of the figure we show the subsample of Bad Charge Collection, SimFiducial events with $E > 3$ keV (to avoid lowest energies issues), where we define Bad Charge Collection as $\Delta E_{Coll}/E_Q < -0.02$. The cause of mismeasurement is the charge collection (as opposed to waveform or reconstruction issues) because the well-behaved region is the same as SimFiducial (compare with Figure 9.3). Overall this is evidence that the SimFiducial requirement does an effective job of removing charge collection mismeasured events.

Figure 10.2 shows the charge collection efficiency of both inner channels, for both the full sample, and the subsample that pass the SimFiducial requirements. The top figures show that the SimFiducial criteria completely removes the incomplete collection tail. The bottom figures are a zoomed-in region with skewed Gaussian fits for the SimFiducial events. These figures show that events in the SimFiducial region are well measured with a typical charge loss of no less than 96%, and that the charge collection efficiency follows a skewed Gaussian distribution with a very narrow resolution. From the fits we know that the charge collection for all SimFiducial events is within 96% and 99%, with a peak value of 98.5% for both sides, an RMS of 0.3% for qi1 and 0.6% for qi2, and a left-hand-side tail. The distribution is not

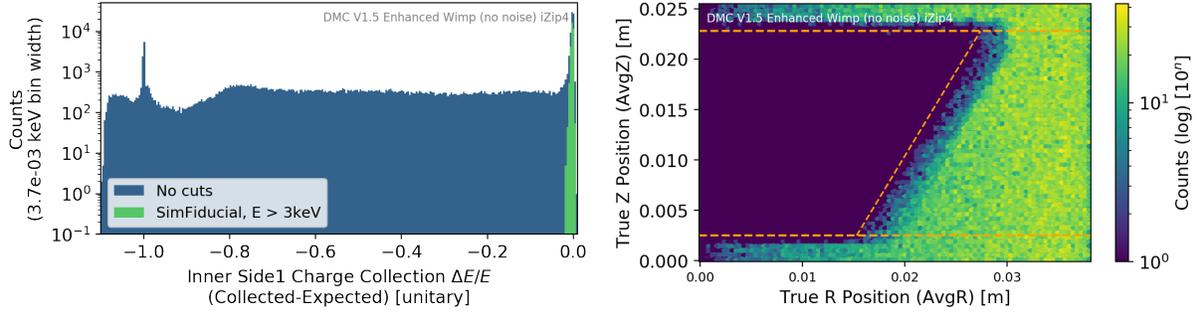


Figure 10.1: This figure shows the energy, as measured by charge collection in two different ways. On the left is the SimFiducial $\Delta E_{Coll}/E_Q$ (side1), compared with the full sample. Note that there are no events that pass the SimFiducial requirements that go beyond -0.02 . On the right is the position dependence of the sample of events with poor charge collection on either side ($\Delta E_{coll}/E_Q < -0.02$, $E > 3$ keV), which shows that the cause of the mismeasurement outside the SimFiducial is poor charge collection.

perfect Gaussian due to charge trapping effects (and that the mean value not being 100% is not a problem as we simply calibrate this out on average), although it is something we still need to follow up on, being a clear next step to work on. Perhaps more interesting is that these distributions have small tails, but on the contrary that the RMS is too narrow (as data resolution is much worse, as shown in Figure 9.12). For now, we can assert that the peak value is in good agreement with the absorption condition of 2% difference between the electrode and the final position (described in Chapter 3.3) so we trust that this result is proof of the simulations performing as expected. Since the distributions are single-peaked and while not pure-Gaussian have such narrow resolutions they also help show that the response is linear inside SimFiducial.

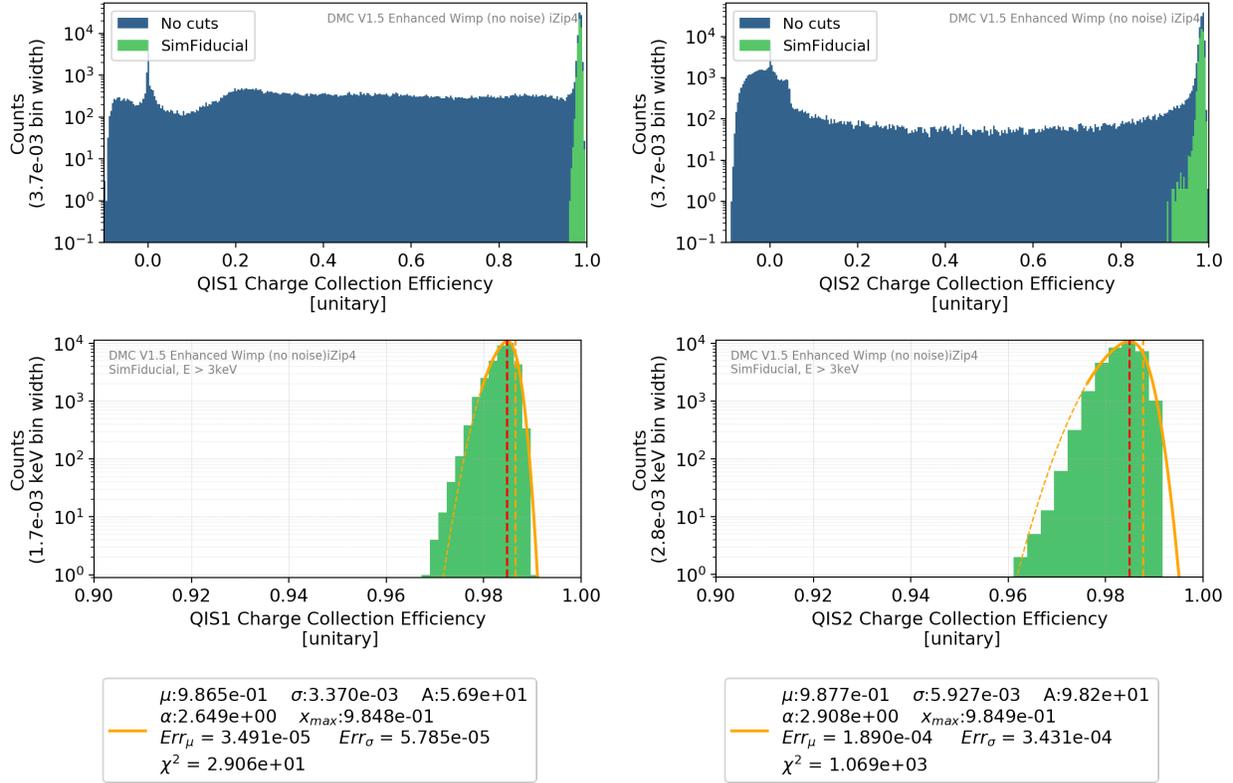


Figure 10.2: Collected charge energy by the electrodes, divided by the expected (maximum) charge energy (we call this charge collection efficiency) in qi1 and qi2, for the full Enhanced WIMP sample for the full set of events as well as the subset which pass the SimFiducial requirements. The top plots show the distributions for the full sample and the SimFiducial subsample. The bottom plots show skewed Gaussian fits for the SimFiducial events. The plots show a zoomed-in version for just the subset of events that pass the SimFiducial requirements showing that they remove any case of incomplete collection (below 96%). Furthermore the fits are consistent with the absorption condition of 2% difference between the electrode and the final position (described in Chapter 3.3) because the peak values are exactly at 98.5%, and have small resolution (0.3% for qi1 and 0.6% for qi2). It is expected that the left-hand-side tail is caused by charge trapping effects.

10.3 Predicted Accuracy and Resolution of Electron Recoil Events

With the simulations calibrated in the same way as data and our measurement methods in place, we now show the accuracy as function of energy of the ¹³³Ba sample after the LT requirements both with and without noise. By using the truth information with and without noise we will be able to differentiate the effects from noise and intrinsic detector

resolution. First we will describe how we are separating our sample by true energy bins to make fits of the mismeasurement and measure the accuracy and resolution as function of energy. Next, we show once more that the response is linear, and that the sample is well calibrated in all energies. Finally we show that noise resolution dominates at low energies but intrinsic detector resolution takes over at high energies, since it also rises linearly with energy. With these results in hand we will do the same for our signal-like WIMP sample in the next section.

To observe the energy dependence of the accuracy and resolution we sub-divide our sample by true energy and then make fits to each. For each part of the energy regime we bin the data and make each bin, denoted by the i -th index with $E_{r,i} \pm 10$ keV, for $E_{r,i}$ from 20 to 380 keV. Next, we make Skewed Gaussian fits to each binned sample, using the same procedure as described above (Figure 9.12). As previously shown, from the skewed Gaussian we use the location of the peak (maximum value) to measure the accuracy and the RMS to measure the resolution.

The accuracy as function of energy for the Barium sample is shown in Figure 10.3 (we will later show that at lowest energies, below 3 keV, there is a systematic mismeasurement which is not noticeable in this scale). From the fact that the mean value of ΔE is zero we see that both sides are well calibrated in the full energy range with linear response of both channels. We note that the accuracy of side2 (holes) while close to zero, has structure going below zero before 250 keV and above zero for higher energies for reasons we don't yet understand. On the other hand they are so small that we are not worried about them. In all cases the sample with noise gives the same answer as the sample without noise. We note that the sample behaves as expected, with linear response of both channels. A second important caveat is that the accuracy of the qimean energy measurement drops with energy at a rate of 0.15%. While this is a small effect with respect to variations, we will show in Section 10.5.3 that this is a function of true Z-position and true energy.

The measurements of the RMS of the sample with and without noise give the intrinsic

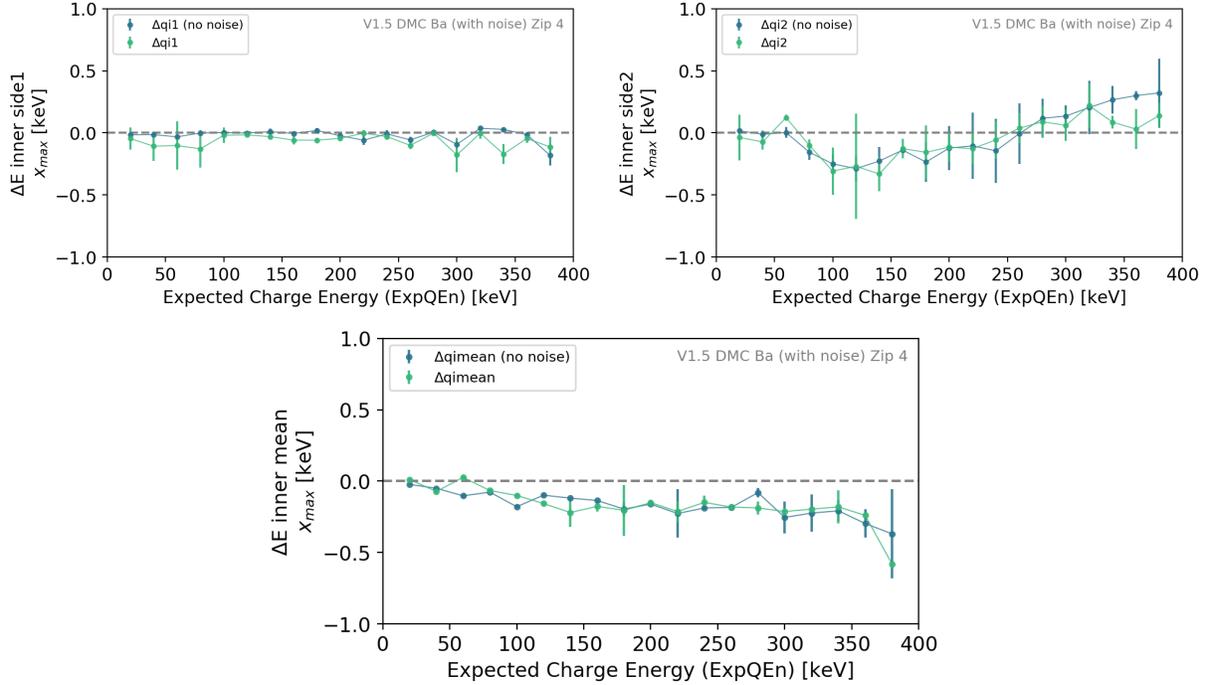


Figure 10.3: The accuracy of the electron recoil energy measurement as determined from our simulated sample of ^{133}Ba events that pass the LT Fiducial requirements. The observed minus expected mean value as function of energy is shown for qi1 (top left), qi2 (top right), and qimean (bottom), with noise (green) and without noise (blue). Each of the results are basically well behaved, linear in energy, and centered at zero. The expectation of a single constant calibration works as expected for the two separate sides. A caveat is that qimean drops at a rate of 0.15% (very small effect with respect to variations and other mismeasurement reasons), which is a function of true Z-position and true energy, described in Section 10.5.3.

detector resolution as well as the combined detector and noise resolution. The results are shown in Figure 10.4, along with the baseline noise resolution for comparison taken from our corrected IRR resolution measurement (Chapter 8.3). Noise RMS dominates at low energies (below ~ 120 keV), but the detector resolution rises roughly linearly with energy and becomes more important at higher energies. From Chapter 8.3 we determined that pure-noise RMS is 250 keV for qimean. Here we see that the intrinsic detector resolution (from no noise qimean) increases at a rate of $0.003 \text{ keV} \cdot E_Q$. The total resolution is the addition of both in quadrature. While this is encouraging, we note that the resolution does not reproduce what is observed in data (see Figure 9.12).

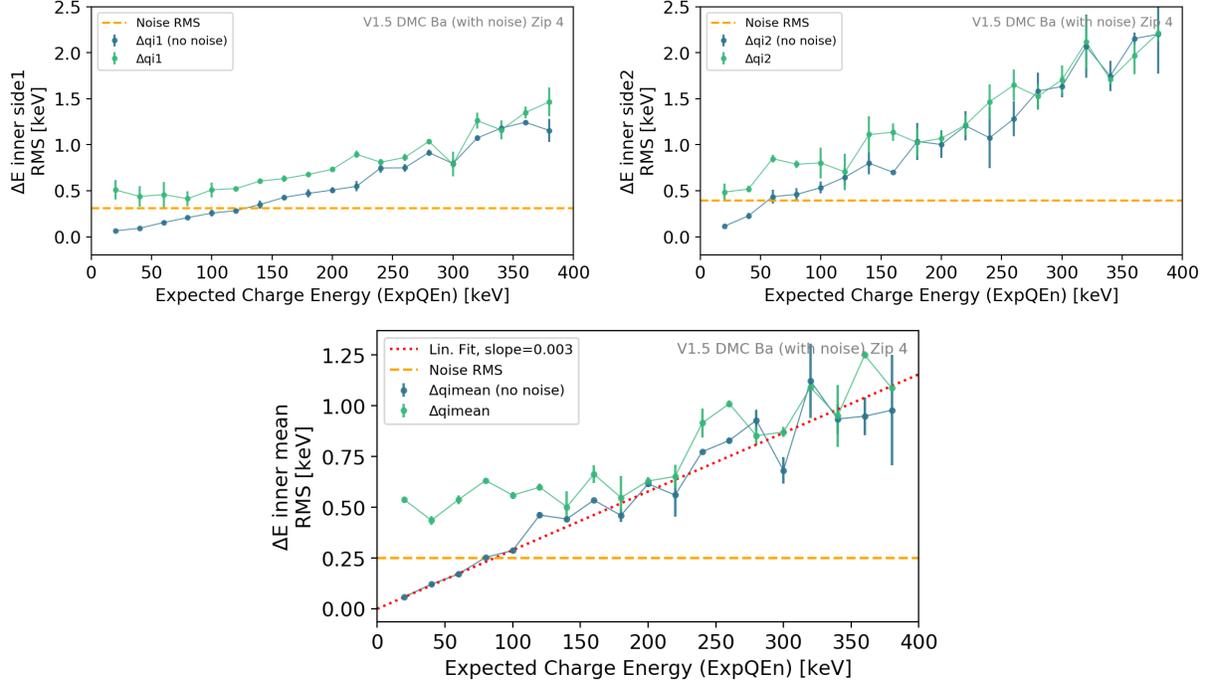


Figure 10.4: The resolution of the electron recoil energy measurement as determined from our simulated sample of ^{133}Ba events that pass the LT Fiducial requirements. The RMS as function of energy is shown for qi1 (top left), qi2 (top right), and qimean (bottom), with noise (green) and without noise (blue), and the simulated noise RMS (dashed orange). Noise dominates below 120 keV, but intrinsic detector resolution rises linearly with energy, at the rate of $0.003 \text{ keV}\cdot E$ (for qimean).

10.4 Predicted Accuracy and Resolution of Nuclear Recoil Events

With our results for electron recoils in hand we move to a study of nuclear recoils, and point out two important effects at the lowest energies. We will replicate the same methodology as in the previous sample but with our signal-like, single-scatter, Enhanced WIMP sample. This energy region is where next generation analyses will be focusing on. In this case, a skewed Gaussian fit is not necessary because the tails are much lower (since the sample is single-scatter), so restricting a simple Gaussian fit to $\mu - 2\sigma$ gives a good measurement of the mean and the RMS. We will now show the results and move on to using the simulations for describing causes of mismeasurement in the next section.

The results are shown in Figure 10.5. On the left we see that the response behaves

as expected, is still linear, and well calibrated, and will follow up in details about the mismeasurement (particularly at the lowest energies) in the next section. The resolution (right of same figure) also behaves as expected, with noise dominating over the full range of this sample, while the detector intrinsic resolution rises linearly (right of same figure).

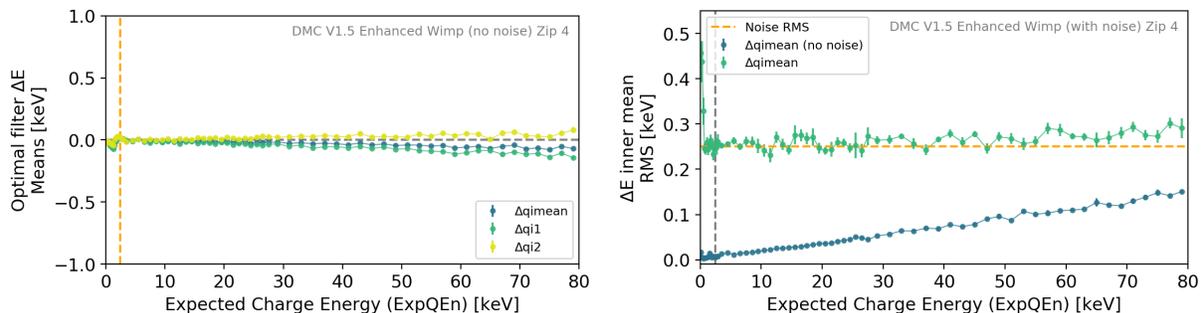


Figure 10.5: Two figures showing the detector response (accuracy and resolution) to single scatter nuclear recoil events from the Enhanced WIMP sample. Left: Single scatter low energy measurements are well behaved, linear and centered at zero, except for lowest energies (< 2.5 keV). Right: RMS behaves as expected, linear response without noise (blue), but dominated by noise (green). Noise causes large variations when energy goes to zero because of pure noise measurements, as described in Section 4.4.

Perhaps the most interesting and important effects are at lowest energies (below 2.5 keV) where we are most sensitive to WIMPs. The first is that the resolution doubles in the sample with noise, and the second that there is an energy dip at roughly 1.4 keV. The resolution doubling is due to the known effects of the optimal filter in pure-noise measurements as shown in Figure 4.4. The second effect, the measurement dip will be inspected closely in the next section and will be shown that it is due to limitations in the waveform digitization.

10.5 Using the Simulation to Study Causes of Energy Mismeasurement

Now that we have shown that the detector responds well when events are in the well-measured portion of the detector we are ready to look in more detail at some of the more interesting sources of mismeasurements that we haven't yet described (and were unexpected).

In this section we describe three sources of mismeasurement from the simulation's insight. They are:

- Digitization effects at low energy (Section 10.5.1)
- While the LT fiducial requirements do a good job of selecting well-measured events, simulations provide more detail about the limitations of this application to SimFiducial (Section 10.5.2)
- The simulation predicts a Z and energy dependent measurement bias even for events that fall in the SimFiducial region (Section 10.5.3)

We describe each in turn.

10.5.1 Lowest Energies Limitations: Noise and Digitization

The first feature we will explain is the mismeasurement of the energy when the recoil energy is below 2.5 keV, as shown in Figure 10.5, and note that the real data searches for WIMPs (described in Chapter 5.5) only started with events with $q_{\text{imean}} > 3$ keV. These are the most interesting energies as they are where the most WIMP events occur (See Figure 1.10). We next show that this effect is due to the optimal filter measurement on events where the minimum digitization amplitude (1 ADC count) is no longer capable of representing the tail of the events. Figure 10.6 shows two zoomed-in versions of the simulation results, with and without noise to see how it affects the measurement. With noise the digitization effect is obscured, and as expected there is a positive bias at lowest energies due to measuring noise (without the 0-delay constraint as shown in Chapter 4.4, and fitting with preference to the right-hand-side of the mean). Without noise (right of figure) all channels show the same shape with the highest mismeasurement at ~ 1.4 keV.

Figure 10.7 shows the charge measurement waveform in ADC counts (equivalent to energy) as a function of time for a low energy event without noise. Superimposed on the event is the the best fit template shape which indicates that the tail is mismeasured by

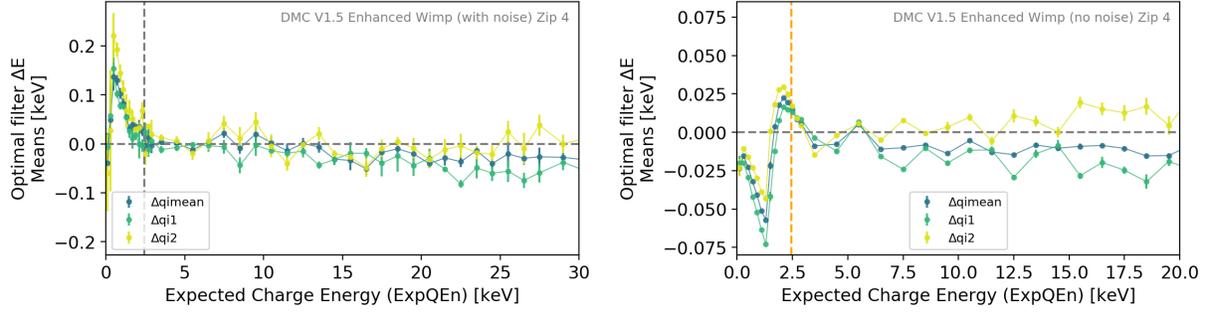


Figure 10.6: Lowest energies show large mismeasurement below $E < 2.5$ keV, both in sample with noise (left) and sample without noise (right). The sample without noise also shows a mismeasurement, an artifact of the pulse digitization.

readout/digitization effects. In particular, we note that within the waveform an important portion of the distribution, in the tail, falls below a single ADC count, which causes a digitization misrepresentation (notice that the tail of the digitized event in blue, is flat-zero, while the template in red is below it). In principle, this issue may be corrected with a template optimized for low energies (one without the tail-feature), but since optimization is not the subject of this thesis we will not pursue it further.

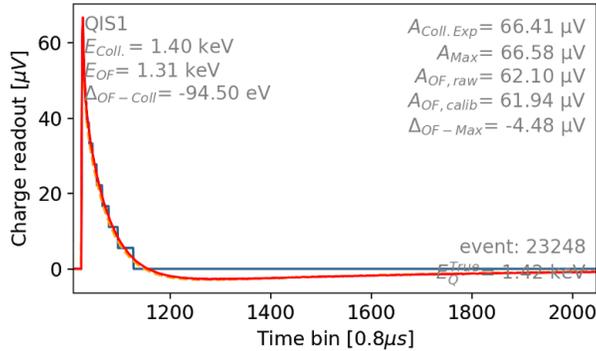


Figure 10.7: This figure shows the charge system readout waveform, in ADC counts as a function of time, for a simulated low-energy event ($E_Q = 1.31$ keV) with no noise to show the problems of the digitization effect on our energy measurement. Superimposed is the best fit template shape. We note the the readout waveform values are so low that the shape from the ADC has been truncated.

10.5.2 LT Fiducial Comparison

While we showed some preliminary results in Chapter 9.4, we now compare how good LT Fiducial is at removing mismeasurement, and at reproducing the SimFiducial event selection. First we observe the effectiveness at removing mismeasured events of LT Fiducial and we will compare it with the mismeasurement from SimFiducial. Next we will show the location of the scatters selected by the LT Fiducial event selection and the location of the mismeasured events that pass the requirements as we compare again with the SimFiducial volume.

Since we know the resolution has a linear energy dependence we use the Enhanced WIMP sample and measure $\Delta E/E_Q$ for both event selections. The plot in Figure 10.8 shows that while LT Fiducial (blue) does a good job at removing the majority of the mismeasured events (see Figure 9.9), there are still many mismeasured events, even some with over 100% mismeasurement which occurs when the scatters with most energy are outside the SimFiducial volume. By construction SimFiducial (yellow) has no events beyond $\Delta E/E_Q < -0.015$.

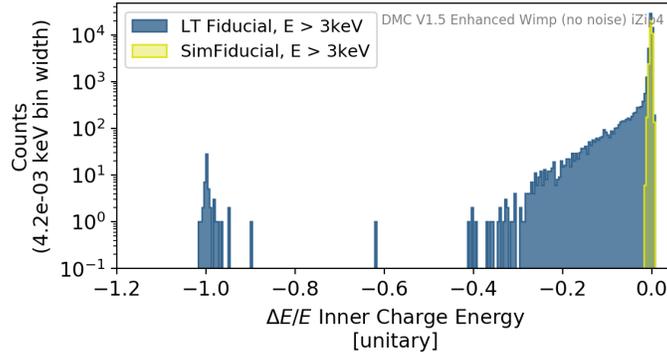


Figure 10.8: The $\Delta E/E_Q$ distribution for the Enhanced WIMP sample for the subsamples of events which pass the LT Fiducial and the SimFiducial requirements. We use samples with $E > 3$ keV to ignore lowest energies limitations. We see that while we already know LT Fiducial does a good job at removing mismeasured events there is a large population of mismeasured events (even with $\Delta E/E_Q$ near -1). SimFiducial does not have this long tail, so making LT Fiducial closer to SimFiducial would represent a major improvement in data analysis.

Now we show the true position of the events that pass the LT Fiducial requirements and compare with the SimFiducial volume. Figure 10.9 (left) shows the true position of events that pass the LT Fiducial requirement with $E > 3$ keV . There is a large population of events which would fail the SimFiducial requirement, as expected. The right hand side of the figure shows the location of the events which pass LT Fiducial but which are mismeasured (as selected with $\Delta E/E_Q < -0.015$). All of the Bad LT Fiducial events are beyond SimFiducial, some even in the edges of the detector. We have shown that the dominant cause of mismeasurement is position dependence of the scatter. This clearly suggests that it may be possible to do a better job at rejecting events beyond SimFiducial with the LT Fiducial criteria.

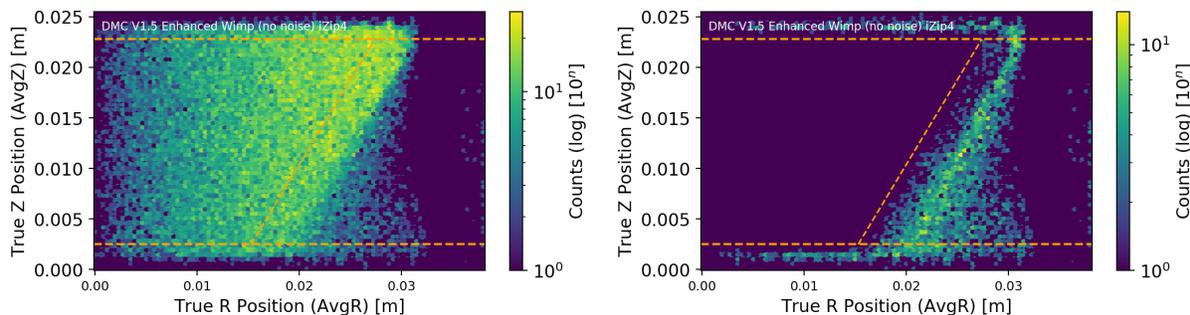


Figure 10.9: Two figures showing the location of energy deposits for the enhanced WIMP sample, but have charge energy mismeasurement. Left: Shows true position of all of the events which pass LT Fiducial requirements. We note that many are outside SimFiducial volume (dashed orange). Right: The true position of events which pass the LT Fiducial requirement but which are poorly measured (selected with $\Delta E/E_Q < -0.015$), their true position is outside SimFiducial.

10.5.3 A Study of the Qimean Measurement Drop as a Function of Energy: Z-Position Dependence

Now that we understand the basic properties of our detectors simulations, we can find the cause of the previously noticed energy drop in qimean as a function of energy as shown

in Figures 10.3 and 10.5. While we do not yet know the full cause of this effect, we point out that the simulation shows that there is a systematic bias in the qimean measurement. This is due to the fact that on each side (qi1 and qi2) there is a bias in the energy measurement as a function of total deposited energy and Z position. Figure 10.10 shows the ΔE for both sides as a function of Z position for a large energy where this effect is pronounced. There is clearly a linear correlation as a function of Z that is partially, but not completely, offsetting. It is possible to do a fit to make corrections in qi1 and qi2 energy measurements and make qimean flat, but it is not clear it is applicable or useful in real data where we do not have the true Z position. We note that this dependence could be useful in analyses for improving a Z-position estimator.

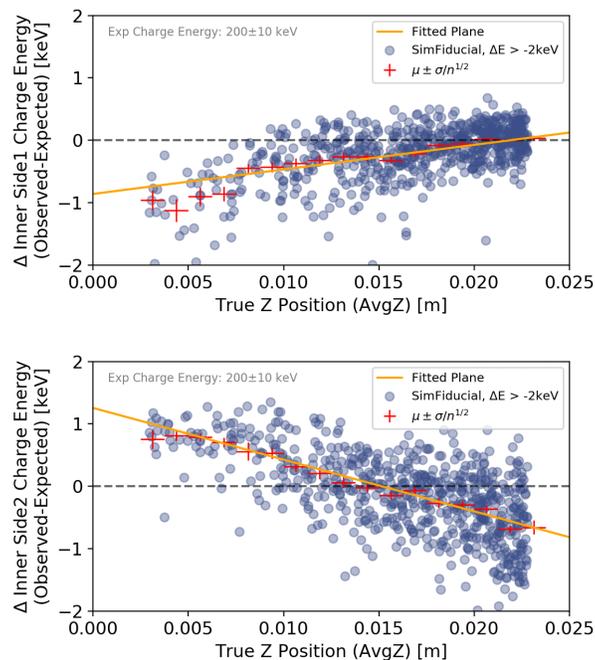


Figure 10.10: This figure shows the amount of energy mismeasurement on both sides of the detector as a function of Z for a large energy sample that pass all the SimFiducial requirements. For both sides there is a linear dependence as a function of Z which partially, but not completely, offsets the full qimean measurement. This explains the qimean measurement as a function of energy shown in Figures 10.3 and 10.5. The qimean energy measurement drops at a rate of 0.15%.

10.6 Next Steps for the Simulations Development

With our studies complete, we turn back to our goals of optimizing our dark matter signal and background model for discrimination in analysis (Chapter 5.6). We next describe some future steps for the simulation development now that our work is complete. They can be thought of as improvements to what we already have, and completing sets of studies that we have begun. The 4 major steps for next developers to do are:

- Tuning or unaccounted effects: Since there are some clear differences between our simulations and the results in data, thus there may be tuning need of the parameters in our detector models, or effects which are simply not modelled and need to be added. These will require follow up on the studies we have done, as well as the next three studies.
- Nuclear recoil studies similar to the ones presented here, but using ^{252}Cf data to compare the detector response to nuclear recoils: Before we can claim that our simulations are ready to be trusted in optimization studies we need to be able to understand and reproduce not only the affects in ^{133}Ba data for electron recoil events, but also with ^{252}Cf data for nuclear recoils.
- Phonon readout simulation studies: While we have done studies with the charge system, we will also need to do a parallel set of studies for the phonon system.
- Simulation of other backgrounds: While we have addressed many of the sources of background in Table 2.1, we will need need to have confidence in the simulations of our other backgrounds. We have already made full simulations of many of them, especially lead, but we have not validated them yet.

We next comment on each in turn.

The difference in the results from data in Chapter 9.5 suggest that there is tuning needed or unaccounted effects that need to be added to our simulations. Before we can declare

victory our simulations need to agree quantitatively with data. Of principle concern is that the resolution in data for the electron readout is significantly worse than the hole readout resolution, as shown in Figure 9.12 and both are different than what is seen in simulations. Setting aside that it may be a measurement problem that can be corrected for in the data or there is a missing effect in the simulation, there are a number of parameters that we expect to control some of the differences. These include trapping parameters, surface reflection and absorption, or additional fluctuation effects in the ionization process (like the Fano factor as described in Ref.[87]). DetectorSim has over 100 parameters to tune, while most of them material specific (and may not need tuning), some are not completely understood. A first guess is charge random-trapping, a known issue in real detectors [66], which we could tune to introduce smearing resolution effects in our simulation. Side-wall trapping is already at maximum, but a study with outer-channels could help us understand if this is an accurate approximation. Ideally additional data taking with the detectors in different configurations might lead the way, but since the experiment has been dismantled we need to think about tuning processes with the existing data. To check them we would make toy samples with highly different trapping parameters first to observe the effects on resolution. We could also think about other calibration samples to determine the intrinsic detector resolution at different energies, for example the Ge activation line (at 10 keV) [51], which would give us another handle for the linearly rising resolution effect. We could also study this using the data taking with the same detectors with a higher bias voltage. Indeed preliminary results from those studies indicate that different calibration constants are needed to properly correct the data.

While we have done measurements and studies for electron recoils from ^{133}Ba , a clear next step is to do the same for nuclear recoils with ^{252}Cf simulations. Intentionally, we have already suggested a third sample: ^{71}Ge activation, which provides a third possibility of comparison of electron recoils with features at lower energies. We have all the tools in hand to study the charge readout for these two samples. So the obvious next step (perhaps in

parallel with tuning) is a full validation of the detector response to both nuclear as well as electron recoil depositions with them. We have initiated the electron recoil validation in the previous chapter with the ^{133}Ba sample, and while we also started a validation for nuclear recoils, we note that all the pieces for comparison to the ^{252}Cf calibration data is in place.

The phonon readout simulation needs to be understood before using the simulations in a dark matter search analysis. Reproducing the same as we have done with the charge readout but for the phonon system is a start. A good validation, after tuning and validating both charge and phonon systems work for both electron and nuclear recoils would be to show that between ^{133}Ba and ^{252}Cf samples we can produce the yield variables as shown in Figure 2.2. With the phonon propagation and readout simulation and the data already in hand we are ready for a study like this.

In addition to validating the calibration samples, we can already think ahead to improve our model of the complicated background sources, for example the neutron background from lead, as described in Chapter 2.1. The tools have already been created and lots of events simulated for many different positions of atoms around the apparatus. With these in hand, and the studies above completed we hope to find new ways of rejecting these events, prior to a sophisticated event selection method, as done previously with a BDT in Ref. [47].

10.7 Summary

With our work complete, and the next steps defined we summarize what we have learned from our simulations. We have now demonstrated that qualitatively the charge energy measurements behave as expected: the energy response is linear, mismeasurement causes are not un-physical, and well-reproduce dominant features from data (like the overall ^{133}Ba spectrum). Quantitatively we have seen that the baseline (pure-noise) resolution dominates at low energies, but the intrinsic detector resolution increases linearly with energy, and even though the resolution is much better than that of data, we show that the simulations are ready to be tuned. We have also given clear steps for what needs to be done next, which range from tuning and validating to understanding the detector response and making corrections

in data. We are well set to do this with our simulations in an iterative process, so we are ready to give final remarks on how the simulations impact an analysis and conclude this thesis in the next chapter.

11. CONCLUSIONS

The search for dark matter in direct detection experiments is a long and difficult process. We have fully described our simulations infrastructure package, our studies that show how far along we are with our simulations, and we have outlined the next steps for simulation development studies. This systematic effort has taken us from modest beginnings to being well on the way to having simulation-based searches which may be what is needed to discover WIMPs from small energy deposition in a detector.

In this chapter we conclude this work providing an overview of using the simulations in future analyses, along with our conclusions and hopes for the future as we look towards upcoming searches, and potentially a major discovery. In Section 11.1 we make comments about how the working simulations could be used to optimize a dark matter search analysis. And in Section 11.2 we summarize the status and what we have learned with this study with an eye towards the future.

11.1 Overview of Optimizing a Dark Matter Search with the Simulations and Clear Examples to Follow Up From This Work

With the full description of our simulations, results, and next steps for development, we have a clear path for how they will benefit future dark matter searches. We first describe how the simulations impact in analyses is an iterative step starting where the previous analysis ended and using them to understand the detectors response. Then we summarize how a dark matter search can be improved, by optimizing the signal to background rejection. We follow with what the fully functional simulations would bring forth with an eye towards optimizing the WIMP searches. Finally we present a set of clear examples to be followed up on as identified by our studies, prior to concluding this work.

Before looking forward, we look back to where we started and combine that with where we are today. In many ways, this simulation work starts where the previous analysis ended. As

discussed in Chapter 5.5 the previous search for WIMPs was very sensitive to low masses, but had a number of events that passed all the selection criteria that were not really understood nor included in the background model during the optimization process. Our primary goal with the simulations is to allow for a full understanding of all the backgrounds so we can use this understanding in an optimized search strategy. We discuss the ideas needed to improve analyses in general with an eye towards simulations development and present an overview of how they will be used in an analysis.

The improvement to an analysis can be thought of as either adding requirements that reduce more background while keeping almost all the signal, or by removing any inefficient requirements to allow a lot more signal and getting rid of the backgrounds that were previously rejected using techniques derived from a deeper understanding. Said differently, we want a better understanding of what the next-dominant problem hurting the signal-to-background rejection really is so we can attack it. A reliable simulation may well lead the way. While the simulations tools today have come a long way, and have taught us a lot, they are clearly not ready for such a task.

The ideal is to have a set of simulations that provide both proper estimates of all the background rates, and what their observables would look like in data and use them to optimize our dark matter searches. If we had such a set of tools we could improve our analyses in two ways: creating new variables and/or correcting old ones for use in analysis, and do a better estimation of the rates for all the backgrounds (including a new ones that need to be added). With both things, simple optimization methods such as determining the expected signal over the background as a function of a cut can be used to optimize the selection criteria. Then again, if we really have variables and background models we trust enough, we could use machine learning techniques or other similar type techniques [47, 88] which require well understood inputs but have been shown to be incredibly powerful when used correctly. These improvements in our analyses decisions are crucial as the experiment moves to sensitivities to lower energy interactions that incorporate more complicated background

sources.

Fully trusting the simulations allows for detector response understanding and energy estimate corrections (as opposed to mismeasurement rejection). We finish this section by giving a few clear examples to be followed up from this work, of how the simulations would help achieve each of these tasks, as we move towards the conclusions:

- Use the BlipFinder to remove mismeasured events from the LT analysis, and the Blip-Corrected PSDs to do better energy measurements of the charge variables. While more understanding of Blips in WIMP-search data is needed this correction might improve our lowest energy measurements
- Develop a template that takes into account ADC truncation effects for low energies to improve the energy measurement
- Use the simulations to understand the cause of resolution smearing in qil data, as shown in Chapter 10.3, and suggest how to improve it
- Improve the LT Fiducial criteria with expectation from the simulation-based fiducial volume (Chapter 9.3)
- Recall from Figure 10.3 that our simulation shows that the accuracy of the measurement oscillates (by up to 0.5 keV) around zero, while we don't fully trust this answer yet and we need to study further, this suggests that our data is miscalibrated (as it appears to be energy dependent), if we corrected this our resolution would improve

11.2 Final Conclusions

With our simulations and results described, and our goals and next steps presented we are ready for our final remarks and conclusions. We first present how the current simulations infrastructure is ready to be tuned to tackle the upcoming challenges and play a key role in the dark matter discovery in next generation experiments. We summarize our accomplishments and results setting the tone for what needs to be studied and developed next for

our simulations to be fully functional. Finally, we conclude our work with emphasis on the bright future that the simulations are leading us to.

The simulations infrastructure has already been used for SuperCDMS Soudan and is ready for use in SuperCDMS SNOLAB. While much of the work to improve the sensitivity for the SuperCDMS SNOLAB program will come from detector upgrades and putting the detectors deeper underground, our work has focused on the long term plan to truly understand the detector response with simulations. While tuning is still needed for Soudan, and will need real data to start the validation process for SNOLAB, we can already see the advantages and improvements it brings forth. Not only because simulation-based analyses in other big search experiments [9, 10] have recently discovered fundamental particles, but we have already shown new understanding we can gain from simulations for our own detectors. We have shown that we can use the simulations to improve our detector response understanding, develop corrections, and improve our event selection criteria, all of which have a direct impact on the sensitivity in a dark matter search as they provide tools to optimize the signal to background discrimination. While we can develop these tools on the current SuperCDMS Soudan data, and aim to improve previous results, we are keeping an eye towards the upcoming data from SuperCDMS SNOLAB, which will require a more robust analysis machinery due to the increased sensitivity to lower energies and higher particle interaction rates. There is much work to be done including new backgrounds, and other models of dark matter which need to be considered.

We have shown a number of accomplishments which tell us that the simulations well-reproduce the expectations, present no known physics and computational issues, and suggest well-motivated causes of mismeasurement, albeit quantitatively they do not agree with all aspects of data. Noise simulations well-reproduce real data noise observations, while charge observations are linear in energy and calibrations work as expected. Current studies show that the intrinsic detector resolution rises linearly with energy and dominates at the highest energies, at the lowest energies, where we have best sensitivity to WIMPS we see that

noise is the dominant issue and we also see a template problem. Our current quantitative comparisons to data show that we have much work to be done because the resolution is far worse in data than predicted by our simulation, but we do not yet know whether this is a simulation problem or poor measurement of the data. It may be that our simulations need tuning, that there may be missing effects in them, or the current energy measurement techniques may just not be good enough to achieve the resolution that is possibly suggested by the simulations. We have shown that we can use the simulations to better understand our event selection criteria, and we already have suggestions in hand that show that current regions of the detector that are poorly measured are still being included in analysis. While we have long known that the main source of mismeasurement is position dependence of interaction in the detector caused by poor charge collection at the electrodes our simulations indicate that there are potential gains from correcting it, like the observed energy (0.15%) effect from Z-position and true energy dependence, or at the lowest energies (below 3 keV) from using new templates to take into account digitization issues of the signal. Currently, the simulation infrastructure is in good shape and ready to be further tuned to reproduce the data as well as otherwise augmented.

We conclude that the future is bright for SuperCDMS in the SNOLAB era as we move to a time when we have simulation-based searches. The next steps for the simulations and their use is now established. This work brings us closer to the day when we can model all the known backgrounds, and makes us well setup for the ones that haven't been important (or we don't know about yet). The plan is clear to use them soon in optimized searches. With our new tools in hand, and a clear path forward, the discovery of dark matter may just be around the corner.

REFERENCES

- [1] I. Newton, *Philosophiae naturalis principia mathematica*. J. Societatis Regiae ac Typis J. Streater, 1687.
- [2] A. Einstein, “The Foundation of the General Theory of Relativity,” *Annalen Phys.*, vol. 49, no. 7, pp. 769–822, 1916.
- [3] F. Zwicky, “Die Rotverschiebung von extragalaktischen Nebeln,” *Helvetica Physica Acta*, vol. 6, 110-127, 1933.
- [4] E. Corbelli and P. Salucci, “The extended rotation curve and the dark matter halo of m33,” *Monthly Notices of the Royal Astronomical Society*, vol. 311, pp. 441–447, Jan 2000.
- [5] M. Markevitch, A. H. Gonzalez, *et al.*, “Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657-56,” *The Astrophysical Journal*, vol. 606, pp. 819–824, May 2004.
- [6] D. J. Eisenstein, I. Zehavi, *et al.*, “Detection of the Baryon Acoustic Peak in the Large-Scale Correlation Function of SDSS Luminous Red Galaxies,” *Astrophys. J.*, vol. 633, pp. 560–574, 2005.
- [7] Planck Collaboration, “Planck 2018 results. VI. Cosmological parameters,” *eprint arXiv:1807.06209*, 2018.
- [8] J. L. Feng, “Dark Matter Candidates from Particle Physics and Methods of Detection,” *Annual Review of Astronomy and Astrophysics*, vol. 48, no. 1, pp. 495–545, 2010.
- [9] F. Abe, H. Akimoto, *et al.*, “Observation of Top Quark Production in pp Collisions with the Collider Detector at Fermilab,” *Physical Review Letters*, vol. 74, pp. 2626–2631, Apr 1995.

- [10] S. Chatrchyan, V. Khachatryan, *et al.*, “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC,” *Physics Letters B*, vol. 716, pp. 30–61, Sep 2012.
- [11] G. P. Kuiper, “The Empirical Mass-Luminosity Relation,” *Astrophysical Journal*, vol. 88, p. 472, Nov 1938.
- [12] Vera C. Rubin and W. Kent Ford, “Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions,” *Astrophysical Journal*, vol. vol. 159, 1970.
- [13] “Coma Cluster.” Web, August 2017. Available at https://www.nasa.gov/multimedia/imagegallery/image_feature_1004.html.
- [14] “Artist’s impression of the expected dark matter distribution around the Milky Way.” Web, April 2012. Available at <https://www.eso.org/public/images/eso1217a/>.
- [15] “Astronomy picture of the day: A Horseshoe Einstein Ring from Hubble.” Web, December 2011. Available at <https://apod.nasa.gov/apod/ap111221.html>.
- [16] T. Treu and R. S. Ellis, “Gravitational Lensing: Einstein’s unfinished symphony,” *Contemp. Phys.*, vol. 56, no. 1, pp. 17–34, 2015.
- [17] Y. D. Hezaveh, N. Dalal, *et al.*, “Detection of lensing substructure using alma observations of the dusty galaxy sdg.81,” *Astrophysical Journal*, vol. vol. 823, 2016.
- [18] D. Toback, *Big Bang, Black Holes, No Math*. Kendall Hunt, 2013.
- [19] “Astronomy picture of the day: The El Gordo Massive Galaxy Cluster.” Web, April 2014. Available at <https://apod.nasa.gov/apod/ap140422.html>.
- [20] “Bullet Cluster: Direct Proof of Dark Matter.” Web, August 2006. Available at https://chandra.harvard.edu/graphics/resources/handouts/lithos/bullet_lithos.pdf.
- [21] E. W. Klob and M. S. Turner, *The Early Universe*. Westview Press, 1994.

- [22] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*. Westview Press, 2016.
- [23] M. Tanabashi, K. Hagiwara, *et al.*, “Review of Particle Physics,” *Phys. Rev.*, vol. D98, no. 3, p. 030001, 2018.
- [24] G. R. Blumenthal, S. M. Faber, *et al.*, “Formation of galaxies and large-scale structure with cold dark matter,” *Nature*, vol. Vol. 311, pp. 517–525, 1984.
- [25] Q. Guo, S. White, *et al.*, “From dwarf spheroidals to cD galaxies: simulating the galaxy population in a Λ CDM cosmology,” *Monthly Notices of the Royal Astronomical Society*, vol. 413, pp. 101–131, May 2011.
- [26] P. De Bernardis, P. Ade, *et al.*, “A flat universe from high-resolution maps of the cosmic microwave background radiation,” *Nature*, vol. 404, pp. 955–959, 4 2000.
- [27] A. Zee, *Quantum Field Theory in a Nutshell*. Princeton University Press, 2003.
- [28] G. Jungman, M. Kamionkowski, *et al.*, “Supersymmetric dark matter,” *Physics Reports*, vol. 267, no. 5, pp. 195 – 373, 1996.
- [29] “Standard model of elementary particles: the 12 fundamental fermions and 5 fundamental bosons.” Web, June 2006. Available at https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg.
- [30] P. J. E. Peebles, “Recombination of the Primeval Plasma,” *Astrophysical Journal*, vol. 153, p. 1, Jul 1968.
- [31] “Build a Universe with the CMB Power Spectrum Analyzer.” Web, May 2019. Available at https://lambda.gsfc.nasa.gov/education/cmb_plotter/.
- [32] D. J. Fixsen, “The Temperature of the Cosmic Microwave Background,” *The Astrophysical Journal*, vol. 707, pp. 916–920, nov 2009.

- [33] M. Aaboud, G. Aad, *et al.*, “Constraints on mediator-based dark matter and scalar dark energy models using $\sqrt{s} = 13$ TeV pp collision data collected by the ATLAS detector,” *Journal of High Energy Physics*, vol. 2019, May 2019.
- [34] J. Ellis and K. A. Olive, *Supersymmetric dark matter candidates*, pp. 142–163. Cambridge University Press, 2010.
- [35] A. Boveia and C. Doglioni, “Dark matter searches at colliders,” *Annual Review of Nuclear and Particle Science*, vol. 68, pp. 429–459, Oct 2018.
- [36] N. Craig and A. Katz, “The fraternal WIMP miracle,” *Journal of Cosmology and Astroparticle Physics*, vol. 2015, pp. 054–054, Oct 2015.
- [37] V. Silveira and A. Zee, “Scalar phantoms,” *Physics Letters B*, vol. 161, no. 1, pp. 136 – 140, 1985.
- [38] R. Essig, J. A. Jaros, *et al.*, “Dark Sectors and New, Light, Weakly-Coupled Particles,” 2013.
- [39] L. Bian, R. Ding, *et al.*, “Two component higgs-portal dark matter,” *Physics Letters B*, vol. 728, pp. 105–113, Jan 2014.
- [40] M. Klasen, M. Pohl, *et al.*, “Indirect and direct search for dark matter,” *Progress in Particle and Nuclear Physics*, vol. 85, pp. 1–32, Nov 2015.
- [41] D. Piras, B. Joachimi, *et al.*, “The mass dependence of dark matter halo alignments with large-scale structure,” *Monthly Notices of the Royal Astronomical Society*, vol. 474, pp. 1165–1175, Nov 2017.
- [42] SuperCDMS Collaboration, “Technical Design Report: Advanced Interleaved Detectors for SuperCDMS Soudan,” *Submitted to DOE and NSF*, 2010.
- [43] J. Lewin and P. Smith, “Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil,” *Astroparticle Physics*, vol. 6, no. 1, pp. 87 – 112, 1996.

- [44] L. M. Widrow, B. Pym, *et al.*, “Dynamical Blueprints for Galaxies,” *Astrophys. J.*, vol. 679, pp. 1239–1259, 2008.
- [45] A. M. Green, “Dependence of direct detection signals on the WIMP velocity distribution,” *Journal of Cosmology and Astroparticle Physics*, vol. 2010, pp. 034–034, oct 2010.
- [46] J. R. Watson, “B.S. Thesis,” *Texas A&M University*, 2016.
- [47] R. Agnese, A. J. Anderson, *et al.*, “Search for Low-Mass Weakly Interacting Massive Particles with SuperCDMS,” *Phys. Rev. Lett.*, vol. 112, p. 241302, Jun 2014.
- [48] R. Agnese, A. J. Anderson, *et al.*, “Projected sensitivity of the SuperCDMS SNOLAB experiment,” *Phys. Rev. D*, vol. 95, p. 082002, Apr 2017.
- [49] D. S. Akerib, C. W. Akerlof, *et al.*, “Projected WIMP sensitivity of the LUX-ZEPLIN (LZ) dark matter experiment,” 2018.
- [50] “SuperCDMS Dark Matter Limit Plotter.” Web, May 2019. Available at <https://supercdms.slac.stanford.edu/dark-matter-limit-plotter>.
- [51] R. Agnese, A. J. Anderson, *et al.*, “Low-mass dark matter search with CDMSlite,” *Phys. Rev. D*, vol. 97, p. 022002, Jan 2018.
- [52] S. Hertel, “Ph.D. Thesis,” *MIT*, 2012.
- [53] A. Anderson, “Ph.D. Thesis,” *MIT*, 2015.
- [54] D.-M. Mei and A. Hime, “Muon-induced background study for underground laboratories,” *Physical Review D*, vol. 73, 053004, 2006.
- [55] S. Al Kenany, J. A. Rolla, *et al.*, “SuperCDMS Cold Hardware Design,” *Journal of Low Temperature Physics*, vol. 167.5-6, pp. 1167–1172, 2012.
- [56] “The Super Cryogenic Dark Matter Search Experiment.” Web, 2015. Available at <http://cdms.berkeley.edu/experiment.html>.

- [57] “Decay chain $4n+2$, uranium series.” Web, December 2008. Available at [https://commons.wikimedia.org/wiki/File:Decay_chain\(4n%2B2,_Uranium_series\).PNG](https://commons.wikimedia.org/wiki/File:Decay_chain(4n%2B2,_Uranium_series).PNG).
- [58] G. Wang, “Phonon emission in germanium and silicon by electrons and holes in applied electric field at low temperature,” *J. Appl. Phys.*, vol. 107, 094504, 2010.
- [59] P. N. Luke, “Voltage-assisted calorimetric ionization detector,” *Journal of Applied Physics*, vol. 64, pp. 6858–6860, Dec. 1988.
- [60] J. Lindhard, V. Nielsen, *et al.*, “Integral Equations Governing Radiation Effects. (Notes on Atomic Collisions, III),” *Kgl. Danske Videnskab., Selskab. Mat. Fys. Medd.*, vol. Vol: 33: No. 10, 1 1963.
- [61] J. F. Ziegler, “The electronic and nuclear stopping of energetic ions,” *Applied Physics Letters*, vol. 31, no. 8, pp. 544–546, 1977.
- [62] “CDMS Detector Monte Carlo Documentation.” Web, February 2012. Available at http://titus.stanford.edu/cgi-test/cvsweb.cgi/CDMS_DetectorMC/CDMS_DMC_Manual.pdf.
- [63] N. W. Ashcroft and N. D. Mermin, *Solid State Physics*. Brooks/Cole, Cengage Learning, 1976.
- [64] C. Kittel, *Introduction to Solid State Physics*. Wiley, 8th ed. ed., 2005.
- [65] “A diagram of the first Brillouin zone of a face-centred cubic (FCC) lattice, with points of high symmetry marked.” Web, May 2008. Available at [https://commons.wikimedia.org/wiki/File:Brillouin_Zone_\(1st,_FCC\).svg](https://commons.wikimedia.org/wiki/File:Brillouin_Zone_(1st,_FCC).svg).
- [66] R. Westbrook, “Hole traps in germanium crystals,” *Nuclear Instruments and Methods*, vol. 108, no. 2, pp. 335 – 338, 1973.
- [67] H. D. Hagstrum, “Auger Electron Ejection from Germanium and Silicon by Noble Gas Ions,” *Phys. Rev.*, vol. 119, pp. 940–952, Aug 1960.

- [68] P. M. Cabrera, B. and *et al.*, “Oblique propagation of electrons in crystals of germanium and silicon at sub-Kelvin temperature in low electric fields,” *eprint arXiv:1004.1233v1*, 2010.
- [69] S. Tamura, J. A. Shields, *et al.*, “Lattice dynamics and elastic phonon scattering in silicon,” *Phys. Rev. B*, vol. 44, pp. 3001–3011, Aug 1991.
- [70] C. Jacoboni, *Theory of Electron Transport in Semiconductors: A Pathway from Elementary Physics to Nonequilibrium Green Functions*. Springer Series in Solid-State Sciences, Springer Berlin Heidelberg, 2010.
- [71] S. Ramo, “Currents Induced by Electron Motion,” *Proceedings of the IRE*, vol. 27, pp. 584–585, 9 1939.
- [72] W. Shockley, “Currents to Conductors Induced by a Moving Point Charge,” *Journal of Applied Physics*, vol. 9, no. 10, pp. 635–636, 1938.
- [73] J. Filippini, “Ph.D. Thesis,” *University of California, Berkeley*, 2008.
- [74] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Prentice-Hall, 1979.
- [75] S. Golwala, “Ph.D. Thesis,” *University of California, Berkeley*, 2000.
- [76] D. Moore, “Ph.D. Thesis,” *California Institute of Technology*, 2012.
- [77] R. Agnese, T. Aramaki, *et al.*, “Results from the Super Cryogenic Dark Matter Search Experiment at Soudan,” *Phys. Rev. Lett.*, vol. 120, p. 061802, Feb 2018.
- [78] E. García-Toraño, V. Peyres, *et al.*, “Simulation of decay processes and radiation transport times in radioactivity measurements,” *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, vol. 396, pp. 43 – 49, 2017.
- [79] L. Zhang, R. Mao, *et al.*, “Fast Neutron Induced Nuclear Counter Effect in Hamamatsu Silicon PIN Diodes and APDs,” *IEEE TRANSACTIONS ON NUCLEAR SCIENCE*, vol. 1, 10 2011.

- [80] J. Sander, “Ph.D. Thesis,” *University of California, Santa Barbara*, 2007.
- [81] J. Allison, K. Amako, *et al.*, “Geant4 developments and applications,” *IEEE Transactions on Nuclear Science*, vol. 53, pp. 270–278, Feb 2006.
- [82] “COMSOL Multiphysics®v.5.2.” Web, 2018. COMSOL AB, Stockholm, Sweden. <http://www.comsol.com>.
- [83] ISO, *ISO/IEC 14882:1998: Programming languages — C++*. Sept. 1998. Available in electronic form for online purchase at <http://webstore.ansi.org/> and <http://www.cssinfo.com/>.
- [84] R. Brun and F. Rademakers, “ROOT: An object oriented data analysis framework,” *Nucl. Instrum. Meth.*, vol. A389, pp. 81–86, 1997.
- [85] “Notes on TES Simulation.” Web, September 2016. Available at <https://confluence.slac.stanford.edu/download/attachments/120202114/notes-tes-simulation.pdf?version=2&modificationDate=1502747212000&api=v2>.
- [86] “CDMS Event Format for Soudan iZip and Hybrid iZip/Mercedes Runs.” Web, September 2012. Available at http://titus.stanford.edu/cgi-test/cvsweb.cgi/Documentation/daq_format/Soudan_iZIP/Soudan_iZIP_event_format.pdf.
- [87] U. Fano, “Ionization Yield of Radiations. II. The Fluctuations of the Number of Ions,” *Phys. Rev.*, vol. 72, pp. 26–29, Jul 1947.
- [88] I. J. Goodfellow, J. Pouget-Abadie, *et al.*, “Generative Adversarial Networks,” 2014.

APPENDIX A

EXCESS OF ZIP4 BLIPS AT $t=0$

In Chapter 6.3, Figure 6.9 (top left) we noted that there is a concentration of Blips at $t=0$, that is half-pulse before trigger time. We have no real understanding about the actual source of these events, but we can inspect them. Upon selecting events with BlipFinder time < 2 time-bins, plotting them, and looking at the average pulse shape (black dashed line) we can notice that the Blip feature wraps around the readout, as shown in Figure A.1. This suggests that these Blips are not an effect occurring during the readout, but rather at the DAQ electronics where something is affecting the data after it has been recorded.

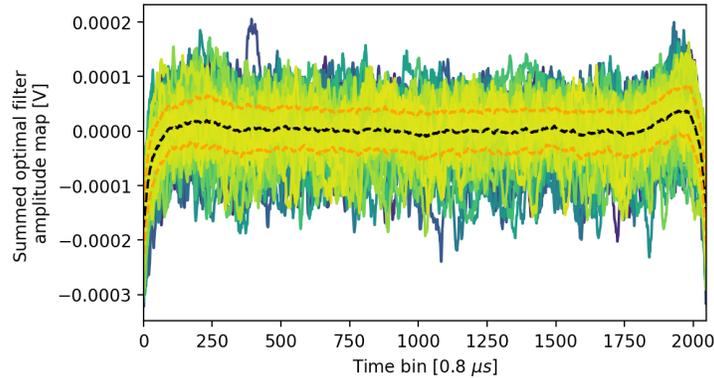


Figure A.1: Waveforms of Zip4 Blips occurring within the first 2 time-bins show that these Blips wrap around the readout, a possible indication that Blips are a DAQ electronics issue.

APPENDIX B

TOWERBLIP RATES

We describe the method and rate of coincidence of Blips across the detectors of the same tower to determine how often they are Single-Detector versus TowerBlips. We begin with an overview of the considerations and describe the method. Next we present our rate characterization results, which justifies, due to the high fraction of TowerBlips, for performing a correction (described in next section).

We have identified that Blips occur in all channels of detectors 4 and 5, and in side2 channels for detector 6. Only correlated channels are considered in the Blip characterization. The BlipFinder method for the characterization is as follows:

1. Use the BlipFinder to look for biggest Blip in one detector of the tower (call this detector A)
 - (a) Add optimal-filtered channels (flipping polarities of anti-correlated channels)
 - (b) Ignore first/last 248 time-bins (avoid structure at the beginning/end, see Appendix A)
 - (c) Obtain BlipFinder time t_A and energy E_A
2. Search for correlated Blip in other 2 detectors independently (call these B and C)
 - (a) Add optimal-filtered channels of each detector one at a time (flipping polarities of anti-correlated channels)
 - (b) Look for biggest Blip in 100 bin search window from t_A ($t_A \pm 50$ time-bins)
 - (c) Obtain the BlipFinder time t_B, t_C and energy E_B, E_C for each detector
3. Estimate the number of correlated Blips in detectors B and C

- (a) Select events with Big Blips in detector A (with large negative Blip energy)
- (b) Calculate the time difference from t_A ($\Delta_B = t_B - t_A$ and $\Delta_C = t_C - t_A$)
- (c) Define thresholds at the edges of the structure formed near $\Delta_{B,C} = 0$, and call events coincident if within such range
- (d) Estimate and subtract the background to obtain fraction of Blips coincident: in B, in C, in B and C, and in neither.

4. Alternate A, B, and C

We have characterized Blips: how often they occur, how often they are correlated with other detectors, and how often they are single-detector Blips. The background subtracted estimates of Blip coincidences are shown in Table B.1. From it we know that most of the Blips occur in all 3 iZips (at least 65%), and only 10% occur in just iZip4. This is a clear justification for developing a Blip correcting algorithm for Blips affecting the iZip4 measurement, identified in iZip5 and 6.

BBZ4		BBZ5		BBZ6	
In 5, but not 6	16.2%	In 4, but not 6	15.5%	In 4, but not 5	0.6%
In both 5 and 6	64.7%	In both 4 and 6	67.0%	In both 4 and 6	71.2%
In 6, but not 5	8.4%	In 6, but not 4	10.3%	In 5, but not 4	14.8%
In neither 5 nor 6	10.7%	In neither 4 nor 6	7.2%	In neither 4 nor 5	13.3%

Table B.1: Results of background subtracted fraction calculations. Note that iZip6 only considers side2 channels. For details on the background subtraction see Appendix C

APPENDIX C

TOWERBLIP COINCIDENCE RATES

We can determine the number of inter-detector coincident blips in our sample. We will expand on details from the methodology described in Appendix B. First we will show the BlipFinder results for each detector (ignoring the first/last 248 bins). Next we will show the BlipFinder results on the other detectors in a 100-bin window around the BlipFinder time of the Blip search of one detector at a time. Finally we will show the background subtractions.

We begin with the BlipFinder results of the independent search in each detector, in Figure C.1. For this search we note that there is no preference for Blips occurring at any particular time, not even for Zip4 BigBlips. We define BigBlips for each detector as events with BlipFinder total energy less than -3.5keV (Zip4), -2.5keV (Zip5), and -1.5keV (Zip6). We will select BigBlips and look at the time-coincidences next.

Now we select the BigBlip events and show the results for the narrowed BlipFinder search performed in the opposite detectors. The search is done within 100 time-bins of the BigBlip time. In particular we are interested in the difference between the BlipFinder search and the time of the BigBlip, the results are shown in Figure C.2, for Zip4 BigBlips (top), Zip5 BigBlips (middle), Zip6 BigBlips (bottom), the left column shows cases when a second detector is coincident, and the right column events where the second detector is not coincident, and the distributions are for the third detector. We arbitrarily define the timing thresholds to determine if they are coincident (in dashed orange), we sub-divide the samples from the results of one detector to look at the distribution in the other one, we estimate the background by taking the average value outside the coincident time, and finally we calculate the background-subtracted number of events. The rates are displayed in the legends of the plots, but need to be multiplied to reproduce the values presented in Table B.1.

We have seen the time-coincident rates, now we look at energy (amplitude) correlation.

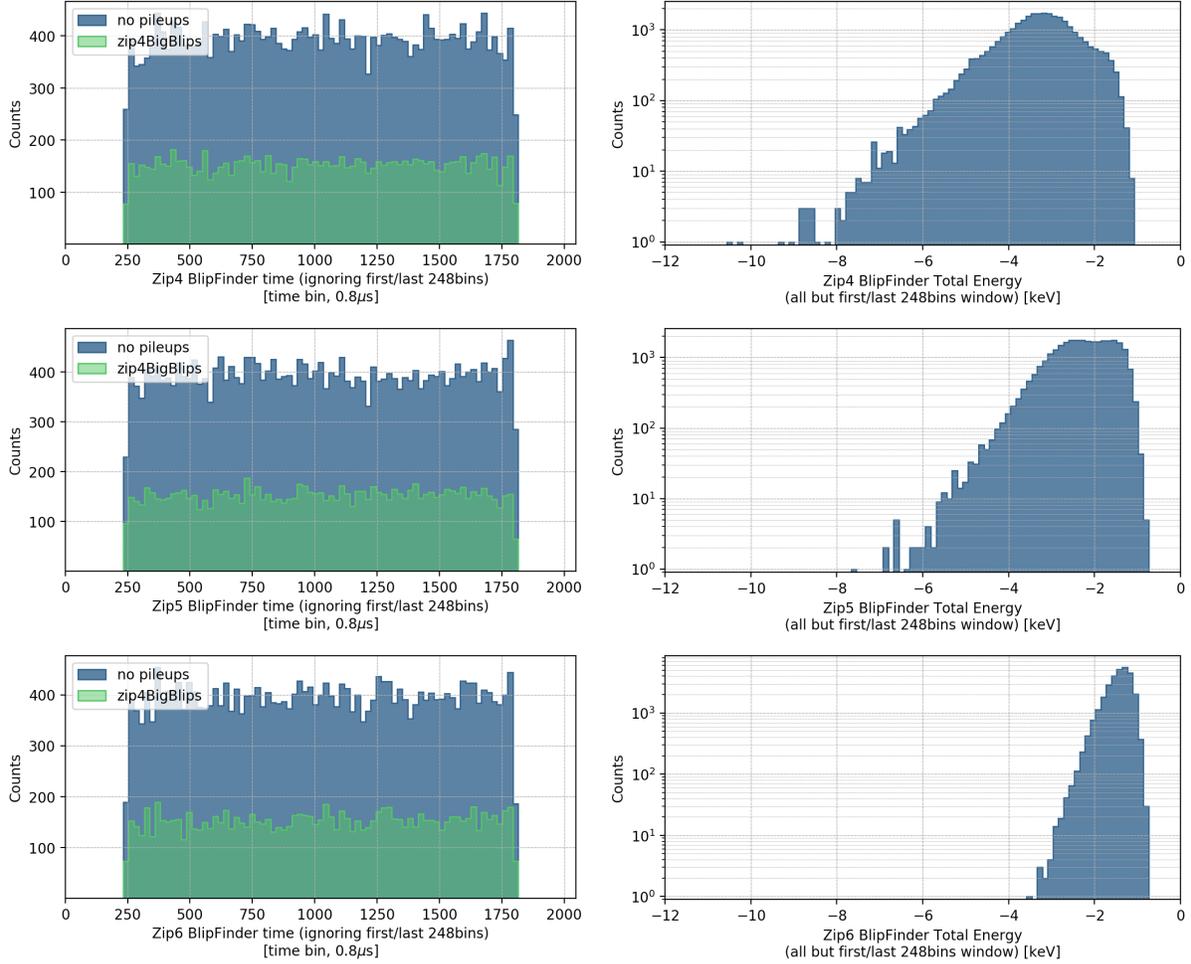


Figure C.1: BlipFinder (ignoring first/last 248 bins), time on the left and energy on the right, for Zips 4 (top), 5 (middle), and 6 (bottom). In none of these independent searches there is a preference for Blips at any time, not even when looking at Zip4 Big Blips. We define BigBlips as events with BlipFinder total energy less than -3.5keV (Zip4), -2.5keV (Zip5), and -1.5keV (Zip6).

We follow a similar approach, but this time narrowed down to just looking at Zip5 and Zip6 Blips within 100 bins of the Zip4 biggest blip (searched in all but first/last 248 time bins). We divide the sample, as done before in Zip4 Big Blips and Zip5 Big Blips. Next we sub-divide the sample by the coincidence definitions as described above (from the difference from the Zip4 Blip time), then we look at the energy of the sub-divided distributions, as shown in Figure C.3. We note that coincident blips always correlate with the negatively biased distributions of energy for all detectors, while not-correlated blips are distributed

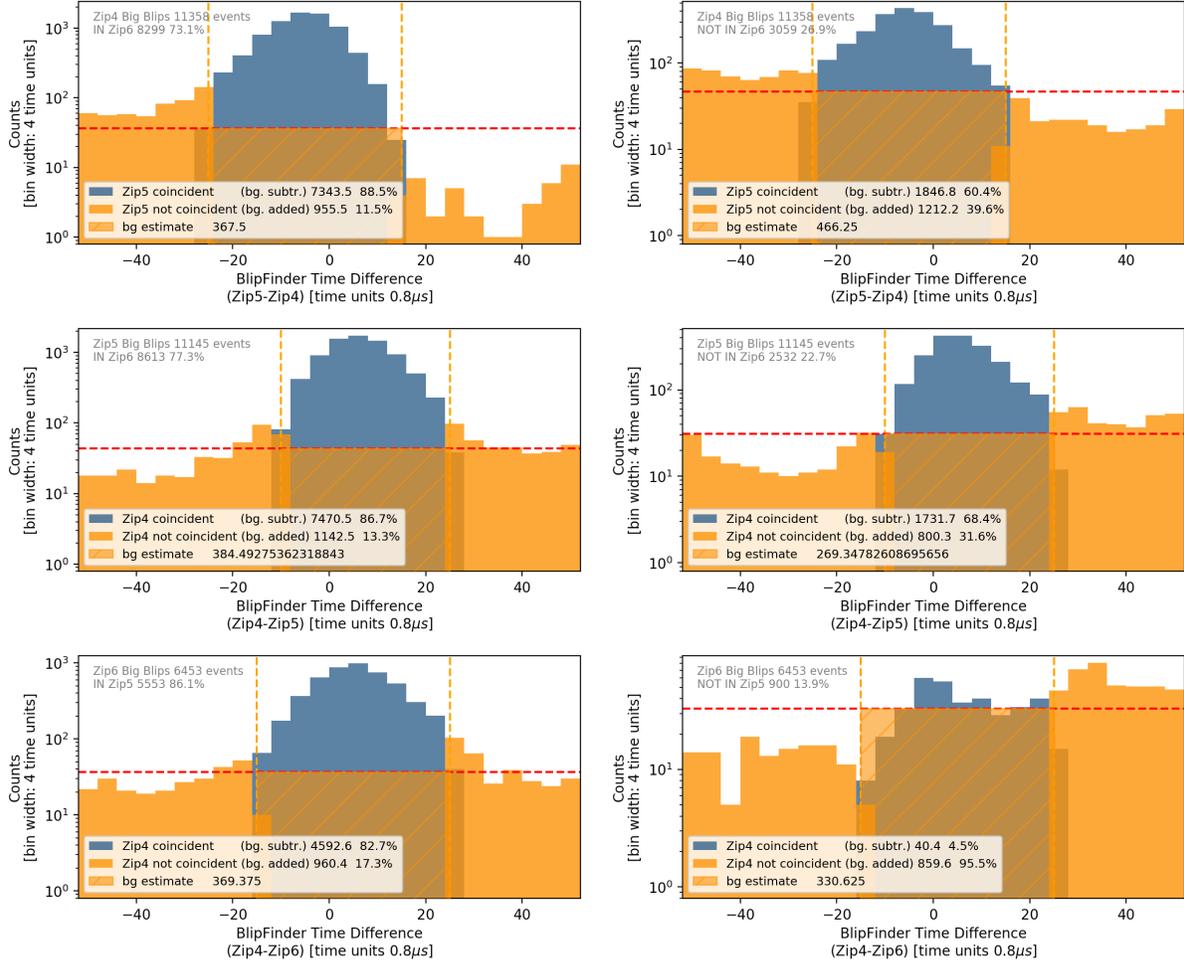


Figure C.2: The difference between the BlipFinder time (performed within 100 time-bins of the BigBlip time), and the time of the BigBlip, for Zip4 BigBlips (top), Zip5 BigBlips (middle), Zip6 BigBlips (bottom), the left column shows cases when a second detector is coincident, and the right column events where the second detector is not coincident, and the distributions are for the third detector. The timing thresholds to determine if they are coincident are in dashed orange. The samples are sub-divided from the results of the second detector (indicated in the top-left label in each plot), to look at the distribution in the third one (indicated in the x-axis label), we estimate the background by taking the average value outside the coincident time (in dashed red), and finally we calculate the background-subtracted number of events. The rates are displayed in the legends of the plots.

closer to zero energy. This is true not only for BigBlips, but also for small ones, where the only difference is that the not-correlated blips distribution is more populated.

We have now shown that Blips can be correlated both in time and energy for TowerBlips. This provides justification for attempting a correction in the Zip4 observation based on Zip5

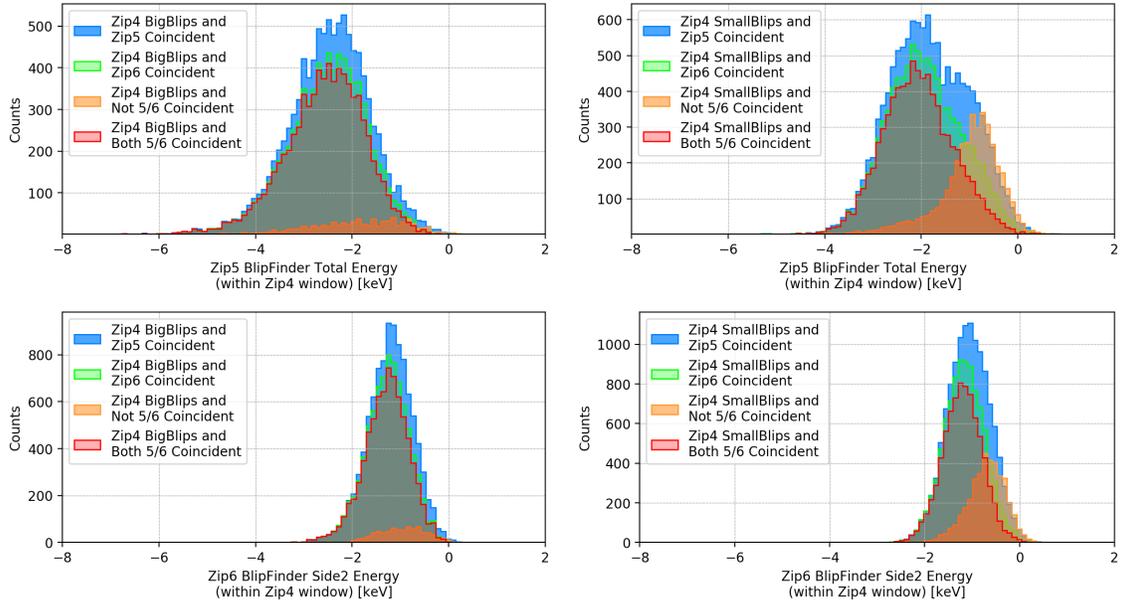


Figure C.3: Zip5 (top) and Zip6 (bottom) BlipFinder energy (searching within Zip4 Biggest Blip time ± 50 bins), for Big Blips (left) and Small Blips (right), showing that the energy is also correlated with coincidence of Blips on another Zip. This is true for both detectors. We note that even from Small Blips there is correlation, the only difference being a higher population of non-correlated Blips (Single-Detector Blips).

and Zip6 information, as performed in Chapter 6.4.

APPENDIX D

ZIP5+6 BLIPFINDER PREDICTED ENERGY FITS

As discussed in Chapter 6.4, we can use the BlipFinder tool on Zip5 and Zip6-side2, added together, to predict the Blip energy in our IRR sample. While the method was described in Chapter 6, here we show the fit values, along with the data and profile plots, to show that the fit well-represents the profile.

We show, in Figure D.1 the fits for qi1, and in Figure D.2 for qi2. The data is subdivided in planes of the Zip5+6 BlipFinder optimal filter total amplitude (uncalibrated energy estimate), from $-600\mu\text{V}$ to $200\mu\text{V}$ in steps of $50\mu\text{V}$. The figures show that the Zip4 0-delay energy estimate (side1 and side2 respectively) can be well described by a Gaussian in Zip4 0-delay energy versus Zip5+6 BlipFinder time with the amplitude being a linear function in Zip5+6 BlipFinder amplitude, as given by:

$$E_{predict}(E_{Z56BF}, t_{Z56BF}) = (A \cdot E_{Z56BF} + B)(C \cdot e^{(-t_{Z56BF} - \mu)^2 / 2\sigma^2} + D \cdot t_{Z56BF}), \quad (\text{D.1})$$

where E_{Z56BF} and t_{Z56BF} are the Zip5+6 BlipFinder energy and time, $\mu = 1024 - 4.9$ time-bins is the trigger time minus the delay observed between Zip5+6 and Zip4 (notice how in Figure C.2 Zips 5 and 6 are coincident but not Zip4), and $\sigma = 7.5$ time-bins (the mean offset and the resolution are the same as the de-weighting function). This 3D fit is our Zip4 0-delay energy prediction from Zip5+6 Blips which can be used to correct the Zip4 measurement, as shown in Section 6.4, Figure 6.15.

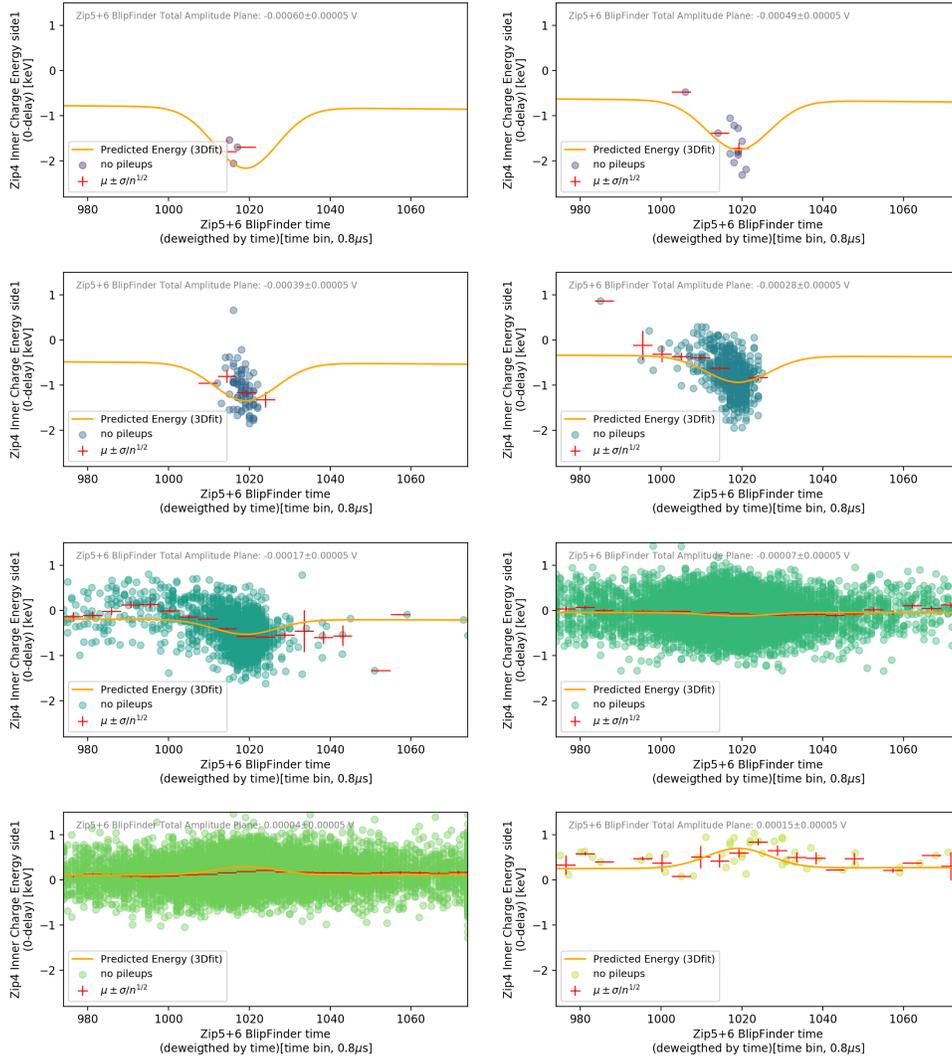


Figure D.1: Zip4 qi1 0-delay energy as function of Zip5+6 BlipFinder time. The data profile is shown in red crosses (mean and with uncertainty on the mean), the 3D fit is shown in orange.

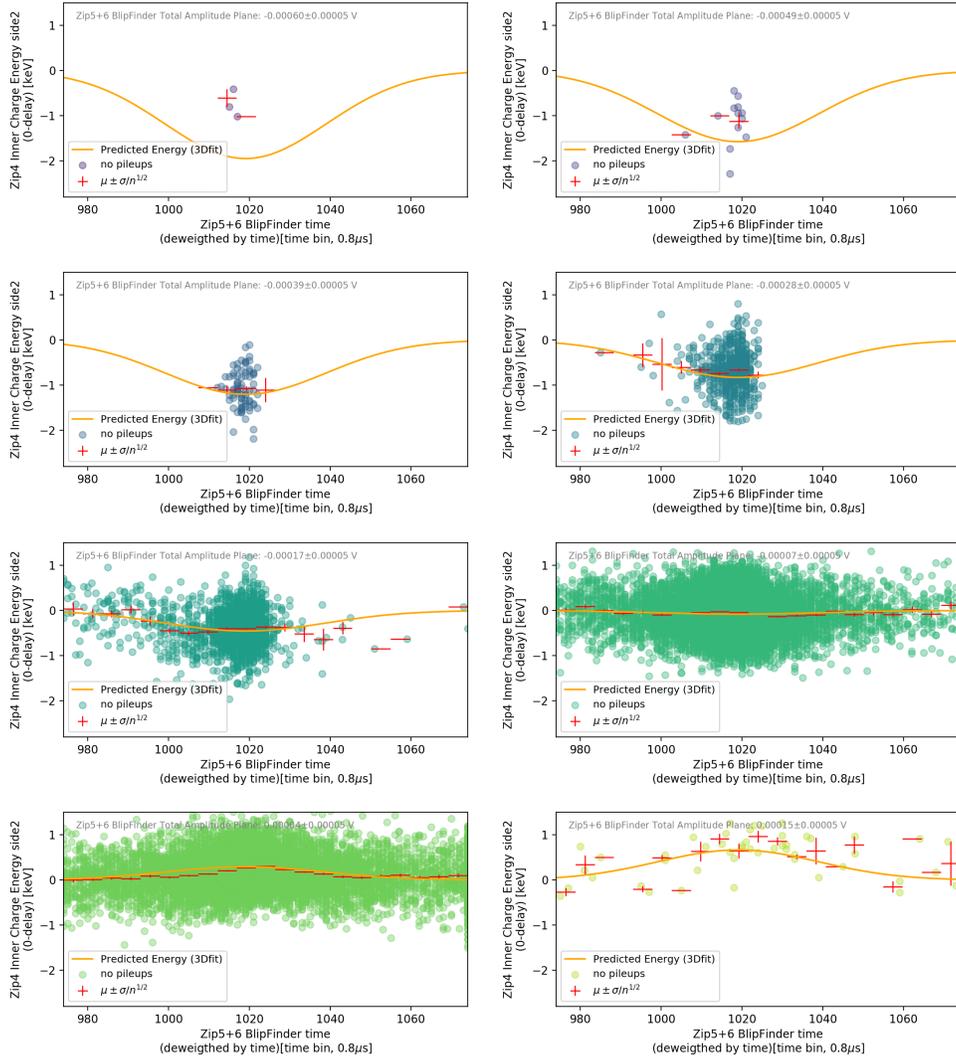


Figure D.2: Zip4 qi2 0-delay energy as function of Zip5+6 BlipFinder time. The data profile is shown in red crosses (mean and with uncertainty on the mean), the 3D fit is shown in orange.