

DEDICATED SEARCHES FOR LOW AND HIGH MASS WIMPS WITH THE
SUPERCDMS SOUDAN IZIP DETECTORS

By

BRADFORD C. WELLIVER JR.

A DISSERTATION PRESENTED TO THE GRADUATE SCHOOL
OF THE UNIVERSITY OF FLORIDA IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

UNIVERSITY OF FLORIDA

2016

© 2016 Bradford C. Welliver Jr.

To everyone that helped get me here

ACKNOWLEDGMENTS

Over the course of the past seven years I have had the opportunity to know and work with an enormous number of people and befriend many of them. This short page is not sufficient to detail my thanks to everyone who has helped me accomplish this work, ranging from collaborators, teachers, friends and family. Everyone that has provided insight, or assistance, or even been up at 3 AM to talk while I was frustrated with a subtle analysis problem, or watching the fridge, I thank you. While I cannot list everyone, there are a few that I would like to thank in particular. My apologies in advance for any omissions, they are unintentional.

The entire SuperCDMS collaboration has provided me with many opportunities to grow as a researcher, and wish to thank them all. To everyone who was on shift with me at Soudan, Rob Agnese, Ray Bunker, Kristi Schneck, Danielle Speller, Tommy Hofer, Jeff Yen, Bob Nelson, and Kunj Prasad, I am grateful for the opportunity. We had many interesting experiences! I wish to thank the Soudan mine crew as well, Jim Beaty, Jerry Meier, Denise Osterholm, Dano DeVaney, and others. Without your work the experiment would not have had as smooth an operation as it did!

Within the collaboration I'd like to thank Matt Pyle as well for providing very insightful explanations of optimal filters and detector response that were very vital in my first few years, and for always replying with enthusiasm and detail. I also wish to thank Lauren Hsu for her attention to detail during her time as analysis coordinator and for her advice on many different aspects of the experiment, as well as general advice on talks and academic life and Ray Bunker as well for providing excellent services as analysis coordinator for the high-mass search and being available to talk about some aspect of the analysis even on Sunday afternoons. I also would like to thank Dan Bauer as operations manager of SuperCDMS Soudan for allowing me to take longer on-site shifts which allowed me to gain valuable insight into the operations of the experiment, and for always being available to call when something inevitably went awry late on a Friday, especially before

a holiday weekend. Additional thanks for finding funding for, and hosting, numerous analysis workshops at Fermilab which proved to be enormously helpful in pushing along numerous analysis tasks, as face to face collaboration proves to be valuable time and again.

I wish to thank both the low-threshold analysis team members, Lauren Hsu, Adam Anderson, Julien Billard, Kristi Schneck, Elias Lopez, and Ray Bunker, and the high-threshold team, consisting of Ray Bunker, Todd Doughty, Brett Cornell, Rob Calkins, Hang Qiu, and others for working on this enormously complicated dataset and for their very insightful conversations about how to handle backgrounds and new ideas for analysis techniques. Other collaboration members such as Bruno Serfass for his technical knowledge and work on data processing, as well as Steve Yellin for managing the SUF cluster I owe thanks to, for without them analysis would be a good deal less feasible. My gratitude is also extended to the UMN team, Prisca, Anthony Villano, Scott Fallows, Mark Pepin, Hassan Chagani, Tommy Hofer, and Allison Kennedy for hosting me on several occasions.

At Florida I had the benefit to work under Tarek Saab who provided an excellent level of management that allowed a lot of freedom to explore solutions to problems on my own, but also was available for discussion if things became stuck. He also was always willing to explore different ideas and also provided me with many great opportunities for further education by suggesting various conferences for me to attend. Also in the group I'd like to thank especially Durdana Balakishiyeva for teaching me an enormous amount on the art of helium leak checking a dilution fridge and for her knowledge of how the fridge operated. She was always willing to be present even late into the day to help close up the fridge and get it running. I also wish to thank her for the many enjoyable coffee breaks during which we brainstormed ways to debug the fridge, or talked about many aspects of the SuperCDMS experiment or other interesting areas of physics, and for her advice on many different subjects. Rob Agnese has also been a great labmate to have, with many

insightful conversations about aspects of computing and detector physics. I would also like to thank Evan Kassof for his work on designing the new gas handling system and many fun anecdotes. A big thanks as well goes to the UF machine shop for help in cutting and welding bits and pieces for our gas handling system, the cryogenic support staff at UF, Greg Labbe and John Graham for helping to find parts and offer suggestions during the construction of the gas handling system, and to the department's pump technician Jay Horton for always be available to help debug some mysterious behavior and offering invaluable wisdom on operations, often with fun stories and enthusiasm for helping restore the lab to a fully operation state.

Finally, without support from my personal life grad school would have been immeasurably more difficult. I wish to thank my friends and fellow grad students Jay Perez, Doug Rank, Sean Goldberg, Evan Thatcher, and Maureen Petterson for commiserating with me and for our occasional cookouts and game or movie nights which proved invaluable for stress reduction! I wish to thank also Michelle Leonard for her longtime friendship and advice, as well as all my other friends. Finally, but not least important, I wish to thank my family for their years of unwavering support as I pursued this work. They were always understanding when I had to travel and miss celebrations or holidays, and always proud of what I was doing. In particular I thank my mom for her constant love and support through all this process which had many ups and downs.

My gratitude to all my friends, family and colleagues, without you I would not have gotten here.

TABLE OF CONTENTS

	<u>page</u>
ACKNOWLEDGMENTS	4
LIST OF TABLES	11
LIST OF FIGURES	12
ABSTRACT	17
CHAPTER	
1 INTRODUCTION	19
1.1 Observational Evidence for Dark Matter	19
1.1.1 Galaxy Dynmaics	20
1.1.2 Large Scale Structure and Weak Lensing	21
1.1.3 Cosmological Observations	22
1.2 Types of Dark Matter	24
1.3 Detection of Dark Matter	27
1.3.1 Collider Production	27
1.3.2 Indirect Detection	28
1.3.3 Direct Detection	29
1.4 Direct Detection Considerations	30
1.4.1 Interaction Rate	30
1.4.2 Halo Dynamics	34
1.4.3 Backgrounds	36
1.4.4 Sensitivity Results	37
2 SUPERCDMS AT SOUDAN	40
2.1 Experimental Setup	40
2.1.1 Shielding	42
2.1.2 Cryogenics	43
2.1.3 Detectors and Towers	44
2.1.4 Special ^{210}Pb implanted detectors	45
2.1.5 Readout Electronics and Data Acquisition	47
2.2 Data Quality and Processing	48
2.3 Data Types	55
2.4 Experiment Time Periods	57
2.4.1 CDMSlite	57
2.4.2 Detector Health	59
2.5 List of Important RQ/RRQs	61

3	IZIP DETECTORS	63
3.1	Physical Properties	63
3.2	Ionization Collection	63
3.2.1	Charge Propagation	64
3.2.2	Ionization Readout	66
3.3	Ionization Pulse Reconstruction	68
3.4	Ionization Energy Calibration	71
3.5	Ionization Noise from Random Triggers	71
3.6	Phonon Collection	79
3.6.1	Transition Edge Sensors	79
3.6.2	Phonon Production and Propagation	81
3.6.3	Phonon Readout	82
3.7	Phonon Reconstruction	85
3.8	Phonon Calibration	86
3.9	Phonon Energy Scales	89
3.10	Phonon Noise from Random Triggers	91
4	BACKGROUNDS AND IZIP DISCRIMINATION	95
4.1	Physical Origins	95
4.2	Background Discrimination	98
4.2.1	Yield	99
4.2.2	Interleaved Sensors and Partitions	100
4.2.3	Ionization Discrimination Ability	103
4.2.4	Phonon Based Discrimination	104
5	DATA QUALITY AND BLINDING CUTS	108
5.1	Analysis Framework and Cut Nomenclature	108
5.2	χ^2 Rejection Cuts	110
5.2.1	Glitch and Low-Frequency Noise Cuts	112
5.3	Other Quality Cutst	113
5.3.1	Event Based Quality Cuts	113
5.3.2	Time Period Quality Cuts	116
5.4	Multiples Cut	120
5.5	Blinding Criteria	127
5.6	High Ionization Noise Cut	131
5.7	KS and QS Test Cuts	138
5.7.1	LT Analysis Tests	140
5.7.2	HT Analysis Tests	141
5.7.3	Time Period Determination	142
5.7.4	Template Generation and Merging	143
5.7.5	Subset Selection	145
5.7.6	Kilt Plots and Thresholds	147
5.7.7	Special Considerations for the Barium Data	149

5.7.8	Cut Results	151
6	LOW MASS WIMP SEARCH	156
6.1	Lowering Thresholds via CROW	157
6.1.1	Software Modeling of the Bandpass	158
6.1.2	Hardware Implementation	160
6.2	LT Data Quality Selection	163
6.3	Pre-Fiducialization and Signal Region	163
6.4	Background Model	167
6.5	Rescaling the pzpart Model	170
6.6	Discrimination Cut Methods	173
6.6.1	Phonon Symmetry Cut	175
6.6.2	Method 2 and the BDT	177
6.7	LT Analysis Results	181
6.7.1	Exposure Weighted Efficiency	181
6.7.2	Unblinding	182
6.7.3	Limit on Low Mass WIMPs	186
7	TOWARDS A HIGH MASS WIMP SEARCH	188
7.1	High Threshold Dataset	189
7.2	Unifying Radial Discriminators	190
7.3	Choice of Discrimination Quantities	193
7.4	Construction of Background Models	199
7.4.1	^{210}Pb Induced Events Model	202
7.4.2	Gamma Model	205
7.5	Phonon Readout Asymmetry Problem	209
8	CONCLUSIONS	219
8.1	Summary	219
8.2	SuperCDMS HT Projections	219
8.3	Importance of a Detector Monte Carlo for Future Work	221
8.4	SuperCDMS SNOLAB	224
	APPENDIX: OPERATING THE DRY DILUTION FRIDGE AT UF	228
A.1	Dry Dilution Fridge	228
A.1.1	Basic Operating Steps	234
A.1.2	Pump Out	234
A.1.3	Pre-cool	235
A.1.4	Empty Pre-cool Line	236
A.1.5	Condensing the Mix	236
A.1.6	Collect the Mix	237
A.1.7	Single Shot Mode	237
A.2	Detector Tests	238

A.2.1	Obtaining Resistances from Noise Spectra and IV Curves	238
A.2.2	Obtaining T_c Values	241
A.3	Fridge Problems	243
A.3.1	Fridge Leaks	245
A.3.2	Failure of the old Gas Handling System	246
A.3.3	A New GHS	253
A.3.4	Adding More Mix and Determination of New Cooling Power	257
A.3.5	Electrical Issues	260
	LIST OF REFERENCES	263
	BIOGRAPHICAL SKETCH	270

LIST OF TABLES

<u>Table</u>	<u>page</u>
2-1 Detector Health List	58
2-2 List of Useful RRQs	62
2-3 List of Common RQs	62
3-1 Ionization OF Flag 3-digit number [xyz]	70
5-1 Singles Efficiencies on Randoms	126
5-2 List of General Blinding Requirements LT	128
5-3 List of General Blinding Requirements HT	128
6-1 LT Analysis Exposures	182
6-2 LT Candidate Events	186
7-1 HT Analysis LiveTimes (pre-unblinding)	190

LIST OF FIGURES

<u>Figure</u>	<u>page</u>
1-1 Multiple Galaxy Rotation Curves	21
1-2 The Cosmic Microwave Background All-Sky Image	23
1-3 The Cosmic Microwave Background Power Spectrum	24
1-4 WIMP Detection Processes	27
1-5 Dark Matter Results from ATLAS	28
1-6 Summary of Direct Detection Experiments	30
1-7 Differential Recoil Rates	34
1-8 WIMP-Search Results Plane	38
2-1 Muon Flux as a function of Depth Underground	41
2-2 Side View of SuperCDMS Soudan Setup	42
2-3 SuperCDMS Soudan Shielding	43
2-4 SuperCDMS Soudan iZIP Towers	46
2-5 Data Access Times	55
3-1 iZIP Phonon and Charge Channel Layout	63
3-2 iZIP Mask and Channel Labels	64
3-3 Electron vs Hole Side Radial Partition	65
3-4 Ionization Readout	66
3-5 Ionization Noise	69
3-6 Ionization Energy Calibration for Outer Channels	72
3-7 Example Noise Blob	73
3-8 Noise Distribution for 1D OF	74
3-9 Noise Distribution for a 2x2 OF without Z Constraint	75
3-10 Noise Distribution for a 2x2 OF with Z Constraint	76
3-11 Ionization Noise Fits	77
3-12 Ionization Noise Through Time	78

3-13	Median Values for Ionization Series-Dependent Noise	78
3-14	Close up view of a QET	79
3-15	Schematic of QET Readout	80
3-16	SQUID Amplifier Circuit	83
3-17	TES Noise Spectra	85
3-18	Relative Calibration Phonon Triangle Plot	88
3-19	Absolute Calibration Phonon	89
3-20	Phonon Noise Distribution	92
3-21	Phonon Noise as a function of Time	94
3-22	Median Values for Phonon Series Dependent Noise	94
4-1	Global Gamma Rate Coadded Over All Detectors	96
4-2	Decay Chain for ^{210}Pb	97
4-3	Ionization Yield	101
4-4	iZIP Electric Field	102
4-5	Ionization Surface Event Trace	103
4-6	Ionization based discrimination planes	104
4-7	Bulk vs. Surface Event Phonons	105
4-8	Effect of Local Saturation on Phonon Energy Partitions	106
4-9	Phonon Symmetry at Low Energy	107
5-1	Phonon and Ionization χ^2 Cuts	111
5-2	Glitch and Low Frequency Noise χ^2 Cuts	114
5-3	Map of Good Events Cut	116
5-4	Map of Livetime Cut	120
5-5	Example Multiples Energy Threshold	122
5-6	Multiples Cut Logic Flow	124
5-7	Singles Efficiency on IT2Z2	127
5-8	High Ionization Noise Cut	137

5-9	Low-Threshold Analysis KS/QS/RS Test Result	140
5-10	KS Test Determined Time Periods for KS Framework	143
5-11	Median KS Significance Scores for Naive Subset Selection in QS Test	146
5-12	Median KS Significance Scores for Better Subset Selection in KS Test	147
5-13	Kilt Plots for KS Tests	148
5-14	KS Threshold Plots	150
5-15	Yield CDFs for Barium Series	151
5-16	Barium Random Sampling for KS Test	152
5-17	Livetime Removed by KS and QS Tests on WS Data	153
5-18	LiveTime Removed by Non-KS Style Cuts	154
5-19	Livetime Removed by KS and QS Tests on Barium Data	155
6-1	Low Mass WIMP Search Detector Selection	156
6-2	Low Mass WIMP Search Time Period	157
6-3	iZIP Phonon Signal to Noise Response to Low and High Frequency Bandpass Cutoffs	159
6-4	Two Dimension Bandpass Filter Optimization	160
6-5	RTF Electronics Board	161
6-6	RTF Electronics Board CROW Labels	162
6-7	Low-Threshold Ionization Fiducialization	164
6-8	Low-Threshold Nuclear-Recoil Band	166
6-9	Low-Threshold Charge Model	167
6-10	Backgrounds in Energy Planes	169
6-11	Simulated Inner Beta Events	170
6-12	Observed vs Simulated pzpartOF	171
6-13	Rescaling pzpartOF in energy bins	172
6-14	Corrected pzpartOF values	174
6-15	Convergence of Optimization Method	178

6-16	LT Method 1 Cuts	179
6-17	LT Method 2 Cuts	179
6-18	LT BDT	180
6-19	LT Exposure Weighted Efficiencies	183
6-20	Observed Events compared to Expectations	184
6-21	Observed Candidate Events in Ionization vs Phonon Energy Plane	185
6-22	Low Threshold Analysis Limit	187
7-1	Ionization Radial Estimators	193
7-2	Ionization Radial Estimators for IT3Z1	194
7-3	Ionization Radial Estimators for IT3Z3	195
7-4	Ionization Radial Estimators in DMC	196
7-5	Yield Plane Discrimination	197
7-6	Radial and Z Partition Plane	198
7-7	Yield vs Radial Partition Plane	199
7-8	Yield vs Z Partition Plane	200
7-9	Ionization 3D Parameter Space	201
7-10	Example of Mapped ^{210}Pb Source Data	203
7-11	Comparison of Yield Distributions for ^{133}Ba and WS Data	208
7-12	Scattered Gamma Leakage in Partition Quantities	209
7-13	Concentrated Gamma Leakage in Partition Quantities	210
7-14	Gamma Leakage in Radial and Z Partitions Prior to Spot Removal	211
7-15	Gamma Leakage in Radial and Z Partitions Prior to Spot Removal	212
7-16	Phonon Channel Hot Spot	213
7-17	Phonon Channel Readout vs Analysis Labels	214
7-18	Yield Sampling Regions Around Detector	215
7-19	Yield Distributions from Sampled Locations in Phonon XY Partition	215
7-20	Side Specific Yield vs Radial Plane with DIB Spot	217

8-1	Projected SuperCDMS HT Limit	222
8-2	iZIP Design for SuperCDMS SNOLAB	225
8-3	Projected Limits for Next Generation Experiments	227
A-1	Pre-cool Circuit	232
A-2	Condensing Circuit	233
A-3	Payload Setup	239
A-4	TES Noise	240
A-5	Ib-Is Curve	242
A-6	I_c vs T to obtain T_c	244
A-7	Leaky Condenser Joint	247
A-8	Leak on Precool Filter	248
A-9	Compressor Pump Degredation	249
A-10	Compressor Pump Diaphragm Damage	250
A-11	Original Gas Handling System	252
A-12	Indicator of Compressor Leak	253
A-13	Construction of the Front Panel	255
A-14	New Gas Handling System	256
A-15	Cooling Power of Reptile	259
A-16	Flow rate and parasitic power of Reptile	261

Abstract of Dissertation Presented to the Graduate School
of the University of Florida in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy

DEDICATED SEARCHES FOR LOW AND HIGH MASS WIMPS WITH THE
SUPERCDCMS SOUDAN IZIP DETECTORS

By

Bradford C. Welliver Jr.

May 2016

Chair: Tarek Saab

Major: Physics

Recent cosmological evidence suggests most of the mass of the universe takes the form of a type of particle that we have not been able to directly detect. Nearly 80 years that have elapsed since the first hints of this dark matter started to appear from astronomers without any direct detection. The high precision era of cosmology and unifying models of particle physics developed in the 20th century have presented us with an exciting mystery at the intersection of these two fields that needs to be solved.

SuperCDMS Soudan operates specialized germanium detectors (iZIPs) that are cooled to milliKelvin temperatures deep underground in the Soudan Underground Laboratory with the hope of detecting a rare collision between dark matter and a nucleus. A search for low-mass dark matter comes with multiple unique challenges since the background discrimination abilities of these detectors becomes less powerful at the low energies needed to probe low-mass dark matter since the signal to noise ratio deteriorates. Using a sophisticated background model via a pulse rescaling technique, SuperCDMS Soudan was able to produce a world leading exclusion limit on low-mass dark matter.

Effort is to extend the analysis to higher masses require long running times during which many aspects of the detectors or the environment can change. Additional challenges are offered by the powerful background discrimination ability of the iZIP. The background distributions are well separated from the signal region, meaning most of the leakage arises from low-probability tails of the background distributions. In the absence of an enormous

dataset, extrapolations from the bulk of the distribution are required. While attempting to obtain a model of gamma induced electron-recoils leaking into the signal region of the detector from high radius a curious asymmetry between the sides of the detectors was discovered potentially indicating an electronics or detector design problem. This thesis describes the physics behind SuperCDMS Soudan, the numerous tasks involved in a low-mass search and the first iZIP array science results with these new detectors, as well as the developments towards a high-mass search result.

CHAPTER 1 INTRODUCTION

One of the largest unsolved mysteries in physics arises at the intersection of cosmology and particle physics. The 20th century saw a so-far unparalleled explosion in our knowledge of the universe at both large and small scales. On the small scale the Standard Model reigns supreme, describing 3 of the 4 known fundamental forces and predicting with astounding accuracy a range of particles and their properties. At the same time, cosmology became ever more quantitative in our understanding of the universe at the largest scales. Driven by the other fundamental force (gravity as explained by general relativity) the century saw the rise of the big bang model, inflationary physics, and modeling of the large scale structure of the universe. Cosmology now has its own version of the standard model, the Λ CDM model [1, 2], where Λ represents dark energy, and CDM is for Cold Dark Matter, two of the most exciting discoveries in cosmology that are not as yet fully explained. Late last century it was discovered that the universe is expanding an accelerating rate [3, 4], leading to the discovery of dark energy. Dark energy is not the only large mystery unearthed by cosmological observations. Beginning in the early 20th century onward, it appeared that an unknown source of extra mass was also influencing the evolution and dynamics of the universe, and that the source of this extra, dark matter, could not be explained by the standard model of particle physics. This dark matter is hypothesized to be a new class of particle (or particles) that do not interact via the strong nuclear force nor with the electromagnetic force (otherwise we would be able to detect it easily). In this chapter I will briefly outline the history and motivation for dark matter as well as the three main approaches towards detection of this mysterious component of our universe. More detailed information can be found in several references [5–7].

1.1 Observational Evidence for Dark Matter

There are three broad observational categories that support the notion that some type of unseen dark matter exists and dominates the mass density of the universe. Below

I will give a brief summary of each of the three, and when taken together they paint a compelling picture that it is indeed some as yet discovered class of particle(s) that give rise to this dark matter, as opposed to a modification of gravity.

1.1.1 Galaxy Dynmaics

The first clues to the existence of dark matter arose from observation of galaxy velocities in clusters and star velocities in galaxies. At the time it was well established how objects in orbit should move based off of Newtonian dynamics. If an object is in orbit around a much larger central mass, then the centripetal force balances out the gravitational force and we can derive the orbital speed as a function of separation distance:

$$v(r) = \sqrt{\frac{GM(r)}{r}} \quad (1-1)$$

The mass distribution as a function of radius can be somewhat complicated, but at distances sufficiently far from the central bulge of a galaxy, the added mass of the luminous disk is negligible, and thus $M(r) \sim \text{const.}$ which means that the orbital velocity should fall as $v(r) \sim r^{-1/2}$. This is certainly what is seen in the orbital motion of planets around the sun. In the 1930's Fritz Zwicky was performing the first observational studies on the Coma Cluster and in particular the motion of galaxies within this cluster. Zwicky measured the radial velocities of galaxies in the cluster using Doppler shifts and proceeded via the virial theorem to compute the total mass of the cluster arriving at a value ~ 500 times larger than would be estimated from the luminous mass alone.[\[8\]](#) He stated that should this observation hold up upon further investigation, this would imply a large component of unseen 'dark matter' dominating the cluster's dynamics.

As observational equipment improved, the study of stellar rotation in individual galaxies as a function of radius could begin in earnest. In the 1970's Vera Rubin studied the orbital velocity of stars in Andromeda and later more high luminosity galaxies [\[9, 10\]](#) and surprisingly found the rotation curves to be fairly flat at high radii, a result that

completely disagreed with standard Newtonian expectations. Since that time this pattern has been confirmed [11, 12] in many other galaxies.




Figure 1-1. Sample of multiple spiral galaxy rotation curves. These show the orbital motion of stars within a galaxy as a function of distance from the central region. At high radii the orbital velocity becomes flat as opposed to the expected $r^{-1/2}$ behavior from Keplerian motion suggesting extra unseen mass. Reprinted by permission from Rubin V. et al. 2001. Annual Review of Astronomy and Astrophysics (Page 149, Figure 4). Annual Reviews.[11].

1.1.2 Large Scale Structure and Weak Lensing

Studies of the large scale structure of the universe show it has a particular distribution that is not uniform.[1, 13] At the largest scales the universe appears to be clumped into filaments along which galaxies cluster. Detailed numerical simulations can try to reproduce this distribution based on known physics of the Λ CDM model. In particular numerical simulations are appealing as they allow one to control the relative abundances of baryonic to non-baryonic matter, as well as the temperature of the non-baryonic matter. Typically these models show that cold (i.e., non-relativistic) dark matter allows for the appropriate distribution of clusters at the largest scales [14, 15].

In addition to large scale structure simulations, weak gravitational lensing surveys can provide clues that dark matter exists [16]. Weak lensing allows one to see regions of space that distort light from background sources without any baryonic matter present and

determine that mass should be present. A striking example of this comes from the Bullet Cluster [17]. Here two clusters of galaxies are in the process of colliding, and the baryonic matter on average gets hotter and remains in a centralized region due to dissipative friction effects. Meanwhile, some form of dark matter proceeds to stream through since it does not strongly interact. The result are lensing maps that suggest over dense regions that are separated from all the baryonic matter and which appear to be otherwise empty space.

1.1.3 Cosmological Observations

At even larger scales the detailed study of the evolution, and structure of the universe as a whole also indicates there must be more mass than is accounted for by standard baryonic matter. The Λ CDM model has many components that support dark matter, ranging from big bang nucleosynthesis (BBN) [18, 19], to baryon acoustic oscillations (BAO) [20], the large-scale structure of the universe [13] and the cosmic microwave background (CMB) [21]. The CMB in particular is a powerful tool to probe the contents of the universe since it contains an imprint of the nascent universe. The universe was extremely hot when it was young, and for the first portion of its existence it was mostly ionized plasma. But as the universe expanded, it slowly cooled. When the universe was approximately 380,000 years old, the temperature had fallen to approximately 3000 K, and hydrogen atoms were able to form making the universe transparent. At this time, photons from the big bang decoupled from these atoms and on average spread throughout the universe. As the universe expanded, these photons were redshifted until today they are on the order of microwave wavelengths with a temperature of ~ 2.73 K. The CMB was first discovered in the 1960's by Penzias and Wilson [22] while trying to remove a mysterious excess of microwave signal from an antenna. To first order the the CMB is isotropic, however small anisotropies are observed with ever finer resolution first with COBE [23], then WMAP [24], and most recently the Planck satellite [21].




Figure 1-2. The all-sky CMB anisotropies. These small temperature variations are the after effects of density fluctuations in the primordial universe. The anisotropies in temperature are very small, on the order of hundreds of μK . Reprinted from the Planck Collaboration. 2015. arXiv:1502.01582v2 (Page 19, Figure 9). arXiv.[25].

The CMB power spectrum can be expanded in spherical harmonics and then analyzed to discern information about the structure and content of the universe, and shows excellent agreement between theory and observation. The first 3 peaks in the power spectrum are the largest and contain useful information about the universe. The first peak provides some information as to the geometry of the universe (i.e., flat vs curved). When the last scattering occurred, certain acoustic waves in the plasma would have just reached maximum compression and so the first peak can tell us the angular scale at which this happened, which in turn is dependent upon the curvature of the universe. In the early universe oscillations of baryonic matter would be influenced by both gravity (causing over-dense regions) and radiation pressure from energetic photons (smoothing the density), however any non-baryonic dark matter that did not interact via electromagnetism would tend to only clump together. As such the ratio of the first and second peaks in the power

spectrum provides information as to the baryonic matter density, and the third peak provides information about the dark matter density (in particular the degree that the 2nd peak is suppressed). Degeneracies exist in the CMB however so alone it cannot provide all the information needed to fix the abundances, however in parallel with the previously mentioned measurements the current best fits for the amount of material in the universe are given as 4.9% matter, 26.8% dark matter, and 68.3% dark energy [21]. This means that dark matter comprises about 5.5 times as much mass as regular matter in the universe (or $\sim 82\%$ of the mass is in the form of dark matter).




Figure 1-3. The CMB power spectrum is displayed with a best fit shown in red, and 1σ error bars. The residuals are in the bottom panel and are very small, showing how well the Λ CDM model agrees with observation. Reprinted from the Planck Collaboration. 2015. arXiv:1502.01589v2 (Page 8, Figure 1). arXiv.[21].

1.2 Types of Dark Matter

There are a wide variety of dark matter candidates in existence [26–32] however some are more favored than others. Generally each candidate type does have a few properties in common. Based on the previously mentioned observational constraints, we have good

reason to expect dark matter to be cold (non-relativistic). Hot dark matter is ruled out from the large scale structure of the universe, as relativistic dark matter would tend to smooth out primordial density fluctuations as the universe expanded. Cold dark matter however would be able to remain clustered together locally despite the expanding universe. This also rules out the one standard model particle that could potentially have fit the bill for dark matter, the neutrino. While the neutrino masses are not yet known, they are known to be very light particles, and since they thermally decoupled from the early universe when it had a temperature of a few MeV, this means they had very high kinetic energies and so would be a form of relativistic dark matter.

Moving only slightly beyond the standard model, a somewhat popular candidate is the axion. Axions are a proposed solution to a standard model problem involving CP violation in the strong force. CP violation is permissible in strong interactions, however thus far no signs of this violation are observed, resulting in what is known as the strong CP problem. Peccei and Quinn [33] proposed a mechanism in which CP violation can be suppressed via the presence of a broken $U(1)$ symmetry which is spontaneously broken. The resulting Nambu-Goldstone boson is the axion. Axions, if they exist, are thought to have masses in the range of μeV - eV . Several searches are underway for these particles, some of which are described in [26, 34, 35].

Moving further beyond the standard model leads to a category of dark matter referred to as Weakly Interacting Massive Particles (WIMPs). Going back to the early universe, these hypothetical massive particles would have been in thermal equilibrium with the rest of the universe possessing an equal rate of pair production and annihilation. Eventually as the universe expanded and cooled production happened less often, and once the horizon distance of the universe was comparable to the annihilation mean free path the annihilation processes halted as well. The result was a “frozen” relic density that we can attempt to determine today. One of the early appeals of this process for WIMPs is that if one has thermally averaged annihilation cross-sections on the order of the weak-interaction

scale ($\langle \sigma v \rangle \sim 10^{-26} \text{ cm}^3/\text{s}$) this could produce the observed relic density for WIMPs (often referred to as the “WIMP-miracle”) in range of 1-1000 GeV/c^2 . This “miracle” does have some disagreement with experimental results for low-mass WIMPs ($m_\chi < 10 \text{ GeV}/c^2$) [36], but the parameter space to search in is fairly large still.

Most WIMP models arise from various supersymmetry (SUSY) extensions to the standard model. A full treatment of SUSY is beyond the scope of this work, but [37, 38] provide a good introduction. Some SUSY models contain a conserved quantity called R-parity, defined as:

$$P_R = (-1)^{3(B-L)+2s} \quad (1-2)$$

Here B, L, and s represent baryon number, lepton number, and spin respectively. Standard model particles are assigned an R-parity of 1, and SUSY particles an R-parity of -1 by convention. If R-parity is indeed conserved, then every SUSY decay must result in a final state that contains other SUSY particles. Eventually these would decay to a lightest supersymmetric particle (LSP), which would then be stable. The simplest SUSY model, the minimal supersymmetric standard model (MSSM) the LSP is taken to be a neutralino (χ), however the discovery of a Higgs mass at $126 \text{ GeV}/c^2$ puts some of the MSSM parameter space into tension. Various extensions to MSSM exist however, such as next-to MSSM (nMSSM) [39] and phenomenological MSSM (pMSSM) [40, 41] which can allow for a variety of WIMP mass ranges, in particular low-mass WIMPs. For the rest of this work, I will be focusing on WIMP dark matter and how it may be detected.

Low-Mass Dark Matter. It is worth a brief mention that in recent years a larger emphasis has started to be placed on low-mass variants of dark matter. This is mostly due to direct-detection experiments, such as DAMA/LIBRA [42], CoGeNT [43], CDMS II Si [44], and CRESST II [45] observing hints of a signal in the low-mass region. Typically the mass-range is around $10 \text{ GeV}/c^2$ and the events seen are near threshold. They remain interesting, but are in tension with other results as will be discussed further in Section 1.3.3. Models of dark matter to explain low-mass WIMPs can exist in the SUSY

extensions as mentioned earlier, and also from asymmetric dark matter (ADM) [28, 46–49]. ADM is interesting theoretically because it proposes that the matter-antimatter asymmetry observed in the universe is related to dark matter, and gives rise to $m_\chi \sim 5\text{--}15$ GeV/c^2 without relying on the WIMP-miracle.

1.3 Detection of Dark Matter

Thus far we have outlined theoretical motivations for dark matter and possible models to explain it, but none of them are useful unless we can experimentally verify dark matter exists. In general there are three ways to do this: WIMP creation in particle accelerators, indirect detection by WIMP-WIMP annihilation products, and direct detection by WIMP-nucleon scattering. The standard way to present experimental results is to use a WIMP-nucleon cross-section vs WIMP-mass plane as seen in Figure 1-8.




Figure 1-4. The three possible WIMP detection methods. Based on some interaction between WIMPs (χ) and standard model particles (ψ) various rotations of the interaction diagram give rise to WIMP creation, WIMP annihilation, and WIMP-particle scattering. In each diagram time is assumed to flow from left to right.

1.3.1 Collider Production

Particle accelerator experiments may be able to produce WIMPs depending on the cross-section of WIMP interactions and the WIMP mass. Since WIMPs are weakly interacting particles, a collider search would not see a WIMP directly, but rather the absence of energy and momentum associated with the creation of the WIMP. Generally collider search experiments will examine missing energy distributions and compare them

to expected standard model background processes. Dark matter searches at particle accelerators are not very competitive at high masses, but are able to probe to much lower masses [50] than direct detection experiments due to energy threshold effects, however some model dependence does enter into these results, via the interaction cross-section, which will control the rate of WIMP creation expected.




Figure 1-5. Spin-independent WIMP limits from the ATLAS collaboration in comparison to several direct-detection experiment results. At low masses collider experiments can probe regions of parameter space that direct-detection experiments cannot. Reprinted by permission from Aad,G. 2015. The European Physical Journal C (Page 299, Figure 12a). Springer Publications. [50].

1.3.2 Indirect Detection

Indirect detection experiments are also possible if WIMP-WIMP annihilation produces SM particles. These searches are performed by observations and typically the signal is an excess of gamma rays produced by the annihilation itself, or a result of other secondary (SM) particles undergoing annihilation. For these types of searches one needs to avoid searching in regions where one might expect an abundance of gamma sources from standard astrophysical sources (e.g., pulsars), and ideally one would like a

high density of dark matter to make annihilation more likely. The Fermi Gamma Ray Space Telescope (FGRST or Fermi for short) is a gamma ray observatory that is well suited to perform these types of searches. Some searches near the galactic center by Fermi have provided hints [51–54] of a gamma ray excess that could be interpreted as a low-mass WIMP of order a few tens of GeV/c^2 , but the center of the galaxy is messy and the expected backgrounds from astrophysical processes are not fully mapped or understood as of yet.

Dwarf spheroidal galaxies (dSphs) provide an excellent alternative target for this type of search however. Based on observations of luminous matter content and constituent motion, they are expected to have a high relative abundance of dark matter and so should be more likely to have annihilation happen. Additionally, dSphs’s are not expected to have an abundance of gamma-ray producing astrophysical sources. Searches from dSphs sources [55, 56] have not turned up any evidence of dark matter, and put constraints on WIMPs below $20 \text{ GeV}/c^2$. It should also be noted that indirect searches also contain some model dependence that influences what products arise from WIMP-WIMP annihilation and with what ratios.

1.3.3 Direct Detection

Direct detection experiments aim to observe the interaction of a WIMP and regular matter in some sort of detector. These types of search experiments look for WIMP-nucleon scattering, and are sensitive to a wide variety of different WIMP models. When a WIMP interacts with the target material, it will deposit some amount of recoil energy, which is measured by the detector. In practice many different recoil energies can be deposited, and the experiment would aim to use both the measured energy and number of events to infer information about the WIMP scattering rate. This rate is dependent on properties of the WIMP and target material and can be used to determine constraints on the mass and interaction cross-section. There are a wide variety of direct detection experiments such as XENON 100 [57], LUX [58], CRESST II [45], CoGeNT [43],

DAMA/LIBRA [42], SuperCDMS [59], and many more. A summary of various approaches to direct detection methods can be found in Reference [60].




Figure 1-6. Several different direct detection experiments organized by the type of energy measurements they perform. Experiments that appear to be between two types of energy measurements perform a simultaneous measurement of the types they are between. This is useful for improving background discrimination. Experiments sensitive to annual modulation are highlighted in blue. Courtesy of Tarek Saab[60].

1.4 Direct Detection Considerations

1.4.1 Interaction Rate

The rate at which dark matter interacts in a particular detector is the result of both astrophysical and particle physics properties. On the astrophysical side, the local density of dark matter will play a somewhat obvious role in the observed event rate. In the standard halo model, the dark matter density profile is assumed to be an isothermal, isotropic spherical halo with $\rho(r) \sim r^{-2}$. The differential density can be expressed as:

$$\rho = \rho_0 f(\vec{v}, \vec{v}_e) d^3v \quad (1-3)$$

Where $\vec{v}$ and $\vec{v}_e$ are the velocity of the WIMP on the target and the velocity of Earth relative to the local WIMP distribution. The constant ρ_0 ($\sim 0.3 \text{ GeV}/c^2/\text{cm}^3$)¹ is the average WIMP density at the region of the detector.[61, 62] The velocity distribution is most often a simple Maxwell-Boltzmann (MB) distribution that has a kinematic cutoff above the escape speed ($v_{esc} \sim 544 \text{ km/s}$):

$$f(\vec{v}, \vec{v}_e) = \alpha e^{-((\vec{v} + \vec{v}_e)/v_0)^2} \Theta(v_{esc} - v) \quad (1-4)$$

Here α is a normalization constant so the probability distribution integrates to unity, and v_0 is the halo root-mean square (RMS) circular speed at the sun ($\sim 220 \text{ km/s}$). Typically when one thinks of a MB distribution, one associates a temperature T with it. In reality the RMS speed term has swept this under the rug so to speak. Recall that for a MB distribution the most probable speed is found from $v_p^2 = 2kT/m$. The RMS speed is then related to this as $v_0^2 = 3v_p^2/2 = 3kT/m$. The differential density can be used to express the differential event rate[63]:

$$\frac{dR}{dE_r} = N_T \frac{\rho_0}{m_\chi} \int_{v_{min}}^{v_{max}} \frac{d\sigma}{dE_r} v f(\vec{v}, \vec{v}_e) d^3v \quad (1-5)$$

Where N_T is the total number of target nuclei ($N_T = N_A/A$), and $d\sigma/dE_r$ is the differential cross-section. The differential cross-section encodes some nuclear physics as it depends on the mechanism that WIMPs interact with various nuclei. Typically the differential cross-section is written in terms of a point-interaction cross-section and a nuclear form-factor

$$\frac{d\sigma}{dE_r} = \frac{\sigma_0}{E_r^{max}} F^2(q) \quad (1-6)$$

¹ There is considerable uncertainty over ρ_0 as its value depends on the parameterization of a WIMP halo one uses as a model, and if one accounts for dark matter streams which may give localized density fluctuations. This factor appears in every experiment's result however so it will only affect the absolute scale, not the relative differences between WIMP-search experiment results.

Where q is the momentum transfer, given by $q = \sqrt{2m_N E_r}$. The usual form factor chosen for spin-independent interactions is a phenomenological Helm form factor that arises from convolving a hard sphere with a gaussian to give a “blurry” sphere:

$$F(q) = \frac{3j_1(qr_0)}{qr_0} e^{-\frac{1}{2}(qs)^2} \quad (1-7)$$

Here j_1 is a spherical Bessel function of the first kind, s is a nuclear skin depth, and r_0 is the effective nuclear radius. The values for r_0 and s are determined based on the particular nucleus under consideration. The point-interaction cross-section for spin-independent WIMP-nucleus interactions can be expressed in terms of a spin-independent cross-section term as follows:

$$\sigma_0 = \frac{4\mu_{\chi,N}^2}{\pi} (Zf_p + (A - Z)f_n)^2 \quad (1-8)$$

The term $\mu_{\chi,N}$ is the reduced mass of a WIMP-nucleus system, f_p and f_n are coupling constants for protons and neutrons respectively, and Z and A are the atomic number and mass respectively. Typically it is believed that $f_p \sim f_n$, in which case an A^2 enhancement arises in Equation 1-8, making the element a detector is constructed out of important. In cases where the coupling constants are allowed to vary, it has been observed that if $f_n/f_p \sim -0.7$ some reconciliation between exclusion limits and closed contours set by different experiments can occur [29, 64]. The value for E_r^{max} can be determined from kinematic considerations. Since the incoming WIMP is a neutral SUSY particle, the recoil energy it deposits can be written as a simple elastic scattering process:

$$E_r = \frac{\mu_{\chi,N}^2 v^2}{m_N} (1 - \cos(\theta)) \quad (1-9)$$

From here we pull out the maximum recoil energy deposited as

$$E_r^{max} = \frac{2\mu_{\chi,N}^2 v^2}{m_N} \quad (1-10)$$

This can be rearranged to give the minimum possible speed with which a WIMP-nucleus interaction can deposit a recoil energy E_r :

$$v_{min} = \sqrt{\frac{E_r m_N}{2\mu_{\chi,N}^2}} \quad (1-11)$$

At this point our expression for the differential event rate can be expressed as:

$$\frac{dR}{dE_r} = \frac{N_T \rho_0 m_N \sigma_0}{2m_\chi \mu_{\chi,N}^2} \int_{v_{min}}^{v_{max}} \frac{f(\vec{v}, \vec{v}_e)}{v} F^2(q) d^3v \quad (1-12)$$

Typically results are expressed in terms of the spin-independent WIMP-nucleon cross-section and not the WIMP-nucleus cross-section. There is a simple relation between the two:

$$\sigma_0 = (Zf_p + (A - Z)f_n)^2 \frac{\mu_{\chi,N}^2}{\mu_{\chi,p}^2 f_p^2} \sigma_{SI} \quad (1-13)$$

Where $\mu_{\chi,p}$ is the WIMP-nucleon reduced mass, assuming the nucleon is a proton. If $f_p \sim f_n$, the expression simplifies greatly:

$$\sigma_0 = \frac{\mu_{\chi,N}^2}{\mu_{\chi,p}^2} A^2 \sigma_{SI} \quad (1-14)$$

Inserting this into the expression for the differential event rate yields:

$$\frac{dR}{dE_r} = \frac{N_T \rho_0 m_N \sigma_{SI}}{2m_\chi \mu_{\chi,p}^2} A^2 \int_{v_{min}}^{v_{max}} \frac{f(\vec{v}, \vec{v}_e)}{v} F^2(q) d^3v \quad (1-15)$$

In Figure 1-7 one can see the impact of the various physical terms inside the differential recoil rate. For high WIMP-masses the A^2 enhancement allows for the recoil rate around a few tens of keV to be much higher in heavier target nuclei. For low-mass WIMPs the story becomes a bit more subtle due to the kinematic terms (Equation 1-10) limiting the maximum possible recoil energy allowed. Higher target masses suppress the maximum recoil energy, and as such the lighter targets have higher rates.

A few other important aspects of the differential recoil rate to note here is that it is generally a featureless (quasi-)exponential shape which means there is no spectral shape one can attempt to search for, and that it need not be constant in time. Since the earth is




Figure 1-7. The differential recoil rates for various target masses are shown for two different WIMP masses. Here it is assumed that the spin-independent cross-section is 10^{-44} cm^2 , $\rho_0 = 0.3 \text{ GeV}/c^2/\text{cm}^3$, and standard assumptions of $v_{esc} = 544 \text{ km/s}$, $v_0 = 220 \text{ km/s}$, and $v_e = 232 \text{ km/s}$. Courtesy of K. Schneck [65]

in motion around the sun, and the sun is in motion around the center of the galaxy, there is an expected “wind” of WIMPs that has an annual modulation. For half of the year the Earth is moving into this wind and interaction rates will increase and for the other half the earth is moving with the wind, lowering the event rate slightly. DAMA/LIBRA is notable for claiming they see this annual modulation [42]. While there is clearly an annual modulation it is not yet ruled out that it could be some background that modulates.

1.4.2 Halo Dynamics

A brief comment about the velocity distribution is worth mentioning here. The standard assumption is that a MB distribution suffices to describe the WIMP halo distribution but this is perhaps too much of an idealization. As discussed in [66, 67], N-body simulations of dark matter halos are not best described by the standard halo model (SHM) of an isotropic, isothermal spherical halo. A popular alternative density profile is the Navarro-Frenk-White (NFW) profile [68] which has a density profile of the

form:

$$\rho(r) = \frac{\rho_0}{\left(\frac{r}{R_s} \left(1 + \frac{r}{R_s}\right)\right)^2} \quad (1-16)$$

Here R_s is a scale length that can be adjusted for different halos. In this particular model the RMS speed is no longer exactly equal to v_0 as was the case for the SHM but rather is only 7/8 as large. However other profile models, such as the Einasto profile² appear to have better agreement between data and models.[69] Further complicating the dark matter halo distribution are local perturbations that can arise from galactic mergers. This can result in streams of dark matter that have a local over-density (or under-density) that could result in a significant modification to the differential recoil rate.

Even in the SHM care must be taken in evaluation of the integral over the velocity distribution. This subtlety is explored in detail in Sebastian Arrenberg's thesis[70], and here I will provide only a summary of the results. It may be tempting to set $v_{max} = v_{esc}$ but this is not quite appropriate. First we can re-write the 3-dimensional velocity distribution integral by evaluating it over all angles. Observe that $d^3v = v^2 dv d(\cos \theta) d\phi$ giving:

$$\int \frac{f(\vec{v}, \vec{v}_e)}{v} d^3v = 2\pi \int d(\cos \theta) \int v e^{-(v^2 + 2vv_e \cos \theta + v_e^2)/v_0^2} dv. \quad (1-17)$$

Following the discussion in [70], imposing a constraint of $|\vec{v} + \vec{v}_e| \leq v_{esc}$ leads to the following form for the maximum WIMP speed:³

$$v_{max}(\theta_g) = \sqrt{v_{esc}^2 - v_e^2 (1 - \cos \theta_g)} - v_e \cos \theta_g \quad (1-18)$$

² $\rho(r) \sim e^{-Ar^\alpha}$

³ To arrive at this expand $|\vec{v} + \vec{v}_e|^2 \leq v_{esc}^2$ and let $v \rightarrow v_{max}$ and take the equality limit.

Where θ_g is the scattering angle in the galactic rest frame. This leads to three cases that the velocity distribution integral can be evaluated under:

$$\begin{aligned}
0 &\leq v_{min} \leq v_{esc} - v_e \\
v_{esc} - v_e &\leq v_{min} \leq v_{esc} + v_e \\
v_{esc} + v_e &\leq v_{min} \leq \infty
\end{aligned} \tag{1-19}$$

The third case trivially evaluates to 0 since the v_{min} is larger than the imposed kinematic constraint that requires scattering speeds to be less than v_{esc} .

For $0 \leq v_{min} \leq v_{esc} - v_e$ the velocity distribution integral yields:

$$\frac{\sqrt{\pi}}{4} \frac{v_0}{v_e} \left(\operatorname{erf} \left(\frac{v_{min} + v_e}{v_0} \right) - \operatorname{erf} \left(\frac{v_{min} - v_e}{v_0} \right) \right) - e^{-v_{esc}^2/v_0^2} \tag{1-20}$$

And for $v_{esc} - v_e \leq v_{min} \leq v_{esc} + v_e$:

$$\frac{\sqrt{\pi}}{4} \frac{v_0}{v_e} \left(\operatorname{erf} \left(\frac{v_{esc}}{v_0} \right) - \operatorname{erf} \left(\frac{v_{min} - v_e}{v_0} \right) \right) - \frac{v_{esc} + v_e - v_{min}}{2v_e} e^{-v_{esc}^2/v_0^2} \tag{1-21}$$

1.4.3 Backgrounds

Backgrounds are a major problem for direct detection experiments since WIMP interactions are expected to be rare and require sensitive detectors. In general, sources of background can come from Compton scattering of photons, cosmic-rays, and environmental sources of radioactivity either in situ or on the experimental apparatus itself. As a result of this many experiments employ active and passive shielding techniques to either veto background events, or attenuate background rates to be as small as possible, and often are situated deep underground to provide an overburden against cosmic rays.

The necessity to eliminate backgrounds can be illustrated in a fairly typical example involving a simple counting experiment with an observed number of signal events S , B background events, an exposure of $\mathcal{E} = MT$, a rate $r = S/\mathcal{E}$, and by observing that the sensitivity is then proportional to the rate (from Equation 1-15). A more complete discussion can be found in [71], with the results summarized here:

- **Background Free:** In the limit of a background free experiment, a 90% Poisson upper limit can be placed on the rate as $r = 2.303/\mathcal{E}$ showing that the sensitivity then is proportional to the exposure, $\sigma \sim \mathcal{E}^{-1}$.
- **Background Subtraction:** If an experiment observes B background events, but has an excellent understanding of the background, it can be subtracted. Since the statistical uncertainty for Poisson statistics goes as $\sigma_B = \sqrt{B}$ and it is expected that the number of background events is directly related to the exposure, we find $\sigma_B \sim \sqrt{\mathcal{E}}$. Since the number of signal events is what is left over from subtracting the background, the uncertainty in the signal will scale in relation to σ_B , indicating that the overall sensitivity for this type of experiment will scale as $\sigma \sim \mathcal{E}^{-1/2}$
- **Background Limited:** The worst state an experiment could be in is to be limited by background. While the statistical uncertainty on backgrounds scale as $\sqrt{B}$ the systematic uncertainties can scale as B itself since the systematic error represents a lack of knowledge about the background itself. Due to this scaling, the sensitivity is limited by B itself and no increase in exposure will improve the sensitivity since the systematic uncertainties will always dominate any statistical uncertainty in the background.

If an experiment can operate in the background-free regime we see it has the best opportunity to have a very sensitive experiment. Even a little background can cause the sensitivity to diminish, and even with background subtraction, larger exposures produce only modest gains compared to the background free state. Typically experiments are able to operate in a background free mode for some initial exposure and eventually become background limited after some time. Experiments can mitigate these to some extent by implementing methods to not only shield backgrounds but reject them as well. For SuperCDMS a dual measurement of ionization and phonon energy provides the ability to discriminate between electron recoils that arise primarily from backgrounds like γ 's, and nuclear recoils which arise from WIMP interactions with the target material.

1.4.4 Sensitivity Results

The results of WIMP-search experiments are presented in a plane that shows the WIMP-nucleon cross-section (usually spin-independent) as a function of WIMP-mass. An example of such a plane is shown in Figure 1-8. There are several aspects of results in this plane that are worth understanding.




Figure 1-8. Spin-independent WIMP-nucleon cross-section vs. mass plane with several experimental results. Closed contours represent discovery claims, whereas open curves are exclusion limits. An irreducible background from coherent neutrino scattering may become an issue in later G2 level experiments and beyond.

In this plane closed contours represent experimental results that appear to have seen evidence for WIMPs. Most experiments however take a more conservative approach and set exclusion curves which rule out the region of parameter space above them. While each curve is different, they do share common properties. At low WIMP masses they all rise rapidly, approaching a vertical asymptote. Likewise at high WIMP masses the curve again starts to increase. To understand why, we can look at the differential rate with respect to energy with a few simplifying assumptions to the integral in order to extract the limiting behavior. If we assume that the kinematic cutoff has been relaxed, and second assume a trivial case that the velocity of the earth relative to the local WIMP distribution is zero

our, the relevant part of the velocity integral from Equation 1–17 becomes:

$$\int_{v_{min}}^{\infty} v e^{-v^2/v_0^2} = \frac{v_0^2}{2} e^{-v_{min}^2/v_0^2} \quad (1-22)$$

With these simplifying assumptions, let's first consider the case where $m_\chi < m_p$. For these low-mass WIMPs both $\mu_{\chi,p} \sim m_\chi$ and $\mu_{\chi,N} \sim m_\chi$ so the differential rate is proportional to:

$$\frac{dR}{dE_r} \sim \frac{m_N}{m_\chi^3} \sigma_{SI} v_0^2 e^{-v_{min}^2/v_0^2} \implies \sigma_{SI} \sim \frac{m_\chi^3}{m_N} \frac{dR}{dE_r} e^{E_r m_N / 2 m_\chi^2 v_0^2} \quad (1-23)$$

However we must impose a constraint on E_r that it cannot be arbitrarily small. Rather, for low-mass WIMPs the smallest recoil energy measurable is going to be the energy threshold of the experiment, E_{th} . With this in mind, for a measured differential rate, the limit on the cross-section will rapidly increase as the WIMP-mass decreases. A low energy threshold will help mitigate this and is necessary to achieve low cross-section limits at low WIMP mass.

In the case $m_\chi > m_N$ we have $\mu_{\chi,N} \sim m_N$ and $\mu_{\chi,p} \sim m_p$. In this case, assuming no arbitrary upper limit on the measured recoil energy, the relationship between the differential recoil rate, cross-section and WIMP mass follows:

$$\frac{dR}{dE_r} \sim \frac{m_N}{m_\chi m_p^2} \sigma_{SI} v_0^2 e^{-v_{min}^2/v_0^2} \implies \sigma_{SI} \sim \frac{m_\chi m_p^2}{m_N} \frac{dR}{dE_r} e^{E_r / 2 m_N v_0^2} \quad (1-24)$$

This expression then shows that the cross-section will scale proportionally to the WIMP mass and is why the limit curves have a more gentle upwards slope after they reach their minimum sensitivity. At high WIMP masses the recoil energy becomes kinematically limited by the (relatively) low target mass of the nucleus. Heavier nuclei aid in searching for high mass WIMPs but as Equation 1–23 hints at, hurts low-mass sensitivity. The converse situation is true as well where light elements will be better suited for a low-mass WIMP search. However in this mode, it is more important to not only reduce background as much as possible, but also push its threshold as far down as possible (which often contradicts the requirement to reduce background), making it a difficult endeavor indeed.

CHAPTER 2

SUPERCDCMS AT SOUDAN

SuperCDMS Soudan is a dark matter search experiment¹ operating at the Soudan Underground Laboratory (SUL, or Soudan for short). SuperCDMS operated 15 germanium detectors called iZIPs (more on these in [3](#))², while using the same cryogenic infrastructure, electronics, shielding and radio-free (RF) room as CDMS II. These components are described in detail in [\[72\]](#), but will be discussed briefly in this chapter.

Soudan Underground Laboratory is located in an old iron mine, in remote northern Minnesota.³ The rock is extremely stable, and the dense overburden provides an excellent site for an experiment that needs to minimize cosmogenic background sources. The experimental cavern is located on Level 27, approximately 713 m below the surface, or 2090 meters water equivalent (m.w.e.) depth. In particular the muon flux is reduced by a factor of 1×10^5 compared to the surface. This is also accompanied by a reduction in secondary products from cosmic rays, most importantly neutrons. Since neutrons have a large mass and are neutral, they appear virtually identical to a WIMP when they interact in our detectors and are a background that must be reduced as much as possible. CDMS started its first run at Soudan in 2003 (Run 118) and has operated there for nearly 13 years, counting both CDMS and SuperCDMS operations. As of November 2015, the fridge is undergoing its final warmup and decommissioning will commence shortly thereafter.

2.1 Experimental Setup

The SuperCDMS detector setup is located inside an RF shielded clean room that is meant to protect against contamination from environmental sources of noise and background. The RF room has an clean anteroom for detector prep and other side

¹ The successor to the Cryogenic Dark Matter Search II (CDMS-II) experiment.

² interleaved Z-sensitive Ionization and Phonon sensors.

³ Which makes for cold winter shifts. A few shifts had -35C wind chills.




Figure 2-1. The muon flux as a function of depth underground is shown here. Soudan is has a muon flux that is approximately 1×10^5 times lower than the surface. The next generation of SuperCDMS will be located at SNOLAB in Sudbury, where the flux is reduced another 3 orders of magnitude.

activities to avoid unnecessary activity inside the RF room itself. To the other side of the RF room, and exposed to the mine itself, is a cryogenic support area (the “cryopad”). Many of the necessary pumps and gas handling control electronics are located here, as well as two external LN₂ cold traps for mixture cleaning. A second floor contains a level called the mezzanine, which contains air handlers for the RF room and the helium reliquifiers. On the other side, still on the second floor, is the electronics room which houses trigger logic boards (TLB) and receiver-trigger-filter (RTF) boards as well as computers which

control data acquisition and the experiment itself. The mezzanine level also has office space for CDMS physicists to work during the 8-10 hour shifts underground.⁴




Figure 2-2. Sideview of the experimental setup at Soudan. Not shown is the office space which spans the mezzanine level from the HVAC to the end of the electronics room and in this view would be out of the page. Courtesy of Jeff Filippini. [71].

2.1.1 Shielding

The outermost layer of shielding for the detectors is the muon veto. These are plastic scintillator panels arranged around the icebox housing such that they do not leave any exposed openings. The muon vetoes allow us to reject events that are coincident with a veto trigger. Internal to the icebox are alternating layers of polyethylene and lead. The outermost polyethylene layer is 40 cm thick and moderates neutron backgrounds to very low energies. The outermost layer of lead is 17.8 cm thick and provides shielding against photons, whereas the innermost layer (4.4 cm thick) is composed of ancient lead recovered from a sunken ship near France. This ancient lead is deficient in ^{210}Pb which has a long half-life, and provides shielding against the outer lead layer. The final inner layer of polyethylene measures 10 cm thick and moderates against very high energy neutrons that were not fully stopped by the other layers of shielding. In order to mitigate high levels of

⁴ A ping-pong table is also located on the mezzanine for the rare moments of free time, or when one needs to clear their head after hunting spurious sources of noise.

ambient radon in the mine cavern the boil-off from the LN₂ dewars is flushed through the gap between the icebox and shielding layers providing one extra type of passive shielding.




Figure 2-3. Views of the SuperCDMS Soudan shielding setup. A.) Picture of the passive shielding elements along with their thickness. B.) Schematic cutaway view of the larger cryogenic system surrounding the icebox, which contains the detectors.

2.1.2 Cryogenics

The detector payload is cooled to a base temperature around 55 mK by an Oxford Instruments 400S dilution refrigerator with a cooling power of 400 μW at 100 mK. The detector assembly is thermally connected to the fridge via the c-stem in order to prevent contamination from the fridge itself. Additional cooling power is provided by an auxiliary Gifford-McMahon cryocooler that is heatsunk to the 4K and 77K stages. This keeps the time between liquid helium refills of the fridge sufficiently long as to not impact the livetime.[73] During the startup of SuperCDMS Soudan data taking with the new iZIP payloads, the time between these cryofills was around 21-22 hours. When a cryofill is initiated by the DAQ data taking must stop due to elevated noise and warming of the detector payload. We took advantage of this pause to perform an LED bake to re-neutralize the detectors, but this all takes away from experimental livetime. A typical cryofill interval lasts approximately 90 minutes, mostly due to the cooldown period.

In order to mitigate both the cost of liquid helium, and reduce interruptions to data taking, significant effort was spent retrofitting a helium recovery and reliquifier system into

the cryogenic system. This system takes helium boil off and liquifies it using a cryocooler on the mezzanine where it is then brought back to the bath through a transfer line. This system has 3 reliquifiers and is able to maintain the bath level for many weeks without the need for additional helium. Unfortunately this system introduced vibrational noise into the experimental setup, causing an increase in low frequency noise and slightly elevated base temperatures. The exact mechanism for this vibrational coupling was investigated but no method was discovered to eliminate the vibrational noise. It did appear that the transfer line itself introduced this excess noise (and this seemed sensible since it was a lengthy pipe), as the base temperature dropped when it was removed, but measurements with accelerometers did not reveal excess vibrations. Additionally the transfer line would occasionally develop an air leak or one of the cryocoolers would fail necessitating a quick turn around to get the system up and running again. When this happened the DAQ would automatically revert back to performing helium fills as needed. Due to the unpredictable nature of these intermittent failures and the fact helium dewars do not last forever, regular, though infrequent, shipments of helium were still required.

An additional source of vibrational noise arose from the Gifford-McMahon cryocooler as well. Over the duration of operations the cold head started to degrade and would cause a distinctive “chirping” and “whooshing” sound each cycle. Accelerometers were able to detect these chirps and were eventually put into the data stream so that they could be used to reject events at low energy during these noisy periods. Excess noise from the cryocooler slowly required raising trigger thresholds to avoid excessive triggering on noise.

2.1.3 Detectors and Towers

The SuperCDMS Soudan payload consists of 15 iZIPs (see Chapter 3) arranged in groups of 3 to form 5 towers as seen in Figure 2-4. Each detector is contained in its own copper housing, and each tower is constructed from ultra-pure copper, with the top and bottom detectors having an additional layer of copper as an endcap. Each detector is biased to $\pm 2\text{V}$ in standard operation mode and the resulting ionization signal readout

is performed by a FET amplifier (Section 3.2.2) and the phonon signal is amplified via a SQUID (Section 3.6.3). These are compacted into a single SQUID + FET (SQUET) card that is then devoted to the read out of half of a detector’s ionization and phonon signals. This is an unfortunate constraint imposed by reusing the CDMS II hardware. CDMS II detectors had 4 phonon channels on a single side and 2 ionization sensors on the opposite for a total of 6 channels. The iZIP contains 4 phonon channels on each side and 2 ionization channels on each side, for a total of 12 channels necessitating two separate readout chains per detector and later in software the detector is reconstructed (more information on the detector channel layout is presented in Section 3.1). Each phonon channel (and analysis quantity) begins with a P and each ionization channel with a Q. The phonon channels on a given side are labeled A-D and the ionization channels are labeled with an I or O to denote whether it is the inner or outer channel respectively. The side on which the channel occurs is indicated by a 1 or 2 for the top or bottom side respectively. As an example, the phonon channel B on side 1 is called “PBS1”, and the outer ionization channel on side 2 is “QOS2”. After the signal is amplified by the SQUET it is sent further upstream via striplines to the e-box through the e-stem, where it will enter the warm electronics portion of the readout process.

Towers are prefixed with “IT” (for iZIP Tower) and numbered 1 through 5. Individual iZIPs contain several different types of labels. They are numbered 1 through 3 in any given tower from top to bottom and prefaced with a “Z” so that the first detector in the first tower would be called IT1Z1. Additionally the detectors have numeric indices denoted by a number of the form 11xx as shown in Figure 2-4. They are also assigned a fabrication identifier that was used during detector testing prior to selection for Soudan.

2.1.4 Special ^{210}Pb implanted detectors

As was discovered during CDMS II, ^{210}Pb is an important source of background to understand as it contributes to surface events (which do not have complete ionization signal collection). These were a dominant source of background leakage in CDMS II and




Figure 2-4. The SuperCDMS Soudan iZIP array is shown here. Each detector is located in its own housing, and rotated by 60 degrees from its neighbor.

influenced the design of the iZIP. The surface event rejection capabilities of the iZIP are discussed further in Section 4.2.2, and the details of the ^{210}Pb decay process can be found in Section 4.1 (or [74]). In order to better study the discrimination abilities of the iZIP on this source of background, two detectors were implanted with a ^{210}Pb source. This source was created by exposing two silicon wafers to a 5 kBq ^{226}Rn source while sealed in an aluminum box. The resulting decay produced ^{222}Rn which decays to ^{210}Pb . The silicon wafers were etched and one implanted on the top surface of IT3Z1 and the other on the bottom surface of IT3Z3 in a uniform fashion, each to a depth of approximately 58 nm. This split between top and bottom side allows for examination of both the electron and the hole collecting sides of the iZIP to be studied. The sources each generate ~ 130 events per hour. Since these detectors contain sources they are not able to be used for low mass WIMP searches and at the very least only about half can be used for a high mass WIMP search.

2.1.5 Readout Electronics and Data Acquisition

The warm electronics part of the readout chain consists of three different steps for filtering and issuing triggers for the DAQ to record.

- FEB: The Front-End-Boards, or FEBs, are located in the RF room adjacent to the e-stem and are connected via 50-wire cables to the detectors.⁵ The FEBs were custom built for CDMS II and contain various control circuitry for the ionization, phonon, LED, charge bias, and SQUID biasing elements of detector control and are stored in racks.
- RTF: The Receiver-Trigger-Filter, or RTF, boards prevent random noise from dominating the data acquisition by filtering and amplifying signals from the FEBs. It is on these boards that one can set a trigger threshold so that only incoming signals above some particular amplitude will be accepted. The original CDMS II RTF boards were optimized for CDMS II style phonon signals and as such have a bandpass filter with a low-high of 900 Hz - 18 kHz. The iZIP pulses are substantially different and contain more power at lower frequencies, and a study was performed to optimize these which is described in more detail in Section 6.1. The actual trigger is done on the summed phonon trace and is referred to as a Plo trigger. Since each iZIP is readout as 2 half-detectors, each iZIP gets 2 RTF boards which allows for two operation modes for triggering. The standard method is referred to as PloOR, in which either RTF board may issue a valid trigger. For some detectors with noise issues a more strict PloAND may be used in which both RTFs must issue a trigger. The PloAND may be applied to only one RTF and have the other one remain in PloOR resulting in a more fine scaled tuning. For diagnostic purposes the RTF boards also contain blinking lights that allow one to gauge by eye how frequently they are triggering.
- TLB: The trigger logic board, or TLB, will accept valid triggers from the RTF or the muon veto system, and additionally can issue random triggers to collect a sample of noise. If a trigger is issued, data is read out around and stored for future analysis, and the trigger system is then rearmed to wait for the next type of trigger to occur. For calibration datasets the trigger rate is quite high, around 20 Hz, but for WIMP search mode the trigger rates are kept as low as possible to increase livetime (typically around 0.3 Hz).

⁵ Great care must be taken not to damage the pins when swapping cables, and one must be sure they lead to the correct detector and FEB. This is easier said than done and a couple of times detectors were misconnected. This was thankfully caught prior to any data taking during continuity checks.

From time to time other aspects of the warm electronics may require intervention to bring back to operation. Some situations where this is necessary include failures of power supply crates, Struck digitizers, or even RTF boards themselves.⁶

After the trigger has been issued data is digitized and recorded in several locations by the DAQ. A small buffer is present in the mine, but otherwise data is copied to a surface control building in Soudan, and sent to Fermilab where it is also stored and processed for data quality. Since the trigger rate is so high, and a lot of information is read out per event, the hardware cannot process calibration data fast enough for data quality purposes if several hours of calibration data are taken in full. Instead only towers in which at least one RTF issued a trigger are read out, which helps keep the DAQ happy and file sizes a bit smaller. This operation mode is referred to as selective-readout (SR) mode.

2.2 Data Quality and Processing

Data is grouped into several different types of discrete collections during the data taking process. The largest is referred to as a “Run” and is generally defined by the operational status of the cryogenics. The next, and most commonly used group, are the data series, or series for short.⁷ A series is determined by when the detectors are biased and lasts until they are grounded and flashed with LEDs to be re-neutralized. Series are labeled with a timestamp that indicates when they began and prefixed with a “1” to denote the data was taken at Soudan (e.g. 1120301_0523 would indicate a series that started on March 1 2012 at 5:23 AM central time). A data series lasts for 3 hours as it

⁶ On several shifts I have had to power down the chain, an apparently risky process as sometimes a power supply crate (which are custom built, and of very finite supply) does not turn back on!

⁷ Sometimes “run” is used interchangeably with “series” by the collaboration. The context is key to differentiate between a “Run” and a “run”.

was determined during commissioning that was the longest time most iZIPs could be operated before suffering charge stability problems.⁸

Internal to each series are smaller blocks called file dumps. Each of these are discrete units that can be processed separately downstream and contain at most 500 raw events. A raw event consists of the readout information for every phonon and ionization channel, trigger information, and numerous other bits of information deemed useful (such as base temperature, bias voltage, accelerometer information, veto status, etc). The first and last dump of every series consists of consecutive random triggers to get a good characterization of the noise on each detector.⁹ The rest of the dumps contain a mixture of triggered events and approximately 10% random triggers throughout. The penultimate dump may not contain 500 events as it will end whenever the DAQ issues a stop command.

As mentioned once we take data it is processed quickly for data quality purposes. As a data series is underway, an online DQ Summary page provides any collaboration member the ability to assess the series' health via several metrics. The DQ page first presents a user with a list of series to choose from, indicating the current LiveTime, EventTime, number of events, data type, and some automatically determined metrics as to whether the trigger rates are ok, and if the phonon offsets are behaving. The user can enter into a series and see various Level 1 or 2 diagnostic plots. I will briefly outline the four Level 1 diagnostic plots, plus the Level 2 plots.

- Noise Plots: Noise PSD plots for each detector (organized by tower) are constructed from the start of series randoms and are a great way to verify all channels are behaving properly. If a channel is about to fail, or has failed already, the PSD will appear noticeably different than the others. Occasionally an actual pulse will sneak

⁸ For calibration runs, data is taken for periods of 25 minutes biased with a brief 5 minute grounding. This repeats for up to 3 hours.

⁹ For the time period prior to 1120418_1051 data was accidentally taken without end of series randoms. Additionally it was discovered there was an unusual DAQ bug with calibration data that was taking end of run randoms in selective readout mode!

into the randoms and cause an alarming looking trace, but generally this is ok. A pulse rejection algorithm is active on all detectors except ones where it causes problems.

- **Detector Plots:** These are generated from very roughly processed data quantities and organized the same way as the Noise plots. They consist of basic plots of distributions of certain useful quantities that can provide a metric as to how well the phonon and ionization channels are performing when recording actual events. Some of the plots show a quantity plotted vs time so that one can see if an unwanted time dependence has shown up.
- **Veto:** The veto plots show the amplitudes of each of the veto panels and all should mostly look the same with low voltage.
- **Trigger Plots:** Two different plots are shown on this page. One is a bar chart for each RTF of a detector showing the number of triggers issued. The other is a chart showing the overall trigger rate as a function of time and is useful for identifying periods of trigger bursts or problems with an RTF triggering too infrequently.
- **Level 2:** These are more reconstructed quantity plots that provide a quick snapshot of how well certain fiducializing quantities are behaving over time. Chiefly these show how our radial and z location estimators behave over time.

Eventually when enough data has been collected a more robust processing is performed using a sophisticated package called `cdmsbats`.¹⁰ This software package reads in raw data files and performs data reconstruction to get energy and timing information. These quantities are referred to as Reduced Quantities, or RQs. A second round of scripts runs through a package called `BatCalib` to construct more sophisticated energy, position, and other event type quantities referred to as RRQs. Specific details about the charge and phonon reconstruction algorithms are described in the next chapter, but underlying each method is the Optimal Filter (OF) process. Specific details about OF theory can be found in the appendix of J. Filippini's thesis [71] or for more detail about improvements since then, in Rito Thakur's thesis.[7] In general an OF allows one to estimate the amplitude of a signal by maximizing the signal to noise ratio. To do this requires a template of the

¹⁰ This processing package has many subsets such as `BatROOT` which produces RQs and `BatCalib` to produce RRQs.

signal and of the noise. Templates for the phonons are derived empirically from a sample of good bulk electron recoil events on a detector by detector basis and are used for every event, and likewise for the ionization signal.¹¹ The noise templates come from the sample of random triggers taken at the beginning or end of each series, and are assumed to be constant throughout the series. The signal (our data) can then be written as the sum of noise and the template in the time domain:

$$s(t) = ap(t) + n(t) \quad (2-1)$$

Here $p(t)$ is the normalized pulse template and a represents the unknown amplitude. In order to quickly determine the best fit, and since the noise is uncorrelated in frequency space, these signals are Fourier transformed and a χ^2 quantity is minimized:

$$\chi^2(a, t_0) = \sum_{\nu} \frac{\left| \tilde{S}_{\nu} - ae^{-2\pi i \nu t_0} \tilde{P}_{\nu} \right|^2}{\tilde{J}_{\nu}} \quad (2-2)$$

Here t_0 is a parameter that denotes the start time of the pulse and is allowed to wander in a small window, and $\tilde{J}_{\nu}$ is the noise power spectrum derived from $n(t)$. The minimum of this χ^2 gives an estimate for the best fit amplitude as:

$$\hat{a}(t_0) = \frac{\sum_{\nu} e^{2\pi i \nu t_0} \tilde{S}_{\nu} \tilde{P}_{\nu}^* / \tilde{J}_{\nu}}{\sum_{\nu} \frac{|\tilde{P}_{\nu}|^2}{\tilde{J}_{\nu}}} \quad (2-3)$$

The quantity $\tilde{P}_{\nu}/\tilde{J}_{\nu}$ can be viewed as a filter that de-weights frequencies with large noise in $\tilde{S}_{\nu}$, which is why this is called an optimal filter. The story is a bit more complicated than just maximizing $\hat{a}$ however given the existence of the time offset. We would like to find the optimal time offset as well and so we must find $\partial\chi^2/\partial t_0 = 0$ as well. Applying the

¹¹ There is considerable room for improvement right here. Not only are the templates themselves derived from summed traces, but they do not capture the sharper peak of surface events. A wide variety of different templates and multiple template OFs are under investigation for future analyses.

derivative we get the following equation:

$$\frac{\partial \chi^2}{\partial t_0} = -2a \sum_{\nu} 2\pi\nu \left(e^{2\pi i \nu t_0} \frac{\tilde{S}_{\nu} \tilde{P}_{\nu}^*}{\tilde{J}_{\nu}} \right) \quad (2-4)$$

This equation is highly nonlinear and has no analytical solution unfortunately. But luckily a convenient computational shortcut exists. Observe that the term in parenthesis contains some similar terms to $\hat{a}(t_0)$ from Equation 2-3. If we take $\partial \hat{a}(t_0)/\partial t_0$ we see

$$\frac{\partial \hat{a}(t_0)}{\partial t_0} = \frac{\sum_{\nu} 2\pi i \nu e^{2\pi i \nu t_0} \frac{\tilde{S}_{\nu} \tilde{P}_{\nu}^*}{\tilde{J}_{\nu}}}{\sum_{\nu} \frac{|\tilde{P}_{\nu}|^2}{\tilde{J}_{\nu}}} \quad (2-5)$$

This now shows that by a nice “coincidence” the partial derivative of the optimal amplitude estimator with respect to the time offset is proportional to the χ^2 derivative as well

$$\frac{\partial \chi^2}{\partial t_0} \sim \frac{\partial \hat{a}(t_0)}{\partial t_0} \quad (2-6)$$

Since the constants don’t matter (we’re looking for the zeros) we can find the minima of the χ^2 by searching for the t_0 that maximizes the amplitude (recall there is an overall minus sign in front of Equation 2-4), and since we have an expression for $\hat{a}(t_0)$ in Equation 2-3 this is quick.

This can be generalized to cases where one has N channels with M templates in which we would like to find the best amplitudes and delay times to fit the ensemble. The process involves solving a matrix equation with cross-terms as described more in [7]. For SuperCDMS Soudan this shows up for the ionization channels in which we attempt to handle the cross-talk between the inner and outer ionization channels on a given face by using a 2x2 OF. The resulting generalization of Equation 2-1 is:

$$\begin{pmatrix} s_i(t) \\ s_o(t) \end{pmatrix} = \begin{pmatrix} Q_i(t) & Q_{io}(t) \\ Q_{oi}(t) & Q_o(t) \end{pmatrix} + \begin{pmatrix} n_i(t) \\ n_o(t) \end{pmatrix} \quad (2-7)$$

Here Q_i and Q_o are the ionization channel templates, and the cross terms represent the expected shape of the crosstalk on the first subscript as a result of the second index (e.g., Q_{io} is the cross talk seen on the inner channel due to the outer channel).¹²

An additional optimal filter is used on the total phonon signal called the Non-Stationary Optimal Filter (NSOF, or NF). Because there are position dependent effects during the initial portion of a phonon pulse which we do not yet have a good handle on, the regular OF algorithm is not as optimal as it could be. The non-stationary OF, developed by Matt Pyle, takes the position dependent effects and encodes them as off-diagonal terms in a covariance matrix between different frequencies, essentially treating the position dependence as non-stationary noise. These are then de-weighted during the computation of the χ^2 which uses the position-independent portion of the phonon pulse instead. The net result is improved energy resolution, but this is slow due to the requirement of inverting the rather large covariance matrix, so we only implement it above a certain (low) energy threshold.

Data Files. I wish to spend a moment here to mention the fashion in which SuperCDMS Soudan stores the processed data files. The bulk of `cdmsbats` is written in C++ and interfaces with ROOT to write the outputs of BatROOT and BatCalib into various ROOT files. A special set of MEX functions are used to load these into MATLAB (the current primary analysis software). The initial step is to select which series to process and group these together into processing lists. Each dump of a given series is processed on a grid at Fermilab, and then combined into common root files based on the series they belong to. These are referred to as “merged” data files. Then each series inside a given processing list is merged together into a “supermerged” file. This mainly keeps the number of ROOT files we need to manage low and for MATLAB minimizes the number of ROOT

¹² In CDMS II this cross-talk effect was about 5% however for the iZIP it has been reduced to $\sim 1\%$.

files one must open when accessing a particular quantity. Inside a given ROOT file there is a directory for either RQs or RRQs depending on the class. This directory contains multiple “trees” of the form zipN or calibzipN (where N is the detector index 1 - 15). Each tree then contains multiple branches, one for each RQ or RRQ that then contains the actual data.

For the low-mass WIMP search that (described in Chapter 6), a very common frustration was the extremely sluggish load time of the data. At the time the blinded WIMP search data contained approximately 23 million events, and the calibration dataset totaled around 100 million events. Accessing a single branch from the multiple ROOT files should not take a very long time, but it would often take a couple minutes to load one RQ or RRQ. When working with multiple quantities this would often mean one would write code, set it to run, and then take a very long coffee break.¹³

In between the low-mass WIMP search and the start of the high-mass WIMP searches, some time was spent diving into the mechanics of how `cdmsbats` created the ROOT files and it was discovered that a subtlety to how we were merging the files resulted in excessive fragmentation. During the initial processing when the raw files are merged together to create a single-series ROOT file, the (unmerged) dump ROOT files are very small. As each branch only contains at most 500 doubles the size is tiny. When ROOT writes to disk it does so by creating “baskets” which roughly indicates specific discrete IO processes. So if a branch has 20 baskets, the contents are located at 20 spots on the disk and require 20 disk IO processes. We used something called a “fast-merge” that essentially performs a raw copy when merging files. So as we progressed from unmerged to merged to supermerged files, the basketsize remained very tiny (32kB) and the number of baskets

¹³ The buffering process in the analysis package, CAP, mitigated this somewhat since after the initial load data was accessed from either memory or MATLAB mat files. But it was clear that since the load time from the mat files was so quick something was wrong with the ROOT files.




Figure 2-5. Data access times as a function of file size are shown for various data types. The integrated time to load a single RRQ for barium data was greater than a few minutes, making cut computation and data analysis very slow.

grew dramatically! After some experimentation (and an opportunity to talk to one of the developers of ROOT, P. Canal, at Fermilab) some code was created to control the number of baskets at the different processing steps. Now when constructing the data files, the merged files will contain a single basket, the supermerged will contain several larger baskets which are then looped over a second time with a fast-merge that also sorts the file by baskets so they are next to each other on disk (based on a branch by branch basis). The net result is that a full barium RRQ that used to take on the order of minutes now takes on the order of 15-30 seconds, a dramatic speed up!

2.3 Data Types

SuperCDMS Soudan did not take WIMP-search data 24 hours a day 7 days a week, even though it would be great for our exposure. Because the detector response is different for different types of recoils we must take calibration data to be able to understand the

full parameter space of the iZIPs. To model backgrounds due to bulk electron-recoils¹⁴ we use a calibration source that emits γ 's with spectral features. Our particular choice is ^{133}Ba of which two small sources are encapsulated at the end of thin metal rods which are inserted into two thin tubes past the shielding,¹⁵ one by each of the stems, allowing us to illuminate the towers. Barium calibration was performed for about 15 hours out of each week and was done during the day when we had mine access.

In order to understand the signal region of parameter space we also took data with a neutron source, ^{252}Cf , periodically. Neutron calibrations were performed much less frequently (typically every few months) due to its ability to activate the germanium in the detectors. The activated germanium increases the event rate due to internal electron capture as the activation products decay, so a few days after each ^{252}Cf calibration are left unblinded to study the activation peaks since it is not suitable for a WIMP-search.

Other special data taking modes occurred periodically. The primary two are veto pulser runs which flash LEDs in the veto panels to ensure they see signal with appropriate amplitude (this was done by the DAQ automatically during a cryofill, or manually once a week when the reliquifiers were operational). The other type of running mode consist of special random triggers which increased the duration of recorded pulses to allow for us to better sample the low frequency noise and were performed once a week generally.¹⁶

¹⁴ Events that scatter off of electrons, and primary type of recoil due to background.

¹⁵ Unfortunately one of these sources eventually fell off due to wear and tear around July 3, 2013. Even more problematic was that the location was inside the shielding. This required a brief warm up and partial removal the veto panels and shielding to retrieve it, but not before many clever attempts to fish it out were tried. This marked the end of Run 133.

¹⁶ The lowest frequency point is related to the inverse of the trace length. Initially the traces were lengthened to 26.2 ms corresponding to a lowest frequency of 38 Hz, but subsequent work extended this to ~ 104.9 ms with a lowest frequency of ~ 10 Hz.

2.4 Experiment Time Periods

During its operation SuperCDMS Soudan had three primary data taking periods. Run 133 (R133) occurred from March 2012 - July 2013, Run 134 (R134) from July 2013 - July 2014, and Run 135 (R135) from September 2014 - November 2015. R133 and R134 contain all of the standard operation WIMP-search exposure. R135 contains a special ^{133}Ba dataset taken with a modification¹⁷ to the RTF boards to allow for more ^{133}Ba statistics in energy ranges of interest. R135 also performed other special tests that were deemed to be too much of a risk to stable operation during WS data taking periods in R133 and R134. The combined datasets have their run numbers joined together such that R133+R134 is c334, and all 3 datasets are c335.¹⁸

Unfortunately during this long duration of running many changes to the detector behavior occurred, the most damaging of which were changes in the functionality of readout channels, both ionization and phonon based. Here I will attempt to give a brief overview of which detectors had problems, and when. For data processing reasons this was very important information to know so the correct OF algorithm could be run and the correct value of the status RQs could be assigned and so for the initial round of data processing for R133 and R134, these had to be tracked down since the information was often scattered in separate email threads and logs. This was collected into an experimental status configuration file in the cdmsbats package.

2.4.1 CDMSlite

An important operation mode that occurred periodically was CDMSlite.¹⁹ This special operation mode takes advantage of the Luke phonon process that occurs when

¹⁷ An energy window for the trigger.

¹⁸ The combined naming scheme was meant to be a shorthand for the combination of 33 and 34. R134+R135 does not fit into this pattern but could be called c345.

¹⁹ CDMS low ionization threshold experiment

Table 2-1. Detector Health List

Detector	Phonon Satus	Ionization Status
IT1Z1	Noisy SQUET on PAS2,PDS2	OK
IT1Z2	Shorts PAS2, PCS2	Shorts QIS1, QOS1
IT1Z2	Short PAS1	Short QIS1, QOS1
IT2Z1	OK	OK
IT2Z2	OK	OK
IT2Z3	Short PBS1, PDS1	OK
IT3Z1	OK	OK
IT3Z2	OK	OK
IT3Z3	Large baseline offsets on PCS1	OK
IT4Z1	OK	Short QIS1
IT4Z2	OK	OK
IT4Z3	OK	OK
IT5Z1	Short PAS1	Short QIS1, QOS1
IT5Z2	OK	Odd Noise QOS1
IT5Z3	Short PCS2	Short QOS1

charged particles drift through a crystal lattice at speeds greater than the speed of sound in the material (discussed in more detail in Section 3.6.2). The Luke phonon energy is proportional to the bias voltage across the crystal and so one would like to increase the bias voltage across the crystal to as large a value as possible to get a very large boost to the signal. Up until breakdown occurs, an increase in voltage will only increase the signal while the noise remains largely unaffected. Above the breakdown voltage the noise will degrade rapidly, greatly reducing the signal to noise ratio. Based on the results of scanning through various bias voltage levels, CDMSlite increased the standard bias from 4V to 70V on IT5Z2 for three running periods. CDMSlite R1 occurred from August through September of 2012[75], CDMSlite R2 occurred from February through July and September through November 2014²⁰, and R3 occurs during the R135 period.

Operating in CDMSlite mode requires grounding one side of the detector and does not allow for collection of the ionization signal via the usual ionization channels.

²⁰ CDMSlite R2 spans R134 and R135. The gap in the middle is due to the warm up to room temperature required for servicing the cryocooler's cold head.

As such only the 4 phonon sensors on one side are readout eliminating the ability for event-by-event recoil type discrimination via yield (see Section 4.2.1 for more information). However the amplification of the Luke phonon signal allows for very low energy thresholds. In CDMSlite R2 the 50% trigger efficiency was pushed down to 75 eV_{ee} and 56 eV_{ee} due to the amplification process allowing for a world leading limit to be placed on low mass WIMPs.[76] Further information on the CDMSlite experiment can be found in Rito Thakur’s thesis which covers in detail the design and analysis for CDMSlite R1.[7]

2.4.2 Detector Health

Here I will list each detector and indicate its health in the full dataset. As Table 2-1 indicates many detectors suffered both phonon and ionization shorts. The shorts themselves come in different degrees of severity for the phonon channels. Some have a high-resistance short which introduces excess noise to the channel but it is otherwise functional. If the short worsens it can become a low-resistance short which makes the channel unusable and can sometimes heat the fridge up. In these cases the channels are disabled. The ionization shorts can also appear either on the bias lines or the FET side requiring either the channel to be grounded or the FET to be deactivated. If the channel has to be grounded this will alter the electric field which can make the detector problematic for analysis.

The time variation in detector status complicates analysis efforts significantly since shorts may occur during time periods in which we do not have much calibration data available to properly characterize how the signal region has shifted as a result of this change in detector status, not to mention the possibility for the loss of background rejection capability depending on what channel becomes unavailable. For example, the loss of an inner ionization channel would eliminate the ability to reject events at high radius on the sidewall of the detector.

- IT1Z1: This detector has a noisy SQUET on PAS2, and to a lesser extent on PDS2. This has been present since the start of the experiment and seems to become particularly prominent during ²⁵²Cf calibration. Due to the noisy behavior of this

channel, IT1Z1 has a special 7-channel summed energy threshold in contrast to other detectors' 8-channel summed energy threshold for blinding.

- IT1Z2: This detector has several shorts on it. The short on QOS1 has existed since the start of the experiment and occurs on the bias line meaning this channel is grounded, producing an altered electric field. The FET was disabled on September 8, 2014. Shortly after this channel also developed a short on the feedback which had its FET disabled on October 10, 2014. QIS1 developed signs of a feedback line short shortly after the start of the experiment and was disabled June 27, 2012 when the FET was disabled.²¹ PAS2 and PCS2 developed high and low resistance shorts respectively on December 12, 2012, potentially related to CDMSlite testing. This required the grounding of PCS2.
- IT1Z3: Several shorts exist on this detector as well. Both QIS1 and QOS1 have had shorts since the start of the experiment, with QIS1 being a bias line short and QOS1's short occurring on the feedback circuit. Both channels had their FETs disabled on June 14, 2012 and were subsequently grounded as a result. Channel PAS1 developed a low resistance short on September 14, 2012.
- IT2Z3: This detector developed shorts on two inner phonon channels since the start of the experiment. PBS1 was a low resistance short, whereas PDS1 was a high resistance short. Both channels are disabled. However PDS1's short turned low resistance between July 15 2013 and September 8 2014 and then became a high resistance short from September 8 2014 onward.
- IT3Z1: While this detector has no shorts present, it is worth noting it has severe sensitivity to low-frequency noise and had some of the highest trigger thresholds of the entire detector array.
- IT3Z3: This detector has large periodic baseline shifts on PCS1 which are often referred to, somewhat incorrectly, as square-pulses. These are not actual pulses but merely shifts in the baseline. The reason for their occurrence is not known but a data quality cut can remove them. As a result this detector also has a special 7-channel summed energy threshold.
- IT4Z1: This detector suffered from an shorted feedback line on QIS1 that started on March 28, 2012. The FET was grounded June 14, 2012 due to the excessive noise causing instability in the base temperature.
- IT5Z1: This detector is one of the worst performing of the whole array. QIS1 has had a feedback line short since March 2, 2012 and QOS1 since April 25, 2012. Both

²¹ An error on my part led to this being classified this as a bias line short with the FET disabled for the R133 processing, however this was caught for subsequent processings.

had their FETs disabled on May 22, 2012. PAS1 developed a low resistance short on May 22, 2012 as well (though its performance had been degrading prior to this, it was merely confirmed on this date). While the other channels did not develop shorts, they were susceptible to intense low frequency noise rendering this detector virtually useless for any WIMP-search analysis.

- IT5Z2: This detector had unusual noise fluctuations on QOS1 since the start. These showed up as extremely high-frequency spikes in the channel readout and contributed to very poor behavior on this channel. In addition, this detector was used as the CDMSlite detector and operated in this mode with a 70V bias reducing its availability for standard WIMP-searches.
- IT5Z3: This detector had a shorted feedback line on QOS1 since the start of the experiment. Additionally PCS2 developed a high resistance short around September 21, 2012 which worsened to a low resistance short on December 12, 2012 (also coincident with CDMSlite testing).

The astute reader will remark that while there are no apparent patterns to which phonon channels are shorted, every single ionization short occurred on side 1 of the detectors. This physically corresponds to the top side of the detector. This suggests possible mechanisms for the shorts are a result of tower debris that fell onto the detectors during installation, however this will require a post-mortem analysis of the detector surfaces once they are removed at the end of Soudan operations.²²

2.5 List of Important RQ/RRQs

Throughout the rest of this work I will be referring to many RQ and RRQs by their shorthand names. Here I will provide a table that lists each one and a brief definition. This list is by no means exhaustive of all the RQ/RRQs that exist, but should cover the ones that occur frequently. Often their names make a logical sense once the naming pattern is understood. Ionization based energy estimators often start with a “q” and phonons with a “p”. A quantity that ends with an “OF” indicates it is derived from an OF energy estimator, and “NF” is for the non-stationary OF. Channels are labeled by side of the detector by 1 or 2 (indicating top or bottom respectively), and channel label is

²² Assuming the shorts do not vanish at room temperature and become irreproducible.

Table 2-2. List of Useful RRQs

RQ/RRQ	Description
p[j][k]OF	Individual phonon channel OF energy
ps[k]OF	Side-trace summed phonon OF energy
psumOF	Summed phonon OF energy
ptNF	Trace summed phonon NSOF energy
qi[k]OF	Inner ionization channel OF energy
qo[k]OF	Outer ionization channel OF energy
qsum[k]OF	qi[k]OF + qo[k]OF
qsummaxOF	max(qsum1OF, qsum2OF)
qimeanOF	(qi1OF+qi2OF)/2
plukeqOF	qsum1OF - qsum2OF + 2min(qsum1OF, qsum2OF)* QIS1bias /3
precoilsumOF	psumOF - plukeqOF
precoiltNF	ptNF - plukeqOF
ysumOF	qsummaxOF / precoilsumOF
ytNF	qsummaxOF / precoiltNF
qrpart[k]OF	qo[k]OF / qsum[k]OF
qzpartOF	(qsum1OF-qsum2OF)/(qsum1OF+qsum2OF)
prpartOF	(pa1OF+pa2OF)/psumOF
pzpartOF	(ps1OF-ps2OF)/(ps1OF+ps2OF)
[j] = A,B,C,D [k] = 1,2	

Table 2-3. List of Common RQs

RQ/RRQ	Description
SeriesNumber	Returns the series number of a particular event
EventNumber	Returns the event number within a series (resets each series)
EventTime	Returns the unix time stamp of a particular event
LiveTime	The time in ms the experiment was armed and waiting for an event
BaseTemp	The base temperature of the fridge in K from IT3's temperature monitor
Empty	Specifies whether a detector was readout for a particular event
P[j][k]status	Returns information on the operational status of phonon channel
Q[m]S[k]status	Returns information on the operational status of ionization channel
P[j]S[k]OFchisq	Optimal Filter χ^2 value
PTNFchisq	Non-stationary optimal filter χ^2 value
QS[k]OFchisq	Optimal filter χ^2 value
Q[m]S[k]OFflag	Indicates which OF routine was run
Q[m]S[k]bias	Returns ionization channel bias voltage
[j] = A,B,C,D [k] = 1,2 [m] = I,O	

often just a letter (“a-d” for phonons, “i” or “o” for ionization). Partition quantities relate to radial, x, y, and z location information and are constructed by taking ratios of energy estimators, and yield quantities start with a “y”. The quantity **plukeqOF** is the estimate for the Luke-Neganov phonon energy and is discussed more in Section 3.9.

CHAPTER 3 IZIP DETECTORS

3.1 Physical Properties

The SuperCDMS Soudan detectors are referred to as iZIPs (interleaved Z-sensitive Ionization and Phonons). Each of the 15 iZIPs at Soudan are composed of ultra-pure germanium crystals (dislocation impurities $\sim 10^{10}\text{cm}^{-3}$) that are cylindrical in shape and measure 76 mm across by 25 mm deep. On each of the faces readout sensors are deposited via photolithography into patterns of interleaved ionization and phonon sensors. The separate layout of the channels is shown in Figure 3-1. Together both of these separate types of measurement sensor are interleaved in an alternating pattern of grounded phonon sensors and biased ionization sensors (the implications of this will be discussed in Chapter 4). The phonon channels are rotated by 60 degrees on each face to break up any degeneracy in x-y position reconstruction.




Figure 3-1. The layout of the 8 phonon channels and 4 ionization channels is shown here. Both types of measurement consist of an outer guard ring sensor to help identify events that occur at high radius.

3.2 Ionization Collection

The SuperCDMS iZIPs collect two types of energy from an interaction. The ionization energy (sometimes referred to as the “charge signal”) represents the energy deposited in the crystal to generate electron-hole pairs and drift them across the detector. At a microscopic level when a particle interacts in the crystal lattice it will deposit some




Figure 3-2. The iZIP is an interleaved detector with both ionization and phonon sensors. The flattened edges are where DIB connections occur for readout. A.) The iZIP mask is show here. The larger blue beaded lines represent the TES devices for the phonon sensors. Between them are thinner lines representing the ionization electrodes. B.) The channel naming scheme for an iZIP shows that the top side is labeled 1 and the bottom side is labeled 2. In normal operations side 1 collects electrons and side 2 holes.

fraction of its energy into the electrons in the crystal. If this energy is high enough it will allow the electrons to dislocate from their original location and drift through the crystal if an externally applied electric field is present. At the same time, a hole is produced with effectively the opposite charge which will make its way towards the oppositely biased face of the crystal. A brief overview of how the ionization signal is readout and how each species of charged particle propagates through the crystal is highly useful since the ionization signal is one of the primary discriminants for a high-mass WIMP search.

3.2.1 Charge Propagation

The behavior of charged particles in the crystal are not as simple as in free space unfortunately. In a crystal lattice electrons behave as if they have a tensor mass[77] as a result of the shape of the electron conduction band. The configuration leads to elliptical valleys aligned along the $[111]$ direction of the crystal lattice. At the low field strengths present in standard operations inter-valley scattering is not favored and an electron will tend to remain in whichever valley it scatters into initially. Coupled with the very

anisotropic tensor mass, this leads to a larger than naively expected spread in the electron radial position as it propagates through the crystal. Comparatively the holes are much more simple in their propagation as they behave as if they have a scalar effective mass, and as a result do not have as large a spread in their radial position as a function of propagation distance. This asymmetry in charge species behavior has large implications for fiducialization in the iZIP. Generally hole collection will have far fewer events shared between the inner and outer electrodes, resulting in a cleaner radial discrimination ability.




Figure 3-3. Example of electron vs hole side radial partition. The electron side has an abundance of events in the shared (middle) region compared to the hole side collection. This is due to oblique propagation effects smearing out the radial position for electron events. It is also the reason that there appear to be relatively fewer events tagged at very high radial quantities when compared to hole side collection.

For the side that collects electrons far more events will be shared between the inner and outer electrodes due to the oblique propagation (by a factor of ~ 10 [77]), and one will also tend to observe a reduced ionization signal as one progresses from the bulk to the outer regions of the detector. This is simply due to oblique propagation again, as some of

the electrons will propagate into the sidewalls of the detector which is not instrumented and surface defects that can cause charge trapping are present.

3.2.2 Ionization Readout

For SuperCDMS at Soudan the electronics readout hardware is recycled from the CDMS-II experiment (for better or worse). These electronics have been described previously in Z. Ahmed's thesis [73] among others. A brief overview is useful however. In CDMS-II a single side of the detector was devoted to two ionization sensors (an inner and outer channel). These were readout on the same detector-interface-board (DIB) as the 4 phonon channels on the top. The iZIP exploits this as well. Each iZIP is treated in some sense as two oZIPs pasted together in that each detector has two DIBs that readout half of the detector (4 phonon + 2 ionization). The readout circuit can be seen in 3-4.




Figure 3-4. Ionization Readout amplifier circuit. Here $R_b = R_{fb} = 40 \text{ M}\Omega$, $C_{fb} = 1 \text{ pF}$, $C_c = 300 \text{ pF}$, $C_{stray} = 75 \text{ pF}$. The detector capacitance is taken as the fiducial channel capacitance and is $C_d \sim 70 \text{ pF}$ [78]. Image taken from [71]

The ionization readout circuit is an amplifier circuit which is configured in a negative feedback mode. As charges move through the crystal and approach the biased electrodes, they produce a voltage signal at the amplifier's output through the feedback. In order to be isolated from a DC bias voltage, V_b , a coupling capacitor ($C_c = 300 \text{ pF}$) is inserted between the feedback circuit and the bias line. This setup has a very beneficial feature

in that the signals that come out have a fixed falltime determined solely by the feedback resistor, R_{fb} and feedback capacitance C_{fb} . To see why this is the case, we should consider the shape of an ionization pulse. It can be parameterized as having a risetime and a falltime. The risetime is dependent upon the physics of charge propagation through the crystal and as such is not unique. However the ionization signal is digitized at a rate of 0.8 μs per sample, thus if the longest time it takes for a charge signal to propagate through the crystal is close to this time, the risetime will appear to be virtually instantaneous. This happens to be true as the drift time across the crystal is on the order of nanoseconds. This is sufficiently smaller than the digitization rate that the risetime of an ionization signal will appear to be virtually instant regardless of where in the detector it originates. This is very useful because the falltime itself is determined solely by the RC constant of the feedback circuit. For SuperCDMS (and CDMS-II) $R_{fb} = 40 M\Omega$ with the parasitic capacitance of this resistor, $C_{fb} = 1$ pF. This results in a falltime of $\sim 40 \mu s$ for each ionization pulse. As a result of this setup, the shape of an ionization pulse is fixed such that the only parameter that varies is its amplitude, which is determined by the amount of charge collected by the sensor. The voltage noise on this circuit comes mainly from the first-stage JFET and from the bias and feedback resistors, but can also be traced to detector leakage current, current noise from the JFET and microphonics. A very detailed study of this circuit was performed by Shutt [79], in which he derived the following relation for the output voltage noise:

$$\epsilon_0^2(\omega) = |A(\omega)|^2 \left\{ \epsilon_{FET}^2 \left[(\omega)^2 (C_d + C_{fb} + C_{stray})^2 + \left(\frac{1}{R_{fb}} + \frac{1}{R_b} \right)^2 \right] + 4k_B T \left(\frac{1}{R_{fb}} + \frac{1}{R_b} \right) + i_{FET}^2 + i_d^2 + i_\mu^2 \right\} \quad (3-1)$$

Here i_{FET} is the current noise due to the JFET, i_d is the detector leakage current, and i_μ is the current due to microphonic noise in the wiring. The term $A(\omega)$ is the frequency

dependent transimpedance of the amplifier and is given by:

$$A(\omega) = \frac{R_{fb}}{1 + j\omega R_{fb}C_{fb}} \quad (3-2)$$

The JFET voltage noise is $\sim 0.5 \frac{nV}{\sqrt{Hz}}$ and is the dominant noise source in the above model. For most all of the data taking at Soudan the base temperature was ~ 50 mK, which supports the notion that the FET and detector current noise terms should be sub-dominant. Microphonic currents can contribute to low frequency noise in the ionization however. In Figure 3-5 a representative noise spectrum is shown from the hole-side readout of one of the detectors. As the figure indicates the noise is well described by the FET voltage term, however a forest of high-frequency peaks is visible and is caused from electronic pick up and microphonic resonances. This particular detector has a lower than typical amount of noise in the low frequency region, however it is still above the model predictions, and likely originates from electronic pickup, microphonics, and possible 1/f noise from the JFET.

3.3 Ionization Pulse Reconstruction

As was mentioned earlier, ionization pulses are digitized and reconstructed via an optimal filter (OF). The old OF algorithms for CDMS II took a shortcut and simply maximized the amplitude of the template to fit the pulse rather than minimizing the χ^2 . During the previous c58 CDMS-II result Lauren Hsu investigated this reconstruction while doing checks of the candidate events. This is described in a bit more detail in [73], but the main outcome for this work is that the OF algorithm now minimizes χ^2 for all events since it was discovered that the minimum χ^2 is not necessarily given by the maximum (positive) amplitude.

Other changes exist as well however and have led to the development of several types of OF algorithms that can be run. A brief overview is worthwhile since this does impact the reconstructed energy resolution of the ionization signal.




Figure 3-5. Ionization Noise. An example of an observed ionization noise spectrum for detector IT3Z1 during February 2013, referenced to the FET gate. An overlay of the FET noise model and Johnson noise components from Equation 3-1 is shown. Here the FET noise was set to $0.5\text{nV}/\sqrt{\text{Hz}}$ to provide agreement with the observed spectrum. The charge signal rolls off at high frequencies due to an antialiasing filter.

- **Pulse Fitting Algorithm:** Two different pulse fitting algorithms exist. The first is a single pulse fit. This is equivalent to the traditional OF fitting routine in which a single readout pulse is fit independently of the others. The 2x2 OF attempts to take into account crosstalk between the two ionization channels on the same side. The 2x2 name means that we have two traces (the inner electrode and the outer electrode) each having two possible templates (the pulse, and the crosstalk).
- **Interpolation:** All signals are digitized by the SuperCDMS DAQ and this can lead to artifacts in the event reconstruction. This was first explored by J. Filippini [71], and further followed up in [80]. Interpolation of the OF time delay significantly improves the spread of the χ^2 fit allowing for better rejection of pileup and other poorly fit events. The effect of amplitude interpolation is less dramatic and seems to have an impact of perhaps 1% on the overall resolution at both low and high energies.
- **Z Time Constraint:** This is a new feature only possible due to the two-sided charge collection of the iZIP. As discussed earlier, there is a maximum time delay between charge carrier arrivals at each face, so the noise performance could be improved if we

Table 3-1. Ionization OF Flag 3-digit number [xyz]

	0	1	2
x	N/A	1D OF	2x2 OF
y	No Interpolation	Interpolation	N/A
z	No z-time constraint	z-time constraint	N/A

constrain the time shift between each side to reflect what is physically possible. Matt Pyle performed an investigation [81] into this and found the gain to be marginal, but present. A time window of $\pm 2\mu s$ was used for the iZIP data from Soudan.

In the SuperCDMS analysis framework, a three-digit number returned by the analysis quantity `Q[m]S[k]OFflag` indicates which of the OF options were active during the event reconstruction process. Here [m] can be either “I” or “O” for inner or outer, and [k] is “1” or “2” for side 1 or side 2. The three-digit number [xyz] is encoded such that “x” is either 1 or 2, and “y” or “z” are either 0 or 1. The various combinations are shown in Table 3-1.

The three OF flags that occurred in this work can be summarized as follows:

- 211: Fully functional ionization readout. Uses 2x2 OF for both sides, interpolation, and Z constraint
- 210: Appears on functional side of detector if opposite side has channel shorts. This side still uses 2x2 OF and interpolation, but the Z constraint is no longer possible.
- 110: Appears on side of detector with channel shorts. Only 1D OF is possible, with interpolation. No Z constraint is possible.

For example if `QIS20Fflag = 211` we would understand that this event was reconstructed using a 2x2 OF, with interpolation and a Z constraint. By default for a healthy detector the `OFflag` was run with all options true. Unfortunately various ionization channel shorts occurred which made it impossible to use all the available OF constraints for every channel. If a detector has ionization channel problems that affect only one side, fully functional side will get a flag value of 210 and the broken side will be assigned 110. As we will see in Section 3.5, this has a very dramatic impact on the OF derived quantity noise resolution. As a shorthand I will refer to these quantities as the `OFflag` from now on.

3.4 Ionization Energy Calibration

Event reconstruction does not provide us with actual energy estimates for the various ionization quantities. It outputs quantities in arbitrary OF units which must be calibrated in a second round of processing (BatCalib) to allow for the creation of analysis quantities in units of energy (keV by default). In order to achieve this a very large calibration dataset is used. Since it has several recognizable peaks, we use two ^{133}Ba sources to calibrate the iZIPs. The sources provide a few hundred thousand events per series that can be used to obtain the necessary statistics to calibrate an energy scale. The 356 keV peak is particularly strong and is the primarily used peak to perform calibration, however due to the nature of the detectors the calibration must be performed in two steps.

The first step is to calibrate the inner electrodes (the fiducial electrodes) on each face. These electrodes generally will have full charge collection and as such it is possible to see several spectral lines. One can then perform a calibration to convert from OF units to calibrated energy units by matching the uncalibrated spectrum peaks to the known values. After this a 2D calibration can be performed to calibrate the outer electrode. We must use a 2D approach since the outer electrode often does not have full charge collection to produce a clean signal. The shared region will be populated and if one plots the outer channel vs the inner channel a diagonal band can be seen representing the shared region, with each end of this band corresponding to the energy of a particular peak (in this case the 356 keV γ line). This allows us to get the correct calibration for the outer electrode and a schematic of the process is displayed in Figure 3-6.

3.5 Ionization Noise from Random Triggers

After calibration is performed the random triggers (“randoms”) can also be used to estimate the OF resolution of the ionization quantities (or more precisely, the zero-energy resolution). There is a subtle distinction worth mentioning here briefly. There are two types of optimal filter classes that are processed. One is called the zero-delay optimal filter (OF0) which constrains the optimal filter delay to be $t_0 = 0$ and reconstructs to




Figure 3-6. Calibration of the outer ionization sensor requires using events that are shared between both electrodes. Using the previously calibrated inner electrode one can identify the 356 keV line as an over-dense region when plotting the uncalibrated outer electrode against the calibrated inner electrode energy. A linear regression along this line is performed and the outer channel calibration is extrapolated by assuming the y-intercept of this line corresponds to 356 keV on the outer electrode. Courtesy of Adam Anderson.

whichever amplitude best fits the noise fluctuation in this time window. As a result the mean of the distribution of OF0 fits will tend to be closer to 0 keV, and provide a less biased estimate of the true experimental noise than the standard OF. Because the OF has the freedom to shift the time delay around in the search window it will seek out and fit the largest noise fluctuation it can find that minimizes the χ^2 . This is a more restricted population as a result and will tend to have a smaller spread than the OF0 type of fit, and will have means that are offset from zero. Since the actual analysis quantities use the OF to reconstruct energies, we do want to know the offset and spread of these noise distributions since they will allow one to set energy thresholds in certain circumstances based on the noise.

If we have two channels that are independently fit with 1D OFs that are designed to select the maximum amplitude (like the original CDMS-II days), the resulting histograms of the noise on each channel should be gaussian with some mean and standard deviation. A plot of one channel vs the other would result in a 2D gaussian. A simple plot of these however would appear to be just a blob of points, and gave rise to the internal nomenclature of a “noise blob” as seen in Figure 3-7 which shows the result from the zero-delay OF.




Figure 3-7. An ionization channel noiseblob constructed from the zero-delay optimal filter (OF0) distributions. The marginal noise distribution for each channel are shown as 1D histograms on their respective axes.

The actual OF however minimizes the χ^2 of the OF fit and does not try to maximize the amplitude. This means that the OF will always be able to fit some sort of pulse in the search window, whether it has a positive or negative amplitude. As a result the simple gaussian distribution is now split into a new shape that will never reconstruct an event to 0 energy. In the case of running this type of processing without any constraint between

two channels we get a distribution that resembles that shown in Figure 3-8 which has four nodes representing the mutually independent regions that the peaks overlap. The marginal distribution shape corresponds to the `OFFlag` value of 110 from earlier (no crosstalk, interpolated, no Z constraint).




Figure 3-8. An ionization channel noise blob constructed from two independent 1D OFs with χ^2 minimization preventing any event reconstructing to 0 energy. The marginal noise distribution for each channel are shown as 1D histograms on their respective axes.

The distribution becomes a bit more interesting if we turn on the crosstalk so a 2x2 OF is performed. Here there is still a bias to avoid fitting an event with zero energy, but there is now a shared region between the peaks since the time delay is set to be the same for both channels in the 2x2 OF. As a result of this and the χ^2 minimization a particular channel will always be able to find some sort of noise peak to best fit during the reconstruction process. This results in a noise ring (or noise donut) as seen in Figure 3-9 which illustrates the shared region and the shape that results for an `OFFlag` = 210 as described above (2x2 OF, interpolation, no Z constraint).




Figure 3-9. An ionization channel noiseblob constructed from a 2x2 OF without Z constraint. This OF forces the time delay for both channels to be the same and the χ^2 minimization results in no event reconstructing to 0 energy resulting in a noise donut. The marginal noise distribution for each channel are shown as 1D histograms on their respective axes.

For the case of a fully functional detector, the 2x2 OF is still run but now the physically motivated Z constraint between both faces is present, which limits difference in the time delays between each face to be within $\pm 2\mu s$ of each other. As a result this greatly decreases the radius of the noise donut since the parameter space within which the OF can work is now much more constrained.

The noise unfortunately varies from series to series and especially over the course of one run to another. For maximal utility the noise distributions can be fit on a series dependent basis using the statistics afforded by the randoms. As the shapes are quite clearly not gaussian it is not as simple as fitting a gaussian to the resultant noise distributions. Instead an approach is followed that depends on the OF flag quantity. If a channel has a fully functional set of ionization channels (OFflag = 211) then it




Figure 3-10. An ionization channel noiseblob constructed from a 2x2 OF with Z constraint. This OF forces the time delay for both channels to be the same and the χ^2 minimization results in no event reconstructing to 0 energy resulting in a noise donut that is much tighter due to tighter constraint on the time delay. The marginal noise distribution for each channel are shown as 1D histograms on their respective axes.

was observed that fitting a double gaussian shape did a decent job at fitting the shape. The positive energy mean and standard deviation were then used to give the OF noise offset and resolution respectively. The `OFFlag` = [210, 110] cases are much less gaussian. For these distributions a localized gaussian approximation near the peaks as seen in Figure 3-11 is used to estimate the ionization offset and noise. A single valued gaussian distribution is fit to a region around the positive energy peak in each case to get the OF offset and resolution. In each case these are binned fits, and in the future an unbinned approach can probably yield better fits. This would also lend itself to a better study of the actual functional forms of the noise shapes, perhaps taking into account the full set of correlations in each case (e.g., the `OFFlag` = 211 case certainly has correlations between

the inner and outer channels on a given side, and also between channels on the opposite side).




Figure 3-11. Ionization Noise Fits based on `OFflag`. A) Double Gaussian fit performed on a fully functional ionization channel. The right most gaussian parameters are used for the noise offset and resolution. B) Peak fit performed on a charge channel that is on a working side of the detector and opposite a side with ionization channel problems. C) Peak fit performed on a working ionization channel that occurs on the same side as a broken channel.

Once these series dependent fits are performed we can see that the noise as a function of time does oscillate from series to series but generally there are regions where the noise is fairly constant around a typical value. Unfortunately the ionization noise is not as low as we would have preferred. In the best cases the resolution is about 250 eV, with some detectors having noise resolutions around 800 eV. The OF noise baseline is also substantially high, ranging from 400 - 800 eV. The noise is influenced by activities at Soudan as well and can sometimes result in a noticeable shift in noise behavior. One such case was after some modifications in the RF room to set up for CDMSlite, the detector ionization noise degraded slightly as can be seen as the discrete jump near the early part of December 2013 in Figure 3-12.

The noise fits are constructed for each of the ionization energy RRQs in the same way and are accessed through a helper function in MATLAB called `noiseblob_q_v53` which will return as outputs the means and sigmas based on nearest previous series number, useful for setting ionization thresholds.




Figure 3-12. Ionization Noise through time for R134. A) The positive energy peak of a fully functional detector. The values plotted here are the mean OF offset due to the tendency to reconstruct on an upward fluctuating noise peak. A 2σ envelope surrounds the mean values based on $\sqrt{N_{points}}$. B) The corresponding standard deviation of the noise fit for this detector as a function of time. Note the discrete jump from ~ 350 eV to ~ 450 eV and the increase in spread over time indicating a slowly degrading noise performance.




Figure 3-13. The ionization noise quantities for both electron side and hole side collection estimators (qsum10F and qsum20F respectively). These quantities vary from series to series and shown here are the median values taken over the whole c334 dataset. The ionization noise offsets are generally near or above 500 eV and resolutions are typically around 400 eV and the hole side quantity appears to have slightly worse noise performance than the electron side. Detector IT1Z3 (1103) does not have a noise offset shown here as it is negative. This particular detector recall has no functional side 1 ionization channels. Detectors 1102, 1103, 1111, 1113 and 1115 have what appear to be good noise resolution but recall these detectors have one or more shorts on side 1, the noise blob is much narrower for a “110” OF algorithm, and that qsum10F is just defined as the working channel in this case, instead of the sum of both inner and outer channels.

3.6 Phonon Collection

In addition to ionization collection and reconstruction, the iZIP measures the heat deposited into the crystal lattice via phonon sensors. The phonon physics of the iZIP are comprehensively discussed in Matt Pyle’s thesis [82], but I will provide a brief outline here of the iZIP phonon physics.

3.6.1 Transition Edge Sensors

Measuring the energy deposited into the crystal in the form of phonons is not as straightforward as was the case with the ionization signal. To accomplish this feat SuperCDMS uses transition edge sensor (TES) devices which are voltage biased and kept in electrothermal feedback with the crystal substrate. The fundamental operating principle of a TES exploits the change in resistance of material as it transitions from superconducting to normal. This transition region often has a very steep step such that minuscule changes in temperature can cause substantial changes in resistance and hence current flow through the circuit.

The TES itself is made of tungsten approximately $1\mu m$ wide sandwiched between aluminum fins to aid in phonon absorption. Together this device is referred to as a QET (quasiparticle-assisted electrothermal-feedback transition-edge sensor), and 458 of these, in parallel, comprise a single phonon channel on the iZIP.




Figure 3-14. High zoom view of a single QET. The W TES is outlined in red, with aluminum fins visible as brighter regions surrounding the TES. Courtesy of Ben Shank.

As the SuperCDMS detectors perform athermal phonon measurement, the Al fins are important in determining how many phonons are detected. When a phonon travels from the substrate to the Al fin it can split a Cooper pair inside the Al so long as it has an energy greater than twice the superconducting gap energy in Al ($2\Delta_{Al} = 0.36\text{meV}$). These energetic quasiparticles diffuse through the Al fin towards the W TES. Along the way they may lose energy by creating more quasiparticles, or phonon emission. Upon reaching the Al-W interface the quasiparticles can enter the TES and create more quasiparticles within the TES. Since the TES T_c is substantially lower in W compared to Al ($\sim 90\text{mK}$ vs. 1.2K) the corresponding necessary energy is also much lower to create quasiparticles. This quickly causes a large number of quasiparticles to exist each with less energy than necessary to diffuse back into the Al fin, minimizing energy losses from uncollected quasiparticles.




Figure 3-15. A particle interaction in the Ge substrate (green) can produce a phonon which travels towards the superconducting Al fin of a QET (grey). If the phonon has sufficient energy it can split up Cooper pair electrons (white with black circles around) into quasiparticles (orange). Since the W TES (blue) has a lower T_c than the Al, these quasiparticles will diffuse through the Al fin and enter the W TES through a bi-layer where the W is atop the Al. Here since $\Delta_{Al} > \Delta_W$ ($\sim 340\mu\text{eV}$ vs $\sim 20\mu\text{eV}$) the quasiparticles become trapped (black circles) in the TES and heat it up causing a sharp increase in resistance.

3.6.2 Phonon Production and Propagation

When an event occurs within the crystal three different types of phonons may be produced subsequently. Primary phonons are produced from the direct recoil of nuclei and electrons as a result of an incoming particle interaction. These phonons contain position information of the event and are created with very high frequencies ($\nu \sim 10$ THz). These phonons will proceed to initially undergo anharmonic decay in which they split into two lower frequency phonons, a process that is extremely frequency dependent with a rate $\Gamma \sim \nu^5$ and will act to very quickly convert the initial phonons to lower frequencies, keeping them confined to the initial interaction site. Also important is isotope scattering which starts to dominate at $\nu \sim 1.6$ THz and occurs at a rate that is also very frequency dependent, $\Gamma \sim \nu^4$. Here phonons can scatter due to the presence of different isotopes within the crystal lattice. In a crystal constructed of absolute pure material of a single isotope this would not occur. Eventually the initial primary phonons will have down-converted to phonons with mean-free paths on the order of the detector size and will transition from a quasi-diffusive propagation to become “ballistic” phonons and eventually become homogenous throughout the crystal.

The second class of phonons that are produced come from the motion of e^-/h^+ pairs through the crystal and are referred to as Luke-Neganov (often just Luke for short) phonons. As described by Luke and Neganov [83, 84] when charged carriers are drifted in a crystal at speeds in excess that of the crystal speed of sound, this excess energy is emitted in a process analogous to Cherenkov radiation. As the charge carriers drift across the crystal with very high speed, the work done by the external electric field to accelerate them further will be shed as Luke phonons making the Luke phonon energy directly

related to the bias voltage across¹ the crystal, assuming full charge collection.

$$E_{Luke} = N_{eh}eV_b = \frac{E_Q}{\epsilon_\gamma}eV_b \quad (3-3)$$

The third type of phonon produced during an event comes from the recombination of the e^-/h^+ pairs when they reach the surface at the electrodes. The initial energy required by an incoming particle to produce and move an e^-/h^+ pair is roughly $\epsilon_\gamma = 3$ eV and consists of the energy to overcome the band gap in germanium as well as losses to primary phonon production. When the e^-/h^+ pairs reach the surface they are not collected, rather the ionization electrodes sense the image charge produced by the charge carriers. The pair eventually down-scatter near the surface and relax to lower energy states releasing the stored band gap energy back to the crystal in the form of recombination phonons,

$$E_{recom} = N_{eh}E_{gap}^{Ge}.$$

Together then these three phonon types contribute to the total phonon energy (in units of total phonon energy, keV_t) detected by a phonon sensor

$$E_t = E_{primary} + E_{Luke} + E_{recom} = E_r + E_{Luke} \quad (3-4)$$

Note that the sum of the primary and recombination phonons is the amount of recoil energy deposited by the event, since the Luke phonon energy occurs from the work done by the E-field. As the phonons propagate through the crystal they eventually will thermalize after about 750 μs leading to uniform phonon power on all phonon channels. This feature is quite useful to use for relative calibration of the phonon channels.

3.6.3 Phonon Readout

The signal from a particular event as described above will be very minuscule, and as such it needs to be amplified. To do this SuperCDMS at Soudan uses the same hardware as CDMS-II in which a SQUID amplifier circuit is used to measure the change in the

¹ One face is biased at +2V and the other at -2V

TES current via the readout circuit shown in Figure 3-16. When an event deposits energy into the crystal and subsequently into the QET, the TES will warm up slightly, causing a sharp increase in the TES resistance. The steepness of the transition region is often characterized by a parameter α :

$$\alpha = \frac{T_0}{R_0} \frac{\partial R}{\partial T} \quad (3-5)$$

Where T_0 and R_0 are the operating temperature and resistance of the TES. As the resistance rises, more current will flow through the shunt resistor ($R_{sh} \sim 20m\Omega$). This change in current through the input inductor will result in a change in magnetic flux through the SQUID array. In turn, this alters the current through the feedback circuit, which is kept in negative feedback via the feedback inductor. The feedback coil has $1/10^{th}$ the number of turns as the input coil and as a result the feedback circuit's response to any change in current through the input side will be larger by a factor of 10, providing amplification.




Figure 3-16. The SQUID amplifier circuit used to readout phonon signals. When the TES is superconducting current flow is steady through the TES branch and only noise fluctuations in the current are read out. When the TES heats up sufficiently due to an event, R_{TES} goes normal reducing the current flow through that branch. Here $R_{sh} \sim 22m\Omega$, $R_{fb} = 1k\Omega$, $L_i = 250nH$, and $R_{TES} \sim 700m\Omega$ (normal). Courtesy of Jeff Filippini. [71].

Since the TES is operated in a voltage biased state, this means that the contribution to Joule heating will decrease as the TES heats up:

$$P_J = \frac{V_{bias}^2}{R_{TES}} \quad (3-6)$$

Recalling that R_{TES} is a function of temperature, we can verify the Joule heating response to changes in temperature:

$$\frac{dP_J}{dT} = \frac{d(V_{bias}^2/R_{TES})}{dT} = \frac{\partial(V_{bias}^2/R_{TES})}{\partial R_{TES}} \frac{dR_{TES}}{dT} = -\frac{V_{bias}^2}{R_0^2} \frac{dR_{TES}}{dT} \quad (3-7)$$

We note that since $dR/dT > 0$ then for positive temperature excursions $dP_J/dT < 0$ as expected. Since the Joule heating decreases as the TES is slightly warmed, this allows the cooling power from thermal contact with the substrate to bring the TES back to an equilibrium point. The quantity V_{bias}^2/R_0 can be thought of as the operating point Joule heating term, but since at the operating point we are in thermal equilibrium with the bath, this is also just P_{bath} .

I will comment here that ETF allows for TES systems to have much larger bandwidths than one gets from just thermal contact. If an object is in thermal contact there is an associated timescale of $\tau_{therm} = C/G$, where C is the thermal heat capacity of the TES and G is the thermal conductivity between the TES and substrate. This is the characteristic fall time in the absence of ETF. When ETF is accounted for ([71, 82]), the falltime is now $\tau_{etf} = \tau_{therm} \frac{1}{1+\mathcal{L}}$. Here $\mathcal{L}$ is a symbol utilized commonly in [85] to denote the ratio $\alpha P_{bath}/GT_0$. As we can see here the characteristic ETF time can be quite short compared to the thermal time scale allowing for a much quicker TES response.

The phonon noise is has many subparts that contribute to it, but the shunt resistor Johnson noise is a dominating factor as it is heat sunk to the still stage which is at about 1 K as opposed to the $\sim 55mK$ TES resistor. As a result the Johnson current noise due to the shunt resistor should be on the order of $\sim 10pA/\sqrt{Hz}$. However observations of the TES noise spectrum at Soudan (3-17) show that at low frequencies the noise is much

higher than we might expect. This low-frequency noise issue is believed to arise from several sources, such as $1/f$ electrical noise and complex vibrational couplings of various cryocoolers to the detector tower assembly the latter of which proved to be problematic for low energy events, requiring a low-frequency noise cut to be developed.




Figure 3-17. TES noise spectra for detector IT4Z2. These noise spectra were created from the first 500 random triggers at the start of this particular data series. At low frequencies we are dominated by various noise sources that couple into the tower assembly. Image generated from online data quality (DQ) summary page.

3.7 Phonon Reconstruction

As with the ionization signals, the phonon signals are also reconstructed using an optimal filter. Recall that the phonon signals are also digitized (4096 bins at $1.6 \mu s$ per bin). Unlike the ionization pulses though, there is no unique phonon template as determined by the electronics. The rising portion of a phonon signal encodes position dependence of the event and the tails have uniform energy after about $750 \mu s$. As such an experimentally determined template is used for all channels on a given detector. This template is selected by averaging many high-quality electron recoil events from the bulk of the detector.²

² This can introduce a systematic bias when fitting nuclear recoil events or even surface events which can have slower or faster rise times respectively due to Luke phonon dynamics.

For the phonons two types of optimal filter are used: the traditional 1D OF and a new non-stationary optimal filter (NSOF, or NF) [86, 87]. The NSOF is a way to remove any position dependence (or at least, minimize it) from the energy reconstruction. The trick is to construct an average total phonon trace template by examining many phonon traces with the same start time. This average template is then subtracted from every single (total) phonon template to create a set of residuals. These residuals physically originate from subtle position dependent effects in the detector, but we could also treat them as non-stationary noise to hide our lack of knowledge about the actual position dependence. The NSOF has very good energy resolution compared to the regular OF, but since it is computationally expensive and de-weights any position dependence it is only run on the total phonon energy to create **ptNF**.

3.8 Phonon Calibration

Unlike the ionization channels we cannot perform as simple a calibration using two channels since we have a total of 8 phonon channels. Calibration is done in a two-step process in which we first generate a relative calibration amongst all 8 phonon channels, and then scale the total summed trace to an absolute energy scale. Since each phonon channel has the same Al fin coverage, and it takes $\sim 750\mu s$ for phonons to thermalize, we expect that the total power in each channel to be the same at this point. Therefore to perform the relative calibration a set of coefficients is determined such that the integrals of the traces at late times will all be equal relative to a particular channel. In c334 the channel PAS1 was chosen to be the channel others were normalized against, however this choice is not special and any working channel could be used. One of the benefits of having multiple phonon channels is that x-y partition information is possible and so one way to judge the accuracy of the calibration is to create a triangle plot. If the channels are properly relatively calibrated then the inner channels should form two overlapping equilateral triangles.

To construct the triangle plot the phonon x and y partitions must be defined. A set of bases can be constructed such that each channel's unit vector bisects each phonon channel resulting in the following definition for the partition quantities:

$$\begin{aligned} \text{pxpart10F} &= \frac{\text{pd10F} \cos(\pi/6) + \text{pb10F} \cos(5\pi/6) + \text{pc10F} \cos(3\pi/2)}{\text{pd10F} + \text{pb10F} + \text{pc10F}} \\ \text{pypart10F} &= \frac{\text{pd10F} \sin(\pi/6) + \text{pb10F} \sin(5\pi/6) + \text{pc10F} \sin(3\pi/2)}{\text{pd10F} + \text{pb10F} + \text{pc10F}} \\ \text{pxpart20F} &= \frac{\text{pd20F} \cos(7\pi/6) + \text{pb20F} \cos(11\pi/6) + \text{pc20F} \cos(\pi/2)}{\text{pd20F} + \text{pb20F} + \text{pc20F}} \\ \text{pypart20F} &= \frac{\text{pd20F} \sin(7\pi/6) + \text{pb20F} \sin(11\pi/6) + \text{pc20F} \sin(\pi/2)}{\text{pd20F} + \text{pb20F} + \text{pc20F}} \end{aligned}$$

A combined pxpart0F and pypart0F can be constructed by adding the numerators of each individual side and then dividing by the total inner channel energy.

After a set of relative calibration constants is created, the phonon traces can be appropriately summed together into a total energy trace. This trace must then be calibrated to an absolute energy scale. Here the traditional approach was to use the fact that for a pure electron recoil the yield is equal to 1, and so $E_r = E_Q$ and thus the sum of the total ionization energy plus Luke phonons should be equal to the total phonon energy:

$$\begin{aligned} E_t &= E_Q + E_{\text{Luke}} \\ \text{pgqOF} &= \text{qsummaxOF} + \text{plukeqOF} \Rightarrow \text{ptNF} \quad \text{for } Y = 1 \end{aligned} \tag{3-8}$$

The yield of an event is simply the ratio of the ionization energy to the total phonon recoil energy, and will be discussed more in detail shortly. Under this assumption, since pgqOF is entirely determined from (calibrated) ionization quantities, we can select an overall constant to enforce the equality from the uncalibrated total phonon quantity. The picture is unfortunately more subtle in the case of the c334 dataset. It was first noticed by Lauren Hsu that the position of the 10.4 keV activation peak in post-neutron calibration data seemed to drift over time by a few percent. Subsequent examinations revealed that the




Figure 3-18. An example of a triangle plot after relative calibration has been performed. The vertex of each triangle corresponds to the center region of a phonon channel based on the iZIP mask geometry. Since S1 and S2 are rotated by 60 degrees relative to one another this explains why the triangles are not perfectly overlapping. In this plot S1 corresponds to the blue triangle and S2 is the red triangle. Courtesy of Adam Anderson.

yield varied dramatically as well over time. The ionization spectral features did not appear to be time dependent, suggesting that the origin lay with the phonon signals. The phonon gain can vary in relation to temperature, and a plausible avenue of influence could be from base temperature fluctuations relative to a particular detector's T_c values. Unfortunately there was no observed correlation between the degree of temperature dependence and detector T_c values.³ There is a very clear temperature dependence to the yield quantity however, so while the cause remains uncertain, this dependence was calibrated out.

³ A very important caveat to this is the fact we only had working temperature measurements for tower 3. It is entirely plausible other towers had different temperatures.




Figure 3-19. Absolute phonon calibration process for ptNF on detector IT4Z3. A.) There is a very strong dependence on base temperature present in the uncalibrated yield quantity. Temperature bins of 1mK in width were created to get mean yield values which are then fit to a quadratic polynomial. B.) After calibrating temperature dependence out and scaling ptNF to be equal to the total phonon energy from the gamma hypothesis the yield appears to be constant as a function of base temperature. Courtesy of Adam Anderson.

A second order polynomial was observed to produce the best fit between base temperature and the uncorrected yield dependence. Figure 3-19 shows an example of this dependence. It is important to note that not every detector (or phonon quantity) shows as strong a temperature dependence, and not every detector has a monotonic slice of the polynomial either. Bizarrely some detectors have a max or minima present! It is not understood at all why this should be so at this time, but there could be one or more underlying temperature dependent phenomena that are manifesting themselves in this fashion.

3.9 Phonon Energy Scales

In Equation 3-3 we expressed the Luke energy in terms of the number of e^-/h^+ pairs created, and in terms of the collected charge energy. Recoils off of nuclei (nuclear recoils, or NRs) produce less ionization energy per keV of recoil energy than recoils off of electrons (electron recoils, or ERs) [88]. This quenching factor can be expressed in terms

of a quantity called yield:

$$Y = \frac{E_Q}{E_r} \quad (3-9)$$

For ERs the yield is constant and is assigned a normalized value of $Y = 1$. Energy estimators that are constructed from this type of recoil are said to have electron-equivalent recoil energies, or keV_{ee} . For a NR event however, the yield now varies as a function of recoil energy and the situation becomes a bit more messy. Consider the total phonon energy equation re-expressed as follows:

$$E_t = E_r + E_{\text{Luke}} \Rightarrow E_r + \frac{E_Q}{\epsilon_\gamma} eV_b \Rightarrow E_r + \frac{Y E_r}{\epsilon_\gamma} eV_b \Rightarrow E_r \left(1 + Y \frac{eV_b}{\epsilon_\gamma} \right) \quad (3-10)$$

One can construct a recoil energy estimator out of the measured total energy E_t from Equation 3-10 either by making assumptions about the recoil type, by measuring the yield in some fashion, or using a yield model.

$$E_r = \frac{E_t}{1 + Y(E_r) \frac{eV_b}{\epsilon_\gamma}} \quad (3-11)$$

If we assume an ER type of event for example, we can get a recoil energy in keV_{ee} where $E_r = E_Q$ from Equation 3-9 and we see that since $V_b = 4V$, our conversion factor is just

$$E_{r,ee} = \frac{E_t}{1 + \frac{4}{3}} = \frac{3}{7} E_t \quad (3-12)$$

If one assumes a NR event, one method to determine the recoil energy is to use Lindhard theory [88] to parameterize the yield and then solve Equation 3-10 for the nuclear-recoil energy, with units of keV_{nr} .

In SuperCDMS Soudan we construct an estimate for the Luke phonon energy from the measured charge energy itself which then lets us solve for the recoil energy directly⁴ If one assumes that the collection efficiencies are equal for all phonons and negligible charge

⁴ In general this can be written with placeholder units of keV_r to denote it is a recoil energy and not total phonon energy. Whether it is keV_{ee} or keV_{nr} depends on the interaction type.

trapping⁵, we can express the Luke phonon energy as follows:

$$E_{Luke} = \frac{2\min(E_{Qh}, E_{Qe})}{\epsilon_\gamma} |eV_b| + \frac{|E_{Qh} - E_{Qe}|}{\epsilon_\gamma} |eV_b| \quad (3-13)$$

Where here E_{Qh} and E_{Qe} are the measured hole side and electron side ionization energies respectively. In this form the ionization signal has been split into a symmetric term and an asymmetric term. For symmetric events (i.e., in the bulk of the crystal) charge carriers will drift to either side and each will produce their own Luke phonons, giving rise to the factor of 2 in the first term of Equation 3-13, and tending to make the other term small. Events near the surface only wind up on one side and so produce an asymmetric signal in which one of the measured energy terms is very small compared to the other. This will cause the symmetric term to be very small now, and the asymmetric term to be larger as expected. Armed with this ionization-signal derived estimate for E_{Luke} , the recoil energy falls out as:

$$E_r = E_t - \frac{2\min(E_{Qh}, E_{Qe})}{\epsilon_\gamma} |eV_b| + \frac{|E_{Qh} - E_{Qe}|}{\epsilon_\gamma} |eV_b| \quad (3-14)$$

In terms of our analysis quantities the recoil energy from our total phonon energy estimator `ptNF` is just: `precoiltnf = ptNF - plukeq0f` where `plukeq0f` is based on the side-summed ionization energy estimators.⁶

3.10 Phonon Noise from Random Triggers

Similar to how random triggers were utilized to obtain ionization (zero-energy) noise fits, we can do the same for each of the primary phonon energy estimators. In many

⁵ Not necessarily true.

⁶ This method has a several Achilles' heels however. For detectors with broken charge channels or altered fields it is likely to not result in a correct estimate of the Luke phonon production. Additionally at low recoil energy the amount of ionization signal produced is greatly reduced, perhaps even indistinguishable from noise. In such cases this is where a yield model, such as the standard Lindhard model or an empirical model as was used in the LT analysis can be used instead. This allows one to compute the yield for a given recoil energy under the assumption that the event is a NR.

senses it is easier to obtain the phonon fits since each optimal filter is a 1D fit. A handy MatCAP⁷ function (`IterFHGauss`) performs a binned, iterative gaussian fit. To do this it will cut away data that occurs greater than 2σ from the mean until the resulting change in the mean and standard deviation is less than some threshold. Once this threshold is reached, a gaussian is fit to the remaining data.




Figure 3-20. Example phonon noise distribution for IT3Z2 and energy estimator `ptOF`. The skew and excess-kurtosis are also shown but are not standard outputs of `IterFHGauss`. The heavy tail on these fits were initially thought to be primarily LF noise events but may be a mixture of these plus an intrinsic bias in the OF routine to select maximum amplitude events from noise.

Every phonon channel fit shows evidence of a heavy tail in the actual distribution, as seen in Figure 3-20 on the right side. During the initial portion of R133 many attempts were made to understand and remove this tail so that proper blinding thresholds in terms of the noise thresholds could be obtained. Early hypotheses were pileup in the random noise triggers or oddly reconstructed events. A loose χ^2 -cut did not improve this situation any, and this motivated the computation of the skew and (excess) kurtosis quantities. These

⁷ MATLAB CDMS Analysis Package

higher-order moments did not shed any light on the behavior of the distributions other than confirm that the heavy-tailed distributions were constant over time. Eventually a low frequency noise template was developed by Kristi Schneck to allow for a χ^2 discrimination against low-frequency noise events, which could be super-imposed on a random trigger and cause a systematic bias in energy. Since these are essentially zero energy traces however, the discrimination power of the cut at such low energy is marginal at best and indeed > than 95% of all randoms pass this cut. Another possibility that may be adding to the shape as well is that the OF algorithm, while attempting to minimize the χ^2 of the fit, generally maximizes the amplitude. Even if the noise is assumed to be perfectly gaussian, using the OF is a biased selection in which we are sampling from the maxima of the underlying noise distribution.⁸

Unfortunately the phonon noise resolutions and offsets varied from one series to another by approximately 20%. On some detectors it was seen that there were clear time dependancies to the noise offsets and means serving as a cautionary flag that a phonon channel's performance was marginal. Towards the end of 2013 the phonon noise was particularly bad in part due to excessive low frequency noise induced by the cryocooler's helium charge running low. In February of 2014 this was recharged and, for a time, resulted in a noticeable reduction in noise until the cold head itself deteriorated to the point that base temperature could not be stably maintained anymore (after an impressive ~ 3 years of operation). These time dependent noise parameters are stored in a helper function `noiseblob_ph_v53` that allow one to request the mean noise and sigma values as a function of series and are useful for setting energy thresholds.

⁸ However even in the zero-delay OF processing the heavy tail is present so this is not sufficient to explain away the tail.




Figure 3-21. Phonon mean noise offset and noise width as a function of time. A.) The mean offset of **ptNF** noise vs time with $2\sqrt{N}$ noise bounds in pink. B.) The noise resolution as a function of time. The sharp decrease around Feb 2014 is due to a recharge on the primary cryocooler.




Figure 3-22. The phonon noise quantities vary on a series to series basis. Shown here are the median noise offset and median noise resolutions from the series dependent values for each detector and the primary total phonon energy estimator **ptNF**. Most noise resolutions are quite good at below 300 eV, but offsets are a bit high on the source detector 1107. When compared to the ionization noise medians it is clear the phonons have both lower noise offsets and better resolution across the board.

CHAPTER 4

BACKGROUNDS AND IZIP DISCRIMINATION

The iZIP is purpose built to detect WIMPs while having excellent background discrimination abilities. This chapter will discuss some of the physical origins for the various backgrounds, and then show a few techniques at our disposal to reject these backgrounds. In CDMS II, pulse shape and timing information were used to attempt background discrimination. In this work however, no timing information is used for discrimination purposes (though with better pulse templates and OF algorithms, timing quantities could possibly be used in future analyses), just energy based quantities.

4.1 Physical Origins

The various backgrounds present in the iZIP come from a few different vectors, even with the immense amount of shielding present. Cosmogenic sources of background are the smallest, with Monte Carlo simulations indicating they are subdominant, with much less than 1 cosmogenically induced neutron expected in the full exposure. This is good because these sources are expected to contribute to a neutron background rate, and neutrons mimic WIMPs in these detectors. All the other sources of background are electron-recoils. As we will see later in this chapter, yield is an excellent way to discriminate between ER vs NR type events, but depending on the background type they can appear to have yield values consistent with NRs. Gamma and X-ray sourced backgrounds arise from external (to the detectors) sources such as the lead shielding or copper housing itself. Typically these photons will Compton scatter inside the detector and produce bulk ERs. A study of the global gamma rate was performed at the start of R133 and showed that the overall global gamma rate was 104.2 ± 1.4 counts/kg/day in an energy range of $10 - 100\text{keV}_{\text{ee}}$ [89], with an average of ~ 2 counts/kg/day/keV at lower energies. The global gamma spectrum can be used with Monte Carlo simulations to attempt to discern the contribution from contamination of various radioisotopes such as ^{40}K , ^{232}Th , ^{238}U from various parts of the fridge assembly.[90].




Figure 4-1. The global gamma rate over all detectors for the SuperCDMS Soudan iZIP array. Various spectral features are present, but not as well resolved as in CDMS II. Courtesy of Peter Redl [91].

Occasionally a photon can photo-eject an electron from the copper housing, or a neighboring detector and this electron (sometimes referred to as an “ejectron” in CDMS jargon) will then interact in the surface region of the detector, however this is not a frequent occurrence. An additional source of background arises after neutron calibration of the detectors. During this process germanium isotopes in the detector can become temporarily activated and produce spectral peaks as they decay due to electron capture. The most prominent occurs around 10.4 keV_r and is a result of K-shell capture, but there is another noticeable peak around 1.3 keV_r that is the result of L-shell capture. The K-shell feature provides a convenient peak to verify the NR energy scale is calibrated properly.

Another set of backgrounds originate from the ^{210}Pb decay chain (shown in Figure 4-2). Due primarily to ^{222}Rn contamination, this long-lived isotope of lead is present on both the detector housings and the detector surfaces themselves unfortunately. The decay products of ^{210}Pb create ERs that occur in the outer surfaces of the detector (both

high radius and high z) and possibly α s with very high energy as well. The high energy α s are easily discriminated against since they possess 5.3 MeV of recoil energy, which saturate our readout channels. Face-side surface events are also easily tagged due to their asymmetry parameter (discussed later in this chapter). These events typically consist of β s and ^{206}Pb nuclei. At high radius radial discriminators can help tag the sidewall events, but there is a strong z-dependence to the radial parameters depending on whether the event is on the electron or hole side of collection due to previously mentioned propagation effects.




Figure 4-2. As ^{222}Rn decays it eventually forms ^{210}Pb which has a comparatively lengthy half-life of 22.3 years, making it effectively a permanent source of background in an experiment. The decay chain starts with ^{210}Pb undergoing β decay. It may either decay directly to ^{210}Bi or to an excited ^{210}Bi state which then relaxes by either ejecting electrons and possibly an X-Ray, or the emission of a γ -ray (this relaxation process happens on ns timescales). The ^{210}Bi then decays via an easily identifiable high-energy β to ^{210}Po . Since the ^{210}Bi is very heavy it has very little recoil energy and depth ($\sim\text{nm}$) as such does not produce a noticeable response in an iZIP. The final decay step is from ^{210}Po to ^{206}Pb via a very high energy α which is easily identified. The ^{206}Pb nuclei does not necessarily deposit the full 103 keV of recoil energy into the detector, but rather an amount that is dependent upon the trajectory of the ^{206}Pb . Additionally since it has such a high mass and low penetration depth it does not produce a large ionization signal, resulting in a severely reduced yield (often on the lower edge of the NR band) making it an important source of leakage at low yield.

Because the ^{210}Pb decay chain is in secular equilibrium the rate of observed α s should provide a direct proxy for the other expected products allowing an overall normalization

of these backgrounds. The α s are chosen simply because they are the easiest to identify due to their high energy. The ^{210}Bi nuclei recoil is extremely difficult to tag due to its low kinetic energy, and the β s all arise from many different parts of the chain and may not deposit the full decay kinetic energy into the crystal. Based on studies of the α rate, the estimated rate, over all functional detectors during the full c334 time period, is estimated to be about 11.4 counts/day/cm².

4.2 Background Discrimination

With the presence of known sources of backgrounds contaminating the environment around the detectors, and the detectors themselves, it is very important to be able to identify background events as distinct from signal events. Recalling the discussion in Section 1.4.3 it is extremely useful to be background free and have as much exposure as possible. The iZIP gives a multitude of parameters with which to identify differences between background and signal events, most of which are very powerful when used with other quantities simultaneously. By setting a final set of discrimination cuts on these parameters in some optimized fashion, one can attempt to minimize background leakage while allowing for as much exposure to remain as possible. This process is ultimately what sets a fiducial volume in the detectors since the outer regions will be heavily populated with background events they are most likely to be removed by cuts in various discrimination quantities. The primary discrimination quantities may be broken down into two main categories: yield and energy partitions. The yield quantity allows for discrimination between ER and NR interactions, and the partition quantities will allow for fiducialization as they correlate roughly with how far into the bulk of the crystal an event occurs.¹

¹ As we will see the partition quantities are often normalized such that 0 is in the bulk and ± 1 corresponds to the outer edges. It is important to keep in mind that a partition value of 0.5 for example does not imply half the detector is removed. Because the inner electrode extends outward in radius to approximately 35 mm and the surface electric field becomes complex only in

4.2.1 Yield

The most useful quantity constructed from the simultaneous measurement of ionization and phonon energy is referred to as the ionization yield, or yield for short. As was discussed in Section 3.8 the yield of an event is just the ratio of the ionization energy to the phonon recoil energy.

$$Y = \frac{E_Q}{E_t - E_{Luke}} = \frac{E_Q}{E_r} \quad (4-1)$$

Events that are pure ER are calibrated to have a yield of 1, and fall within an ER band as defined by ^{133}Ba calibration data. Events consistent with NR are designated as such based on a NR band as fit from ^{252}Cf calibration data taken periodically. Both bands are extremely well separated over most of the energy range used for WIMP searches, however it is clear that at lower energy the bands flare and eventually overlap. This is a combination of both the inverse dependence on the recoil energy and a diminishing signal to noise ratio from the ionization signal. This can be seen by computing the uncertainty on the yield measurement.

$$\sigma_Y^2 = \left(\frac{\partial Y}{\partial E_Q} \sigma_{E_Q} \right)^2 + \left(\frac{\partial Y}{\partial E_r} \sigma_{E_r} \right)^2, \quad \frac{\partial Y}{\partial E_Q} = \frac{1}{E_r}, \quad \frac{\partial Y}{\partial E_r} = -\frac{E_Q}{E_r^2} = -\frac{Y}{E_r} \quad (4-2)$$

Note that $\sigma_{E_r}^2 = \sigma_{E_t}^2 + \sigma_{E_{Luke}}^2$, and recall we determine the Luke phonon energy entirely from the charge signal, very roughly² $\sigma_{E_{Luke}} \sim \sigma_{E_Q}$. Substituting this into the equation for σ_Y we observe:

$$\sigma_Y^2 = \left(\frac{\sigma_{E_Q}}{E_r} \right)^2 + \left(-\frac{Y}{E_r} \right)^2 \left(\sigma_{E_t}^2 + \sigma_{E_Q}^2 \right) \quad (4-3)$$

the first mm of depth, the partition values can be close to zero for the vast majority of the detector's physical volume.

² This is assuming that we have full charge collection, and even then is a rather large assumption! But for this argument it is a sufficient assumption to illustrate the ionization noise will dominate the yield.

Cleaning this up and taking the square root we find the expression for the yield uncertainty:

$$\sigma_Y = \frac{1}{E_r} \sqrt{(1 + Y^2) \sigma_{E_Q}^2 + Y^2 \sigma_{E_t}^2} \quad (4-4)$$

This equation implies that at a given recoil energy we should expect the yield measurement to have a larger uncertainty for ER (where $Y = 1$) and a somewhat tighter distribution for a NR event, and that at very low recoil energy a flaring of the bands due to the $1/E_r$ dependence. This is indeed what is observed in the iZIP detectors as Figure 4-3 shows. It is also quite clear that if an event has incomplete charge carrier collection, the ionization signal, and hence the yield will be reduced. This is a hallmark of a surface event or a very high radius event, which require a different type of quantity to tag.

4.2.2 Interleaved Sensors and Partitions

Thus far the the yield has been described as a powerful, method to allow for discrimination between backgrounds and signal like events in the iZIP detectors. However the iZIP has a few new methods available to it that make it a much more versatile dark matter detector than the old oZIP technology used in CDMS-II. We have a significant background that exists where yield is not very useful due to events that occur near the surfaces of the detectors, especially in the $\sim 10\mu m$ dead-layer. These suffer from reduced ionization collection due to crystal imperfections leading to charge trapping, and additionally at high radial distances from oblique propagation of electrons into the sidewalls. These events are the infamous background population referred to as “surface events” (SE) that proved to be rather difficult to tag. Their origin is seen to mostly arise from ^{210}Pb decay daughters, but can also occur from electrons ejected from adjacent detectors or detector housings that interact in the surface layer of the detectors.

Here the iZIP’s design allows for very easy identification of these surface events due to two very useful design features:

- Each face is outfitted with ionization sensors, allowing the collection of energy from drifting electrons and holes across the crystal




Figure 4-3. The ionization yield as a function of recoil energy is a “classic” SuperCDMS plot. This shows the discrimination power of the iZIP against ERs (blue) and NRs (green). Here the ER events (with yield ~ 1) originate from a ^{133}Ba γ calibration source, whereas the NR events are from a ^{252}Cf neutron calibration source. The distributions from each type of calibration can be fit to gaussian bands as a function of recoil energy and produce bands (shown as solid lines in lighter color) that define the signal region. At low energies the yield bands flare and event by event discrimination of ER and NR events becomes impossible. Blue events in the signal region in this plot are surface events that occur at the edges of the detector and have incomplete charge collection.

- Each biased ionization electrode is interleaved with a grounded phonon sensor

The second item in particular is crucial to discrimination against SE. Far away from the surface regions of the detector the potential has a uniform and linear gradient resulting in a uniform electric field in the bulk of the crystal. However near each face a clam-shell like pattern emerges resulting in a complex electric field as seen in Figure 4-4. Combining the two features mentioned above results in the following:

- Bulk Events: Events in the bulk occur in the region of uniform electric field. As a result e^-/h^+ pairs will drift to opposite faces of the crystal. The ionization sensors will record the signal from both types of charge interaction. The result is a roughly symmetric ionization signal.

- **Surface Events:** Events near the surfaces have a different story. Due to the complex electric field charge propagation across the crystal does not occur. One species of charged particle will drift to the appropriately biased electrode, while the other will follow the field lines towards the grounded phonon sensor. The result here is an absence of significant ionization signal on the opposite face.




Figure 4-4. iZIP Electric Field. Here we see a view of the electric field of an iZIP. Notice that in the bulk region the field lines become uniform, however within ~ 1 mm of the surface the interleaved ionization and phonon sensors (biased and grounded respectively) give rise to a more complex field.

As seen there is a clear difference in ionization signal for bulk and surface events.

This is very suggestive in that we can create an analysis quantity to discriminate based on the asymmetry as follows:

$$\text{qzpart} = \frac{E_{Q1} - E_{Q2}}{E_{Q1} + E_{Q2}} \quad (4-5)$$

This quantity will be +1 for events near surface 1, 0 for events in the bulk, and -1 for events near surface 2. This quantity is mostly independent of energy except at low energy, at which point the reduced signal to noise ratio causes the z-partition value to flare.

In a similar fashion the ratio of energy collected on the outer channel over the total energy on a particular side is used to construct a radial partition quantity, which is 0 for a bulk event and 1 for a sidewall event. Unlike the z-partition, there are two radial partition




Figure 4-5. Example of a surface event from the ionization signal. The largest trace (red) has been vertically offset for visibility. Each trace represents the sum of the inner and outer ionization channels for the particular side (Side 1 in blue, Side 2 in red). The maximum signal energy is seen to be on Side 2 for this event, and is used to determine the ionization energy (48.62 keV). A bulk event would have a similar amplitude between the two sides.

quantities available, one for each side. As will be seen in Chapter 7, this quantity has a great deal of dependence on the z-symmetry value.

4.2.3 Ionization Discrimination Ability

As mentioned earlier, since the iZIPs represent a new detector technology it was decided by the collaboration to implant two of the detectors with a ^{210}Pb source to collect statistics on surface event rejection. One detector, IT3Z1, had its top surface implanted for study of electron collection side discrimination, and the other, IT3Z3, had its bottom side to study rejection for hole collection side discrimination. After several months of collecting data, totaling ~ 900 live-hours, an analysis of the iZIP discrimination abilities was performed showing the rejection of surface events from the bulk at a rate better than 1.7×10^{-5} , and bulk ERs from NRs at a rate better than 5.9×10^{-7} above $8 \text{ keV}_{\text{nr}}$, and with a ionization fiducial volume of approximately 30%. [74]




Figure 4-6. Primary iZIP ionization discrimination planes demonstrating surface event rejection capabilities. a) The ionization energy collected on Side 2 is plotted against that collected on Side 1. Events that are symmetric occur around the line $y = x$ in this plane and a generous region is tagged in blue. Surface events occur primarily on one side or the other. In this case the detector shown here has a ^{210}Pb source on S1, hence the large number of events on that face. The band near 46 keV that penetrates into the symmetric region is due to the 46.5 keV γ line in the ^{210}Pb decay chain. b) The yield vs energy plane with the asymmetric events tagged in red. While many ER events show up in the NR band region, they are all tagged as asymmetric as expected since these are surface events. Very low yield recoils from ^{206}Pb are well separated from β 's due to ^{210}Pb and ^{210}Bi decays which show up in mid-yield. The black vertical line represents the $8 \text{ keV}_{\text{nr}}$ threshold, and the hyperbolic solid black line is a $2 \text{ keV}_{\text{ee}}$ ionization threshold. c) View of the ionization yield vs ionization z-symmetry parameter. This shows very clearly that the signal region (green) is very well separated from not only ERs (y-axis) but also from surface events (x-axis).

4.2.4 Phonon Based Discrimination

Ionization is not the only tool the iZIP provides to discriminate between background and signal. Indeed at low energies, since the ionization signal is quenched, it becomes far less useful since the signal will become very similar in amplitude to the noise. Since the phonon signals do not suffer from this quenching effect, and the phonon channels have lower noise than the ionization channels, background discrimination based on phonon information can be very important at low energy. Thanks to the layout of the phonon channels the iZIP has the ability to provide radial as well as z asymmetry discrimination as well. At high energies ($> 20 \text{ keV}$) ionization discrimination usually performs better

thanks to the clean signal afforded by the rapid speed at which the electron-hole pair drifts across the crystal. The phonons by contrast, diffuse throughout the crystal and as a result partition information is not as pronounced. This is also compounded further by the optimal filter template used. Recall that for phonon reconstruction a single template based on bulk ERs was utilized. There is however a pulse-shape difference between events in the bulk and those that occur near the surface as can be seen in Figure 4-7 which shows two examples for a 60 keV event. This shape primarily is due to the presence of prompt and Luke-Neganov phonons that are produced at the surface. These phonons occur on short time scales and contribute to the sharper rise-time, whereas the slower ballistic phonons diffuse throughout the crystal resulting in the typical rounded pulse shape. As a result of this template mis-match the OF energy is often under-reported for these side-split phonon signals, diminishing the separation of the phonon z-symmetry quantity.




Figure 4-7. Bulk vs. Surface Event Phonons. A) A bulk event with 60 keV of phonon energy. These events have a rounded peak. B) A surface event with 60 keV of phonon energy. Notice the sharper peak for the event at the surface.

In addition to pulse-shape systematics, the phonon z-symmetry is impacted by another phenomenon called local saturation. Local saturation occurs for surface events which deposit enough energy in a localized region of the detector to cause some of the individual TES sensors in a given phonon channel to saturate. The net effect is an underestimation of the phonon energy from that region, reducing the total reported energy

for the side of the detector with the surface event. As a result, `pzpart` merges into the bulk as energy increases as shown in Figure 4-8.




Figure 4-8. The phonon z-partition (blue) is shown as a function of total phonon energy compared to the ionization z-partition (red) for IT3Z1, which has an implanted ^{210}Pb source on its top side. Several differences in the two z-partition estimators are immediately visible. The spread between surface and bulk events is much wider for most energies in the ionization case since phonons diffuse throughout the entire crystal even for a surface event, whereas e^-/h^+ pairs do not. Also evident is the effect of local saturation on the phonon distribution. The bulk band is roughly centered around a z-partition of 0, however the surface event band can be seen to decrease as a function of energy. The produces roughly a factor of 2 reduction in z-partition estimator at high energy compared to the low energy region.

Application of both ionization and phonon symmetry cuts can be very useful in catching events that leak through an ionization symmetry cut (Figure 4-9). Since phonon information is more robust down to lower energy, radial and z partitions in phonons are excellent choices to use as fiducializers for low-mass WIMP searches and will be discussed further later in this work in the context of a low-mass WIMP search.




Figure 4-9. An example of how phonon symmetry can outperform ionization symmetry discrimination at low energy. A) The phonon symmetry plane, showing the total phonon energy collected on Side 2 vs Side 1. Symmetric events should occur near the line $y = x$ in this plane. Events in blue are tagged as symmetric events via ionization symmetry whereas events in red are tagged as asymmetric based on ionization. Solid black lines show a rough phonon symmetry cut distinguishing asymmetric vs symmetric events via phonon information. At high energies an asymmetry in ionization and phonon signals roughly agree but below about 10 keV a distinct population of blue events occur in what is clearly an asymmetric band. These are events that are mis-tagged as symmetric according to ionization signal but which are actually asymmetric. A high energy outlier occurs around 35 keV on S1. B) The same events are plotted in ionization yield vs recoil energy. The solid black hyperbolic line is a $2 \text{ keV}_{\text{ee}}$ ionization threshold and the vertical solid black line is an $8 \text{ keV}_{\text{nr}}$ recoil threshold. The high energy outlier has a yield that is dangerously close to the NR band. While it leaks through an ionization symmetry cut, it can be easily removed with a phonon symmetry cut.

CHAPTER 5

DATA QUALITY AND BLINDING CUTS

No experiment takes perfectly usable data 100% of the time even if it is a goal that experimentalists strive for. Factors outside of the experimenter’s control such as environmental noise, random failures or just intrinsic detector response can create periods of time or types of data that are not useful for a physics analysis. To avoid having a massive dataset populated with questionable events, we use data quality cuts so that only the best quality events from periods of time that are well-behaved will be used for a WIMP-search analysis. Additionally, it is important to blind the dataset so that while tuning cuts that discriminate based on criteria that select WIMPs one does not accidentally bias themselves into selecting what they believe a WIMP should look like. If done properly, one can tune cuts to select against non WIMP-like events and then unblind (or, “open the box”) and see what the final cuts select. There are a great many different cuts that are developed in the course of an analysis and in this chapter I will describe some of the many important data cuts implemented in SuperCDMS Soudan for both the low-mass and high-mass WIMP search efforts, as well as a brief mention of the software implementation for the analysis effort.

5.1 Analysis Framework and Cut Nomenclature

SuperCDMS Soudan primarily uses MATLAB for data analysis purposes, however ROOT and Python are also used on occasion. The actual data is stored in ROOT files but a set of custom MEX files allows analyzers to load data into a MATLAB environment. Many years ago in the CDMS II era, a special analysis package was developed called CAP¹ which aimed to help manage memory and prevent accidental use of incorrect or stale data through buffering. RRQs and cuts were written in this language and managed

¹ CDMS Analysis Package

by the FCCS² which tracks dependencies and revision dates to avoid situations where a user might modify a function but use stale buffered data. Much of the core of this code was written in turn of the century MATLAB code which eventually caused numerous slowdowns by the time the SuperCDMS Soudan data was being analyzed due to changes in MATLAB and in the size of the datasets being analyzed.³ Eventually a small team consisting of Todd Doughty, Danielle Speller, Bruno Serfass, and myself worked on modernizing CAP in 2014, resulting in MatCAP. The full details of this new analysis package are beyond the scope of this work, but it has resulted in a significantly faster, less memory intensive, and more flexible analysis framework for the future.

The underlying nomenclature is generally the same in MatCAP as it was in CAP. All cuts that are developed begin with the letter “c” and their name gives an indication of the type of event they select. So for example a cut with the name “cGoodEvents” would be defined internally to return to the analyzer a logical mask⁴ that would select good events, and cut out bad events. For some cuts it makes more sense to define them in terms of things one wishes to remove. For example, a cut named `cBadPulseShape` might internally follow a simple set of logic to select events with a bad pulse shape. The analyzer would then perform a logical NOT on this mask to select all events that do not have bad pulse shapes. In practice we cannot make cut names contain plain text so often they are abbreviated (e.g., `cGoodEv` for the good events cut). Additional information is used via suffixes to the cut names. To keep track of whether a cut depends on the `cdmsbats` production version cuts in this work may be suffixed with “_v53” to denote it

² Function Code Control System

³ The old FCCS tree grew to an enormous size where it would take half an hour to regenerate, a process that had to be performed whenever a user modified anything. This made cut development cumbersome.

⁴ A vector of true or false values.

depends on production version 5.3. If the cut is invariant to the production version they can be suffixed with “_133” to denote a run dependence.⁵ Finally if a cut was specific to the low-mass WIMP search it would receive a suffix of “_LT” and cuts specific to the high-mass search “_HT”. An example of a cut is `cTriggeredEvent_133_HT`. This cut tells us that it selects events that issue a trigger, its logic depends on the particular run and not `cdmsbats` production, and it is specific to the HT analysis.

With the development of MatCAP one additional new feature that exists are the storage of cut decisions in common external files. Previously in CAP (which was used in the LT analysis) cuts were always computed on the fly in the analysis framework. They were stored in the buffer system once computed, but this was not very portable. In MatCAP cuts written in MATLAB obtain one extra final suffix at the very end (“_dev”) to denote they are a “development” version. These are periodically computed for all the data types on a computer located at FNAL and saved to ROOT files. MatCAP then constructs a special cut namemap, similar to the data namemap to allow users to call cuts from the pre-computed ROOT files by simply calling the cut without the “_dev” suffix. This saves time in application of the cuts as they do not need to be computed even when starting fresh.⁶

5.2 χ^2 Rejection Cuts

One of the simplest and powerful sets of cuts that can be constructed are so-called χ^2 cuts. These examine the χ^2 of the OF fit as a function of reconstructed amplitude and can reject events that have particularly poor fits. In SuperCDMS Soudan we used such a cut

⁵ It turns out that there are no cuts that have “_133” and “_134” versions. Instead the cut handles differences from run to run internally.

⁶ Though practically, there are still bugs to work out. On more than one occasion the ROOT cuts have differed from the development cuts in insidious ways. Sometimes cuts are written in odd ways that cause different behavior when run on individual datasets compared to multiple ones simultaneously, or cuts will somehow fail silently during generation and give default values of false.

for the ionization and phonon signals. For the SuperCDMS Soudan dataset, the primary χ^2 cuts are defined on the `qsum#OF` quantities⁷ and on `ptNF` and `ptOF`. An example of the cut behaviors is shown in Figure 5-1 including an example of the type of event they reject.




Figure 5-1. The behavior of the ionization and phonon based χ^2 cuts is seen here. Both exhibit the same characteristic flaring upwards as energy increases. A.) An example of a phonon χ^2 distribution with boundary fit by a quartic polynomial and the type of pulse this cut rejects. All data points below the green line are passed by this cut. B.) An example of an ionization χ^2 distribution and an anomalous pulse shape it rejects. All events in red are passed by this cut. Courtesy of Adam Anderson [92] and Todd Doughty [93] respectively.

⁷ Since we had many S1 ionization shorts, on detectors where this happened it was computed for whichever channel was working on that side

One striking feature of the χ^2 distributions is that they flare with energy. To zeroth-order this may not be expected, however it is thought [94] to arise from a combination of effects. In the ionization case the belief in CDMS II was that it was a result of the fact that the templates are constructed from a finite number of sampled traces each with their own noise which is then scaled when fitting a template to an event of some high energy, and also as a result of the finite start-time resolution of each fit. These appeared to predict a quadratic fit would be most appropriate which did seem to be the case. The interpolation scheme employed here reduces the energy dependence of the ionization χ^2 dramatically but nonetheless it still exists. In the phonon case this effect is exaggerated more so since unlike the ionization pulse shape there is no well-defined characteristic shape. Additional systematics arise from the fact there are pulse-shape differences between surface, bulk, ER and NR events. Since the phonon templates are derived from bulk ER events they may have slightly worse performance on NR events, but this appears to be a marginal effect. Luckily the total pulse shape OF's appear to reduce this somewhat. Each cut was constructed by binning in energy and then selecting points equivalent to $\sim 3\sigma$ above the mean χ^2 in that bin. In the phonon case a quartic polynomial fit was performed, and a quadratic was used for the ionization.

5.2.1 Glitch and Low-Frequency Noise Cuts

A χ^2 cut is also useful for discriminating against other type of events that are abnormal. As long as one has an OF template, one can construct a χ^2 quantity for each template and then make a cut that rejects events that are better fit by one of these unusual noise traces. Two primary classes of abnormal event occur regularly enough that it is worth expending the computational time to run an additional OF over each event to examine their event type. One particular type is referred to as a “glitch” event. For glitch events often multiple detectors will simultaneously trigger with extremely sharp spikes on

all readout channels.⁸ This cut is tuned such that $\sim 99\%$ of randoms pass at low energy, and at higher energies a quadratic fit in $\Delta\chi^2$ was created using 2 keV energy bins from neutron calibration data in the $\Delta\chi^2$ plane.[95] A similar approach was taken for the low-frequency noise cut. These low frequency noise events were very problematic for the low-mass analysis since they tend to reconstruct from the regular OF approaches to low energy and thus constituted a significant contamination initially. Much work was done by Kristi Schneck to develop templates for these events, and more information can be found in her thesis [65]. The net approach for the LF noise cut passed $\sim 95\%$ of randoms at low energy and followed a quadratic shape at higher energies. Further details are described in Figure 5-2.

5.3 Other Quality Cuts

There are a large number of quality cuts that have been developed for the SuperCDMS Soudan iZIP detectors. I will briefly discuss a few of them below before describing in depth a few other cuts that remove certain types of data, and the blinding cut.

5.3.1 Event Based Quality Cuts

This class of data quality cut selects events to pass or fail on an event by event basis typically by examining information related to the pulse shape or timing. The χ^2 cuts mentioned previously are examples of this type of cut.

- **Pre-Pulse:** Recall from Chapter 3 that the raw traces contain information about the channel prior to the actual triggering pulse. This pre-pulse information can be used as a quality metric to remove excessively noisy events. In much the same way that the random noise traces were used to form a series dependent characterization of the noise, a similar technique can be applied to the pre-pulse standard deviation RQs and a cut devised based on the pre-pulse standard deviation value.[97]. For this work the threshold is set such that only pre-pulse deviations greater than 5σ away from the mean value are removed. This type of cut is generated for each phonon

⁸ This is a good way to discriminate in a non-template based way. If a large number of phonon and ionization triggers are issued simultaneously, it is likely to be a glitch and can be rejected outright.




Figure 5-2. The existence of glitches and low frequency noise that could not be removed by other quality cut metrics necessitated the development of special OF templates for each type of event so a χ^2 discriminant could be developed. A.) The glitch χ^2 cut is formed by setting a cut boundary in a $\Delta\chi^2$ vs energy plane. At low energies of a few keV the cut is flat and passes 99.7% of randoms. At higher energies it is defined by a quadratic function based on californium calibration data. Events below the magenta line are tagged as glitches. Courtesy of Yu Chen [95]. B.) The low-frequency cut plane. Similar to the glitch $\Delta\chi^2$ plane this cut is flat at low energies and passes approximately 95% of randoms. At higher energies a quadratic band is also defined from californium calibration data, until 30 keV where the cut is again flat. Events below the green line are tagged as low frequency noise. Courtesy of Kristi Schneck [96]. C.) Example of the glitch OF template compared to the regular OF phonon template. D.) Example of a particular low-frequency noise template compared to a regular OF phonon template. Each detector's LF noise template is different.

and ionization channel and are contained within the cuts `cPstd_v53` and `cQstd_v53`. Both of these cuts appear to have a very high passage of livetime of $\sim 99\%$

- **Good Phonon Start Time:** The cut `cGoodPStartTime_v53` addresses a small bug in the Prodv5-3 data release (R133 data). The OF search window did not extend to times long enough prior to the global phonon trigger to catch some low-energy phonon pulses on a few detectors. In subsequent data processing releases the phonon OF window was adjusted from $[-100, +200]\mu s$ to $[-200, +100]\mu s$ relative to the global trigger, and the ionization OF search window was allowed to float relative to the total OF delay. The lower bound of this ionization OF window was set by eye. This cut can be used to reject pile-up and to reject total OF delay greater than the maximum edge of the ionization OF search window, effectively removing events at the high end of the phonon OF search window.
- **Triggered Events:** Not every event on a detector has an associated trigger associated with it because we do not have 100% trigger efficiency down to arbitrarily low energy. As such an event below trigger threshold in one detector can be read out if it is coincident with an above threshold trigger in another detector, and as a result can make computing the trigger efficiency a bit tricky. The cut `cTriggeredEvent_133_HT` is meant to impose the constraint that an event also have a trigger bit set in whichever detector it belongs to. This cut accomplishes this by issuing a query, via some java code written by Joel Sander, to a data quality database to pull accurate trigger settings. This is important because they have changed so often, human error can easily result in a mis-classification when writing up code, and the problem is compounded more when one remembers each detector is readout as two half-detectors, each of which can be in the PloOR or PloAND triggering mode. For the LT analysis a human created version of the cut existed, `cTriggeredEvent_133_LT`, which did not query the database and it did indeed have some disagreement with the output of this (correct) cut, however it did not impact the LT analysis results. This cut is not inside any umbrella cut and must be called explicitly.
- **Saturation:** For events that arise from extremely energetic particles, such as α 's, the energy deposited is enough to saturate the readout channels. The reconstruction algorithm will still run on these events and return non-sensical values. The cuts `cPsat_133` and `cQsat_133` remove events where at least one of the readout channel ADCs were saturated.
- **Square Pulse:** This cut removes events on the two detectors that have anomalous baseline shifts. These are PDS1 for IT2Z1 and PCS1 for IT3Z3. When viewing the raw trace this looks like a square pulse, but it appears in actuality to be a large offset in the baseline noise that gives it this appearance. The cut `cSquarePulse_v53` is actually composed of 3 internal cuts. One determines if the phonon signal integral is more than 2σ below the mean aggregate value. Another examines the distribution of the baseline value itself and issues a cut if the baseline level is outside the typical

value (this was tuned by eye it seems). The final component examines the time it takes the pulse to fall to 20% of its maximum value and tags an event if this is 0. This is an especially strange case since a fall time of zero would imply an instantaneous drop, however it is an artifact of the algorithm which outputs a 0 if it failed or found a negative max amplitude. This could be likely if the baseline offset is large and negative and then rose near the trigger window, making it seem like an inverted phonon pulse perhaps.




Figure 5-3. A directed graph of the good events cut `cGoodEv_v53`. This cut is broken into 3 main sub-cuts generally by whether the cuts remove events based on pulse reconstruction characteristics or by other criteria that impact the quality of data during the time period.

5.3.2 Time Period Quality Cuts

These cuts remove periods of time, or whole series from consideration based on the behavior of event distributions as a function of time. Many of these are typically what are referred to as livetime cuts since they will remove livetime and do not judge events based on whether they are reconstructed properly or not.

- **Bad Series:** Some series occur during bad periods of time, or were flagged as rather questionable in the offsite DQ page, but processed anyways. The cut `cBadSeries_133` removes specific series that occur during either times of large environmental noise at Soudan, or where there are clearly documented abnormal

behavior in detectors. A particularly prominent block of time removed was August 2012 when a large amount of ionization noise was present due to a bad AC filter at Soudan. This cut removes only about 1% of the total livetime.[98]

- **Bad OF Resolution:** When the OF quantity is computed, it also outputs a measure of the resolution of the OF. If the noise template for a given series has a high amplitude then this can produce a very large resolution. This cut will reject series where the total phonon OF resolution or the total outer phonon channel OF resolution are greater than 25% higher than the mean. Since noise varies with time, this cut also relies on time blocks (the same ones used for the multiples and blinding cut) when computing the mean OF resolution. The outer channels are included since they were important for radial discrimination in the LT analysis. This cut is ultimately a livetime cut since it removes series, and removes up to $\sim 5\%$ of the raw livetime for some detectors.[99]
- **Base Temperature:** Another important data quality cut is `cGoodBaseTemp_v53`. This cut selects only periods of time when reliable information on the fridge base temperature is known, and when the base temperature matches those with which the phonon energy scale is calibrated against. Ideally this would not remove much of anything, however occasional hardware errors with the LR700 resistance bridge, or control computer resulted in gaps in the temperature record, and in some cases the base temperature was elevated beyond the point the phonon energy scale was calibrated against. This cut is also a livetime cut and removes only about 7% of the total raw livetime. [100]
- **Trigger Bursts:** Occasionally the trigger rate can become extremely large within a data series due to transient noise pickup.⁹ Since most of the data is bound to be due to noise during this period there is no reason to keep it and so `cTrigBurst_133` removes periods of time during which an a large number of triggers occur relative to the typical number for a given series. This cut is also a livetime cut and removes $\sim 4\%$ of the raw livetime.[101]
- **Good Bias Time:** Occasionally the detectors may exceed the accepted 3 hours between periods of flashing with LEDs to restore neutralization. This can happen either on purpose (for special studies) or through glitches in the DAQ that cause a skipped flash. The cut `cGoodBiasTime_133` will remove series with greater than 183 minutes since the last flash, or in the case of calibration data, an additional requirement will remove the series if it has been more than 25 minutes since the last pause to ground the detectors.

⁹ Evidence suggests this is related to the health of the cryocooler, which can induce a significant amount of vibration into the towers causing excessive noise triggers.

- **Finalized Phonon Settings:** The earliest portion of the dataset contains some series in which IT1Z3, IT2Z1, IT2Z3, and IT4Z1 had not yet finished QET and SQUID tuning. Any series before 1837 on March 14, 2012 are removed for these detectors via `cFinalPhononSettings_133`.
- **Trigger Errors:** On occasion the DAQ will suffer an error that results in a nonsensical state due to a missing bit of information (such as which detector issued a trigger). In such cases an error bit is set to true in the DAQ and `cErrMask_133` will allow removal of these events.
- **GPS Timing:** SuperCDMS Soudan uses a GPS for precise timing information. On occasion a glitch can occur in which the GPS timing information disagrees with the recorded EventTime (unix time). The cut `cBadGPStime_133` removes events in which the discrepancy between these two timing metrics is larger than 17 seconds.[\[102\]](#)
- **NuMI Coincidence:** The NuMI experiment is active on occasion and can result in high energy neutrinos passing through the detectors. This cut removes time periods in which the NuMI beam is active. In general this is not an occurrence that is expected to be significant in this work, especially since the NuMI beam was only active for short periods of time.[\[102\]](#)
- **Trigger Enable Bit:** Another error mode that can arise very infrequently is the absence of a trigger enable bit in the data stream. This results in 0 livetime being recorded for the event and as such it is removed by `cNoTrigEnable_v53`.
- **DC Offset:** Occasionally a SQUID can lose its lockpoint or may have reduced sensitivity due to trapped flux. This results in phonon pulses with reduced amplitudes at best or no phonon pulses in the worst case. This state of affairs is often associated with an abnormal DC voltage offset in the phonon channel and provides a handle with which to reject such time periods. The cut `cGoodDCOffset_v53` will only allow events to pass in which the channel DC offset is within $\pm 25\%$ of a mean value. Interestingly while most detectors only lose 3% of livetime as a result, IT4Z3 loses 12% of its raw livetime due to this cut. While livetime is not as useful a metric in calibration data as WS data due to the high trigger rates, it is worth noting that IT5Z3 loses a significant portion of its ^{252}Cf livetime, $\sim 70\%$, for the August 22, 2012 neutron calibration series.
- **Bad LED:** For a few detectors the status of the LEDs used for neutralizing can cause instability in the ionization signal. In particular for IT3Z2 and IT4Z2, if all of the 4 LEDs for that detector are non-functional, the data should be removed (this logic does allow for one or more LED failures to be permissible). Interestingly for IT3Z3, if the LEDs are on the detector appeared to perform worse. This cut therefore

removes data if all four LEDs or LED 2,3 and 4 are on. For data taken with just LED 1 on, no action is taken. The cut `cBadLED_v53` is used to handle these cases.¹⁰

- **Stable Trigger Thresholds:** This particular cut exists because of the high variability in noise on each detector. Over the course of the experiment the trigger thresholds on the RTF boards had to be adjusted many times, both up and down. This is problematic when the duration of time between trigger threshold adjustments is not long enough to accumulate enough statistics to characterize the trigger efficiency. For the LT analysis the cut `cStableTrigThresh_133_LT` was used to remove these time periods, and for the HT analysis `cStableTrigThresh_133_HT` was used.
- **Random Trigger χ^2 :** This cut is a successor to the RS test (described in Section 5.7) used in the LT analysis. One feature that was observed during the χ^2 cut development was the presence of multiple bands in the χ^2 distributions, often taking the form of a population with abnormally low χ^2 in the fit.¹¹ It is believed that periods of poor electronic noise can allow such situations to occur and as such should be removed. The RS tests aimed to do this but proved to require a large amount of by hand tuning that would not be appropriate for the larger HT dataset. By examining the χ^2 distribution of randoms it was found that series with abnormally low χ^2 could be removed effectively. The cut `cGoodRandomChi2_v53_HT` removes approximately 4-10% of the total livetime across various detectors.^[103]
- **Bias Voltage:** During certain periods of time the voltage on some detectors was altered for special studies or to take reverse bias data. Typically the source detectors have the most bias voltage flips, however IT5Z3 also has its bias flipped during periods of time in R134 and R135 when IT5Z2 was in CDMSlite mode.¹² The cut `cBiasPM_133` enforced standard polarity operation in the LT analysis, and `cWSBias_133_HT` for the HT analysis.
- **CDMSite Mode:** Because of the fashion in which CDMSlite detectors are tagged in the data a special cut was required to avoid periods of time during which CDMSlite was active on a particular detector. This mainly impacts IT5Z2, however IT2Z1,

¹⁰ Given the some of the unusual behavior of IT3Z3 over time it may not have remained true that only using 1 LED was optimal.

¹¹ Ideally it should come out to be around 4095 for phonons, and 4095 for ionization channels since the 2x2 charge OF χ^2 is computed by adding the individual channel χ^2 for a given side. This is just one less than the number of bins.

¹² It was discovered that the high voltage on IT5Z2 was distorting the field lines on IT5Z3 causing a the yield to be reduced over time. Operating in reverse bias mode alleviated this problem somewhat, but the actual field strength on IT5Z3 would not be the standard ± 2 V, leading to its exclusion from the HT analysis.

IT4Z1, and IT5Z3 also have some periods of time during which CDMSlite testing occurred on them. The cut `cCDMSlite_v53` removes these series.




Figure 5-4. A directed graph of the good events cut `cLiveTime_v53_HT`. This cut has many dependancies that remove data based on environmental factors, not necessarily related to the data quality itself, however some data quality cuts are present (such as the phonon χ^2 cut).

5.4 Multiples Cut

Since WIMP-nucleon interactions are expected to be very rare, it is even more unlikely that a scenario will occur where a WIMP will scatter off of a nuclei in one detector and then scatter off the nuclei in another detector. This allows for another important type of discrimination cut to be developed that will reject events that show a signature consistent with having multiple detector scatters, or put another way, we would only consider an event a potential WIMP if only a single detector showed evidence of an event above threshold. This discrimination between multiples and singles requires an energy threshold for each of the detectors to be specified. In its simplest form, the logic in a multiples cut is as follows: For a given event on detector D (the detector of interest), examine the energy on all other detectors (the vetoing detectors), and if any of the other

detectors has energy above some threshold, consider the event a multiple candidate between detector D and the other detectors that were above threshold. The application of an energy threshold on D is needed as well to indicate whether or not the event is itself just noise or a physics event. In the past for CDMS-II the energy threshold for other detectors was set as 4σ above the mean zero-energy noise. However as we saw in Section 3.10, the phonon noise distributions from zero energy randoms is both non-Gaussian, and not time invariant. One can circumvent the non-Gaussian nature by relying on the CDF of the noise distribution: simply set the threshold such that some percentage of the randoms occur below the noise threshold. For this work, the threshold was set such that 99.87% of randoms are below this threshold.¹³ An typical non-calibration series has approximately 4000 events within it, with about 1300 of these being random triggers. This would result in there being approximately 1 random trigger above the threshold, but clearly with abysmal statistics. To ensure the tail of the randoms distribution is somewhat better represented and to keep a time-varying threshold, series were assigned to time blocks (i.e., bins) based on their noise performance and a set of parameters devised by Elias Lopez. Time blocks are formed from clusters of series based on the following set of criteria:

- A threshold is constructed per series based on a combination of the noise mean and resolution
- A cluster of series is formed such that it contains 20 series to allow for sufficient statistics in the tail
- A time block is formed by grouping consecutive clusters
- The difference between the max and min threshold in the proposed time block must be less than 20% of the minimum
- The series boundary of a time block is set at the series that maximizes the difference in threshold between two consecutive clusters

¹³ If the noise distribution was identical to a Gaussian distribution, this would be equivalent to setting a threshold 3σ above the mean.

- Up to two series may be removed from a cluster, if they are clear outliers

Once the time blocks are formed, the noise means and sigmas were recomputed on a time block by time block basis since one cannot simply average over the series by series noise means and resolutions as computed earlier. Once the time block dependent noise means and sigmas are in hand, the prescription for generating the threshold as described earlier is performed.




Figure 5-5. A single time block for detector IT2Z1 with the threshold as determined by requiring 99.87% of all randoms to occur below. To show the behavior of the tail of the distribution, the plot is constructed as $\log_{10}(1 - \text{CDF})$. Shown in blue is the expected behavior for a perfect Gaussian distribution. The threshold is indicated by the vertical magenta line, and in this case is set such that the threshold is 6.2σ above the noise distribution mean, corresponding to an energy threshold of ~ 1.5 keV.

All of the thresholds are computed for the quantity `psumOF` on the time block basis¹⁴ with a couple of modifications for two detectors. For IT1Z1 the thresholds are based off

¹⁴ This in retrospect makes life more difficult, since the analysis uses `ptNF` and `precoiltnf`, which behave slightly differently than `psumOF`.

of a 7-channel `psum0F` which has the individual channel energy `pa20F` subtracted due to excessive noise, and likewise (but for a different type of noise), `IT3Z3` has individual channel energy `pc10F` subtracted away. Further complicating matters is that `IT5Z2` was operated in the special `CDMSlite` mode occasionally during which times only a 4-channel `psum0F` quantity and threshold, could be computed. At each step of the process the appropriate `psum0F` quantity was used in computation of the noise means and resolutions.

Since the multiples cut would be used as part of the blinding cut as well, a few other considerations are required for its construction. As will be discussed shortly, only single scatters were blinded, and all multiples were kept unblinded for examination. This philosophy in part is why the above requirements are fairly strict: If one has a very strict definition of what a multiple scatter is, the logical inverse of this selection criteria provides a very loose selection of singles, and it is generally better to slightly over-blind rather than under-blind. This also led to a few other requirements on what a multiple scatter event can be in addition to an energy threshold requirement. The event must also not be tagged as LF noise in any detector, and it must not be a glitch event either. A flow chart of the multiples selection logic is shown in Figure 5-6 indicating it is a fairly involved process.

The implementation of this cut (called `cPmultTight_blind_v53`) proved to be quite challenging due to the sheer amount of edge cases that occurred from Run to Run, and detector to detector. Due to the sheer number of iterations required the code was eventually optimized to be as quick as possible without using an enormous amount of memory by relying on a series of logical ORs rather than accumulating a giant matrix of decisions.¹⁵ When referring to events categories, multiples are meant to indicate those events passing this cut, and singles are those events failing this cut.

¹⁵ The actual cut can output the decision matrices if desired, which proved useful for debugging purposes




Figure 5-6. The logic flow of the multiples selection cut is shown here, where the decision process starts at the top. The region in the grey box should be interpreted as a loop over each of the vetoing detectors. One can view this as accumulating a set of true/false decisions in an $N_{event} \times 14$ matrix. An event is only considered a multiple scatter candidate if, on a vetoing detector, it is above threshold, not a LF noise event and not a template based glitch. After obtaining the multiple candidacy decision for each of the vetoing detectors, a logical OR is taken across this matrix to produce a single event by event set of true/false decisions. A final logical AND with a trigger based glitch cut removes any events from consideration that are not really multiple scatters. The final output then is a set of event by event decisions on whether an event is a multiple scatter.

After the cut was constructed a check on its efficiency was performed.[\[104\]](#) Since actual single scatters in the signal region are blinded, a proxy is used. The random triggers essentially represent a population of zero-energy singles. As such the passage fraction of the logical inverse of the cut (`-cPmultTight.blind.v53`) on randoms should be very high. The expected efficiency depends on whether the noise on the detectors is independent or correlated. We can consider extreme cases to provide an upper and lower bound on the expected efficiency (or passage fraction in this case) of this cut on the randoms. Since the threshold is set such that the fraction of randoms above threshold is $F = 1 - 0.9987$, if the noise across all detectors is perfectly correlated, then the expected fraction of randoms above threshold can at most be just $P = F$, representing the case where if one detector is above threshold for an event, all detectors are above threshold. The corresponding singles efficiency is then $1 - P = 0.9987$. Conversely if the noise is perfectly independent across all 14 vetoing detectors, the expected number of randoms being tagged as a multiple can be computed as $P = 14F(1 - F)^{13}$, giving a singles efficiency of $1 - P = 0.9821$.

The situation is more subtle however since the multiples cut also depends upon the glitch and LF noise cut efficiency on randoms. Generally it is seen the glitch cut passes $\sim 99\%$ of the randoms, however the version of the LF noise cut implemented in `cPmultTight.blind.v53` has only a $\sim 95\%$ pass rate on randoms (meaning, it can misidentify a random as LF noise about 5% of the time). Thus, ignoring the glitch cut efficiency, an expected range of randoms passing the singles selection criteria is dependent upon whether both the randoms and LF noise are independent or correlated amongst themselves. Recall that to be a multiple an event must be above threshold and not LF noise, thus if we let the probability of being tagged as a LF noise event be $G = 0.05$, the probability that on a single detector a random is tagged as a multiple is $P_1 = F(1 - G)$. If this is totally independent across all detectors, then following the same logic as earlier, we can arrive at the probability of observing 1 multiple out of 14 detectors

Table 5-1. Singles Efficiencies on Randoms

	Fully Independent Threshold	Fully Correlated Threshold
Fully Independent LF Noise	0.9830	0.9994
Fully Correlated LF Noise	0.9830	0.9988

as $P = 14P_1(1 - P_1)^{13}$, yielding a singles efficiency of $1 - P = 0.9830$. If we go the other direction and assume the noise behavior and LF noise tagging behavior are perfectly correlated, then if one detector is flagged as a multiple, all will be, resulting in a singles efficiency of $1 - P_1 = 0.9988$. There exist cross-cases as well, where just the threshold behavior is perfectly correlated but LF noise is independent, or where the thresholds remain independent but the LF noise is maximally correlated. In the first cross-case if any detector is above threshold all will be, and the probability of being a multiple depends upon none of the 14 detectors being LF noise in this case. This gives a probability of being tagged as a multiple of $P = F(1 - G)^{14}$ and a singles efficiency of $1 - P = 0.9994$. Similarly if the threshold behavior is independent and the LF noise behavior is totally correlated, we require just 1 detector to be above threshold and no LF noise to be present. The probability of being tagged as a multiple is therefore $P = 14(1 - F)^{13}(1 - G)$, giving a singles efficiency of $1 - P = 0.9830$. This behavior is summarized in Table 5-1 but indicates that the range of efficiencies should occur somewhere between 98.3% to 99.9% where we see the behavior of the randoms relative to the noise threshold is the primary driver for the efficiency, with changes in LF noise having small influences. The typical efficiency over the dataset considered in this work is about 99% for all detectors with only a few series having lower efficiencies due to noisy time periods. This is suggestive that as hoped the noise is fairly independent across detectors.

The efficiency in the actual data is computed by taking the ratio of randoms tagged as not multiples to the total number of randoms, and is computed on a series-by-series basis. In general the passage fraction of a particular cut C on some general population defined by cut S is just given by the sum of events that pass over the total, since events that pass selection criteria are represented with a 1 and those that fail a 0 this allows one

to sum up the cut vector and compute the passage fraction:

$$Pf = \frac{\sum C \wedge S}{\sum S} \quad (5-1)$$




Figure 5-7. The singles passage fraction on detector IT2Z2 for R134 is shown on a series by series basis. Series with large drops in passage fraction are typically due to excessive noise on one of the vetoing detectors and when examining all 15 detectors would show up as having a less severe drop in efficiency for the series compared to others. Here the average passage fraction is about 98.9%, a typical value for all detectors across the c334 dataset.

5.5 Blinding Criteria

Before analysis can begin in earnest, one of the first tasks that must be performed is to blind the data to avoid introducing any bias when tuning discrimination cuts. This proves to be a technically challenging process, especially for never before used detectors. One does not want to over-blind the data, as then useful sidebands could be obscured. At the same time, it is worse to under-blind and accidentally allow yourself to be biased. A perfect blinding where only signal is masked is not possible (if it was, then we would already know how to select the signal!) but a set of conservative blinding requirements were created after study on data from calibration sources and randoms. This allowed development and tuning of certain data quality metrics and thresholds. Since SuperCDMS Soudan planned to perform two blinded analyses, for low and high mass WIMPs, a dual

Table 5-2. List of General Blinding Requirements LT

Cut	Definition	Comments
Not Template Glitch	<code>cGlitch1_v53 = false</code>	
Not Trigger Glitch	<code>cGlitch_133 = false</code>	
Not LF noise	<code>cLFnoise1_v53 = false</code>	
No Veto	Veto Timing $\notin [-250]\mu\text{s}$	
Energy Limit	<code>ptNF < 10(1 + 0.078V_b) keV</code>	Lindhard model, 13.1 keV
Minimum Energy	<code>psum0F > E_{mult}</code>	Multiples threshold
Not Multiple Scatter	<code>cPmultTight_v53_blind = false</code>	
LT Detectors	1101,1104,1105,1111,1112,1114,1115	

Table 5-3. List of General Blinding Requirements HT

Cut	Definition	Comments
Not Template Glitch	<code>cGlitch1_v53 = false</code>	
Not Trigger Glitch	<code>cGlitch_133 = false</code>	
Not LF noise	<code>cLFnoise1_v53 = false</code>	
No Veto	Veto Timing $\notin [-250]\mu\text{s}$	
Upper Energy Limit	<code>precoilsum0F < 150 keV</code>	
Radial Inner	<code>qi[k]0F $\geq \mu + 2\sigma$</code>	Time Dependent Threshold
Radial Outer	<code>qo[k]0F $\leq \mu + 8\sigma$</code>	Logical OR with Inner
Z Symmetry S1	<code>qsum10F $\geq \mu + 2\sigma$</code>	Time Dependent Threshold
Z Symmetry S1	<code>qsum20F $\leq \mu + 8\sigma$</code>	Logical OR with Previous.
Z Symmetry S2	<code>qsum20F $\geq \mu + 2\sigma$</code>	Time Dependent Threshold
Z Symmetry S2	<code>qsum10F $\leq \mu + 8\sigma$</code>	Logical OR with Previous.
NR Band	$\bigvee_{k=1}^3 \text{ysum0F} \in 3\sigma_{NR}$	Exist in any 3σ NR band
NR Threshold	<code>precoilsum0F > 0</code>	
Minimum Energy	<code>psum0F > E_{mult}</code>	Multiples threshold
Not Multiple Scatter	<code>$\neg$cPmultTight_v53_blind</code>	
Not in Exception Period	<code>$\neg$cDoNot_blind_v53</code>	

[k] = 1,2

blinding scheme was constructed understanding that the low-mass WIMP search would occur first.

In addition to the above requirements, certain parts of the dataset were kept open for cross-checking purposes as they were not suitable for a low-mass WIMP search. These include all random triggers, and events 10 days after any neutron calibration in 2012 and 20 days after a neutron calibration in January 2013. For the high-mass WIMP search the blinding criteria were a bit more numerous as were exceptions.

The list of exceptions to the high-mass search blinding criteria mainly cover special study modes on a few detectors and compose the contents of the `cDoNot_blind_v53` cut:

- Anything before March 2, 2012 is considered commissioning data
- Random triggers
- Events from IT3Z1 and IT3Z3 prior to June 1, 2012 (used for iZIP discrimination study[74] and proposals)
- Events from detectors operating in reverse bias mode ($\mp 2V$ instead of $\pm 2V$)
- Events from CDMSlite mode detectors
- Any data collected 2 days after a neutron calibration. Internal activation makes the data unsuitable for a WIMP search
- Events from IT3Z1 when it had S1 grounded for a special study (Aug. 17-21, 2012 and Aug. 30-31, 2012)
- Events in the Low-Mass WIMP search region as described above
- Events during the CDMSlite R2 testing in December 2013
- Any reverse bias data during R134 that was not taken on IT3Z1, IT3Z3 or IT5Z3
- Any data taken 3 days after extended neutron calibrations in February and May of 2014

Each of the blinding requirements are composed as individual cuts that are joined via logical AND to obtain a detector specific blinding decision. The ultimate event based blinding decision is then computed by taking a logical OR over every detector's blinding decision so that every event that is flagged by any detector as worthy of being blinded is tagged. Once the blinding cut is in place, a few cross-checks are performed on neutron calibration data to ensure the blinding cut is not behaving in an unusual fashion (e.g., it does not remove a narrow energy window, or cause any odd gaps in partition planes). Typically it is useful to plot the various energy planes that go into making the blinding decisions and show the effects of the total blinding cut on each detector, the full detector specific blinding decision for a particular detector, and the individual blinding cut

decisions in their appropriate planes. After the checks are complete, the unblinded data files are used to generate new files that do not contain the blinded events and distributed for widespread use to the analyzers.¹⁶

Blinding Error. Since data continued to be collected while the low-mass search was underway, the blinding process was to happen twice (once for R133 data and then for R134). However during the initial checks for the R134 blinding scheme an alarming issue was brought to attention due to some discrepancy between the MATLAB based multiples cut and the version implemented in the blinding framework. Since MATLAB is where most of the primary analysis is performed and where cuts originate, these are the accepted versions. However, the blinding logic was written in Python and interfaced with the data files via a macro. It turns out that an error was introduced in the Python based version of the multiples cut. Recall that for each vetoing detector, an event must be above the energy threshold and not be a LF noise or glitch event. In the Python version, it was improperly coded such that each vetoing detector merely had to be above energy threshold, and that the primary detector not be LF noise or a glitch. This mis-match reduced the singles efficiency on randoms by a fairly substantial amount to $\sim 94\%$ and made the series to series variation a lot broader. In essence it implied that the about 6% of our singles efficiency was accidentally left unblinded for R133.

This lead to the decision to ultimately re-blind the R133 data files at the same time as the R134 data files (to ensure similar blinding schemes for both datasets) and to create a replica of the blinding framework in CAP/MatCAP, which was then used to cross-check the Python version. While constructing the CAP version of the blinding cut I discovered our radial fiducial volume was actually a bit tighter than it was intended to be due to another coding issue. The original intent was to have an event be in the S1 radial fiducial

¹⁶ The unblinded files are prefixed as “bg_restricted” and have their file permissions set so that only a special user account can read them. The blinded files are prefixed with “bg_permitted”.

volume OR the S2 radial fiducial volume, instead the logical OR was coded as a logical AND making the fiducial volume slightly tighter.¹⁷ This was ultimately decided to be left as is and incorporated into the CAP cut as well. At the end of much effort, the two frameworks agreed and the data was (re)blinded.

This did not turn out to be as bad as it could have been for a few reasons. The first is that while this was discovered after the LT analysis was finished, the blinding cut used for the low-mass dataset was written in CAP entirely and so used the correct cuts. Thus the LT analysis was properly blinded. The second is that the R133 HT data was not heavily used during the timeframe that the LT analysis was underway and so not much of a substantial amount of work was performed with it that could have potentially introduced bias against signal. The third is that the 6% accidental unblinding of the signal efficiency is itself very small anyways.

5.6 High Ionization Noise Cut

Periods of high ionization noise are not good candidates for a WIMP search since it can cause unexpected yield values and lead to mis-identification of events at low energy. For example a surface event during a particularly noisy period might wind up with a symmetric ionization symmetry value, and slightly higher total ionization energy estimator and appear like a good NR event. The series-by-series noise information generated can be used as a way to discriminate against periods of high noise on the various ionization channels. For this study events with a phonon energy less than 4σ above the mean noise in `psum0F` were considered, and the glitch and ionization pre-pulse cuts were applied. This selection guarantees that we have low phonon energy so that any even with a high ionization energy is likely to be noise that does not arise from easily identified sources.

¹⁷ This happened because in the Python code the cut actually checks that an event is NOT an outer channel event on S1 OR NOT an outer event on S2, but when negating one has to recall “A or B” becomes “NOT A and NOT B”.

But “high ionization energy” is a relative term that needs to be defined a bit more rigorously. In CDMS-II the prescription was to take a simple moving average filter of the ionization energy as a function of `EventTime` and remove time periods that had a moving average some number of σ above the noise distribution mean (adjusted for the number of points in the moving average). An additional set of logic joins consecutive periods of high noise if they are in close proximity as measured by event index, and then time, and finally a buffer before and after each window is generated as extra precaution since the moving average does tend to smear the actual distribution slightly.

As was quickly discovered this method was not directly applicable to iZIP data however, due to the new OF algorithms used. The old CDMS-II data relied upon an OF that simply maximized the amplitude and as a result the noise distributions were Gaussian-like, lending themselves to simple $\mu + n\sigma$ style thresholds when a moving average is taken. However, recall the noise distributions in this work are bi-modal and fairly symmetric, so a moving average is going to tend to collapse the central value of the moving average towards zero. Since the time dependent means of the noise distributions are centered on the positive energy peak, this would cause very few events in the moving average to be above the threshold unless a very small (or potentially negative) multiple of the time dependent noise width σ was used.

Suppose the charge noise is distributed according to a normal distribution X with mean μ and variance σ^2 (as was the case for CDMS-II). Let us take a sample of noise from N such distributions and average them together. We can compute the expected value of the mean and variance of this average as follows:

$$E[\overline{X}] = E\left[\frac{X_1}{N} + \frac{X_2}{N} + \dots + \frac{X_N}{N}\right] \quad (5-2)$$

This can be further simplified:

$$\frac{1}{N} (E[X_1] + E[X_2] + \dots + E[X_N]) = \frac{1}{N} (N\mu) = \mu \quad (5-3)$$

Following this we can proceed to compute the variance $V[\bar{X}]$. We have:

$$V[\bar{X}] = V\left[\frac{X_1}{N} + \frac{X_2}{N} + \dots + \frac{X_N}{N}\right] = \frac{1}{N^2} (V[X_1] + V[X_2] + \dots + V[X_N]) \quad (5-4)$$

Here we have neglected covariance terms by assuming each X_i is independent. This assumption states that the noise from one distribution cannot influence the past or future occurrence of noise. This should be a safe assumption. Continuing to simplify equation 5-4 we have

$$\frac{1}{N^2} (NV[X]) = \frac{\sigma^2}{N} \quad (5-5)$$

Thus thresholds for set by a N point moving average over the charge noise for a single peaked gaussian distribution will be of the form

$$E_{th} = \mu + m \frac{\sigma}{\sqrt{N}} \quad (5-6)$$

Where m is some number of standard deviations above the mean we wish to have the threshold set. This is appropriate for the old CDMS II ionization noise and for any 1D OF phonon channels.

An alternative proposal (suggested by Tarek Saab) is to take the moving average of the square of the charge noise. Doing this will tend to exaggerate large energy values and diminish small energy values, and avoid negative energy all together. However the type of threshold will change dramatically. In essence we are now sampling from N distributions of squared normally distributed quantities, X_i^2

Doing this will require us to calculate $E[X^2]$ which is just the second moment for a normal distribution.

$$E[X^2] = \sigma^2 + \mu^2 \quad (5-7)$$

Armed with this we can now proceed to compute the expectation value of the mean of the squares

$$E[\bar{X}^2] = E\left[\frac{X_1^2}{N} + \frac{X_2^2}{N} + \dots + \frac{X_N^2}{N}\right] \quad (5-8)$$

This can be broken up similarly as in 5-3 into:

$$\frac{1}{N} (E[X_1^2] + E[X_2^2] + \dots + E[X_N^2]) = \frac{1}{N} (N(\sigma^2 + \mu^2)) = \sigma^2 + \mu^2 \quad (5-9)$$

Next to compute $V[\overline{X^2}]$ we require the expression for variance in terms of expectation values:

$$V[X] = E[X^2] - E[X]^2 \quad (5-10)$$

Substituting into 5-10 we obtain:

$$V[\overline{X^2}] = V\left[\frac{X_1^2}{N} + \dots + \frac{X_N^2}{N}\right] = \frac{1}{N^2} (V[X_1^2] + \dots + V[X_N^2]) \quad (5-11)$$

This necessitates carrying out the computation of $V[X^2] = E[X^4] - E[X^2]^2$ which requires the fourth moment for a normal distribution, $E[X^4] = \mu^4 + 6\mu^2\sigma^2 + 3\sigma^4$. From this it would follow that

$$V[X^2] = \mu^4 + 6\mu^2\sigma^2 + 3\sigma^4 - (\sigma^2 + \mu^2)^2 = 2\sigma^4 + 4\mu^2\sigma^2 \quad (5-12)$$

Plugging this into 5-11 we arrive at:

$$V[\overline{X^2}] = \frac{1}{N^2} (NV[X^2]) = \frac{2\sigma^4 + 4\mu^2\sigma^2}{N} \quad (5-13)$$

This suggests that if we wish to set the charge high noise cut using a N point moving average on the square of a normal distribution we should set a threshold of the form:

$$E_{th}^2 = (\sigma^2 + \mu^2) + m \left(\sqrt{\frac{2\sigma^4 + 4\mu^2\sigma^2}{N}} \right) \quad (5-14)$$

It is critical to notice the method discussed above should be valid only if the underlying distribution we are squaring, X is normally distributed!

The c334 noise blobs are, for better or worse, decidedly not unimodal distributions. They have a distinct twin peak feature in all three processing modes present. The moving average threshold should attempt to account for this bimodal shape. In the absolute most simple case we can approximate the distribution as a mixture model of two gaussian

distributions $Z = pX_1 + (1 - p)X_2$ where p is the mixture coefficient. Let us specify that $E[X_1] = \mu_1$ and $E[X_2] = \mu_2$ with $V[X_1] = \sigma_1^2$ and $V[X_2] = \sigma_2^2$. In practice the means of the noise blobs are fairly close to being opposite in value and the variances are not terribly different between each side. Further for more simplicity let's assume the mixture is evenly distributed, or $p = \frac{1}{2}$, and $\mu_1 = -\mu_2$ and $\sigma_1 = \sigma_2$. Following the example for a two-component mixture model in [105] we can compute the expectation value of this distribution:

$$E[Z] = p\mu_1 + (1 - p)\mu_2 \rightarrow 0 \quad (5-15)$$

And the variance can be computed as:

$$V[Z] = E[Z^2] - E[Z]^2 \quad (5-16)$$

The second moment for a bimodal distribution is a bit complicated but simplifies greatly under the assumptions made above:

$$E[Z^2] = p(\sigma_1^2 + \mu_1^2) + (1 - p)(\sigma_2^2 + \mu_2^2) \rightarrow \sigma^2 + \mu^2 \quad (5-17)$$

In reality this expression could be modified by a covariance term, $cov(X_1, X_2)$. In practice for the charge noise blob the covariance between the two gaussians that form the mixture model is a difficult quantity to determine. If the population of positive amplitude values is independent of the negative amplitude values the covariance can be safely ignored, and there is not much reason to expect a positive or negative noise should impact one or the other on a given channel given the relatively low amplitude and the assumption of causality made earlier.

This allows us to finish with the computation of $V[\bar{Z}]$:

$$V[\bar{Z}] = V\left[\frac{Z_1}{N} + \dots + \frac{Z_N}{N}\right] = \frac{1}{N^2}(NV[Z]) \rightarrow \frac{V[Z]}{N} = \frac{\sigma^2 + \mu^2}{N} \quad (5-18)$$

This suggests that if one sets a threshold based on a moving average of N points from a bimodal distribution we should follow:

$$E_{th} = m\sqrt{\frac{\sigma^2 + \mu^2}{N}} \quad (5-19)$$

Where m is the number of “sigmas” above the approximately zero mean of the moving average filter. The assumption of a 50-50 split may not be correct (however to first order one would think we should get positive and negative fluctuations fairly equally). There will also be a systematic if the bimodal index B is close to 2. For distributions to be considered as separate the bimodal index must be larger than 2.

$$B = \left(\sqrt{2} \frac{|\mu_1 - \mu_2|}{\sqrt{\sigma_1^2 + \sigma_2^2}} \right) \quad (5-20)$$

Luckily it appears that $B > 2$ for fully functional detectors, and will be larger for the 210 and 110 processing modes due to the more pronounced gap between the peaks.

The biggest flaw to this method is the assumption that the distribution is drawn from a mixture of Gaussians. Due to the nature of the OF noise distributions (drawn from extreme values of the actual gaussian noise distribution) there is a tendency for the OF noise blobs to be non-Gaussian in their mixture distributions. This is especially true for the 210 and 110 noise blob fits.¹⁸ This hints that this method will perform best for charge channels that are working perfectly and perform worse for 210 and worse still for 110 fits.

Using the threshold as determined from 5-19 one can now properly apply a moving average filter to the ionization noise as a function of time and select for periods of high noise. For the cut I set $m = 5$ and used a 20-point moving average. This seemed sufficiently large to avoid transient spikes in noise and sufficiently small to capture obvious periods of large noise fluctuations. If any consecutive periods of high noise occurred within 5 events of each other, they were merged together, and subsequently if any two consecutive

¹⁸ The 110 filters are pretty convincing Gumbel distribution curves.

noise periods occurred within 10 minutes of each other these were also joined together (to protect against periods where the noise may still be rather high but not quite high enough to exceed threshold). Finally a buffer extending to 3 minutes in real time on either side of these noise periods is applied to obtain the final selection criteria for a given channel.




Figure 5-8. An example of the selection criteria for high ionization noise on IT2Z3 for R133. The grey points represent the actual events in the data, and green points result from the application of a 20-point moving average. The horizontal red lines represent the energy threshold to be considered high ionization noise. They are time dependent and based off of the time varying ionization noise blob means and sigmas. Blue points exist for time periods where the moving average exceeds the threshold. This particular detector had a very large period of high noise from September through December 2012.

The final cut, `cQhighnoise_v53`, will remove all events within the noise windows for all 4 of the individual charge channels, but only for periods of time where a given channel is fully functional. For example, if charge channel 1 is non-functional during one particular period of time that it flags as high noise, this period will not be used, though another functional detector may coincidentally remove the same period. Since this cut removes time periods it is considered a livetime cut and it is important that it not remove an excessive portion of the livetime. By computing the amount of raw livetime that passes this cut it was found that for most detectors the amount removed is about

$\sim 1 - 2\%$. [106] Some detectors, like IT2Z3 had closer to 5% removed due to excessive high noise triggering. This cut was computed for all 3 of the main data types.

5.7 KS and QS Test Cuts

One important class of cuts that are conceptually simple, but hard to implement, are cuts based on performing a two-sample Kolmogorov-Smirnov test (KS test) on the data. This is a type of statistical test that can be used to compare two distributions to test whether they are sampled from the same underlying distribution and at what significance. [107]. The two-sample KS test is a generalization of the one-sample test that is used to test the similarity of two different distributions. The KS test compares the CDFs of two distributions, F and G, and computes a KS-statistic by finding the maximal (vertical) separation between the CDFs bin by bin:

$$D_{n,m} = \max(|F_n - G_m|) \quad (5-21)$$

From here the significance can be computed via the following infinite series:

$$Q = 2 \sum_{n=1}^{\infty} (-1)^{n-1} e^{-2n^2 \lambda^2}, \quad \lambda = \left(\sqrt{N_{eff}} + 0.12 + \frac{0.11}{\sqrt{N_{eff}}} \right) D_{n,m} \quad (5-22)$$

If F contains N points and G contains M points, $N_{eff} = NM/(N + M)$. Here we see that the sum should converge quickly, unless $\lambda = 0$, in which case $Q \equiv 1$. The value of Q represents the approximate 2-sided P-value for the KS test and is what is used to determine how well two distributions match. The closer to 1 Q is, the more likely both distributions are manifestations of the same underlying distribution (we can see that as D gets smaller so should λ).

If two distributions are very dissimilar the significance will be very small, so setting a cut on this value is a natural choice. In practice though this is super complicated. The choice of what distributions to test, what data to select for these distributions, how to group the data to form distributions, and how to set the threshold are all points of difficulty. Naively we would like to use as much good quality data, and test each series

individually against every other series and then reject the lowest fraction of the data based on its typical KS score compared to all others. But this is challenging because not every series will have enough good quality data to be tested, and with ~ 4800 data series, a full comparison matrix would involve performing 4800^2 tests, multiplied by the number of RQ/RRQs being tested and by the number of each iZIP. Choosing just 10 RQ/RRQs to test in this fashion would lead to over 3 billion tests being performed.¹⁹ Further sabotaging this approach is the very fact that the behavior of the iZIPs changes over time due to noise and other environmental effects (channel shorts, different runs, etc).

This can be mitigated somewhat by instead choosing a small subset of characteristic samples to use as good templates, and then picking a smaller subset of this subset of the highest scoring templates. These better templates could then be used to test every series, reducing the output significance matrix from N^2 to something more manageable. From the CDMS-II analyses the general prescription was to select this first pass template by forming a histogram of the livetime of each series, and then selecting some fraction of the total number of series around the typical livetime value to use as “cherry-picked” series. This subset was then KS-tested amongst all the series and some fraction of the subsets that had the highest average KS significance across all series were then selected²⁰ to be the representative templates to be compared against every series in a second round of KS testing. From this second round, the mean KS-significance was again computed across all series, and a first pass cut removed all series 3σ below the average aggregate KS significance value, though this was often then tuned by hand.

¹⁹ Each test appears to take about 1×10^{-4} s, and performing all 3 billion would take over a week. Parallelization can help this somewhat but not enough.

²⁰ Very cherry-picked in old jargon.

5.7.1 LT Analysis Tests

With the LT analysis Ray Bunker converted the old framework to be compatible with iZIP style detectors and discovered one would have to sometimes merge adjacent series together to get enough statistics to test (typically at least 10 events per sample). This was not seen in the CDMS-II analyses because their data series lasted longer than the 3-hour duration of a typical iZIP series. It was also discovered that due to the time variation, 3 separate time periods had to be selected to test amongst themselves. Another new test was to make a separate set of KS tests for ionization quantities only, so-called QS tests, and for the behavior of the randoms, the so-called RS test. Each of these was meant to be sensitive to a different aspect of the data and verify the data used was self-consistent. The net result of this cut was the removal of approximately 20 days worth of raw livetime from the LT dataset.




Figure 5-9. The amount of livetime removed by successive application of the RS, KS, and QS tests. Only about half of the livetime removed by the RS test was not already removed by other cuts. The KS test removes a maximum of 5 days more livetime on IT1Z1 and substantially less on other detectors, with the QS test not removing much extra livetime at all. Courtesy of Ray Bunker.

5.7.2 HT Analysis Tests

One of the major drawbacks to this cut is that the inherited framework was very clunky, and required user intervention to make separate code files per RQ/RRQ being tested, as well as to manually code the individual time blocks everywhere. In addition the code was very slow to run, though part of this was due to the bad ROOT file problem described earlier, most of it was simply un-modernized MATLAB code from the early CDMS-II days. The net time to run the KS scripts on the LT dataset was measured in a few days. For the HT analysis, we wanted to run KS/QS tests as well, but this posed a serious problem. The dataset was enormous compared to the LT search, and certainly had much more than 3 time periods within it. This required a significant amount of time completely overhauling the KS testing framework to streamline the memory access times and minimize the need for user modification to scripts. A comprehensive description of the underlying KS testing framework is beyond the scope of this work, but the net speedup was such that from start to finish the entire HT WIMP-search dataset could be tested in about a day. The general flow of the KS testing is however worth describing since it has a non-trivial impact on the HT livetime. For the HT analysis the following quantities were used for the KS tests: `ytNF`, `qzpart0F`, `pzpart0F`, `prpart0F`, `prxypart0F`, `prdel1WK`, `prdel2WK`, `ptRiseTime`, and `ptFallTime`.²¹ The delay quantities radial quantities give some sensitivity to changes in the noise environment of the detectors, however the radial partition quantities would prove not very useful. For the QS tests `qzpart0F`, `qrpart10F`, and `qrpart20F` were used. The RS test was retired for the HT analysis, instead using the `cGoodRandomChi2_v53_HT` cut developed by Hang described earlier.

²¹ The quantity `prxypart0F` is a radial partition constructed from the 6 inner phonon channels. One constructs x and y partitions and then takes the square root of the sum of these. The `prdel[k]WK` quantities are the time delay between the 20% risetime of maximum energy phonon channel and the 20% risetime of the radial phonon channel on side [k] as determined by walking along the pulse. The `ptRiseTime` and `ptFallTime` quantities encode the total time it takes the total phonon trace to rise and fall respectively.

5.7.3 Time Period Determination

With the RRQs selected for testing purposes, the next major step in the KS testing framework is to determine a list of time periods. In an ideal case there is just a single time period for every detector and quantity. In the worst case, every RRQ for every detector could have its own time periods. Alas, the SuperCDMS Soudan dataset falls into this latter category and so this informed the overhaul of the framework to allow this most general time period entry. Actually selecting the time periods in a reasonable fashion required some thought. It turned out to not be sufficient to simply examine detector status flags since subtle variations in quantities come and go. Since a time period is defined as a point during which the distribution changes and can result in a lower KS significance (or higher KS statistic), I decided to use that as a metric to determine the time periods. A very rough selection cut was placed on each of the RQs to select only events in a typical value range. From here the time periods are determined based on the following logic:

- Starting with the earliest time, form two adjacent windows T and T_2 containing N points and perform a KS-test.
- If the KS-statistic is below a threshold K , merge the two windows, $T \leftarrow [T, T_2]$, and select the next N events as the new T_2 . Proceed in this fashion until either a test statistic is above K , or one reaches the end of the data.
- If the KS-statistic is larger than K we need to verify if it is a temporary spike or a real change. Select a window, T_3 of N points beyond T_2 , and compare T with T_3 to get K_{13} and T_2 with T_3 to get K_{23} . If $K_{13} < K$ consider T_2 a temporary spike, and merge the 3 time windows together into a new T and form a new T_2 . If however T_2 and T_3 are more similar, then a time period boundary has been identified. Set the time boundary to be at the first event time in T_2 , then discard the old accumulated window T and continue the process by setting $T_2 \rightarrow T$ and continuing this walking process until the end.

For this work I constrained the KS statistic threshold to be 0.23, and required N to be the maximum of 2000 events or 10% of the total number of events, and imposed an additional overriding constraint that at least 250 series be contained within a time window, and if the start of a new Run occurred in a time window that was flagged as

being a new time period, to set the time period boundary at the start of the Run. This proved to be a very robust way to detect time periods across most of the RRQs for detectors, but was not perfect. On occasion the method would select too many time periods, or just miss time periods. Luckily some of the RRQs should have the same time periods for physically motivated reasons (the radial quantities should be somewhat related, as should the yield and any large changes in either ionization or phonon behavior). In such cases selected time periods were combined to be common between them. Ultimately some manual adjustment was still required to catch a few known changes in channel performance, but significantly less so than if the whole process was performed manually. This automated process should be useful for future KS testing.




Figure 5-10. The result of using a KS test to determine time periods for the more sophisticated KS testing process on IT2Z2 for the phonon z partition. The algorithm correctly identified a region of time where the distribution's behavior changed between R133 and R134. The `EventTime` RQ on the x-axis is just unix-time.

5.7.4 Template Generation and Merging

Once the time periods are constructed, the series templates can be generated. This consists of the application of the good data quality selection criteria as well as a series of loose fiducial volume cuts on the ionization quantities. To avoid storing everything

in memory the information is stored in various HDF5 (h5) files on disk, organized into directories based on detector number, and then each file is named such that it indicates the RRQ and which test (KS or QS) it belongs to. The contents are simply the passing events organized into by-series groups and stored as a variable that indicates which series the data belongs to. Additionally information is saved regarding the livetime for a given series, as well as the total livetime (the sum of the livetime). Once the template files are generated, the merging process occurs and completes two goals: it breaks data into the necessary time periods per detector per RQ/RRQ, and it ensures a merged template contains enough data to be meaningfully tested. For the latter constraint, user definable parameters control the number of events to be stored in a merged template ($N = 10$ for this work) and the total number of series allowable within a group ($S = 8$). Additionally the grand total livetime (GTL) of a series (the sum of the livetime across all detectors) is used as a metric for rejecting series with zero livetime whatsoever. The logic used to merge series together into separate groups follows the following test cases:

- Case 1: If the number of events in the series is greater than the required minimum and the GTL is greater than zero, this series can close its own group. Proceed to the next series and start here again.
- Case 2: If the minimum event number requirement has failed but the GTL is greater than zero, add this series to a group and proceed to Case 3 for subsequent series until the group can be closed.
- Case 3: Here we are growing the number of series in a group until the number of events reaches the desired minimum. Each new series can be appended to the group as long as the GTL is greater than zero (otherwise it goes to Case 4). We do not want groups to contain a large number of series that are separated in time so a user definable maximum number of series is imposed. If we reach the desired minimum number of events before hitting this cap, the group is considered complete and we proceed to the next group. If however we reach the maximum number of series without satisfying the group event number constraints, the very first series added to this group is flagged for rejection. Before rejecting the series, a check is performed on the previous group to see if the series can be appended to it (a back-merge) without violating the maximum series limit. If it can, it is added and the previous group information is updated. If it cannot, then the series is rejected and put into the fail list. Now the current group contains $S - 1$ members so this process repeats until we

can close out the group. Once we have successfully met the group requirements, the group is closed and we start over at Case 1 for the next series.

- Case 4: Any series that contains 0 GTL, or gets rejected from back-merging finds itself in the fail list. These series are not KS tested since they either contain too few events, or no livetime anyways.

Once the merging lists and associated metadata are created, the actual events are written to merged template files, stored in a fashion similar to the originals except now broken up by each specific time period, and now instead of referring to discrete elements of a template as “series” we refer to them as “groups” since there may be more than one series present. The group starts at 1 for every RQ/RRQ and detector and time period, so when comparing all time periods of the same RQ/RRQ on the same detector one must recall there is no unique “group 1” (unless you create a running group index).

5.7.5 Subset Selection

After merging has occurred the actual subset selection process can begin. The first round is the naive subset selection which uses the typical livetime contained in the detector and RRQ specific time period as a metric for selecting some amount of templates. Since each time period can have a variable number of groups, unlike the old CDMS-II style KS tests, I allowed the naive subset number to be somewhat variable. For this work I required the number of groups chosen for the naive subset to be 75% of the total number of groups if that value was between 200 - 400 groups, otherwise no more than 400 groups, and no fewer than the minimum of 200 or all groups (in case the 75% requirement resulted in less than 200 groups present). For most cases this resulted in 400 groups being selected as naive subsets. These were then compared to every group in the appropriate time period to construct a KS significance matrix, from which the median score of a naive subset across all groups was computed. This differs from all the previous analyses which used a mean at this step. The KS significance distribution is very non-Gaussian with a clear heavy tail. Using the mean is thus sensitive to outliers causing some naive subsets that are more representative of most of the groups to have a lower score than they should.

Once the median was taken, a better subset was chosen from among the naive subset by picking the highest scoring members. For this work the range required a “better subset” to have 36% of the total number of naive subset groups if it was between 80 and 120 groups, otherwise no more than 120 groups and no fewer than the minimum of 80 or all (similar selection logic to the naive subset number).




Figure 5-11. This is a diagnostic plot for IT3Z3’s QS test (a KS test on ionization quantities). Shown here are the median KS significance scores for each of the naive subsets. All time periods are plotted here as a running group index. The existence of 3 groups is easy to discern by the step like jumps in point clustering. Horizontal bars indicate the KS significance threshold determined for selecting better subset members. All groups above this threshold are used as the better subset and all groups below are not. This style of plot is useful for verifying no systematic changes occur within a group and that better subset members tend to be clustered above the exclusion threshold. The KS test version of this plot is far more cluttered.

Once the better subset groups are selected a set of self consistency checks is performed to verify they are of good health. This is accomplished by performing a KS test amongst themselves and by aggregating the entire better subset into a single master template, and then running a KS test between each better subset member and the master template.




Figure 5-12. Example diagnostic plot for better subset selection for the KS test on IT3Z3. In this plot a KS test was run between each member of the better subset and the median KS significance computed. A healthy better subset should have no large dips towards low KS significance values in any of the groups. If this does occur it suggests the better subset groups are not fully consistent. There is some evidence here that `prpartOF` is suffering some issues due to its fluctuations. A similar plot is generated from comparison against the master template and is interpreted in a similar fashion. Additionally the behavior should be the same between both types of plots, any discrepancy could be a sign of bad templates.

5.7.6 Kilt Plots and Thresholds

Once the better subsets are all created and verified, the main KS testing process takes place in which each of the better subsets are compared against all the groups for a particular RRQ and detector. From this big matrix the median KS score of a particular group across all the better subsets is taken to be that group's representative KS significance. It is on this quantity that a threshold is created and cut upon to remove poorly performing groups. This step in the analysis provides cross-checks on the output of the KS test via the use of so-called "kilt plots". These are basically colormap plots of the KS significance matrix, and provide a quick way to diagnose the performance of the better

subset templates and time period selection, and how useful the RRQ being tested happens to be based on how much variation occurs in the KS significance values.




Figure 5-13. A colormap plot of the KS significance matrix is useful for diagnosing behavior within a time block (period), often referred to internally as a kilt plot. Based on the distribution of the lower KS scores as a function of subset groups and all groups one can infer the utility of the RRQ and even the time period definition. If the time period is incorrect, the better subset will be likely to contain template members from each time period. This will result in a large checkerboard like pattern of reasonable KS significance values (blue) and failing values (light yellow). Another sign of this is if one mis-matched time period has no representatives in the better subset population, a large horizontal band of failing values will be present. In each plot the better subset series represents the first series present in a particular group, however for these plots the axes values are not as important as the behavior of the plot itself. A) Shown here is an example kilt plot showing some questionable behavior near the start of the run for this particular detector. Since the light yellow region is not a solid block it does not indicate a time period mismatch, merely a sign of questionable detector behavior near the start of R133. The rest of the time represents fairly stable behavior for this particular RRQ. B) An example of a quantity that is not useful for KS testing purposes. There is simply too much variation from group to group and no better subset template performs well always. This behavior for `prxyartOF` and `prpartOF` was present for all detectors and time periods unfortunately, but a loose threshold can still be set to cut the worst offenders.

After a few iterations on testing to adjust slightly mis-aligned time periods a tentative threshold to cut on the KS significance is generated by rejecting the lowest 1% of

scores above a minimum acceptable KS significance of 1×10^{-4} (both user adjustable parameters). The behavior of these thresholds is then judged based on an examination of the several factors: where it cuts on a the KS significance vs number of tested events plane, the amount of livetime removed, and what the distributions that just fail and just pass the threshold cut look like compared to the master distribution. All this is a process that requires a pair of eyes unfortunately and can get tedious after a while.²² Once the thresholds are set, these are incorporated into a KS and QS cut that removes series that are members of failing groups, as well as series that failed the merge step.

5.7.7 Special Considerations for the Barium Data

Unlike the LT analysis the HT version also ran the KS testing framework on the ^{133}Ba calibration data. This required a slight modification because the KS significance values were absurdly low, and somewhat numerically untrustworthy, initially (on the order of 10^{-100}). This can be explained by recalling that Q in Equation 5–22 depends on an exponential term such that the first term in the sum goes as $Q_1 \sim e^{-2N_{eff}D^2}$. If N_{eff} is somewhat large and D is not extremely small, the very first term in series is going to start off very small and subsequent contributions will themselves be very small. Since the ^{133}Ba calibration datasets are much larger compared to WS data, the typical value for N_{eff} was $\mathcal{O}(10^3)$ resulting in very small significance values since there was an unfortunate amount of series to series variation.

To get around this issue, a series of random samples from each distribution being compared were taken and these random subsets were tested. The number of points in each random sample is itself randomly chosen from between 20 and 100 points and differs for each of the two sub-samples being compared. This process is repeated multiple times to

²² At some point in the analysis towards this cut, the framework output approximately 2GB worth of histograms to examine...approximately 23,000! Thankfully the other planes are more useful than direct comparison of each failing distribution but a substantial fraction were examined.




Figure 5-14. When determining if the KS significance threshold is appropriate to set a cut on these two plots provide excellent diagnostic information as to whether the threshold needs to be raised or lowered. A) A scatter plot of the median KS significance as a function of the number of events in the group. Generally a distinct population of low scoring outliers shows up in these plots that one wishes to ensure are excluded from the central bulk distribution. The vertical red line indicates the minimum acceptable number of events a group should have to be tested which may be different than the minimum number allowed by merging. The blue horizontal line represents the KS significance threshold below which groups would be removed. If points have a blue + sign over them this indicates they are also removed by another RRQ's KS test. B) The livetime in days removed as a function of threshold. Most distributions have a somewhat linear relationship in this log-log plane however RRQs with unusual behavior, such as `prpartOF` may have distinct changes in slope.

build up a distribution of KS significance scores between the two. If the two distributions are close to being the same, the KS significance will have a distribution that is very slightly biased towards high significance values. As the distributions become less similar the mode of the significance score distribution will start to move towards ever lower values. As such the median of this distribution is chosen to represent the KS significance result between the two original distributions. The downside of course is that the time it takes to generate the test results will scale up with the number of random samples one takes. For this work 25 random samples were taken for each comparison which provided enough




Figure 5-15. This plot shows each of the ^{133}Ba calibration series yield CDFs. Since these are what the KS test compares when generating the test statistic, it is seen quite clearly that there is a significant spread to the yield which will result in values of D that are not small enough to compensate for the large effective number of points.

to get a decent fix on the median, however the time delay resulted in an extra day of processing.²³

5.7.8 Cut Results

At the end of the analysis the ultimate quantity that we are interested in is how much livetime this cut removes. Too aggressive of a KS threshold will remove too much livetime to be useful, so it is a somewhat fine balance that needs to be struck. As long as the distributions that are just above threshold does not look extremely different from the distribution just below the threshold the cut is probably set at a decent level. Otherwise

²³ Most likely 25 is just barely enough and 100 such tests would be more optimal.




Figure 5-16. Result of the random sampling method between two different Barium datasets for the yield RRQ. A) The CDF plane with maximal separation indicated by the black horizontal line. The KS significance is shown as 4.138×10^{-43} . B) The result of 200 KS tests on random sub-samplings of each distribution. The median KS significance value has shifted to something on the order of 10^{-3} which is a bit more trustworthy.

there may be a lot of variation around where the cut is set and some examination may be required. For the ^{133}Ba calibration data the livetime is a less meaningful quantity, so the effect on the number of good bulk ERs (defined as events passing the good events cut and in a 3σ ER band) was chosen. The results of the KS/QS tests on the WS data are shown in Figure 5-17. The starting livetime for each detector is slightly different due to the application of a WS data period cut that excludes some time when the two source detectors (7 and 9) were unblinded, and when IT5Z2 (14) was in CDMSlite mode. The application of the livetime cut prior to the KS and QS cuts removes an unfortunately high amount of livetime, averaging about 20 – 40% on the HT detectors.²⁴ The KS tests remove an additional 5 – 10% of livetime. The QS tests do not remove much additional

²⁴ Recall exposure is the product of the mass and livetime so removing this much livetime diminishes our exposure proportionally.

livetime, just an additional few percent. The net result application of these two cuts gives us the total livetime available for the HT analysis on each of the detectors. It is a reduction of about 30 – 50% of the initial raw livetime but it appears that there is no mis-tuned cut that is responsible for eliminating a large chunk of livetime, as seen in Figure 5-18.




Figure 5-17. The resulting livetime, in actual days and percentage, that survives after the application of the KS and QS cuts on the WS data are shown here. Each detector's raw livetime is determined by a data period cut that defines the actual time periods a detector can be used for a WIMP search. The KS and QS cuts remove approximately 20% of the livetime that is not already removed by previous livetime cuts.

The results of the barium KS/QS analysis are shown in Figure 5-19. A variable number of good bulk ERs (GBERs)²⁵ are present on each detector owing to both data quality and source exposure issues. The T4 detectors typically had a larger exposure to the ¹³³Ba sources and as such have more statistics to work with. The initial tuning of the cuts removed approximately 60% of all the GBERs mostly due to a few RRQs

²⁵ Defined by selecting events between 45 and 150 keV_r and within a 3 σ ER band.




Figure 5-18. The cutflow diagram for the component cuts in the HT analysis livetime umbrella cut, not including the KS style cuts. Approximately 2 months of data are removed by the data period cut initially, corresponding to the end of R134 when noise was extremely high due to poor cryocooler performance. However afterwards no single cut removes a large fraction of the total livetime, with each detector appearing slightly different than the one shown (IT2Z2). Note the y-axis starts at 360 livedays.

which had their thresholds lowered significantly to allow more passage.²⁶ As a result of the tuning approximately 20% of the GBERs are removed by both these cuts. IT5Z2 (14) has a significant amount of data removed by the QS tests unlike every other detector. This is due to to both the radial partition quantities interestingly, with the side 2 version removing more. This appears to be mainly the result of unusual detector behavior at the start of R135 that is not removed by the good events selection cut so it is good the QS tests catch this!

²⁶ These were the radial partition quantities primarily as well as the yield and to a lesser extend the phonon z partition.




Figure 5-19. The amount of good bulk ERs in a 3σ ER band are shown after application of the KS and QS cuts since livetime is not very meaningful in calibration data. The cut thresholds were tuned such that about 20% of the good ERs are removed on the HT detectors. The significant reduction on IT4Z1 (10) is due to the KS test on qzpart0F but is not important since this detector is not a HT detector due to its bad ionization channel.

CHAPTER 6 LOW MASS WIMP SEARCH

The low-mass WIMP search (or the low-threshold (LT) analysis) was the first iZIP analysis with multiple detectors.[59] The dataset was chosen to be a time period in which we had the best detector performance, as measured by trigger threshold stability, at the time.¹ This wound up being a period of time from October 2012 through the end of July 2013, constituting a total exposure of 577 kg-days. Additionally only the best performing detectors were utilized for this analysis, as seen in Figure 6-1. Initially the detector IT1Z3 (1103) was included in the subset due to its low phonon noise despite having broken ionization and phonon sensors, but it quickly proved difficult to understand this detector and so it was dropped. The overall energy range was chosen to be between 2 - 13.1 keV_t.²




Figure 6-1. A subset of the best performing detectors used in the LT analysis (shaded). These were found across several different towers and chosen based on the phonon resolutions as determined by the zero-energy noise fits.

¹ This is important to have enough statistics to characterize the trigger efficiency accurately.

² Because the ionization signal is quenched at such low energies, the ionization signal is dominated by noise and consequently so is the yield. Assuming a standard Lindhard model this range works out to be $\sim 1.6 - 10$ keV_{nr}.




Figure 6-2. The time period in the context of the full SuperCDMS WIMP search period. At the time the analysis occurred data did not exist beyond mid 2013.

6.1 Lowering Thresholds via CROW

Following the discussion on direct detection results in 1.4.4, recall that for low-mass WIMPs the limit one can set on the cross-section quickly deteriorates as roughly:

$$\sigma_{SI} \sim e^{E_{th}m_N/2m_\chi^2v_0^2} \quad (6-1)$$

This illustrates that if one can push down the threshold of the experiment (E_{th}) the runaway asymptotic behavior in the sensitivity limit can be pushed to lower WIMP mass. This can be accomplished by a reduction of environmental noise, lowering detector noise, and ensuring the hardware triggering is optimized for the signal one is measuring. If the intrinsic energy resolution of a detector is too high, then one will quickly lose ability to differentiate signal from noise, let alone identify signal from background. Similarly, if the hardware that issues triggers misses most of your signal, then at low energies it will also trigger mostly on noise events which will reduce the operational livetime, requiring one to increase the trigger threshold and effectively raise the energy threshold of the experiment.

The RTF boards from SuperCDMS were reused from CDMS II and as part of their design include 2^{nd} order low-pass and high-pass filters combined to form a bandpass filter. When the incoming phonon signals from all 4 channels are fed into the RTF board, it will sum the signals, and then send this trace through the bandpass filter. If the filtered

signal amplitude is above the user-defined threshold, the incoming trace is deemed a valid triggering event. The bandpass range was designed and optimized for the CDMS II phonon signals and set to span the range of 900 Hz - 18kHz. The iZIP phonon traces are significantly different than the old CDMS II style traces however. Because of the reduced Al fin coverage the phonon traces are longer in duration, and significantly more information is carried in the tails of the phonon pulses (recall they all become equal in power around $750\mu s$). This strongly suggested that the original bandpass was set too high and would exclude some of the phonon signal power.

6.1.1 Software Modeling of the Bandpass

Naively one expects the signal amplitude to be maximized by as broad a bandpass filter as possible, with the lowpass cutoff very low, and the highpass cutoff arbitrarily high. However one cannot arbitrarily set the bandpass filter this way, as to do so will invite contamination from noise, which has its own power spectral density (PSD). The solution is to perform an analysis that attempts to maximize the signal-to-noise ratio as a function of bandpass window. An initial approach to explore the behavior of each filter independently is explained further in Figure 6-3.

The best approach to optimize this would be to simultaneously vary the bandpass windows and then locate the maximum signal to noise in this 2D space. In order to do this with enough frequency resolution at the lowest frequencies, special long-duration random noise traces were taken. As mentioned earlier in Section 2.3 these special modes doubled the number of bins per trace from 4096 to 8192, by decreasing the sample rate from 625kHz to 312.5kHz (effectively increasing the time duration of a trace by a factor of 4 to 26.2ms). This long mode run allowed the lowest frequency to be measured by the noise PSD to decrease from 152 Hz to 38 Hz with a resolution of 38 Hz per step. Several noise PSDs were recorded at different times and they were averaged together to produce the characteristic trace seen in Figure 6-3 part C.




Figure 6-3. Pictured here are the responses of a phonon template to variations in one parameter of the old bandpass filter (referred to as “current filter” at the time). A) The signal to noise response to the location of the low-pass filter cutoff. The best cutoff occurs not at the lowest frequency possible due to the presence of noise in the signal. The high-pass filter cutoff was kept fixed at 18kHz. B) The signal to noise response to the location of the high-pass filter cutoff. The phonon signal is expected to rise to a certain amplitude quickly as the cutoff increases and then plateau as less and less power is available at very high frequencies. The maximum is thus controlled by the high frequency part of the noise. The low-pass cutoff is kept fixed at the old value of 900Hz. C) The average noise PSD for this detector with the old bandpass filter values of 900Hz - 18kHz applied.

Having the data in hand, the final step to do is to create a model of the bandpass circuit in software. The MATLAB software license at Stanford includes a filtering toolbox which allowed this circuit to be modeled with relative ease. Because the low-pass and high-pass filters each have 2 reactive elements (capacitors) they are 2^{nd} order filters, so they were simulated together as a 2^{nd} order butterworth bandpass filter using MATLAB’s `butter` function. The outputs of this function are the transfer function coefficients, which encode information on the frequency response. By feeding these coefficients into MATLAB’s `freqz` function along with the frequencies at which to sample and the cut-off frequency (determined from the data sampling rate), the complex-frequency response is returned. This can then be applied to the signal and noise PSDs in the frequency domain by multiplication to obtain the filtered signal and noise.

The results from this optimization suggested two possible new bandpass ranges: 300Hz - 3kHz, or 600Hz - 3kHz. The 600Hz low-pass filter value was slightly more conservative against low-frequency noise but was not optimal for several detectors.




Figure 6-4. The results of simultaneous variation of the bandpass filter cutoff values on the signal to noise ratio, normalized such that the maximum is 1. The bandpass cutoffs are denoted by LF for the low-pass and HF for the high-pass filters. The current s/n refers to the relative (compared to the maximum) signal to noise for the old bandpass filter values and is located at the asterisk. A) For a detector with a clean noise spectrum (i.e., no prominent features) a clean looking blob is defined. This figure shows that the best signal to noise is contained in a window at much lower values than the original filter, which had a signal-to-noise ratio only 42% of optimal. B) In cases where spectral features are present at lower frequencies, distorted features are present in the parameter space. Notice as well the maximum signal to noise ratio shifts to higher values of LF. Despite the presence of such spectral features, the old signal to noise is still poor compared to the new location.

Conversely, the 300Hz value would be more prone to low-frequency noise issues but displayed two properties that weighed in its favor: Many detectors had higher signal-to-noise ratios than 600Hz, and the filtered signal amplitudes were $\sim 33\%$ higher. Thus it was decided to implement a modification to RTF boards that changed the bandpass window from 900Hz - 18kHz to one of 300Hz - 3kHz.

6.1.2 Hardware Implementation

Late in December 2012 while on shift with Ray Bunker, modification of the RTF boards began. This required powering off the electronics readout system and removal of the RTF boards to perform a removal and soldering of new, very small capacitors and resistors to the boards. This work began on the last day of my shift onsite so I was not

While testing the modifications of the first detector to undergo modifications (IT5Z2), an acronym for the modification work was hit upon: CDMS RTF bandpass Optimization Work, or CROW for short. This led to an attempt to create a label for the RTF boards so future shifters would be able to quickly identify what boards have been modified. A subsequent image was eventually found online that took place of the hand-drawn versions.

161




Figure 6-6. Labels to indicate which RTF boards underwent CROW modifications. The original green sticker was my attempt at drawing a crow. Courtesy of Ray Bunker.

40 – 50%), others were only marginally better. IT4Z2 had to be specially modified to a 900Hz - 3kHz bandpass filter due to unexpectedly large low-frequency noise dominating the trigger rate.³ Overall the CROW modifications appear to have successfully allowed for a reduction of trigger thresholds on some of the detectors that were going to be used for the LT analysis. Unfortunately as 2013 continued the LF noise would become worse over time on most detectors resulting in the need for somewhat higher trigger thresholds to be set.

³ This was one of the detectors the study suggested would be better off with a 600Hz - 3kHz bandpass

6.2 LT Data Quality Selection

Armed with the blinded dataset, a set of quality cuts were used to ensure the data being examined were of sufficient quality that we were sure they useable for a WIMP-search analysis. Many of these data quality cuts are described in detail in Chapter 5, but not all are used for the LT analysis. I will very briefly summarize the types of cuts used below, and refer the reader back to Chapter 5 for more in-depth descriptions of the types of cuts.

- Pulse shape: Each of the pulse shape cuts described earlier were used in this analysis in the LT form is applicable. This included the χ^2 , pre-pulse standard deviation, glitch template, LF noise template, triggered events, good phonon start time, and saturation cuts.
- Time Periods: Most of the time period cuts described earlier were used as well, in their LT form if applicable. Important cuts include the bad OF resolution, base temperature, trigger burst, good bias time, phonon settings, DC offset, stable trigger threshold, and bad series cuts.
- KS/QS/RS Cust: A event must also belong to a series that passed the LT version of the KS, QS, and RS test cuts. As a reminder these cuts perform a 2-sample Kolmogorov-Smirnov test on certain RRQ distributions compared to a representative set of distributions and reject series which deviate dramatically.
- NuMI: An event must not occur within $200\mu\text{s}$ of a neutrino beam dump by the NuMI experiment in case of any unlikely neutrino-sourced scatters.
- Post-Cf: An event must also not belong to a series within 20 days of a neutron calibration in order to avoid the period of time in which a detector is activated.
- Multiples: An event must not be a multiply-coincident scatter with any other detector (i.e., it must be a single-detector scattering event).

6.3 Pre-Fiducialization and Signal Region

In addition to these types of quality cuts, the LT analysis also utilized a set of somewhat strict pre-fiducialization cuts on ionization quantities. At these low energies the ionization quantities are very close to their noise distributions, but can still be used to remove the most egregious surface events to avoid contaminating the signal region with events that are simply due to poor charge collection. A set of radial and z-symmetry

cuts were created that reject events at high radius as determined by requiring strict requirements on the inner electrode, and which deposited the majority of their energy between both faces.




Figure 6-7. Cartoon schematic of how the ionization fiducialization was performed for the LT analysis. A) The radial cut was set so that events with very low inner electrode energy were excluded if they had high outer electrode ionization. An additional requirement was that the then the inner electrode not share much energy with the outer electrode. Towards the noise donut, the cut was allowed to encompass the outer envelope before becoming flat to avoid the outer channel. B) In the face-face symmetry plane an event must occur along the symmetric band where side 1 energy is approximately equal to the side 2 energy. At the lowest energy values the noise donut is enveloped by the fit. Courtesy of Todd Doughty.

In addition to removing the clear outlier regions of the detector, the signal region needs to be defined as well. Since yield is not a useful discriminating technique at low energies (recall it flares out asymptotically, and that we have poor ionization signal resolution) the signal region was defined in a plane determined by the phonon total energy `ptNF` and an ionization energy estimator called `qimeanOF`. As the name implies, this quantity was the simple mean of the two inner-electrode ionization energy estimators. The choice of using `qimeanOF` arises from several considerations. The first is that the `qsumOF` quantities contain outer channel information in them, which will introduce unnecessary noise into the energy estimates, so this along with the strict inner fiducialization implies we should avoid anything with the outer channels. This still leaves different ways to utilize

the inner channels. One possible method is to just use the maximum inner channel as the estimator. If a given channel has signal S_i and noise σ_i then the signal to noise ratio for this quantity can be set as $R_{qimax} = S/\sigma$. But let's consider the other possible method, by taking the mean of the two sides. In the bulk (recall, we implement a strict ionization symmetry requirement), to first order, the signal strengths for the two channels should be fairly equal, and as the channels are on opposite sides of the detector, their noise should be roughly uncorrelated.⁴ Since signals add directly and uncorrelated noise will add in quadrature, we can write the signal to noise of `qimeanOF` as $R_{qimean} = (S_1 + S_2)/\sqrt{\sigma_1^2 + \sigma_2^2}$. If $S_1 \sim S_2$ and $\sigma_1 \sim \sigma_2$, which is not a bad approximation, then the signal to noise ratio of `qimeanOF` becomes $\sqrt{2}S/\sigma$. Thus `qimeanOF` should have roughly a factor of $\sqrt{2}$ better signal to noise than simply choosing the largest energy value.⁵ This quantity does have a problem with surface events since the opposite side will have no signal the net result will be an underestimate of the ionization energy by about a factor of 2, resulting in surface events appearing to leak into the signal region.

Using neutron calibration data from ^{252}Cf , we binned the total phonon energy and fit the `qimeanOF` data in each bin to a gaussian distribution. The median, and 1-3 σ contours were fit to a functional form derived from Dave Moore's thesis [109]:

$$f(E_t) = \alpha_1 \frac{E_t^{\alpha_2+1}}{1 + \alpha_3 E_t^{\alpha_2}} + \alpha_4 \quad (6-2)$$

⁴ This does appear to be the case. During the start of R133 I performed a study to examine correlations in noise amongst all the channels and detectors, and found that there is very little correlation between opposite side charge channels [108].

⁵ It also has a single-peaked noise distribution which makes it much easier to work with when setting bands than a bimodal quantity.

This functional form was ultimately chosen as it can curve at low total phonon energy to account for OF energy bias⁶ and it becomes linear at higher energy which is expected from Lindhard theory.[88]




Figure 6-8. The nuclear recoil band for IT2Z2 is shown here derived from the parameterization approach, with 68% and 95% confidence level contours shown. Similar bands were constructed for each of the LT detectors from Californium datasets. Courtesy of Julien Billard [110].

An empirical model of the ionization energy vs phonon energy was also developed for this analysis, allowing one to convert from measured total phonon energy to recoil phonon energy. This was required since the estimated Luke phonon energy from ionization alone on an event-by-event basis would suffer from the poor ionization signal to noise. The model chosen by Julien eventually took a functional form shown in Equation 6–3,

⁶ Recall the OF is going to tend to select positive noise fluctuations when reconstructing rather than negative ones

which was chosen such that again it behaves linearly at high energy and at low energy it is able to curve, and a constant term can account for possible OF bias, and proved numerically stable. With this model one was able to convert from total phonon energy to recoil phonon energy for each of the LT detectors.

$$f(E_p) = \alpha_1 + \alpha_2 E_p + 10^{\alpha_3} \text{erf} \left(-\frac{E_p}{10^{\alpha_4}} \right) \quad (6-3)$$




Figure 6-9. The empirical charge model used in the LT analysis for IT2Z1. The best fit line is in red, with the mean `qimeanOF` from a particular energy bin shown in magenta. Lindhard theory prediction is in dashed and is in generally good agreement at higher energy. At low energies deviations are observed on either side of the Lindhard line, but no more than $\sim 50\%$ for energies above 2 keV_t . Courtesy of Julien Billard [111].

6.4 Background Model

With good quality data selected, and a signal region defined, the last ingredient needed for an analysis is a way to characterize the background. For this analysis a data-driven background model was used. This was needed due to a lack of statistics at low energy, the fact the signal region itself was blinded, and that energy resolution effects reduce the effectiveness of the iZIP background discrimination methods (which are quite powerful at high energy via the ionization signals). A well developed description

of the background model is presented in a thesis by Adam Anderson.[112] In general the background model samples data from outside the LT analysis region at much higher energy where background discrimination is sufficient to identify regions of parameter space where particular backgrounds should exist. For example, a low-yield surface event population can be selected by tagging events between the bottom of the ER band and the top of the NR band, in the energy range of 50 - 120 keV_t which have very asymmetric face-face ionization signals. This population of events then constitutes a set of template events. The raw traces of these template events then have their amplitudes scaled down to a lower energy value that is randomly sampled, but with weights derived from an energy input file which reproduces the energy spectrum of the background in question (derived from Monte Carlo simulations). The scaled down pulse is then injected with experimental noise derived from a random trigger⁷ to create a simulated event. This “fake-event” is then run through the standard `cdmsbats` processing chain (but not every RQ/RRQ is computed) to obtain event information for the simulated pulse. In this way we can create a rich population of events that share the energy and partition relations for a particular background, and rely on energy resolution effects to smear any systematic energy dependencies away.⁸

The background models are then normalized to the observed populations in the data: The 1.3 keV L-shell activation model was scaled to the observed rate of the 10.4 keV K-shell activation line in the (unblinded) post-Cf data period, Compton scatter γ s were scaled to the observed rate in the ER band (fit similarly to the NR band, but

⁷ Specifically, an in-run random, not one of the 500 start or end of run randoms, to avoid any potential systematic effects the exist at the start or end of a series.

⁸ Indeed, at low energies most of the partition quantities do suffer from loss of discrimination ability due to energy resolution smearing out the shape of the distribution. This suggests this type of data-driven model will not be as useful in a high-mass WIMP search, where one cannot ignore position and energy dependent variations.

with a different functional form), and the lead-sourced events were normalized based on observed rates in the post-Cf data sample as well. Each model, and the quantities chosen to simulate are based on the expected manner in which leakage into the signal region can occur as outlined in Figure 6-10.




Figure 6-10. The location in energy, and sources of the expected primary backgrounds in the LT analysis. Externally sourced gammas can Compton scatter in the detectors and at low energy leak into the signal region. The ionization energy vs phonon energy plane provides a primary discrimination plane for this background. Events originating from the ^{210}Pb decay chain arise as surface events due to ^{206}Pb nuclei or β 's interacting with the detector. These bleed into the signal region significantly in the ionization energy vs phonon energy planes, but can be eliminated with radial and z-symmetry fiducialization. Internal activation is also a problem as it shows up in a generally localized region of the energy plane. Since it has a well defined width and shape this can be removed by selecting a contour around it. Note that in each tower 2 detectors are end caps and have a layer of copper opposite a face as well. The 3 simulation figures are courtesy of Adam Anderson.

6.5 Rescaling the pzpart Model

Once the background model was generated various cross-checks were performed to verify it accurately represented the expected iZIP response. For many of the discriminators the simulations performed quite well in reproducing the expected distributions. For example in Figure 6-11, the side 2 inner-electrode ionization energy vs side 1 inner-electrode energy is plotted for a simulated dataset that is meant to represent inner-electrode surface events (originating from ^{210}Pb β 's). We would expect the distribution in the side-sum ionization plane to occur along the axes, and for the ionization energy to be reduced somewhat between the ER and NR bands.




Figure 6-11. Example plots from the pulse simulation methodology, for inner ionization electrode surface events. A) The energy distribution for surface events in the face-face plane should be aligned along each axis, which is seen here quite clearly. B) In the q_{imeanOF} vs p_{tNF} plane surface β 's have a reduced ionization energy compared to ERs, but still have more than the typical NR event.

While these distributions are well behaved, the phonon partition quantities did not initially agree with observed behavior. In particular the phonon z-partition quantity was narrower than observed in background data, and it was relatively energy independent. In reality, `pzpartOF` has some energy dependence due to local TES saturation in phonon channels as described in Section 4.2.4. Since the template events used to create samples

for the pulse simulation were sampled at high energy, the reconstructed simulated **pzpart0F** values are centered around the lower values associated with higher energies and do not accurately represent the low energy region.




Figure 6-12. A comparison of the phonon z-partition estimator **pzpart0F** for real data and simulated data illustrating the effect of local saturation. A) Observed energy dependence of **pzpart0F** in actual data. The ionization based symmetry flag tags many events (red) in the outer bands but not all. Bulk events as defined by **pzpart0F** are roughly in the central region bordered by the two inner solid green lines. The outer envelope of the surface event bands are shown in light blue. B) The same distribution as seen from the simulated data. The solid lines are transposed from part A to illustrate the problem caused by local saturation. High energy events have much narrower **pzpart0F** bands than at low energy.

To fix this, an energy dependent rescaling method was devised. Using the available sideband data, the **pzpart0F** distributions for both the data and simulation were split into two halves around the mean value to select each peak separately. Four energy bins in **ptNF** were used: 3.8 - 7 keV, 7 - 13 keV, 13 - 20 keV, and 20 - 40 keV. In each of these energy bins the mean of the data **pzpart0F** distribution was found by fitting a gaussian distribution, after which the the simulated **pzpart0F** values were multiplied by small increments and a gaussian fit was performed to obtain the new mean of the simulated data. This continued until the fit means of the data and simulation were within an acceptable tolerance. An example of this rescaling is shown in Figure 6-13 for two energy bins.




Figure 6-13. Two energy bins are shown with the original simulated `pzpart0F` distribution (blue dashed line) along with the actual observed (red) and new rescaled distribution (green). A gaussian was fit to the each peak of the observed distribution to obtain the mean value of `pzpart0F` in this energy bin. A) The rescaling process for energy bin 2 (7 - 13 keV). The rescaled peak shifts to match the observed means of the distribution. B) Same process for a higher energy bin (13 - 20 keV) again showing good agreement.

A small systematic error is introduced into this since if we have a set of randomly distributed data X with mean μ and variance v and we multiply the mean such that $\mu' = \alpha\mu$, then the variance of the distribution becomes $v' = \alpha^2v$. In each case however the scaling factor was relatively small making the widening of the `pzpart0F` peaks negligible.⁹ Once a set of scale factors are determined at the specified energy bins, they are interpolated across the entire energy range using a shape preserving cubic spline to obtain a rescaling vector that will act on each simulated `pzpart0F` value. This leads to a gap in the distribution in the noise wall, so the interpolated function is matched smoothly to an error function that drives the re-scaled simulated distribution towards the unscaled version in the noise wall (such that at 0 keV the scaling factor would be 1). The net result of this is shown in Figure 6-14 where the original simulated distribution is overlaid with the new rescaled distribution. The rescaled quantity was named `pzpart0FSIMcorr`. A similar rescaling was performed by Adam Anderson for the simulated phonon radial quantity, but was done so in an energy-independent fashion.

6.6 Discrimination Cut Methods

In order to obtain a final result, a set of discrimination cuts are required that should only allow signal-like events to pass. In this analysis several different methods were employed and the one that had the best predicted sensitivity was chosen as the method to use on the blinded data. Each of the methods used four quantities as their primary discriminants: `ptNF` as the primary measure of event energy, `qimean0F` as the primary measure of ionization energy and to provide limited recoil-type information, `pzpart0F` to provide rejection power against surface events on the detector faces, and `prpart0F` to

⁹ Since the actual `pzpart0F` bands broaden at low energies due to energy resolution effects, this broadening effect may actually be less problematic since the scale factors will themselves get larger at lower energies as well.




Figure 6-14. The rescaled `pzpart0F` energy distribution is shown (purple) against the original distribution (green). The rescaled quantity is in better agreement with observed surface events as tagged by ionization, and demonstrate the flaring effect observed in actual data. The bands are also shifted outward as expected, and merge into a noise wall at the lowest energies.

provide rejection power against events at high radius near the sidewall of the detector, as outlined in Section 6.4.

The first method, called Method 1 (M1) relied on a series of serially tuned rectilinear cuts to remove each background. The order of operations for cut tuning was chosen to be proceed as follows:

- NR band: First the NR band was tuned, and a notch was inserted into it to avoid the 1.3 keV activation region due to the L-shell capture background. This notch appears in the 3σ wide NR band and effectively removes the L-shell background.
- Phonon Radial cut: Next the phonon radial cut was tuned against the lead-sourced decay chain events for high radius. To mimic the signal, a set of good neutron calibration data was re-weighted to mimic a WIMP energy spectrum, which was

then efficiency-corrected for each of the numerous quality and threshold cuts. This was repeated separately for 4 different WIMP masses: 5,7,10, and 15 GeV/ c^2 to ultimately construct four different cuts. The final discrimination cut on this quantity was defined as the logical OR over all 4 mass-dependent cuts. Each individual cut itself was tuned in separate energy bins by allowing a cut boundary to vary and observing the proposed cut leakage from the background model vs the efficiency on signal, and a 90% Poisson upper limit on the count was used to construct an upper limit on the expected cross-section. The optimal cut for a given mass was thus the cut that minimized this expected cross-section.

- **Phonon Z Symmetry Cut:** This was the last serial cut to be tuned and was performed in a similar fashion to the radial cut. It used the lead-sourced surface events from low-radius templates. Like the radial cut, the signal model used here was also re-weighted neutron calibration data, and the same four WIMP masses were used to compute 4 separate cuts on `pzpartOF` which were then combined by taking a logical OR. The `pzpart` quantity has two cut boundaries however, unlike the radial quantity, and so both sides were allowed to vary independently of the other. At each energy bin both cut boundaries were varied and a background leakage and signal efficiency were computed and used to obtain the 90% Poisson upper limit as well.

6.6.1 Phonon Symmetry Cut

I will describe the results of the phonon symmetry cut in a bit more detail here. The optimization technique relies on using the background models and their appropriate weights. The background model weights are normalization constants such that the model will reproduce the observed spectrum in the data, and sum up to the total expected number of background events. Thus if we create a cut on the background model and apply this cut to the weight vector, the number of leakage events passing the cut will simply be the sum of the weight vector with the cut points removed. The signal model relies on a similar process to compute the efficiency, except here the weight vectors are defined such that the integral of any particular bin will provide the spectrum-averaged efficiency (SAE). The SAE is basically an energy-dependent weighting of the efficiency by the WIMP-recoil spectrum. If we multiply the spectrum-averaged-efficiency by the experiment's exposure $\mathcal{E}$

we get the spectrum-averaged-exposure.¹⁰

$$SAE = \frac{\int_{E_{th}}^{\infty} \frac{dR}{dE} \epsilon(E) dE}{\int_0^{\infty} \frac{dR}{dE} dE} \quad (6-4)$$

Where $\epsilon(E)$ is the energy dependent efficiency of the various pre-selection cuts, triggers, and analysis threshold cuts. With this a series of proposed discrimination cut boundaries can be computed in energy bins of **ptNF** and then the observed leakage and efficiency used to compute 90% Poisson upper limit on the cross-section, which is computationally easy compared to running multiple instances of the optimum interval code[113]. To do this, first the total number of observed leakage events for a proposed cut boundary is determined, N_{bg} . Next the expected 90% CL upper limit, $\bar{\mu}$, can be computed via:

$$\bar{\mu} = \sum_{i=0}^{\infty} \mu(i) \frac{N_{bg}^i e^{-N_{bg}}}{i!} \quad (6-5)$$

Where $\mu(i)$ is the 90% CL Poisson upper limit on the mean number of events if one observed only i events. This can be found by solving for the value of μ that makes the Poisson CDF = 0.1:

$$\mu \mid e^{\mu} \sum_{j=0}^i \frac{\mu^j}{j!} = 0.1 \quad (6-6)$$

Thus for every term in the sum in Equation 6-5, a different $\mu(i)$ is determined, and this is then multiplied by the probability of actually observing i events with an expectation of N_{bg} . By repeating this and summing one obtains a weighted average of the 90% CL upper limit, which is thus the expected 90% CL upper limit. This can be computed fairly rapidly as the values of $\mu(i)$ can be pre-computed to a sufficiently large value of i and then fed into a truncated sum for Equation 6-5 since at sufficiently large values of i the factorial term will drive subsequent terms to ever diminishing values. If we then divide this by the number of expected WIMPs given the cut efficiency under consideration, this gives us the

¹⁰ Which also has the abbreviation SAE

expected 90% CL upper limit on the cross-section. The expected number of WIMPs can be computed by choosing an arbitrary cross-section for the mass you are optimizing for (e.g., a 5 GeV/ c^2 WIMP) and computing the WIMP rate spectrum in the energy range under consideration.

An additional complication for this process is that the data are often discrete enough that an optimizer can get stuck if it chooses a boundary not suitably populated by data. To counteract this a weighted KDE was constructed for the signal and background models using the weight vectors already on hand. The KDEs were initially constructed via MATLAB's built-in `ksdensity` function. The resulting KDEs are normalized so that an integral under them is equal to 1, so one can renormalize this integral by the sum of the weight vector. Then, if one integrates the KDE on an interval of $[a, b]$ as defined by a proposed cut boundary, this will give the fraction of total leakage events contained within. Since the built-in function was a bit too slow, a manual integration method was used that took a sum of the product of the weights with the CDF of the distribution which proved much faster. Initially the optimization algorithm used was MATLAB's built-in `fminsearch` but this had very poor convergence due to the leakage being very small and flat. A variety of optimizers were explored, from `patternsearch` which does a direct-search, to simulated annealing, which took a sufficiently long time as to be useless. Once the KDE integral method was developed, `fmincon` could be used which uses a gradient-descent method and as a result is very quick, and never differed significantly from the results of `patternsearch`. The result of all 3 serially tuned cuts is shown in Figure 6-16.

6.6.2 Method 2 and the BDT

Another method used, called Method 2 (or M2), attempted to do the same thing as M1 simultaneously. This method used the signal and background models for each type of background and developed a 3D KDE [116] in order to avoid issues with sparse data becoming even more likely in this higher dimensional space. For the optimization on the




Figure 6-15. A direct scan of the sensitivity parameter space as a function of variable cut boundary positions. This method shows the resulting 2D map for the numerical optimizer determined minimum and the direct search minimum. The two are in very close agreement for this particular case. The top-side cut had a much more clear minimum to fall into, but the bottom side had a relatively flat variation which most likely accounts for the disagreement between the direct search and minimizer. In any event this difference is small.

qimeanOF quantity, it was shifted to a normalized value by subtracting the mean NR band value in some bin and dividing by the width of the NR band. The chief advantage of M2 over M1 is that it can take advantage of correlations in the discrimination quantities that are not possible in the serial method. This can allow for cuts to be loosened in some planes for significant exposure gains while tightening in others without loss of exposure (but this isn't guaranteed to happen).

The final method was a boosted-decision-tree (BDT) approach. The BDT is a machine learning algorithm that attempts to classify events as signal like or background like. Like the cuts based methods, the BDT used the same 4 discriminator quantities, and was run on the same 4 WIMP masses as the other methods. The BDT outputs a single discrimination parameter with which a cut can be set. It is normalized such that a score




Figure 6-16. The sequential series of Method 1 cuts tuned for the LT WIMP search. Similar plots exist for each LT detector, and for each of the 4 WIMP masses analyzed A) The notched NR band cut is shown here. The 1.3 keV activation events are selected in cyan, and as such anything within the circle is excluded from the NR band, making the band effectively dip under the activation region. Courtesy of Kristi Schneck [114]. B) The results of tuning the phonon radial discrimination cut are shown here. This was a single boundary formed by excluding events at high radius. Courtesy of Adam Anderson [115]. C) The result of tuning the phonon z-symmetry cut. This cut allowed the top and bottom side selections which were not constrained to have the same absolute value in `pzpartOF` boundary.




Figure 6-17. The results of the Method 2 optimization on the discrimination quantities. In each plane the cuts may appear to be worse than the M1 cuts, but they can take advantage of correlations in quantities to have more powerful rejection by allowing more leakage in one quantity to gain much more exposure while other quantities can be tightened without loss of exposure. Courtesy of Adam Anderson [115]

of -1 is entirely background like and +1 is entirely signal like. The output of the BDT is shown in Figure 6-18 which shows the excellent discrimination against background.




Figure 6-18. Example of a boosted decision tree output on the combined LT dataset, and optimized for a $10 \text{ GeV}/c^2$ WIMP. The various components of the background model show up as the colored histograms all with mostly negative BDT score. A hypothetical WIMP signal (in grey) is clearly different than the background. Courtesy of Julien Billard.

The BDT cut boundary was done by simultaneously varying it for each detector while trying to optimize the sensitivity. This approach was shown to produce a better overall limit than setting the boundary on each detector independently. The BDT method was determined to be the best performing discriminator after comparing each method's expected sensitivity using the background model. This was the assumption going into the analysis, but it was important to use the other methods as a cross-check since no WIMP-search analysis had been performed with a BDT prior to this. The sensitivity curves were generated using the optimum interval (OI) method [113] with standard halo and form factor assumptions. Using the background model, poisson fluctuations around the expected number of events were used to instantiate multiple test-experiments that were fed into the OI method to produce limit curves. The median of these curves was used as the best guess result. None of the 3 methods ever crossed in the sensitivity plane; if one

was better performing than another for a particular WIMP mass, it remained better for all WIMP masses.

6.7 LT Analysis Results

6.7.1 Exposure Weighted Efficiency

In order to determine the final limit of the analysis several steps must be completed along the way. The total efficiency of the experiment as a function of energy must be characterized. This is typically computed in a staged fashion in which the efficiency of the data quality cuts is computed, then in addition any other fiducialization cuts are added in, and finally the discrimination cut is added for a final efficiency curve. To measure the efficiency of data quality cuts a simulation method using randoms plus a pure pulse template to provide high statistics. This is valid since the data quality cuts are primarily meant to reject events with poor pulse shape reconstruction and so should pass events with good underlying pulse-shapes compared to the templates. By smearing a template with the noise from random triggers, it is then easy to measure what fraction are rejected by the cuts. In addition to data quality cuts a set of pre-selection cuts were used. These included the ionization fiducial cuts, as well as the NR band cut. To determine the efficiency of these, neutron calibration data was used as a proxy for signal events and the passage rate was computed from the data itself. The BDT was chosen as the final discrimination method and its efficiency was also computed using neutron calibration data. Unfortunately the neutron calibration does not produce events that are perfectly like a WIMP. In the ^{252}Cf calibration datasets, some γ s are produced as well, and neutrons can multiply scatter, whereas WIMPs are unlikely to do so. These multiple scattering effects were corrected out by using a Geant4[117] simulation of the neutron interactions to determine a correction factor, and the γ s were subtracted out by using the ^{133}Ba calibration dataset, normalized to the rate of γ s arising during a neutron calibration to obtain an energy dependent subtraction factor.[118].

Table 6-1. LT Analysis Exposures

Detector	Mass [kg]	Livetime [days]	Exposure [kg-days]
IT1Z1	0.6095	131.6	80.21
IT2Z1	0.5974	138.8	82.92
IT2Z2	0.5913	136.8	80.89
IT4Z2	0.5974	146.3	87.40
IT4Z3	0.5943	141.0	83.80
IT5Z2	0.6065	136.3	82.67
IT5Z3	0.5913	134.0	79.23
Total	4.188	964.8	577.1

The efficiency is then weighted by the exposure of each detector. Recall the exposure is just the product of the detector mass and the total livetime of the detector, $\mathcal{E} = MT$.

The total livetime is detector dependent and is determined, generally, by how much survives time period cuts and cuts that remove series as a whole (in essence, cuts that do not reject based on specific event characteristics, with the exception of the pre-pulse baseline cut which also removes the livetime for that event). For this analysis the detector masses and livetimes were fairly comparable for most detectors and are listed in Table 6-1. Once the detector dependent exposures are computed the energy dependent analysis efficiency curves are weighted by this exposure and combined as shown in Figure 6-19.

6.7.2 Unblinding

The LT analysis opened the box on February 3 around noon Pacific time. Adam Anderson initiated the opening process with myself following at the same time as an independent cross-check to verify what we saw was indeed the same. Upon unblinding the dataset, we observed 11 events passing the BDT cut in 4 of the 7 detectors used for this analysis. The candidate events all appeared to be good looking events with the exception of the lowest energy event on IT2Z2 ($\text{ptNF} = 2.09$ keV) which had a boxy pulse shaped that resembled a noise fluctuation. The observed event counts were compared to the background model expectations and a p-value was computed for each detector assuming a Poisson distribution for fluctuations. The expected total number of background events was




Figure 6-19. The analysis thresholds are shown here as a function of energy and level of cut application in sequence. The data quality cuts (thin black line) typically have very high efficiencies across all energy ranges. The first notable hit comes from the application of the trigger efficiency and energy thresholds (grey line). These were derived from the Barium calibration dataset[119]. Steps in this curve occur due to time variation of the trigger thresholds and are detector dependent. The application of the fiducial and NR band cuts (blue line) is where largest efficiency drop across high energy region occurs. This primarily is the result of the strict ionization selection cuts. The BDT cut (red line) then heavily removes events at low energy as at such low energy most events are hard to distinguish from background and tagged as such. The final efficiency curve itself is smoothed with a cubic spline to remove a jig-saw pattern that would emerge from computation in energy bins. See [59] for more information.

found to be $6.07^{+1.14}_{-0.82}$ events. A breakdown of each detector's observed counts compared to the background model is shown in Figure 6-20.

Immediately it was seen that IT5Z3 appears to be particularly troubling given the unlikely fluctuation observed, and the energy at which the events occurred at. These events are all at higher energies than all the other observed events in the signal region, as can be seen in Figure 6-21.




Figure 6-20. Expected backgrounds for each detector derived from the background models on each detector are shown with a 1σ background fluctuation band in red. The range of background fluctuations possible to still have $p > 0.05$ is shown in blue. IT5Z2 and IT5Z3 especially are anomalously high. In particular IT5Z3 had a background expectation of $0.13^{+0.06}_{-0.04}$, giving 3 observed events a $p = 4 \times 10^{-4}$

IT5Z3 is not a fully functional detector. It has both an ionization and phonon short that occurred within it.¹¹ The most relevant short is the outer ionization channel on side 1, which means it has no way to radially fiducialize on that particular side, and that the E-field is slightly different. In retrospect it seems likely the background model was not properly created for this particular detector due to our lack of understanding of this short. Because the outer ionization channel is shorted and grounded the electric field will be fundamentally different in that region of the detector, allowing high radius events to travel into the inner region. A full study of the implications of this is an active area of research in the collaboration, and in particular amongst the Detector Monte Carlo group. Kristi

¹¹ In the analysis the phonon short resulted in the detector being treated as two separate detectors, with a pre and post-short model.




Figure 6-21. The locations of each of the observed candidate events is shown as a function of ionization and phonon energy. Candidate events are color and shape coded by detector. The 3 anomalous events on IT5Z3 are in yellow stars. Grey events represent events that pass all cuts prior to the BDT. The colored contours are the 95% confidence contours for a 5,7,10,15 GeV/c^2 WIMP in order of lightest to darkest contour.

Schneck has explored some of the effects of this in her thesis [65], which seems to confirm the initial suspicions regarding the background model.

The 11 candidate events themselves appear to be well distributed in time with no small period of time containing a plurality of the events. A summary of the event dates, energies, which mass-optimized BDTs the event passed, and whether they passed the sequential cut tuning discrimination method (M1) or the simultaneous cut tuning method (M2), is listed in Table 6-2. As the table reveals the candidate events from the primary BDT method passed the other attempted discrimination cuts as well, with the exception of the 2.41 keV event on IT2Z1 which did not pass either. Lest one think the other methods did comparably well, Method 1 passed a total of 66 events (most on IT5Z3 with 29) and Method 2 passed 61 events (with IT5Z3 having 21 events, but IT2Z1 having 20

Table 6-2. LT Candidate Events

Detector	Date	Energy [keV _t]	BDT Passed? [GeV/ <i>c</i> ²]	Passed Other Method?
IT2Z1	May 05, 2013	2.30	5	M2
IT2Z1	Jun 19, 2013	2.41	5	None
IT2Z2	Nov 06, 2012	2.09	5	M1,M2
IT2Z2	Feb 23, 2013	3.24	7	M1
IT5Z2	Oct 05, 2012	7.18	7,10,15	M1,M2
IT5Z2	Nov 21, 2012	2.15	5	M1,M2
IT5Z2	Jun 15, 2013	3.56	7	M2
IT5Z2	Jun 27, 2013	2.76	5,7	M1,M2
IT5Z3	Oct 12, 2012	9.79	15	M1,M2
IT5Z3	Feb 23, 2013	11.99	15	M1
IT5Z3	Apr 24, 2013	8.74	15	M1,M2

as well), confirming that the BDT approach did indeed do a much better job than a cut based approaches in this case.

6.7.3 Limit on Low Mass WIMPs

With the box opened, the final limit curve, with systematic and statistical uncertainties folded in, was computed, again using the OI method [113]. The standard halo assumptions and form factor are used again: $\rho_\chi = 0.3 \text{ GeV}/c^2/\text{cm}^3$, WIMP velocity of 220 km/s, $v_{\text{earth}} = 232\text{km/s}$, galactic escape speed of 544 km/s, and a Helm form factor as discussed in Chapter 1.

The SuperCDMS LT result is in tension with most of the other low-mass searches by a large margin over the mass range presented here, and excluded a new region of spin-independent parameter space in the range of 4 - 6 GeV/*c*² at the time. Very recently new results from CDMSlite have pushed lower their threshold down to 70 eV_{ee}, remaining below the SuperCDMS LT curve until around 6 GeV/*c*². [76] Effective field theory arguments provide some room for more favorable interpretations between different target experiments, as discussed in [65, 126], but not for similar target materials.




Figure 6-22. The final limits from the R133 LT analysis are shown here in black with the 68% and 95% confidence level systematic uncertainties shown in grey around the black curve. The pre-unblinding expected sensitivity is shown in dark green (68%) and light green (95%) assuming no signal. The final limit is within the expected range, except at high WIMP masses due to the unusual events on IT5Z3, which occur at high recoil energy. The closed contours are 90% confidence levels for DAMA/LIBRA [120] (tan), the CDMS II Silicon result [44] (blue), CoGeNT [43] (yellow), and a 95% CL is shown for CRESST II [45] (pink). The exclusion limits shown are 90% confidence-level curves for CDMSlite[75] (solid dark red), XENON10 S2-only[121] (dashed dark green), CDMS II Ge low-threshold[122] (dashed-dotted red), LUX[58] (solid green), CDMS II Ge[123] (dotted dark red), EDELWEISS low-threshold[124] (dashed orange), and CDMS II low-energy likelihood[125] (dashed dark red).

CHAPTER 7

TOWARDS A HIGH MASS WIMP SEARCH

In addition to a low-mass WIMP search, SuperCDMS Soudan is able to search for high-mass WIMPs ($m_\chi > 20\text{GeV}/c^2$). In this mass range the recoil energies are much higher and the iZIP's powerful ionization discrimination abilities can be exploited to their full extent. However a search at higher masses comes with its own unique challenges not faced by the LT analysis. This kind of search cannot use a pulse simulation background model since at these high energies position dependent effects on the pulse-shape will not be smeared away like they were when scaling down to low energy. Additionally, even if this was a method that could be tried, it would be preferred to sample events from an energy range outside the analysis window. However at such high energies, local saturation effects and lack of statistics makes this infeasible. A high-mass search also needs a lot of exposure to be able to set a low limit on the cross-section, which translates to a long running time in the context of this work. Whenever an experiment runs for an extended duration there are always the risk that environmental noise, or detector performance can change for the worse, producing undesirable time variations in the experiment. As I discussed in Chapter 2, there were many such time dependent variations in this experiment. Nevertheless the iZIPs are very good at separating background from signal at the energy ranges we are considering for this analysis ($\sim 6 - 150 \text{ keV}_r$) which makes the real challenge the extrapolation of the low-probability background-distribution tails into the signal region. As the analysis is still underway what will follow is a description of the progress towards the completion of the HT analysis including the description of an unforeseen problem with the iZIP response that was discovered when constructing a model of the gamma leakage. Further details on the HT analysis will be forthcoming in Todd Doughty and Brett Cornell's theses.[\[127, 128\]](#)

7.1 High Threshold Dataset

The HT dataset encompasses the entire c334 time period for WS mode and the c335 data period for calibration data, excepting blocks of time due to special operating conditions on a few detectors (such as reverse bias mode or CDMSlite), and the last two months of R134 due to excessively high noise and unstable base temperature. The detectors used in the LT analysis have already been unblinded up to a recoil energy of about 10 keV_{nr} and as such this energy range is not used for the HT analysis on those detectors, otherwise the energy range used for the HT analysis will push to as low an energy as possible before background leakage becomes dominant and an upper recoil energy of 150 keV_{nr} (approximately 232 keV_t). The upper limit is based on kinematic considerations for the recoil rate expected in a germanium target, and on the fact that at such high energies our neutron calibration source does not have a high number of statistics.

The signal region is defined by the NR band in the yield vs recoil energy plane. As mentioned earlier the data for this comes from specific periods of time when a ²⁵²Cf neutron source is used. The yield distribution in a lightly fiducialized region of parameter space are fit to a gaussian distribution in energy bins of `precoiltnf` to obtain the band means, which are fit to a power law of the form $\mu = \alpha E_r^\beta$. The band widths arise from fitting a gaussian to the residual distribution and have an empirically derived functional form [129]:

$$\sigma = \frac{\sqrt{aE_r^b + c}}{E_r} \quad (7-1)$$

With the completion of the KS and QS test cuts described in the previous chapter, the livetime of the HT dataset can also be computed for each detector. This is an important quantity since it contributes to the exposure that sets the sensitivity of the experiment. The detector masses and livetimes (pre-unblinding) for the HT dataset are listed in Table 7-1.

Table 7-1. HT Analysis LiveTimes (pre-unblinding)

Detector	Mass [kg]	Livetime [days]
IT1Z1	0.6095	285.65
IT2Z1	0.5974	317.20
IT2Z2	0.5913	322.29
IT2Z3	0.5791	313.06
IT3Z1	0.6034	234.40
IT3Z2	0.5913	307.70
IT3Z3	0.6095	242.57
IT4Z2	0.5974	322.76
IT4Z3	0.5943	347.43
IT5Z2	0.6065	144.80

In the list of detector livetimes the two source detectors have less total livetime due to periods of time when they were operated in reverse bias mode and the first several months of operation when they were kept unblinded for special studies. IT5Z2 has such low livetime due to it being the CDMSlite detector, during which time it cannot be used for a HT WIMP search. There is an uncertainty associated with each livetime due to the manner in which the recorded livetime is rounded in the data stream. This effect is extremely small (on the order of 9 - 18 seconds) when compared to the magnitude of the measured livetime. As described earlier about 40-50% of the raw livetime is removed by various livetime quality cuts. This is a somewhat large amount, however no particular set of cuts indicates it is removing an excessive amount of livetime. Indeed, each cut appears to be tuned properly and is well motivated for removal of questionable time periods, indicating the HT exposure just happens to have a high rate of questionable time periods within it, most likely arising from old electronics, and old hardware introducing excessive noise and trigger bursts.

7.2 Unifying Radial Discriminators

The ionization radial discriminators are powerful quantities for rejecting events that occur near the sidewalls where poor charge collection occurs. However as mentioned back in Section 3.2.1 the electron side collection suffers from a much larger incidence of shared events than the hole side due to the nature of electron propagation in the

germanium crystal. Some optimization methods, such as a BDT, can handle this just fine, by using both radial quantities as inputs, however when making plots, or making cuts based discriminators, if one can minimize the number of quantities to optimize over, then it is a good idea to try. A simple average of the two radial quantities is not quite appropriate since it will grossly underestimate the radial distance near the faces. Consider a surface event that has a high radial position on S1. The total energy collected on S2 will be very small, such that the signal is consistent with noise resulting in a radial quantity that can be potentially very small. The average then would be in the shared region instead of clearly near the outer edge. This example suggests that there is an important z dependence to the radial estimators. Depending on where we are in the detector, one radial quantity or another is likely to perform better.

The simplest thing to do is to take a z -weighted mean of the radial quantities. This approach uses the value of `qzpart0F` to create a weight vector such each side's radial quantity has a weight of 1 on the appropriate face, and a weight of 0.5 in the middle of the detector. This approach does improve radial estimation in the shared region, however appears to broaden the high radius peak in some cases worse than the electron side collection alone. This may be due to a combination of effects: at low energies the z partition value may reconstruct to random values, and likewise at sufficiently high z partition values, the radial quantity may fluctuate to such a high value such that even with a z dependent weight it still is sufficient to interfere with a sensible value.

Another method is to instead use some physical motivation for the z partition weight vector. Instead of using a full linear function, a ramp function is somewhat more appropriate. Since the hole side collection is expected to perform much better at high radii due to the lack of oblique propagation, the ramp function de-weights the electron-side collection for z symmetry values less than 0. In the electron side, the holes are still expected to perform very well up until one approaches the surface of the detector. To account for this the ramp function simply uses the z partition value as the weight in

this region. Thus at a z partition value of 0.5 the two radial estimators will be equally averaged together, as opposed to the simple linear z weight described above where this occurs at a z partition value of 0. A histogram of some of the various ionization radial estimators is shown in Figure 7-1. The simple average approach gives some marginal improvement in the middle of the shared region but introduces a much larger uncertainty in high radius (as indicated by the broadened peak). This is due to the SE effect described earlier. A simple linear z -weighted average improves the situation significantly but still has a less well-defined radial peak than S2 collection only. By using a ramp-function weight such that the radial quantity used is only the hole based estimator for the hole side, and a linear weight between the two for the electron side, a dramatic improvement is reached across all radial partition values. For partition plane plots in the rest of this work, radial quantities will be plotted with this ramp weighted, or half-linear, partition quantity, called `qrpartOF_zhalf`.

The behavior of the various radial partitions with surface events can be seen by examining the two ^{210}Pb source detectors. The distributions for these two detectors are shown in Figures 7-2 and 7-3. In both detector cases we can see the side with the source implanted generally has the better single side radial estimator, and in both cases the half-linear z weighted quantities perform best over the other combined estimators. The source detectors show the problem that arises from a simple mean between the two radial partition values for SEs.

Finally, although the DMC is still under development, the ionization propagation physics is sufficiently mature that some insight can be gleaned from simple simulations. In Figure 7-4 we can see that the behavior of the combined estimators performs as one might expect with the half-linear (ramp) z weighted partition estimator roughly correlating with flat cuts in true radius.




Figure 7-1. A comparison of various ionization radial partition estimators. The side specific estimators are shown in blue and red. In general, hole side collection (S2) has better performance on radial discrimination than electron side (S1) due to oblique propagation and charge trapping effects. However at high radius and high z values the side specific estimators can become better estimators for the specific side they occur on. Various methods to introduce a z-dependence are shown with the quantity derived from using a ramp-function weighting where the bottom side of the detector uses just the hole side estimator seen as having the best behavior in the shared region.

7.3 Choice of Discrimination Quantities

As was briefly discussed in Chapter 4, the ionization signal provides very powerful and useful background discrimination quantities via the ionization signal alone. The HT energy ranges are typically well above the threshold where the ionization signal becomes consistent with noise, making it primarily useful for this type of analysis. The primary discriminators utilized for the HT analysis are the ionization yield (`ytNF`), the ionization z symmetry (`qzpart0F`), ionization radial estimator (`qrpart10F`, `qrpart20F`, or `qrpart0F_zhalf`), with recoil energy given by `precoil1tNF`. The choice of these discrimination quantities is based on the origin of backgrounds in the iZIP detectors




Figure 7-2. The radial estimator behavior for IT3Z1 is shown here. This detector had the ^{210}Pb source implanted on the electron collection side. As can be seen in this figure the S1 estimator (blue) performs better than the S2 estimator (red) due to the presence of a large number of S1 surface events that result in very low total S2 energy. Application of a z-weighted radial quantity does improve bulk radial estimation but performance starts to degrade near the high radius portion of the detector where oblique propagation of electrons causes incomplete collection to become important.

(discussed earlier in Chapter 4) and the generally cleaner separation of background from signal in the ionization channels compared to the phonon signal.¹

While the discrimination quantities taken independently can provide discrimination in a limited sense, there are relationships amongst the discrimination quantities that make a simultaneous combination even more useful. The ionization yield vs recoil energy is one of the standard planes that is often used to show the discrimination of ERs vs NRs, as

¹ That is not to say the phonon quantities will be ignored. A very simple “sandbox” cut can be applied to remove events with physically unrealistic phonon partition values. For example, an event with a phonon z partition value of 2 would indicate something unusual.




Figure 7-3. The radial estimator behavior for IT3Z3 is shown here. This detector had the ^{210}Pb source implanted on the hole collection side of the detector. As is evident here, the hole side only estimator has excellent performance at bulk and shared values. In the shared regions a small feature is present in the average and z-weighted average estimators around a radial partition value of 0.5. This is due to the presence of S2 surface events that result in a noise dominated S1 radial estimator. The ramp z weighted estimator performs well across the whole range, though has a slightly broader bulk peak than the S2 only selector due to S1 surface events.

has been mentioned previously. This plane is unable to separate SEs from NRs due to the reduced ionization yield that arises from incomplete charge collection. In WS data mode the ^{210}Pb source detectors show this very clearly, as is seen in Figure 7-5. In non-source detectors the ^{210}Pb backgrounds are still present but at a substantially reduced rate.

At the detector faces ionization asymmetry gives excellent discrimination ability from surface events, but as one moves further out to high radius, incomplete charge collection causes a diminished ionization signal as well which can cause reduced ionization yield allowing for leakage into the NR band. The radial and z partition quantities can break degeneracies arising from surface events and events at high radius. As seen in Figure 7-6,




Figure 7-4. The true R and true Z in meters is plotted for a small detector Monte Carlo simulation, where the bottom face of the detector occurs at true $Z = 0$. Each subplot has a colormap associated with the titled radial estimator. Thus, in the first subplot titled “S1 Only”, the electron radial partition quantity value determines the color assigned. We can see for the electrons one once starts to enter the middle region of the detector oblique propagation causes incomplete S1 ionization energy collection resulting in reduced radial estimation. As expected the Ramp Z Weighted estimator has the most linear behavior of the combined estimators.

this plane cannot distinguish between between ER and NR events as evidenced by the overlap shown along the $qzpart0F = 0$ axis.

These two planes fail to separate all background types on their own suggesting a combination is necessary. The partition quantities are fairly energy independent, except at very low energy where the signal becomes comparable to the noise. The yield vs radial and yield vs z partition planes are shown in Figures 7-7 and 7-8 respectively. Each of these shows that they also cannot quite capture all the background components on their own, however does indicate that the ionization yield vs z-symmetry plane is quite powerful.




Figure 7-5. Yield as a function of recoil energy provides powerful discrimination of bulk ER vs NR events. These ER events are primarily induced via external γ rays that Compton scatter within the detector. Very rarely an electron can be photo-ejected from the surrounding copper housing and interact in a detector as a surface event β . Calibration data from ^{133}Ba (blue) and ^{210}Pb (red) show that while bulk ER separation from NRs is very high, surface events from ^{210}Pb β s and ^{206}Pb nuclei can leak into the NR band. These surface events are not easily discriminated against in this plane.

These marginal planes that don't quite capture everything suggest that ultimately the most power can arise from consideration of a higher dimension parameter space constructed out of the discrimination quantities as no single quantity nor pair of quantities, captures enough information to discriminate against the regions of the detector where backgrounds are likely to enter. A 3D space constructed out of each of the ionization quantities is the simplest such space one can construct. This ignores some of the energy dependence present in the yield quantity which will tend to cause bleeding of the quantities together at the lowest energies, but at sufficiently high energy this becomes less of a problem. The ionization discrimination space shown in Figure 7-9 underlies the importance of understanding the tails of the background distributions. In this space the bulks of the background distributions are well separated from the bulk of the signal region, however leakage enters the signal region in different fashions for each of the backgrounds, often as a function of radius. At high radius leakage into the signal region occurs from all parameters, however at lower radial values, it becomes evident that the β SEs remain




Figure 7-6. The origin and discrimination of ^{210}Pb derived backgrounds is shown here using IT3Z1, with its ^{210}Pb implanted source on S1. The ^{210}Pb decay chain produces β s that do not penetrate into the bulk of the detector, and can also create ^{206}Pb nuclei that also have reduced yield. The ionization z-partition gives a measure of depth in the detector, whereas the radial partition gives an estimate relative distance towards the sidewall. In the plane formed by these two partition quantities, plotted for the detector with a ^{210}Pb source on S1, surface events on the top face are clearly separated from the bulk (near the origin). Likewise the radial estimator helps to split sidewall events which can consist of ^{210}Pb decay products but also importantly, γ s at high radius. This plane however cannot discriminate between ER and NR events as evidenced by the overlap of ^{133}Ba and ^{252}Cf calibration events at the origin. Here all the ^{252}Cf events are selected to be those that occur in the 3σ NR band based on the yield vs recoil plane.




Figure 7-7. The ionization yield vs radial partition plane for IT3Z1 shows some background separation power. This plane can separate ERs from NRs and is useful for separation of high radius populations of events, however significant surface event leakage still occurs into the signal region (around the origin). In this plane the leakage originates from ^{210}Pb decay products as is typical for surface events. Also revealed in this plane is an indication that at high radius γ leakage into the NR band becomes significant as shown by the high radius “finger” of ^{133}Ba events that droop down into yield values consistent with NRs. This is not terribly surprising since charge trapping and anisotropic propagation of charge carriers into the sidewall result in reduced ionization collection. The data for ^{252}Cf events is selected such that all events occur in the 3σ ionization yield NR band.

a complex source of leakage entering from both the yield and z-partition “directions”.

We also know from Figure 7-5 that the yield quantity will have even better separation of backgrounds at higher energies than at the lower end of the energies plotted here.

Unfortunately I cannot generate a 4D plot that can encapsulate this multidimensional behavior!

7.4 Construction of Background Models

Since there is not a fully mature and verified iZIP detector Monte Carlo accessible for this analysis, the calibration datasets are required to model the backgrounds. Further




Figure 7-8. The ionization yield vs z-symmetry partition plane is the final 2D distribution plane that can be created (ignoring combinations of partition quantities with recoil energy). This data is taken from IT3Z1 (with a ^{210}Pb source on S1) and in this plane one can easily separate ERs from NRs from SEs and which suggests it is one of the more powerful discrimination planes. This plane is unable to discriminate against high radius events which accounts for some of the leakage into the signal region in this plane. At high radius the both the ionization yield and ionization z-symmetry values can be suppressed due to incomplete charge collection. In the former case it is simply due to the direct dependence of yield on ionization energy. In the latter's case it can arise from the difference in the propagation of electrons vs holes. Since electrons have anisotropic propagation in the crystal, a larger fraction will fail to be collected at more bulk like z values than will be the case for holes. This lack of complete collection from electrons will result in a lower S1 ionization energy producing a more symmetric appearing partition value. The ^{252}Cf events are selected such that they occur in the 3σ ionization yield NR band.




Figure 7-9. A 3D space formed from three ionization yield discriminators in the recoil energy range of 8 - 150 keV_{nr} for IT3Z1 (with ^{210}Pb source on S1). Calibration data from ^{252}Cf are selected to occur in the 3σ yield band. In this plane we see clear separation of backgrounds from the signal region (near the origin). The bulk of the background populations are also clearly seen as localized regions in this parameter space and the pathways for leakage can be identified visually. For γ s we can see that leakage towards the signal region is primarily at high radius. For ^{206}Pb leakage primarily occurs via incorrect z-symmetry values at higher radius, though the main population at low radius has some leakage. Other ^{210}Pb induced β s have a more complex leakage from all three ionization quantities. At high radius one can see the result of incomplete charge collection as there tends to be much larger population away from the bulks of the background distributions, underlying the importance of optimizing the radial discriminant.

complicating the background modeling is the fact that the HT analysis cannot generate a background model using the same technique as the LT analysis. Recall that analysis used a region of parameter space outside the blinded analysis region to sample from and assumed any position dependent effects were washed out by the OF resolution at the simulated lower energies. In the HT analysis there is not a statistics rich region at higher energies from which to sample from, nor can position dependent effects simply be ignored. This means that the HT analysis must accurately use the calibration datasets and extrapolate the tails into the signal region. Further complicating matters is the fact that only two detectors have ^{210}Pb source exposure to accurately populate the SE background in the iZIP parameter space, and none has a sidewall calibration source.

7.4.1 ^{210}Pb Induced Events Model

To get around this limitation, the data from the two source detectors used as a proxy for the non-source detectors. IT3Z1 informs the electron collection side and IT3Z3 is used for the hole side collection. An energy smearing procedure can be used to map the source detector information to the other detectors. To obtain a smeared energy E' , simply take the quantity from the source detector, E_s and adjust the offset from the noise distribution mean, then add some random noise to it by sampling from a normal distribution $\mathcal{N}$ with 0 mean and a standard deviation equal to the difference in the source and mapped detector.

$$E' = E_s + \mu - \mu_s + \mathcal{N}(0, \sigma - \sigma_s) \quad (7-2)$$

An important caveat occurs in the situation that the source detector has a larger noise resolution than the detector being mapped to. In that case the model is constructed such that the smeared distribution is no worse than the original source distribution (i.e., one performs the shift to account for the offset in the means, but does not add a small bit of noise). This process is performed for each of the base energy estimators, from which the useful discrimination parameters (yield, partitions) are reconstructed from their

definitions. Each detector has a map for S1 and S2 constructed separately from each of the source detectors.




Figure 7-10. Data from IT3Z3 (blue) is smeared with noise on a non-source detector (red) to map ^{210}Pb behavior. A) The reconstructed ionization yield vs recoil energy plane shows the smeared data exists in the correct location with the mid-yield β and low-yield ^{206}Pb bands evident. B) A histogram of the yield shows agreement between the source and smeared yield quantities with a slight positive bias in the smeared data. The large peak corresponds to ^{210}Pb induced β s and at low-yield, the ^{206}Pb recoils. Courtesy of Todd Doughty.[130]

With mapped source detector data available for each detector, each model needs an appropriate normalization to be meaningful. This is accomplished via the use of the α rates as discussed in Chapter 4.

The sidewall background contributions are a bit harder to quantify. Since ^{210}Pb induced events on the sidewall will tend to not penetrate deep into the detector, they will occur at high radius, where charge trapping effects will dominate. Since these quantities will tend to have reduced yield it is important to understand how the leakage varies as a function of radial estimator, and whether it will be a subdominant background component or not. Unfortunately this is a rather hard background to derive from the available data since we do not have a sidewall source to perform calibration with. To somewhat alleviate this problem, a special calibration run was taken with an iZIP at the UCB test facility

with a 40 Hz ^{210}Pb source on the detector housing. This source was generated by exposing a detector housing to radon source in an analogous fashion to the ^{210}Pb sources at Soudan. Unfortunately several problems exist with this calibration run that make a determination a bit tricky. The detector used² had a noisy outer ionization channel on side 2, reducing the resolution of the radial estimator from that side. Additionally the neutron source used did not generate enough statistics to create a useful NR band for this detector. Instead a NR band from Soudan was used. This appeared to give decent agreement but may not be a fully optimal setup. The `cdmsbats` software is also not compatible with the output of the UCB test facility, meaning not every algorithm used at Soudan was available for this study, in particular the NSOF was not implemented.³ Finally, since this is a test facility there was also a high rate of cosmogenic neutrons bombarding the detector, inflating counts of events in the signal region. Despite these complications the study was able to make a somewhat tenuous claim that the sidewall background is non-negligible but manageable with the leakage on the test facility detector dropping to below 1 event with a radial fiducial cut of about 40%.[\[131\]](#) Further study is needed to ensure the proper normalizations were used and that the neutron background is properly accounted for, as this is a very large systematic uncertainty that significantly reduces the strength of the leakage mentioned above. However, even with this in mind it is not clear how to produce a background model of these events without resorting to some sort of simulation. It is possible that a physical argument can be used to claim that since β s penetrate less into the detector than γ s, if we successfully remove the γ sidewall leakage, the sidewall β s should be removed as well, however this remains to be verified.

² Named G55.

³ But that may be changing shortly.

7.4.2 Gamma Model

The γ background model originated from the ^{133}Ba calibration data. There are a few systematic differences between ^{133}Ba and WS data that must be accounted for. Chief amongst them are the energy spectra differences. In WS mode there tend to be more low energy Compton scatters, whereas in ^{133}Ba data the spectrum is fairly flat, with some spectral features that will be present as well owing to the nature of the source. These are effects that must be taken into consideration when deriving a model for γ leakage. A method to accomplish this is to generate a 1D KDE of the quantity of interest in both datasets and create a set of weights that will transform the ^{133}Ba KDE into the WS KDE. This will allow proper relative weighting of events when sampled from the calibration dataset for use in a leakage based discrimination method or even a BDT method. The overall normalization however is very important as well since it ultimately controls where the final cut boundaries occur. In order to accomplish this one needs estimates for the number of single-detector scatter events in the signal region. If the data were not blinded this would consist of simply counting the number of such events.⁴ Since the data region is blinded such a simple method is not possible. There are a few methods that were used to try and estimate the number of single-scatter γ s in the NR band that rely on information in the sidebands to extrapolate. In order to compensate for the systematic differences in the energy spectra, the normalization estimations can be constructed by an appropriate binning in energy. In the following descriptions I will adopt the following notations:

- The number of single scatters, multiple scatters, or a rate will be denoted by S , M or R respectively
- Superscripts will denote the source of the number, either WS or Ba

⁴ Indeed in CDMS II something like this was performed. That analysis had access to a previous analysis and was able to scale the observed count by the livetime to obtain an estimate. As this is the first iZIP analysis we could not do that.

- Subscripts will denote whether the count originates from observing the sideband (SB) or the signal region in the NR band (NR)

A set of light fiducialization cuts were applied before estimation to avoid contamination from SEs by requiring the z-partition quantity to occur between ± 0.8 (except on the source detectors where the upper or lower limit is set to 0 to avoid the entirety of the source side). An additional set of cuts were applied to require that both side specific radial estimators be no larger than 0.95 and that events occur below the lower edge of the 4σ ER band in **ytNF**. The signal region is defined events in the 3σ NR band as defined by **ytNF** and in the blinding cut region. This last requirement is necessary since the blinding requirements occur with quantities that are slightly different than the standard analysis quantities and as such there are non-trivial overlaps in the different parameter spaces.⁵ This blinding cut requirement is applied to the ^{133}Ba data as well when selecting the NR band region. Finally, the sideband was defined as events that are not in the 3σ **ytNF** NR band and not in the blinding cut NR band either (note that this is not the simple logical inverse of the signal region).

- Method 1a: This method utilizes the rate of singles to multiples in the WS data sideband to extrapolate into the WS signal region. If there is a difference between the rate of singles and multiples in the sideband compared to the signal region then this method will have some systematics. However given that the data is entirely WS in origin this should be minimized. A statistical uncertainty can arise due to a low number of statistics in the WS NR band.

$$S_{NR}^{WS} = \frac{S_{SB}^{WS}}{M_{SB}^{WS}} M_{NR}^{WS} \quad (7-3)$$

- Method 1b: Similar to Method 1a, this method generates a rate of singles to multiples from the ^{133}Ba calibration data, which generally has more statistics than the WS data. However a new source of systematic uncertainty enters as the singles to multiples ratio can behave differently in the ^{133}Ba SB compared to the WS data as a whole. The same limitation applies for the statistical uncertainty in the WS NR

⁵ This is the complication I alluded to earlier when discussing the blinding cut.

region.

$$S_{NR}^{WS} = \frac{S_{SB}^{Ba}}{M_{SB}^{Ba}} M_{NR}^{WS} \quad (7-4)$$

- Method 1c: This method utilizes the rate of singles to multiples in the ^{133}Ba NR band region. This may avoid any systematics due to differences in SB rates of singles to multiples, but may still suffer from any systematic differences in ^{133}Ba data compared to WS data as a whole.

$$S_{NR}^{WS} = \frac{S_{NR}^{Ba}}{M_{NR}^{Ba}} M_{NR}^{WS} \quad (7-5)$$

- Method 1d: This method utilizes only single-scatters thus avoiding any potential systematics in singles to multiples rates. It does introduce a potential source of systematic discrepancy if the rate of singles in each region is appreciably different between the two data types however. Statistical uncertainties are a bit better behaved due to the larger number of statistics in the WS SB compared to the WS NR band region.

$$S_{NR}^{WS} = \frac{S_{NR}^{Ba}}{S_{SB}^{Ba}} S_{SB}^{WS} \quad (7-6)$$

- Method 1e: This method attempts to apply a correction factor α to Method 1d. The correction factor is generated by using replacing contents of Method 1d with multiple scatters to attempt to compute an extrapolated number of WS multiples in the NR band. Since we have access to this number, the correction factor is just the ratio of the observed count to the predicted count. In addition to the systematics that may occur due to the different data types in both the multiples or singles, uncertainty can occur due to behavioral differences between multiples and singles in the SB and NR band regions entirely.

$$S_{NR}^{WS} = \alpha \frac{S_{NR}^{Ba}}{S_{SB}^{Ba}} S_{SB}^{WS} \quad (7-7)$$

Each of the methods above come with their own sources of systematic and statistical uncertainties. Generally Method 1a is expected to have minimal systematic uncertainties since the data types are the same, but Method 1d appears to also suffer from minimal systematic uncertainties and minimizes the statistical uncertainties as well.

Initially the first 3 methods were explored only and they indicated an unexpectedly large number of leakage events should be present (on the order of 10 - 200 depending on the detector). This was a bit alarming given the past CDMS-II values were an order of magnitude or two lower. Since the γ leakage is expected to mainly enter from high radius, and the iZIP is $\sim 3x$ as thick as the older CDMS-II detectors, naively one might expect




Figure 7-11. The yield distributions for ^{133}Ba and WS data are smoothed by generating a KDE and normalized such that the area integrates to unity. An energy range cut on the recoil energy shows only data between 30-150 keV_r. The singles-detector scatters are shown in the thick lines with multiple-detector scatters in lighter line weights. There appears to be general agreement between the yield distributions for the singles between both the WS and ^{133}Ba data with only small fluctuations at the edge of the selection region. The WS data has a severe deficit around the NR band region since the WS data is blinded. The multiples appear to be in less agreement, though this behavior differed from detector to detector much more so than the agreement between singles.

a proportional increase in the leakage all things being equal. The reasons behind this large leakage are discussed in the next section but once it was removed, the estimated gamma leakage dropped to more reasonable values of around 2 ± 1 across the whole energy range. Generally the leakage tends to manifest itself in two different ways from detector to detector as indicated by the ^{133}Ba calibration plots shown in Figures 7-12 and 7-13. The case where the leakage appears scattered in partition values is much more troublesome than the case where it is concentrated at high radius as it will be much harder to eliminate without losing signal efficiency.




Figure 7-12. Shown here are the gamma leakages into the signal region in ^{133}Ba calibration data. Data points with blue circles are single scatters that occur inside the 3σ NR band and the blinding cut NR band. A) Scatter plot of the S2 vs S1 radial quantities indicating the leakage is scattered across all values of radial estimators and not originating from one side in particular nor at high radius. There is a slight bias for events to occur where one quantity is high and the other is low indicating potential SE behavior. B) The same data points in the yield vs z-partition plane. Here the leakage events have a clear bias towards S1 based on the z partition value, but the values are not large enough to be a clear SE.

7.5 Phonon Readout Asymmetry Problem

As mentioned in Section 7.4.2, the initial attempts to construct a gamma model ran into an unexpected problem in the form of a very large number of leakage events. Figure




Figure 7-13. The second type of gamma leakage behavior is shown here in ^{133}Ba calibration data. Blue circles indicate events that occur in the 3σ NR band and the blinding cut NR band. A) The S2 vs S1 radial quantities are plotted indicating most of the leakage is concentrated at high values of both quantities which indicates it will be relatively easy to remove. B) The ionization yield vs z partition value for the same data points. The leakage like events occur primarily at bulk z partition values near 0 indicating they would be impossible to remove in this plane without significant loss of signal. This underlies the importance of proper radial fiducialization.

7-14 shows the initial leakage type events in the ^{133}Ba data as having a markedly different behavior than what is shown in Figure 7-13 for the same detector.

In order to cross check these results I created a plot of the phonon z vs radial partition parameters to verify the leakage occurred in roughly the same location as indicated by the phonon information (bias towards high radius and S1). As seen in Figure 7-15 this was not the case.

Given that the energy ranges selected here were well above the threshold at which the ionization signal becomes unreliable, this discrepancy between phonons and ionization was quite unexpected and ultimately prompted the creation of more phonon partition plots that revealed a serious problem.




Figure 7-14. Initial attempts at extrapolating the gamma leakage from ^{133}Ba calibration data. Blue circles indicate events that occur in the 3σ NR band and the blinding cut NR band. A) The radial S2 vs S1 quantity is shown here with a significant amount of leakage extending far into the bulk region of the detector, including a dense region at high radius. B) The same data shown in the yield vs z partition plane indicate the data has a cluster of leakage at the bulk and some leakage is biased towards S1.

These planes form the triangle plots used for the relative calibration of phonon channels discussed earlier but could indicate if there were specific “hot-spots” on some of the detectors in the x-y plane. Figure 7-16 shows the startling result of this investigation.

This result is startling because the exact same over-density of data appears not only in both WS and ^{133}Ba calibration data, but it appears in the same general region in phonon x-y partition space for every single detector! Clearly something is amiss. The first guess might be that there is contamination on one side of each tower assembly, however this falls because as mentioned in Chapter 2, each detector is rotated relative to its predecessor going down a detector tower. Thus any given phonon channel is not going to face the same side of a tower across different detectors. There is another coincidence that suggests this is a more subtle problem. Recall as well that while detectors are split in the analysis amongst the electron collection side and the hole collection side (or, face to face), for electronic readout they were split into two half detectors that bisect the detector at right angles to the analysis scheme. This difference is shown in Figure 7-17.




Figure 7-15. Example of discrepancy in ionization and phonon partitions during initial attempt at constructing gamma leakage model plotted using ^{133}Ba calibration data. Blue circles represent single scatters inside the 3σ NR band. The two ionization planes indicate a positive z partition (S1) whereas the phonon plane suggests a negative z partition (S2). A) Ionization z partition vs S1 radial shows a clear abundance of S1 leakage at high radius down to fairly low radial partition values. B) The same data plotted as ionization z partition vs S2 radial estimator showing a slightly different behavior but still focused on S1 and with leakage from high radius down to very low radius. C) The same data plotted in phonon z vs radial partition planes. If the data behaved well one would expect to see a concentration of blue circles at high radial partitions at positive z partition values. Instead one sees this cluster occurring at negative z partition values suggesting a S2 phenomena.




Figure 7-16. The phonon x-y partition planes on S1 and S2 for IT2Z2 shown with WS data. Blue circled points are single scatters in the 3σ NR band. There is a very clear and markedly dense concentration of single-scatters and NRs in each plane corresponding to **pc10F** on S1 and in the shared region between **pd20F** and **pb20F** on S2, though typically closer towards **pd20F**.

This unexpected behavior required some investigation to eliminate trivial (but problematic) issues such as data reconstruction bugs or actual position dependence of yield around the detector.⁶ A quick generation of the phonon partition quantities from the basic energy values seemed to produce the same results and a cursory review of the relevant `cdmsbats` code did not reveal any obvious bugs⁷ that could influence the phonon x-y partition values.

To examine the position dependence of the yield, several regions of the x-y partition plot were selected and the yield distributions compared. The sample selection regions are shown in Figure 7-18 with the yield histograms of each region shown in Figure 7-19.

⁶ Position dependence of yield is not unheard of and was a problem with the CDMS-II detectors. Due to fabrication issues there was a y-dependence in the yield. This was corrected for iZIP fabrication.

⁷ Except for the phonon angle partition `pthetaPartOF` which incorrectly computed the quantity as the arctan of x/y and set all 3rd quadrant values to 0. This did not result in the creation of the spot however.




Figure 7-17. The layout of channels for electronic readout purposes is based on which DIB the channel belongs to and splits the detector in half across the crystal. The analysis labels however split the detector by each face. The spot events all occur on DIB 2 based on the phonon x-y partition plots.

As is seen several regions are selected to try and mimic the DIB 2 layout, including an equivalent DIB 1 region, however with the exception of the spot, the yield appears to have similar distributions around the detector. Only in the spot do we have an excess of mid to low yield events.

As there was no position dependence to the yield distributions, save for the specific spot, it became a bit of a mystery as to what the origin of these events were. One possibility is that the particular DIB 2 readouts are, for some reason, uniformly different from DIB 1 readout hardware. Luckily we have access to reverse bias datasets for ^{133}Ba . If this is a phenomena that arises from DIB 2 being fundamentally different, the spot should




Figure 7-18. Various yield sampling regions around the detector. The data selection is defined as the union of events that exist within similar colored boxes on S1 and S2. The anomalous spot on DIB 2 occurs in the green box. A similar DIB 1 spot selection is defined by the red box.




Figure 7-19. The yield distributions from each of the sampled regions in phonon x-y partition space. The histograms are normalized such that the area under each curve is unity. All of the locations around the detector have similar yield behaviors with the exception of the DIB 2 spot which has an abundance of mid to low yield events. A similar behavior was seen in other detectors.

remain in reverse bias mode. But it was rather likely that if it was some unknown detector response related to the iZIP design, the spot would flip to DIB 1 in reverse bias mode. Unfortunately the amount of reverse bias ^{133}Ba data available is not very large compared to the normal bias dataset and as such the signature of the spot was very weakly seen in reverse bias mode. It did appear that there was evidence that the spot did indeed shift in reverse bias mode which suggests that this was a previously unseen “feature” of the iZIP design.

A final cross-check plot that can be performed is to create side-specific yield quantities and plot their behavior as a function of the side-specific radial partition. Recalling that the ionization yield has the maximum of the ionization energy in the numerator, the side-specific yield contains just the side specific ionization energy

$$\text{yt_1} = \frac{\text{qsum10F}}{\text{precoilNF}}, \quad \text{yt_2} = \frac{\text{qsum20F}}{\text{precoilNF}} \quad (7-8)$$

This split proved to be rather useful into identifying the source of the underlying physics behind the spot as it shows that it appears to be a problem with the hole collection. As Figure 7-20 shows the hole side collection has a tail of low-yield events that leak into low radial partition values and as such show up as excessive gamma leakage.

The exact physical origin of the spot is not yet known and represents a potential defect of either the readout hardware or iZIP design in the worst case. It could be that the asymmetric number of inner phonon channels on the electron side and hole side as a function of DIB gives rise to this behavior (DIB 1 contains 2 electron side inner phonon channels and 1 hole side, and the reverse is true for DIB 2) but it is not clear why this should matter, and given what is known about anisotropic propagation of electrons one




Figure 7-20. The side-specific yield distributions as a function of the side radial estimator for ^{133}Ba calibration data. The blue circled points represent single scatters in the 3σ NR band. Purple squares are the subset of single scatter NR events that occur roughly in the DIB 2 spot region. A) The electron side (S1) yield vs radial partition quantity. In this plane we see the majority of NR single scatters and spot events occur at fairly high radius and mid-yield. B) The side specific yield and radial partition for the hole collection side (S2). Unlike the electron specific side, we see here that the single scatter NR events form a tail that leaks to low yield values as radial parameter decreases suggesting a problem with hole collection.

might expect the electron side to suffer from this, not the hole side.⁸ It is also possible that the electric fields along the DIBs are slightly unusual which could result in excessive collection of charge carriers nearby, but without more careful study this is uncertain. This remains an area of investigation for the detector Monte Carlo group and in future results from simulations it will be interesting to look and see if this DIB 2 spot exists, perhaps providing a reason for its existence.

Regardless of the lack of knowledge as to why the spot exists, several features of it allow for it to not seriously impact the HT analysis. The rate of spot events appears to be

⁸ The future iZIP design for SuperCDMS SNOLAB have 6 phonon channels per side, 1 outer radial, 1 central circular, and 4 quadrants of channels that are more symmetric for hole vs electron side. It's unlikely the new design would split the detector in two like was done at Soudan anyways.

constant through time, and the spot appears to behave roughly similar in ^{133}Ba and WS data⁹ meaning the spot could be removed from both datasets without a loss of important background distribution features. Brett Cornell has developed and tuned by hand a spot removal cut [132] that removes the spot based on the intersection of a cut defined by the phonon x , y , and θ distributions. This cut knocks out the spot event regions in the phonon x and y partition planes and appears to have about a 90% signal efficiency meaning we do not suffer a huge loss by removing these very unusual events from consideration.

⁹ Nothing has been said of ^{252}Cf data yet. In ^{252}Cf data the spot is much harder to see due to the abundance of NR events that drown out any spot like signature. There is little reason to suspect it does not exist only in ^{252}Cf mode while existing in the other two primary modes.

CHAPTER 8 CONCLUSIONS

8.1 Summary

In this work I have described the SuperCDMS Soudan experiment and the versatility of the iZIP detectors. With a suitable background model the low-mass WIMP search analysis pioneered the use of a boosted decision tree in a direct detection experiment and set a limit that was much better than a traditional cuts based approach as was taken past searches with CDMS II. This machine learning approach will be attempted again for the high-mass WIMP search analysis and is expected to become a useful technique in future analyses beyond SuperCDMS Soudan. The low-mass analysis also demonstrated that phonon based energy partitions can be used successfully as primary discriminators when ionization energy techniques fail, providing a promising method for performing future lower mass searches and even cross-checking ionization partitions at higher energies. Additional work on the high mass search has shown the iZIPs have excellent background discrimination abilities, able to clearly separate the bulk of different background populations from the signal region. I have also shown however that there is an apparent anomaly due to the presence of excessive low to mid yield events clustered on one side of a detector that requires further investigation and detector simulation. I will conclude this work by outlining a projection for the HT analysis results, describing the importance of a detector simulation, and give a brief overview of what lays ahead for SuperCDMS at SNOLAB.

8.2 SuperCDMS HT Projections

While the HT analysis is still underway a projected range of limits are possible to construct based on a set of conservative and optimistic assumptions about the detector performance. In order to generate a projected limit one must have some estimate for the exposure, which recall depends on the total sensitive mass times the livetime. The total sensitive mass is a function of how much mass exists in the fiducial volume, $\mathcal{F}$ of

the detector. If the fiducial volume removes 50% of the detector's signal efficiency, then the sensitive mass is just 50% the total detector mass.¹ Additionally, no quality cut is 100% efficient at passing signal thus the exposure takes an efficiency hit when factoring in quality cut efficiency, Q_{eff} , which itself can depend upon energy. Additional efficiency hits come from the trigger efficiency (expected to be nearly 1 for most of the HT energy range), and analysis threshold. Thus a quick estimate for exposure may be given as:

$$\mathcal{E} = \mathcal{M} * LT * Q_{eff} * \mathcal{F} \quad (8-1)$$

The most optimistic projection for the analysis operates under the assumption that 13 detectors were used (IT1Z3 which has no S1 bias voltage, and IT5Z1 which has bad noise excluded), about 375 livedays per detector, with quality cut efficiency of $Q_{eff} = 0.95$ and a fiducial volume $\mathcal{F} = 0.65$ that was flat with respect to energy. This results in an exposure of $\mathcal{E} \sim 1800$ kg-days. The energy thresholds are determined by examination of the energy spectra for each detector and chosen to occur where the trigger efficiency is nearly 100%. In the aggressive case this is around 4 keV, and in a more conservative case around 5 keV. For the detectors used in the LT analysis, the conservative case sets the threshold as above 13.1 keV (~ 10 keV_{nr}). The conservative approach keeps only the detectors with fully functional ionization sensors, a smaller fiducial volume of $\mathcal{F} = 0.6$ that itself varies as a function of energy. This latter requirement is more realistic since at low energy the ability to discriminate between signal and background diminishes as was discussed in the LT analysis. The functional form rolled the efficiency off towards 0 at 0 keV_{nr}, with 50% of the total fiducial volume around 8 keV_{nr}. This mainly affects the projected sensitivity at low mass values and does not significantly alter it at higher masses since the fiducial volume becomes flat above about 20 keV_{nr} in this estimate. This yields an exposure

¹ This clearly is assuming that the full detector is sensitive in the limit of 100% fiducialization. This is probably not true very near the sidewalls but to first order is a safe approximation.

estimate of $\mathcal{E} \sim 1100$ kd-days. Using the optimum interval method these exposures can be used to generate projected limits as shown in Figure 8-1 under the assumption of 0 leakage events.

Since the projections were made the course of the analysis has shown that the conservative case is closer what will be realized. Only the functional detectors are being used, and the typical livetimes per detector range from about 270-320 days, a bit lower than estimated 375 days. Of course the final result will depend significantly on how the fiducial volume varies as a function of energy, and what the maximum value takes, as well as whether or not any events leak into the signal region or not.

8.3 Importance of a Detector Monte Carlo for Future Work

It's also worth highlighting the importance of a mature Detector Monte Carlo (DMC) to future analysis efforts. As was seen for the LT analysis, having access to a background model allowed for the use of a machine learning technique to set a very strong limit on low mass WIMP-nucleon interactions. This model was not based on a pure detector simulation but rather, empirically derived by sampling from regions at high energies outside the scope of the analysis. This fundamentally limited both the number of statistics for all backgrounds (the rarest ones most of all) and introduced some small degree of systematic uncertainty due to position dependent effects at higher energy.

Such a model would be very useful for a high mass WIMP search as well since the leakage into the signal region is dominated by the low-probability tails of the background distributions. Unfortunately there is no clean way to produce a data-driven background model like the LT analysis was able to get away with. Instead we must rely on the ^{133}Ba calibration data to stand in as a proxy for WS γ backgrounds, and the ^{210}Pb source detectors as proxies for every other detector's SE event background (both bulk and high radius). These introduce a variety of systematic effects ranging from energy spectra discrepancies between calibration and WS data set to detector-dependent differences between the two source detectors and the non-source detectors.




Figure 8-1. The projected range for the SuperCDMS Soudan HT analysis. A conservative case using fully functional detectors only with a 60% fiducial volume forms the upper edge of the projected reason, whereas a more aggressive limit is formed by excluding only the two worst detectors and maintaining a 65% fiducial volume bounds the lower edge. The difference in exposure is only approximately a factor of two resulting in a similar difference between the two cases in the projected limit. While the limit would not set a world leading exclusion curve, it will represent a world-leading exclusion curve for a Ge target. The closed contours are 90% confidence levels for DAMA/LIBRA (gold) [120] and CDMS-II Si (blue) [44]. The exclusion curves represent 90% confidence levels for CDMSlite Run 2 (red) [76], SuperCDMS LT (dark red) [59], CDMS-EDELWEISS combined (orange) [133], ZEPLIN III (purple) [134], XENON 100 (torquoise) [57], LUX (green)[58], and the SuperCDMS HT projection in the shaded salmon colored region [135]. The yellow area under the dashed curve on the bottom is where coherent neutrino scattering would be expected to contribute 1 background event.[136]

A well developed detector Monte Carlo would allow for the generation of a high number of statistics for all energy ranges and ability to clearly identify which backgrounds leak into the signal region in which ways. This would help alleviate the systematic discrepancies in both types of background model used for SuperCDMS Soudan as well. As both analyses have shown as well, there are situations where detectors perform in fashions that are unexpected. In the LT analysis this arose with the performance of the broken detector IT5Z3 with its shorted outer ionization channel. Some effort to study the detector after unblinding have led to results that do strongly suggest the background model was improperly constructed due to this short [65] leading to the occurrence of the 3 passing events in that analysis. In the HT analysis, besides the fact that the background model will have to rely on extrapolations and rescaling data from calibration datasets, 5 full detectors were rejected from use due to their shorted channels and lack of understanding on how they behave, something which may have been remedied with a simulation. Additionally the unexpected spot is a big mystery that a robust detector simulation could help to address. Based on the behavior of the spot as a function of detector bias, it doesn't seem like the origin is due to an electronic readout issue, but rather some previously unseen deficiency in hole collection. A mature detector simulation could provide some answers to this and many other questions that will probably need to be addressed for the upcoming SuperCDMS SNOLAB project.

Currently work is progressing on writing a condensed matter package for GEANT4[117] called G4CMP that will allow for simulations of solid state detectors such as the iZIP. This work has been ongoing for a few years and is nearing completion, after which it will need to be verified against observed iZIP performance and will be described in detail by an upcoming thesis from Rob Agnese [137]. Currently e^-/h^+ physics is implemented [138] and work is underway on implementing a TES simulation routine. An older DMC does exist (written in MATLAB) however it is somewhat limited and slower than the new one is anticipated to be. One of the biggest issues it suffers from is that it cannot reproduce is

local saturation effects of the TES (much like the LT background model) which as a result causes phonon partitions from simulated data to be not quite realistic, the effects of which was seen in Chapter 6. A kludge, referred to as the bulldozer algorithm, does partially fix this but is physically unrealistic. Kristi Schneck’s thesis contains more information on the MATLAB DMC and its utility in understanding IT5Z3[65] in the LT analysis, and in CDMSlite by revealing issues that arise due to electric field lines at high radius connecting with the detector sidewall, leading to a reduced Luke phonon production for high radius events and in effect broadening expected activation lines. Eventually a fully mature GEANT4 DMC should allow for a much deeper understanding of the iZIP and allow for much improved analysis results.

8.4 SuperCDMS SNOLAB

SuperCDMS has been approved to continue operations with a larger scale experiment at SNOLAB. This is expected to start in the next few years and should allow for dramatic increase in reach compared to Soudan for several reasons. SNOLAB is about 2km below ground which provides a significant reduction to cosmogenic backgrounds over Soudan. Additionally much greater care is being exercised in ensuring the materials used in the design of the cryostat and in the detector housings remains as low-activity as possible and free of background contamination. This latter bit should help reduce the ^{210}Pb background significantly. Currently the design scope for SuperCDMS SNOLAB calls for ~ 50 kg of iZIP detectors. These detectors will be a mixture of Ge and Si much like in CDMS II. Additionally, one whole tower (~ 5 detectors) will be devoted to high voltage CDMSlite style detectors that will probe to much lower WIMP masses and cross-sections by a few orders of magnitude over current limits. R&D efforts towards improving the TES design are underway with the hope that they can be created with much lower T_c values that will improve the intrinsic resolution of the sensor. Great care is also being taken to learn from Soudan and minimize any vibrational resonances that might appear in the fridge design eliminating LF noise and allowing one to have a lower energy threshold.

The iZIP design itself has also been revamped from SuperCDMS Soudan’s 12 channel design. Currently the new iZIP masks maintain two ionization channels, but increases the number of phonon channels to 6 per side as seen in Figure 8-2. This design is symmetric in the number of phonon channels per face per readout side and so may mitigate the unusual asymmetric “spot” discussed earlier, but only if the asymmetry in the Soudan iZIPs plays a role. The new detectors are also larger than the Soudan versions, resulting in more target mass per detector.




Figure 8-2. Current iZIP design for SuperCDMS SNOLAB. This design has an inner phonon channel with 4 middle region channels that can help provide x-y position information. An outer ring shaped phonon channel is also present to provide discrimination of events at high radius.

The SuperCDMS SNOLAB project is expected to be heavily focused on the search for lower mass dark matter candidates as it can provide leading sensitivity results in this mass range that other large scale experiments will not be able to due to their heavy target nuclei mass. Nonetheless, SuperCDMS SNOLAB will provide world leading limits in the high mass region for Si and Ge targets, which may be very important from a complementarity perspective. Recent work has suggested that in an effective field

theory (EFT) framework, the standard spin independent cross-sections do not capture all the ways in which dark matter may interact with the standard model.[65, 126] It is possible that in some EFT frameworks dark matter interactions may be enhanced and simultaneously suppressed for different target nuclei. Since the largest high mass WIMP searches will be dominated by liquid xenon experiments, having a somewhat competitive experiment with other materials should prove useful and provides motivation for SuperCDMS SNOLAB not focusing exclusively on low mass dark matter. These next generation experiments will come tantalizingly close to a so-far unobserved process involving neutrinos referred to as coherent neutrino scattering. This elastic scattering of neutrinos, if observed, will be an important discovery in its own right, but hitting this background before observation of a dark matter signal complicates such searches enormously since coherent neutrino scattering will become a source of background that cannot be eliminated since there is no known way to shield against neutrinos. If this situation arises, directional detection experiments will become even more necessary.

The SuperCDMS SNOLAB project is an exciting prospect for the future and future analysis efforts with it will draw on lessons learned from the SuperCDMS Soudan experiment as described in this thesis. In particular the use of machine learning techniques like the BDT will be allow for better discrimination between background and signal by taking advantage of the multidimensional parameter space present in the iZIP, and the use of the KS test framework will allow for rapid and consistent rejection of statistically dissimilar periods of data. The examination of phonon partition plots has proven to be an excellent cross-check on the ionization signal and should prove to continue to be a useful source of information. With the multitude of large scale dark matter search experiments expected to come online in the next few years, we should soon start to see new regions of parameter space ruled out, potential discovery in the low or high mass regions and perhaps new physics from neutrino scattering! The next decade looks to be an exciting one that




Figure 8-3. Projected limits (dashed lines) in the spin-independent cross-section vs mass plane are shown with current limits (solid). SuperCDMS SNOLAB is expected to reach regions of parameter space at low mass that other experiments cannot and set world leading limits in this region. Also shown is the predicted cross-section in which coherent neutrino scattering will become an unavoidable background.

will hopefully provide answers to the nature of this mysterious and significant component of the universe.

APPENDIX

OPERATING THE DRY DILUTION FRIDGE AT UF

In addition to data analysis efforts, some of the early work performed for SuperCDMS was to operate detector test facility at the University of Florida. This was accomplished via the use of a cryogen-free dilution fridge (often called a dry-dilution fridge) nicknamed “Reptile”. The standard testing process was to assemble a single CDMS II style detector into a tower and to cool down and observe the detector’s noise behavior and TES transition temperatures, T_c , via the use of a prototype Detector Control and Readout Card (DCRC). Unfortunately a set of hardware failures resulted in the necessity of designing and building a new gas handling system (GHS) and led to the end of the detector test facility operations for the remainder of SuperCDMS Soudan. In this appendix I shall briefly describe the dry dilution fridge, how we determined the transition temperatures of detectors, and the rebuild process.

A.1 Dry Dilution Fridge

Reptile is a Vericold cryogen-free dilution fridge with a rated cooling power of $100\mu\text{W}$ at 100 mK and a typical base temperature of around 15 mK. The cooldown time from room temperature to base is approximately 36-42 hours depending on payload. Cryogen-free means that unlike other cryogenic fridges² this one does not require the use of a liquid nitrogen shield and liquid helium bath in order to reach mK temperatures. Instead this fridge uses a powerful pulse tube cooler to bring the fridge to 4K while a $^3\text{He}/^4\text{He}$ mixture circulates in a pre-cool line. This pre-cool line is heat-sunk to the pulse tube cooler at several stages (referred to as the 77K and 4K plates) and incoming mix is chilled and then circulates to the lower stages of the cryogenic fridge (the still, cold plate, and mixing chamber). The cooled mix absorbs heat from these stages, and passes back out into a compressor pump where it is then re-injected to start the process over again. The

² Such as the one at Soudan.

pressure in the pre-cool line is stepped down at certain temperatures in order to reduce the thermal mass and allow continued cooling of the fridge components. Once mixing chamber reaches approximately 10 K, the pre-cool line is evacuated and condensing begins. In this process mix is run through the various upper level stages of the fridge much like in the pre-cool phase, except it then proceeds to enter the mixing chamber itself and the still. Both of these are pumped on allowing evaporative cooling to drive the temperature down over time until eventually liquid helium accumulates. Below a certain temperature that is dependent upon the ratio of ^3He to ^4He (X) a spontaneous phase separation of the isotopes occurs. This separation is what allows for the cooling power of the fridge to occur and is described in detail in [139, 140]. The important takeaway is that as the temperature decreases towards some critical value, the two isotopes will separate into a concentrated phase (rich in ^3He) and a dilute phase (poor in ^3He). The concentrated phase will float on top of the dilute phase, and as the temperature drops, the concentrated phase will approach 100% ^3He . However the dilute phase will asymptotically approach a concentration of about 6.6% ^3He even at 0 K. If one pumps the dilute phase up into a still, that is pumped on, the much lower vapor pressure of ^3He will cause it to be preferentially removed from the still. Since the mixture here is now too poor in ^3He , some of it will enter the dilute phase in the mixing chamber from the concentrated phase on top, attempting to restore the stable concentration of 6.6%. This process requires energy which is taken from the surrounding environment thus providing the cooling power. The ^3He that is pumped away in the still is recirculated and goes through several sets of heat sinks and impedances before re-entering the mixing chamber in the concentrated phase. Often the mixing chamber is constructed with a very specific mixture ratio in mind such that any deviations in this ratio will alter the boundary of the dilute and concentrated phases and can result in sub-optimal cooling power at best, or never reaching a suitable base temperature at worst (this is especially so if the mix is too poor in ^3He). If the mixture is too rich in ^3He the phase boundary will wind up occurring not in the mixing chamber but potentially in

the still line (i.e., there will be so much ^3He that the dilute phase boundary is pushed up into the still line). If the mix is too poor in ^3He , then the phase boundary can potentially occur in the condenser line leading towards the mixing chamber. Both of these result in poor cooling performance since the mixing chamber often has a large surface area inside of it. Additional problems can occur if there is too little mixture. In this case the still may run dry as there is not enough liquid to fill it. This would cause the still to run hot and reduce circulation rate. Conversely if there is too much mixture the still may be overfull. This would introduce a large heat load at the still stage due to the presence of so much liquid helium and increased still exhaust pressures (not good for a turbopump). Each of these problem cases can be diagnosed as follows:

- Too Much Mixture: Still temperature around 1.1 K, high still exhaust pressures and high base temperature. This is an easy case to diagnose. Solution would be to collect some of the mix and not allow it to condense.
- Too Little Mixture: Still temperature is high (potentially higher than 1 K) with very low pressure on still exhaust due to reduced ^3He evaporation rate. If possible, activate still heater and observe if flow rate changes. If it does not change this likely means you have too little mixture in your system.
- ^3He Amounts: A single shot method can be used here to collect ^3He by interrupting circulation after all the mix has been condensed. If the mixing chamber temperature decreases after collecting some ^3He this may indicate the mix is too rich since the mixing chamber will experience significant cooling only once the phase boundary exists within it. If the mixing chamber temperature does not decrease at all this would indicate too little mixture is present. Note that even if the mixture ratio is correct, once single shot mode is initiated the mixing chamber will cool initially. This is due to the reduced heat load that occurs from interrupting the circulation of ^3He back into the mixing chamber. This should occur very shortly after initiation of single shot mode, so if the decrease is observed at some time afterwards that is the indication of too much mixture.

Most of this process for Reptile is controlled by automatic software which adjusts valves on the gas handling system and is generally free from user intervention, but it is still a very wise idea for an operator to periodically check that the cooldown proceeds at an acceptable rate, and that the pressures in various lines remain within acceptable values since the software cannot respond to all contingencies. In Figures [A-1](#) and [A-2](#) schematic

representations of the gas handling system are shown. Pumps are represented by large circles with diagonal lines that taper in the direction of pump flow. Valves are prefixed with a V followed by a number, with pneumatic valves controllable by software depicted as circular and manual valves as squares. Pressure transducers are labeled with an M followed by a number and are depicted as circular. The regulator valve V9 is circular with a smaller circle inside. Special access valves are shown as triangles joined at the vertex, and special access ports with KF flanges are denoted as KF with a number. The mix is stored in the keg, which has a volume of approximately 50 L, and the boundary where the mix leaves the gas handling station and enters the fridge is denoted by a shaded dashed line. The pre-cool and condensing circuits inside the fridge are indicated as impedances and are in reality quite complex.

The temperatures inside the fridge are measured via thermo-resistors that are read out via a 4-wire measurement on a Lakeshore 370 resistance bridge. There are Vericold supplied thermometers at the 77K, 4K, Still, and mixing chamber stages³. Additionally up to 6 extra thermometers exist that were purchased afterwards, referred to as UF 1-6. These sensors were externally calibrated by Lakeshore but had to be re-calibrated from time to time ourselves, with some of the small wires breaking. These were repaired by the in-house electronics shop and required calibration afterwards as well.

In the aftermath of the catastrophic failure, described later, Rob Agnese built a custom data plotter that was more responsive than the Vericold supplied one, and written in National Instrument's LabVIEW language. This VI allowed one to plot any of the sensors as a function of time, or of each other, and output the measurements to a csv file⁴ for easy analysis. Several alarms were built in as well that would email or text in

³ One sensor is inside the mixing chamber and the other is outside it.

⁴ The Vericold data was also logged and is where the VI reads data from. It had the option of converting to a csv as well.




Figure A-1. The pre-cool circuit for Reptile is shown here in a schematic. The red line represents the path mix takes from the storage keg into the pre-cool lines. During the time that V10 and V12 are opened, V2 is closed. The blue path represents the typical circulation path of the mix during pre-cool. The compressor pump (C) starts off at about 2.8 bar of pressure at room temperature and drops periodically, finishing at around 0.2 bar at the end of pre-cooling.




Figure A-2. During the condensing phase of operation, mix is removed from the keg and flows through the red path. This path is connected to a regulator valve (V9) that maintains a specific pressure on sensor M1 (2.2 bar). Otherwise mix circulates along the blue path. It exist the fridge through the still line which passes through a turbo-gate value (TG) into a turbo pump (TP) which is backed by a fore-pump (FP) that sends it into the compressor (C) before it re-enters the fridge. Once the mix is condensed the contents of the blue lines are mainly ^3He .

the event that certain temperatures or pressures met specified conditions that indicated troublesome behavior (e.g., very high pressure on the compressor, or a rate of change in pressure that was too steep.)

A.1.1 Basic Operating Steps

Below I will describe the basic procedure for several common operating modes for the fridge. In practice many of the steps can be done by computer by selection of the appropriate step name, but it is worth knowing how to perform each process manually if need be. In addition valve and pressure checks should always be performed prior to and after each run.

A.1.2 Pump Out

- Begin by ensuring the external vacuum can is secured and on. Verify the external turbopump cart is connected to the copper line but valved off. Ensure the fridge top valve to the pump out line is open.
- In the hallway prepare the large blue Kinney pump for activation by opening the bleed valve and turning on chilled water. Flip on-switch to activate.⁵ Promptly close the bleed valve and open the silver swinging valve. Then very slowly open the black rotary valve to expose the pump to the vacuum space. Verify oil level remains within the indicator and is not overflowing.
- Once the pump is open all the way monitor the pressure on MV.
- Once MV reaches about 0.05 mbar activate the turbopump cart if it is not already on. It is very important not to open the turbo cart to the vacuum lines until you have closed the big blue pump. Typically this involves quickly closing the black rotary valve then moving to quickly open the turbopump cart to the vacuum space so as to avoid a spike in pressure due to outgassing.
- Once the turbopump is set, turn off the blue pump by first closing the silver swing valve, then flipping the switch off. Once off, open the bleed valve until the hissing stops and close it again.

⁵ It will often sound like a horrendous bit of screeching but this is fine as long as it passes quickly.

- At this stage you can turn the blue pump chilled water off, and wait for the pressure on MV to read below 5×10^{-3} mbar.

A.1.3 Pre-cool

- Once sufficient pressure has been reached on MV you can begin pre-cooling. To do this, ensure the chilled water is on for the pulse tube cooler.
- Create a new log file in the control software and ensure it is reading all valves and sensors.
- Verify no open lines exist to the outside world in the gas handling system and that no mix is present in the condensing circuitry by examination of M21 and M22.
- Open the keg valves V28, V29, and V13 and V14 and ensure the reading on M3 is sensible (if it is very low or very high something may be very wrong) before proceeding.
- To engage a pre-cool operation, activate the cryocooler.⁶
- Once the cryocooler is on and chirping, activate the compressor pump and open V7 so that M1 indicates the pressure output by the pump.
- Open V10 and V12 to allow mixture into the pre-cool lines. Let the compressor charge to 2.8 bar. Do not exceed 3.2 bar. If you do overshoot, quickly close V12 and open V8 to dump some back into the keg then close V8 and reopen V12.
- Once M1 is at a satisfactory value, close V12 and V10 and open V1 and V2 to begin circulation.
- Once a specified temperature is reached, reduce the pressure in the pre-cool line. Initially this can be done by closing V1 and V2 and opening V8 and V10 and quickly closing V8 to bleed off mix. Once the pressure in the keg is equal to or greater than the pressure desired in the pre-cool line one will need to ensure the tubogate valve and V31 are open and very quickly toggle V35 to bleed some mix into the still line. From here mix can be returned to the keg via the use of the forepump but beware it should never have more than 1 bar on it.
- Continue in this fashion until the cooldown sensor reads approximately 10 K.

⁶ The pulse tube cooler

A.1.4 Empty Pre-cool Line

- When 10 K (or cooler) is reached, one can begin pumping out the pre-cool line. Since mixture will be condensing we do not want any liquid helium in this part of the plumbing as it will act as a heat load.
- Close V1 and V2 and open V36 and V35. This will let mix start emptying into the still exhaust line. The mix can be either pushed into the keg via V11 and V10, or behind the compressor via V11 and V12 as long as M1 does not exceed 2.8 bar. If M1 gets too high, close V11 and V12 temporarily, and open V8 and V10 to dump the excess into the keg. Close V10 and V8 and reopen V11 and V12.
- Once the pressure on M21 falls below 1 mb, verify V31 is open and activate the chilled water for the still turbopump and then the pump itself to let it work on the pre-cool circuitry until M21 bottoms out and then 5 minutes past.
- At this point shut the turbo pump off and wait for it to cycle down. Shortly after it is powered off close V35 and V36.
- This is a tricky stage as if it takes too long the fridge may warm up enough that condensing will take a very long time. If this happens, reset an appropriate pre-cool stage and try to get the temperature lower again.

A.1.5 Condensing the Mix

- One is now ready to condense. Open V40 and select the regulator valve V9 and turn it to regulate mode. This will ensure the valve opens whenever M1 drops below the programmed pressure.
- During this time ensure the forepump is on. The gas should follow the blue path in Figure A-2 while circulating.
- Eventually enough mix will begin to condense that the regulator valve will remain open and the pressure on M1 will be below the initial threshold amount. At this point wait until the still temperature drops to around 1 K and activate the turbo pump. This will help cool the still stage and set up circulation of the ^3He .
- Proceed to activate the still heater to around 500 - 1000 μW of power to encourage circulation.
- From here on out it is a matter of time until M3 bottoms out and the mixture is fully condensed. If this process takes more than a day there is likely a problem either due to a thermal short somewhere or a blockage in the flow impedance that is preventing proper circulation. A blockage in the impedance will show up as a low M21 pressure. This can be tested by examining the amount of time it takes gas to flow from one side of the impedance to the other.

A.1.6 Collect the Mix

- To collect the mixture one needs to interrupt circulation. First deactivate the turbo pump and the cryocooler. Set the still and chamber heaters to maximum. From here there are several paths to choose.
- One is to close V3, V9 and V40 and open manual V6. This will allow mix to drain from the fridge to the keg. Once the turbo pump has spun down V23 can be opened as well. Ensure the forepump pressure never exceeds 1 bar and the compressor pump never exceeds 2.8 bar.
- An alternative path is to close V7, V12, V9 and V40 and open V10 to allow mix to simply go into the keg, again ensuring the turbopump is off.
- In either case one will likely wind up with some mix left in the still lines that cannot be removed by the forepump due to its maximum pressure. In this case close V10, and open V12 and V7 to allow the compressor pump to pull some of the mix. If M1 becomes greater than M3, close V11 and V12, and open V8 to bleed this excess into the keg. Close V8 and reopen V11 and V12 repeating as necessary.
- If one is attempting to proceed with an expedited warm-up, the introduction (carefully) of dry nitrogen gas into the vacuum space can help. This should only happen above 60 K to avoid condensing the gas. Repeated fills and pump outs can help keep warm transfer gas in the system. Only do this step if there is a reason to warm up fast.
- It is not possible to return all of the mix to the keg since some of it will remain in parts of the plumbing, but generally you wish to minimize the amount of mix left in the fridge itself and in the still exhaust lines.
- Once the mix is collected, turn off heaters, and ensure all valves are closed and all pumps and chilled water lines are off.

A.1.7 Single Shot Mode

- Occasionally a special test mode called a single shot is required. This can be useful for determining the amount of ^3He present in the circulation. This mode starts off just like a mix collection.
- Turn off the cryocooler but keep the turbopump on.
- Close V9 and V12 initially and open V10 so that mix returns to the keg.
- At this point the gas entering the keg is ^3He and it should increase linearly with time. Once M1 bottoms out, close V7 and V3 and continue observing.

- If any spikes in temperature occur right away one will need to restart the compressor and circulate while things settle. It is not recommended to re-open V9 as that would just re-condense the ^3He .
- If no spikes occur and the pressure on M3 is plateauing, a sudden spike on the still temperature could indicate that the ^3He has run out leaving just ^4He behind.
- One should then start to see a jump in the pressure on M3 as ^4He starts to collect at a more rapid rate.⁷
- From here one can either restart circulation and condensing to perform other measurements or proceed to collect mix. In either case be mindful of the turbopump pressure on M21 and M22 and shut the turbo down if necessary.

A.2 Detector Tests

When the fridge reaches base temperature, one can perform various tests of a detector to study its behavior (as well as that of the fridge itself). The detector is typically mounted into a payload similar to Figure A-3 and heatsunk to the mixing chamber stage. Continuity checks are performed during cooldown to ensure no shorts develop prior to reaching base temperature via a breakout box that allows one to record resistances between various parts of the readout circuitry.⁸ Assuming all is well there are several types of measurements that can be performed.

A.2.1 Obtaining Resistances from Noise Spectra and IV Curves

Using a DCRC card the TES and SQUID noise could be characterized. Once the SQUID was properly locked, the noise spectrum could be observed by reading out each channel for a specified length of time. The control software for the DCRC was able to translate the readout from the detector into the noise PSD on the fly and dump an averaged file for further analysis in MATLAB. In principle there are two operation modes that noise spectra can be obtained in: open loop and closed loop mode. In open loop mode the feedback resistor and inductor of the TES circuit (See Figure 3-16) are not

⁷ Assuming the cryocooler is off.

⁸ The QETs are wire-bonded to the DIB pins and may break.




Figure A-3. A detector setup in Reptile. Visible here are several aspects of the fridge and payload assembly. The mixing chamber sits on top the mixing chamber plate and is offset towards the rear. The detector tower assembly contains a detector within the detector housing. On the side a DIB extender connects to the bottom of the tower where a SQUET card is connected. The SQUET will be connected via the stripline, which is heatsunk at multiple stages up to room temperature for detector control and readout.

connected to the output of the amplifier. This allows one to measure the SQUID noise directly. In closed loop mode, the feedback elements are present and we can measure the total noise spectrum, which is comprised of a SQUID noise term, $1/f$ noise and TES Johnson noise. An example of this spectrum is shown in Figure A-4.

The measurement of the noise spectrum can actually be used to work backwards to values for the TES, shunt and inductor values. In Figure A-4 two of the equations for the TES circuit are shown, one for the roll-off frequency and one for the TES noise.

$$i_n = \sqrt{\frac{4k_B (T_{TES}R_{TES} + T_{sh}R_{sh})}{(R_{TES} + R_{sh})^2}} \quad (\text{A-2})$$




Figure A-4. Measured TES noise spectra for phonon channel B on a CDMS II detector in closed loop mode. During this run the fridge was encased in a μ -metal shield with an iron core around the DCRC power supply cord. The red curve is the observed noise spectrum. The black line is a best fit parameterization of $1/f$ noise, TES Johnson noise, a SQUID noise term and a 120 Hz noise peak. All noise terms are in units of $\text{pA}/\sqrt{\text{Hz}}$. The theoretical values for the roll-off frequency f_0 and TES current noise i_n (or N_{TES}) are shown via the green dashed line and used the measured temperatures and assumed given values for the resistors and inductor. Shown in blue dashed line is the actual measured noise and roll-off curve.

$$f_0 = \frac{R_{TES} + R_{sh}}{2\pi L} \quad (\text{A-3})$$

If we make the reasonable assumption that our temperature sensors record accurate values for the TES and shunt resistor temperatures (at the mixing chamber level and the still level respectively) then we have 3 unknowns in these two equations: R_{TES} , R_{sh} and L . One final equation is needed and this arises from taking an IV curve (sometimes called an IbIs measurement). This measurement consists of slowly sweeping a current (I_b) through the TES circuit and observing the output signal voltage V_s . The relationship between V_s and I_s is given via the following equation:

$$I_s = \frac{V_s}{NGR_f} \quad (\text{A-4})$$

Here N is the turn ratio between the TES and feedback inductors, G is the gain of the amplifier, and R_f is the feedback resistance. We can write the relationship between I_s and I_b as follows:

$$I_s = \frac{R_{sh} + R_p}{R_{sh} + R_{TES} + R_p} I_b \quad (\text{A-5})$$

Where R_p is parasitic resistance. This lets us solve for R_{TES} :

$$R_{TES} = \left(\frac{I_b}{I_s} - 1 \right) (R_{sh} + R_p) \quad (\text{A-6})$$

In conjunction with Equations A-2 and A-3, Equation A-6 gives us a third equation without introducing any new unknowns (assuming the parasitic resistance is negligible) since I_b and I_s are measured via the IV curve. These equations can be solved to then determine best fit values for R_{TES} , R_{sh} and L .

A.2.2 Obtaining T_c Values

Besides knowing the resistance of the TES in normal and superconducting mode, we also want to know the critical temperature of the superconducting transition for each of the phonon channels. A good detector will have uniform T_c values across all channels. If there is a gradient then position dependent effects will become very significant and if the




Figure A-5. Observed I_b - I_s curve for a CDMS II detector on phonon channel C at 33.55 mK. Starting from right to left, at very large currents the TES is driven normal and so a linear relationship exists between the output voltage and the input bias current. As the bias current gets smaller, the TES starts to transition towards a superconducting state and enters the biased region. Here the cooling power of the substrate is comparable to the Joule heating in the TES, and it is in this region that one can operate the TES at any point along its transition. As the current drops towards the I_c the TES will quickly cool and go superconducting resulting in an abrupt change on the I_b - I_s curve (a process called snapping). Once the input bias current drops below the critical current, I_c for the TES, it is in a superconducting state (SC) with a very small resistance (non-zero due to parasitic resistance). Again the relationship returns to a linear dependence. This process repeats as the current continues to increase in magnitude but with opposite sign.

gradient is large enough it may not be possible to operate all channels simultaneously! There are two methods one can use to measure the critical temperatures. One method is simply to slowly increase the heat on the mixing chamber and observe the resistance of the phonon channels. Eventually as the mixing chamber (and hence detector) warms up, the detector channels will transition to normal. One can then slowly cool down and observe the transition in the other direction and obtain an average value for the critical temperature.

The other method is to determine the critical temperature by extrapolating where the critical current goes to zero. This method requires performing an IV trace at different temperature intervals as well. At each temperature step one records the critical current from the IV curve (the point where one transitions from the SC region to the bias region) and slowly builds up I_c as a function of T. From here one can perform a best fit to determine at which temperature $I_c = 0$.

A.3 Fridge Problems

While the fridge is quite easy to operate when it works properly, it suffered from many hardware failures that proved to be quite costly. Adding to difficulties is that Vericold ceased to exist when it was purchased by Oxford Instruments, effectively eliminating technical support and requiring all repairs to be handled in-shop. A recurring type of hardware failure were the presence of helium leaks in various soldered joints on the fridge itself. Leaks in the gas handling system are problematic since they result in the loss of mix, and can sometimes produce thermal shorts depending on the temperature stage they occur at. If the mix is condensed the leak will preferentially allow ^3He to leave the system, which is both costly and over time can alter the mix ratio. Other problems included a catastrophic pump failure that led to a loss of mix and the need to rebuild the gas handling system, and electrical issues.




Figure A-6. The critical current for the 4 phonon channels for a CDMS II style detector. Each I_c was measured at specific temperatures by taking an IV curve once the mixing chamber temperature stabilized. A best fit line is generated for each channel and solved for the x-intercept. Each of the 4 best fit line equations are presented in order of channel A - D. The best fit lines give T_c values of 84.05 mK, 85.85 mK, 83.16 mK, and 84.00 mK for channels A - D. In this particular case we can see that Channels C and D had a lot more variation in their critical current values, a sign that the channels were perhaps experiencing excessive noise.

A.3.1 Fridge Leaks

One particular leak occurred in the condensing line input to the fridge, as is shown in Figure A-7. This particular leak was discovered when checking up on the fridge after a catastrophic event I will discuss shortly occurred. Other leaks occurred fairly regularly at various joints similar to the ones shown in Figure A-7 and A-8. A substantial leak (more of a crack) appeared in the precool filter shortly after the leak in Figure A-8 was repaired. This required the machine shop to construct an entirely new filter for us since the leak was so bad. Prior to this, an even larger leak appeared on a heat exchanger on the precool line at the interface of a copper to stainless steel weld.⁹ This also required the construction of a new heat exchanger. Eventually after reflowing several joints and replacement of the heat exchanger and filter, the fridge side leaks appeared to stop, though there are plenty of suspicious looking joints that remain.¹⁰

The general leak testing process consists of pumping out the vacuum space of the fridge and connecting a helium leak checker to the copper vacuum pump out line. Once a suitably small baseline is present (typically $< 5 \times 10^{-10}$ mbar L/s) a small bit of mix is circulated at room temperature through each of the fridge circuits sequentially, usually with the pre-cool lines checked first. If the leak rate is observed to climb in conjunction with a particular circuit having mix inside of it, this is a sure sign that there is a leak somewhere. From here as much mix is collected into the keg as possible, and the rest is pushed into the plumbing leading to the compressor pump and up to V7. The goal here is to empty the fridge and still exhaust lines of as much mix as possible. Once this is done and one is absolutely certain the keg is closed, the external turbopump cart

⁹ The machine shop crew remarked the busted heat exchanger leak was so big it made a great aquarium aerator, a fact confirmed by submerging it in water.

¹⁰ Worryingly, the still has several suspect joints where superfluid will exist which may allow for a superleak to develop at some point down the road.

(which has a decent roughing pump to back it) is connected to one of the external access ports (typically V41) and pumped for a few days to reduce the helium background. This may be accompanied by a few flushes with dry nitrogen gas to help clean out the helium. Afterwards a leak checker is connected and we open up the fridge (i.e., remove the vacuum cans) and methodically spray pipes and joints with small bursts of helium from a compressed gas cylinder with a small nozzle. This process must be done slowly since the signal will not instantly reach the leak checker and often requires 2 people (one in the back hallway to monitor the leak checker, and one to spray). Eventually one will hone in on a suspect joint in this fashion. This is not always the case however and other variations of leak hunting may be explored, such as sealing suspected joints in thick plastic with copper tape and repeating the other mix circulation style test. In cases where the leak is so big that a suitable background cannot be reached in the plumbing, a reverse pressure test can be performed in which the plumbing is pressurized with dry nitrogen and a liquid that forms bubbles readily can be applied to various joints.

A.3.2 Failure of the old Gas Handling System

An unusual situation developed at one point during cooldown run 34. This particular run had a payload of 4 tungsten chips with the goal that we would attempt to measure their transition temperatures via a 4-wire measurement of their resistances. These W chips had known T_c values and could be used to diagnose any parasitic noise issues. Additionally a round of thermometry calibration was due to occur. On Sunday August 22, 2010 I logged in remotely to check the status of the fridge and discovered it was in a very poor state indeed. Instead of being at a base temperature of around 33mK the mixing chamber was around 15 K, with the compressor line at 9 bar briefly with a good chunk of mix back in the keg. Needless to say this was a very bad state of affairs for the fridge. Still exhaust pressures had risen high enough to trip the turbopump off. This 9 bar event (9BE) in retrospect was the start of a persistent decline in fridge performance and much later appears to be related to odd glitches in the Lakeshore bridge behavior. It appears



Figure A-7. Visual appearance of a connection that had a helium leak. The green is a result of acidic flux that was never properly cleaned during fridge installation. Over time it is believed this corroded the connection to the point that a leak appeared. This was repaired by the machine shop re-flowing the solder. This particular line is on the condensing circuit. Courtesy of Durdana Balakishiyeva.

that for some as yet discovered reason, occasionally the Lakeshore Bridge will enter a state where it cannot find the correct resistance range to measure a thermometer. It will then dwell on the sensor cycling through every resistance range it has instead of moving on to a new channel. During this time it winds up dumping a significant amount of power into the thermometer (recall it is a resistor fundamentally) which in turn causes a significant amount of heating. This behavior appears to occur on any of the standard thermometers, however at the upper stages of the fridge the excess power is easily handled. If such a spike occurs in the mixing chamber however, it appears that it can lead to a runaway boil-off of the mix. This had not been observed before because it seems to take more than




Figure A-8. Shown here are the locations of suspected large leaks (around 1×10^{-6} mbarL/s) on the input lines to a small circulation filter on the precool line (heatsunk to the 77K stage). The solution was to remove the section of pipe and have the machine shop replace it with a new copper sleeve and solder it which seemed to fix it. Courtesy of Durdana Balakishiyeva.

90 hours of operation for the spikes to occur in the mixing chamber level (that is, it is an infrequently occurring phenomena.)

After the 9BE it was observed that the compressor pump was making an unusual sound however running proceeded. Eventually it was pointed out that the compressor pump had trouble charging the pre-cool line above 2 bar when it should have been at 3.6 bar. It was eventually found that the compressor pump had been degrading for some time, and more noticeably after the 9BE. This pattern is shown in Figure A-9. Eventually during a post-mortem analysis I determined that the total time from room temperature to 170K¹¹ did not increase significantly if the starting pressure was reduced from 3.6 bar to

¹¹ The first temperature step where the pressure on M1 drops below 3 bar.

2.8 bar.¹² Unfortunately during the compressor pump's deterioration, and most probably as a result of the 9BE, it suffered an internal diaphragm failure that spewed rubber debris inside the GHS!




Figure A-9. The average pressure of the compressor pump as measured at pressure sensor M1 over the course of multiple runs. Notable events are highlighted with vertical bars. A) The initial pressure charge as a function of run number is shown here. The value should be 3.6 bar for this step however we see that from Run 25 onward the pressure started to degrade. After the 9BE the pressure could not rise above 2 bar initially. The sharp rise after Run 50 indicates when the compressor was repaired. B) The average pressure during the second temperature step is shown. The default value for this phase should also be 3.6 bar. More noticeably after the 9BE the compressor pump struggled to keep an elevated pressure during this phase.

Once it was discovered that the compressor pump rubber had entered some of the circulation lines, it was decided to investigate how far the debris made it into the GHS. Depending on where the debris made it, this could result in damage to the valves at best, and a permanent blockage to the flow impedances at worst. It was discovered that the large external charcoal trap prevented the rubber debris from going towards the fridge, however the rubber contaminated the entire GHS from the compressor pump towards the keg and appears to have made it to the keg (so there may be debris inside it) which was

¹² After consulting with the manufacturer of the compressor pump (KNF) it was discovered the pumps should not exceed 3.2 bar, hence the choice to be safe and go to 2.8 bar.



Figure A-10. Shown here are the damaged compressor pump diaphragm and the rubber debris it shed on the compressor head. This type of rubber debris would go on to contaminate the entire GHS.

a surprising situation. It was decided at this point to clean the entire GHS of this rubber debris. The original GHS is shown in Figure [A-11](#). It is was a very slick and compact design, connected with Swagelok compression fittings. However much like Pandora's box, once it was opened it would prove to never be possible to close it. The cleaning process consisted of flushing out each and ever pipe and valve with acetone, then ethanol, followed by some deionized water, and more ethanol after. Dry nitrogen gas was then blown through each segment to remove moisture and volatiles. Small pipe cleaners were also used to attempt to manually dislodge rubber debris.¹³ During the cleaning process the pneumatic valves were examined for damage to their seals and it was discovered that several contained scratches on their metal diaphragms. Replacement diaphragms

¹³ The resulting runoff was an extremely dark suspension of ultra-fine rubber debris that did not settle even after a week of being left alone.

were ordered and replaced, however it was quickly discovered that the valves had to be closed to a very precise torque or the rubber o-rings inside would become damaged and leak. This led to the necessity to replace several pneumatic valves entirely. It was also discovered that the system could not be closed without introducing a tremendous amount of strain in certain locations, and furthermore, several connections were over-tightened both before and after the disassembly process. This resulted in damage to the ferrules used in the compression fitting. Since the ferrules could not be replaced without replacing the component they were attached to this presented a problem. A solution that arose was to epoxy the joints together and seal them that way.¹⁴ Very worrisome was that the primary valves at the keg were also damaged and so slathered in a generous coating of epoxy, but this is not a permanent solution. An electrical short also developed on the control for V1 during this time that required significant effort in creating a map of the wiring control unit before it could be replaced.

Once these issues were complete, the lines were pumped out for approximately a week and the mix circulated through the filters to clean it. Run 51 was begun on Thanksgiving Day (November 24) 2011 nearly 8 months after the cleaning process started. Unfortunately this run did not proceed smoothly. Prior to the start of the run the leak checker being used had a filament burn out, however up until this point there appeared to be no significant leaks in the fridge nor the epoxied GHS so the run was approved to start. The fridge never reached base temperature, instead plateauing at a few hundred mK before commencing a slow rise. Additionally the still exhaust was at a very low pressure. It was thought that a blockage formed in the impedance lines, a not unreasonable hypothesis given that this was the first run after having the entire GHS open

¹⁴ Ultimately this would wind up consuming 2 weeks of time and resulting in the epoxying of nearly every single connection due to the presence of leaks.




Figure A-11. The original gas handling system (aka Pandora's box) is shown here. It was very compact due to the presence of many flyover connections and used compression fittings to connect joints together. The system was under a great deal of strain as it turns out and would prove impossible to put back together entirely.

to atmosphere for cleaning.¹⁵ Eventually the run was ended and upon collection, was very alarmed to discover the keg was reading 0.76 bar instead of the typical 0.98 bar. It was very clear about 1/6 of the mix had been lost to a leak. After obtaining a working leak checker and a temporary keg that we filled with ultra-high purity ^4He , several tests were performed that eventually revealed the compressor pump was the point of failure again. It had developed a new, and very massive leak at the output line on the pump head.

¹⁵ Recall from earlier however too little mix can produce a similar result.

This was also epoxied to eliminate the leak. Further post-mortem analysis indicated that compared to previous runs, the drop in pressure on M1 was significantly higher during cooldown as shown in Figure A-12. During Run 59 it was observed the pressure on M1 was decreasing much too fast, a sign of another leak developing. Combined with the numerous leaks appearing on the GHS, and the sheer absurdity of having to epoxy the entire GHS, it was finally decided we should construct an entirely new GHS from scratch.



Figure A-12. The pressure on the compressor as a function of upper stage temperature. During a normal run the pressure should stay relatively flat during cooldown, with only modest decreases due to the temperature of the gas cooling slightly. When the mixture was lost, this rate is significantly higher than it should be. Run 59 shows the return of this leak.

A.3.3 A New GHS

The design of the new GHS commenced around April of 2012 with a few constraints present. At the suggestion of the post-doc, Durdana Balakishiyeva, it was decided to switch from using compression fittings to variable compression ring (VCR) fittings.

This would allow for a significantly easier ability to swap out damaged components. An undergraduate student, Evan Kasoff, with experience using SolidWorks drew up schematics based on design requirements in order for the machine shop to work with. Even though we were designing this GHS from scratch, several constraints were present. The new GHS had to fit within the rack the old one did, and it had to contain the same pneumatic valve topology as the old GHS since the proprietary control software could not be altered.¹⁶ It was decided early on that it would be best to allow the GHS to take up much more space in the rack for the sake of ease of repair and construction. Additionally a new, more powerful compressor pump was purchased from KNF that had a double-head construction as opposed to the single head of the original. A manifold panel was constructed to allow the compressor to be bypassed if necessary. Additionally the emergency check valves were replaced with values set to lower thresholds on OPV1 (connecting the compressor output to the keg). The old GHS was set at 8 bar¹⁷ and the new one at 5.5 bar. The original design called for a front and rear panel both made of stainless steel. However design tolerances were not taken into account in the SolidWorks schematics and as such during construction it proved very difficult to connect pieces. Eventually Tarek brought a rectangular piece of wood in for us to use as the front panel. This allowed for us to do a mock assembly the front panel and mark on the wood where holes should be drilled when the system was closed, as seen in Figure [A-13](#). After many weeks of work by everyone in the lab and machine shop, the new GHS was finally assembled and mounted into the rack as shown in Figure [A-14](#). Once assembled the GHS was leak tested and several leaks were discovered and were addressed as appropriate by either tightening or replacing a VCR, or having the machine shop weld new parts. Along the way it was discovered that the new pressure transducers were gauge pressure

¹⁶ Someday this will need to be fixed since the machine is over 12 years old at this point.

¹⁷ But apparently was somewhat sluggish to trip given the 9BE.

transducers and not absolute pressure transducers. As the software could not be altered to account for this, and absolute pressure transducers of the type we required did not exist, we had to have a custom adapter created to connect the old transducers, which luckily do not leak. The regulator valve V9 could also not be replaced which posed more of a problem since a leak was discovered on it. Upon opening it appeared that the O-ring had cracked for some reason. An exact replacement was not able to be procured as it was a metric sized ring, but a black #13 BUNA-N type o-ring was put in its place and seemed to fix the leak.



Figure A-13. Picture of myself creating holes in the front panel for connections that will run towards the back panel. This particular hole would allow V7 to be connected to M1 and the compressor pump input. Courtesy of Kristen Zych.

Eventually the GHS itself was deemed leak tight and after a few patches on the quick-connects that hook up to the rear panel were epoxied, the GHS appeared




Figure A-14. Picture of the new gas handling system fully assembled. It is not as compact as the old GHS, but substantially easier to diagnose and repair. A) Front panel view. The new front panel is about 2.5x as tall as the original GHS. B) Side view showing flyover pipes leading to the rear panel which connect to pumps and the fridge. Various pressure sensors and pneumatic control switches are seated in this section.

leak tight.¹⁸ A few weeks were spent purging the GHS and fridge plumbing of air via continuous pumping and a few purges with dry nitrogen gas. Once the system was deemed ready, the external keg of ^4He only was connected to the new GHS via KF1 and a liquid nitrogen cold trap in between. At 2100 on December 11, 2012, nearly 8 months since the start of the new GHS project, and 10 months since the last attempt at cooldown, Run 60 was initiated. The run reached a successful base temperature of approximately 800 mK on the still with just ^4He and no leaks appeared.

A.3.4 Adding More Mix and Determination of New Cooling Power

After the repair of the GHS, the next step was to restore the proper amount of mixture to the system to allow cool down to base temperature. A small sample of the original mix was injected into the system bringing the keg up from 0.76 bar to 0.77 bar. Afterwards about 10 L of ultra-pure ^4He were added to the system in another run to allow for enough mix to start filling the still.¹⁹ With this accomplished the fridge cooled to a base temperature of approximately 35 mK, however did not have much cooling power ($\sim 60 \mu\text{W}$ at 100 mK). At this point three challenges existed. We needed to determine what the current mixture ratio of $^3\text{He}/^4\text{He}$ was, we needed to procure some ^3He , and finally we needed to determine the actual original concentration of mixture. It appears this last detail was never recorded during the fridge initial setup, and with Vericold no longer in existence this number appeared lost to the ages. A single shot mode was run on the ^3He poor mix and seemed to suggest approximately 11.4% of the mix was ^3He after the leak and addition of more ^4He .

Recalling that the phase separation occurs at different temperatures as a function of the concentration, the single shot run cooldown profile was examined to see if there was

¹⁸ These would eventually be replaced by cutting the thin stainless steel tubing and replacing with new ferrules.

¹⁹ This brought the keg pressure up to about 0.98 bar.

any sign that the mixing chamber temperature sensor indicated when the phase separation occurred. In principle it should show up as a slight decrease in the mixing chamber temperature since the formation of the phase boundary should start to alter the cooling power in the mixing chamber. Unfortunately no clear cut pattern was found. A kink in the temperature around 500 mK existed for most runs prior to the loss of mix, which works out to approximately 21% concentration, however this may be coincidence. In the run immediately after the injection of ^4He a kink also exists at this temperature, and while the cooldown profile was different beyond this temperature, one would expect the addition of ^4He to cause the phase boundary to appear at lower temperatures. Unfortunately no single shot appears to have been recorded in the early days of the fridge so the true concentration is unknown. Eventually some ^3He was procured and 2.85 L of ^3He at STP was injected into the system bringing the concentration up to 18.5% where it remains as of this writing. Ultimately the cooling power of the fridge may be the deciding metric for the proper mix ratio

The cooling power of the fridge can be determined by application of heat to the mixing chamber and observing what the resting temperature is as a function of applied heat. This value will depend on the flow rate as well, which is controlled via the still heater. Generally the more power on the still heater, the greater the flow rate. A higher flow rate can bring about a greater cooling power, up until the point the excessive flow of warm ^3He back into the mixing chamber introduces too much of a heat load. In the past measurements have used a still heater power of 1000 μW typically, however here 900 μW were used to to keep the base temperature low. After taking measurements at several chamber heater values, a quadratic fit in temperature is then applied and one can determine the cooling power at 100 mK by use of this function. The results from Run 69 are shown in Figure [A-15](#).




Figure A-15. Plot of the internal mixing chamber temperature sensor as a function of chamber heater power. A best fit quadratic is applied to the data (dashed line) with 95% confidence intervals shown (solid line). The value of the cooling power at $T = 0.1$ K is applied, and the uncertainty in the fit propagated to give an estimated cooling power of $63.1 \pm 0.5 \mu$ W at 100 mK. This is less than previously measured cooling powers of 100μ W at 100 mK. Note that for this data sample the still heater was on at 900μ W.

The flow rate of ^3He and parasitic power can be computed by making use of the following equation from Pobell[140]:

$$P = \frac{dn_3}{dt} (H_{3,d}(T) - H_{3,c}(T)) - P_p = 84 \frac{dn_3}{dt} T^2 - P_p \quad (\text{A-7})$$

Here dn_3/dt is the ^3He flow rate, $H_{3,d}$ is the enthalpy of ^3He in the dilute phase and $H_{3,c}$ is the enthalpy of ^3He in the concentrated phase, the difference being approximately $84T^2 \text{ JK}^{-2}\text{mol}^{-1}$. Making use of this we can pick off the flow rate and parasitic power by

plotting the mixing chamber power as a function of the square of the temperature sensor and fitting a line to the resulting data. This is shown in Figure A-16. The flow rate was observed to be $76.6 \pm 0.4 \mu\text{mol/s}$ with a parasitic power of $1.0 \pm 0.2 \mu\text{W}$. Past values of the circulation rate at $1000 \mu\text{W}$ were on the order of $74 \mu\text{mol/s}$ so this value is not too far out of line. The parasitic power is also consistent with past measurements. The cooling power is not much changed from prior to the addition of the ^3He suggesting that we do indeed still have too little and the phase boundary is not at the optimal location for this mixing chamber. As more ^3He is added care will have to be exercised that the total amount of mixture does not get too large either, or the still could become too full.

A.3.5 Electrical Issues

With most of the cryogenics up and running attention turned towards attempting to produce measurements of T_c again in a W chip setup. Unfortunately the critical current observed in each of the 4 W chips was well below $1 \mu\text{A}$ at base temperature of 25 mK, inconsistent with the expected transition temperatures up around 80 mK. It was eventually discovered that large amounts of parasitic noise were causing warming of the W chips as the outer edges of the mixing chamber plate appear to be at a higher temperature than they are reporting. This suggests that the temperature sensors being used (2 of the UF sensors) required calibration as they were most likely reporting temperatures about 20 mK cooler than in reality. After using a spectrum analyzer, and taking rough capacitance measurements around the fridge table top, it was discovered that there were multiple ground loops present, and that a large portion of the noise was coincident with the cryocooler being on. The ground loop issue is somewhat related to the fashion the electronics rack is connected to the fridge and the lack of a true ground and has yet to be solved. A suggestion by the electronics shop was to use a star-ground pattern in which we connect separate parts on the fridge top to a common point and then (via a heavy grounding braid) connect that common point to a ground. This did appear to help somewhat as base temperature fell approximately 5 mK however the noise in the




Figure A-16. A linearized plot of the mixing chamber heater vs the mixing chamber temperature recording. A best fit line is plotted through the data points (dashed line) with a 95% confidence interval on the fit parameters shown (solid line). Making use of Equation A-7 the flow rate and parasitic power are derived.

W chips remained. This was eventually traced to the inverter for the 3-phase rotary valve used by the pulse tube cooler. The inverter sends a signal both down and upstream in the electronics and when it is on, one can pick up the signal in multiple wall outlets throughout the lab! A high-frequency filter was constructed for this device many years ago by collaborators at Fermilab, however the electronics shop here recommended several modifications. The proposed plan now is to encase the inverter in a metal box with a special long lead filter to help reduce the noise spikes. This will however be work for a future graduate student or post-doc.

The UF dilution fridge has had a troubled history with multiple leaks and hardware failures. However through the many hours spent searching for leaks and attempting to reconstruct the mix ratio, and aiding in construction of the new gas handling system, it has proven to be very informative on the workings of these devices and provided excellent practice for debugging multiple issues. I believe that once the the inverter issue is addressed and a proper grounding scheme implemented, and perhaps suspicious joints reflowed, this fridge can be up and running again for future studies.

LIST OF REFERENCES

- [1] V. Springel, C. S. Frenk, and S. D. M. White, [Nature](#) **440**, 1137 (2006).
- [2] J. A. Frieman, M. S. Turner, and D. Huterer, [Ann. Rev. Astro. Astrophys.](#) **46**, 385 (2008).
- [3] S. Perlmutter, G. Aldering, G. Goldhaber, *et al.*, [Astrophys. J.](#) **517**, 565 (1999).
- [4] A. G. Riess, A. V. Filippenko, *et al.*, [The Astronomical Journal](#) **116**, 1009 (1998).
- [5] R. Caldwell and M. Kamionkowski, [Nature](#) **458**, 587 (2009).
- [6] M. J. Rees, [Phil. Trans. Roy. Soc. Lond.](#) **361**, 2427 (2003).
- [7] R. B. Thakur, *The Cryogenic Dark Matter Search Low Ionization-Threshold Experiment*, [Ph.D. thesis](#), University of Illinois, Urbana-Champaign (2015).
- [8] F. Zwicky, [Astrophys. J.](#) **86**, 217 (1937).
- [9] V. C. Rubin and W. K. Ford, Jr., [Astrophys. J.](#) **159**, 379 (1970).
- [10] V. C. Rubin, N. Thonnard, and W. K. Ford, Jr., [Astrophys. J.](#) **225**, L107 (1978).
- [11] Y. Sofue and V. Rubin, [Annual Review of Astronomy and Astrophysics](#) **39**, 137 (2001).
- [12] A. Klypin, H. Zhao, and R. S. Somerville, [The Astrophysical Journal](#) **573**, 597 (2002).
- [13] M. Tegmark *et al.* (SDSS), [Astrophys. J.](#) **606** (2004), 10.1086/382125.
- [14] A. A. Klypin, S. Trujillo-Gomez, and J. Primack, [The Astrophysical Journal](#) **740**, 102 (2011).
- [15] V. Springel, S. D. M. White, *et al.*, [Nature](#) **435**, 629 (2005).
- [16] Y. Meillier, [Ann. Rev. Astro. Astrophys.](#) **37**, 127 (1999).
- [17] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky, [The Astrophysical Journal Letters](#) **648**, L109 (2006).
- [18] B. Henning and H. Murayama, (2012), [arXiv:1205.6479 \[hep-ph\]](#) .
- [19] J. Beringer, J. F. Arguin, *et al.* (Particle Data Group), [Phys. Rev. D](#) **86**, 010001 (2012).
- [20] W. J. Percival, B. A. Reid, *et al.*, [Mon. Not. Roy. Ast. Soc.](#) **401**, 2148 (2010).
- [21] P. Ade *et al.* (Planck Collaboration), (2015), [arXiv:1502.01589 \[astro-ph.CO\]](#) .
- [22] A. A. Penzias and R. W. Wilson, [Astrophys. J.](#) **142**, 419 (1965).

- [23] G. F. Smoot, C. L. Bennett, *et al.*, [Astrophys. J. **396**, L1 \(1992\)](#).
- [24] G. Hinshaw, D. Larson, *et al.*, [The Astrophysical Journal Supplement Series **208**, 19 \(2013\)](#).
- [25] R. Adam, P. Ade, *et al.* (Planck Collaboration), (2015), [arXiv:1502.01582v2 \[astro-ph.CO\]](#) .
- [26] L. J. Rosenberg and K. A. van Bibber, [Physics Reports **325**, 1 \(2000\)](#).
- [27] G. Jungman, M. Kamionkowski, and K. Griest, [Phys. Rep. **267**, 195 \(1996\)](#).
- [28] R. R. Volkas and K. Petraki, [Int. J. Mod. Phys. A **28**, 1330028 \(2013\)](#).
- [29] J. L. Feng, J. Kumar, D. Marfatia, and D. Sanford, [Phys. Lett. **B703**, 124 \(2011\)](#).
- [30] N. Arkani-Hamed, D. P. Finkbeiner, T. R. Slatyer, and N. Weiner, [Phys. Rev. D **79**, 015014 \(2009\)](#).
- [31] A. Berlin, P. Gratia, D. Hooper, and S. D. McDermott, [Phys. Rev. D **90**, 015032 \(2014\)](#).
- [32] J. L. Feng, [Ann. Rev. Astro. Astrophys. **48**, 495 \(2010\)](#).
- [33] R. D. Peccei and H. R. Quinn, [Phys. Rev. D **16**, 1791 \(1977\)](#).
- [34] L. D. Duffy and K. van Bibber, [New Journal of Physics **11**, 105008 \(2009\)](#).
- [35] S. Asztalos *et al.*, [Nucl. Instrum. Meth. A **656**, 39 \(2011\)](#).
- [36] G. Steigman, B. Dasgupta, and J. F. Beacom, [Phys. Rev. D **86**, 023506 \(2012\)](#).
- [37] S. P. Martin, (2011), [arXiv:9709356v6 \[hep-ph\]](#) .
- [38] G. L. Lane, ed., *Perspectives on Supersymmetry*, Advanced Series on Directions in High Energy Physics, Vol. 18 (World Scientific, 1998).
- [39] M. Maniatis, [Int. J. Mod. Phys. A **25**, 3505 \(2010\)](#).
- [40] C. F. Berger, J. S. Gainer, J. L. Hewett, and T. G. Rizzo, [Journal of High Energy Physics **2009**, 023 \(2009\)](#).
- [41] M. Cahill-Rowley, R. Cotta, A. Drlica-Wagner, S. Funk, J. Hewett, A. Ismail, T. Rizzo, and M. Wood, (2013), [arXiv:1305.6921 \[hep-ph\]](#) .
- [42] R. Bernabei *et al.*, [Eur. Phys. J. C **67**, 39 \(2010\)](#).
- [43] C. E. Aalseth *et al.* (CoGeNT Collaboration), [Phys. Rev. D **88**, 012002 \(2013\)](#).
- [44] R. Agnese *et al.* (CDMS Collaboration), [Phys. Rev. Lett. **111**, 251301 \(2013\)](#).
- [45] G. Angloher *et al.*, [Eur. Phys. J. C **72**, 1971 \(2012\)](#).

- [46] A. Falkowski, J. Ruderman, and T. Volansky, [J. High Energy Phys. **1105**, 106 \(2011\).](#)
- [47] D. E. Kaplan, M. A. Luty, and K. M. Zurek, [Phys. Rev. D **79**, 115016 \(2009\).](#)
- [48] T. Cohen, D. J. Phalen, A. Pierce, and K. M. Zurek, [Phys. Rev. D **82**, 056001 \(2010\).](#)
- [49] K. M. Zurek, (2013), [arXiv:1308.0338](#) .
- [50] G. Aad *et al.*, [The European Physical Journal C **75**, 299 \(2015\), 10.1140/epjc/s10052-015-3517-3.](#)
- [51] L. Goodenough and D. Hooper, ArXiv e-prints (2009), [arXiv:0910.2998 \[hep-ph\]](#) .
- [52] T. Linden, D. Hooper, and F. Yusef-Zadeh, [Astrophys. J. **741**, 95 \(2011\).](#)
- [53] D. Hooper and T. R. Slatyer, ArXiv e-prints (2013), [arXiv:1302.6589 \[astro-ph.HE\]](#) .
- [54] L. Goodenough and D. Hooper, (2009), [arXiv:0910.2998 \[hep-ph\]](#) .
- [55] A. Drlica-Wagner *et al.*, [The Astrophysical Journal Letters **809**, L4 \(2015\).](#)
- [56] Fermi-LAT Collaboration, (2015), [arXiv:1503.02641v2 \[astro-ph.HE\]](#) .
- [57] E. Aprile *et al.* (XENON100 Collaboration), [Phys. Rev. Lett. **109**, 181301 \(2012\).](#)
- [58] D. Akerib *et al.* (LUX Collaboration), [Phys. Rev. Lett. **112**, 091303 \(2014\).](#)
- [59] R. Agnese *et al.* (SuperCDMS Collaboration), [Phys. Rev. Lett. **112**, 241302 \(2014\).](#)
- [60] T. Saab, (2012), [arXiv:1203.2566v1 \[physics.ins-det\]](#) .
- [61] J. Bovy and S. Tremaine, [The Astrophysical Journal **756**, 89 \(2012\).](#)
- [62] F. Nesti and P. Salucci, [Journal of Cosmology and Astroparticle Physics **2013**, 016 \(2013\).](#)
- [63] F. Donato, N. Fornengo, and S. Scopel, [Astroparticle Physics **9**, 247 \(1998\).](#)
- [64] K. Hamaguchi, S. Liew, T. Moroi, and Y. Yamamoto, [Journal of High Energy Physics **2014**, 86 \(2014\), 10.1007/JHEP05\(2014\)086.](#)
- [65] K. Schneck, *Search for Low-Mass Dark Matter with SuperCDMS Soudan and Study of Shorted Electric Field Configurations in CDMS Detectors*, [Ph.D. thesis](#), Stanford University (2015).
- [66] A. Green, [Modern Physics Letters A **27**, 1230004 \(2012\), <http://www.worldscientific.com/doi/pdf/10.1142/S0217732312300042> .](#)
- [67] S. Ahlen *et al.*, [International Journal of Modern Physics A **25**, 1 \(2010\), <http://www.worldscientific.com/doi/pdf/10.1142/S0217751X10048172> .](#)

- [68] J. F. Navarro, C. S. Frenk, and S. D. M. White, *Astrophys. J.* **462**, 563 (1996), [astro-ph/9508025](#) .
- [69] A. A. Dutton and A. V. Macciò, *Monthly Notices of the Royal Astronomical Society* **441**, 3359 (2014), <http://mnras.oxfordjournals.org/content/441/4/3359.full.pdf+html> .
- [70] S. Arrenberg, *Searching for Dark Matter with the Cryogenic Dark Matter Search Experiment*, *Ph.D. thesis*, Zurich U. (2011).
- [71] J. P. Filippini, *A Search for WIMP Dark Matter Using the First Five-Tower Run of the Cryogenic Dark Matter Search*, *Ph.D. thesis*, University of California, Berkeley (2008).
- [72] D. Akerib *et al.* (CDMS Collaboration), *Phys. Rev. D* **72**, 052009 (2005).
- [73] Z. Ahmed, *A Dark-Matter Search Using the Final CDMS II Dataset and A Novel Detector of Surface Radiocontamination*, *Ph.D. thesis*, California Institute of Technology (2012).
- [74] R. Agnese *et al.* (SuperCDMS Collaboration), *Appl. Phys. Lett.* **103**, 164105 (2013).
- [75] R. Agnese *et al.* (SuperCDMS Collaboration), *Phys. Rev. Lett.* **112**, 041302 (2014).
- [76] R. Agnese *et al.*, ArXiv e-prints (2015), [arXiv:1509.02448 \[astro-ph.CO\]](#) .
- [77] B. Cabrera, B. Sadoulet, *et al.*, (2010), [arXiv:1004.1233v1 \[astro-ph.IM\]](#) .
- [78] S. Hertel, “[izip5 capacitances](#),” Internal CDMS Electronic Note [Online] (2012).
- [79] T. Shutt, *A Dark Matter Detector Based on the Simultaneous Measurement of Phonons and Ionization at 20 mK*, *Ph.D. thesis*, University of California, Berkeley (1993).
- [80] B. Welliver, “[Investigating sub-bin interpolation’s effects on amplitude](#),” Internal CDMS Electronic Note [Online] (2011).
- [81] M. Pyle, “[Charge optimum filters with physically consistent time constraints](#),” Internal CDMS Electronic Note [Online] (2012).
- [82] M. Pyle, *Optimizing the design and analysis of cryogenic semiconductor dark matter detectors for maximum sensitivity*, *Ph.D. thesis*, Stanford U. (2012).
- [83] P. Luke, J. Beeman, F. Goulding, S. Labov, and E. Silver, *Nucl. Instrum. Meth. A* **289**, 406 (1990).
- [84] B. Neganov and V. Trofimov, *Otkrytye. Izobreteniya* **146**, 215 (1985).
- [85] K. D. Irwin and G. C. Hilton, *Cryogenic Particle Detection*, edited by C. Enss, Transition-edge sensors, Topics in Applied Physics, Vol. 99 (Springer-Verlag, 2005).

- [86] M. Pyle and B. Serfass, “[Initial attempts to create high quality pt estimators without position correction](#),” Internal CDMS Electronic Note [Online] (2010).
- [87] R. Thakur, “[Notes on optimal filter theory](#),” Internal CDMS Electronic Note [Online] (2012).
- [88] J. Lindhard, V. Nielsen, M. Scharff, and P. Thomsen, *Integral Equations Governing Radiation Effects* (Munksgaard i komm., Copenhagen, Denmark, 1963).
- [89] R. H. Nelson, “[R133 gamma rates](#),” Internal CDMS Electronic Note [Online] (2012).
- [90] E. Lopez Asamar, “[Measurement of the radioactivity levels in supercdms soudan using the gamma-ray spectrum](#),” Internal CDMS Electronic Note [Online] (2015).
- [91] P. Redl, “[R133 gamma rate and leakage estimations](#),” Internal CDMS Electronic Note [Online] (2013).
- [92] T. Doughty, “[Charge chi square cut \(v5-3\)](#),” Internal CDMS Electronic Note [Online] (2013).
- [93] A. J. Anderson, “[Update to phonon chi-square cut \(cpchisq_v53\)](#),” Internal CDMS Electronic Note [Online] (2013).
- [94] S. Golwala, *Exclusion Limits on the WIMP-nucleon Elastic-Scattering Cross Section from the Cryogenic Dark Matter Search*, [Ph.D. thesis](#), University of California, Berkeley (2000).
- [95] Y. Chen, “[Glitch cut update for r133 low threshold analysis](#),” Internal CDMS Electronic Note [Online] (2013).
- [96] K. Schneck, “[R133: Updated low-frequency noise cut](#),” Internal CDMS Electronic Note [Online] (2013).
- [97] T. Doughty, “[R133 prepulse cut](#),” Internal CDMS Electronic Note [Online] (2012).
- [98] S. Upadhyayula, “[cbadseries_133 cut extension for r134](#),” Internal CDMS Electronic Note [Online] (2014).
- [99] R. Agnese, “[Phonon ofres cut \(r133 + r134\)](#),” Internal CDMS Electronic Note [Online] (2013).
- [100] R. Agnese, “[R133 cgoodbasetemp_v53 part 2](#),” Internal CDMS Electronic Note [Online] (2013).
- [101] J. Wilson, “[Trigger burst cut, r133 and r134](#),” Internal CDMS Electronic Note [Online] (2015).
- [102] J. Morales, “[Numi time coincidence cut for run 134](#),” Internal CDMS Electronic Note [Online] (2015).

- [103] H. Qui, “Charge/phonon random χ^2 -rq check for r135 data,” Internal CDMS Electronic Note [Online] (2015).
- [104] B. Welliver, “R134 cpmulttight_blind.v53 update,” Internal CDMS Electronic Note [Online] (2014).
- [105] T.-H. Kim and H. White, *Finance Research Letters* **1**, 56 (2004).
- [106] B. Welliver, “R133 high charge noise bg,” Internal CDMS Electronic Note [Online] (2014).
- [107] F. J. Massey, *Journal of the American Statistical Association* **46**, 68 (1951).
- [108] B. Welliver, “R133 noise correlations,” Internal CDMS Electronic Note [Online] (2012).
- [109] D. Moore, *A Search for Low-Mass Dark Matter with the Cryogenic Dark Matter Search and the Development of Highly Multiplexed Phonon-Mediated Particle Detectors*, *Ph.D. thesis*, California Institute of Technology (2012).
- [110] J. Billard, “Updated nuclear recoil bands for lt analysis,” Internal CDMS Electronic Note [Online] (2013).
- [111] J. Billard, “Nuclear recoil energy scale for lt analysis,” Internal CDMS Electronic Note [Online] (2013).
- [112] A. J. Anderson, *A search for light weakly-interacting massive particles and applications to neutrino physics*, *Ph.D. thesis*, Massachusetts Institute of Technology (2015).
- [113] S. Yellin, (2007), [arXiv:0709.2701](#) .
- [114] A. Anderson, “R133 - effect of 1.3 kev line on post-cf sensitivity,” Internal CDMS Electronic Note [Online] (2013).
- [115] A. Anderson, “Optimization of cpin.v53_lt (v3),” Internal CDMS Electronic Note [Online] (2013).
- [116] A. Anderson, “Discrimination method 2: Simultaneous cut optimization,” Internal CDMS Electronic Note [Online] (2013).
- [117] S. Agostinelli *et al.* (Geant4), *Nucl. Instrum. Methods A* **506**, 250 (2003).
- [118] E. Lopez Asamar, “Estimation of the acceptance of the fiducial volume cuts in the supercdms-soudan low-threshold wimp search,” Internal CDMS Electronic Note [Online] (2014).
- [119] B. Loer, “Trigger efficiencies, october 2012 - july 2013 low threshold ba dataset,” Internal CDMS Electronic Note [Online] (2013).

- [120] C. Savage, G. Gelmini, P. Gondolo, and K. Freese, [JCAP **0904**, 010 \(2009\)](#).
- [121] J. Angle *et al.* (XENON10 Collaboration), [Phys. Rev. Lett. **110**, 249901\(E\) \(2013\)](#).
- [122] Z. Ahmed *et al.* (CDMS Collaboration), [Phys. Rev. Lett. **106**, 131302 \(2011\)](#).
- [123] Z. Ahmed *et al.* (CDMS Collaboration), [Science **327**, 1619 \(2010\)](#).
- [124] E. Armengaud *et al.* (EDELWEISS Collaboration), [Phys. Rev. D **86**, 051701 \(2012\)](#).
- [125] R. Agnese *et al.* ((SuperCDMS Collaboration)), [Phys. Rev. D **91**, 052021 \(2015\)](#).
- [126] K. Schneck, B. Cabrera, *et al.*, [ArXiv e-prints \(2015\)](#), [arXiv:1503.03379v1 \[astro-ph.CO\]](#) .
- [127] T. Doughty, Ph.D. thesis, University of California, Berkeley (2016 (anticipated)).
- [128] B. Cornell, Ph.D. thesis, California Institute of Technology (2016 (anticipated)).
- [129] R. Calkins, [“Nuclear and electron recoil band fits for the ytnf-precoiltnf yield plane,”](#) Internal CDMS Electronic Note [Online] (2015).
- [130] T. Doughty, [“Lead source surface background model,”](#) Internal CDMS Electronic Note [Online] (2015).
- [131] T. Doughty, [“Sidewall source test,”](#) Internal CDMS Electronic Note [Online] (2015).
- [132] B. Cornell, [“Dib2 spot cut,”](#) Internal CDMS Electronic Note [Online] (2015).
- [133] Z. Ahmed, D. S. Akerib, E. Armengaud, and others. (The CDMS and EDELWEISS Collaborations), [Phys. Rev. D **84**, 011102 \(2011\)](#).
- [134] D. Akimov *et al.*, [Physics Letters B **709**, 14 \(2012\)](#).
- [135] T. Doughty, [“Ht analysis projected limits,”](#) Internal CDMS Electronic Note [Online] (2014).
- [136] F. Ruppin, J. Billard, E. Figueroa-Feliciano, and L. Strigari, [Phys. Rev. D **90**, 083510 \(2014\)](#).
- [137] R. Agnese, Ph.D. thesis, University of Florida (2017 (anticipated)).
- [138] R. Agnese *et al.*, [Journal of Low Temperature Physics **176**, 930 \(2014\)](#).
- [139] J. C. Wheatley, [J. Phys. Colloques **31**, C3 \(1970\)](#).
- [140] F. Pobell, in [Matter and Methods at Low Temperatures](#) (Springer Berlin Heidelberg, 1996) pp. 120–156

BIOGRAPHICAL SKETCH

From a young age Brad Welliver has been curious about the world around him and how it functions, in particular the very large and very small. From the age of 15 he set his sights on being a physicist and read every book he could get his hands on that dealt with space, becoming particularly fascinated with quantum physics. Brad was admitted to Florida Institute of Technology in 2003 where he spent the next 4 years taking almost every course he could in the physics department and many in the mathematics department. His first experience with research occurred under Dr. Marc Baarmand where he helped construct and test small photo-calibration boxes for the CMS experiment. Brad graduated cum laude in May of 2007 with a B.S. in Physics and a second B.S. in Astrophysics. In 2008 Brad was accepted for graduate studies in physics at the University of Florida where he joined Dr. Tarek Saab's research lab, and the SuperCDMS collaboration in 2009. Much of Brad's early graduate career was focused on detector testing and characterization as well as other hardware related tasks such as assisting in the rebuild of a gas handling system for the lab's dilution fridge. Starting in December of 2010, he took his first of many trips to the remote town of Soudan, Minnesota where the SuperCDMS experiment was located. While this trip only lasted 5 days, subsequent ones over the course of the next 4 years would add up to about 106 days of total time on site. In 2011 Brad became more active in data analysis efforts in the SuperCDMS collaboration, studying many aspects of the new iZIP detectors and gave his first conference talk at the 2011 Phenomenology Conference in Madison, Wisconsin. Brad gave talks subsequent years at the April APS conferences in Atlanta, Denver, and Savannah. In 2013 Brad was selected by SuperCDMS to give a talk at the international Recontres de Blois in France where he gave an overview of the SuperCDMS experiment as well as the then recently published CDMS II Si results. Brad became much more involved in the data analysis aspect of the collaboration from 2012 onward, joining the low-mass analysis team and worked on numerous aspects of the analysis and PRL publication.

He was selected as a co-team later for the high mass search along with Todd Doughty and proceeded to work on this analysis project discovering an unusual detector response feature of the iZIPs along the way. He successfully defended his dissertation in December of 2015 and received his Ph.D. in Physics from the University of Florida in 2016.