

Detector Simulation and WIMP Search Analysis for the Cryogenic Dark Matter Search Experiment

by

Kevin McCarthy

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2013

© Kevin McCarthy, MMXIII. All rights reserved.

The author hereby grants to MIT permission to reproduce and to
distribute publicly paper and electronic copies of this thesis document
in whole or in part in any medium now known or hereafter created.

Author.....

Department of Physics

February 12, 2013

Certified by.....

Enectali Figueroa-Feliciano

Associate Professor

Thesis Supervisor

Accepted by.....

Chairperson

Chair

This doctoral thesis has been examined by a Committee of the
Department of Physics:

Professor Enectali Figueroa-Feliciano
Thesis Supervisor
Professor of Physics

Professor Jocelyn Monroe
Member, Thesis Committee
Professor of Physics

Professor Steven Nahn
Member, Thesis Committee
Professor of Physics

Detector Simulation and WIMP Search Analysis for the Cryogenic Dark Matter Search Experiment

by

Kevin McCarthy

Submitted to the Department of Physics
on February 12, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Astrophysical and cosmological measurements on the scales of galaxies, galaxy clusters, and the universe indicate that $\sim 85\%$ of the matter in the universe is composed of *dark matter*, made up of non-baryonic particles that interact with cross-sections on the weak scale or lower. Hypothetical Weakly Interacting Massive Particles, or WIMPs, represent a potential solution to the dark matter problem, and naturally arise in certain Standard Model extensions.

The Cryogenic Dark Matter Search (CDMS) collaboration aims to detect the scattering of WIMP particles from nuclei in terrestrial detectors. Germanium and silicon particle detectors are deployed in the Soudan Underground Laboratory in Minnesota. These detectors are instrumented with phonon and ionization sensors, which allows for discrimination against electromagnetic backgrounds, which strike the detector at rates orders of magnitude higher than the expected WIMP signal.

This dissertation presents the development of numerical models of the physics of the CDMS detectors, implemented in a computational package collectively known as the CDMS Detector Monte Carlo (DMC). After substantial validation of the models against data, the DMC is used to investigate potential backgrounds to the next iteration of the CDMS experiment, known as SuperCDMS. Finally, an investigation of using the DMC in a reverse Monte Carlo analysis of WIMP search data is presented.

140.23 kg-days of WIMP search data from the silicon detectors in the CDMSII experiment is also analyzed. The resulting upper limits on the WIMP-nucleon cross-section are higher than those published by other experiments at all WIMP masses, and the lowest limit on the WIMP-nucleon cross-section is $1.07 \cdot 10^{-42} \text{ cm}^2$ at a mass

of $60 \text{ GeV}/c^2$. These results do provide new and interesting constraints at WIMP masses $\leq 40 \text{ GeV}/c^2$ and cross sections from 10^{-42} - 10^{-39} cm^2 , a region in which some WIMP search experiments have claimed evidence for a WIMP signal, which other experiments claim to have ruled out.

Thesis Supervisor: Enectali Figueroa-Feliciano

Title: Associate Professor

Acknowledgments

The work presented in this dissertation benefited enormously from the work and instruction of others, and my time in graduate school has been enriched by an incredible group of coworkers, friends, and family.

First and foremost, I would like to thank my PhD advisor, Prof. Enectali Figueroa. Your enthusiasm for science and deep understanding of experimental physics are inspirational and were fundamental to my success and development as a scientist. You gave me and all your students great freedom in our research while making sure that your office door is always open to us when we need guidance.

Dr. Steven Leman also deserves my sincere gratitude. I count myself lucky to have had you as a mentor while working with you on the DMC and on ever-so-frustrating cryogenic experiments, and I miss Friday evening beers with you and the group.

The MIT CDMS group is an awesome group of people who manage to balance hard work and dedication with friendship and fun outside the lab. Scott, I greatly enjoyed sharing the CDMS experience almost from Day 1 with you. Adam, and Julien, you have all been great to work with, and I've learned a great deal from each of you.

Being a member of the CDMS collaboration has been a fantastic graduate school experience, with unbelievable camaraderie among groups from far-flung institutions. It was outstanding to discuss detector physics with Matt Pyle and Kyle Sundqvist. My thanks to Joseph Kiveni for his immense help during the c58 Si analysis. Zeesh, Dave, and Sunil, thank you for hosting a fun and extremely productive summer at CalTech. I can't name everybody here, but I do extend my gratitude to all of the grad students, postdocs, and professors who took the time to help me throughout

my graduate school career.

Mom, Dad, and Sheila, thank you for your love and for your unending willingness to listen when I needed to vent. I owe my past and future accomplishments to the encouragement you've provided throughout my entire life.

My deepest gratitude and love to Angela: without your sincere support, patience, and understanding, and the occasional motivational fire under my feet, I would never have made it. Thank you.

Contents

1	The Dark Matter Problem	36
1.1	Introduction	36
1.2	Λ CDM	37
1.2.1	The Friedmann-Robertson-Walker metric	39
1.2.2	The Energy Budget of the Universe	41
1.2.3	The Dark Matter Fraction	47
1.3	The Weakly Interacting Massive Particle Hypothesis	55
1.3.1	Properties of a generic WIMP	56
1.3.2	The WIMP as a thermal relic	58
1.3.3	Beyond Standard Model candidates	59
1.3.4	Problems for the WIMP dark matter hypothesis	61
1.4	Conclusions	64
2	Experimental Searches for Dark Matter	66
2.1	Introduction	66
2.2	Production at Colliders	67
2.3	Indirect Dark Matter Detection	70
2.3.1	Gamma rays	70

2.3.2	Cosmic rays	70
2.3.3	Neutrinos	71
2.4	Direct Dark Matter Detection	71
2.4.1	Event Rates - Astrophysics	73
2.4.2	Event Rates - Particle physics	77
2.4.3	Detector Media	78
2.4.4	Detection Channels	84

3 Physics of the CDMS Detectors and the CDMS Detector Monte

Carlo		87
3.1	Introduction	87
3.2	Basic design of the CDMSII and SuperCDMS Detectors	88
3.3	Physics of Electron and Nuclear Recoils	92
3.3.1	Gamma Interactions	92
3.3.2	Electron Interactions	95
3.3.3	Nuclear Recoil Interactions	97
3.3.4	Modeling ERs and NRs in the DMC	97
3.4	Phonon Physics	101
3.4.1	Phonon Basics	102
3.4.2	Phonon Interactions	103
3.4.3	Phonon Creation, Propagation and Interactions in the DMC .	107
3.5	Phonon Sensing: Quasiparticle traps and Transition Edge Sensors . .	109
3.5.1	Quasiparticle traps	109
3.5.2	Transition Edge Sensors	113
3.6	Charge Carrier Physics	117
3.6.1	Charge Carrier Propagation and the effective mass tensor . . .	118

3.6.2	Emission of Phonons by Charge Carriers	121
3.6.3	Charge Measurement	122
3.6.4	Charge Carrier Physics in the DMC	125
4	Validation and Tuning of the CDMS DMC	133
4.1	Phonon absorption and losses in the mZIP-style detectors	134
4.1.1	The mZIP detector	134
4.1.2	mZIP Experiment, Data Selection and Processing	135
4.1.3	Simulations	140
4.1.4	Results	141
4.2	Phonon-quasiparticle downconversion, anharmonic decay, and TES critical temperature in the iZIP3	146
4.2.1	The iZIP3 detector	146
4.2.2	iZIP3 Experiment, Data Selection and Processing	148
4.2.3	Simulations	151
4.2.4	Results	152
4.3	Phonon-QP downconversion and diffusion in the iZIP4	164
4.3.1	The iZIP4 detector	164
4.3.2	iZIP4 Experiment, Data Selection and Processing	165
4.3.3	Simulations	168
4.4	Performance of current DMC mZIP, iZIP3, and iZIP4 models	172
4.5	Validation of Charge Propagation	183
4.5.1	Basic Charge Propagation Validations	183
4.5.2	Charge Carrier-Surface Interactions	186
5	MISS events	195

5.1	Anatomy of a MISS event	196
5.2	MISS rate estimation	200
5.2.1	Back of the envelope expectation	202
5.2.2	Estimation from GEANT4	203
5.2.3	Raw MISS event rate	206
5.2.4	Raw Event Rate from SNOLAB simulations	208
5.3	MISS events in the signal region	210
5.3.1	Back of the envelope estimate	211
5.3.2	Estimation using the DMC	212
5.3.3	Expected MISS rate in the signal region	222
5.4	Phonon-based MISS discrimination	224
5.5	Conclusion	231
6	Phonon pulse matching	235
6.1	Introduction	235
6.2	Data Selection	236
6.3	Pulse Matching Method	237
6.4	Side-summed pulses	238
6.4.1	Example pulse matches	240
6.5	Matching all 8 phonon pulses	244
6.5.1	Example pulse matches	244
6.6	Future work	250
7	WIMP Search Analysis using the CDMSII silicon detectors	251
7.1	Introduction	251
7.2	Description of the Experimental Setup	253

7.2.1	Shielding	253
7.2.2	Cryogenics	254
7.2.3	Cold hardware	256
7.2.4	Warm Electronics	263
7.2.5	DAQ, Data Storage, and Data Processing	265
7.3	Data processing and calibration	266
7.3.1	Data processing	267
7.3.2	Energy Calibration	270
7.3.3	Position and Energy Dependent calibration	274
7.4	Blinding	286
7.5	Data Quality	288
7.5.1	Detector Performance Cuts	288
7.5.2	Total Exposure	300
7.5.3	Event-by-Event Cuts	301
7.6	Physics Cuts	308
7.6.1	The ionization yield bands	309
7.6.2	Charge Fiducial Volume Cut	310
7.6.3	Scintillator Veto Cut	312
7.6.4	Single-Detector Scatter	313
7.6.5	Energy thresholds	315
7.7	Surface Event Rejection Cut	316
7.7.1	Calibration Sample	317
7.7.2	The χ^2 cut	317
7.7.3	Nuclear Recoil Consistency Cut	321
7.7.4	Wimp Search/Calibration Systematics	321
7.7.5	Leakage Estimation	321

7.7.6	Spectrum-Averaged Exposure	326
7.7.7	Optimization of the Rejection Cut	330
7.7.8	Simultaneous Optimization of the Nuclear Recoil Bands . . .	331
7.7.9	Final rejection cut	333
7.8	Conclusion	334
8	WIMP Search Results from the CDMSII silicon detectors	339
8.1	Introduction	339
8.2	Pre-unblinding	340
8.3	Final Leakage Estimates	342
8.3.1	Surface Event Leakage	342
8.3.2	Bulk Gamma Leakage	349
8.3.3	Cosmogenic neutrons	349
8.3.4	Radiogenic neutrons	351
8.4	Unblinding	352
8.5	Checking the events	362
8.5.1	Basic checks	362
8.6	Constraints on WIMP parameter space	364
8.6.1	Spin-Independent constraints	364
8.6.2	Spin-dependent constraints	366

List of Figures

1-1	Energy budget of the universe: baryonic matter, dark matter, and dark energy. Figure from the NASA/WMAP Science team	38
1-2	Angular power spectrum of CMB anisotropies. The fit corresponds to the best-fit Λ CDM parameters. Figure from [10]	43
1-3	Correlation function times the square of the separation distance vs. the comoving separation distance. The acoustic feature is the spike at $\approx 100h^{-1}$ Mpc. The colored lines show the expected curves for models with $\Omega_m h^2 = 0.12, 0.13$, and 0.14 in green, red, and blue, respectively, and a purely CDM model ($\Omega_b = 0$, so no acoustic peak))with $\Omega_m h^2 = 0.105$ in magenta. Figure from [12]	44
1-4	(a) Magnitude vs. redshift of Type Ia supernovae from the Calan/Tololo Supernova Search and the Supernova Cosmology Project in Panel, and (b) residuals of the measured data with respect to the expected curve for a $\Omega_m = 0.28, \Omega_\Lambda = 0.72$ universe. Expected curves for various values of Ω_m, Ω_Λ in dashed lines for flat cosmologies and solid lines for $\Omega_\Lambda = 0$ cosmologies; the exact values of Ω_m, Ω_Λ are indicated outside the upper right corner of the figure. Figure from [13] and contains data from [14]	46

1-5	BBN predictions of the relative abundances of light elements (relative to that of hydrogen) versus the baryon density of the universe, with measurements. The widths of the curves indicate the uncertainties in the predictions. The vertical extent of the boxes indicate measured elemental abundances, and the horizontal extent is set by the range in which measurement and predictions are in agreement. The D and ^3He components are in agreement with a baryon density of the universe, Ω_b , ≈ 0.04 . The heavier ^4He and ^7Li indicate lower values, and the authors mention poorly understood systematic errors in these measurements and in relating these measurements to the primordial populations. This figure compiles measurements and predictions from a number of researchers, please see [16].	49
1-6	Set of galactic rotation curves. Note that the velocity of stars does not decrease at high radius, but rather asymptotes to a constant, indicating an interior mass function $M(r) \propto r$. As the luminous interior mass grows much more slowly than this, these curves imply contributions from a significant dark matter component [21].	51
1-7	The black squares show the measured galactic rotation curve of NGC 3198 with a fit. The contributions from the measured luminous component of the galaxy and the proposed dark matter component are also shown in the lines labeled ‘disk’ and ‘halo’, respectively. [16]. . .	52

1-8	Left: Optical image of the Bullet Cluster 1E 0657-558 from the Magellan observatory. The white bar indicates 200 kpc at the distance to the cluster. Right: X-ray image of the Bullet Cluster from the Chandra X-ray Observatory. The contour lines illustrate the lensing parameter κ , with the outer contour corresponding to $\kappa = 0.16$ and interior lines increasing in steps of 0.07. κ is directly proportional to the surface mass density of the gravitational lens, indicating that the majority of the mass is concentrated away from the regions of high plasma density [23].	54
1-9	Illustration of some theoretically-motivated dark matter particle candidates. The y-axis illustrates the interaction cross-section with normal matter, and the x-axis indicates the expected mass range of the particle. Figure from [25].	57
1-10	Illustration of freezeout for a 100 GeV dark matter particle. The y-axis indicates the comoving number density Y (left) and the resulting thermal relic density Ω_{DM} (right). The x-axis shows the temperature (bottom) and the corresponding time (top). The black line traces the history of a 100 GeV particle with the correct relic density, and the shaded regions around this line vary the dark matter particle's annihilation cross-section by factors of 10, 10^2 , 10^3 . The dashed line indicates the density profile of a particle in thermal equilibrium. Figure from [26].	60
2-1	Inferred 90% CL ATLAS limits on the spin-independent WIMP-nucleon cross-section vs WIMP mass. Figure from [43]	69

2-2	Energy-integrated WIMP interaction rate vs. the low energy threshold of a dark matter detector for Xe and Ge, assuming perfect detector efficiency, a WIMP mass of 70 GeV, and a spin-independent cross-section $\sigma_{SI} = 10^{-44} \text{ cm}^2$. Figure from [52]	76
2-3	Recent limits and regions of interest in the WIMP mass and cross-section plane from a variety of direct dark matter search experiments. Figure from [70]	86
3-2	Yield-based discrimination of electron recoils and nuclear recoils: Bulk electron recoils are shown in red, in a band around a yield of 1, and bulk nuclear recoils are shown as blue circles in a reduced, energy-dependent yield band. Both bands are indicated in black solid lines. However, surface electron recoils, shown as black crosses, exhibit reduced yield and can contaminate the WIMP signal region; the surface-event “band” is shown as a black dash-dotted line. The blue and magenta lines represent recoil energy and total charge thresholds, respectively.	91
3-3	Gamma photoelectric, Compton, and pair production cross-sections vs. gamma energy in Ge and Si. Fig. adapted from [71]	94
3-4	Electron range versus energy in germanium and silicon. Fig. adapted from [71]	96
3-5	Predicted and measured yield vs. energy curves for nuclear recoils in Si and Ge detectors.	100
3-6	Illustrative R vs. T curve for a thin film TES with a transition temperature of 100 mK, transition width of 1 mK, and normal resistance R_n	113

3-8	Number of collected charge carriers (triangles, in thousands) and total phonon signal (circles) from 59.5 keV photons in germanium vs. electrode bias. Devices were 1 cm thick crystals, so the x-axis can also be read as applied field in V/cm. The line represents the expected total phonon signal including Neganov-Luke phonons. Figure from [104]. .	124
4-2	Circuit used to bias and read TES channels in G9F, UCB Run 362 .	138
4-3	Inner and outer X-Y partition plots from DMC simulations of the 6-channel mZIP. The sequence of figures illustrates the effects of varying the phonon transmission coefficient at the Ge-Al interface on the instrumented surfaces. Six transmission coefficients, 0.01, 0.03, 0.1, 0.15, 0.2, and 0.3, are presented, and the coefficient of phonon absorption on the non-instrumented sidewall of the detector is held constant at 0.001	143
4-4	Inner and outer X-Y partition plots from DMC simulations of the 6-channel mZIP. The sequence of figures illustrates the effects of varying the phonon absorption coefficient at the non-instrumented germanium sidewall. Six absorption coefficients, 0.001, 0.003, 0.01, 0.03, 0.1, and 0.3, are presented, and the coefficient of phonon absorption at the instrumented Ge-Al interface is held constant at 0.03	144
4-5	Inner and outer X-Y partition plots from DMC simulations the 6-channel mZIP and Run 362 data. The best matching simulations are shown, and the parameters of these best matching simulations were $(A_{surf}, A_{side}) = (0.1, 0.03)$, $(0.15, 0.01)$, and $(0.15, 0.03)$. In the bottom right corner, the data is shown.	145

4-7	Illustration of the phonon pulse calibration process, aligning the late-time tails of the phonon pulses. Outer channel pulses are scaled by a factor of four to account for the difference in total area	150
4-8	Illustration of the time to position independence. The red line indicates the point at which all pulses are within 10% of the mean pulse, indicating that the phonon signal no longer contains position information.	153
4-9	Histograms of the fraction of energy in the eight phonon channels from data (black solid line) and four different simulation runs (red dashed line). The smaller outer channels each collect $\sim 5\%$ of the energy, while the larger inner channels are distributed around 20%. Increasing $\mathcal{L}$ and increasing γ_A both increase the degree of partitioning (the breadth of the histogram), as expected.	154
4-10	Decay time histograms from data (black solid line) and four different simulation runs (red dashed line). Increasing $\mathcal{L}$ from 4 to 8 has little effect, as most phonons at late times have downconverted to $\sim 3E_{gap}$, but reducing $\mathcal{L}$ to 2 lengthens the decay time. Reducing γ_A further increases the decay time of the pulses at $\mathcal{L} = 2E_{gap}$, though at higher $\mathcal{L}$ values this effect disappears.	155
4-11	Histograms of the slowest 10%-40% phonon pulse risetimes at the three T_c values tested: 50, 55, and 60 mK. Data is shown in the black solid line, and DMC results are presented as a red dashed line. It is clear that the pulses rise more rapidly as T_c increases. 55 mK provides the best match, though it a value somewhere between 50 and 55 mK would likely improve the agreement.	156

4-12	Histograms of the fastest 10%-40% phonon pulse risetimes in data (black solid line), and DMC (red dashed line). The mild variation among the extreme cases presented here (the lowest and highest values of both $\mathcal{L}$ and γ_A are presented) indicates that TES dynamics probably dominate the behavior of this parameter.	158
4-13	Histograms of the time until phonon signals become position-independent (i.e., when all 8 phonon traces converge) in data (black solid line), and DMC (red dashed line). The observed position independence time increases with increasing γ_A and decreasing T_c , as expected, though the predictions from the best matching overall parameter set (upper left) are in disagreement with the data by $\sim 10\text{-}20\%$	159
4-14	Histograms of the peak phonon power and the in data (black solid line), and DMC (red dashed line). The mild variation among the extreme cases presented here (the lowest and highest values of both $\mathcal{L}$ and γ_A are presented) indicates that TES dynamics probably dominate the behavior of this parameter.	160
4-15	Sum of squared residuals between measured and simulated histograms for various pulse shape quantities; darker colors indicate better agreement. Each box bounded by red lines presents a single variation of T_c , $\mathcal{L}$, and anharmonic decay constant, and the 6 individual rectangles within each box correspond to the 6 pulse shape parameters under investigation, as indicated in the key below the figure. There are 75 DMC variations presented here, and the squared residuals for each pulse shape parameter are normalized such that the sum over the 75 variations is 1.	162

4-16	Overall match between data and simulation, measured by combining the sum of squared residuals of all 6 pulse shape parameters presented in Fig. 4-15.	163
4-19	Rate of interaction between phonons and quasiparticles versus temperature. The different lines correspond to different phonon energies. The important consideration for the DMC is to model $\tau_B(E)$, which is the rate at which phonons break Cooper pairs into quasiparticles. .	169
4-20	χ^2 between measured and simulated phonon pulse timing parameter histograms (10%, 40% 70%, and 100% points on the rising edge of the pulse, and the 80% point on the falling edge) versus time for quasiparticle diffusion from the Al fin to the TES. A broad minimum is evident between 15 and 30 μ s.	171
4-21	Comparison of iZIP4 phonon parameters from data (left) and simulation (right)	174
4-22	Comparison of iZIP4 phonon parameters from data (left) and simulation (right)	175
4-23	Comparison of iZIP3 phonon parameters from data (left) and simulation (right)	178
4-24	Comparison of iZIP3 phonon parameters from data (left) and simulation (right)	179
4-25	Comparison of mZIP phonon parameters from data (left) and simulation (right) (except bottom hists, with data and sim on the same plot)	182
4-26	Fractional error in Luke phonon energy produced by DMC vs. applied drift field. The expected Luke phonon energy is simply the number of charge carriers times the applied potential (drift field * 2.54 cm). . . .	184

4-27	Fractional error in total phonon energy produced by DMC vs. applied drift field. The expected total phonon energy is defined as the recoil energy plus the expected Luke phonon energy.	184
4-28	Electron drift velocities vs. field. DMC is the solid blue line. Measured data is from [107]. Comparison theoretical predictions are from [108, 109, 110]. Whether a detector is n- or p-type does not affect the DMC's charge propagation algorithms in its current iteration, which clearly represents a direction for future work.	185
4-29	Hole drift velocities vs. field. Measured data is from [107]. Comparison theoretical predictions are from [108, 109]. Whether a detector is n- or p-type does not affect the DMC's charge propagation algorithms in its current iteration, which clearly represents a direction for future work.	185
4-30	Validation of simulated oblique electron propagation. The blue dots indicate the final electron positions at $z = 0$, while the red line outlines the expected electron propagation 'cone' at $.825 \text{ cm} (1.27 \text{ cm} * \tan(33 \text{ degrees}))$. The population of events outside of the cone are attributed to stochastic diffusion, while those inside the cone are attributed to both diffusion and intervalley transitions.	187
4-31	Histogram of electron propagation angle relative to the applied field in the Z direction. The expected propagation angle of 33 degrees is clearly visible in the peak. Angles less than 33 degrees are a consequence of intervalley transitions and electron diffusion.	188
4-32	Yield of bulk ERs (blue) and surface events (black) in data (left) and DMC (right). The DMC simulation also includes bulk nuclear recoils (green).	191

4-33	χ^2 goodness-of-fit between the surface event yield band means vs. energy in data and simulation in electron-collecting mode (left) and hole-collecting mode (right). Each filled rectangle corresponds to a single DMC parameter set of initial carrier kinetic energies (y-axis) and surface energy barriers (x-axis). The barrier energies are defined as fractions of the initial energy; when the initial kinetic energy is 0, the barrier heights are 3, 11, 25, 45, 70, and 101 meV. Darker colors indicate better matches, and the color scale is logarithmic to increase contrast.	192
4-34	Means and widths of the fitted surface event yield bands in hole-collecting (ionization and phonon rails biased to -2V and 0V, respectively) and electron-collecting mode (+2V/0V) in 2 keV energy bins from 10-80 keV. The results of fits to the data are plotted in blue, and the DMC is shown in red.	194
5-2	Measured ionization yield as a function of the ratio of the energies of the two scatters in a MISS event. The yield reaches a minimum at 0.5 when the two scatters are of exactly equal energy. Events between the two dashed vertical lines would pass the ionization symmetry cut and be deemed bulk events. This figure shows that MISS events that pass the symmetry cut will naturally exhibit event yields near the upper edge of the nuclear recoil band at approximately 0.4. If the ionization from one or both of the scatters is mis-measured for any reason, these events could populate the WIMP signal region.	201

5-3	Background gamma spectrum in the CDMSII experiment and co-added simulated spectra. The data is shown in the solid blue line, and the red dashed line represents the simulated spectra, co-added with the weights shown in Table 5.1	205
5-4	Background gamma spectrum in a GEANT4 simulation of the SNO-LAB setup (solid blue) performed by Mark Pepin and the spherical source G4 simulation (red dashed) with 76.2 mm diameter x 25.4 mm thick detector. The SNOLAB sim utilized expected contamination levels in the shielding and support materials in the experiment, while the spherical sim used the spectrum described in Table 5.1	209
5-7	Charge symmetry cut used in the MISS analysis. The X and Y axes are the total ionization energy measured on Side 1 and Side 2, respectively. The dense band of events at 45° demonstrates that the MISS candidate sample contains a large fraction of ‘normal’ bulk events. Blue events pass the symmetry cut, and red events fail.	217
5-8	Histogram of MISS event yields from events passing the ionization fiducial volume cuts ($Q_{in,AND} \& Q_{sym}$). The MISS candidate sample contains large fraction of true bulk events with yields around 1. The expected secondary population of well-collected true MISS events around 0.5 in yield is also evident. The tail of yields below the secondary peak arises due to MISS events with incomplete ionization collection in one or both of the surface scatters.	220

5-9	Nuclear recoil band cut applied to MISS candidates passing ionization fiducial volume cuts; the X axis is the event's recoil energy, and the Y axis is the measured yield. Cuts placed at 2 and 3 σ are visible as the solid blue and dashed black lines, respectively. Red events fail both cuts, black events pass the 3 σ cut, and blue events pass both the 2 and 3 σ NR band cut.	221
5-11	10-70% risetimes of the side-summed phonon pulses for bulk ERs (blue), bulk NRs (green), and MISS background events.	227
5-12	Sum (x-axis) and absolute difference (y-axis) of the 10-70% risetimes of the 2 side-summed phonon pulses for bulk ERs (blue), bulk NRs (green), and MISS background events. A cut on the sum of the two risetimes is placed at 32.5 μ s and indicated by the red dashed line. This cut removes $90^{+3.8}_{-3.1}\%$ of the MISS events while retaining $98.5^{+0.7}_{-1.0}\%$ of the bulk NRs (quoted errors are binomial statistical errors).	228
5-13	Sum of the 10-70% risetimes (x-axis) vs event energy (y-axis) of the 2 side-summed phonon pulses for bulk ERs (blue) and bulk NRs (green) in g48 R400 data.	229
5-14	Distance between event x-y partition positions on side 1 and side 2 of the detector.	232
6-1	Results of matching side-summed phonon pulses from electron and nuclear recoils from data to electron and nuclear recoils from DMC. Electron recoils from data are shown in blue, nuclear recoils are shown in green.	239
6-2	Recoil energy dependence of the rejection parameter for electron recoils (left) and nuclear recoils (right)	240

6-3	Best-matching DMC pulses for a correctly identified electron recoil (top) and nuclear recoil (bottom), zoomed on the first 700 μs of the event.	241
6-4	Best-matching DMC pulses for a mis-identified electron recoil. Top: the best-matching DMC pulse, which is a simulated nuclear recoil. Bottom: the best matching simulated electron recoil.	243
6-5	Results of matching all 8 phonon pulses from electron and nuclear recoils from data and DMC. Electron recoils from data are shown in blue, nuclear recoils are shown in green.	245
6-6	Recoil energy dependence of the rejection parameter for electron recoils (left) and nuclear recoils (right) when matching all 8 phonon pulses.	246
6-7	Example pulse matches from a correctly identified ER (left) and NR (right)	247
6-8	Example pulse matches from two correctly identified ERs, showing that the DMC does not seem to match strongly peaked pulses as well as less ‘peaky’ pulses.	248
6-9	Example pulse matches from two correctly identified low-energy NRs, showing that this method can find good matches and classify events in the presence of significant noise.	249
7-1	Cross-sectional schematic of the CDMSII shielding layout. The dilution refrigerator is shown to the right, and the electronics readout hardware is on the left. Figure from [121]	255

7-2	Cross-sectional schematic of the CDMSII cryogenic system. The innermost layer in black is the base temperature stage, and the detectors are deployed inside of the black can in the nested rectangular structure in the middle of the figure. Moving outward, two intermediate (between the base and the LHe bath) temperature stages, the 50 mK and 600 mK stages, are shown in dark and light blue. Outside of these are the LHe bath, the LN bath, and the room temperature stage (green, yellow, red, respectively.)	257
7-4	CDMSII detector holders, with placeholders occupying the detector space. The DIB cutout is visible on the empty placeholder standing on its side. Courtesy: Dennis Seitz.	260
7-5	Picture of a CDMSII tower, with temperature stages, readout wiring, and side coax card connections labeled. Courtesy: Dennis Seitz. . . .	261
7-6	Picture and schematic layout of a SQUET card. Courtesy: Dennis Seitz.	262
7-7	CAD rendering of a fully assembled detector stack: detectors, towers, and SQUETS. At right is a cross-sectional view, with one ZIP removed. Courtesy: Patrick Wikus.	263

7-8	Illustration of the walking algorithm applied to four phonon pulses from a single event. The first row shows the charge trace, and the derived start time of the charge signal is shown as a solid red line across the lower four rows. These four rows show the four phonon traces; the full trace is on the left and a zoomed image around the onset of the pulse is shown at right. Unfiltered traces are shown in black, and filtered traces are shown in blue. The 20% risetime of each phonon pulse is indicated by a red X. The channel A trace highlights the poor resolution of this algorithm at low energy. Image courtesy of Scott Hertel [131]	269
7-9	Inner and outer charge signals: uncalibrated (left) and cross-talk corrected with preliminary calibration (right). Images courtesy of Kyle Sundqvist [132]	270
7-10	The 356 keV line of ^{133}Ba is not evident in Si detectors alone. Therefore, Si detectors are calibrated using events that also scatter in an adjacent detector (either Ge or Si). Image courtesy of Kyle Sundqvist [132]	272
7-11	‘Shrimp’ plot (<code>rde1</code> vs. <code>rppart</code>) from Detector T2Z4, Channel A. In this space, the foldback degeneracies evident in partition and delay alone are broken, and almost all outer events appear at the head of the shrimp.	277
7-14	Electron recoil band after first-pass correction in Run 126, Detector T2Z4. Events in red pass the first-pass yield cut, while those in blue fail it.[139]	284

7-15	Comparison of uncorrected (left) and position-energy corrected variables (right) in Detector T2Z4. Top: Yield vs. energy of gammas from WIMP search data. Bottom: 20% phonon risetime (<code>pdel</code>) vs. energy of gamma events from WIMP search data	285
7-16	KS test ‘kilt plots’ for Detector T2Z1 in Run 125. The columns represent the thirty randomly chosen comparison series, and the rows represent the >300 WIMP search series taken during Run 125. Blue indicates a good match, and red indicates a poor match. The onset of poor detector performance is obvious towards the end of the run. Courtesy: Scott Hertel [148]	295
7-17	Rate of low-yield events ($y < 0.8$) vs. time in Detector T2Z4, Run 126. Series shown in black pass the neutralization cut, red series fail criterion 1, blue series fail criterion 3, and cyan series failed more than one criterion. Error bars hidden to improve the legibility of the figure.[149]	296
7-18	Rate of zero-ionization events ($y < 0.1$) vs. time in Detector T2Z1, Run 411. The onset of Helium films towards the end of the run is clearly visible. The cut is placed at the red dashed line. X’s indicate data series already cut by other data quality cuts, while squares indicate data series considered good before application of <code>cHeFilm_c58</code> . Error bars hidden to improve the legibility of the figure. [150]	298
7-19	Setting of the <code>cBadResTight_c58</code> cut. Points indicate the amplitude resolution of all 6 channels of detector T1Z4 vs. time in Run 125. All resolutions are normalized to the run mean for that channel, and any series in which one of the channel resolutions increases above 1.25 (red dashed line) is cut [153].	299

7-20	Charge optimal filter goodness-of-fit vs. ionization energy in Detector T4Z3, Run 125. The black solid line indicates the cut used in this analysis, while the red dashed line indicates a previous definition of the cut, defined prior to some updates to the charge optimal filtering routine. Courtesy: Jianjie Zhang [158]	303
7-21	Baseline noise vs. time in the four phonon channels of Detector T2Z1, Run 125. Blue events pass <code>cPstd.c58</code> , while red events fail. The cyan lines indicate the limits of the cut. Courtesy: David Moore [160] . . .	305
7-22	Baseline noise vs. time in the two charge channels of Detector T2Z1, Run 125. Blue events pass <code>cQstd.c58</code> , while red events fail. The cyan lines indicate the limits of the cut. The effect of cryocooler microphonic pickup in the outer channel of this detector is clear. Courtesy: David Moore [160]	306
7-23	Example ‘cross-detector pileup’ event failing the <code>cGoodPStartTime</code> cut. The charge search window is shown in magenta, and the phonon search window is shown in green. Courtesy: Jeff Filippini [161]. . . .	308
7-24	The yield bands for electron recoils (red) and nuclear recoils (green) vs. event energy. The mean yields of the ER and NR bands are presented as dashed lines, and the $\pm 2\sigma$ limits are indicated by solid lines.	310
7-25	Plot of <code>qo</code> vs. <code>qi</code> showing the effect of the charge fiducial volume cut, Detector T4Z3, Run 125. The limits of the cut are presented as green solid lines. The polynomial fits extend to <code>qi</code> $\sim$ 5 keV, below which the cut is widened	311

7-26	Spectrum of trace amplitudes from a single veto scintillator panel, with the hardware threshold indicated by a green dashed line and the software threshold indicated by a red dashed line. Courtesy: Matt Fritts, [166]	313
7-27	Fits of the mean <code>pdel</code> and <code>pminrt</code> in Detector T2Z4. Neutrons are shown in green, phonon-side surface events in blue, and charge-side surface events in red. [171]	319
7-28	Fits of the elements of the covariance matrix between <code>pdel</code> and <code>pminrt</code> in Detector T2Z4. Neutrons are shown in green, phonon-side surface events in blue, and charge-side surface events in red. [171]	320
7-29	$\chi_b^2 - \chi_n^2$ vs. recoil energy of neutrons (green) and surface events (blue) in Detector T2Z4. [171]	320
7-30	Timing distributions of side-tagged betas on interior detectors in barium calibration data in R125-128. [172]	322
7-31	Weighting factors $s_{e,f}^{(i)}$ for the 6 energy/face bins. Interior detectors are shown in blue, and endcap detectors are shown in red after correction for the lack of phonon-side tagging. The mean weighting factors across the 6 interior detectors and and error on the mean are shown as a red band. [172]	325
7-32	Estimated and fit leakage vs. rejection cut position in all Detectors in recoil energy bins: 7-20 keV (red), 20-30 keV (blue), 30-100 keV(green). The cut position is offset by 10 to allow negative cut values to appear on the log scale x-axis. Because the cut will be placed at low expected leakage, the fit is only required to be good at leakages below ~ 0.5 events.	328

7-33	Measured and fit Spectrum-Averaged Exposure to a $60 \text{ GeV}/c^2$ WIMP vs. rejection cut position in all detectors in recoil energy bins: 7- 20 keV (red), 20-30 keV (blue), 30-100 keV (green). The cut position is offset by 10 to allow negative cut values to appear on the log scale x-axis. The fit is only required to be good at values of the rejection parameter >10	329
7-34	SAE vs. L in all energy bins and detectors for a $60 \text{ GeV}/c^2$ WIMP. .	332
7-35	$\frac{dSAE}{dL}$ vs. L in all energy bins and detectors for a $60 \text{ GeV}/c^2$ WIMP. Curves for individual bins are cut off to only show regions where $\frac{d^2SAE}{dL^2} < 0$	333
7-36	Cumulative distribution of yields from -2σ to $+2\sigma$ of the NR band for neutrons (blue solid) and surface events (black dashed). The surface event distribution is heavily weighted towards higher yields.	334
7-37	The rejection cut is illustrated by the magenta line; if no line is visible in a particular detector/energy bin, then the rejection cut was set to infinity in that bin. The calibration betas are shown in blue; those events that pass the neutron consistency cut are circled in red.	336
7-38	Nuclear recoil efficiency (right axis) and total exposure (left axis) vs. recoil energy of the c58R Si analysis.	337
7-39	Sensitivity of the c58R Si analysis (the limit that would result if no events are seen in the signal region). Results from other direct dark matter searches are included for comparison.	338
8-1	Posterior distributions of the expected leakage from surface events, with statistics only (blue) and statistics + systematics (red).	347

8-2	Estimating the bulk gamma leakage. Left: Fraction of bulk gammas expected to leak into the signal NR band, averaged over all detectors. Middle: Histogram of all WIMP search gammas that pass all cuts except the timing cut (Q-in, single, veto-anticoincident, data quality, thresholds, etc.). Right: Integrated expected leakage vs. energy. . . .	350
8-3	Yield (yic) vs. recoil energy (pric , keV) of all good low-background events from Runs 125-128 in the 8 Si detectors of this analysis. Events in red pass the neutron consistency cut, and events with black crosses pass the rejection cut. The signal NR band is outlined in green, the -3σ edge of the ER band is in blue, the charge threshold cut is shown in magenta, and the recoil energy threshold is the black dashed line. The width of these lines indicate variations in the band/threshold definitions across runs and series.	355
8-4	All good low-background events from Runs 125-128 in the 8 Si detectors of this analysis. Events in red pass the neutron consistency cut, and events with black crosses pass the rejection cut. The y-axis is ‘normalized yield’, the yield expressed in terms of number of σ s from the NR band mean. The x-axis is the rejection parameter $\chi^2_\beta - \chi^2_n$. The black box outlines the signal region, with the three dashed lines at the left edge indicating the rejection cut in the three energy bins. .	356
8-5	Phonon delay vs. phonon risetime in the two detectors containing the three WIMP candidates. The top row shows the neutrons and barium surface event calibration sets, along with the candidate events. The bottom row shows the neutrons, the candidate events, and all other signal region singles in these two detectors.	358
8-6	Phonon and charge traces of Candidate 1	359

8-7	Phonon and charge traces of Candidate 2	360
8-8	Phonon and charge traces of Candidate 3	361
8-9	Experimental upper limits (90% confidence level) for the WIMP-nucleon spin-independent cross section as a function of WIMP mass. We show the limit obtained from the exposure analyzed in this work alone (<i>black dots</i>), and combined with the CDMS II Si data set reported in [181] (<i>blue solid line</i>). Also shown are limits from the CDMS II Ge standard [182] and low-threshold [183] analysis (<i>dark and light dashed red</i>), EDELWEISS low-threshold (<i>orange diamonds</i>) [184], XENON10 S2-only [185] (<i>light dash-dotted green</i>), and XENON100 [186] (<i>dark dash-dotted green</i>). The filled regions identify possible signal regions associated with data from CoGeNT [187] (<i>magenta</i> , 90% C.L., as in- terpreted by Kelso <i>et al.</i> including the effect of a residual surface event contamination described in [188]), DAMA/LIBRA [189, 190] (<i>yellow</i> , 99.7% C.L.), and CRESST [191] (<i>brown</i> , 95.45% C.L.) experiments. 68% and 90% C.L. contours for a possible signal from these data are shown in blue and cyan, respectively. The asterisk shows the maxi- mum likelihood point at (8.6 GeV/c ² , 1.9×10^{-41} cm ²).	368
8-10	Limit on the spin-dependent WIMP-neutron cross-section from this analysis. Results from other direct dark matter searches are included for comparison.	369

List of Tables

4.1	Known values of parameters of the TES biasing and readout circuit used in Eq. 4.1	139
4.2	Relevant parameters of iZIP4 simulation used to make Figures 4-21 and 4-22. All parameters not listed are the default values available in the iZIP4 DMC template in the SuperCDMS CVS repository.	173
4.3	Relevant parameters of iZIP3 simulation used to make Figures 4-23 and 4-24. All parameters not listed are the default values available in the iZIP3 DMC template in the SuperCDMS CVS repository.	177
4.4	Relevant parameters of mZIP simulation used to make Figures 4-25. All parameters not listed are the default values available in the mZIP DMC template in the SuperCDMS CVS repository.	181
5.1	Weights for each of the simulated sources in the radiogenic background simulations. Weights were derived from a least-squares fit to the data spectrum.	206

5.2	Percentage of MISS candidates in the energy range 10-100 keV passing various fiducial volume cuts. The center column presents results from the ionization mapping method, and the right column presents results from full DMC simulations. All quoted errors are statistical. Though the two methods produce results in strong qualitative agreement, the differences are much larger than the statistical errors, indicating clear (and expected) systematic differences between the two methods. . . .	218
5.3	Summary of MISS leakage through a set of charge-based fiducial volume, nuclear recoil yield band, and multiple detector scatter cuts. . .	222
7.1	Total livetime (in days) for each detector in the c58 Si analysis	301
7.2	Factors entering the computation of the expected number of NRSSs in the c58 silicon WIMP search data. The observed number of NRSSs in c34 (col. 2) is multiplied by the ratio of c58 livetime to c34 livetime (in kilogram-days, columns 4 and 3, respectively) to produce the expected number of NRSSs in c58 (column 5). Errors in column 5 are $\sqrt{N}$. The expected number of NRSSs in T4Z3 is also corrected to account for the change in low-energy threshold between Runs 125 and 126. . . .	324
8.1	Observed and expected counts of NRSSs failing the timing cut	341
8.2	Contamination levels in the experimental materials and corresponding expected Si neutron rates.	352
8.3	Information about the events in the signal region in this analysis. . .	354

Chapter 1

The Dark Matter Problem

1.1 Introduction

Much of the history of physics has been driven by investigations of the fundamental constituents of the universe. This drive led to the development of the periodic table, particulate descriptions of materials, and Maxwell's theory of electromagnetic radiation. In the early 20th century, physicists elucidated and advanced the quantum description of atoms, and investigations of high-energy particle collisions and radioactive decay eventually culminated in the Standard Model of particle physics, which describes with unprecedented accuracy the nature of and interactions between all particles known to date.

Concurrently, astronomers studied the nature of the universe on the largest observable scales. Over the last century, a wealth of astrophysical and cosmological observations have led to the development of a standard model of big bang cosmology known as Λ CDM.

However, Λ CDM states that the majority of the matter content of the universe

does not appear to be composed of the stable particles making up the everyday world, but of a mysterious form of *dark matter*. Figure 1-1 shows the energy budget of the universe, which is apportioned into the baryonic matter that makes up the everyday world (5%), “cold” (nonrelativistic) dark matter (23%), and an unknown form of dark energy (72%) [1]. The Standard Model of particle physics provides no candidate particles with the right properties and abundances to fulfill the dark matter quota. This thesis focuses on efforts to directly detect interactions at low momentum transfer between nuclei and the particles that make up the dark matter component of the universe, and this chapter will review the astrophysical observations and arguments for the existence of dark matter, its inferred properties, and the numerous candidates that have been proposed to compose the dark matter.

Rather than taking the historical approach by reviewing, in chronological order, the decades of work that have coalesced to form the basis of modern cosmology, I believe it is much clearer to begin with this modern picture of the universe and relate the appropriate evidence in a top-down approach.

1.2 Λ CDM

The Λ CDM model represents the current understanding of the history and composition of the universe on the largest scales. A detailed discussion of the historical development and the entire observational and theoretical background of Λ CDM is beyond the scope of this discussion, but it is necessary to briefly cover the current state of knowledge and the relevant data providing the scientific underpinnings of the work comprising this thesis.




Figure 1-1: Energy budget of the universe: baryonic matter, dark matter, and dark energy. Figure from the NASA/WMAP Science team

1.2.1 The Friedmann-Robertson-Walker metric

One of the most astonishing insights from Einstein’s General Relativity was that the universe itself can be treated as a dynamic object that evolves over time. This revelation was completely at odds with the concurrently accepted assumption of a “static” universe, and Einstein famously posited a cosmological constant to prevent this dynamic behavior in his model of the universe. However, by the late 1920s, observations of distant galaxies by Edwin Hubble and other scientists had definitively demonstrated that the universe is expanding. Alexander Friedmann, Georges Lemaître, Howard Percy Robertson and Arthur Geoffrey Walker independently developed the most general time-dependent metric describing a large-scale isotropic and homogenous universe, which is now commonly known as the Friedmann-Robertson-Walker (FRW) metric.

$$ds^2 = c^2 dt^2 + a^2(t) \left(\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right) \quad (1.1)$$

where c is the speed of light; t represents time; k takes the values $\{-1, 0, 1\}$ to indicate whether the universe exhibits negative curvature, flatness, or positive curvature, respectively; $a(t)$ is a time-dependent scale factor indicating the radius of curvature of the universe; and r and Ω are unitless, comoving spherical coordinates.

However, the FRW metric by itself makes no predictions about the dynamics governing the evolution of the universe. To study the dynamics of the universe, the FRW metric must be coupled to the Einstein field equations to produce the Friedmann equations:

$$H^2 = \left(\frac{\dot{a}}{a} \right)^2 = \frac{8\pi G\rho}{3} - \frac{kc^2}{3} + \frac{\Lambda c^2}{3} \quad (1.2)$$

$$\dot{H} + H^2 = \left(\frac{\ddot{a}}{a}\right) = -\frac{4\pi G}{3} \left(\rho + \frac{3p}{c^2}\right) - \frac{kc^2}{3} + \frac{\Lambda c^2}{3} \quad (1.3)$$

where H is the Hubble parameter quantifying the rate of expansion of the universe, G is Newton's gravitational constant, and ρ and p are the time-dependent energy densities and pressures, respectively, of the contents of the universe. The possible contents usually included in a general model of the universe are as follows:

1. *Matter* refers to all non-relativistic matter, including ordinary baryonic matter and any cold dark matter. Matter provides no pressure and an energy density that scales as $a(t)^{-3}$.
2. *Radiation* encompasses relativistic contents of the universe, including photons, light neutrinos, and any other relativistic radiation. Radiation contributes pressure in direct proportion to its energy density ($p = \rho c^2/3$), which dilutes as $a(t)^{-4}$ – three powers from the volumetric expansion of space and one power from cooling as the universe expands.
3. *Dark Energy* refers to a cosmological constant-like energy density inherent to space. Dark energy exerts negative pressure in proportion to its energy density ($p = -\rho c^2$), which is entirely independent of $a(t)$.
4. *Curvature* refers to a potential energy inherent to large-scale spatial curvature. The energy density of curvature decreases with the square of the radius of curvature – $\rho \propto a(t)^{-2}$.

Utilizing the above energy-pressure relationships and scale-factor dependences and re-writing energy densities as fractions of the critical density of the universe $\Omega_x = \rho_x/\rho_c$, $\rho_c = \frac{3H^2}{8\pi G}$, the Friedmann equations can be recast into the more clear and useful form:

$$\left(\frac{H}{H_0}\right)^2 = \Omega_R a(t)^{-4} + \Omega_M a(t)^{-3} + \Omega_k a(t)^{-2} + \Omega_\Lambda. \quad (1.4)$$

Note that the relative energy densities of the components of the universe are not constant over time in a dynamic universe: in an expanding universe, the radiation energy density dilutes away more rapidly than the matter energy density, which in turn dilutes away more rapidly than the curvature energy density, while the dark energy density remains constant.

1.2.2 The Energy Budget of the Universe

A great deal of observational cosmology has focused on determining the large-scale structure of the universe, the evolutionary history of the scale factor and the energy ‘budget’ of the universe, i.e., measuring the values of the various Ω_x . Here, we will focus on three powerful and very different observations/techniques that combine to provide convincing constraints on the energy densities of the components of the universe: angular anisotropy in the Cosmic Microwave Background, Baryon Acoustic Oscillations, and Type Ia supernovae (SNe Ia).

Cosmic Microwave Background

The Cosmic Microwave Background (CMB) is thermal radiation that uniformly fills the universe at a blackbody spectrum at a temperature of 2.725K. The CMB was first predicted as a consequence of Big Bang cosmology in the late 1940’s in early studies of Big Bang Nucleosynthesis [2, 3, 4]. The discovery of the CMB is credited to Arno Penzias and Robert Wilson [5] in 1965, though earlier studies are now known to have accurately measured the CMB temperature through its effects on absorption lines [6] and directly through radio astronomy [7, 8]. The existence of the CMB and its strong

degree of homogeneity even at the largest scales is among the strongest evidence for Big Bang cosmology (along with the cosmological redshift and the abundances of the light elements).

Though the CMB is largely homogenous [9], the anisotropies in the CMB temperature (on the order of $1 : 10^{-5}$) reveal a wealth of information about the structure and content of the universe and can be used to place stringent bounds on the parameters of the Λ CDM model. The peak positions, peak shapes, and ratio of the powers in even and odd peaks are sensitive to the curvature of the universe, the baryon matter density, and the total matter density (among other parameters). Figure 1-2 shows recent measurements of the CMB temperature angular anisotropy spectrum from the WMAP experiment [1, 9, 10]. The fit corresponds to the best-fit 6-parameter Λ CDM model, which provides the values $\Omega_b = .0449 \pm 0.0028$, $\Omega_c = 0.222 \pm 0.026$, and $\Omega_\Lambda = 0.729 \pm .029$, where $\Omega_b, \Omega_c, \Omega_\Lambda$ refer to the baryonic matter, dark matter, and dark energy fractions of the critical density, respectively [10].

Baryon Acoustic Oscillations

The term Baryon Acoustic Oscillations refers to measurements of the length scale of baryonic matter clustering in the large-scale structure of the universe. This length scale of matter clustering in the current universe is related to the speed of acoustic wave propagation through baryonic matter in the early universe [11]. Density fluctuations in the plasma composing the early universe gave rise to acoustic oscillations due to the competing effects of gravity, which drew matter to clump in overdense regions, and radiation pressure, which tended to disperse overdense regions. At decoupling, around 377,000 years after the Big Bang, the photon pressure driving these acoustic oscillations would have disappeared as the universe cooled sufficiently to al-




Figure 1-2: Angular power spectrum of CMB anisotropies. The fit corresponds to the best-fit Λ CDM parameters. Figure from [10]

low hydrogen to form, imprinting an 'acoustic horizon' that should be evident as both a fluctuation scale in the CMB and as a correlation in the distance between overdense regions of matter (i.e., the matter 2-point function) in the modern universe. The length scale of this correlation provides a measurement of the matter component of the universe.

Figure 1-3, from [12] shows the correlation function times the square of the separation distance vs. the comoving separation distance. The acoustic feature is the spike at $\approx 100h^{-1}$ Mpc. The colored lines show the expected curves for models with $\Omega_m h^2 = 0.12, 0.13$, and 0.14 in green, red, and blue, respectively, and a purely CDM model ($\Omega_b = 0$, so no acoustic peak) with $\Omega_m h^2 = 0.105$ in magenta. The data allow a measurement of the total matter content of the universe: $\Omega_m = 0.273 + 0.123(1 + \omega_0) + 0.137 \times \Omega_K \pm 0.025$, where ω_0 is the effective value




Figure 1-3: Correlation function times the square of the separation distance vs. the comoving separation distance. The acoustic feature is the spike at $\approx 100h^{-1}$ Mpc. The colored lines show the expected curves for models with $\Omega_m h^2 = 0.12, 0.13$, and 0.14 in green, red, and blue, respectively, and a purely CDM model ($\Omega_b = 0$, so no acoustic peak)) with $\Omega_m h^2 = 0.105$ in magenta. Figure from [12]

of the equation of state parameter ω in the range $0 < z < 0.35$ ($\omega = -1$ in a dark energy dominated universe, and z is the redshift parameter) [12].

Type Ia supernovae

Type Ia supernovae are generally believed to be the explosions of white dwarf stars. The mass of white dwarf stars is bounded above by the Chandrasekhar limit of about 1.4 solar masses. Above this mass, electron degeneracy pressure is insufficient to support the star against its own gravity, and the star collapses. Often the star

collapses to form a neutron star, but in a Type Ia supernova, the increased pressure and density re-ignites the thermonuclear fusion processes in the star. The carbon and oxygen content of the star is burned in a matter of seconds, producing sufficient energy to trigger a supernova. Type Ia supernovae occur when a carbon-oxygen white dwarf star accretes sufficient matter to exceed the Chandrasekhar limit (1.38 solar masses). This mechanism ensures that the progenitor stars are all of equal mass, and so the luminosity profiles of Type Ia supernovae are remarkably uniform, making them excellent astronomical ‘standard candles’.

Because Type Ia supernovae are bright enough to be observed at cosmological distance scales, they can also provide a means of measuring cosmological parameters. With a sample of Type Ia supernovae at various distances, the expansion rate of the universe and its rate of change over time can be directly observed by comparing the absolute luminosities observed on earth and the cosmological redshift of the spectral features of the supernovae. Figure 1-4 (from [13], contains data from [14]) shows the magnitude vs. redshift of Type Ia supernovae from the Calan/Tololo Supernova Search and the Supernova Cosmology Project in Panel (a), and the residuals of the measured data with respect to the expected curve for a $\Omega_m = 0.28, \Omega_\Lambda = 0.72$ universe in Panel (b). Expected curves for various values of Ω_m, Ω_Λ are plotted in dashed lines for flat cosmologies and solid lines for $\Omega_\Lambda = 0$ cosmologies; the exact values of Ω_m, Ω_Λ are indicated outside the upper right corner of the figure. For a cosmology with $\Omega_{total} = 1$, the data indicate that $\Omega_m = 0.28^{+0.09}_{-0.08}$ (stat) $^{+0.05}_{-0.04}$ (syst.) and are strongly inconsistent with a $\Lambda = 0$ universe [13].

Combining the results discussed above, the energy budget of the universe can be determined to be $\Omega_m = 0.26 \pm 0.01$ and $\Omega_\Lambda = .74 \pm 0.03$, with negligible contributions (today) from radiation and curvature.




Figure 1-4: (a) Magnitude vs. redshift of Type Ia supernovae from the Calan/Tololo Supernova Search and the Supernova Cosmology Project in Panel, and (b) residuals of the measured data with respect to the expected curve for a $\Omega_m = 0.28, \Omega_\Lambda = 0.72$ universe. Expected curves for various values of Ω_m, Ω_Λ in dashed lines for flat cosmologies and solid lines for $\Omega_\Lambda = 0$ cosmologies; the exact values of Ω_m, Ω_Λ are indicated outside the upper right corner of the figure. Figure from [13] and contains data from [14]

1.2.3 The Dark Matter Fraction

With the knowledge that matter accounts for roughly a quarter of the total energy density of the universe, I now turn to a review of studies investigating the components of the matter fraction and indicate the existence of a large non-baryonic contribution to the matter content, known as dark matter.

Big Bang Nucleosynthesis

Big Bang Nucleosynthesis (BBN) was first proposed in the late 1940s to describe the creation of heavy nuclei (i.e., heavier than hydrogen) from fusion of free protons and neutrons in the first few minutes of the universe [3, 4, 2, 15]. At these early times, the extreme temperature and density conditions of the universe facilitated the creation of deuterium (D), helium 3 and helium 4 (He-3 and He-4), and lithium along with unstable isotopes such as tritium and low-Z beryllium. The theory underlying BBN provides a detailed description of the production of heavier elements in a hot, dense soup of baryons in an expanding, cooling universe, and observation of the relative abundances of these elements today thus provides information about the initial density of baryons (protons and neutrons) in the universe.

Fig. 1-5, from [16], shows BBN predictions of the relative abundances of light elements (relative to that of hydrogen) versus the baryon density of the universe. The figure compiles measurements and predictions from a number of researchers, please see the references and caption of Fig. 32 in [16] for the list. The widths of the curves indicate the uncertainties in the predictions. The vertical extent of the boxes indicate measured elemental abundances, and the horizontal extent is set by the range in which measurement and predictions are in agreement. The D and ^3He components are in agreement with a baryon density of the universe, $\Omega_b, \approx 0.04$. The

heavier ^4He and ^7Li indicate lower values, and the authors mention poorly understood systematic errors in these measurements and in relating these measurements to the primordial populations [16]. Despite the discrepancy between elements, the value of $\Omega_b \approx 0.04$ compared with other measurements indicating $\Omega_m \approx 0.24$ implies that 5/6 of the matter content of the universe is in non-baryonic forms.

Galaxy Cluster Dynamics

Dark matter was first proposed in 1933 by Fritz Zwicky based on his work measuring the dynamics of galaxies in the Coma cluster [17]. Zwicky used the virial theorem to relate the mass of the cluster to the velocity dispersion (measured through the Doppler effect) of the galaxies comprising the cluster, and found that the cluster mass measured in this way was in strong disagreement with mass estimations from the luminosities of galaxies in the cluster. To resolve this discrepancy, Zwicky proposed that a large fraction of the cluster mass was contained in a non-luminous component, which he termed ‘dark matter’. Though Zwicky’s initial proposal of a dark to luminous matter ratio on the order of 500 has since been shown to be inaccurate as measurement techniques have evolved, more sophisticated studies of galaxy clusters with mass measurements made through gravitational lensing continue to support the dark matter scenario [18].

Galactic Rotation Curves

The line-of-sight velocities of stars in spiral galaxies can be measured through Doppler shifts of known emission lines. These Doppler shifts can be measured at various radii from the galactic center, and correcting the data for the inclination of the galaxy to the observer’s line of sight provides the ‘galactic rotation curve’, the rotational




Figure 1-5: BBN predictions of the relative abundances of light elements (relative to that of hydrogen) versus the baryon density of the universe, with measurements. The widths of the curves indicate the uncertainties in the predictions. The vertical extent of the boxes indicate measured elemental abundances, and the horizontal extent is set by the range in which measurement and predictions are in agreement. The D and ${}^3\text{He}$ components are in agreement with a baryon density of the universe, Ω_b , ≈ 0.04 . The heavier ${}^4\text{He}$ and ${}^7\text{Li}$ indicate lower values, and the authors mention poorly understood systematic errors in these measurements and in relating these measurements to the primordial populations. This figure compiles measurements and predictions from a number of researchers, please see [16].

velocity of stars as a function of radius from the center of the galaxy. Identifying the gravitational acceleration

$$a_g \propto \frac{GM(r)}{r^2} \quad (1.5)$$

with the centripetal acceleration in uniform circular motion

$$a_c = \frac{v^2}{r} \quad (1.6)$$

where G is Newton's gravitational constant, r is the radius of the test particle, v is the rotational velocity of the star, and $M(r)$ represents the amount of mass interior to the radius of the test particle, we find the relationships

$$v^2 \propto \frac{GM(r)}{r} \quad (1.7)$$

or alternatively

$$M(r) \propto \frac{v^2 r}{G} \quad (1.8)$$

Thus, the rotational velocity curve can be used to measure the mass distribution of a galaxy. A separate estimate of the (luminous) mass can be made using the mass-luminosity relationships for the stars comprising the galaxy. If the luminous mass were the entire mass of the galaxy, the prediction is that the rotational velocity would rise rapidly near the galactic center, and then fall off as $r^{-\frac{1}{2}}$ as the interior mass function $M(r)$ asymptotes to a constant at the luminous edge of the galaxy.

Observations of the rotational velocity curves of a number of spiral galaxies began in the 1950s [19], though it was Vera Rubin's work in the 1970s [20] that renewed interest in Fritz Zwicky's proposals of dark matter from the early 20th century. Rubin's work has since been extended to include a large number of galaxies, see




Figure 1-6: Set of galactic rotation curves. Note that the velocity of stars does not decrease at high radius, but rather asymptotes to a constant, indicating an interior mass function $M(r) \propto r$. As the luminous interior mass grows much more slowly than this, these curves imply contributions from a significant dark matter component [21].

[21]. A collection of galactic rotation curves from [21] can be seen in Figure 1-6, and Figure 1-7 [22] shows the contributions of the luminous disk and the dark halo for NGC 3198 specifically. Rather than falling off at high radii, the velocity curves roughly asymptote to a constant for all stars outside of the central dense region, out to the furthest observable stars. According to Eq. 1.8, this corresponds to an interior mass function $M(r) \propto r$. This observation naturally leads to the proposal of a large non-luminous mass contribution to the mass of galaxies, and the relative fraction of dark-to-luminous matter falls in line with the ratios revealed by other techniques.




Figure 1-7: The black squares show the measured galactic rotation curve of NGC 3198 with a fit. The contributions from the measured luminous component of the galaxy and the proposed dark matter component are also shown in the lines labeled ‘disk’ and ‘halo’, respectively. [16].

The Bullet Cluster

The Bullet Cluster (1E 0657-558) is composed of two colliding galaxy clusters whose gaseous (baryonic) and dark components have separated due to the collision [23]. The cores of the two subclusters passed each other ≈ 100 Myr ago. In a cluster collision, the stellar components of the clusters (localized in the galaxies) pass through each other in an essentially collision-free manner, but the intracluster plasma, which makes up the majority ($> 85\%$) of the total luminous matter, acts as a collisional fluid and experiences ram pressure that substantially slows the relative motion of the plasmas from the two subclusters. The dark matter components of the two clusters, like the stellar components, are expected to be effectively collisionless and pass through each other.

Measuring the X-ray luminosity of the Bullet Cluster reveals the mass and spatial distribution of the intracluster plasma and demonstrates that the intracluster plasmas of the two subclusters indeed slowed significantly during the collision due to strong hydrodynamic interactions; see the right panel of Fig. 1-8. The stellar mass content can be directly measured from the spatial luminosity distribution in the optical band (left panel of Fig. 1-8) and is significantly smaller than the plasma mass content. Therefore, the total mass distribution should track the plasma distribution in the absence of any dark matter content, but should track the collisionless stellar mass distribution if the subclusters do indeed have large dark matter components (as dark matter is also collisionless).

The gravitational microlensing technique measures the spatial distribution of the total mass independent of the nature of the mass (i.e., whether it is dark or luminous). Gravitational microlensing measurements of the mass distribution of the Bullet Cluster (contours in Fig. 1-8) have revealed that the regions of highest to-




Figure 1-8: Left: Optical image of the Bullet Cluster 1E 0657-558 from the Magellan observatory. The white bar indicates 200 kpc at the distance to the cluster. Right: X-ray image of the Bullet Cluster from the Chandra X-ray Observatory. The contour lines illustrate the lensing parameter κ , with the outer contour corresponding to $\kappa = 0.16$ and interior lines increasing in steps of 0.07. κ is directly proportional to the surface mass density of the gravitational lens, indicating that the majority of the mass is concentrated away from the regions of high plasma density [23].

tal density are strongly separated from the intracluster plasma and are significantly more massive than the stellar content alone (as measured through optical luminosity), thereby providing extremely convincing evidence for the existence of dark matter. Moreover, these conclusions are independent of hypotheses about the nature of the dark matter. The observations of the Bullet Cluster convincingly demonstrate the existence of a collisionless dark matter fluid present in these clusters with a ratio of approximately 6:1 between the dark and luminous components of the clusters.

1.3 The Weakly Interacting Massive Particle Hypothesis

The converging experimental evidence for the existence of a non-luminous, non-baryonic form (or forms) of matter that dominates the ‘normal’ baryonic component of the universe has led to numerous theoretical proposals about the nature of the dark matter, as the Standard Model provides no particles with the necessary properties. These proposals include (but are not limited to):

- Massive Compact Halo Objects (MACHOs): non-luminous objects such as black holes, neutron stars, dwarf stars, etc., composed of normal baryonic matter and thus disfavored by BBN considerations (Sec. 1.2.3).
- Axions: axions arise from the Peccei-Quinn mechanism, which was postulated as a possible solution to the strong CP problem in QCD [24]. Axions are hypothesized to have very low masses and cross-sections with normal matter, and are currently the subject of active experimental searches.
- Weakly Interacting Massive Particles (WIMPs): The term WIMP describes a class of particles that are heavy, stable, and interact only through the weak nuclear force or through channels with interaction cross-sections no larger than the weak scale. A number of theories of beyond the Standard Model physics provide WIMP candidates.
- Finally, theories such as Modified Newtonian Dynamics (MOND) that modify the behavior of gravity at large scales have been proposed to account for the above observations without requiring a novel form of dark matter

Figure 1-9, reproduced from [25], summarizes the mass and cross-section ranges corresponding to some of the new particles proposed to solve the dark matter problem. It is beyond the scope of this discussion to address each of these proposals in detail, and the reader is referred to [26] for a thorough discussion. As this thesis deals with the search for WIMP dark matter, we will turn our attention to the arguments for and against the WIMP hypothesis.

1.3.1 Properties of a generic WIMP

It is important to note that the term WIMP does not refer to a specific proposed particle, but a class of previously unobserved particles with particular properties. A generic Weakly Interacting Massive Particle by construction exhibits the inferred properties of the dark matter component of the universe; that is, WIMP candidate particles are

- *Dark*: Obviously a dark matter candidate must be dark, and WIMP particle dark matter fulfills this requirement by being completely electrically neutral and thus completely inert to electromagnetic radiation.
- *Heavy*: WIMP candidate particles span a range of masses depending on the detailed parameters of the particular theory giving rise to the candidate, but the masses tend to be in the GeV–TeV scale when the WIMP is considered to be a thermal relic of the Big Bang 1.3.2, and this mass range is consistent with a particle that could provide the necessary mass while remaining undetected in collider experiments while also being ‘cold’ (slow) enough to form substructure.
- *Weakly Interacting*: Strong interactions (not necessarily the strong nuclear force) in the dark sector could drive dark matter clustering on scales inconsis-




Figure 1-9: Illustration of some theoretically-motivated dark matter particle candidates. The y-axis illustrates the interaction cross-section with normal matter, and the x-axis indicates the expected mass range of the particle. Figure from [25].

tent with gravitational lensing observations.

- *Stable*: A WIMP must be stable against decays to known particles on the timescales of the age of the universe, or the decay products would be observed today, and observations of the universe on different time scales (e.g., CMB vs. Bullet Cluster) would provide different values for the WIMP abundance.

1.3.2 The WIMP as a thermal relic

Another attractive feature for WIMP dark matter is that the general range of WIMP masses and cross-sections allows for the current dark matter density to arise naturally in the thermal history of the universe. This convergence of GeV-scale masses, weak-scale cross-sections, and the observed dark matter density through thermal production and freeze-out in the early universe is often referred to as the WIMP miracle.

When the temperature of the early universe exceeded the mass of the WIMP, WIMPs would have been in thermal equilibrium, that is, the WIMP annihilation and creation rates ($\chi\bar{\chi} \Rightarrow X\bar{X}$ and $X\bar{X} \Rightarrow \chi\bar{\chi}$ respectively, where χ represents the WIMP and X represents generic SM particles) would have been equal. Thus, WIMPs were present in the same densities as normal particles. As the universe expanded and cooled, the average energy of particles decreased, and thus WIMP production from collisions of lower-mass particles became exponentially less probable. In this phase, the WIMP creation rate fell rapidly to zero, while the WIMP annihilation rate was dictated by the number density of WIMP particles, i.e., $\chi\bar{\chi} \rightarrow X\bar{X}$. Continuing forward in time, as WIMP annihilation proceeded without WIMP creation and the universe continued expanding, the number density of WIMPs decreased until the probability of two WIMPs interacting to annihilate also fell to essentially zero, i.e.,

$\chi\bar{\chi} \rightarrow X\bar{X}$. At this point, the WIMPs are said to have ‘frozen out’, and the remaining number of WIMPs will henceforth remain essentially constant as long as the WIMP is stable.

Figure 1-10 illustrates the density history of a 100 GeV dark matter particle during the early universe. The y-axis indicates the comoving number density Y (left) and the resulting thermal relic density Ω_{DM} (right). The x-axis shows the temperature (bottom) and the corresponding time (top). The black line traces the history of a 100 GeV particle with an annihilation cross-section that reproduces the correct relic density today, and the shaded regions around this line vary the dark matter particle’s annihilation cross-section by factors of 10, 10^2 , 10^3 . The dashed line indicates the density profile of a particle in thermal equilibrium. The “WIMP miracle” refers to the fact that a WIMP of order 100 GeV with interactions on the order of the weak scale produces a thermal relic with the correct relic abundance.

For a more detailed mathematical discussion of the relic WIMP density, I refer the reader to [27, 26].

1.3.3 Beyond Standard Model candidates

The Standard Model of particle physics contains 17 observed fundamental particles: 6 quarks, 3 charged leptons, 3 neutral leptons (the neutrinos), 4 gauge bosons (γ , W, Z, and the gluon), and the Higgs boson. None of these particles meet the four requirements of a WIMP discussed above. Of the particles listed above, only the γ , the neutrinos, the electron, and quarks confined to a proton exhibit stability. Of these, the γ is massless, the electron and proton are electrically charged, and the neutrinos are too light and too ‘hot’ to allow for structure formation. Therefore, the Standard Model does not currently provide any WIMP candidates.




Figure 1-10: Illustration of freezeout for a 100 GeV dark matter particle. The y-axis indicates the comoving number density Y (left) and the resulting thermal relic density Ω_{DM} (right). The x-axis shows the temperature (bottom) and the corresponding time (top). The black line traces the history of a 100 GeV particle with the correct relic density, and the shaded regions around this line vary the dark matter particle's annihilation cross-section by factors of 10, 10^2 , 10^3 . The dashed line indicates the density profile of a particle in thermal equilibrium. Figure from [26].

However, taking a broader view, the Standard Model really provides no candidate dark matter particles at all. Despite its overwhelming success in describing an incredibly wide range of experimental results, the lack of a dark matter candidate is one of a number of theoretical and phenomenological issues indicating the existence of physics beyond the Standard Model. These phenomena include the observed matter/antimatter asymmetry, the strong CP problem, the hierarchy problem, and the existence of neutrino mass. These issues have led theorists to investigate allowable extensions to the Standard Model, and many of these extensions naturally provide a WIMP candidate. For example, supersymmetric extensions to the Standard Model would necessitate the existence of supersymmetric partner particles for each of the known particles. There exist regions in the supersymmetric parameter space in which the lightest supersymmetric partner is heavy, electrically neutral, weakly interacting, stable against decays to other supersymmetric particles by virtue of being the lightest, and stable under R-parity against decays to Standard Model particles. Other example ‘Beyond the Standard Model’ WIMP candidates include sterile neutrinos and the lightest Kaluza-Klein particle. For further information about theories beyond the Standard Model and their relation to dark matter, the reader is referred to [27, 26, 28].

1.3.4 Problems for the WIMP dark matter hypothesis

Cold dark matter is well-motivated at the largest scales of cosmology, but tensions between the WIMP dark matter hypothesis and observations arise when considering structure formation on the scale of individual galaxies. Cosmological simulations of galaxy formation within the Λ CDM framework with weakly interacting dark matter produce some results that seem to be at odds with observations. The study of

galaxy formation is a complex and active field of research, and these considerations are issues for the overall CDM paradigm and not WIMPs in particular, but it is worth mentioning them here.

The Cuspy Halo Problem

The Cuspy Halo Problem refers to the discrepancy between the galactic dark matter halo density profiles predicted by Λ CDM galaxy formation simulations and the density profiles inferred from galactic rotation curves [29, 30]. Galaxies are believed to form through a process of hierarchical clustering [31], in which large galaxies are formed via mergers of many smaller objects, with the initial density fluctuations arising during the Big Bang. Hierarchical clustering simulations tend to produce dark matter halos that are significantly denser near the core and have steep density profiles at outer radii. The Navarro-Frenk-White (NFW) profile, proposed as a universal profile for dark matter halo densities, predicts densities $\rho_{DM} \propto r^{-1}$ in the inner core and $\rho_{DM} \propto r^{-3}$ at high radii [32].

In contrast, measured galactic rotation curves of disks galaxies reveal that the rotational velocities of stars asymptote to a constant velocity at high radii, and as the relative density of luminous matter is small in these regions, this observation implies $\rho_{DM} \propto r^{-2}$. In the inner regions of disk galaxies, the rotational velocity profile depends on both the luminous and dark components, and the term associated with the dark matter component is found to rise roughly linearly with radius, which implies that the interior mass function $M(r) \propto r^3$, or ρ_{DM} is constant. The entire picture is somewhat more complicated than presented above, and more recent research has focused on studying CDM halos in dark matter dominated dwarf galaxies with low surface brightnesses. Many authors have proposed solutions that involve changes to

the simple cold dark matter model that may provide insight into the source of the problem (see [33, 34, 35] for examples). Galaxy formation simulations remain an active field of research, and the Cuspy Halo Problem awaits resolution.

The Missing Satellites Problem

Semianalytic and numerical treatments of structure formation in the Λ CDM paradigm indicate that dark matter halo substructure should survive the merging and accretion processes responsible for the formation of galaxies, just as galactic-scale substructure survives these processes in the formation of galaxy clusters. The surviving substructure should take the form of dwarf galaxy satellites of larger galaxies, and in particular, the Milky Way would be expected to have $\mathcal{O}(10^2 - 10^3)$ satellite dwarf galaxies with observable luminosities [36, 37, 38]. The actual number of observed satellites of the Milky Way is significantly less, of $\mathcal{O}(10)$, and this discrepancy has become known as the Missing Satellite Problem. However, as the technology utilized in near-field cosmology observatories has improved in recent years, the number of observed satellites has steadily increased [39, 40]. Many of these newly discovered satellites have extremely low surface brightnesses, explaining why they escaped detection until recently, as well as large velocity dispersions that indicate that they are extremely dark matter dominated [41, 42]. Furthermore, these observations provide empirical evidence for the likely existence of a large number of satellites that have yet to be discovered. It remains to be understood why these satellite galaxies have such faint luminosities – impairment of stellar formation due to reionization has been proposed as an explanation – but the Missing Satellites Problem does not appear to be quite as troubling as it was just ten years ago.

Another issue that arises on the galactic scale is known as the angular momentum

problem; CDM galaxy formation implies too much baryonic matter with low angular momentum to support the observed disk structure of spiral galaxies. This problem may be resolved if the baryonic and dark matter components of the galaxy are not required to have the same initial momentum distributions. Other open issues include whether Λ CDM can properly predict the number of galaxies, the merger history, galactic luminosities, morphologies, etc. The current situation can be summarized by stating that the Λ CDM model performs extremely well at the largest scales of the universe, but whether it can accurately describe galactic-scale phenomena remains to be determined.

1.4 Conclusions

In conclusion, a model of the universe containing $\approx 5\%$ baryonic matter, $\approx 23\%$ dark matter, and $\approx 72\%$ dark energy reproduces a wide range of phenomena on cosmological scales, including the nature of the anisotropies in the cosmic microwave background, the abundance of baryonic matter, and the scale of baryon acoustic oscillations. The ratio of dark to luminous matter is also revealed in the behavior of galactic clusters and the rotation curves of individual galaxies. Finally, the model of dark matter as a cold, collisionless fluid is strongly supported by gravitational lensing measurements of the Bullet Cluster merger. However, unresolved problems arise when the detailed formation histories of galaxies are considered within the Λ CDM framework.

A number of candidate explanations for the nature of the dark matter have been proposed. Many theories of physics beyond the Standard Model offer a natural dark matter particle in the form of a heavy, stable WIMP. In the next chapter, we will turn to the experimental efforts to directly or indirectly detect and constrain the

nature of WIMP dark matter.

Chapter 2

Experimental Searches for Dark Matter

2.1 Introduction

The astrophysical evidence for dark matter has spawned a wide variety of experimental efforts to investigate the nature of the dark matter. These experimental efforts have searched across the spectrum of dark matter candidates, and in this chapter, we will focus on the experiments designed to search for WIMP dark matter. These experiments can be broadly classified into three categories:

- *Production*: Dark matter candidate particles can be sought among the products of high-energy particle collisions at accelerator experiments. The exact details of the production channel are unknown, but as dark matter candidates must be stable and tend to have low cross-sections with normal matter, a produced dark matter particle would generally escape the detector and reveal its presence through large amounts of missing energy.

- *Indirect Detection:* Though the cosmological annihilation rate of WIMP dark matter fell to essentially zero at freeze-out (see Sec 1.3.2), the density of relic dark matter particles distributed throughout the universe may be locally enhanced due to gravitational binding, e.g., at the centers of galaxies. At these higher densities, dark matter particles and anti-particles may be able to annihilate in sufficient number for decay products to be detected.
- *Direct Detection:* Dark matter particles may be able to directly interact with normal matter through non-gravitational channels; WIMPs, for example, are a candidate dark matter particle that would be able to scatter from weakly-charged SM particles through the weak force. Direct detection experiments seek to extract dark matter-nucleon interactions through a background of ‘normal’ electromagnetic and neutron radiation.

2.2 Production at Colliders

Discovery and characterization of new particles has traditionally been the domain of particle accelerators. If WIMPs couple to nucleons, as they must for direct detection experiments to observe them, then they can also be produced in collisions of high energy nucleons at colliders. Direct detection experiments are capable of discovering WIMP dark matter and constraining the mass and cross-section of the particle, though these constraints are not independent of the assumed astrophysical density and velocity distributions. However, detailed study of the nature of the dark matter and its couplings to Standard Model particles is more suited to collider experiments.

The Large Hadron Collider (LHC) at CERN is currently the highest energy collider in the world, colliding two proton beams at a center-of-mass (CM) energy of

7 TeV. Therefore, if the WIMP mass is at or below the hundreds of GeV - TeV scale, it should be accessible to production at the LHC. The accessible WIMP mass does not reach 7 TeV because the 7 TeV of kinetic energy is shared among the partons (quarks and gluons) making up the protons in the beam. Unlike most direct detection experiments, for which a handful of events would be sufficient to claim discovery, study of dark matter at the LHC will require a large number of produced WIMPs for a number of reasons. The LHC beam energy is sufficiently high that collisions must be treated as collisions between partons (quarks and gluons) rather than between the protons themselves. The momenta of the incoming partons are unknown, which complicates reconstruction of any individual event. The exact details of a dark matter search at the LHC also depend on models for how a particular WIMP candidate couples to Standard Model particles. It is often assumed that the interactions between the SM particles and the WIMPs are mediated by higher-mass unknown particles, which are integrated out as intermediate states in computing effective WIMP production cross-sections. Thus, discovering a dark matter particle and extracting the details of the WIMP production mechanism will require statistical treatment of a relatively large number of WIMP production events. Finally, it is worth noting that because the individual quarks and gluons involved in the collisions tend to carry a small fraction of the total momentum of the parent nucleon, the frequency of WIMP production at the LHC decreases as the WIMP mass increases. However, I would be remiss to not state that the details mentioned above are complications rather than show-stoppers, and the ATLAS and CMS collaborations have published competitive WIMP production limits [43, 44], which are presented in Figure 2-1.

The proposed International Linear Collider (ILC) or Compact Linear Collider (CLIC) should also prove to be a valuable tool in characterizing WIMP dark matter. These competing proposals are both electron-positron colliders with proposed CM




Figure 2-1: Inferred 90% CL ATLAS limits on the spin-independent WIMP-nucleon cross-section vs WIMP mass. Figure from [43]

energies of 500 GeV-1 TeV (ILC) or 3 TeV-5 TeV (CLIC). Electron-positron colliders preclude many of the complications present in hadron colliders as the colliding particles are fundamental. Any particle with a mass below one-half of the CM energy is available for production. Importantly, the CM energy is known exactly and can be scanned, allowing for precision measurements unavailable at hadron colliders. Please see [45, 46] for recent analyses of dark matter-nucleon cross-section limits at the Tevatron, and [47, 48, 49] for treatments of the prospects for WIMP discovery and characterization at the LHC and ILC.

2.3 Indirect Dark Matter Detection

Indirect dark matter detection experiments search for products from annihilation of WIMPs and anti-WIMPs. The self-annihilation cross-section must be small for thermal relic WIMPs to have the proper cosmological dark matter density. Also, given that annihilation requires a WIMP and an anti-WIMP, the success of indirect searches hinges on the WIMP either being a Majorana particle or not exhibiting the matter-antimatter asymmetry observed in the ‘bright’ matter sector. The absolute rate of annihilations scales with the square of the local density, and so indirect dark matter searches are particularly sensitive to the distribution of dark matter. There are three primary channels exploited by indirect dark matter searches: gamma rays, cosmic rays, and neutrinos.

2.3.1 Gamma rays

High-energy gamma rays produced in WIMP annihilations retain directional information over large distances and could exhibit sharp spectral features such as a monoenergetic line or a spectrum with a sharp cutoff (depending on the details of the annihilation process).

2.3.2 Cosmic rays

WIMP annihilation could also produce high energy charged particles, known as cosmic rays. Searches for dark matter annihilation in cosmic rays tend to focus on finding excesses in the antimatter-to-matter ratio at high energies, as most cosmic rays are normal matter particles that have been accelerated by cosmic sources. However, these ‘normal’ cosmic rays can still produce high energy antiparticles in secondary in-

teractions, so the cosmic antiparticle background must be well understood to claim a dark matter signal from cosmic rays. Because cosmic rays are charged particles, their tracks are bent as they propagate through interstellar magnetic fields, and cosmic rays thus do not provide strong directional information. Positrons are a relatively small component of the cosmic ray flux. WIMP annihilation is expected to produce electrons and positrons in equal number (either directly or through decay chains beginning with higher-mass annihilation products), and thus the positron fraction is the quantity of interest for dark matter searches in the positron channel. However, WIMP annihilation is obviously not the only process by which e^+e^- pairs can be created, and it is quite difficult to rule out other sources of any unexpected rise in the positron fraction at particular energies.

2.3.3 Neutrinos

Due to their weak interaction cross-sections, neutrinos produced in WIMP annihilations retain directional information and should escape the production source without interaction, raising the possibility of searching for annihilation products from WIMP overdensities captured in the sun or earth (sources effectively closed to searches through gamma rays, which would be absorbed soon after production). The high energies of neutrinos produced in WIMP annihilation should allow them to be distinguished from neutrinos produced in normal nuclear interactions. [50]

2.4 Direct Dark Matter Detection

Direct dark matter detection experiments aim to measure dark matter particles from the astrophysical halo of our galaxy scattering off of target nuclei. The primary

difficulty faced in direct dark matter detection is the extreme rarity of the dark matter-nucleon interaction as compared to the rate of known particles scattering off of valence electrons or nuclei in the detector. These experiments must be able to either discriminate, veto, or screen out backgrounds from γ and β particles scattering from valence electrons in the target and α particles and neutrons scattering from the target nuclei, and must reject these backgrounds to ever-increasing levels as the target mass scales up. Detectors must also have energy detection thresholds on the order of keV. A wide variety of experimental techniques have been developed to separate the dark matter signal from the host of backgrounds. A generic WIMP dark matter particle is expected to interact with normal matter primarily through recoiling off of atomic nuclei rather than electrons. CDMS, among other experiments, utilizes detectors that produce measurably different signals when a particle scatters from a nucleus (nuclear recoil, or NR) or from an electron (electron recoil, or ER). These experiments can discriminate against the primary radiogenic backgrounds on earth on an event-by-event basis but must still shield the target from NR backgrounds. Event-by-event background rejection is not the only choice, however. Some experiments search for an overall excess in event rate with an energy spectrum consistent with an astrophysical WIMP source, and others seek to exploit the astrophysical properties of the dark matter halo by searching for annual modulation in the recoil rate (as the earth oscillates within a ‘dark matter wind’) or by measuring the directionality of recoils within the target mass. In this section, I will review the computation of event rates in a direct detection experiment, the basic physics of the detectors used by various collaborations, and some of the search strategies employed in direct detection experiments.

2.4.1 Event Rates - Astrophysics

The rate of WIMP-nucleon scattering in an earthbound detector depends in general on the WIMP's intrinsic properties, its local density and velocity distributions, and the properties of the detector (energy threshold, atomic mass, etc.). The derivation of the event rate presented here follows the discussions in [27]; see also [51].

To first order, the rate of elastic WIMP-nucleon scatters per unit detector mass should scale with the elastic WIMP-nucleon cross-section σ , the mean WIMP speed relative to the detector $\langle v \rangle$, the local WIMP number density n ($= \frac{\rho_\chi}{m_\chi}$, where ρ_χ and m_χ are the WIMP density and mass, respectively), and inversely to the mass of the target nucleus m_N (as higher atomic mass means fewer target nuclei per unit detector mass). At a more detailed level, we must consider that the velocities of astrophysical WIMPs will be distributed according to some distribution function $f_1(v)$, and that the differential cross-section $\frac{d\sigma}{d|q|^2}$ will depend on the relative velocities as well.

With these considerations, we write the differential rate per unit detector mass as:

$$dR = \frac{\rho_0}{m_\chi m_N} v f_1(v) \frac{d\sigma}{d|q|^2} d|q|^2 dv \quad (2.1)$$

$\frac{d\sigma}{d|q|^2}$ can be written in terms of a form factor $F(|q|)$ through the relationship (Eq. 7.12 in [27])

$$\frac{d\sigma}{d|q|^2} = \frac{\sigma_0}{4\mu^2 v^2} F^2(|q|) \quad (2.2)$$

The total momentum and energy transferred to the nucleus in a scatter are given by the equations:

$$|q|^2 = 2\mu^2 v^2 (1 - \cos \theta) \quad (2.3)$$

$$Q = \frac{|q|^2}{2m_N} = \frac{\mu^2 v^2}{m_N} (1 - \cos \theta) \quad (2.4)$$

where θ is the scattering angle in the center-of-momentum frame, and μ is the reduced mass of the WIMP-nucleus system ($\mu = \frac{m_\chi m_N}{m_\chi + m_N}$). From Eq. 2.4, the minimum velocity v_{min} required to transfer energy Q to the detector occurs at $\theta = \pi$ and is equal to

$$v_{min} = \sqrt{\frac{Q m_N}{2\mu^2}} \quad (2.5)$$

With these equations in hand, we can rewrite the differential rate as

$$\frac{dR}{dQ} = \frac{\sigma_0 \rho_0}{2m_\chi \mu^2} F^2(Q) \int_{v_{min}}^{\infty} \frac{f_1(v)}{v} dv \quad (2.6)$$

Information regarding the astrophysical model for the WIMP velocity distribution is contained in the integral in Eq. 2.6, and we can rewrite a general equation for the differential scattering rate as

$$\frac{dR}{dQ} = \frac{\sigma_0 \rho_0}{\sqrt{\pi} v_0 m_\chi \mu^2} F^2(Q) T(Q) \quad (2.7)$$

where v_0 is the speed of the solar system around the galactic center (and thus presumably with respect to the WIMP halo), and

$$T(Q) = \frac{\sqrt{\pi} v_0}{2} \int_{v_{min}}^{\infty} \frac{f_1(v)}{v} dv \quad (2.8)$$

This formulation of the differential event rate is useful in that it separates the

rate into a prefactor, the factor $T(Q)$, which contains the astrophysical information about the WIMP velocity distribution, and the form factor $F^2(Q)$, which captures the details of the nuclear physics in the WIMP-nucleus scattering. The observed event rate in a detector can be obtained by integrating Eq. 2.7 from the lower energy threshold of the detector to the upper energy threshold (if any).

When computing expected WIMP-nucleus scattering rates, a Maxwellian speed distribution is commonly assumed:

$$f_1(v) = \frac{4v^2}{v_0^3\sqrt{\pi}} e^{-\frac{v^2}{v_0^2}} dv \quad (2.9)$$

Solving the integral in Eq. 2.8 gives

$$T(Q) = e^{-\frac{v_{min}^2}{v_0^2}} \quad (2.10)$$

Plugging Eqs. 2.10 and 2.5 back into Eq. 2.7 reveals a differential rate that decreases exponentially with the transferred energy:

$$\frac{dR}{dQ} = \frac{\sigma_0 \rho_0}{\sqrt{\pi} v_0 m_\chi \mu^2} F^2(Q) e^{-\frac{Q m_N}{2\mu^2 v_0^2}} \quad (2.11)$$

Equation 2.11 emphasizes just how crucial the lower energy threshold of a dark matter detector is to the prospects of discovery. Figure 2-2 presents the energy-integrated WIMP interaction rate vs. the low energy threshold of a dark matter detector for Xe and Ge, assuming perfect detector efficiency, a WIMP mass of 70 GeV, and a spin-independent cross-section $\sigma_{SI} = 10^{-44} \text{ cm}^2$.




Figure 2-2: Energy-integrated WIMP interaction rate vs. the low energy threshold of a dark matter detector for Xe and Ge, assuming perfect detector efficiency, a WIMP mass of 70 GeV, and a spin-independent cross-section $\sigma_{SI} = 10^{-44} \text{ cm}^2$. Figure from [52]

2.4.2 Event Rates - Particle physics

This section is a brief presentation of the characteristics of σ_0 ; [53], [27], or [51] are suitable references for more in-depth treatments. The spin-independent (scalar) WIMP-nucleon cross-section in the limit of zero momentum transfer can be written as:

$$\sigma_{SI}^0 = \frac{4}{\pi} \mu^2 [Z f_p + (A - Z) f_n]^2 \quad (2.12)$$

where μ is again the reduced mass of the WIMP-nucleus system; A and Z are the atomic mass and number, respectively; and f_p and f_n are factors for proton and neutron scattering, respectively. This coherent addition of the scattering amplitudes of individual nucleons causes the spin-independent cross-section to scale like A^2 under the assumption that $f_p \sim f_n$, which is true for a wide range of WIMP models. Relaxation of this assumption can drastically change how different experiments compare with each other, especially when $f_p \approx -f_n$, in which case cancellations appear depending on the isospin content of the nucleus; see [54] for more.

In the spin-dependent scattering case, we have

$$\sigma_{SD}^0 = \frac{32(J+1)}{\pi J} G_F^2 \mu^2 [a_p \langle S_p \rangle + a_n \langle S_n \rangle]^2 \quad (2.13)$$

where J is the nuclear spin, and $\langle S_p \rangle$ and $\langle S_n \rangle$ are the expectation values of the proton and neutron spins, respectively. These factors must be obtained from nuclear spin structure calculations. In general, $a_p \approx a_n$.

The form factor $F^2(q)$ used to set spin-independent WIMP-nucleon cross-section

constraints in this work is the Helm form factor [51]

$$F(q) = 3 \frac{j_1(qr_n)}{qr_n} e^{-q^2 s^2/2} \quad (2.14)$$

where s is a measure of the nuclear skin thickness and r_n is the nuclear radius. When computing limits in this work, I take the following from [51]: $r_n^2 = c^2 + \frac{7}{3}\pi^2 a^2 - 5s^2$, $c = 1.23A^{1/3} - 0.6$ fm, $a = 0.52$ fm, and $s = 0.90$ fm.

Because $a_p \approx a_n$, momentum-dependent effects do not easily factor out of the differential spin-dependent cross-section, and so three spin-dependent form factors are necessary; those used in this work can be found in [55].

2.4.3 Detector Media

Liquid Nobles

The use of liquid noble elements for particle detection was first proposed in 1970 in [56], in which the authors proposed a two-phase (liquid/gas) time projection chamber for tracking ionizing particles. In two-phase liquid noble detectors, the dense liquid phase is used as the target material. When a recoil occurs in the liquid, the energy of the recoil is divided into primary scintillation light and ionization channels. The ionized electrons are then drifted in an applied electric field to the surface of the liquid and are then stripped from the liquid surface into the gas using a relatively higher electric field. These electrons are then accelerated in a third field applied in the gaseous region to produce proportional scintillation (or secondary scintillation), which provides a measurement of the total ionization signal. The details about the relative strength of the three fields, the electron extraction efficiency from the liquid surface, the secondary light yield per electron, etc. vary depending on the particular

liquid noble gas employed.

Single-phase liquid noble detectors, in which only the primary scintillation signal is measured, are also viable dark matter detectors. ER/NR discrimination in these detectors is provided by the shape of the scintillation pulse in time. The primary scintillation arises from the formation of excited diatomic molecules, which emit a photon when breaking back into two atoms. These diatomic molecules can exist in a triplet or a singlet state, and the decay time constants of these two states can be measurably different. ERs and NRs produce the triplet and singlet state dimers in different ratios, allowing the scintillation light profile to be used to discriminate the two recoil types.

Though originally proposed as tracking detectors, dual-phase liquid noble detectors offer a number of properties that make them good dark matter search detectors. The dual-phase operation of noble liquid detectors allows very strong rejection of background ERs through the ratio of primary to secondary scintillation; pulse shape/timing information can also be used for ER/NR discrimination. With high scintillation yields, the energy thresholds of noble liquid detectors can reach the keV scale. Finally, the position reconstruction available in dual-phase liquid noble TPCs enables the definition of an active fiducial volume, which provides self-shielding against external radiation backgrounds.

The primary disadvantage associated with most liquid noble detectors is the presence of radioactive elements that may be distributed throughout the target volume. Natural xenon contains no long-lived xenon radioisotopes but does contain radioactive ^{85}Kr . Argon-based detectors face the challenge of removing the long-lived radioisotope ^{39}Ar . Natural neon is non-radioactive, but its low mass reduces its sensitivity in searches for high-mass WIMPs.

Dark matter searches using a variety of liquid noble elements have been demon-

strated (indeed, at the time of this writing, the XENON100 experiment holds the world's leading limit on the WIMP-nucleon cross-section at WIMP masses above $8 \text{ GeV}/c^2$), and scaled-up experiments with fiducial masses on the order of 1 ton are currently planned, including the CLEAN experiment using neon; the ArDM, WARP, and DEAP experiments using argon, and the XENON, Zeplin, LUX, and XMASS experiments using xenon. For a recent, albeit brief, review of current liquid noble dark matter detection experiments, please see [57]

Cryogenic Semiconductors

Crystalline semiconductors are characterized by the existence of a band of forbidden electron states between the valence band, in which electrons are bound to lattice nuclei, and the conduction band in which electrons are free to propagate. This forbidden band produces a number of interesting properties, including strong temperature dependence of the conductivity (in certain temperature regimes) and the ability to drastically affect the material's conductivity through the addition of a relatively small number of dopant atoms. In intrinsic (undoped) semiconductors at cryogenic temperatures, there are effectively zero free charge carriers because $kT \ll E_g$, where k is Boltzmann's constant, T is the temperature, and E_g is the difference in energies between the highest-energy valence state and the lowest-energy conduction state. The semiconductor thus behaves as an insulator, and a relatively large electric field can be applied across the semiconductor without inducing a current (unless the field is sufficiently large to induce breakdown).

Radiation incident on a semiconductor can transfer energy to bound electrons, elevating them to the conduction band, and if an electron (or hole) can traverse the crystal volume without trapping or recombining with a hole (electron), then this ion-

ization signal can be measured by ionization sensors attached to the semiconductor, making cryogenic semiconductors useful radiation detectors.

Cryogenic semiconductor detectors include crystalline Si, Ge, or CaWO_4 detectors cooled to cryogenic (sub-K) temperatures, giving them excellent energy resolution and thresholds. Because the heat capacity of a semiconductor scales as T^3 , deposition of small amounts of energy ($\mathcal{O}(\text{keV})$) into large crystals can produce measurable heat signals either through athermal phonons (CDMS, CRESST) or thermal phonons (Edelweiss). Because the crystals employed are highly pure and thus have low intrinsic carrier densities, the charge carriers are frozen out at cryogenic temperatures, and an ionization signal can act as a second energy measurement channel. Additionally, CaWO_4 is a natural scintillator, and the scintillation light provides another measurement channel for experiments employing CaWO_4 crystals (CRESST). As with liquid nobles, the two complementary energy measurements provide a method of discriminating NR signals from ER backgrounds on an event-by-event basis through the ratio of the ionization/scintillation to phonon energies of an event.

Solid Scintillators

Alkali metal halide crystals (e.g., NaI, CsI, CaWO_4) scintillate when exposed to ionizing radiation, and doping with other elements, commonly thallium, can increase the scintillation efficiency. Radiopure solid scintillators coupled to light guides and photomultiplier tubes can achieve recoil energy thresholds corresponding to just a few scintillation photons and have been applied in a wide variety of fields. The DAMA/LIBRA experiment [58] experiment is probably the most well-known solid scintillator dark matter search, utilizing 25 solid NaI(Tl) scintillator crystals with a total mass of $\sim 250\text{kg}$ to search for dark matter interactions. DAMA/LIBRA achieves

a lower energy threshold of 2 keV electron equivalent (the amount of scintillation light produced by a 2 keV electron recoil), and the 25 separate crystals can act as multiplicity vetoes, as dark matter is expected to only interact in a single detector. Without ER/NR discrimination, DAMA/LIBRA and other experiments like it rely on the time-dependence of the event count rate and energy to search for dark matter. Photon pulse shape discrimination is a potential means of distinguishing electronic and alpha backgrounds from the nuclear recoil signal (see [59] for an example by the KIMS collaboration using CsI(Tl) crystals).

Ionization Detectors

Ionization-only detectors based on semiconductor or liquid noble technology have primarily been used to search for neutrinoless double beta decay but can also be used to search for WIMP-nucleon interactions. I separate these into their own category because ionization-only detectors do not allow strong ER/NR discrimination (though pulse-shape discrimination may be possible) and thus rely on detecting a dark matter signal using the background-subtracted recoil spectrum or a time-dependent signal. The COGENT [60] experiment is one example of an experiment utilizing this type of detection channel, and other experiments primarily designed for neutrinoless double-beta decay search, such as MAJORANA [61], could also probe for astrophysical dark matter.

Bubble Chambers

Standard tools in particle physics for decades, the bubble chamber has been repurposed for dark matter search by the COUPP [62], PICASSO [63], and SIMPLE [64] experiments. A liquid target is superheated and compressed, and when sufficient

energy density is deposited in the target, a bubble is nucleated. The conditions of the chamber are set such that electromagnetic backgrounds are incapable of producing sufficient energy densities to nucleate bubbles, rendering these detectors blind to the electromagnetic backgrounds that must be actively discriminated against in many other detector media. A WIMP recoiling in the detector would be expected to produce a single bubble, while neutrons exhibit some degree of multiplicity; the neutron signal can be moderated by passive shielding, low-radioactivity materials, and an active muon veto to tag cosmogenic muons that can produce neutrons in the detector space. Transducers record the acoustic signal from bubble nucleation, and the frequency content can be used to distinguish nuclear recoils from alpha particles that recoil in the detector.

Gas Detectors

Unique among these detection media, gaseous detectors offer the ability to measure the direction of the recoiling nucleus in a WIMP-nucleon scatter, as recoiling particles deposit energy over macroscopic length scales due to the lower density of the target material relative to liquid/solid detectors. The relevance of directional information is explained in Section 2.4.4. The DMTPC experiment [65] utilizes CF_4 gas as a target material and CCDs, PMTs, and an electronic amplification mesh to measure scintillation and drift electron signals. A prototype run at the surface has demonstrated the ability to recover the energy, direction and ‘head/tail’ sense of nuclear recoils. Because the energy and track length observed in a recoil are correlated with the species of the incident particle, electromagnetic and alpha backgrounds can be strongly suppressed through a cut requiring the energy and length to be consistent with a recoiling C or F nucleus. The inclusion of fluorine also provides sensitivity to

spin-dependent interactions.

2.4.4 Detection Channels

Zero-background

The conceptually simplest dark matter detection technique is to shield or discriminate backgrounds to a level such that less than 1 background event is expected and to consider all remaining signal events to be WIMP candidates. In practice, this goal is achieved by employing detectors with strong ER rejection and placing them into environments with extremely low rates of neutrons incident on the detectors. As detailed above, a variety of technologies offer very strong ER rejection, and the incident neutron flux can be reduced to extremely low levels by locating the search experiments in deep-underground laboratories, surrounding them with active veto systems and passive shielding, and constructing the experimental apparatus using radiopure materials.

Annual Modulation

Because the WIMP-nucleon scattering rate and spectrum depends on the relative velocity of the WIMP and the target, the orbit of the earth around the sun should result in a modulation of the WIMP-nucleon scattering rate with a period of one year. In effect, the earth travels through a WIMP ‘wind’ directed against the direction of rotation of the sun around the galactic center. The rate of WIMP-nucleon scatters is then maximized when the earth’s orbital velocity is antiparallel or ‘into’ the WIMP wind and minimized when the orbital velocity is parallel to the WIMP wind. The detection of an annually modulated signal in dark matter detectors has been reported by both the DAMA [66] and COGENT [67] experiments. However, these results are

in tension with null results reported by other experiments [68, 69].

A similar argument can be made for a daily modulation in the WIMP-nucleon interaction rate. A detection of daily modulation in the WIMP-nucleon event rate would require significantly larger event rates than are accessible by current experiments given the current limits on the WIMP-nucleon cross-section. However, because this modulation would be relative to the sidereal day rather than the earth day, backgrounds that could mimic the daily modulation signal would necessarily be astrophysical rather than earthbound in origin.

Directional Search

The motivation for directional detection searches is similar to that for annual modulation search - because the earth is traveling through a WIMP wind, the distribution of incoming WIMP velocities should be strongly directed against earth's velocity around the galactic center. It is also to be expected that the direction of incoming WIMPs relative to the experimental setup is similarly modulated both annually and diurnally. Potential backgrounds that could exhibit diurnal or annual modulation (e.g., due to fluctuations in temperature, humidity, etc.) can be expected to be independent of the earth's velocity relative to the WIMP wind. Thus, observation of a WIMP signal in a directionally-sensitive detector should provide clear evidence that the signal is in fact a detection of astrophysical dark matter.

Current State of Direct Dark Matter Detection

Figure 2-3 presents a set of current limits and regions of interest in the WIMP mass and cross-section plane from a variety of direct dark matter search experiments. Tension between different experiments is evident, as the best upper limits from

XENON100, CDMSII, ZEPLIN-III, and EDELWEISS rule out the closed contours from the CRESST-II, DAMA, and CoGeNT results under ‘standard’ hypotheses about the dark matter distribution: a Maxwellian astrophysical velocity distribution and a spin-independent, isospin-independent WIMP-nucleon cross-section. The theoretical community has proposed a number of hypotheses that can bring different experiments into better agreement. It is certainly an exciting time to be working in direct dark matter search!




Figure 2-3: Recent limits and regions of interest in the WIMP mass and cross-section plane from a variety of direct dark matter search experiments. Figure from [70]

Chapter 3

Physics of the CDMS Detectors and the CDMS Detector Monte Carlo

3.1 Introduction

The CDMS Detector Monte Carlo (DMC) was developed primarily by CDMS collaborators at MIT, UC Berkeley, and SLAC, and grew as an extension to Steve Leman’s doctoral thesis at Stanford, entitled “Development of Phonon-Mediated Transition-Edge-Sensor X-Ray Detectors for use in Astronomy”. I view my role in the development of the DMC from a phonon and TES simulator to an integrated tool for modeling CDMS detector physics, from GEANT4 event distributions as input to CDMS Analysis Package signals as output, as perhaps my most important contribution to the CDMS collaboration. The DMC aims to provide an integrated CDMS event simulation package, combining phonon and charge propagation simula-

tions with sensor models and automated conversion of outputs into CDMS Analysis Package format. In this chapter, I will discuss the physics of the CDMS detectors and describe the corresponding physical models utilized by the DMC.

It is important to note that the DMC currently treats only germanium-based detectors. Adding descriptions of silicon-based detectors is a priority target for future DMC development. In the following, I describe the physics of the CDMSII and SuperCDMS detectors, usually discussing both Si and Ge when appropriate. When discussing the implementation of these physics in the DMC, it is important to keep in mind that the DMC only models germanium, and any numbers that appear in the text of these discussions can be assumed to apply only to germanium-based detectors.


3.2 Basic design of the CDMSII and SuperCDMS Detectors

The CDMSII detectors, known as Z-sensitive Ionization and Phonon (ZIP) detectors, are cylindrical germanium and silicon ‘pucks’ 1 cm thick and 3” in diameter. Concentric ionization collection lines are patterned on one of the flat surfaces of the detector, with an inner electrode covering $\approx 85\%$ of the detector surface and an outer electrode covering the remainder. On the opposite surface are four phonon sensing channels, each containing 1036 Transition Edge Sensors (TESs), which are wired in parallel. Phonons are collected from the crystal by superconducting aluminum fins (quasiparticle traps) that overlap with the TESs but cover significantly more detector surface area. A drift field is applied across the detector by maintaining a voltage difference between the ionization sensors on one surface and the phonon sensors on the other (-3V for Ge, -4V for Si detectors). Figures 3-1(a) and 3-1(b) present a

picture and a schematic of the CDMSII detectors, respectively. The detectors were deployed approximately 1/2 mile underground in the Soudan Underground Laboratory from 2001-2009; for more details of the experimental setup and data analysis, see 7.



(a) Figure 3-1: Picture of a CDMSII ZIP detector. The phonon sensor side is shown in the image.



(b) Schematic of a CDMSII ZIP detector, illustrating the 4 TES channels on the top and the 2 ionization channels on the bottom.

The primary discriminator between ER backgrounds and NR signal events is the ratio between the measured ionization signal and the recoil energy of the event measured using the phonon signal, a quantity known as the ionization yield. Figure 3-2 presents the ionization yield vs. recoil energy of ERs and NRs and clearly shows that these two populations are strongly separated in this space. The physics underlying the different ionization yields of ERs and NRs is described in Sec. 3.3. However, Figure 3-2 also contains a population of events marked in black that exhibit yields below those of well-measured ERs, some of which fall into the NR signal band. The charge carriers generated in an ER initially have high kinetic energies, and when an ER takes place sufficiently near one of the two surfaces, some of these carriers can overcome the applied field to be collected at the ‘wrong side’, thus reducing the measured

yield. These surface events represented the primary background in the CDMSII experiment. Though surface events can be discriminated from true NRs using phonon pulse timing information, this discrimination comes at a significant cost in NR signal efficiency, and the surface event background represented a prohibitive hurdle to scaling CDMSII ZIP detector technology to the exposures targeted in next-generation direct dark matter searches.

The SuperCDMS detectors are based upon the same basic idea as the CDMSII detectors, but are fabricated with interleaved phonon and ionization sensing lines on each side of the detector, and are known as interleaved ZIPs (iZIPs). In this detector, the drift field is applied by biasing the ionization sensors on opposite sides of the detector to opposite voltages (+2V and -2V on the two sides, for example) relative to the phonon sensors, which are biased with very low voltages and thus act as grounds. This layout produces a crucial difference in the structure of the drift field: in the bulk, charges will drift to ionization sensors on opposite sides of the detector, but near the surfaces (≈ 2 mm), the interleaved biased ionization rails and grounded phonon rails create a complex, scalloped local drift field that directs one species of charge into the ionization sensor and the other species into the phonon rail on the same side. This layout creates an effective veto for scatters that occur near one of the two surfaces, as the ionization sensors on the opposite surface will register no charge signal. This surface event veto solves the primary background problem observed in the CDMSII experiment, thus paving the way for cryogenic germanium detector technology to continue to remain competitive in the field of direct dark matter detection as experiments push forward to total exposures on the scale of ton-years. Pictures and layout illustrations of the iZIP detectors are shown in Figs. 4-6(b), 4-6(a), 4-17(b), and 4-17(a).




Figure 3-2: Yield-based discrimination of electron recoils and nuclear recoils: Bulk electron recoils are shown in red, in a band around a yield of 1, and bulk nuclear recoils are shown as blue circles in a reduced, energy-dependent yield band. Both bands are indicated in black solid lines. However, surface electron recoils, shown as black crosses, exhibit reduced yield and can contaminate the WIMP signal region; the surface-event “band” is shown as a black dash-dotted line. The blue and magenta lines represent recoil energy and total charge thresholds, respectively.

3.3 Physics of Electron and Nuclear Recoils

The primary discrimination power of the CDMS detectors arises from fundamental differences in the interactions of electromagnetically-interacting backgrounds (γ and β particles) and electrically neutral particles (presumably WIMPS, though neutrons also represent a background) with the detector material.

3.3.1 Gamma Interactions

Exchange of energy between incident gamma particles and the detector crystal proceeds through three primary channels: photoelectric absorption, Compton scattering, and e^+e^- pair production. The interaction probability between an incident gamma is generally parametrized using an exponential distribution:

$$P(x)dx = \frac{1}{\lambda}e^{-\frac{x}{\lambda}} \quad (3.1)$$

where λ is known as the penetration depth and has units of distance. The value of the penetration depth at a particular γ energy is fundamentally related to the cross-sections of the three interaction channels mentioned above, which are described below in more detail.

The photoelectric effect occurs when an incident photon is absorbed by an electron in the material, and the electron receives all of the energy of the incident photon. Some material-dependent amount of this energy is dissipated in liberating the electron from the atom; this is known as the work function and tends to be on the order of 1 – 6 eV. Higher-energy photons can liberate inner shell electrons as well as outer shell electrons, and the liberation of inner shell electrons can result in a secondary photon or electron; this phenomenon is known as the Auger effect. Any remaining

energy becomes the kinetic energy of the freed electron. The cross-section of photoelectric absorption decreases rapidly as E_γ increases, but dominates in germanium in the < 100 keV gamma energy range.

Compton scattering describes the transfer of energy from a photon to an electron through an elastic scattering process. Similar to the photoelectric effect, the kinetic energy of the electron is equal to the energy transferred from the photon minus the binding energy of the material. In germanium, the Compton cross-section becomes comparable to the photoelectric cross-section around 100 keV and is the dominant interaction at gamma energies ranging from a few hundred keV to a few MeV (above which pair production dominates).

Pair production occurs when a high-energy photon ($E_\gamma > 2m_e = 1.022$ MeV) in matter converts into an electron-positron pair. This process is forbidden in vacuum by momentum conservation: a rest frame can always be found in which the e^+e^- pair will have zero net momentum, which is not possible for a single γ in vacuum. The cross-section for pair production increases as E_γ increases. However, because the WIMP cross-section generally sharply decreases with recoil energy, a WIMP search analysis will usually place an upper energy cut well below the threshold for pair production, and this process does not contribute significantly to our electromagnetic background. Fig 3-3 shows the cross-sections for these three processes vs. the γ energy in silicon and germanium.

The electrons freed from atomic binding through any of these interactions with a photon are then free electrons in the material and are rapidly stopped by the processes detailed in Sec. 3.3.2.




Figure 3-3: Gamma photoelectric, Compton, and pair production cross-sections vs. gamma energy in Ge and Si. Fig. adapted from [71]

3.3.2 Electron Interactions

Electrons at energies below a few MeV interact in matter primarily through Coulomb interactions with other electrons in the material, resulting in ionization and excitation of atomic electrons in the material. At energies above a few MeV, which is significantly above the energy range of interest to a WIMP dark matter search, bremsstrahlung processes become important as well. Due to the large difference in the masses of the electron (511 keV) and nuclei (multi-GeV scale), the electron interacts almost exclusively with the electrons in the material. The cross-section of these electron-electron Coulomb interactions are sufficiently high that the range of electrons below 100 keV is tens of microns or less (see Fig. 3-4) and falls rapidly with decreasing energy. The new electrons freed by the interactions with the incident electron then shed their own kinetic energy in the same processes [71]. The electrons can also shed kinetic energy to the crystal lattice in the form of vibrational excitations known as phonons.

After the dust settles, incoming β or γ radiation in a semiconductor detector ends up creating electron-hole pairs in a number directly proportional to the recoil energy. The average energy required to create an electron-hole pair E_{eh} in cryogenic germanium is 3 eV (3.96 eV in Si). However, the forbidden bandgap of germanium is only .75 eV (1.17 eV in Si), and the newly created electron-hole pairs rapidly shed the difference in these two energies, 2.25 eV per charge carrier pair (2.80 eV in Si), to the crystal lattice as phonons. Thus, an electron recoil of energy E will produce $N_{eh} = \frac{E}{E_{eh}}$ electron-hole pairs and shed $\approx \frac{3}{4}E$ of the recoil energy into the phonon system. This energy-independent (at least in the energy range of interest) constant ratio of energy partitioning into the charge carrier and phonon systems can be used to distinguish electromagnetic backgrounds from the expected nuclear recoil signal,




Figure 3-4: Electron range versus energy in germanium and silicon. Fig. adapted from [71]

which will be discussed in Sec. 3.3.3.

3.3.3 Nuclear Recoil Interactions

While gammas and betas transfer energy to the detector crystal through electromagnetic interactions with the electrons in the crystal, other particles, including neutrons, WIMPs, and α particles will deposit energy primarily through electromagnetic or nuclear interactions with the atomic nuclei making up the crystal lattice. If the lattice binding energy is less than the kinetic energy transferred to the recoiling nucleus, then the nucleus will free itself from the lattice and propagate through the crystal. From this point, we can use the theory of Lindhard et al. [72], which describes the transfer of energy of a propagating ion to the electronic and phonon systems of a medium of the same Z and A as the ion. Unlike the previously described case of electron recoils, the ratios of the energies in the charge and phonon systems to the initial recoil energy are energy-dependent for nuclear recoils. A simple argument for this comes from considering the processes by which energy from a propagating ion can be transferred to the lattice ions and the electrons in a material[73]. The propagating ion interacts with the lattice ions through Rutherford scattering in a screened Coulomb potential to first order, with $\frac{dE}{dx} \propto v^{-2}$. In interactions with the bound electrons, the interaction rate is determined by the number of states accessible to the electrons, which scales as v , and thus $\frac{dE}{dx} \propto v$. Thus, the fraction of ionization energy per unit recoil energy increases as the energy of the nuclear recoil increases.

3.3.4 Modeling ERs and NRs in the DMC

The CDMS Detector Monte Carlo does not model the passage of particles through the detector media. Rather, input events are generated using a GEANT4 model that

describes the experimental geometry [74] and the appropriate radioactive sources: the calibration sources employed by the CDMS experiment (including ^{133}Ba for high-energy γ s, ^{252}Cf for neutrons, ^{109}Cd for β s and low-energy γ s, ^{210}Pb for β s) and known radiogenic isotopes in various materials used to construct the experiment (including ^{40}K , ^{60}Co , and a series of radionuclides from the ^{232}Th and ^{238}U decay chains). The GEANT4 simulations produce distributions of event positions, energies, and types (ER vs. NR) in the detector, which are used as input to the DMC.

Furthermore, the DMC does not model the stopping of the high-energy electrons or ions released in scatters with the incident particle. Rather, at the site of each scatter, a corresponding amount of energy is partitioned among the charge carrier and phonon systems. In the case of electron recoils, one electron-hole pair is created per 3 eV of recoil energy; the user can enable or disable fluctuations in the number of electron-hole pairs, and Fano statistics is employed if fluctuations are enabled. 0.75 eV of this energy is utilized to promote electrons and holes across the forbidden bandgap of germanium, and the remaining 2.25 eV of recoil energy is partitioned among direct phonon creation and the kinetic energies of the charge carriers; any significant amount of kinetic energy (above $\approx 90 \mu\text{eV}$, when the electron velocity exceeds the speed of sound in Ge) will be rapidly dissipated into phonons created through carrier-lattice interactions described in Sec. 3.6. It is worth noting here that for computational reasons, not all of these charge pairs needs to be simulated; scaling factors can be employed to effectively simulate a smaller number of “super-electrons” and “super-holes”, each of which represents multiple real electrons and holes (a corresponding scaling can be applied to the phonon system as well).

In the case of nuclear recoils, the fractions of the recoil energy that are partitioned into the charge carrier and phonon systems depend on the energy of the recoil (approximately equal to the kinetic energy of the ion displaced in the interaction).

Lindhard theory describes the partitioning of energy into the phonon and charge systems, and [72, 51] provide a simplified equation describing the ionization yield (the ratio of energy given to charge carriers to the total recoil energy) as a function of recoil energy:

$$y(E_R) = \frac{kg(\epsilon)}{1 + kg(\epsilon)} \quad (3.2)$$

where

$$\epsilon = 11.5E_R(keV)Z^{-7/3} \quad (3.3)$$

$$k = 0.133Z^{2/3}A^{-1/2} \quad (3.4)$$

$$g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon \quad (3.5)$$

where Z and A are the atomic and mass numbers, respectively, of the detector medium. Figure 3-5 shows the yield vs. recoil energy curves predicted by this parametrization for Si and Ge ionization detectors along with some measured yield curves. The disagreement between measurement and Lindhard theory has been the subject of much study within the collaboration, as they potentially indicate an issue with the reconstruction of the nuclear recoil energy scale. A paper treating this phenomenon is in preparation. The DMC uses the Lindhard functional form to determine the energy deposited into the charge system, $y(E_R) * E_R$, and creates 1 charge pair per 3 eV of energy into the charge system - as before, .75 eV of this is bandgap energy and 2.25 eV is partitioned between carrier kinetic energy and phonon creation. The remaining energy, $E_R * [1 - y(E_R)]$, is entirely utilized for phonon creation.

Thus, the input to the DMC consists of a set of locations and energies of ER




Figure 3-5: Predicted and measured yield vs. energy curves for nuclear recoils in Si and Ge detectors.

or NR scatters, and the simulation truly begins after placing initial populations of charge carriers and phonons at each of the scatter locations. The upcoming sections will review the physics of phonon and charge propagation and interactions with the sensors along with descriptions of how the DMC models these physical processes.

3.4 Phonon Physics

Phonons are propagating, normal-mode vibrational excitations within a periodic atomic lattice. Phonons in the DMC can be created in multiple ways. ‘Prompt’ phonons are created at the location of electron or nuclear recoils - ERs create a number of high energy charge carriers that rapidly shed the 2.25 eV difference between the e-h pair creation energy (3 eV) and the bandgap energy (0.75 eV), while NRs create phonons directly through screened Coulomb interactions between the lattice atoms and the propagating ion dislocated by the recoil. Phonons can also undergo a process known as anharmonic decay in which a single phonon decays to two phonons of lower energies. ‘Luke-Neganov’ [75] phonons are created as electrons and holes propagate through material in an external electric field; when a carrier’s velocity exceeds the speed of sound in the material, Luke-Neganov phonons are created in a process similar to Cherenkov radiation. Finally, ‘recombination’ phonons may be created when electrons and holes are collected at the ionization electrodes patterned on the detector surfaces, returning the .75 eV bandgap energy to the crystal. This section discusses the physics of phonon propagation in the crystal; interactions with the instrumented surfaces of the detector will be described in Sec. 3.5.

3.4.1 Phonon Basics

Phonons are quantized normal-mode vibrational excitations in periodic lattices. A basic derivation of phonon energies and eigenstates in a 3-D, N-atom harmonic lattice follows from a relatively straightforward extension of the standard quantum harmonic oscillator and is not reproduced here; interested readers can refer to Chapter 23 and Appendix L of [76]. Anisotropies in the geometry of the atomic lattice can cause phonon phase velocities to depend on the direction of propagation and the polarization mode of the phonon. In general, the phonon dispersion relation can be written as

$$\rho\omega^2\epsilon_\mu = \sum_\tau \left(\sum_{\sigma\nu} c_{\mu\sigma\nu\tau} k_\sigma k_\nu \right) \epsilon_\tau, \quad (3.6)$$

where ρ is the crystal's mass density, ω is the phonon frequency, ϵ_μ is a component of the polarization vector ϵ , $c_{\mu\sigma\nu\tau}$ is the stiffness tensor, and k_σ is a component of the phase velocity vector $\mathbf{k}$.

All indices run over x, y, z . The definition of $c_{\mu\sigma\nu\tau}$ in terms of the stress and strain tensors, in addition to symmetry, energy, and torque considerations, allow $c_{\mu\sigma\nu\tau}$ to be reduced to C_{IJ} , where I, J range from 1 to 6 and are mapped to (i, j) , (l, m) , respectively, as follows: $1 = xx$, $2 = yy$, $3 = zz$, $4 = yz$, $5 = xz$, $6 = xy$. Furthermore, the symmetries of a cubic lattice reduce the number of stiffness constants further, to a set of only 3: $C_{11}(= C_{22} = C_{33})$, $C_{12}(= C_{13} = C_{23})$ and $C_{44}(= C_{55} = C_{66})$, with all remaining $C_{IJ} = 0$. The remaining eigenvalue equation determines the phonon phase and group velocities through the standard relations $v_\phi(\theta, \phi) = \frac{\omega(\theta, \phi)}{k}$ and $v_g = \frac{d\omega(\theta, \phi)}{dk}$, and lead to interesting phenomena such as non-spherical phase velocity surfaces and strong phonon focusing; see [77, 78].

3.4.2 Phonon Interactions

After the initial creation mechanism, phonons undergo three primary types of interactions while propagating through the crystal lattice: isotope scattering, anharmonic decay, and absorption/reflection at the detector surfaces. Both isotope scattering and anharmonic decay exhibit strongly frequency-dependent rates; anharmonic decay is the more common process at high frequency, while isotope scattering dominates at lower frequencies. It is worth noting here that we often utilize a conceptual distinction between two regimes of phonon propagation: quasidiffusive propagation and ballistic propagation. Quasidiffusive propagation describes the propagation of high-energy phonons, which frequently undergo one of these two scattering processes and thus follow effectively random walks with frequent changes in direction (and energy, in the case of anharmonic decay). Ballistic propagation, on the other hand, describes the travel of phonons whose energies are sufficiently low that they are likely to traverse the detector basically along a straight line, i.e., they are unlikely to undergo either interaction within the time to travel across the detector. Separating phonon behavior into these two regimes thus does not represent any fundamental behavior, rather, the frequency separating quasidiffusive and ballistic propagation is set by the geometry of the detector. Because anharmonic decay creates two low-energy phonons from one high-energy phonon, and because energy is constantly removed from the phonon system by sensors patterned on the surface, many phonons propagate quasidiffusively immediately after a recoil, but after a short time, the majority of phonons in the crystal will propagate ballistically, and the crystal will rapidly thermalize.

Isotope Scattering

Isotope scattering occurs when a propagating phonon encounters an isotopic atom with a different mass number than the bulk medium. These isotopes represent imperfections in the crystal lattice, referred to as mass defects, and the presence of these defects can alter the direction of propagation and the polarization state of incident phonons. Because naturally occurring germanium is a mixture of 4 stable isotopes (^{70}Ge , 21.23%, ^{72}Ge , 27.66%, ^{73}Ge , 7.73%, ^{74}Ge , 35.94%) and 1 long-lived isotope (^{76}Ge , 7.44%, decays through double β decay with a $1.78 * 10^{21}$ year half-life) [79], isotope scattering plays a major role in determining the mean free path of propagating phonons.

Isotope scattering in germanium depends only on the frequency of the phonon. In [80], the scattering rate is predicted to be

$$\Gamma_{B,Ge} = 3.67 * 10^{-41} [s^3] \nu^4 \quad (3.7)$$

The isotopic scattering rate for individual phonons is

$$\gamma \propto \frac{|\vec{e}_{\lambda} \cdot \vec{e}_{\lambda'}|^2}{\nu_{\lambda'}^3} \quad (3.8)$$

where λ and λ' represent the initial and final phonon, $\vec{e}$ is the polarization vector, and ν is the phonon frequency. Thus, isotope scattering allows transitions between longitudinal and transverse polarization modes, thus determining the equilibrium phonon mode distribution.

Anharmonic Decay

Phonons are also able to decay into lower-frequency phonon modes through perturbative anharmonicities in the lattice; that is, the lattice atoms do not actually act as perfect harmonic oscillators and thus propagating phonons do not represent perfect normal modes of the lattice system. As would be expected, the deviation from harmonicity, and thus the rate of decay of phonons into lower-energy modes, is strongly dependent on the frequency of the excited phonon mode [81, 82].

$$\Gamma_{A,Ge} = 1.61 * 10^{-55} [s^4] \nu^5 \quad (3.9)$$

The constant prefactor in $\Gamma_{A,Ge}$ represents a mode-averaged anharmonic decay constant, as the transverse anharmonic decay rate is negligible compared to that of the longitudinally polarized phonons.

With these two rates in hand, and with the mode-averaged speed of sound in Ge of 5.4×10^5 cm/s, we can outline the behavior of phonons over time. Comparing the constant prefactors in Eqs. 3.7 and 3.9, we can see that isotope scattering dominates anharmonic decay at all frequencies below ≈ 100 THz, but anharmonic decay still plays a very important role, as each decay of a phonon into two lower-energy phonons significantly increases the isotope scattering path length. The path lengths and scattering/decay rates vary extremely rapidly with frequency, but we can divide the phonon lifetime into a few regimes. First, any initial, high-energy phonons with $\nu \geq 2$ THz have exceedingly short path lengths ($\frac{v}{\Gamma_B} \leq 10\mu m$), so phonons scatter very often and travel very short distances from the interaction point before downconversion ($\Gamma_A \geq 5\mu s^{-1}$). Below 2 THz, the path length of isotope scattering grows rapidly, and the phonon ‘ball’ expands quasidiffusively (random walk with growing step sizes) outward from the recoil location, with anharmonic

downconversion becoming increasingly infrequent. Finally, at a frequency around 0.3 THz, the path length of isotope scattering increases to the cm-scale, on par with the detector’s dimensions, and the rate of anharmonic decay has decreased well below the rate of phonon absorption at the surfaces, so phonons below these frequencies are said to propagate ballistically.

Interactions with detector surfaces

When phonons strike the surfaces of the detector, they are either reflected away or absorbed at the surface. At the instrumented surfaces, phonons can be collected into aluminum quasiparticle traps, which will be described further in Sec. 3.5.1, at which point they are removed from simulation. However, most of the detector surfaces are not instrumented, and phonons can be reflected from the surfaces a number of times before absorption. Thus, after some time has elapsed since the initial recoil ($\mathcal{O}(100)$ μs), information about the initial position of the events has effectively been lost, as most of the phonons have downconverted to energy scales at which propagation is ballistic and have undergone numerous reflections from the detector surfaces. At this point, the detector bulk is filled with a basically uniform ‘gas’ of phonons, and the rate of phonon absorption at all instrumented surfaces is well-described by a featureless falling exponential. The decay constant of this exponential is useful in determining the reflection coefficient at the Ge-aSi-Al surface, as will be described in Sec. 4.1.

3.4.3 Phonon Creation, Propagation and Interactions in the DMC

Phonon creation in the DMC is simple: all recoil energy that is not apportioned to charge carriers in the form of gap energy or initial kinetic energy is used to create $N = \frac{E}{E_{deBye}}$ phonons. The gap energy is always 0.75 eV per charge pair in Ge, and the kinetic energy initially given to the charge carriers is a tunable parameter. In electron recoils, 2.25 eV of recoil energy per carrier pair remains after the, and the default is generally to give this entire energy to the carrier kinetic energy and allow phonons to be created in carrier-lattice interactions. As nuclear recoils exhibit reduced charge production, even if the carriers receive 2.25 eV per pair of kinetic energy, a great deal of recoil energy remains and is used to create initial phonons. The initial phonon frequency distribution in the DMC is a delta function at the Debye frequency of 2 THz; any distribution of high-energy phonons would do, as the process of anharmonic decay causes the phonons to rapidly downconvert to lower energies, thus destroying any information about the initial high-energy distribution. The initial phonon mode distribution is simply the equilibrium distribution, as isotope scattering is also frequent at high phonon frequencies. Both of these simplifications rely on the fact that the detector itself is much larger than the ‘quasidiffusive ball’ of initial phonons, and so these downconversion and mode equilibration processes occur on shorter timescales than phonon absorption at the instrumented surfaces.

Determining the group and phase velocities of each phonon by solving Equation 3.6 at each step in the event simulation is computationally expensive; rather, the equation is solved when the DMC is initialized for a large set of unit wavevectors $\mathbf{k}$ finely spaced in θ and ϕ , from which lookup tables of group velocities and polarization vectors are constructed. Utilizing the known phonon wavevectors, the group

velocity and polarization vectors can then be taken directly from the table rather than computed in each iteration.

Anharmonic decay is a three-body process, and combining energy-momentum conservation with the anisotropic nature of the phonon dispersion relation creates a computationally expensive problem when dealing with the large numbers of phonons handled in the DMC. However, isotropic approximations have been computed and are utilized in the DMC; please see [82, 78] for details. Even the two-body process of isotope scattering is computationally expensive due to the necessity of conserving energy and momentum when changing direction and phonon polarization, and this computation is only performed during the initial stages of phonon evolution; an isotropic scattering process with rate given by Eq. 3.7 is used after 30 iterations of the DMC.

Finally, phonons in the DMC experience only three interactions during propagation: anharmonic decay, isotope scattering, and reflection/absorption at the detector surfaces. The rates of anharmonic decay and isotope scattering are easy to compute from Eqs. 3.7 and 3.9, and given the group velocity of each phonon it is simple to compute the time until the phonon will hit each detector surface (top/bottom, radial surface, flats). Thus, there is no need to employ a constant or energy-dependent phonon timestep; rather, we can directly compute each phonon's lifetime until its next interaction, propagate each phonon the appropriate distance, and then determine the initial state of the phonon in the upcoming time-step, i.e., determine the new wavevector after isotope scattering, the frequencies and wavevectors of the daughter phonons after anharmonic decay, the new wavevector after reflection from a detector surface, or remove the phonon from simulation and save for use in TES sensor simulations after absorption at a surface.

3.5 Phonon Sensing: Quasiparticle traps and Transition Edge Sensors

The phonon sensors patterned on the surfaces of the CDMS detectors are composed of aluminum fins connected to tungsten transition edge sensors (TESs); the specific dimensions of these elements varies depending on the detector architecture. The superconducting aluminum fins act to collect phonons, which can break the aluminum Cooper pairs to create quasiparticles (unpaired electrons in the superconductor). These quasiparticles then diffuse through the aluminum into the tungsten TES. A TES is composed of a superconducting film patterned onto a substrate and biased such that the film is held in the transition between the normal and superconducting phases in thermal equilibrium, i.e., the Joule heat applied to the TES is equal to the heat lost to the substrate. The resistance-temperature curve in the normal-superconducting transition is extremely steep at the transition temperature, and thus a minuscule change in temperature causes a measurably large change in the resistance of the sensor. Each phonon sensing channel (4 per surface) is composed of a large number of these sensors wired in parallel to cover a large area on the detector face.

3.5.1 Quasiparticle traps

The superconducting aluminum quasiparticle traps collect phonons generated in an event (prompt phonons, Neganov-Luke phonons, and recombination phonons) that have sufficient energy to break a Cooper pair to create quasiparticles. This minimum energy is equal to twice the superconducting gap of aluminum ($2\Delta_{Al} = 360\mu\text{eV}$), which is significantly larger than the energy of thermal phonons at the ≈ 50 mK

operating temperature of the ZIP detectors ($k_B T = 4.3 \mu\text{eV}$). Thus, the only phonons collected in the quasiparticle traps are high-energy, athermal phonons generated by recoils and propagating charge carriers in the detector. Furthermore, BCS theory predicts a direct relationship between a material’s critical temperature T_c and its superconducting gap (the energy difference between the bound Cooper pair state and the lowest single-electron state) of $2\Delta = 3.5k_B T$ [83, 84], and because the tungsten TESs have significantly lower critical temperatures than aluminum ($\approx 100 \text{ mK}$ vs. 1.2 K), they thus have a correspondingly lower superconducting gap ($2\Delta_W \approx 25 \mu\text{eV}$). Consequently, quasiparticles excited in the Al fins that diffuse to the interface with the W TES can enter the TES, but will then rapidly lose energy through shedding sub-gap phonons and will be unable to return to the Al.

QPTs in the DMC

In the DMC, phonons propagating in the detector crystal can be absorbed upon contact with an instrumented surface. Rather than incorporate the details of aluminum fin placement on the surfaces of the detector, the probability of phonon absorption into a quasiparticle fin is modeled using the known geometric coverage of Al on the surface of the detectors, as well as an energy-independent probability of transmission from the crystal to the aluminum. The transmission probability represents a free parameter that must be fit using the data; see Chapter 4. It is possible that this reflection parameter can be derived from acoustic mismatch theory [77], though our results indicate that the DMC’s reflection coefficient must be varied by a factor of two for different detector architectures to match the data; this discrepancy could possibly arise from variations in the Ge surface polishing and aSi/Al deposition processes.

Upon the absorption of a phonon into an aluminum quasiparticle trap, the DMC

has three available models to describe the processes of Cooper pair breaking and phonon downconversion.

1. The first and simplest model is just to input the entire phonon energy into the TES.
2. The second model imposes an upper limit on the amount of phonon energy absorbed in terms of Δ_{Al} ; up to $N\Delta_{Al}$ of energy is absorbed into the TES (where N is a tunable parameter), and the rest of the phonon energy, if any, is returned to the crystal as a lower-energy phonon.
3. The third option uses a quasiparticle cascade model developed in [85, 86, 87, 88]. The model is described in detail below.
 - The initial phonon that enters the aluminum can interact to break a quasiparticle or can propagate through the aluminum and back to the interface to be re-emitted into the germanium.
 - If a Cooper pair is broken, then two quasiparticles are created, and their energies are determined using energy conservation and the PDF

$$P(E) = \left(1 + \frac{(2\Delta_{Al})^2}{E(\Omega - E)}\right) \rho(E)\rho(E - \Omega) \quad (3.10)$$

where Δ_{Al} is the minimum quasiparticle energy, Ω is the phonon energy and $\rho(E)$ is the quasiparticle density of states function:

$$\rho(E) = \frac{E}{\sqrt{E^2 - (2\Delta_{Al})^2}} \quad (3.11)$$

- Individual quasiparticles with energies above $3\Delta_{Al}$ can emit phonons with the energy distribution

$$P(\Omega) = \Omega^2 \rho(E - \Omega) \left(1 - \frac{(2\Delta_{Al})^2}{E(E - \Omega)} \right) \quad (3.12)$$

- Quasiparticles with energies below $3\Delta_{Al}$ are assumed to decay to Δ_{Al} through emission of sub-gap phonons, which are not modeled.
- Phonons created with $E > 2\Delta_{Al}$ can break a second Cooper pair or can return to the crystal, while phonons with $E < 2\Delta_{Al}$ are simply removed from simulation.
- The cascade ends when all quasiparticles have $E < 3\Delta_{Al}$ and all phonons have been returned to the crystal or removed from the simulation due to insufficient energy to break another Cooper pair. The number of quasiparticles created is counted, and each contributes Δ_{Al} to the TES.

The inherent assumption in this model is that the quasiparticle lifetime is much longer than the quasiparticle recombination time and the timescale of absorption into the TES, though there is an associated efficiency factor of 10% employed in the DMC to prevent TES saturation from occurring at excessively low energies compared to the observed saturation scale in the detector (few hundreds of keV). This factor encapsulates loss processes, which may include a second cascade process like that described above but occurring in the tungsten making up the TES. It is also assumed that the thermal phonon density is sufficiently low that interactions between thermal phonons and quasiparticles in the traps can be ignored.




Figure 3-6: Illustrative R vs. T curve for a thin film TES with a transition temperature of 100 mK, transition width of 1 mK, and normal resistance R_n

3.5.2 Transition Edge Sensors

Transition Edge Sensors (TESs) are quantum calorimeters based on thin superconducting films biased so that the sensor remains in the normal-superconducting transition; see [89] for a review. The resistance around this bias point is an extremely steep function of the TES temperature (see Fig. 3-6), allowing highly accurate measurements of energy depositions into the TES or attached radiation collectors. These extremely sensitive sensors produce the actual measured signals that enable the detection of the athermal phonons produced by recoils and drifting carriers in the germanium crystals.

The basic calorimeter consists of an absorber at temperature T with an attached thermometer, a weak thermal link to a bath at temperature T_b , some input power P

from the thermometer or other environmental loads such as radiation, and the input energy E to be measured. The equation for the absorber temperature measured by the thermometer is

$$C \frac{dT(t)}{dt} = P - P_{link}(T(t), T_b) + E(t) \quad (3.13)$$

where P_{link} is the power delivered from the absorber to the bath through the weak thermal link and C is the specific heat of the absorber material. Assuming that $E(t) = E\delta(t)$ and $P_{link}(T(t), T_b)$ has the form $G(T(t), T_b)$, the solution to this equation is

$$T(t) = \frac{E}{C} e^{-t\frac{G}{C}} + \left(\frac{P}{G} + T_b \right) \quad (3.14)$$

Thus, an energy deposition into the absorber material produces a temperature rise of $\frac{E}{C}$, and the temperature exponentially decays back to its quiescent state with a time constant $\frac{C}{G}$.

Describing the behavior of a real TES requires a bit more detail than the idealized calorimeter treatment above. I reproduce a couple of basic results below, and a much more complete treatment can be found in [89]. First, the steepness of the R-T curve around the bias point is described using the dimensionless parameter α :

$$\alpha \equiv \frac{T}{R} \frac{\partial R}{\partial T} \quad (3.15)$$

The TES sensors on the CDMS detectors are held in voltage-bias, and the input power is thus the Joule heating $P(T) = \frac{V^2}{R(T)}$. In voltage-bias mode, if the temperature of the device increases, then $R(T)$ increases as well, so the bias power decreases, allowing the TES to cool back to its quiescent state. Similarly, if the temperature decreases, the input power increases. In this biasing mode, the TES is said to be in

negative electro-thermal feedback (ETF). More concretely, we can derive the change in the input Joule heating given a temperature change:

$$\frac{dP}{dT} = \frac{d}{dT} \frac{V^2}{R(T)} = -\frac{V^2}{R(T)^2} \frac{dR}{dT} = -\alpha \frac{P}{T} \quad (3.16)$$

We can also return to Equation 3.13 and add the details of the TES biasing scheme and the nature of the weak thermal link. In the TESs, the electrons are heated by the input Joule power and by quasiparticles from the aluminum fins. At these temperatures, the electron system is weakly thermally coupled to the phonon system, which is strongly coupled to the bath. The relevant thermal link is thus the link between the electron system at temperature $T(t)$ and the phonon system at temperature T_b , and the thermal power flow through this link is $P_{link} = K(T^n - T_b^n)$, where the values of K and n characterize the thermal link [89]. We now have

$$C \frac{dT(t)}{dt} = \frac{V^2}{R(T)} - K(T(t)^n - T_b^n) + E\delta(t) \quad (3.17)$$

where $n = 5$ for electron-phonon coupling. Expanding this equation to first order in ΔT , and using the relation that at the bias point $\frac{V^2}{R(T)} = K(T(t)^n - T_b^n)$, we have

$$C \frac{d\Delta T}{dt} = -\frac{V^2}{R(T)^2} \frac{\partial R}{\partial T} \Delta T - nKT^{n-1} \Delta T + E\delta(t) \quad (3.18)$$

Now, defining $G = nKT^{n-1}$ and using $P = \frac{V^2}{R(T)}$ and the definition of α , we can write

$$\frac{\Delta T}{dt} = -\left(\frac{\alpha P}{TC} + \frac{G}{C}\right) \Delta T + \frac{E}{C} \delta(t) \quad (3.19)$$

which again has a simple exponential solution but with an effective time constant

$$\tau_{eff} = \frac{\tau_0}{1 + \frac{\alpha P}{TG}} = \frac{\tau_0}{1 + \frac{\alpha}{n}(1 - (\frac{T_b}{T})^n)} \quad (3.20)$$

Thus, the recovery time of the TES can be sped up significantly. In the CDMS detectors, the ETF recovery time is on the order of 50-200 μs , depending on the details of the detector layout [90, 91], which allows the decay time of phonon pulses in the CDMS detectors to provide a window into the underlying phonon physics (see Section 4.2).

TESs in the DMC

The CDMS phonon sensors are composed of a large number of individual TESs wired in parallel, and each channel of parallel TESs is treated as a single large TES modeled as a network of resistive nodes. Each TES is modeled as a set of nodes with serial resistive interconnections, and parallel interconnections are present only at the first and last nodes of the individual TESs. A minimum of two nodes per TES is required, though more are necessary to capture TES phase separation [92]. The interconnect resistances versus temperature and current are modeled using the equation:

$$R_{i,j} = \frac{R_{max} - R_{min}}{2} \left\{ 1 + \tanh \left[\frac{T_{i,j} - T_c \left(1 - \frac{|I_{i,j}|}{I_c} \right)^{n_{sc}}}{T_w} \right] \right\} \quad (3.21)$$

and $R_{i,j} = R_{max}$ when $I_{i,j}$ is sufficiently large to drive the TES normal. In Eq. 3.21, R_{max} is the TES's normal state resistance, R_{min} is the TES's superconducting state resistance (set to a very small but nonzero value), T_c is the temperature at the midpoint of the superconducting-normal transition, T_w parameterizes the width of the transition, I_c is the critical current, and $n_{sc} \approx 2/3$ [88, 93].

The system is solved by assigning each node a resistance that is the average of the resistances of the TES connections to neighboring nodes, and Kirchoff's Current Law allows the construction of a matrix equation relating the vector of node voltages $\vec{V}$, the resistances of the interconnects, and the vector of voltage boundary conditions $\vec{B}$: $\mathbf{G}\vec{V} = \vec{B}$, where the elements of G are determined as described in [88].

The thermal power inputs to each TES node can be broken down into the Joule heating, the phonon power input, the weak thermal link to the bath, and thermal diffusion between adjacent nodes. Joule heating is described simply by $P_J = \frac{V^2}{R_{i,j}}$, where V is determined from the node voltages around interconnect i, j . The power to the bath is calculated using $P_{link} = K(T^5 - T_b^5)$, where K is the product of the effective node volume and an electron-phonon coupling parameter. The phonon input power comes from the quasiparticle traps, as described above; phonons generate quasiparticles, each of which contributes $\epsilon \times \Delta_{Al}$ of thermal energy to the TES, where ϵ is an efficiency parameter, measured to be $\sim 10\%$ through a rough comparison of the recoil energy scale at which TESs become saturated (driven normal). As mentioned in Sec. 3.5.1, phonons can be absorbed anywhere on the instrumented surfaces, and the phonon energy distribution is parsed up by absorption position/time and delivered to individual TESs. Finally, thermal diffusion in the DMC is modeled as a 1-D diffusion process, i.e., there is no thermal diffusion between physically separated TESs wired in parallel. Please see [88] for more detail on thermal diffusion modeling in the DMC's TES simulator.

3.6 Charge Carrier Physics

As noted in Sec. 2.4.3, cryogenic semiconductors make useful radiation detectors because all charge carriers are bound (frozen out) in equilibrium but can be elevated

to the conduction band through interactions with incident radiation. The resulting electron-hole pairs can be drifted in an externally applied field and be measured using ionization sensors in contact with the semiconductor crystal. Charge carrier propagation in semiconductors is described using an effective mass that captures how charge carriers accelerate in the crystal in the presence of an electric field and how they interact with the lattice nuclei. In this section, I review the physics of charge carrier propagation in semiconductor crystals with a focus on germanium and describe the implementation of this physics in the DMC.

3.6.1 Charge Carrier Propagation and the effective mass tensor

The Newtonian concept of a scalar mass that relates force and acceleration vectors cannot be directly applied to charge carrier propagation in a crystal under an applied electric field (if it could, then applying any field should always produce a proportional current, as in the Ohmic conduction regime). Though conduction electrons and the corresponding holes are free to propagate through the crystal, they are still subject to periodic forces from lattice ions and bound electrons, and thus do not necessarily respond by accelerating in position-space in direct proportion to and collinear with the applied field. The applied electric field instead acts to induce a derivative in the carrier's ‘crystal momentum’

$$\dot{\vec{p}}_{crystal} = q\vec{E} \quad (3.22)$$

where q is the charge of the particle, $\vec{E}$ is the electric field, and $\vec{p}_{crystal}$ is related to the wavevector of the carrier through

$$\vec{p}_{crystal} = \hbar\vec{k} \quad (3.23)$$

Because the carrier is propagating in the presence of periodic lattice forces from valence electrons and nuclei, the dispersion relation $\epsilon(\vec{k})$ is generally not quadratic or isotropic. The acceleration of the particle can be derived using Eqs. 3.22 and 3.23 along with the standard relation $\epsilon = \hbar\omega$ and the group velocity $\vec{v}_g = \frac{d\omega}{d\vec{k}}$

$$a_i = \frac{q}{\hbar^2} \vec{E} \cdot \left(\frac{d}{d\vec{k}} \frac{d\omega}{dk_i} \right) \quad (3.24)$$

which leads to the definition of an effective mass m^* [94, 95]

$$\left(\frac{1}{m^*} \right)_{i,j} = \frac{1}{\hbar^2} \frac{d^2\omega}{dk_i dk_j} \quad (3.25)$$

Finally, near a minimum at $\vec{k}_0$, a quadratic but still anisotropic form can be used to approximate the dispersion relation


$$\epsilon(\vec{k}) \approx \epsilon(\vec{k}_0) + \frac{\hbar^2}{2} \left(\frac{k_x^2}{m_x} + \frac{k_y^2}{m_y} + \frac{k_z^2}{m_z} \right) \quad (3.26)$$

In germanium, the conduction band minima (henceforth referred to as valleys) lie at $\vec{k}_i$ along the eight $\langle 111 \rangle$ crystal directions (see Figures 3-7(a) and 3-7(b)), resulting in a dispersion relation of the form


$$\epsilon(\vec{k}) = \frac{\hbar^2}{2} \left(\frac{\delta k_1^2}{m_{\parallel}} + \frac{\delta k_2^2}{m_{\perp}} + \frac{\delta k_3^2}{m_{\perp}} \right) \quad (3.27)$$

where $m_{\parallel} = 1.58m_e$, $m_{\perp} = 0.081m_e$, $\delta\vec{k}$ represents the k vector relative to the valley minimum $\vec{k}_i$, direction 1 lies parallel to $\vec{k}_i$, and directions 2 and 3 are perpendicular to direction 1 and each other. As shown in [96], the orientation of the electron valleys and the extreme difference between $m_{\parallel}$ and $m_{\perp}$ can cause electrons to propagate in directions oblique to the applied electric field; in particular, for fields

applied in the $\langle 100 \rangle$ direction as in the CDMS detectors, the electron drift velocity is oriented along four directions, each approximately 33 degrees from the $\langle 100 \rangle$ direction.



(a) Electron and hole valley structure of germanium, showing the electron L-valley at $\langle 111 \rangle$, the Γ -valley at $\langle 000 \rangle$, and the X-valley at $\langle 100 \rangle$. Figure from [97].



(b) Brillouin zone of germanium, with L-valleys depicted along $\langle 111 \rangle$ axes. The bands are described as ellipsoids with $\frac{m_\perp}{m_\parallel} \approx 19.5$. Figure from [97].

Holes, on the other hand, do propagate collinear with the field, as the conduction band maximum is at $\langle 000 \rangle$. The hole dispersion relation is anisotropic, and holes can enter one of two hole conduction bands degenerate at $\langle 000 \rangle$ (the heavy and light hole bands). However, in the DMC we assume a scalar effective mass $m_h^* = 0.35m_e$. Because the ionization sensors of the CDMS detectors are sensitive only to radial position, the oblique propagation of the electrons is of critical importance, but the details of collinear hole propagation are probably of less importance in simulations of the CDMS detectors. Furthermore, previous simulations [96] as well as our own have demonstrated that this scalar effective mass model does accurately reproduce

the drift velocity of holes at low fields (< 3 V/cm).

3.6.2 Emission of Phonons by Charge Carriers

As a charge carrier drifts through a semiconductor under an external electric field, it emits phonons in a process analogous to Cherenkov radiation when the carrier velocity exceeds the speed of sound (specifically, the velocity of longitudinal phonons) in the medium; these phonons are known as Luke-Neganov phonons [75]. In equilibrium, the energy acquired by the charge carriers as they propagate in the field is balanced by the energy emitted as phonons. Below, we review the physics of Luke-Neganov phonon emission in the case of a scalar effective mass.

Emission of phonons by charge carriers is an energy- and momentum-conserving process, so we can write

$$\vec{k} - \vec{k}' = \vec{q}; \epsilon - \epsilon' = \hbar\omega \quad (3.28)$$

where $\vec{k}$ and $\vec{k}'$ are the initial and final carrier momenta, respectively, ϵ and ϵ' are the initial and final carrier energies, respectively, and $\vec{q}$ and ω are the phonon momentum and energy, respectively, and are related through the dispersion relation $\omega = v_L q$, where v_L is the velocity of longitudinal phonons. Letting θ be the angle between $\vec{k}$ and $\vec{q}$ and ϕ be the angle between $\vec{k}$ and $\vec{k}'$, energy-momentum conservation leads to

$$k'^2 = k^2 + q^2 - 2kq\cos(\theta) \quad (3.29)$$

$$q = 2(k\cos(\theta) - k_L) \quad (3.30)$$

$$\cos(\phi) = \frac{k^2 - 2k_L(k\cos(\theta) - k_L) - 2(k\cos(\theta) - k_L)^2}{k\sqrt{k^2 - 4k_L(k\cos(\theta) - k_L)}} \quad (3.31)$$

Fermi's Golden Rule can be used to derive the phonon emission rate and the

angular distribution (see [98] for details):

$$1/\tau = \frac{1}{3} \frac{v_L}{l_0} \frac{k}{k_L} \left[1 - \frac{k_L}{k} \right]^3. \quad (3.32)$$

$$P(k, \theta) d\theta = \frac{v_L}{l_0} \left(\frac{k}{k_L} \right)^2 \left(\cos(\theta) - \frac{k_L}{k} \right)^2 \sin \theta d\theta \quad (3.33)$$

where $0 \leq \theta \leq \arccos(k_L/k) < \pi/2$, $l_0 = \frac{\pi \hbar^4 \rho}{2m^3 C^2}$, $k_L = mv_L/\hbar$, and C is the deformation potential constant, which describes the strength of electron-phonon coupling.

Because the holes have an approximately scalar mass near the valence band maximum, the above treatment can be applied directly to emission of Luke phonons by holes. In the case of electrons, the above treatment is not immediately applicable, but we can perform a (non-unitary) transformation into a space in which the effective mass of the electrons can be treated as constant. This transformation is termed a Herring-Vogt transformation [99, 100]. Because the transformation is non-unitary, the magnitude of the phonon wavevector q must be re-scaled after the wavevector undergoes the reverse transformation, to ensure energy conservation. More details on the Herring-Vogt transform will be presented in 3.6.4.

3.6.3 Charge Measurement

The ionization signal is measured through the image charge induced on the aluminum electrodes patterned on one (ZIP) or both (iZIP) sides of the detector. The Shockley-Ramo theorem [101, 102] states that the image charge induced on an electrode by a charge carrier created at initial position x_i and moving to a final position x_f is

$$\Delta Q = -q[V(x_f) - V(x_i)] \quad (3.34)$$

where ΔQ is the induced charge, q is the moving charge, and V is known as the weighting potential or the Ramo potential, and is defined as the potential field in the system under the following boundary conditions: the selected electrode is at 1 V, all other electrodes are at 0 V, and there are no charges. Interestingly, the observed signal from a single moving carrier is independent of the actual potential applied to the electrode and of the presence of other static or moving charges, except through the corresponding effects on x_f and x_i . In the SCDMS detectors, the charge carriers are created in a charge ‘droplet’ at the site of the interaction. If no drift field were applied or if the drift field is not sufficiently strong to overcome self-shielding effects in the droplet, these carriers can diffuse and recombine with each other or can fall into mid-gap ‘traps’ caused by local ionization centers. However, when a sufficiently strong drift field is applied, the electrons and holes in the droplet separate and propagate until they reach the appropriate electrode. Because a hole drifting in one direction induces the same signal as an electron drifting in the opposite direction, the final signal observed in a perfectly collected event (i.e., x_f is at the electrode) is simply equal to $\Delta Q = \pm ne_0$, where n is the number of charge pairs, e_0 is the electron charge, and the sign of the induced charge depends on the sign of the potential applied to the electrode. Charge collection vs. applied field in germanium was investigated in [103] and it was found that charge collection is effectively complete at fields above ~ 200 V/m, as shown in Fig. 3-8. The SCDMS detectors are operated well above this point, and so charge collection can be expected to be complete, though effects from accumulated space charge appear over time. Events that occur very near the surface also exhibit reduced charge collection due to an increased local trap density or charge carriers being collected by the ‘wrong’ electrode. Finally, it should be mentioned that the drift velocity of carriers in germanium is sufficiently high that all carriers are collected within a couple of samples at the 800 ns sampling frequency




Figure 3-8: Number of collected charge carriers (triangles, in thousands) and total phonon signal (circles) from 59.5 keV photons in germanium vs. electrode bias. Devices were 1 cm thick crystals, so the x-axis can also be read as applied field in V/cm. The line represents the expected total phonon signal including Neganov-Luke phonons. Figure from [104].

employed in CDMSII and SCDMS, and so all ionization pulses are highly self-similar, with a rapid rise lasting a couple of samples, and a fall time of $\sim 40\mu s$ determined by the feedback impedance of the charge amplifier circuit.

3.6.4 Charge Carrier Physics in the DMC

Charge creation

Charge carriers in the DMC are created only at the point of the initial scatters. The cascade process by which a single high-energy electron or ion produces multiple electron-hole pairs is not modeled, nor are any carrier-carrier interactions during simulation (the exception being a special serial mode in which the field of trapped carriers is added to the applied external field, producing an effective interaction between charge carriers trapped in previous iterations and those propagating in the current iteration. This mode is used to model time-dependent trapping of space charge).

Charge propagation and the Herring-Vogt Transform

The anisotropy of the electron conduction band is treated by transforming from the $(\vec{x}, \vec{y}, \vec{z})$ coordinate system into a new system with unit vectors $(\vec{e}_{\parallel}, \vec{e}_{\perp 1}, \vec{e}_{\perp 2})$, where $\vec{e}_{\parallel}$ is aligned with the electron band minima along the $\langle 111 \rangle$ directions (the specific direction of $\vec{e}_{\parallel}$ is determined by the corresponding electron's valley), and the 2 $\vec{e}_{\perp}$ unit vectors lie perpendicular to $\vec{e}_{\parallel}$. The Herring-Vogt transform is then applied in the $(\vec{e}_{\parallel}, \vec{e}_{\perp 1}, \vec{e}_{\perp 2})$ space to make the electron conduction band isotropic; this transformation is non-unitary and is given by

$$\mathcal{T} = \begin{pmatrix} \sqrt{\frac{m_{eff}}{m_{\parallel}}} & 0 & 0 \\ 0 & \sqrt{\frac{m_{eff}}{m_{\perp}}} & 0 \\ 0 & 0 & \sqrt{\frac{m_{eff}}{m_{\perp}}} \end{pmatrix} \quad (3.35)$$

where the effective electron mass is defined as $\frac{3}{m_{eff}} = \frac{1}{m_{\parallel}} + \frac{2}{m_{\perp}}$. In the Herring-

Vogt space, the equation of motion is

$$\hbar \frac{d\vec{k}^*}{dt} = e\vec{E}^* \quad (3.36)$$

The electron momenta are propagated according to Eq. 3.36, and the corresponding position space steps are determined by transforming the momenta using the back transform $\mathcal{T}^{-1}$. The phonon scattering equations described in Sec. 3.6.2 are applied in the same way to holes in the $(\vec{x}, \vec{y}, \vec{z})$ space and electrons in the Herring-Vogt space, as the isotropic sound speed v_L is unchanged in the transformation. The phonon and electron momenta and directions are solved for in the Herring-Vogt space and transformed back to $(\vec{x}, \vec{y}, \vec{z})$ space for position propagation. Because of the non-unitarity of the Herring-Vogt transform, phonon energy is not conserved in the transformation, so the magnitude of the final phonon momentum vector in the Herring-Vogt space must be used to renormalize the phonon momentum vector in the $(\vec{x}, \vec{y}, \vec{z})$ space. The angular distributions of the outgoing phonons as well as energy conservation are thus simultaneously maintained.

Because charge carriers are continuously being accelerated by the applied field, the phonon emission rate of each carrier is constantly changing, which imposes an upper limit on the time step that the simulation can use without producing unphysical results. On the other hand, it is inefficient to employ a constant time step, as the derivative of the emission rate with respect to energy is also energy-dependent, and thus imposing an absolute maximum time step suitable for high-energy carriers would result in low-energy carriers propagating on simple parabolic paths for a large number of steps. The DMC's charge transport algorithm has evolved multiple times; here, I will describe the Herring-Vogt transform [100] used to model the anisotropic and only the two most recently employed algorithms, which were utilized

in all validation studies of the charge transport system presented in 4.5.

I will refer to the first of the two charge propagation algorithms as the *first-order* charge propagation algorithm. In this algorithm, the maximum possible sufficiently small timestep for carrier propagation is determined empirically by running the Monte Carlo and finding the maximum observed carrier momentum at a given applied field. The following two relations were determined using this method

$$\vec{k}_{max,el} = 13 \times k_L |\vec{E}|^{1/3} \quad (3.37)$$

$$\vec{k}_{max,h} = 6.8 \times k_L |\vec{E}|^{1/3} \quad (3.38)$$

The time step is set to one-half of the mean phonon-emission time at the maximum momentum in the carrier's local field, given by substituting k_{max} into Eq. 3.32. Testing longer time steps has demonstrated that Luke phonon energy conservation is violated at an unacceptable level ($> 1\%$) when the time step is increased.

The first-order method has a major deficiency in that the current carrier momentum is ignored entirely, resulting in a large number of time steps in which low-energy carriers are accelerated along simple parabolic paths. When carrier velocities are below v_L , phonon emission is impossible, and so it would be computationally beneficial to utilize larger timesteps; however, the fact that the carrier's local field changes along its path must be considered. Because of these computational disadvantages, the first-order charge propagation algorithm represents the dominant contributor to total run time when it is used with the phonon MC and the TES/FET simulators, even when simulating < 100 charge carrier pairs. The *second-order* charge propagation algorithm deals with these disadvantages and makes it possible to run the DMC with thousands of simulated charge pairs without dominating the DMC run time. To do this, the scattering rate is sampled at the current charge momentum and local

field and also at a proposed future momentum and local field. The random timestep is then sampled from the distribution

$$\ln N = -(c_0 t - c_1 t^2/2) \quad (3.39)$$

where c_0 is the inverse of the scattering time τ_0^{-1} at the current momentum, and c_1 is determined from a linear interpolation of the current and future scattering times, $c_1 = \frac{1}{t_1 - t_0}(\tau_1^{-1} - \tau_0^{-1})$ [88]. The functional form is motivated in [88]; in short, this distribution corresponds to sampling from an exponential distribution, with the scattering rate $\tau(t)$ expanded to first order. Empirically, it is found that the time step $\Delta t_0 = t_1 - t_0$ can be increased by an average of 15 or 20x over that used in the first-order method for electrons and holes, respectively. With this method, we can propose an initial Δt_0 larger than τ_0 and determine a probable time to the next Neganov-Luke phonon emission. The carrier positions and momenta can be iterated using the equations $x_1 = x_0 + v\Delta t_0 + 1/2a\Delta t^2$ and $v_1 = v_0 + 1/2(a(x_0) + a(x_1))\Delta t$, and the present and future scattering rates are determined from v_0 and v_1 . If the proposed timestep Δt_0 is longer than the Neganov-Luke scattering time drawn randomly from Eq. 3.39, then the charge is propagated the full distance, and otherwise the charge propagates to the next Neganov-Luke phonon emission event and emits a phonon with energy determined using Eq. 3.33.

Electric Field Modeling

The electric potentials and Ramo potentials were solved using the COMSOL Multiphysics finite element software package, which provides a 3-dimensional mesh of nodes with associated electric potentials. The electric potential away from these nodes is interpolated from the node potentials. For the ZIP and mZIP detectors,

which both have high electrode coverages on both surfaces and no interleaving issues, the boundary conditions were set such that the entire top and bottom face of the detector were at constant potential. For the iZIP detectors, the complex E-field geometry near the surfaces is crucial to include in the model, and so it is necessary to include an appropriate electrode geometry.

The electric field lookup during carrier propagation can become the dominant contributor to computation time if care is not taken. However, a number of simplifications can be made when the mesh is generated and when the results are utilized by the DMC:

1. The electric field in the iZIP configuration is relatively constant within a constrained area away from the top, bottom, and radial surfaces. The size of the 3D mesh can thus be significantly reduced by defining a set of boundaries in R and Z with an associated constant electric field.
2. Because the potential values within a mesh tetrahedron are linearly interpolated from the stored values at the vertices, the electric field within a tetrahedron is necessarily constant. To reduce the number of field lookups, the DMC keeps track of which tetrahedron each charge is currently in and only recomputes the local field when a charge leaves one tetrahedron and enters another.
3. The density of the mesh is highest near the thin ($\sim 8 \mu m$) ionization rails, as the potential varies most rapidly in these regions. To reduce the number of mesh points in these regions, the model and boundary conditions are iterated as follows. A high-density mesh is generated using the correct electrode width and applied voltage. A surface of constant voltage surrounding the ionization rails

is determined. Then, the electrodes and corresponding boundary conditions are replaced by this surface of constant voltage, which is shaped like a semicircular tube with a significantly larger radius than the electrode. The model is solved again with the larger ‘electrode’ and the new boundary conditions, and node points defining the original electrode voltages are inserted into the mesh. The number of mesh points is reduced because the voltage outside of the surface of constant voltage varies much more slowly than the voltage in the immediate vicinity of the electrodes; accurate voltages near the electrodes are sacrificed, but the fields in this region are sufficiently large that charges reach the electrode rapidly after entering this region.

Interactions between Carriers and the Detector Surface

Charge carriers that strike a detector surface are absorbed under two conditions. The first is that the carrier has reached an appropriately biased ionization electrode, at which point it is collected by the ionization electrode and removed from simulation, and recombination phonons are emitted at the point where the carrier entered the ionization electrode. The second condition is that the carrier’s momentum perpendicular to the surface is high enough for the charge to inject into the surface layer. The density of charge traps at the surface is assumed to be sufficiently high that the charge is immediately trapped, and recombination phonons are again emitted from the point of contact with the surface. If neither condition is met, the charge is reflected diffusively away from the surface and continues propagating until one of the two conditions is met.

The second absorption condition described above deserves some further discussion. On very short timescales after a recoil, charge carriers are generated as a ‘ball’

of hot electrons and holes, and their high momenta and self-shielding effects allow them to propagate quasi-diffusively for a short time before they shed their kinetic energies into the phonon system and the drift field separates carriers of different sign. In surface events, some of these generated high-energy carriers can back-diffuse into the opposite electrode, thus reducing the measured yield. An amorphous silicon layer approximately 50 nm thick is deposited on the surfaces of the germanium crystal to act as a barrier against this back-diffusion, as the bandgap of amorphous Si (≈ 1.7 eV) is significantly larger than that of germanium (≈ 0.75 eV). The inclusion of the amorphous silicon layer has been observed to increase the yield of surface events [105]. The DMC also predicts that charge carriers generated in the bulk of an iZIP could first contact the detector surface away from the ionization electrodes due to diffusion and oblique propagation, but the data show that bulk ERs and NRs tend to exhibit very good charge collection. This model of surface absorption effectively assumes the existence of a boundary layer between the bulk semiconductor crystal and a non-conductive surface layer. The final fitted value for the momentum threshold is well below the full 1 eV difference in the bandgap of the two materials (Sec. 4.5), however, such a model can be justified by band bending near the surface of the detector and the existence of surface states with a momentum barrier. A more complete model would investigate a field-dependent barrier along with asymmetry between the electron and hole surface boundaries, however, a simple model with a constant and symmetric barrier is able to reproduce the behavior of surface events (see Sec. 4.5).

Readout

As mentioned in 3.6.3, ionization pulses are highly self-similar. Thus, rather than modeling the readout circuitry, it is simpler to utilize template ionization pulses averaged from a large number of well-measured high-energy ionization pulses. These template pulses are scaled according to the induced charge signal determined from the final carrier positions and the simulated Ramo potentials. Recent updates to the DMC have allowed charge tracks, rather than just initial and final carrier positions, to be stored, and these charge tracks can be used along with a realistic model of the charge amplifier to compute the expected ionization signal over time. The pulses produced by this method are still very linear, but small variations as charges move into/out of the Ramo potentials of particular electrodes can be reconstructed. However, this new option has not been employed in any of the results in this thesis.

Chapter 4

Validation and Tuning of the CDMS DMC

The CDMS DMC contains a number of physical and tunable parameters that determine the detailed behavior of the phonon, charge, and TES systems. For the CDMS DMC to be a useful tool, it is important to validate its performance against those of real detectors to verify that detector behavior is consistent with DMC predictions using measured values of the physical parameters and to tune any detector-specific tunable parameters to describe the real detector being tested. This chapter will review the work performed to validate the physics implemented in the DMC and to tune the physical models used to describe the CDMS detectors. Validation of the DMC progressed in parallel with development and refinement of the physics described in the DMC, and the version and details of the DMC models used in each of the validation steps described below will be described along with the results.


4.1 Phonon absorption and losses in the mZIP-style detectors

The earliest DMC validation and tuning studies utilized data from a test detector known as the mZIP to investigate the rate of phonon absorption from the Ge crystal into the aluminum quasiparticle traps and overall phonon losses using the partitioning of energy into the phonon sensing channels, the decay time of the phonon pulses, and the absolute phonon energy collected relative to the recoil energy of the event (calibrated to known lines in the employed radioactive sources). When these tests were performed, the DMC did not yet include descriptions of phonon-quasiparticle downconversion in the aluminum quasiparticle traps, nor did it include anisotropic electron propagation in the charge carrier system.


4.1.1 The mZIP detector

A photo of the mZIP (Mercedes ZIP, named due to the similarity between the inner channel layout and the iconic Mercedes-Benz logo) detector and accompanying schematic illustration are shown in Figures 4-1(a) and 4-1(b). The mZIP detector utilized in these studies was a Ge crystal 1" thick and 3" in diameter. The detector architecture included 6 phonon sensors on the phonon-collecting surface: 3 inner sensors, each covering 120 degrees of the detector and extending from 0" - 1.3" in radius (so that each covers 1/4 of the total surface area), and 3 outer sensors whose boundaries are rotated 60 degrees relative to those of the inner sensors, each covering 1/12 of the total surface area. The 3 outer sensors were intended to be wired as a single channel but was divided as shown in Fig 4-1(b) for these experiments. The aluminum coverage on the phonon channel side was 43.1%, 37% of which was active

(QETs) and 6.1% of which was passive (wiring/ionization grid). The opposite side of the mZIP has two concentric ionization channels with 10% total aluminum coverage; phonons can be collected in these rails but go unmeasured. In the experiments compared with the DMC, the ionization sensors were biased at -6V relative to the phonon sensors.



(a) Figure 4-1: Picture of mZIP detector (G10F) in housing, with phonon side shown. The quasiparticle traps and ‘Mercedes’ channel layout are clearly visible. Photo courtesy of Paul Brink



(b) Schematic of an mZIP detector, wire in 6-channel configuration with 3 inner and 3 outer phonon channels on top. The inner and outer ionization channels are visible on the bottom

4.1.2 mZIP Experiment, Data Selection and Processing

Experiments

An SCDMS mZIP detector was exposed to an ^{241}Am source to produce an ER population in the detector. The ionization channels were biased at -6V relative to the phonon sensors, meaning that the total Luke phonon energy emitted by charge carriers was roughly twice the recoil energy. A position-dependent delay parameter was defined using the relative 20% phonon pulse risetimes (the time for the phonon pulse

to reach 20% of its maximum) in the three inner channels. The 2D (xdelay, ydelay) histogram was ‘flattened’ (i.e., an equal number of events from each bin in [xdelay, ydelay]) to produce a roughly position-uniform set of data events to compare with DMC (roughly uniform because it is known that the inter-channel delay quantities start to ‘fold back’ at high radius).

Data Quality

mZIP data was selected from experimental runs at the UC Berkeley detector testing facility (Run 362 [106]). Detector calibration is described in Sec. 4.1.2. For this analysis, only events with recoil energies between 20 and 100 keV were selected; the lower energy threshold is used to remove events with low signal-to-noise, and the upper energy threshold prevents the inclusion of events in which the TES sensors become saturated (enter a fully normal resistive state). The data quality cuts employed were:

- cQstd_G9F and cPstd_G9F: These two cuts enforce an upper limit on the magnitude of the baseline charge and phonon noise, respectively, measured in the 512 pre-pulse bins.
- cChi2_G9F and cPChi2_G9F: These two cuts enforce upper limits on the magnitude of the χ^2 of the fit between the measured charge and phonon pulses, respectively, with their corresponding optimal filter templates. These cuts primarily help to remove ‘pileup’ events, events in which more than one recoil occurs within the sampling window (quite common at the surface-based testing facilities), as the fit between the measured pulses and templates will be quite poor if multiple recoil events are present in the data.

- `cGoodTimingLoose_G9F`: This cut enforces limits on the difference between the measured start times of the charge and phonon pulses.
- `cdel_consist_G9F` and `ctri_part_G9F`: These cuts effectively ensure that the relative phonon pulse delays and energy partitioning among the 3 inner phonon sensors are within reasonable ranges.

Fitting pulse decay times

Fitting the phonon pulse decay times is a relatively straightforward procedure. As mentioned in Sec. 3.4.2, at late times, the phonons have downconverted to sufficiently low energies that they propagate ballistically rather than quasi-diffusively, and so the rate of phonon power input to the TES sensors is set by the dimensions of the crystal, the phonon propagation velocity in Ge, and the reflection coefficient at the Ge-aSi-Al surface.

Absolute energy calibration and phonon energy partitioning

The absolute phonon energy deposited into each of the TESs can be determined from the measured current pulse using the known parameters of the TES circuit and measurement system. The QET biasing circuit is shown in Fig. 4-2. Analysis of Fig. 4-2 provides the following equation for determining the TES current I_s from the measured ADC trace:

$$I_s = \frac{N_{ADC}}{R_{FB} \times N_1/N_2 \times Q \times G} \quad (4.1)$$

where N_{ADC} is the measured ADC bin, R_{FB} is the feedback resistance, N_1/N_2 is the turns ratio in the inductive coupling shown in Fig. 4-2, Q is the resolution of the




Figure 4-2: Circuit used to bias and read TES channels in G9F, UCB Run 362

ADC in bins per Volt, and G is the gain of the output driver amplifier. Expanding $P_{TES} = V_{TES} \times I_s$ about the bias current I_{s0} provides the following equation for the change in TES power provided by the biasing circuit (which is equal to the phonon power input to the TES)

$$\Delta P(\Delta I_s)_{I_s=I_{s0}} = (R_{sh}I_b - 2(R_{sh} + R_P)I_{s0})\Delta I_s - (R_{sh} + R_P)\Delta I_s^2 \quad (4.2)$$

All parameters in these two equations are either known (see Table 4.1.2) or measured (I_s) with the exception of I_{s0} . I_{s0} must be estimated using the known bias current I_b and curves of I_s vs. I_b measured at the beginning of the run, or by using the known normal resistance of the TESs and the observed current when the TES channels are driven completely normal in high-energy events caused by cosmogenic radiation. The determination of I_{s0} represents a significant systematic in this calibration method due to uncertainties in R_P and R_{sh} . Using the two methods to determine I_{s0} produced results differing by an average of 5% in the 3 inner channels and 25% in the outer channels, though because the outer channels have 3x fewer

R_{FB}	1.2 k Ω
N_1/N_2	10
Q	1024 V ⁻¹
G	10
R_{sh}	~ 20 m Ω
R_P	~ 10 m Ω

Table 4.1: Known values of parameters of the TES biasing and readout circuit used in Eq. 4.1

TESs and are thus biased at lower currents, similar absolute differences in the estimated I_{s0} results in significantly larger percentage discrepancies. Because of this systematic, in most cases (including future validation studies later in this chapter), ΔP is assumed to be exactly linear with ΔI_s in the energy range of interest, and the phonon signals are calibrated by comparison with the ionization signal, which is calibrated using known gamma lines from a radioactive source. However, using Eq. 4.2 allows the determination of the absolute phonon power into the TESs, while assuming $\Delta P \propto \Delta I$ and calibrating to the observed charge signal does not.

Phonon energy partitioning among the channels of the CDMS detectors is usually presented as scatter plots of the X and Y energy partitioning, which are essentially the first geometric moments of the total energy in each channel. For the mZIP detector, these parameters are computed as follows:

$$x_{part,inner} = \cos\left(\frac{\pi}{6}\right) * P_D + \cos\left(\frac{5\pi}{6}\right) * P_B + \cos\left(\frac{3\pi}{2}\right) * P_C \quad (4.3)$$

$$y_{part,inner} = \sin\left(\frac{\pi}{6}\right) * P_D + \sin\left(\frac{5\pi}{6}\right) * P_B + \sin\left(\frac{3\pi}{2}\right) * P_C \quad (4.4)$$

$$x_{part,outer} = \cos\left(\frac{\pi}{2}\right) * P_A + \cos\left(\frac{7\pi}{6}\right) * P_E + \cos\left(\frac{11\pi}{6}\right) * P_F \quad (4.5)$$

$$y_{part,outer} = \sin\left(\frac{\pi}{2}\right) * P_A + \sin\left(\frac{7\pi}{6}\right) * P_E + \sin\left(\frac{11\pi}{6}\right) * P_F \quad (4.6)$$

where P_i indicates the fraction of the total phonon energy absorbed in channel i . The arguments of the sine and cosine functions are the angles between the center of a given channel and the major flat of the Ge crystal (taken to be the x-axis). The delay parameters mentioned in 4.1.2 are computed using much the same formulas, replacing the phonon energies with pulse risetimes. The 2D $x_{part} - y_{part}$ distributions strongly reflect the sensor layout of the detector, and are often termed ‘box plots’ (ZIP, iZIP3) or ‘triangle plots’ (mZIP, iZIP4); the reason behind the name is obvious in Fig. 4-3. The extent of the distributions measures the amount of ‘position-dependent’ phonon energy vs. the total phonon energy; the earliest absorbed phonons, especially the Neganov-Luke phonons, provide information about the location of the incident event, but after $\sim 200 \mu s$, the phonons are essentially a uniform gas filling the detector bulk and no information about the event position remains.

4.1.3 Simulations

DMC model

In this round of validation, only the first of the phonon-quasiparticle downconversion models described in Sec. 3.5.1 was implemented, so all incident phonon energy

was absorbed. Furthermore, the TES simulator was not fully coupled to the phonon simulator yet in these simulations; the energy partitioning distributions were determined directly from the absorbed phonon energies and positions, and the experimental TES exponential decay times were compared with the absorbed phonon energy trace. Oblique electron propagation had also not yet been included in the DMC at the time of this analysis. Monoenergetic electron recoils were simulated on a 3D grid within the detector, with a grid spacing of 2mm. The grid extended through one fourth of the detector (+x, +y, all z), and the signals from ERs in the other four quadrants were generated by appropriately reflecting the final phonon positions about the central planes through the detector.

The purpose of this validation investigation is to determine the phonon absorption parameters at the surfaces and at the detector walls that allow us to reproduce phonon energy partitioning and phonon pulse decay time parameters. Thirty-six variations of the DMC were run. The phonon absorption coefficients on the non-instrumented sidewalls (denoted A_{side}) were set to (0.001, 0.003, 0.01, 0.03, 0.1, 0.3), and the phonon transmission coefficients on instrumented surfaces, denoted (A_{surf}), were set to (0.0232, 0.0696, 0.232, 0.348, 0.464, and 0.696). The numbers seem a bit arbitrary, but multiplied by the aluminum surface coverage of 43.1% on the phonon sensor side, these numbers correspond to (0.01, 0.03, 0.1, 0.15, 0.2, 0.3), and I will use these numbers throughout this section for readability.

4.1.4 Results

For comparison with the energy partitioning data, the phonon energies in the DMC were parsed according to the final phonon absorption location and summed. Phonon partitioning plots are presented in Figs. 4-3 and 4-4. The inner channel energy

partitioning is plotted in blue and is distributed in a triangular shape, while the outer channel energy partitioning is presented in magenta and appears in a crosshair-like pattern. The triangular shape of the inner channels indicates that a significant amount of the position-dependent energy can be shared between the primary channel and a secondary channel, while in the outer channels, the crosshair indicates that usually one or even none of the channels receives a significant amount of position-dependent energy. The effects of changing the phonon absorption coefficient of the sidewalls while holding the phonon transmission probability at the instrumented surfaces constant is illustrated in Fig. 4-4, while Fig. 4-3 presents the effects of changing the absorption coefficient of the instrumented surfaces while holding that of the sidewalls constant. In both cases, it can be seen that increasing the absorption coefficient increases the degree of position-dependent energy partitioning, which is reasonable as the event position information becomes more diluted as phonons have more opportunities to bounce from the surfaces. The energy partitioning plots of the data and the three best matching simulations are presented in Fig. 4-5. The ‘best-match’ was determined by eye, as Figs. 4-3 and 4-4 demonstrate that the partition plot changes drastically as the absorption coefficients are varied. The parameters of the best matching simulations were $(A_{surf}, A_{side}) = (0.1, 0.03)$, $(0.15, 0.01)$, and $(0.15, 0.03)$. The partitioning plots from each of these simulations match the data well in terms of the extent of the partition triangles.

Comparing the pulse decay times should be taken with a grain of salt at this stage. The mZIP phonon pulse decay times are approximately $300 \mu s$, which is short enough that the TES ETF time could still play a role in slightly extending it. Because the TES simulator was not utilized in this study, we compare the phonon energy pulse decay time and find the best matches for the simulation runs with $(A_{surf}, A_{side}) = (0.03, 0.03)$, $(0.1, 0.001)$, $(0.1, 0.01)$, and $(0.1, 0.03)$ with decay




Figure 4-3: Inner and outer X-Y partition plots from DMC simulations of the 6-channel mZIP. The sequence of figures illustrates the effects of varying the phonon transmission coefficient at the Ge-Al interface on the instrumented surfaces. Six transmission coefficients, 0.01, 0.03, 0.1, 0.15, 0.2, and 0.3, are presented, and the coefficient of phonon absorption on the non-instrumented sidewall of the detector is held constant at 0.001




Figure 4-4: Inner and outer X-Y partition plots from DMC simulations of the 6-channel mZIP. The sequence of figures illustrates the effects of varying the phonon absorption coefficient at the non-instrumented germanium sidewall. Six absorption coefficients, 0.001, 0.003, 0.01, 0.03, 0.1, and 0.3, are presented, and the coefficient of phonon absorption at the instrumented Ge-Al interface is held constant at 0.03




Figure 4-5: Inner and outer X-Y partition plots from DMC simulations the 6-channel mZIP and Run 362 data. The best matching simulations are shown, and the parameters of these best matching simulations were $(A_{surf}, A_{side}) = (0.1, 0.03)$, $(0.15, 0.01)$, and $(0.15, 0.03)$. In the bottom right corner, the data is shown.

times of approximately 400, 250, 250, and 200 μs , respectively.


4.2 Phonon-quasiparticle downconversion, anharmonic decay, and TES critical temperature in the iZIP3

The next major round of DMC validation investigated the effects of including a simple phonon-quasiparticle downconversion model in the DMC (the second model described in Sec. 3.5.1) along with the effect of varying the anharmonic decay rate constant. The detector utilized for comparison in this validation study was one of the new (at the time) iZIP-style detectors (Detector G3D, UCB Test Facility Run 380), which has phonon sensors on both the top and bottom surfaces. The superconducting critical temperatures (T_c) of the TES channels on the two sides were markedly different from each other, allowing for additional tests of the TES model's ability to reproduce side-specific differences in the phonon pulse shapes caused by the T_c difference.


4.2.1 The iZIP3 detector

An image of the iZIP3 detector mask layout and accompanying schematic illustration are shown in Figs. 4-6(a) and 4-6(b), respectively. This detector was also fabricated using a Ge crystal 1" thick and 3" in diameter. The detector architecture included 4 phonon sensors on both the top and bottom faces: 2 inner sensors, each covering 180 degrees of the detector and extending from 0"-1.3" in radius (37.5% of the total surface area each), and 2 outer sensors whose boundaries are rotated 90 degree

relative to those of the inner sensors, each covering 12.5% of the total surface area. The sensor boundaries on the bottom side are also rotated 90 degrees relative to those on the top side, as shown in Fig. 4-6(b). The total aluminum coverage on each surface is 5.3%. The ionization rails are interleaved between the phonon collection rails, and separated into two concentric ionization channels. In the experiments that were compared with the DMC, the ionization channels on the two sides were biased at $\pm 2V$ relative to the phonon sensors.



(a) Figure 4-6: Illustration of the mask design of the iZIP3 detector. The thicker, ‘dashed-looking’ lines are the phonon channels, and each ‘dash’ is an individual QET, pictured to right of the image. The thinner lines are the ionization rails.



(b) Schematic of an iZIP3 detector, illustrating the 4 TES channels on both sides. The inner ionization channel extends throughout the three inner phonon channels, and the outer ionization channel occupies the same space as the outer phonon channel.

4.2.2 iZIP3 Experiment, Data Selection and Processing

Experiment

An SCDMS iZIP3 detector was exposed to a ^{133}Ba source at the UC Berkeley Testing Facility (Run 380) to produce an ER population throughout the detector. ^{133}Ba decays with a half-life of 10.53 years and emits a γ with an energy of 356 keV. The ionization sensors on the top and bottom surfaces were biased to -2V and +2V, respectively, and the phonon sensors act as grounds. The TES sensors of this particular iZIP3 detector (g3d) exhibited a high T_c of approximately 105 mK on one side of the detector, but a significantly lower T_c , near the achievable substrate temperature in the refrigeration system, on the other side. Because the substrate could not be cooled to well below T_c , actually measuring the T_c value is somewhat problematic. Maintaining a constant (low) current flux through the TESs and sweeping the substrate temperature using a heater provided a T_c value of 46 mK. A second measurement was performed by sweeping the bias voltage applied to the TES and converting the current to a corresponding Joule heating. As long as the current is swept slowly, the Joule heating input power can be equated to the output power to the substrate ($K(T^5 - T_b^5)$ in Eq. 3.17), and this method provides a T_c measurement of 60 mK. The discrepancy between these two measurement techniques provides an opportunity that the DMC may be able to resolve.

Data Selection

First, events with hole-side charge signals consistent with recoil energies of 125-175 keV were selected; this selection ensures high signal-to-noise in the phonon signal, but the energies are still sufficiently low that the TES sensors should not saturate. Because the penetration depth of a 356 keV γ in germanium is approximately 2 cm,

the source produces electron recoils throughout the 3"x1" detector, but the event distribution is non-uniform in position. Similar to the mZIP data selection process described in Sec. 4.1.2, the relative delays between the 20% risetimes in the two inner phonon channels on each sides were used as proxy variables for the event position, and the 2D (xdelay, ydelay) histogram was ‘flattened’ to produce a roughly position-uniform set of data events to compare with DMC.

Data quality cuts similar to those described in Sec. 4.1.2 were developed for the iZIP3 detector by other data analyzers and used as preselection cuts here. Surface events were excluded from the dataset using an ionization symmetry cut requiring the ionization signal measured on the electron side to be within 0.75 - 1.1 times the signal measured on the hole side. The asymmetry reflects an empirically observed asymmetry in the signals observed on the two sides, but can be justified by the fact that electrons are slightly easier to trap, and the oblique propagation of electrons can cause some fraction of the signal to be lost to the sidewalls. Events in which either outer ionization sensor registered any amount of energy inconsistent with noise were rejected.

Data Processing

Phonon pulse calibration in the iZIP3 is significantly simpler than in the mZIP. The observed phonon pulse falltimes in the iZIP3 detector are approximately 900 μ s, which is significantly longer than the ETF falltimes of the TESs themselves. Knowing that the power is linear with the current to first approximation (see Eq. 4.2; the ΔI_s^2 term is approximately 1% for the known values of the other parameters) and that the phonon power input per unit area to the TES sensors is highly uniform at late times, we can simply ‘align’ the tails of the phonon pulses after $\sim 500 \mu$ s




Figure 4-7: Illustration of the phonon pulse calibration process, aligning the late-time tails of the phonon pulses. Outer channel pulses are scaled by a factor of four to account for the difference in total area

as shown in Fig. 4-7. The current in the outer phonon channels is multiplied by a factor of four to account for the difference in coverage area between the inner and outer channels. Finally, both data and simulation phonon pulses are zero-phase filtered with a fourth-order Butterworth filter at 250 kHz; calling the filter eighth-order would be more accurate, as the zero-phase filtering process involves filtering the pulse forward and then backward using the same fourth-order filter to eliminate any phase artifacts. The high cutoff frequency of the filter ensures that only high-frequency noise is strongly affected by the filter.

4.2.3 Simulations

To prepare for this round of validations, a set of simulations similar to those presented in Sec. 4.1.3 was performed to determine an appropriate set of absorption coefficients; these iZIP3 simulations are not presented here. The iZIP3 metal coverage is 6.4%, about a factor of 7 lower than the 43.1% of the mZIP due to the need to interleave ionization sensors between phonon sensors. Comparing the phonon partitioning and decay times of the mZIP and iZIP3 simulations with the respective experiments provided best matches when A_{side} was .001 and A_{surf} was .1 for the mZIP and .015 for the iZIP3. Given the known aluminum coverage of the two different detectors, these values correspond to a phonon Al/Ge transmission coefficient of approximately 23%. This Al/Ge transmission coefficient and the phonon sidewall loss term were held constant in all of the simulations performed below.

In this set of simulations, the second of the phonon-quasiparticle downconversion models described in Sec 3.5.1 was implemented. In this model, a phonon that is absorbed into the aluminum fins deposits a maximum discrete amount of energy corresponding to a particular number of broken Cooper pairs (i.e., E_{gap} , $2E_{gap}$, etc.); this maximum amount of absorbed energy is labeled $\mathcal{L}$. If the incident phonon energy minus $\mathcal{L}$ is greater than E_{gap} , a lower-energy phonon is returned to the crystal and can be absorbed later, and if the remaining energy is less than E_{gap} , then any remaining energy is simply removed as the produced phonon would be incapable of breaking Cooper pairs. Finally, if the incident energy is less than $\mathcal{L}$ but greater than E_{gap} , the phonon is completely absorbed. This model is a simplification of the more realistic cascade model implemented later. DMC simulations were run with $\mathcal{L} = 2, 3, 4, 6$, and 8 times E_{gap} .

The effect of the low T_c on Side 2 of the detector also allows for tests of the

TES model implemented in the DMC. The dependence of the TES response time on T_c in Eq. 3.20 effectively disappears when T_c is significantly larger than the bath temperature T_b ($\tau_{ETF} \sim \frac{1}{1-(\frac{T_b}{T_c})^5}$), but as T_c approaches T_b , the TES response rapidly slows down. The T_c of the Side 2 TESs in the DMC simulations was varied among 50, 55, and 60 mK to determine the effects of T_c on the TES pulse shapes.

Finally, the effects of varying the anharmonic decay constant, γ_A , on the phonon pulse shapes and energy partitioning is investigated. The value $1.61 * 10^{-55} s^4$ presented in Eq. 3.9 can be found in the literature, and the values $9.05 * 10^{-56}$, $5.37 * 10^{-56}$, $2.86 * 10^{-56}$, and $1.61 * 10^{-56}$ are also investigated. Reducing the anharmonic decay constant effectively keeps the phonons in the quasidiffusive transport regime at higher energies for a longer period of time and should thus produce longer phonon pulse risetimes.

4.2.4 Results

A variety of pulse shape parameters were identified, and the following six were selected for their sensitivity to the variable DMC parameters: the energy fraction partitioned into the inner and outer TES channels, the phonon pulse decay times, the fastest and slowest 10%-40% risetimes among the 8 phonon pulses, the ratio of the peak instantaneous power to the average power in the primary channel, and the time when position sensitivity is lost (judged by the time at which all 8 phonon pulses are within 10% of the mean phonon pulse; illustrated in Fig 4-8). These parameters are not completely independent, but each provides a separate window into the effects of the varied DMC parameters. Seventy-five individual DMC variations were performed, and the distributions of these 6 observables from certain illustrative ‘corner cases’ with the extreme values of T_c , $\mathcal{L}$ and γ_A will be presented alongside




Figure 4-8: Illustration of the time to position independence. The red line indicates the point at which all pulses are within 10% of the mean pulse, indicating that the phonon signal no longer contains position information.

the best matching set of parameters: $\mathcal{L} > 4E_{gap}$, $T_c = 55$ mK, and $\gamma_A = 1.61$ or $0.905 \times 10^{-55} s^4$ (these values are justified in Figs 4-15 and 4-16, which are discussed later).

The energy partitioning fraction is sensitive to $\mathcal{L}$ and γ_A but relatively insensitive to T_c , which is expected because T_c is a TES parameter that does not affect phonon absorption (though different amounts of energy loss from the TES to the substrate does produce some slight differences). Fig. 4-9 shows the measured and simulated energy partitioning histograms for the best matching parameter set, and the sets $(\mathcal{L}, T_c, \gamma_A) = (2E_{gap}, 55 \text{ mK}, 1.61 \times 10^{-55} s^4)$, $(8E_{gap}, 55 \text{ mK}, 1.61 \times 10^{-55} s^4)$, $(8E_{gap}, 55 \text{ mK}, 1.61 \times 10^{-56} s^4)$. The energy partitioning into the outer channels, which occupy one-fourth of the detector area coverage, forms the low peak around 0.05, while the inner channels are the higher distribution around 0.2. The inner channel distribution is doubly-peaked because the primary inner channels (small peak at higher partitioning, around 0.22) are lumped with the secondary inner channels (peak




Figure 4-9: Histograms of the fraction of energy in the eight phonon channels from data (black solid line) and four different simulation runs (red dashed line). The smaller outer channels each collect $\sim 5\%$ of the energy, while the larger inner channels are distributed around 20%. Increasing $\mathcal{L}$ and increasing γ_A both increase the degree of partitioning (the breadth of the histogram), as expected.

around 0.2 and tail down to approximately 0.17). Increasing $\mathcal{L}$ and increasing γ_A both increase the degree of partitioning (the breadth of the histogram), as expected.

The phonon pulse decay time is sensitive primarily to $\mathcal{L}$, and is also sensitive to γ_A at the lowest value of $\mathcal{L}$. At long timescales, information about γ_A becomes washed out as phonons downconvert to the ballistic regime, and τ_{ETF} is significantly shorter than the decay time even at the lowest T_c value. Figure 4-10 presents the decay time histograms at 55 mK and $(\mathcal{L}, \gamma_A) = (4E_{gap}, 1.61 * 10^{-55} s^4)$, $(8E_{gap}, 1.61 * 10^{-55} s^4)$, $(2E_{gap}, 1.61 * 10^{-55} s^4)$, and $(2E_{gap}, 1.61 * 10^{-56} s^4)$. Increasing $\mathcal{L}$ from 4 to 8 has little

effect, as most phonons at late times have downconverted to $\sim 3E_{gap}$, but reducing $\mathcal{L}$ to 2 lengthens the decay time. Reducing γ_A further increases the decay time of the pulses at $\mathcal{L} = 2E_{gap}$, though at higher $\mathcal{L}$ values this effect disappears. The DMC decay time histograms are substantially less wide than those measured in data. The data does contain noise that was not modeled in the DMC in these simulations, and the data histograms may be further widened due to mild non-linearities in the responses of the physical ADCs used in the experiment. These non-linearities were characterized, and later investigations included correction terms to linearize the ADC response when comparing data with DMC.




Figure 4-10: Decay time histograms from data (black solid line) and four different simulation runs (red dashed line). Increasing $\mathcal{L}$ from 4 to 8 has little effect, as most phonons at late times have downconverted to $\sim 3E_{gap}$, but reducing $\mathcal{L}$ to 2 lengthens the decay time. Reducing γ_A further increases the decay time of the pulses at $\mathcal{L} = 2E_{gap}$, though at higher $\mathcal{L}$ values this effect disappears.

The slowest 10%-40% risetimes occur only on the low- T_c side of the detector due to its reduced τ_{ETF} , making this observable the most strongly sensitive to the value of T_c . Varying γ_A while holding T_c constant has a slight effect, of order 10%, and varying $\mathcal{L}$ does not produce an observable effect. Figure 4-11 presents histograms of the slowest risetimes at the three different T_c values and $(\mathcal{L}, \gamma_A) = (4E_{gap}, 1.61 \times 10^{-55} s^4)$, and it is obvious that the as T_c increases, the risetimes become faster, as expected due to the increasing τ_{ETF} .




Figure 4-11: Histograms of the slowest 10%-40% phonon pulse risetimes at the three T_c values tested: 50, 55, and 60 mK. Data is shown in the black solid line, and DMC results are presented as a red dashed line. It is clear that the pulses rise more rapidly as T_c increases. 55 mK provides the best match, though it a value somewhere between 50 and 55 mK would likely improve the agreement.

In contrast, the fastest 10%-40% risetimes are fairly independent of T_c , as they

appear exclusively on the high- T_c side of the detector, and are mildly dependent on γ_A and $\mathcal{L}$. Matching these fast risetimes is more difficult than any of the previous parameters, as the risetimes of these high- T_c TESs depend on precise modeling of electrothermal dynamics, which may include phase separation (not modeled at all here). However, reasonable agreement is achieved, and we can still study how these risetimes vary as the phonon propagation and absorption physics are changed. Fig. 4-12 presents the fastest 10%-40% risetimes for the best matching case and for $(\mathcal{L}, \gamma_A) = (8E_{gap}, 1.61 \times 10^{-55} s^4)$, $(2E_{gap}, 1.61 \times 10^{-56} s^4)$, $(8E_{gap}, 1.61 \times 10^{-56} s^4)$, all at $T_c = 55$ mK. Interestingly, increasing $\mathcal{L}$ slows down the fastest risetimes, even though more energy ought to be absorbed at these early times, indicating that the TESs are exhibiting a slew rate. Similarly, decreasing γ_A speeds up the fastest risetimes, though one would expect phonon absorption to slow down due to increased quasidiffusive propagation. However, the variation among these corner cases is quite small, indicating that untuned TES dynamics are probably dominating the measurement.

The time to position independence exhibits a degenerate dependence on T_c and γ_A , as it is influenced by both the phonon system and the recovery time of the low- T_c TESs, and is only mildly sensitive to the value of $\mathcal{L}$. As γ_A increases, the phonons remain quasidiffusive for a longer period of time, retaining information about the initial event position as the phonon gas expands slowly. Decreasing T_c also extends the time until the observed phonon signals converge, because the low- T_c TESs have a slower response time as evident in the different timescales in Fig. 4-11. Fig. 4-13 shows histograms of the position independence times at $\mathcal{L} = 4$ and at $(T_c, \gamma_A) = (55 \text{ mK}, 1.61 \times 10^{-55} s^4)$, $(50 \text{ mK}, 1.61 \times 10^{-55} s^4)$, $(50 \text{ mK}, 1.61 \times 10^{-56} s^4)$, and $(60 \text{ mK}, 1.61 \times 10^{-56} s^4)$. The expected behavior with γ_A and T_c is evident, though the histograms at the best overall parameter set (upper left) disagree by $\sim 10\text{-}20\%$.

The peak instantaneous phonon power normalized by the mean power is a mea-




Figure 4-12: Histograms of the fastest 10%-40% phonon pulse risetimes in data (black solid line), and DMC (red dashed line). The mild variation among the extreme cases presented here (the lowest and highest values of both $\mathcal{L}$ and γ_A are presented) indicates that TES dynamics probably dominate the behavior of this parameter.

sure of the pulse ‘peakiness’. Being primarily dictated by phonon propagation and absorption physics, this metric is strongly dependent on $\mathcal{L}$ and γ_A and only and weakly dependent on T_c . Figure 4-14 presents histograms of the peak instantaneous power over the mean power at $T_c = 55$ mK and $(\mathcal{L}, \gamma_A) = (4E_{gap}, 1.61 * 10^{-55} s^4)$, $(4E_{gap}, 1.61 * 10^{-56} s^4)$, $(2E_{gap}, 1.61 * 10^{-55} s^4)$ and $(8E_{gap}, 1.61 * 10^{-55} s^4)$. The results show that pulses become less peaky at high γ_A , which is expected as the phonons remain quasidiffusive longer, preventing a large influx of phonons at early times. As $\mathcal{L}$ increases, the pulses become more peaky, which is also expected as the high-energy early phonons are absorbed more efficiently at high $\mathcal{L}$. The best overall parameter set




Figure 4-13: Histograms of the time until phonon signals become position-independent (i.e., when all 8 phonon traces converge) in data (black solid line), and DMC (red dashed line). The observed position independence time increases with increasing γ_A and decreasing T_c , as expected, though the predictions from the best matching overall parameter set (upper left) are in disagreement with the data by $\sim 10\text{-}20\%$.

is presented in the upper left, and exhibits fairly good agreement, though with a less prominent peak and a slightly higher tail than the data, indicating increased peakiness. This observation may be related to un-tuned TES physics and is reasonable given the faster 10-40% risetimes observed in Fig. 4-12.

The metric used to compare the measured and simulated distributions of each of the tested parameters is a simple least-squared residual (R) between the histogrammed distributions, which are normalized to sum to 1 to account for differences




Figure 4-14: Histograms of the peak phonon power and the in data (black solid line), and DMC (red dashed line). The mild variation among the extreme cases presented here (the lowest and highest values of both $\mathcal{L}$ and γ_A are presented) indicates that TES dynamics probably dominate the behavior of this parameter.

in the number of events between data and simulation. The 6 observables studied each exhibit different dependences on different sets of the DMC parameters, and these results can be combined to determine the more favored regions of the DMC parameter space. Figure 4-15 presents the sum of squared residuals for all parameter sets and all 6 observables. The x axis is γ_A , and the y-axis is $\mathcal{L}$; the y-axis is repeated 3 times for each of the 3 T_c values. The R value for each of the 6 observables is bounded in a red rectangle for each set of input parameters. Inside each red rectangle are 6 colored cells indicating the value of R , and darker colors indicate lower R values and thus a better match. The legend below the figure indicates how the 6 cells corresponding

to the 6 observables are arranged within each bounding box. To get a better sense of the best matching overall parameter sets, these residuals are combined them into a parameter $\mathcal{R} = 1 - \Pi(1 - R)$. A map of $\mathcal{R}$ versus the DMC input parameters is shown in Fig. 4-16. This map clarifies the results quite well, and indicates that the preferred DMC parameters are: $\mathcal{L} \geq 4$, $T_c=55$ mK, and $\gamma_A = 1.61 * 10^{-55}$ or $9.05 * 10^{-56} s^4$. The value $T_c=55$ mK is between the two measured values of 60 and 46 mK. Prior measurements [81] have found $\gamma_A = 1.61 * 10^{-55} s^4$, and this value provides good agreement in this study, so this value for γ_A has become the standard value utilized in all future DMC studies. This simple phonon-quasiparticle downconversion model was eventually replaced by a more sophisticated phonon-quasiparticle cascade model, and the favored $\mathcal{L}$ values and the known 300 nm thickness of the aluminum QP traps are in agreement with a ~ 700 nm phonon-QP interaction length [85].




Figure 4-15: Sum of squared residuals between measured and simulated histograms for various pulse shape quantities; darker colors indicate better agreement. Each box bounded by red lines presents a single variation of T_c , $\mathcal{L}$, and anharmonic decay constant, and the 6 individual rectangles within each box correspond to the 6 pulse shape parameters under investigation, as indicated in the key below the figure. There are 75 DMC variations presented here, and the squared residuals for each pulse shape parameter are normalized such that the sum over the 75 variations is 1.





Figure 4-16: Overall match between data and simulation, measured by combining the sum of squared residuals of all 6 pulse shape parameters presented in Fig. 4-15.

4.3 Phonon-QP downconversion and diffusion in the iZIP4


The final major round of DMC validation investigated the effects of including the full Kaplan phonon-quasiparticle downconversion model in the DMC (the third model described in Sec. 3.5.1) and estimating the timescale of quasiparticle diffusion from the QPTs into the TESs. The detector utilized for comparison in this validation study is the iZIP4 detector (Detector G48, UCB Test Facility Run 400), which has phonon sensors on both the top and bottom surfaces. Unlike the iZIP3 presented in Sec. 4.2, the T_c s of the TES channels on the two sides were not markedly different from each other, with channels on both sides exhibiting T_c values from 90-100 mK.

4.3.1 The iZIP4 detector

A photo of the iZIP4 detector and accompanying schematic illustration are shown in Figs. 4-17(a) and 4-17(b), respectively. This detector was fabricated using a Ge crystal 1" thick and 3" in diameter. The detector architecture included 4 phonon sensors on both the top and bottom faces: 3 inner sensors, each covering 120 degrees of the detector and extending from 0"-1.3" in radius (25% of the total surface area), and 1 outer sensor covering a full 360 degrees for the remaining 25% of the total surface area. The sensor boundaries on the bottom side are rotated 60 degrees relative to those on the top side, as shown in Fig 4-17(b). The total aluminum coverage on each surface is 6.4%. The ionization rails are interleaved between the phonon collection rails, and separated into two concentric ionization channels. In the experiments that were compared with the DMC, the ionization channels on the two sides were biased at +/-2V relative to the phonon sensors.



(a) Figure 4-17: Image of the iZIP4 detector (detector name G51), with rules to show scale. The arrangement of the TES channels is clearly visible, though the interleaved ionization lines are too thin to be seen in the image.



(b) Schematic of an iZIP4 detector, illustrating the 4 TES channels on both sides. The inner ionization channel extends throughout the three inner phonon channels, and the outer ionization channel occupies the same space as the outer phonon channel.

4.3.2 iZIP4 Experiment, Data Selection and Processing

Experiment

An SCDMS iZIP4 detector was cooled down at the UC Berkeley Testing Facility (Run 400) with a collimated ^{109}Cd source included in the detector mounting apparatus. The ^{109}Cd source emits electrons and photons at 22 (83.4%), 25 (13.7%) and 88 keV (3.6%); photons are emitted exactly at these energies, while electrons can be emitted with any kinetic energy up to these energies. Electrons and photons at these energies have very short penetration depths, providing a large surface event sample. The bias state of the detector was varied over the course of the run, and for this validation we compare the surface event yield bands in data series taken with the detector biased at -2/0 V (‘hole-collecting’ mode) and +2/0 V (‘electron-collecting’ mode). The detectors were also exposed to external ^{133}Ba and ^{252}Cf sources to provide bulk ER

and NR samples.


Data Selection

Data quality cuts similar to those described in Sec. 4.1.2 were developed for the iZIP4 by other data analyzers and were applied as preselection cuts. A recoil energy window of 10-200 keV was used (10-90 keV for the surface events, as the ^{109}Cd source only emits up to 88 keV). Data was selected only from data sets run in -2/0 V bias on the top side (the side with the ^{109}Cd source) and +2/0 V bias on the opposite side. Bulk ERs and bulk NRs were selected using a combination of ionization symmetry, radial fiducial volume, and ionization yield cuts, and surface ERs were identified by having ionization signals observed only on the top side of the detector and reduced yields (as discussed in Sec. 4.5).


Data Processing

Because the 8 channels each cover the same area, the relative calibration of the phonon pulses is rather simple. In any event, regardless of event position, all channels should receive some minimum amount of position-independent energy. Thus, if the fraction of energy deposited into each channel is histogrammed, each channel should exhibit a high-energy peak (events in which that channel was the primary channel) and a low-energy peak (non-primary events), with the exception of the outer channel, as applying the ionization fiducial volume cut ensures that the outer channel is never primary. The low-energy peaks can be fit to Gaussian functions (windowed to prevent the high-energy events from affecting the fit) and aligned. This process is illustrated in Figs. 4-18(a) and 4-18(b).

Post-calibration, phonon pulse shape parameters can be extracted. Obviously,



(a) Figure 4-18: iZIP4 fractional phonon energy by channel, barium data, pre-calibration



(b) iZIP4 fractional phonon energy by channel, barium data, post-calibration

the relative calibration does not affect individual pulse shapes, but extracting the pulse shape parameters of side-summed and total phonon pulses requires proper relative channel calibration. To extract these timing parameters, phonon pulses are zero-phase filtered with a 50 kHz, fourth-order Butterworth filter. The timing parameters are extracted using a ‘walking’ algorithm - beginning at the peak of the pulse, the algorithm walks down the rising or falling edge until it finds the first point crossing some pre-defined threshold (e.g., the first bin in which the phonon signal is below 70% of the maximum). Once this bin is found, the crossing point is estimated by interpolation between the first bin below threshold and the adjacent above-threshold bin. Clearly, this process is sensitive to the signal-to-noise ratio, and at low energies, noise will push the rising-edge timing parameters to later times and the falling-edge parameters to earlier times.

4.3.3 Simulations

QP Cascade

The subject of phonon-quasiparticle interactions in superconductors is treated in [85]. The article treats the timescales of a wide variety of physical processes, including quasiparticle (QP) recombination, QP branch mixing (between QP states above and below the Fermi momentum), QP scattering, phonon scattering, and phonon Cooper pair breaking. The most current DMC model incorporates only the rate of Cooper pair breaking by supragap phonons and the energy distributions of the resulting QPs and any new phonons emitted by the excited QPs. Future models ought to improve on this aspect of the DMC, particularly by describing QP diffusion to the TESs, which will likely require implementing descriptions of the QP scattering and recombination rates. Figure 4-19 presents the phonon Cooper pair breaking rate (y-axis) vs. the normalized temperature T/T_c in a superconducting material near thermal equilibrium from [85]. The labeled lines present the pair-breaking rate at various phonon energies, described in units of the material-specific superconducting gap. The rate in Figure 4-19 is normalized by a material-specific constant T_0^{ph} , which is equal to 242 ps in aluminum. The T_c of aluminum is 1.2 K, so at the ~ 50 mK operating temperature of the SCDMS detectors, the pair-breaking rate is effectively temperature-independent, and we need only consider the energy dependence (indeed, Δ is itself temperature-dependent near T_c but effectively temperature-independent in the detector's operating regime).

The quasiparticle interaction rate is parametrized in the DMC as the rate at $E_\phi = 2\Delta$ (where E_ϕ denotes the phonon energy) plus a linear energy-dependent term: $\tau^{-1}(E_\phi) = \tau^{-1}(2\Delta) + A(E_\phi - 2\Delta)$. The tests performed in this section used $\tau(2\Delta) = 242, 200, 150, 100, 50$ ps and $A = 0, 0.3, 0.5, 1, 1.5$. Changing the QP




Figure 4-19: Rate of interaction between phonons and quasiparticles versus temperature. The different lines correspond to different phonon energies. The important consideration for the DMC is to model $\tau_B(E)$, which is the rate at which phonons break Cooper pairs into quasiparticles.

interaction rate also affects the decay times of the pulses, so the reflectivity of the Al-aSi-Ge interface is adjusted to compensate for this effect and keep the decay times of the simulated pulses at $750 \mu\text{s}$, in agreement with the data.

It is clear that the observed phonon partitioning does not vary incredibly strongly as these parameters are changed, and with no compelling indication that any particular set of parameters is significantly better than any others, we choose the theoretically predicted phonon pair-breaking time constant of 242 ps, and ignore any slope in this rate for simplicity, as the slope has a relatively minor effect.

QP Diffusion Time

With the phonon pair-breaking rate set, the next step is to test the mean time for the created quasiparticles to propagate from the Al fins to the W TESs. The DMC introduces a delay between quasiparticle creation and introduction of the energy to the TES, and the delay is randomly drawn from an exponential distribution with mean time τ_{QP} . This delay does not affect the partitioning at all, but does affect the timing structure of the measured phonon pulse. In these tests, the delay time was set in steps of $3 \mu\text{s}$ from 0 to $30 \mu\text{s}$, and $35 \mu\text{s}$ is also included.

A set of timing parameters must be selected for comparison against data. The walking algorithm used to extract raw pulse timing parameters was described in Sec. 4.3.2, and because there is some trigger ‘jitter’ associated with the real data, it is best to use differences between timing parameters rather than raw values. For this comparison, I chose the 10%, 40% 70%, and 100% points on the rising edge of the pulse, and the 80% point on the falling edge. All sets of differences that can be generated with this set were tested using the primary phonon pulse on each side, the secondary pulses on each side, the sums of the 4 phonon pulses on each side (the




Figure 4-20: χ^2 between measured and simulated phonon pulse timing parameter histograms (10%, 40% 70%, and 100% points on the rising edge of the pulse, and the 80% point on the falling edge) versus time for quasiparticle diffusion from the Al fin to the TES. A broad minimum is evident between 15 and 30 μs .

‘side-summed’ phonon pulse), and the sum of all 8 phonon pulses (the total pulse), and bulk ERs, surface ERs on both sides, and bulk NRs were individually compared against data.

Figure 4-20 reveals a fairly broad minimum around 15-30 μs , and 18 μs is chosen for now as the standard for simulations utilized in the rest of this thesis.

4.4 Performance of current DMC mZIP, iZIP3, and iZIP4 models

I conclude the phonon physics validation section of this chapter with a set of plots meant to serve as a demonstration of what the DMC is and is not currently good at reproducing. I begin with comparing the G48 iZIP4 detector data taken in UC Berkeley Run 400 against DMC simulations using the parameters shown in Table 4.4. Any parameters not listed in the table are assumed to be default values available in the DMC template in the SuperCDMS CVS repository, many of which are tuned and validated in the preceding work in this chapter.

Figures 4-21 and 4-22 show a series of figures with data on the left compared against simulation on the right. The first two rows of figures show side-specific partition plots, similar to those shown in the mZIP comparison at the start of this chapter (Figures 4-3 and 4-4). Bulk ERs, Bulk NRs, and surface events on both Side 1 and Side 2 are shown. The DMC correctly reproduces the scale and shape of the partition plot, as well as the fact that same-side surface events are strongly partitioned while opposite-side surface events are weakly partitioned. Noise is not included in the simulations, which may explain why the simulations look substantially cleaner than the corresponding data plots.

The third row in this set of figures is the partitioning and relative delay between Side 1 and Side 2 of the detector (relative delay defined as the time between the 20% risetime of the total phonon pulse on Side 1 and that on Side 2). The DMC predicts substantially more partitioning and more temporally-separated pulses than are observed. This phenomenon is the subject of ongoing study within the DMC group.

Al coverage, top	0.064
Al coverage, bot	0. 064
Al absorption prob.	0.442
Al interaction length	720 nm
TES Parameters, top	Un-tuned default iZIP4 template
Quasiparticle diffusion time	18 μ s

Table 4.2: Relevant parameters of iZIP4 simulation used to make Figures 4-21 and 4-22. All parameters not listed are the default values available in the iZIP4 DMC template in the SuperCDMS CVS repository.

The next set of figures shows the delay plots on Side 1 and Side 2. The delay plots are constructed similar to the partitioning plots (Eq. 4.3), with the 20% risetime of the phonon pulse in each channel taking the place of the fractional phonon energy. Essentially, this forms an event position estimate by looking at the relative phonon arrival times in each of the channels. The DMC produces a slightly wider distribution than the data, but this parameter is strongly dependent on TES dynamics, and the agreement is satisfactory. The surface events on Side 1 in the data are produced by localized sources, whereas they are spread throughout the surface in the DMC, so the DMC does not exhibit the Side 1 delay localization for Side 1 surface events observed in the data. Finally, the 10-70% risetimes of the total phonon pulses on Side 1 and Side 2 are plotted against each other, and the DMC accurately reproduces this distribution for each event class, though with less spread than is present in the data.




Figure 4-21: Comparison of iZIP4 phonon parameters from data (left) and simulation (right)




Figure 4-22: Comparison of iZIP4 phonon parameters from data (left) and simulation (right)

The next set of figures (4-23 and 4-24) show a series of figures from iZIP3 simulations and data from the G3D detector taken at UC Berkeley in Run 380. Data is again on the left, and simulation is on the right. Table 4.4 provides the parameters of the simulation. The iZIP3 detector was not included in the final rounds of validation, so it is interesting to check agreement between data and sim using parameters tuned on the iZIP4 detectors. The first two rows show the “inner” and “outer” partition plots; energy partitioning among the 4 inner channels (2 top, 2 bottom) and the four outer channels. The DMC generally correctly reproduces the scale and shape of the inner partition plot, though the Side 2 surface event partitioning is somewhat stronger than in the data. The outer partition plots are somewhat difficult to compare; they are fairly featureless and clustered around (0,0), though the surface events do exhibit somewhat stronger partitioning that is reasonably well reproduced.

The third row in this set of figures again shows the partitioning and relative delay between Side 1 and Side 2 of the detector (relative delay defined as the time between the 20% risetime of the total phonon pulse on Side 1 and that on Side 2), and emphasizes the issues noted in the iZIP4 comparion. The DMC predicts substantially more partitioning and more temporally-separated pulses than are observed

The next set of figures shows the inner and outer delay plots, and the DMC again produces a slightly wider, more well-defined distribution than the data. Side 1 and Side 2 (which provide the Y and X inner delay and the X and Y outer delay, respectively) have visibly different ranges in the delay plot, caused by the difference in TES T_c on the two sides, and the DMC reproduces this feature. Finally, the 10-70% risetimes of the total phonon pulses on Side 1 and Side 2 are plotted against each other. The DMC accurately reproduces the shape of the distribution for each event class, but overall predicts faster Side 1 risetimes and slower Side 2 risetimes than are observed.

Al coverage, top	0.053
Al coverage, bot	0. 053
Al absorption prob.	0.442
Al interaction length	720 nm
TES T_c , top	105 mK
TES T_c , bot	55 mK
Quasiparticle diffusion time	18 μ s

Table 4.3: Relevant parameters of iZIP3 simulation used to make Figures 4-23 and 4-24. All parameters not listed are the default values available in the iZIP3 DMC template in the SuperCDMS CVS repository.




Figure 4-23: Comparison of iZIP3 phonon parameters from data (left) and simulation (right)




Figure 4-24: Comparison of iZIP3 phonon parameters from data (left) and simulation (right)

Finally, we turn to the mZIP. The mZIP has not been tuned at all since the nascent stages of DMC development, when the first tests of phonon absorption physics were performed, so observing good agreement here would be a solid confirmation of much of the work performed on other detector architectures. The set of figures in 4-25 show data from mZIP3 simulations (parameters in Table 4.4) on the right and data from the G9F detector taken at UC Berkeley in Run 363. No surface event tagging is available in the mZIP, so we compare only ERs and NRs. The first row shows the partition plots, which are in good agreement. The middle row shows the delay plots, which are slightly wider in data than in DMC, particularly for neutrons. It is interesting that in both data and DMC, ‘caustics’ outlining the channel boundaries are visible in the delay plot. Finally, I look at two ER/NR discriminators, the 10-40% risetimes of the fastest phonon pulse and of the summed phonon pulse. The DMC reproduces the range of the bulk ER histograms, but the DMC histogram’s shape is markedly weighted towards faster risetimes than the data. For NRs, the DMC distribution of the fastest risetime is shifted by $\sim 1\text{-}2\ \mu\text{s}$ relative to the data, and the DMC also fails to reproduce the slow tail in the risetime of the total phonon pulse seen in data. These delay parameters are strongly dependent on TES dynamics, which have not been tuned in these simulations, so the rough agreement is not totally unexpected.

Al coverage, top	0.39
Al coverage, bot	0.10
Al absorption prob.	0.30
Al interaction length	720 nm
TES Constants	default in template
Quasiparticle diffusion time	18 μ s

Table 4.4: Relevant parameters of mZIP simulation used to make Figures 4-25. All parameters not listed are the default values available in the mZIP DMC template in the SuperCDMS CVS repository.




Figure 4-25: Comparison of mZIP phonon parameters from data (left) and simulation (right) (except bottom hists, with data and sim on the same plot)

4.5 Validation of Charge Propagation

4.5.1 Basic Charge Propagation Validations

Luke Phonon Energy conservation

The primary method of separating electron and nuclear recoils in the ZIP detectors is through the ratio of the measured ionization and phonon energies. To determine the recoil energy of an event, the Luke phonon population produced as the charge carriers drift across the detector must be subtracted from the total phonon signal. Thus, it is crucial to ensure that the microscopic charge carrier propagation simulation conserves energy in the emission of Luke phonons so that the total predicted phonon signals of ERs and NRs match the expected signals under a given drift field. Figures 4-26 and 4-27 show the fractional errors in the Luke and total phonon signals versus applied drift fields over a range from 0.04-5 V/cm (0.1-12.7 V applied across a 2.54 cm detector). The fractional error here is defined as the phonon signal predicted by DMC divided by the expected signal, minus 1. The total expected phonon energy is the recoil energy plus the Luke phonon energy, and the Luke phonon energy is simply the number of charge pairs produced times the traversed voltage ($E_{Luke} = \frac{E_{recoil}}{E_{eh}} * V$), and the drift field is the applied voltage divided by 2.54 cm. Though the Luke phonon energy becomes underpredicted by nearly 10% at very low drift fields, the total phonon energy predicted by the DMC is within 0.5% at all applied drift fields.

Charge Carrier Drift Velocities

Because of the sampling rate employed in the CDMS experiment, the detectors are not particularly sensitive to the drift velocities of the charge carriers; the sampling rate is 1.25 MHz (800 ns sampling time), while the measured drift velocities over




Figure 4-26: Fractional error in Luke phonon energy produced by DMC vs. applied drift field. The expected Luke phonon energy is simply the number of charge carriers times the applied potential (drift field * 2.54 cm).




Figure 4-27: Fractional error in total phonon energy produced by DMC vs. applied drift field. The expected total phonon energy is defined as the recoil energy plus the expected Luke phonon energy.

the range of 1-4 V/cm drift fields that have used in various CDMS detectors are greater than 10 km/s, meaning that charges cross the 1 cm or 2.54 cm-thick detectors within just one or a couple of samples. However, comparing the DMC's predicted drift velocities against measurements performed in specially-instrumented CDMS detectors at test facilities provides a useful test of the DMC's charge propagation physics. Figures 4-28 and 4-29 depict the DMC's predicted electron and hole drift velocities vs. applied drift field against theoretical predictions and measurements from [107, 108, 109, 110, 111]. The theoretical basis of the DMC's charge propagation algorithm is taken from [100].




Figure 4-28: Electron drift velocities vs. field. DMC is the solid blue line. Measured data is from [107]. Comparison theoretical predictions are from [108, 109, 110]. Whether a detector is n- or p-type does not affect the DMC’s charge propagation algorithms in its current iteration, which clearly represents a direction for future work.




Figure 4-29: Hole drift velocities vs. field. Measured data is from [107]. Comparison theoretical predictions are from [108, 109]. Whether a detector is n- or p-type does not affect the DMC’s charge propagation algorithms in its current iteration, which clearly represents a direction for future work.

Oblique Electron Propagation

As discussed in Sec. 3.6.1, the anisotropic structure of the electron conduction band in germanium produces the non-intuitive behavior that individual electrons do not necessarily drift parallel to the applied field. The minimum energy electron states are located in the 4 ellipical L valleys along the $\langle 111 \rangle$ directions in k-space, and the effective electron mass is a factor of 19.5 larger in the directions perpendicular to $\langle 111 \rangle$ than in the parallel direction [109]. Thus, if a field is applied in the $\langle 100 \rangle$ direction, the electrons will propagate within the 4 L valleys along oblique directions

approximately 33 degrees from the applied field. The Herring-Vogt transform described in [99] should appropriately handle the anisotropic electron mass tensor, so verifying the electron propagation angle is an important test that this transform was implemented correctly. Figure 4-30 shows the final positions of simulated electrons in the x-y plane after propagating from an event located at (0, 0, 1.27 cm) to the surface at $z = 0$. The blue dots indicate the final electron positions at $z = 0$, while the red line outlines the expected electron propagation ‘cone’ at .825 cm ($1.27 \text{ cm} * \tan(33 \text{ degrees})$). The population of events outside of the circle are attributed to stochastic electron diffusion, while the population well inside of the circle are attributed to both diffusion and intervalley electron transitions. Figure 4-31 presents a histogram of the electron propagation angle calculated from the final x-y positions and the known 1.27 cm of drift in the Z direction. This histogram exhibits a clear peak at the expected propagation angle of 33 degrees.

4.5.2 Charge Carrier-Surface Interactions

Electron recoils near the surfaces of the detector are known to exhibit reduced ionization yields. This reduced ionization yield has generally been explained [90, 105] as the result of ‘back-diffusion’ of wrong-sign charge carriers into the local electrodes, as the initial charge ‘ball’ formed in the region near the initial recoil consists of charge carriers with relatively high kinetic energies before they shed their initial momentum into phonons on distance scales from $\mathcal{O}(1 - 10)\mu\text{m}$. An enhanced density of charge traps near the surfaces has also been proposed as a possible mechanism. The DMC’s ability to model the low-yield dead layer in the detectors is crucial to its utility as an analysis tool, as these low-yield electron recoils represent the primary background source in CDMSII-style ZIPs and mZIPs and can also fake NRs in the iZIP detectors




Figure 4-30: Validation of simulated oblique electron propagation. The blue dots indicate the final electron positions at $z = 0$, while the red line outlines the expected electron propagation ‘cone’ at .825 cm ($1.27 \text{ cm} * \tan(33 \text{ degrees})$). The population of events outside of the cone are attributed to stochastic diffusion, while those inside the cone are attributed to both diffusion and intervalley transitions.

in comparatively rare event topologies involving multiple electron recoils at both surfaces of the detector (see 5). Because the DMC creates all charge carriers upon initialization rather than modeling the entire creation and evolution of the charge carrier population as a cascade process, the initial conditions of the carriers and the criteria for absorbing a carrier at the detector surface play critical roles in modeling the surface event yield, and this section discusses tests of how these parameters affect the surface event yield and comparison with data from an iZIP4-style detector.




Figure 4-31: Histogram of electron propagation angle relative to the applied field in the Z direction. The expected propagation angle of 33 degrees is clearly visible in the peak. Angles less than 33 degrees are a consequence of intervalley transitions and electron diffusion.

Experiment

The data was taken from the same experiment described in Sec. 4.3.2 (Detector g48, R400). Data quality cuts were essentially the same, with the exception that surface events were selected from data series taken in both the -2/0 V and +2/0 V bias states on the side of the detector with the ^{109}Cd source.

Simulations

The initial conditions of the charge carriers in the DMC include only the initial positions and momenta of the carriers. In the studies discussed here, all charge carriers

are created precisely at the recoil position, and the directions of the initial carrier momentum are randomized, so the only initial condition varied is the magnitude of the initial carrier momentum. The average energy required to create an electron-hole pair in Ge is 3 eV, but the bandgap is only 0.75 eV (Sec. 3.3.2), and so the carriers can be initialized with up to 2.25 eV of kinetic energy per pair (the DMC will appropriately reduce the number of initial phonons as the initial carrier momentum is increased).

The DMC defines three criteria for absorption when a charge carrier reaches one of the detector surfaces. First, a charge is absorbed if it contacts the surface at an electrode. The DMC is not ‘aware’ of the electrode geometry, so in practice the carrier is absorbed when its local electric potential is approximately equal to the extremal values in the potential map at the detector surface (electrons are absorbed at the potential maxima, holes at the minima). Second, the DMC allows the implementation of a momentum barrier at the surface. A carrier with sufficient momentum normal to the surface can overcome this barrier and be absorbed away from an electrode, while low-momentum carriers will bounce away from the surface. Finally, a small probability of random absorption (1 in 10^5) is implemented so that if the first two conditions are unmet, the DMC will still eventually terminate with all charges absorbed. The yield of surface events in the DMC is determined by the height of the surface barrier and the initial kinetic energy of the carriers.

In the set of simulations used to match the observed surface event yields, the initial kinetic energy was partitioned equally to holes and electrons, and the surface energy barriers for both species were also equal. No dispersion was modeled in the magnitude of the initial momenta. The initial momenta given to the carriers were $(18.5, 13.875, 9.25, 4.625, 0) \cdot 10^8 m^{-1}$ and $(32, 24, 16, 8, 0) \cdot 10^8 m^{-1}$ to the holes. Due to the differing effective masses of the holes and electrons, these values result in equal

kinetic energies for the two species (1.125, 0.63, 0.28, 0.07, and 0 eV, respectively). Though the condition for absorption at the surface is related to $\vec{k} \cdot \hat{n}$, I will refer to the initial kinetic energy and the energy barrier at the surface, as all momenta were chosen to correspond to equal kinetic energies for the two charge species. The energy barriers at the top and bottom surfaces were set as fractions of the initial carrier energy from the set $E_{barrier} = (0.49, 0.56, 0.64, 0.72, 0.81, 0.90, 1) * E_{initial}$; in the case where the initial kinetic energy is 0, the tested barriers corresponded to 3, 11, 25, 45, 70, and 101 meV.

For each of the input parameter sets described above, the DMC was run with an input file consisting of 2000 events from G4 simulations of a ^{133}Ba gamma source, 2000 events from a ^{252}Cf neutron source (also emits gammas), and 500 events on the top and bottom surfaces from a ^{109}Cd beta and gamma source.

Results

Gammas from the barium events were used to calibrate the ionization signals to the known event energies, and then the TES signals were calibrated to produce a yield of 1 for electron recoils. These calibrations are then applied to the full simulation set to produce yield versus energy plots like that shown in Fig. 4-32. The reduced neutron yield is clearly visible in green below yields of 0.5 (variance in the nuclear recoil ionization yield and TES noise were turned off for these simulations, explaining why the band is quite ‘tight’). The reduced surface event yields are also clearly visible.

The beta yield bands in data and sim were fit to Gaussians in 2 keV energy bins from 10-80 keV in both hole-collecting and electron-collecting mode. A χ^2 goodness-of-fit statistic was used to compare the data and simulation, $\chi^2 = \sum_i \frac{(\bar{y}_{i,data} - \bar{y}_{i,sim})^2}{\bar{y}_{i,sim}}$, where i labels the energy bins, and $\bar{y}_i$ is the mean beta yield in energy bin i . The




Figure 4-32: Yield of bulk ERs (blue) and surface events (black) in data (left) and DMC (right). The DMC simulation also includes bulk nuclear recoils (green).

‘bump’ in mean yield around 18-22 keV is due to the presence of a γ line at 22 keV, which has a penetration depth of $\sim 65 \mu\text{m}$ and produces a population of events that ‘stretch’ from the surface event yield band up to the bulk electron recoil yield band, raising the mean measured yield in this energy region.

The landscape of χ^2 values versus initial carrier kinetic energies and surface barrier energies is shown in Fig. 4-33. To select a best-matching parameter set for each case, the barrier energies do not need to be the same on the two sides, however, obviously the initial kinetic energy given to the charge carriers in DMC cannot depend on the bias state of the detector. Thus, 70 meV is chosen as the best match for the initial carrier kinetic energy. The χ^2 minima around barrier energy fractions of 0.72-0.81 (~ 50 -56 meV) in electron-collecting mode and 0.9-1 (~ 63 -70 meV) in hole-collecting mode are then taken as starting points for further refinement in which the promising barrier energies were scanned with much finer steps.

Figure 4-34 presents the mean and sigma of the fitted surface event yield bands from 10-80 keV in 2 keV energy bins on the hole-collecting (ionization and phonon




Figure 4-33: χ^2 goodness-of-fit between the surface event yield band means vs. energy in data and simulation in electron-collecting mode (left) and hole-collecting mode (right). Each filled rectangle corresponds to a single DMC parameter set of initial carrier kinetic energies (y-axis) and surface energy barriers (x-axis). The barrier energies are defined as fractions of the initial energy; when the initial kinetic energy is 0, the barrier heights are 3, 11, 25, 45, 70, and 101 meV. Darker colors indicate better matches, and the color scale is logarithmic to increase contrast.

rails biased to -2V and 0V, respectively) and electron-collecting sides (+2V/0V) for the best parameter set: initial kinetic energies set to 70 meV and barriers of 66 and 54 meV on the sides biased at -2V/0V and +2V/0V, respectively. Strong agreement with the data is observed in all cases except for the width of the yield band on the hole-collecting side, which is a bit too low at high energies (at energies <20 keV, the noiselessness of the DMC causes the width of DMC fits to be lower than that of data fits on both the electron- and hole-collecting sides). This set of parameters is used as the standard conditions in the subsequent studies in this thesis.

Future improvements to the charge models

The charge modeling could benefit from a number of improvements:

1. The changes in band structure near crystal surfaces or in regions of high field are not modeled.
2. The surface ‘barriers’ implemented to reproduce surface event yields are not directly tied to a physical model of available electron/hole states in the surface region.
3. While energy is conserved in the DMC, the charge carriers and phonons begin simulation with completely random momenta; it would be interesting to determine whether the direction of the momentum transferred by the recoiling particle affects the eventual signals at all.
4. All charge carriers are generated exactly at the interaction point, while in reality the recoiling electron or nucleus does exhibit some range. This is a low-priority item, but worth mentioning.




Figure 4-34: Means and widths of the fitted surface event yield bands in hole-collecting (ionization and phonon rails biased to -2V and 0V, respectively) and electron-collecting mode (+2V/0V) in 2 keV energy bins from 10-80 keV. The results of fits to the data are plotted in blue, and the DMC is shown in red.

Chapter 5

MISS events

The iZIP detectors were developed to eliminate the primary background faced by the CDMSII experiment: electron recoils that occur very near the detector surfaces and thus exhibit reduced ionization yield. However, as next-generation direct dark matter searches plan to scale the overall exposure by orders of magnitude, from hundreds of kg-days to multiple ton-years, other event topologies that could fake a nuclear recoil signal are important to anticipate. With an effective detector simulation tool in place, we can not only anticipate the types of events that could fake a NR signal, we can estimate the rate of these events as well as propose and test possible discrimination techniques. One such possible background event topology has been termed a Multiple Internal Surface Scatter, or MISS event. A MISS event occurs when a single gamma recoils at both the top and bottom surfaces of the detector. In this type of event, the ionization sensors on each surface will only be sensitive to the ionization signal from the nearby recoil. However, the phonon sensors will collect the phonons created in both scatters, and thus the observed event yield will be suppressed. If the two recoils transfer relatively equal momenta, then the detectors could observe approximately


equal ionization on the two sides, and so this event would defeat the surface event veto based on the requirement of equal ionization signals on the two sides of the detector.

5.1 Anatomy of a MISS event


Figs. 5.1(a-c) illustrate the event topologies of bulk ERs, surface ERs, and MISS events. Without some discussion, it is not immediately obvious why MISS events should present a problem as long as all of the ionization signal is collected. Thus, to elucidate why MISS events present a problem, below I present a calculation of the measured ionization and phonon signals from a bulk event, a surface event, and a MISS event. In this calculation, it is assumed that both the top and bottom electrodes fully collect the charge released in the initial scatter. The phonon and ionization signals are assumed to be noiseless and perfectly calibrated such that a normal well-measured event produces a measured yield of 1 and thus is easily distinguished from the low-yield signal population of bulk nuclear recoils.

First, I define the terms used in the calculation below.


- $+V$ and $-V$ are the ionization sensor potentials, assumed to be equal in magnitude and sign.
- Q_1 and Q_2 are the energies measured in the ionization channels on side 1 and side 2, respectively, perfectly calibrated to be equal to E for an ER in which all carriers are collected
- $Q_{sym} = \min(Q_1, Q_2)$ will represent the symmetric portion of the ionization signal



(a) Illustration of a bulk event in the iZIP detector. Oblique electron propagation in germanium is depicted. Opposite-sign charge species propagate in the bulk drift field, and an ionization signal is observed in both the top and bottom electrodes.



(b) Illustration of a surface event in the iZIP detector. The complex field geometry at the surface causes both electrons and holes to be collected at the top surface. The top electrode observes an ionization signal, while the bottom electrode sees no signal.



(c) Illustration of a MISS event in the iZIP detector. Both charge species from each scatter are collected and observed only by the nearest electrode. The observed phonon signal will measure the total energy of both scatters, but each ionization electrode will measure only the energy of the local scatter, resulting in a reduced ionization yield.

- $Q_{asym} = |Q_1 - Q_2|$ represents the asymmetric portion of the ionization signal
- Turning to the phonon measurements, P_T represents the actual phonon energy measured in an event. P_T has two components: P_R and P_{Luke}
- P_R is the exact energy of the recoil, equal to the prompt phonons created at the interaction site and the recombination phonons created when the charge carriers are collected and release the electron-hole gap energy as phonons. P_R is calculated from P_T through $P_R = P_T - P_{Luke}$.
- P_{Luke} is the phonon energy produced as charge pairs drift across the detector and is equal to $N\Delta V$, where ΔV is the potential difference traversed by the

charges. P_{Luke} is computed from the measured ionization signal as $2|V|\frac{Q_{sym}}{E_{eh}} - |V|\frac{Q_{asym}}{E_{eh}}$. The factor of 2 difference between the symmetric and asymmetric charge signals arises because electrons and holes created in the bulk propagate to +V and -V, respectively, while near the surface they propagate to +V and 0 or 0 and -V, respectively, depending on which surface the scatter took place at.

- Finally, the yield is measured as $y = \max(Q_1, Q_2)/P_R$.

The measured signals for a bulk event of energy E , assuming no noise and proper calibration using known lines from a γ source, are thus:

$$Q_{sym} = Q_1 = Q_2 = E$$

$$Q_{asym} = 0$$

$$P_T = E + 2VQ_{sym}$$

$$P_R = P_T - 2VQ_{sym} = E \tag{5.1}$$

$$y = \max Q_1, Q_2 / P_R = 1 \tag{5.2}$$

The ionization sensors on both sides measure an ionization signal consistent with energy E . The Luke phonon contribution is appropriately estimated, and the recoil energy of the event is thus properly measured by subtracting the Luke phonon contribution from the total phonon signal. The yield of the event is equal to 1.

For a surface event, the analysis is again simple:

$$\begin{aligned}
Q_{sym} &= 0 \\
Q_{asym} &= E \\
P_T &= E + VQ_{sym} \\
P_R &= P_T - VQ_{sym} = E \tag{5.3} \\
y &= \max Q_1, Q_2 / P_R = 1 \tag{5.4}
\end{aligned}$$

Again, the appropriate ionization signal is measured, the appropriate Luke phonon energy is subtracted off, and the recoil energy of the event is accurately measured. The yield of the event is once again equal to 1. However, for a MISS event, the analysis reveals that the measured yield can be significantly reduced. Consider an event that deposits a total energy E , half of which is deposited near the top surface and half of which is deposited in the bottom surface.

$$\begin{aligned}
Q_1 &= E/2 \\
Q_2 &= E/2 \\
Q_{sym} &= E/2 \\
Q_{asym} &= 0 \\
P_T &= E + 2VQ_{sym} \\
P_R &= P_T - 2VQ_{sym} \tag{5.5} \\
y &= \max Q_1, Q_2 / P_R = \frac{E/2}{E} = 0.5 \tag{5.6}
\end{aligned}$$

The generalization of the above analysis to a case in which the energies deposited

at the two surfaces are different or some extra energy is deposited in the bulk region of the detector is trivial. Figure 5-2 illustrates the observed yield, in the case of perfect measurements, as a function of the ratio of the energies of these two scatters. The two dashed lines define the region where the event will pass a wide ionization symmetry cut, in which the ionization signals on the two sides are within 25% of each other. The yield reaches a minimum at 0.5 when the energies are equal, which is still above though dangerously close to the nuclear recoil signal region at high energies; these events could fall into the nuclear recoil band if noise adversely affects the energy measurements. Additionally, very near-surface events in the ‘dead layer’ can still exhibit poor ionization collection in the iZIPs, which would further reduce the yield and allow MISS events to droop into the nuclear recoil signal region.

5.2 MISS rate estimation

Though any individual γ , especially at low energies, is unlikely to exhibit this particular scattering topology, a future CDMS-style dark matter experiment can be expected to target an overall exposure on the order of several ton-years. In the remainder of this chapter, I will use a target exposure of 4.5 ton-years and a target radiogenic background rate of 1500 events per kg-year in the 10-100 keV energy range, which was derived from a GEANT4 simulation performed by Mark Pepin at UMN of the SNOLAB experimental setup with radiogenic contaminations distributed throughout the various shielding layers [112]; though as the simulations and screening material contamination estimates evolve, this number has most recently been quoted as 900 events per kg-year [113]. The radiogenic backgrounds of the current experimental setup include a number of high-energy gammas from a variety of sources: ^{40}K , ^{60}Co , ^{232}Th , and ^{238}U . At energies above ≈ 400 keV, the penetra-




Figure 5-2: Measured ionization yield as a function of the ratio of the energies of the two scatters in a MISS event. The yield reaches a minimum at 0.5 when the two scatters are of exactly equal energy. Events between the two dashed vertical lines would pass the ionization symmetry cut and be deemed bulk events. This figure shows that MISS events that pass the symmetry cut will naturally exhibit event yields near the upper edge of the nuclear recoil band at approximately 0.4. If the ionization from one or both of the scatters is mis-measured for any reason, these events could populate the WIMP signal region.

tion depth is on the order of or longer than the scale of the iZIP detectors, and the observed radiogenic gamma background includes a number of gammas in this high energy range. Any of these incident gammas could easily deposit tens of keV at one surface and still have sufficient energy to penetrate the 1" detectors and scatter again at the opposite surface. In this section, we use a simplified GEANT4 [74] geometry to estimate the overall rate of MISS event candidates and the DMC [114, 115, 116] to estimate the rate at which these events could leak into the signal region.

5.2.1 Back of the envelope expectation

Before beginning simulation, it is useful to get a prior estimate of the expected MISS event rate from simple, back-of-the-envelope calculations and reasonable assumptions. This exercise can provide some confidence in the results of the simulation (if they turn out to be in agreement with the estimation) and allow for some sanity checks of the MISS rates found at each step of the analysis if the results and expectations differ strongly.

The majority of the MISS event population can be expected to arise from higher-energy ($\gtrsim 400$ keV) decay lines from the radiogenic background sources, as argued above. Thus, to first order we can restrict ourselves to high-energy gammas, with penetration depths longer than the thickness of the germanium crystals, and assume that these gammas are equally likely to scatter anywhere in the detector volume. Conservatively defining the surface region to be 2 mm deep, we expect about $(2 \text{ mm}/25.4 \text{ mm})^2 = 0.64\%$ of the gammas to scatter in both the top and bottom surface. I then divide this by a factor of two to account for the angular distribution of incoming photons - though events in which the first scatter is near the (radial) center of the detector have a large angular phase space that takes them through the

opposite surface ($\approx 70\%$, given a 1.5" radius and a 1" thickness), events that scatter further out (more likely as the area grows with r^2) have a much more constrained angular phase space. After applying this factor of two, I expect $\approx 3 * 10^{-3}$ MISS events per ER at all energies .

Now I add the constraint that the two scatters must have recoil energies below 100 keV, which is the highest energy scatter considered in the standard CDMS analysis. From the GEANT4 sims, we find that $\approx 11\%$ of all scatters are below 100 keV. We also constrain the sum of the energies of the two gammas to be below 100 keV, so another factor of two gives us an expected fraction of $0.11^2 = 0.0121$ of MISS candidates at energies below 100 keV. Combining this with the rate of MISS candidates above, we expect $0.003 * 0.0121 = 3.6 * 10^{-5}$ MISS events in 10-100 keV per ER at all energies. Because only 11% of all ERs lie in the range of 10-100 keV, we expect $3.3\text{e-}4$ MISS events per ER in the range 10-100 keV, and with the expected event rate of 1500 ERs in 10-100 keV per kg-year, we have ≈ 1 MISS event candidate per 2 kg-years.

5.2.2 Estimation from GEANT4

To estimate the rate of MISS events in iZIP detectors, GEANT4 was used to simulate radiogenic gammas from the sources known to be present in the CDMSII experiment: ^{40}K , ^{60}Co , ^{232}Th , and ^{238}U . The first two elements decay directly to stable daughters, while the latter two represent the beginnings of long decay chains, of which the primary γ -contributing daughters are ^{214}Pb , ^{212}Pb , ^{214}Bi , ^{212}Bi , ^{208}Tl , ^{234}Pa , and ^{228}Ac . .

The first step in the study was to use CDMSII data to extract a realistic input spectrum for the GEANT simulations. Because the actual experimental setup of a

future ton-scale experiment is not currently available, I began with a simple GEANT4 geometry in which a 1"-thick, 3"-diameter germanium cylinder (to model a CDMSII ZIP detector) was placed at the center of a spherical shell 5 cm in radius. Gammas were emitted from random points on the spherical shell within a collimated cone centered on the detector with an opening half-angle of 47 degrees. The opening angle of 47 degrees was chosen to ensure that gammas originating at any point on the source sphere could illuminate any point on the detector. The collimated cone was used as a compromise: allowing the radioactive source to emit gammas in all directions would result in a very low efficiency of detector hits per simulated gamma, while forcing all gammas to shine directly radially inward onto the detector would produce unrealistic correlations between the gamma event position and direction. 5 million events of each of the individual radioactive sources listed above were simulated, along with a non-isotope specific source that emitted gammas with a uniform energy spectrum from 1-500 keV, meant to approximate the flux of radiogenic gammas that scatter elsewhere in the experimental setup before reaching the detectors. Comparing the simulation results with data from CDMSII allowed us to normalize the radiogenic sources relative to one another to create a combined source that effectively models all radiogenic backgrounds. Fig. 5-3 shows the recoil spectrum from data as a solid blue line and the co-added recoil spectra from the various sources as a red dashed lines. The weights assigned to each of the simulated sources are shown in Table 5.1. The resulting combined source was then used to expose the larger, 1"-thick iZIP that will be used in the SuperCDMS experiment.




Figure 5-3: Background gamma spectrum in the CDMSII experiment and co-added simulated spectra. The data is shown in the solid blue line, and the red dashed line represents the simulated spectra, co-added with the weights shown in Table 5.1

Source	Weight (of 1)
^{228}Ac	0.2607
^{212}Bi	0
^{214}Bi	0.2007
^{60}Co	0.0614
^{40}K	0.0457
^{234}Pa	0.0770
^{212}Pb	0.0165
^{214}Pb	0.0524
^{208}Tl	0.1859
<i>Uniform</i>	0.0998

Table 5.1: Weights for each of the simulated sources in the radiogenic background simulations. Weights were derived from a least-squares fit to the data spectrum.

5.2.3 Raw MISS event rate

A 1" thick, 3" diameter germanium crystal was exposed to a total of 7.25 billion radioactive decays from the combined source described above. When running these simulations, only events that scattered once at both surfaces were recorded. Thus, some normalization was required to translate the rate of observed MISS candidates into a corresponding exposure. I simulated 2.5 million events while recording all output and found that the overall electron recoil rate (ER events / number of simulated decays) with this geometry and source was 48.66%, and the rate of ERs in 10-100 keV was $5.6 \pm 0.03\%$. Applying these rates to the 7.25 billion events simulated in the MISS simulation, a total of 3.52×10^9 ERs and $(406 \pm 2) \times 10^6$ ERs in 10-100 keV were simulated, for a total exposure of 266 ton-years. I define a MISS candidate

event to be any event containing scatters within 2 mm of both the top and bottom surfaces, and a total of 9766836 MISS candidates, 265424 of which were in 10-100 keV, were observed in this simulation, corresponding to a MISS candidate rate per ER of 2.8×10^{-3} , and 6.6×10^{-4} in 10-100 keV.

We can also estimate the rate of MISS events involving two unrelated but coincident radiogenic gammas interacting in opposite surfaces of a detector within a sufficiently short timeframe to be indistinguishable. Let λ be the event rate in a single detector within the energy range of interest, and let w be the time window within which two separate electron recoils are indistinguishable. Then, the time between two ER events, δt in that detector is exponentially distributed,

$$P(\delta t; \lambda) = \lambda e^{-\lambda \delta t} \quad (5.7)$$

and given that a single electron recoil has occurred, the probability of a second one occurring within the time window w is

$$P(\delta t < w; \lambda) = \int_0^w P(\delta t; \lambda) d\delta t = 1 - e^{-\lambda w} \quad (5.8)$$

The mass of each detector is 600 g, so given an event rate $\lambda = 1500$ events/kg-year (as discussed in Sec 5.2.4), we have $1500 \text{ events/kg} - \text{year} * .6 \text{ kg/detector} * 1 \text{ year} / 365 \text{ days} = 2.5$ events per detector per day. Taking $w = 25 \mu\text{s}$, the rate of double-gamma events per single-gamma event is 7.2×10^{-10} , drastically below the single-gamma MISS rates estimated from the BOE calculations ($\approx 3 * 10^{-3}$) and the simple simulation ($\approx 6.8 * 10^{-4}$) described above. Also note that we have ignored the geometric constraint that the two gammas must actually interact in opposite surfaces of the detector. This rate may be amplified due to certain isotopes that decay to extremely short-lived daughter isotopes, producing an enhanced rate of

double-gamma events, but the two gammas would still have to be emitted within the small angular phase space that would allow them to both interact in one detector. Therefore, we choose to ignore this class of event in the coming analysis.

5.2.4 Raw Event Rate from SNOLAB simulations

The estimated event rate of 1500 events per kg-year was derived from GEANT4 simulations of the proposed SCDMS SNOLAB setup using estimated contamination levels in the shielding materials. It's thus important to check that the simple GEANT4 simulation I performed is comparable with the results of these SNOLAB simulations. The SCDMS SNOLAB simulation included detectors 100 mm in diameter and 33.3 mm thick, but the spherical source simulations used the 76.2 mm x 25.4 mm detector, as models of these detectors have been tested and validated. Figure 5-4 presents the spectra from the two simulations, and shows that they are in reasonable agreement along much of the spectrum. The SNOLAB simulation predicts more events at high energy, as expected given the larger form factor of the detectors, and fewer in the 10-100 keV range. The disagreement is likely due to my use of a radiogenic source that emits uniformly from 1-500 keV to model gammas that scatter in the shielding, and the assumption of uniformity in this energy range is rather unrealistic as gammas at the low end of this energy range are rather unlikely to make it through the shielding to the detector. The spherical source simulations predict a factor of 3 too many ERs in the 10-100 keV range, and because I extract the MISS rate by determining the number of MISS events per ER and then assuming 1500 ERs per kg-year, this comparison indicates that the final MISS rate per ton-year I determine should include the possibility of a x3 systematic error due to overestimation of the 10-100 keV ER rate.




Figure 5-4: Background gamma spectrum in a GEANT4 simulation of the SNOLAB setup (solid blue) performed by Mark Pepin and the spherical source G4 simulation (red dashed) with 76.2 mm diameter x 25.4 mm thick detector. The SNOLAB sim utilized expected contamination levels in the shielding and support materials in the experiment, while the spherical sim used the spectrum described in Table 5.1

5.3 MISS events in the signal region

With a large sample of approximately 5 million MISS candidate events, we can utilize the DMC to simulate the observed phonon and ionization signals and apply the standard analysis cuts utilized in the definition of the CDMS WIMP search signal region. Essentially, this section measures the rate at which MISS events ‘fake’ nuclear recoil signals in the absence of any cut specifically designed to discriminate MISS events from normal bulk ERs/NRs. The data analysis cuts include:

1. The Q-inner cut: Events near the radial surface of the detector are cut by demanding that no ionization signal (above noise) was measured in the outer ionization channels. Because there are inner and outer ionization channels on both sides of the detector, this cut can be defined with only one of the two sides of the detector, or as an AND (neither outer ionization rail measures a signal) or OR of the two one-sided Q-inner cuts.
2. The Q-symmetry cut: Events near the top or bottom surfaces will register ionization signals only in the nearby ionization channel, while events in the detector bulk are observed in ionization channels on both sides.
3. The NR band yield cut: Nuclear recoils create fewer free charge carriers per unit energy than electron recoils.
4. The Singles cut: WIMP events should only be observed in a single detector, while many gamma and neutron background events can deposit energy into multiple detectors.

5.3.1 Back of the envelope estimate

Continuing with the back of the envelope calculation in Sec. 5.2.1, we can also estimate the MISS passage ratio of the basic fiducial volume cuts. The outer ionization veto covers the outermost 4 mm of the detector, and since both of our scatters occur in the surface region, we can ignore lateral charge carrier diffusion and just use the relative areas of the two channels. The detector has a radius of 38.1 mm, so $\frac{34^2}{38^2} = 80\%$ of the radial area is covered by the inner electrode. Since we have two scatters, and thus two chances to fail the outer ionization veto, we expect a MISS passage fraction of 64%.

The ionization symmetry cut normally acts to distinguish surface events from bulk events, but for MISS events it constrains the two scatters to be of roughly equal energy. I define the symmetry cut to require the ionization energies on the two sides to be within 25% of each other. For a MISS event, I then expect this cut to cut out 75% of the phase space (we’ve already constrained the sum of the two scatters, so this just amounts to requiring the lower energy scatter to be between 75 – 100% of the higher-energy scatter) and thus this cut should be 25% efficient. The total fiducial volume efficiency for MISS events should thus be $\approx 16\%$.

While MISS events exhibit naturally reduced yields, a well-measured MISS event still produces ionization yields above 0.5, which is $> 3\sigma$ above the mean nuclear recoil yield throughout the 10-100 keV recoil energy range. Thus, to enter the nuclear recoil band, one of the two scatters in the MISS event must take place within the detector ‘dead layer’, the region of low ionization yield near the detector surface. The dead layer is approximately 10 μm -thick, so the probability of one of our two scatters (already constrained to be within 2 mm of the surface) to be in the dead layer is $2 * 10 \mu m / 2 \text{ mm} = 1\%$.

The single-detector cut is more difficult to estimate from basic considerations. CDMSII measured a total single scatter rate (single detector triggers/all triggers) of 13.6%. Simulations of the CDMSII experiment with the known radiogenic backgrounds indicate a single-scatter rate of 18.5%. We would expect that MISS events are more likely to trigger a second detector, for two reasons. First, all MISS events must be somewhat directed along the Z direction. Because the CDMS experiment deploys detectors stacked on top of one another in towers, gammas that cause MISS events should be more likely than normal radiogenic background gammas to trigger a second detector within the same tower. Second, the low-yield MISS events may be even more likely to trigger an adjacent detector, because ERs very near the surface of the detector can eject a β particle from the detector surface, which would be registered in the detector immediately above or below the detector in which the MISS event occurred. 10 – 20% seems a relatively safe, conservative estimate given the rates measured in previous iterations of the experiment.

Applying all of these cuts in sequence, along with the expected MISS rate per ER from Sec. 5.2.1, the expected rate of single-detector NR fakes from MISS events is $\approx 3 - 6 \times 10^{-8}$ per electron recoil background event in 10-100 keV.

5.3.2 Estimation using the DMC


The DMC represents an ideal tool for quantifying the cut efficiencies estimated in the previous section. I used two separate methods to do this, which provides some natural consistency checks on the results. The first method was to use the DMC to generate maps of the observed ionization signals in each of the four channels vs. the position of a single recoil in the detector. Each of the multiple scatters making up a MISS event was then assigned a set of ionization signals, which could be summed to

predict the total observed ionization signal for the MISS event. The second method was to run a phonon-free simulation of all MISS events to extract the ionization signal.


Ionization Mapping Method

To generate maps of the ionization signals in each of the four charge channels vs. the event position, an input file to the DMC was generated consisting of a grid of monoenergetic electron recoils. 47139 recoils were simulated. The spacing in the bulk of the detector was $100\text{ }\mu\text{m}$ in the radial direction and 2.5 mm in the Z direction. The grid spacing was reduced at high radius to capture the rapid change in the ionization signals near the Q-inner/Q-outer boundary, and the grid spacing was also decreased near the top and bottom surfaces to capture the reduction in yield that occurs near the surfaces. Figures 5-5(a) and 5-5(b) show the two maps of ionization signal vs. position on the electron side (Side 2). Full charge collection on the inner/outer sensor is indicated in dark red, while an observed signal of zero is indicated in dark blue.

The advantage of this method is speed. The maps need only be generated once and can be used with future GEANT simulations in the same detector type. However, this method presents a couple of disadvantages as well. The first concern is the grid spacing of the maps. The maps must be generated with sufficiently tight grid spacing to capture the variation in the ionization signals in certain regions of the detector: the surface region and the ‘shared charge’ region in the boundary between the Q-inner and Q-outer sensors. Because I am concerned with MISS events, which are constrained to interact near the surface, the latter region is not an overwhelming concern - the detector regions to which the inner and outer sensors are sensitive are relatively well constrained because the charge do not propagate far enough to



(a) Ionization signal observed in inner charge channel, side 2, versus radial and Z position of the event. Dark blue indicates no observed signal, and dark red indicates full collection (normalized to 1).



(b) Ionization signal observed in outer charge channel, side 2, versus radial and Z position of the event. Dark blue indicates no observed signal, and dark red indicates full collection (normalized to 1). The effect of oblique electron propagation is visible in the light blue band indicating partial collection: further from the sensor (higher Z), the channel is sensitive to regions further inward in radius.

undergo significant lateral diffusion. However, the region near the surface is of primary concern, and as we will see in comparison to the full simulation of each MISS candidate event, there is evidence that the grid near the surface may need to be ‘tightened’ if these maps are to be used in future studies. The second disadvantage is that the maps were generated on a two-dimensional grid in R and Z. In reality, the detector is not symmetric in θ , both due to the layout of the ionization sensors and due to anisotropic electron propagation, though comparison with the results of the full simulation indicate that this effect is relatively small.

Full simulation of all events


When the second-order charge propagation algorithm was implemented (Sec. 3.6), it became feasible to perform a large number of charge-only simulations in a relatively short time (< 5 minutes per event as opposed to > 20 minutes per event). To get around the issues with the ionization mapping method presented above, I also applied a window cut on the recoil energy of the events (10-100 keV, as the analysis in Sec. 5.1 indicates that the recoil energy of MISS events is still properly measured) and simulated the charge carriers from all events passing this cut in the DMC. As will be shown below, reasonable agreement is demonstrated between the two methods, but the results of the full simulation are deemed to be more accurate than the results of the ionization mapping technique, due to the disadvantages described above.

Qinner cut


The Qinner cut acts as a veto against events at high radius, and is generally set to remove all events with energy above noise in the outer ionization ring. Because the DMC simulations do not include any noise on the ionization channels, the side-specific Qinner cuts were set to remove all events that deposited at least 2 keV of ionization energy in the outer ring. This number was chosen to be representative of where the Qinner cut is likely to be placed in the analysis; the measured ‘energy’ of noise taken with random triggers in the experiment tends to be Gaussian with σ on the order of 300-500 eV, and a qinner cut placed 4σ above the noise is fairly conservative.

The qinner cuts on Sides 1 and 2 are illustrated in Figures 5-6(a) and 5-6(b). Table 5.2 presents the passage rate of MISS events for the side-specific Q-inner cuts as well as the AND and OR versions of the cuts; the two columns refer to the two

methods of ionization signal simulation. The two different methods predict very similar values and confirm the BOE estimations reported in Sec. 5.3.1.



(a) Illustration of Qinner cut on Side1 (-2V/0V bias). X and Y axes are the ionization signals in the inner and outer sensors, respectively. Blue dots indicate events passing (i.e., Qouter S1 signal consistent with noise).



(b) Illustration of Qinner cut on Side2 (+2V/0V bias). X and Y axes are the ionization signals in the inner and outer sensors, respectively. Blue dots indicate events passing (i.e., Qouter S2 signal consistent with noise).

Q-symmetry cut

The ionization symmetry cut acts to distinguish surface events from bulk events by requiring that the charge signals measured in the top and bottom ionization sensors are consistent with a well-measured bulk event. For the purposes of this analysis, I define the ionization symmetry cut to require that the two signals be within 25% of each other, and add an offset of 2 keV so that the cut does not close completely at low energies. The exact definition of the cut is $-2 \text{ keV} + 0.75 * Q_1 < Q_2 < 2 \text{ keV} + 1.33 * Q_1$. This choice is not based on the ionization symmetry cut used in experimental iZIP runs, which must be varied somewhat from detector to detector to account for each detector's unique noise, bulk charge carrier trapping,




Figure 5-7: Charge symmetry cut used in the MISS analysis. The X and Y axes are the total ionization energy measured on Side 1 and Side 2, respectively. The dense band of events at 45° demonstrates that the MISS candidate sample contains a large fraction of ‘normal’ bulk events. Blue events pass the symmetry cut, and red events fail.

and neutralization characteristics. In the idealized DMC situation, however, this simple cut serves just fine. Figure 5-7 presents the ionization signals on the two sides of the detector for the sample of MISS candidate events, with the blue events passing the Q-symmetry cut and red events failing it; Table 5.2 presents the MISS passage rates from the ionization mapping method and the full DMC simulation.

Cut	MISS pass % (Q maps)	MISS pass % (full sim)
$Q_{in,hole}$	80.01 ± 0.15	79.31 ± 0.16
$Q_{in,elec}$	62.41 ± 0.19	65.24 ± 0.18
$Q_{in,OR}$	92.94 ± 0.1	92.79 ± 0.10
$Q_{in,AND}$	49.48 ± 0.19	51.77 ± 0.19
Q_{sym}	41.54 ± 0.18	43.40 ± 0.18
$Q_{fid,OR}$	39.72 ± 0.18	41.45 ± 0.19
$Q_{fid,hole}$	34.12 ± 0.19	36.18 ± 0.18
$Q_{fid,AND}$	20.99 ± 0.16	23.04 ± 0.16

Table 5.2: Percentage of MISS candidates in the energy range 10-100 keV passing various fiducial volume cuts. The center column presents results from the ionization mapping method, and the right column presents results from full DMC simulations. All quoted errors are statistical. Though the two methods produce results in strong qualitative agreement, the differences are much larger than the statistical errors, indicating clear (and expected) systematic differences between the two methods.

Nuclear Recoil Yield Band Cut

Figure 5-8 presents a histogram of the yield of MISS events that pass the ionization fiducial volume cuts ($Q_{in,AND}$ was used in this figure). As can be seen, the wide MISS candidate depth of 2mm from either surface provides a large fraction of true bulk events with yields around 1. The expected secondary population of well-collected true MISS events around 0.5 in yield is also evident. The tail of yields below the secondary peak arises due to MISS events with incomplete ionization collection in one or both of the surface scatters. The nuclear recoil band cut is illustrated in Fig. 5-9. The nuclear recoil yield band cut employed here was modeled on the NR band

cut developed in test facility studies of the iZIP detector g48. The detector was exposed to neutrons from a ^{252}Cf source, and the observed nuclear recoil yields were fit to Gaussians within energy bins. The means and standard deviations from the Gaussian fits were then fit to the following energy dependent functions:

$$\mu(E_r) = A_1 * E_r^{A_2} \quad (5.9)$$

$$\sigma(E_r) = \sqrt{B_1^2 * E_r^{B_2} + B_3^2 / E_r} \quad (5.10)$$

where E_r is the recoil energy of the event in units of keV, $A_1 = 0.2075$, $A_2 = 0.1060$, $B_1 = 0.0267$, $B_2 = 2.1379$, and $B_3 = 0.3681$. These functional forms were chosen simply because they have been used within the collaboration since before I joined. Cuts placed at 2 and 3 sigma are visible as the solid blue and dashed black lines, respectively, in Fig. 5-9.

Multiple Scatters Cut

Because the simulation I performed using a spherical source around a single detector is incapable of modeling gammas that scatter in multiple detectors, I use the simulations performed by Mark Pepin to determine the efficiency of the multiple scatter cut for MISS events. We can expect MISS events to have a greater than average probability of scattering in multiple detectors, as the incidence angles of gammas that create a MISS event are constrained by the topology of the event, and these gammas are more likely than others to scatter again in the detector above or below the detector in which the MISS event occurred. Thus, the upper limit on the multiples ratio (multiple detector triggers / all triggers) of MISS events can be assumed to be the multiples ratio of all low-background gamma events in the SCDMS setup currently deployed in Soudan. A multiple scatter is defined as an event that creates




Figure 5-8: Histogram of MISS event yields from events passing the ionization fiducial volume cuts ($Q_{in,AND}$ & Q_{sym}). The MISS candidate sample contains large fraction of true bulk events with yields around 1. The expected secondary population of well-collected true MISS events around 0.5 in yield is also evident. The tail of yields below the secondary peak arises due to MISS events with incomplete ionization collection in one or both of the surface scatters.




Figure 5-9: Nuclear recoil band cut applied to MISS candidates passing ionization fiducial volume cuts; the X axis is the event’s recoil energy, and the Y axis is the measured yield. Cuts placed at 2 and 3 σ are visible as the solid blue and dashed black lines, respectively. Red events fail both cuts, black events pass the 3 σ cut, and blue events pass both the 2 and 3 σ NR band cut.

a MISS event in 10-100 keV in one detector while depositing at least 2 keV of energy into any other detector in the array. With this definition, the best estimate for the multiple scattering rate of MISS events in the SNOLAB detectors is $7.54^{+1.40}_{-1.24}\%$

5.3.3 Expected MISS rate in the signal region

With all of these cuts in hand, the expected MISS rate in the signal region can be determined.

MISS per ER (10-100 keV)	$(6.54 \pm .006) * 10^{-4}$		
Cut	$Q_{in,OR}$	$Q_{in,Hole}$	$Q_{in,AND}$
Q Fiducial Volume	$41.4 \pm 0.2\%$	$35.2 \pm 0.2\%$	$23.0 \pm 0.1\%$
Yield: 2σ NR band	$.133 \pm .024\%$	$.126 \pm .022\%$	$.144 \pm .029\%$
Single Detector Trigger	$7.54^{+1.4}_{-1.2}\%$		
MISS leakage per 100K candidates	$4.15^{+.64}_{-.60}$	$3.34^{+.62}_{-.58}$	$2.50^{+.68}_{-.65}$
Rate per ton-yr (1500 ERs/kg-year)	$0.04^{+.0062}_{-.0058}$	$0.033^{+.0061}_{-.0057}$	$0.024^{+.0065}_{-.0062}$

Table 5.3: Summary of MISS leakage through a set of charge-based fiducial volume, nuclear recoil yield band, and multiple detector scatter cuts.

MISS rate sensitivity to cut definitions

The exact fiducial volume cuts that will be used in iZIP analyses have not been finalized yet, so it is important to check the sensitivity of the MISS rate to the strictness of these cuts. To describe the strictness of the fiducial volume cuts, I parameterize the Q_{inner} and $Q_{symmetry}$ cuts in terms of angles θ_{in} in the $Q_{in} - Q_{out}$ plane and θ_{sym} the $Q_1 - Q_2$ plane, respectively, and assume that the noise thresholds


are set at ± 2 keV. The Qinner cut can then be defined as

$$-2keV < Q_{outer} < 2 + \tan(\theta_{in}) * (Q_{inner}) \quad (5.11)$$


and the Qsymmetry cut is defined as

$$-2keV + (1 - \tan(\theta_{sym})/2) * Q_1 < Q_2 < 2keV + \frac{1}{(1 - \tan(\theta_{sym})/2)} * Q_1 \quad (5.12)$$

This functional form provides a cone in the $Q_1 - Q_2$ plane with opening angle $\theta_{sym}/2$. The Qinner cut shown in 5-6(a) corresponds to $\theta_{in} = 0$, and no Qinner cut at all is applied at $\theta_{in} = \pi/2$. The Qsymmetry cut shown in 5-7 corresponds to $\theta_{sym} \approx \pi/5$. $\theta_{sym} = 0$ forces the charge signals on the two sides to be equal to within ± 2 keV, and $\theta_{sym} = \pi/4$ corresponds to applying no ionization symmetry cut at all. Figures 5-10(a) and 5-10(b) present the final expected rate of MISS events in the 2σ nuclear recoil yield band as colors in the $\theta_{sym} - \theta_{in}$ plane when using the AND and OR versions of the Qinner cut, respectively. The point chosen for the nominal MISS rate determination is marked with a red X. When the Qinner AND cut is employed, the MISS NR fake rate does not vary drastically over a wide range of angles. The point chosen for this analysis is shown as a red X; the rate found at this point was .024 events per ton-year, but it's clear that a wide region of cut angles are available without changing this rate by more than a factor of 2. The most rapid variation is seen at the far right of the plot, where the Qin angle has increased to 90 degrees and thus no Qinner cut is applied at all. In contrast, when the Qinner OR cut is used (shown in 5-10(b)), the Qsymmetry cut acts as the only cut that constrains both charge sensors, and the MISS rate becomes very sensitive to the definition of Qsymmetry.



(a) MISS rate per ton-year in the 2σ nuclear recoil yield band (color) vs. θ_{sym} (Y-axis) and θ_{in} (X-axis), described in Eqs. 5.11 and 5.12. The red X indicates the point chosen for the analysis, which lies in a region where the MISS rate is relatively stable to widening of the cuts. The rate increases drastically in the furthest right column as events contained entirely in Q_{outer} are allowed at $\theta_{in} = 90^\circ$



(b) Same as 5-10(a), using Q_{inner} OR rather than Q_{inner} AND. When the OR cut is used, the MISS rate becomes very sensitive to the symmetry cut, as the $Q_{symmetry}$ cut acts as the only cut constraining both charge sensors.

5.4 Phonon-based MISS discrimination

The preceding sections determine the rate of MISS events passing a set of standard analysis cuts: the recoil energy, fiducial volume, yield, and single detector trigger cuts. The final result indicates that MISS events are unlikely to pose a problem for iZIP-based dark matter searches until these experiments achieve exposures of multiple ton-years. Furthermore, the γ -emitting isotopes also act as neutron generators; though the DMC has not been applied to the expected radiogenic neutron background, the results of current GEANT4 simulations indicate an expectation of approximately 1 single-scatter neutron hit in 10-100 keV per 10 million ERs in the same energy range (without fiducial volume cuts applied). The analysis in the pre-

ceding sections found a MISS event rate of nearer 1 per 100 million ERs in this energy range (after fiducial volume cuts). Thus, the rate of nuclear recoils from real neutrons is currently expected to be higher than the fake NR rate from MISS events by up to 1 order of magnitude, depending on the fiducial volume cut efficiency on the radiogenic neutron background. This observation may indicate a natural stopping point to this analysis.

However, the DMC does represent an idealized realization of the detectors, and unmodeled phenomena could conspire to raise the observed MISS event leakage rate. The most obvious is simply noise, as the DMC does not currently model noise in the ionization channels; implementing the capacitive and electronic noise of the ionization readout system is a priority issue for future DMC developers. On the other hand, the MISS rate increases with increasing energy, and noise is a much larger problem at low energies, so the effect of noise on the overall MISS rate may be relatively small. Effects due to accumulated space charge within the detector can also strongly affect the ionization readout, and these effects were also unmodeled in this analysis. Finally, while the charge carrier-surface interaction model does reproduce surface event yields very well for low-energy betas, a future validation of this model against data with a higher-energy beta source will be necessary to ensure that the depth of the low-yield surface region is accurately reproduced. For these reasons, it is useful to continue the analysis to determine whether MISS event backgrounds can be discriminated from true NR signals through other means. It is also interesting to consider the performance of these cuts vs. energy as noise becomes an issue near the 10 keV analysis threshold. The DMC simulations of the MISS events, bulk NRs, and bulk/surface ERs utilized in this section thus had TES thermal and current noise enabled to capture the degradation in the phonon pulse resolution at low energies. These noise sources were disabled throughout the validation studies presented in 4

to simplify the comparisons, but the DMC’s ability to describe TES noise has been presented in [91].

Applying the loosest set of cuts ($Q_{in,OR}$ and 3σ nuclear recoil bands) leaves a total of 319 MISS events available for analysis. An obvious direction to improve this study would be to perform a much larger GEANT4 simulation to increase this number to at least a few thousand. However, these 319 events should be sufficient to provide some preliminary ideas for discrimination cuts, which is quite sufficient at this stage; fine-tuning this sort of rare event rejection cut depends heavily on the real background conditions and realized exposure (e.g., whether optimal limits/discovery potential is achieved by maximizing total signal exposure or by targeting a maximum allowed background leakage).

Figures in Sec. 4.4 demonstrated that the 10-70% risetimes of the side-summed phonon pulses is a strong 2-D discriminator of ERs and NRs. Fig. 5-11 presents the 10-70% risetimes of the side-summed phonon pulses for bulk ERs (blue), bulk NRs (black), and the MISS background events, with TES noise included. The NRs are very well-separated from the other two populations, with the majority of the MISS events appearing around the most rapidly rising bulk ERs. Strong discrimination should be achievable simply by cutting on the sum of the two risetimes, but a minority of the MISS events exhibit rapid risetimes on one side and very slow risetimes on the other and will thus survive this cut. Fig. 5-12 shows the risetimes in a more natural space for cut-setting, with the sum of the two side-summed risetimes on the x-axis and the absolute difference between the two on the y-axis. It is clear that an appropriate cut on the sum of the two risetimes will prove effective, and this cut is set by eye at $32.5 \mu s$, removing $88.1^{+3.4}_{-4.0}\%$ of the MISS events while retaining $98.7^{+0.6}_{-0.7}\%$ of the bulk NRs (quoted errors are binomial statistical errors). The remaining $\sim 10\%$ of the MISS events appear in a region of this parameter space that is also occupied




Figure 5-11: 10-70% risetimes of the side-summed phonon pulses for bulk ERs (blue), bulk NRs (green), and MISS background events.

by both NRs and single-sided surface ERs. The expected MISS rate is sufficiently low that sacrificing significant nuclear recoil acceptance to remove these MISS events is unlikely to be necessary.

While a sample of MISS events from real data is unavailable, the DMC predicts that $\sim 90\%$ of the MISS events appear in the same region of timing parameter space as the bulk ERs, and thus the performance of this timing cut can be tested in data by comparing bulk ERs and bulk NRs. Figure 5-13 presents the sum of the two side-specific 10-70% risetime parameters versus energy from iZIP4 data. The energy dependence of the timing parameter is quite clear, but luckily both the ER and NR timing parameters exhibit a similar uptick in timing parameter at low energies. The figure also presents a cut, shown in red. This cut was set by binning in energy and




Figure 5-12: Sum (x-axis) and absolute difference (y-axis) of the 10-70% risetimes of the 2 side-summed phonon pulses for bulk ERs (blue), bulk NRs (green), and MISS background events. A cut on the sum of the two risetimes is placed at $32.5 \mu s$ and indicated by the red dashed line. This cut removes $90^{+3.8}_{-3.1}\%$ of the MISS events while retaining $98.5^{+0.7}_{-1.0}\%$ of the bulk NRs (quoted errors are binomial statistical errors).




Figure 5-13: Sum of the 10-70% risetimes (x-axis) vs event energy (y-axis) of the 2 side-summed phonon pulses for bulk ERs (blue) and bulk NRs (green) in g48 R400 data.

then fitting histograms of the timing parameter of the bulk ERs to Gaussians. I then took the means and sigmas of the fits, and fit $\mu + 2\sigma$ versus energy to the functional form $RT_{sum} = \frac{A}{E-E_0} + B$, where E is the event energy and A, B, E_0 are fit parameters. With this cut, $99.2 \pm 0.2\%$ of the bulk ERs are removed, and $96.2^{+1.0}_{-0.8}\%$ of the bulk NRs are retained. Because the DMC predicts that $\sim 90\%$ of the MISS events share the some region in timing parameter space as the ERs, this cut can be assumed to remove this population of the MISS events.

While the side-specific timing does provide good discrimination of MISS events from nuclear recoils, the MISS events appear basically indistinguishable from bulk ERs or single-sided surface ERs in this parameter space. A way to separate MISS events from both bulk ERs and NRs could prove useful in constraining the overall

rate of MISS events in both γ calibration data and WIMP search data. With phonon sensors available on both the top and bottom surfaces, we can construct the position-dependent energy partitioning on both sides. The x/y energy partitioning on the top and bottom side of the detector ought to produce similar position estimates for a single-scatter bulk ER or NR. However, in a MISS event, the two scatters in the top and bottom surfaces of the detector can be totally uncorrelated in x/y position. The two-sided energy partitioning plots shown in 4.4 shows that a single-sided surface event produces very strong partitioning on the near side and very weak partitioning on the opposite, and thus a large fraction of the MISS events may exhibit significant differences in Side1-Side2 partitioning position. Figure 5-14 presents histograms of the distance between event x-y partition positions on side 1 and side 2 of the detector ($\Delta_r = \sqrt{(x_{part,1} - x_{part,2})^2 + (y_{part,1} - y_{part,2})^2}$) for bulk ERs, bulk NRs, and the background MISS events. The histogram for each event type is independently normalized to 1. At first glance, this parameter looks fairly unpromising. The MISS events are certainly distributed differently than the bulk NRs and ERs, but the NR distribution has a significant tail throughout the region where the MISS events lie. However, Fig. 5-15(a) demonstrates that much of this tail turns out to be due to the lowest energy NRs. In terms of a background rejection cut, this doesn't help, as WIMPS are expected to exhibit an exponentially falling recoil energy spectrum, while the MISS events exhibit a rising spectrum with energy (see Figure), so a cut that only works at high energies helps remove MISS events but does very little for the spectrum-averaged WIMP exposure. On the other hand, this parameter could prove useful in MISS identification. When sufficient low background or γ calibration data has been collected, the analyzers can look for MISS events in the main MISS yield band around 0.5 for an event population with a rising energy spectrum and high values Δr_{part} . Figure 5-15(a) also shows a proposed MISS identification cut, defined

using the same method as the timing vs. energy cut. This cut removes $66 \pm 5\%$ of the MISS events while retaining $96.0^{+1.2}_{-1.4}\%$ of the neutrons, so it is not a particularly good background discrimination cut. Furthermore, all events failing this cut already fail the timing cut proposed above, so it seems that this parameter's only role may be for MISS identification. Figure 5-15(b) shows Δr_{part} vs. energy for bulk ERs and bulk NRs from the data. There are some bulk ERs and NRs that have high Δr_{part} at high energy, but these can be separated from MISS candidate events based on their ionization yields, so this parameter may still prove useful in identifying MISS events. For this analysis, I have restricted the phonon simulations to the NR-band MISS events, since these are the events to be most worried about, but simulating the events in the 'mid-yield' band and determining ways to estimate the MISS event rate from mid-yield, high-energy, high- Δr_{part} events may be a promising project.

5.5 Conclusion

Pending experimental measurement of the real γ rate at the experimental site, it is safe to say that the fake nuclear recoil rate from MISS events should not provide a major source of backgrounds until germanium iZIP-style detectors are used in experiments reaching multi-ton-year scale exposures. For reference, the proposed SCDMS SNOLAB experiment currently aims to accumulate a total exposure on the order of 500 kg-years. This study relieves the very real concern that MISS event backgrounds could limit the achievable sensitivity of iZIP-style detectors in the near future. The final expected MISS event rate without any MISS discrimination cuts is ≤ 0.1 fake nuclear recoils per ton-year of exposure, a rate that cannot be totally discounted but may be dominated by the real nuclear recoil rate from radiogenic neutron backgrounds. Furthermore, DMC simulations predict that phonon signal






Figure 5-14: Distance between event x-y partition positions on side 1 and side 2 of the detector.



(a) Δr_{part} vs. recoil energy for bulk ERs, bulk NRs, and MISS events, from iZIP4 DMC. A proposed MISS identification cut is shown in red.



(b) Δr_{part} vs. recoil energy for bulk ERs and bulk NRs from g48 R400. A cut similar to that shown in 5-15(a) may be useful for identifying MISS events at higher energies, where the bulk ERs and NRs are fairly well constrained, and the events at high Δr_{part} can be identified as ERs or NRs and separated from MISS events based on the ionization yield.

parameters ought to be able to reduce this rate by $\sim$ one order of magnitude without sacrificing true nuclear recoil acceptance, and also indicate the possibility of directly estimating the MISS event rate from high-energy, high- Δr_{part} events between the electron recoil and nuclear recoil yield bands. Further studies ought to include:

1. Further exploring the idea of identifying mid-yield, high-energy, high- Δr_{part} MISS events once the true fiducial volume cuts are known and can be applied to DMC results, and quantifying scaling factors between this class of MISS events and the fake NR MISS background.
2. Data analyzers could combine single-sided surface ERs in the data to create a proxy MISS event sample and study promising parameters for MISS identification and rejection.

3. GEANT4 simulations of ^{133}Ba gamma calibration runs in combination with the DMC could be used to determine how much barium calibration would be required to identify a sufficiently large population of MISS events to study their behavior in the actual detectors. I suspect that this may be quite a large amount of data as the penetration depth of 356 keV gammas from the barium source is $\sim 1''$, reducing the probability of 2 isolated surface scatters when compared to the MeV-scale radiogenic background gammas. The physical location of the barium calibration source relative to the detectors also induces geometric effects on the expected MISS rate that could drastically reduce or enhance the MISS event rate if displaced radially or vertically, respectively, from the detector volume.
4. Exploring other types of possible multiple-scatter backgrounds. In particular, a population of zero-yield events was observed in CDMSII [68] and is thought to arise from gammas that scatter at the radial wall of the detectors, with the carriers being trapped before they can propagate to the outer ionization sensor. One of these high-radius scatters in conjunction with an ER in the bulk of the detector is an example of another type of possible multiple scatter background.
5. The two MISS rejection/identification cuts proposed in this chapter were unable to remove a subset of $\sim 10\%$ of the MISS events, which occupied a region of the timing parameter space shared by both surface ERs and bulk NRs. Further study of these events could prove useful in further limiting the MISS event background.

Chapter 6

Phonon pulse matching

6.1 Introduction

Another test of the CDMS-DMC in this thesis involves directly matching the sets of TES pulses observed in real data events with sets of pulses from events simulated in the Detector Monte Carlo. Matching TES pulses observed in data to simulated pulses requires not just accurate phonon physics but also a well-tuned description of the TES parameters, which is described in [117, 118]. The primary goal of the DMC is for it to be sufficiently well-validated for use as a tool to support the analysis. There are many potential uses of the DMC, e.g., evaluating cut efficiencies accurately without dealing with multiple scatters, pileup, etc. that can occur in real data, or predicting rates of rare backgrounds as done in Chapter 5. One idea that is investigated in this chapter is the direct matching of phonon pulses from data to DMC-simulated pulses to classify an event as a bulk electron recoil, a surface electron recoil, or a nuclear recoil.

6.2 Data Selection

In this demonstration, I test the DMC’s ability to match pulses from the iZIP4 detectors utilizing the data from the G48 detector taken during Run 400 at UC Berkeley (the same dataset used in 4.3. Data was selected using the following cuts:

- All events are subject to `cGoodEv_g48`, `cS1Neg2V_g48`, `cQin_and_g48`, `cQsym_g48`, and $-10 \text{ keV} < \text{precoiltnf} < 200 \text{ keV}$. `precoiltnf` is the recoil energy derived using a non-stationary phonon optimal filter developed by Matt Pyle and Bruno Serfass [119].
- NRs are selected using `cNR_uncorr_g48`.
- ERs are selected using `cER_uncorr_g48`.
- No events from the ^{109}Cd sources used to generate surface event populations in the data are included in this test; these events are assumed to be removed by `cQsym_g48`.

With these definitions, events excluded by the existing fiducial volume cuts (Q-inner and Q-symmetry) are not selected, so we are testing only electron vs. nuclear recoil discrimination in the fiducial volume of the detector.

For future users who may wish to review this work, the simulation used for pulse matching is identified by series number 51120604_0066. This simulation was performed with the latest version of the DMC, v4.1.4, using the best match parameters derived in Chapter 4 and also used for the MISS event studies in Chapter 5. The input file to the simulation is a set of nuclear recoils and electron recoils randomly selected from GEANT4 simulations of the event distributions from the ^{252}Cf and

^{133}Ba sources used during the run. The simulated events are distributed throughout the detector and have energy spectra consistent with the source spectra.

The final sample utilized in this study comprises 8431 bulk ERs and 1889 bulk NRs from data compared against 1000 ERs and 1000 NRs from the DMC.

6.3 Pulse Matching Method

The pulse matching method that I employed was the simplest that still provided a proof of concept of this technique: data and simulation pulses are compared based on the sum of squared deviations. The baseline TES current is subtracted, and then the data and simulation pulses are first filtered using a fourth-order 50 kHz Butterworth filter to mitigate noise. The pulses are normalized by the integral of all 8 pulses in the event to ensure that the energy of the event does not affect the matching. Future iterations of this technique could easily improve it to extract a best fit for a relative scaling factor between the data and the simulation, and this scaling factor would be an estimator of the energy of the event. In this investigation, I do not make any attempt beyond a simple low-pass filter to account for the known noise spectrum of the data, and so it is unlikely that the pulse matching technique would perform better than optimal filtering as an energy estimator. However, optimal filtering does generally assume a constant signal shape, so implementing a bank of optimal filters that combined the DMC's ability to reproduce the variations in signal shape vs. position and energy with the measured noise spectra could potentially combine the best features of both the DMC and the currently employed optimal filters. Nevertheless, the primary goal of this study is to determine whether events can be classified as electron or nuclear recoils, not to estimate the recoil energy of the event, so I leave this problem as a suggested avenue of future research.

No weighting is applied in the sum of squared errors (as I said, I pursued the simplest method that worked). In practice, early parts of the phonon pulse provide more information in classifying the event than late times do; as shown in Sec. 4.2, the phonon pulses eventually tend to featureless falling exponentials as the athermal phonon system equilibrates throughout the detector bulk. Bins prior to the start of the DMC pulse are excluded from the sum of squared errors to prevent noise fluctuations from biasing the result. The relative start time between the DMC pulse and the data pulse is allowed to float within a $[-20, +20]$ bin window around the measured trigger time. Because the trigger is issued when the phonon signal exceeds a certain absolute threshold, the trigger should generally be issued earlier in high-energy recoils than in low-energy recoils, so allowing the relative start time of the two pulses to float is necessary to get a good match.

For each event, I record the sum of squared errors of the match to all 2000 simulated ERs and NRs. The sum of squared errors between the event and it's best matching simulated electron recoil (nuclear recoil) is termed χ_γ^2 (χ_n^2). The proposed event classifier is simply $\chi_\gamma^2 - \chi_n^2$; negative values of this parameter indicate a better match to a simulated ER, and positive values indicate a better match to a simulated NR.

6.4 Side-summed pulses

Because the risetimes of the side-summed pulses exhibited very good ER-NR discrimination in Figure 5-11, I start this study by matching only the side-summed pulses from data and DMC. Figure 6-1 presents histograms of $\chi_\gamma^2 - \chi_n^2$ for the electron and nuclear recoils from data and demonstrates that some amount of discrimination is possible. A cut placed at $\chi_\gamma^2 - \chi_n^2 = 0$ provides ER rejection of $\sim 83\%$ with an NR




Figure 6-1: Results of matching side-summed phonon pulses from electron and nuclear recoils from data to electron and nuclear recoils from DMC. Electron recoils from data are shown in blue, nuclear recoils are shown in green.

acceptance of $\sim 86\%$.

Of course, much greater electron recoil rejection is required given the relative rates of electromagnetic backgrounds and the potential dark matter signal, but this high rejection is provided by the ionization yield. Pulse-shape discrimination techniques used to remove the remaining electromagnetic background tend to require $\sim 99\%$ electron recoil rejection, and this pulse matching technique provides $\sim 50\%$ nuclear recoil acceptance at 99% ER rejection. However, far more detailed studies that include events modeling the potential backgrounds to the SuperCDMS experiment must be performed before this technique could be implemented in practice. Finally, at the recoil energies studied in this analysis (>20 keV), the discrimination parameter appears to be independent of the recoil energy, as shown in Figure 6-2.




Figure 6-2: Recoil energy dependence of the rejection parameter for electron recoils (left) and nuclear recoils (right)

6.4.1 Example pulse matches

Well-identified events

Figure 6-3 presents examples of pulse matches that correctly identified an electron recoil (top) and a nuclear recoil, zoomed in around the first 700 μs of the pulses. Because all pulses are normalized, the y-axis of these figures is related to the instantaneous power through the TES but the scale of the axis is not meaningful. After ~ 300 -500 μs , the pulses become featureless exponentials, so all pulses tend to be well-matched after this time. The quality of the matches in Figure 6-3 is quite encouraging, indicating that the DMC does a good job of reproducing the overall pulse shape and that the classification is based on high-quality event matches rather than relying on finding the best match among a set of poor matches.




Figure 6-3: Best-matching DMC pulses for a correctly identified electron recoil (top) and nuclear recoil (bottom), zoomed on the first 700 μs of the event.

A mis-identified event

It may also be instructive to look at the pulses from a mis-identified event. Figure 6-4 shows the best matches to a mis-identified electron recoil. The best matching simulated pulse is a nuclear recoil, presented in the top panel, and the bottom panel presents the best matching simulated electron recoil. It turns out that neither simulation is a particularly good match to the event, which is actually a good result; it would be very worrisome to the prospects of this type of analysis if this ER were very well-matched by a simulated NR.

The simulated NR that provides the best match does seem to match the risetime of the pulse quite well, but fails to peak as strongly or as early as the data pulse. On the other hand, the best matching simulated ER rises much more rapidly and peaks more strongly than the data pulse. It would take too much space here to scan through a large set of mis-identified events, but this one alone provides an indication that the sample of simulated pulses does not cover the full range of possible pulse shapes. This observation is mirrored in the bottom row of Figure 4-22, which plots the 10-70% risetime of side-summed phonon pulses from data and DMC simulations; the data covers a wider range of values in this parameter space than the DMC does.




Figure 6-4: Best-matching DMC pulses for a mis-identified electron recoil. Top: the best-matching DMC pulse, which is a simulated nuclear recoil. Bottom: the best matching simulated electron recoil.

6.5 Matching all 8 phonon pulses

Having demonstrated encouraging results by matching side-summed phonon pulses only, I next attempt to discriminate electron and nuclear recoils by matching all 8 pulses in an event. I restrict the dataset to include only 2000 electron recoils, as the computational time required for this matching analysis is fairly high and grows with the number of pulses being matched; application of this DMC-based analysis to the WIMP search dataset would certainly require some optimization of the code. Figure 6-5 presents the discrimination capabilities when matching all pulses, and interestingly the discrimination power is worse than in the side-summed case; a cut at 0 provides ER discrimination of 70% with NR acceptance of 90%, and obtaining 99% ER rejection requires an NR acceptance of only 35%.

This loss in discrimination power could arise from a number of factors. First, performing this pulse matching on the set of all 8 pulse shapes necessarily involves searching a higher-dimensional space than matching the set of 2 side-summed pulse shapes; it is quite likely that more than 1000 simulated ERs and NRs are required to provide good coverage of this space. Second, summing the pulses should increase the signal-to-noise ratio; summing 4 pulses increases the signal by a factor of 4 while only increasing random noise by a factor of 2. This hypothesis can be checked by looking for energy-dependence in the rejection factor, as low-energy pulses have lower signal-to-noise ratios. Figure 6-6 shows that the rejection parameter does not seem to be particularly energy-dependent in the range of energies considered.

6.5.1 Example pulse matches

First, I present a couple of good pulse matches. The best pulse matches are very good matches to all 8 pulses of the event, as shown in Figure 6-7. Again, because all




Figure 6-5: Results of matching all 8 phonon pulses from electron and nuclear recoils from data and DMC. Electron recoils from data are shown in blue, nuclear recoils are shown in green.

pulses are normalized, the the y-axes in these figures are in arbitrary units of power, so the scale is not meaningful. These matches are quite good, which is a great demonstration of the DMC's performance. The ability to reproduce the features of all 8 phonon pulses from a single event is a strong corroboration of the DMC validation work in Chapter 4.

Next, I searched for pulse features that the DMC had more trouble reproducing. Figure 6-8 presents some example DMC matches to more strongly peaked ER pulses. Though the events are correctly classified as electron recoils, the best match does not reproduce the shapes of 'peaky' phonon pulses quite as well as the less strongly peaked pulses in Figure 6-7, The shape of the pulse peak is strongly dependent on the phonon physics, the quasiparticle downconversion model, and the TES models, so this is the most difficult section of the pulse to simulate correctly.

Finally, Figure 6-9 presents the results of pulse matching to two low-energy nu-




Figure 6-6: Recoil energy dependence of the rejection parameter for electron recoils (left) and nuclear recoils (right) when matching all 8 phonon pulses.

clear recoils, which exhibit fairly substantial noise. Due to the reduced ionization, the phonon signals from nuclear recoils do not benefit from ‘Luke amplification’ as much as electron recoils at similar energies. Still, these noisy events were correctly identified as nuclear recoils. The pulse matching routine did not correctly identify the start time of the pulses in the right panel. As noted in Section 6.3, the matching routine only searches for the best match within a $[-20, +20]$ bin window around the trigger time, and the best match for this pulse was found at the early edge of this window. Widening the window would have produced a better fit, but also increases the required computational time.




Figure 6-7: Example pulse matches from a correctly identified ER (left) and NR (right)




Figure 6-8: Example pulse matches from two correctly identified ERs, showing that the DMC does not seem to match strongly peaked pulses as well as less ‘peaky’ pulses.




Figure 6-9: Example pulse matches from two correctly identified low-energy NRs, showing that this method can find good matches and classify events in the presence of significant noise.

6.6 Future work

This work can be followed up on in a number of ways.

1. Simulating and matching more events would allow a more thorough investigation of whether the DMC is capable of reproducing the full array of phonon pulse shapes exhibited by the experiment.
2. Following up on (1), the code should be rewritten and either optimized for speed or parallelized for grid computing. Matching $\mathcal{O}(10^3)$ simulated events to $\mathcal{O}(10^3)$ data events takes a few hours with side-summed pulses and $\sim$ half a day using all 8 pulses as currently written. While sufficient for a proof of principle, this is clearly unsustainable for future development of this method. Parallelization of this task should be simple, as each pulse match is an independent calculation.
3. Rather than working in the time domain, it would be interesting to utilize the DMC to produce optimal filter templates. In this way, the measured noise of the experiment can be properly accounted for in the matching scheme, and recoil energy estimation from the pulse matches would be possible.
4. In his thesis, Stephen Smith [120] develops a method for simultaneously determining energy and 1-D event position using optimal filters with known variations vs. position. Given an starting position estimate, this method determines an energy estimate and a proposed $\delta\vec{x}$ to a better position estimate. The entire space of possible matches to the signal shape does not need to be searched. As the positions of simulated recoils are perfectly known, this method could potentially be adapted as a means of determining the energy and 3-D position of recoils in the CDMS detectors.

Chapter 7

WIMP Search Analysis using the CDMSII silicon detectors

7.1 Introduction

The preceding chapters of this thesis have focused almost exclusively on the physics of the germanium ZIP and iZIP detectors. However, the CDMSII payload included eleven silicon ZIPs in addition to 19 germanium ZIPs. The basic detector design described in 3.2 is the same for the two types of detectors. Recent DMC development aims to incorporate descriptions of phonon and charge physics in silicon, but the primary focus of the DMC, most detector development and WIMP search analyses are the germanium detectors. One advantage of germanium is that its higher density (5.33 g/cm^3 as opposed to 2.33 g/cm^3 for silicon) provides a higher total exposure (measured in kg-days) for the same detector form factor and livetime. Another advantage is that the ‘benchmark’ spin-independent, isospin-independent WIMP-nucleon cross-section (σ_0 in Eq. 2.2) scales like A^2 , so germanium (mean A of 72.6)

achieves significantly greater WIMP sensitivity per exposure than silicon (mean A of 28.1). However, the silicon detectors provide their own set of advantages that motivate this analysis. First, the reduced mass of the WIMP-nucleus system also enters the interaction cross-section, improving the sensitivity of silicon at low WIMP masses. Second, many dark matter models include a different coupling to protons f_p and to neutrons f_n . As dark matter limits and claimed signals from various experiments (namely COGENT, CDMSII, XENON100, and DAMA/LIBRA) have come into increasing tension in recent years, these models have gained increasing attention as a means of reconciling limits and signals claimed by experiments using different target masses [54]. For a particular element, if $f_n/f_p \sim Z/(Z - A)$, that element’s sensitivity to WIMP-nucleon interactions is reduced by several orders of magnitude. Isotopic silicon has a $Z/(Z - A)$ of almost exactly -1, while those of Ge (COGENT, CDMSII Ge), Xe (XENON100), Na and I (DAMA/LIBRA) range from -0.9 to -0.7, meaning that the results of a Si-based WIMP search can provide new meaningful constraints on these isospin-dependent dark matter models. In this chapter, I present the results of a WIMP search analysis using data from 8 of the 11 silicon detectors deployed in the CDMSII experiment.

A couple of notes regarding nomenclature and formatting used in this section: individual CDMS detectors are referred to by their tower and their position within the tower (i.e., T2Z3 is the third detector from the top in Tower 2). The set of silicon detectors considered in this WIMP search analysis is T1Z4, T2Z1, T2Z2, T2Z4, T3Z3, T4Z1, T4Z3, and T5Z3. The names of quantities and cuts generated during the CDMS analysis pipeline will be denoted in **this font**. The runs analyzed herein (a run is a single continuous cooldown of the CDMSII cryogenic system) are Runs 125-128, which are also termed c58 (where ‘c’ stands for combined). Within a run, data-taking periods are divided up into individual ‘data series’. Series are

continuous periods of data-taking in WIMP search or calibration mode, and can be interrupted for refills of the cryogenic liquids in the dilution refrigerator, for automatic ‘flashing’ of the detectors with LEDs to neutralize accumulated space charge, or to insert/remove a calibration source. ^{133}Ba gamma calibration occurs three times weekly, and ^{252}Cf neutron calibration occurs infrequently (to minimize neutron activation of the detectors and surrounding materials).

7.2 Description of the Experimental Setup

The CDMSII experiment was deployed at the Soudan Underground Laboratory in the Soudan Mine in northern Minnesota. The laboratory is located on the 27th level of the mine, 2341 feet underground. The overhead rock provides 2090 meters water equivalent of cosmic ray shielding, and a number of rare event search experiments have been conducted in the Soudan laboratory, including Soudan 2 and COGENT. The CDMSII experiment was commissioned in phases, with an initial run (Run 118, Oct 2003 - Jan 2004) of 6 detectors (1 tower) [121, 122], a subsequent run (Run 119, March 2004-August 2004) of 2 towers [123, 124], and a set of six final runs (Runs 123-128) with the full payload of 5 towers [125, 126]. A single run corresponds to a single continuous cool-down of the dilution refrigerator; individual runs were ended for installation of new detectors or cryogenic maintenance.

7.2.1 Shielding

At the underground deployment site, the CDMSII experiment is located within an RF-shielded clean room (called the RF room). The RF room contains the actual detectors, the shielding, the cryogenics, and much of the electronics. Multiple layers

of different shielding materials are employed to ensure low-background operation of the CDMSII detectors. The cosmic ray muon rate is greatly reduced but still non-zero in the laboratory, and these cosmic ray muons can induce particle showers that may include neutrons. Thus, the outermost shielding layer consists of a set of plastic scintillating panels and photomultiplier tubes used to tag incident muons and veto any events observed by the detector within a $[-185 +20] \mu\text{s}$ window in time around the scintillation event. A set of LEDs and optical fibers is also attached to the panels to test the photomultiplier response during each cryogen fill, when the detector readout is disabled. This veto represents the only active component of the CDMSII shielding. The next shielding layer is 40 cm of polyethylene, which serves to shield low-energy incident neutrons from reaching the detectors at sufficient energies to recoil within the WIMP search energy range. Inside of the polyethylene lies 18 cm of lead to shield incident gammas, and inside of this layer is another 4.5 cm of ancient lead, which was recovered from the ballast of a sunken ship and shields the natural radioactivity of the outer lead shield. Finally, the innermost shielding layer is another 10 cm of polyethylene to moderate neutrons generated in the lead shielding layers. A schematic illustration of the shielding layout is presented in Fig. 7-1.

7.2.2 Cryogenics

Inside of the shielding layers lies the CDMSII “Icebox” where the detectors are deployed. An Oxford Instruments 400S ^3He - ^4He dilution refrigeration unit provides 400 μW of cooling power at 100 mK. The fridge itself is not constructed of low-radioactivity materials and therefore is operated outside of the shielding layers, as shown in Fig. 7-1. Figure 7-2 presents a cross-sectional schematic of the CDMSII cryogenic system. The refrigerator is on the left, the icebox is in the middle, and




Figure 7-1: Cross-sectional schematic of the CDMSII shielding layout. The dilution refrigerator is shown to the right, and the electronics readout hardware is on the left. Figure from [121]

the electronics readout stem (“E-stem”) is on the right. The innermost layer, shown in black, is the base temperature stage, which reaches 10 mK in the absence of a thermal load but is operated at a higher temperature due to loads from the readout electronics; the detectors are deployed within the black can in the icebox. Moving outward, two intermediate (between the base and the LHe bath) temperature stages, the 50 mK and 600 mK stages, are shown in dark and light blue. Outside of these lie a liquid helium bath (4.2 K), a liquid nitrogen bath (77 K), and the room temperature stage (green, yellow, red, respectively). The fridge is coupled to the icebox through a series of concentric copper tubes, and the cans comprising the icebox are made of low-radioactivity copper. Not pictured is a Gifford-McMahon cryocooler, which was added when CDMSII was expanded from 2 towers to 5 towers and provides extra cooling power at the bath stages (1.5 W at 4.2 K, 40 W at 77 K) to reduce liquid cryogen consumption.

Monitoring and control of the CDMSII refrigeration system is performed by a Moore Advanced Process and Automation Control System (APACS). This system records the pressures, temperatures, flow rates, valve states, etc. throughout the cryogenic system, issues alarms when certain parameters exit user-dictated ranges, directs daily automated refills of the LHe and LN baths, and provides remote control of the system and remote access to historical cryogenic data.

7.2.3 Cold hardware

Detectors

I described the CDMSII ZIP detectors in Sec. 3.2, but will reiterate this description here for the reader, as the preceding chapters of this thesis have covered a variety of detector architectures. The CDMSII detectors are cylindrical germanium and silicon





Figure 7-2: Cross-sectional schematic of the CDMSII cryogenic system. The innermost layer in black is the base temperature stage, and the detectors are deployed inside of the black can in the nested rectangular structure in the middle of the figure. Moving outward, two intermediate (between the base and the LHe bath) temperature stages, the 50 mK and 600 mK stages, are shown in dark and light blue. Outside of these are the LHe bath, the LN bath, and the room temperature stage (green, yellow, red, respectively.)

‘pucks’ 1 cm thick and 3” in diameter (corresponding to ~ 240 g of germanium and ~ 105 g of silicon). Concentric aluminum ionization collection lines are patterned on one of the flat surfaces of the detector, with an inner electrode covering $\approx 85\%$ of the detector surface and an outer electrode covering the remainder. On the opposite surface are four phonon sensing channels, shaped like four 90° pie slices, each channel containing 1036 Transition Edge Sensors (TESs) wired in parallel. Phonons are collected from the crystal by superconducting aluminum fins (quasiparticle traps) that overlap with the TESs but cover significantly more detector surface area. A drift field is applied across the detector by maintaining a voltage difference between the ionization sensors on one surface and the phonon sensors on the other (-3V for Ge, -4V for Si detectors). Figures 7-3(a) and 7-3(b) present a picture and a schematic of the CDMSII detectors, respectively.



(a) Picture of a CDMSII ZIP detector. The phonon sensor side is shown in the image.



(b) Schematic of a CDMSII ZIP detector, illustrating the 4 TES channels on the top and the 2 ionization channels on the bottom.

The ZIP detectors are clamped in place inside of hexagonal containers made of low-radioactivity copper. When deployed in the experiment, the lids of each can are removed and the detectors are stacked one on top of the other with a space of

2 mm between each detector to form towers of 6 detectors. The detector assembly is shown in Fig. 7-4, with cylindrical placeholders where the detector would sit. Detector interface boards (DIBs) provide the first connection between the detector and the readout chain. One side of the DIB is wire-bonded to the detectors, and the other is connected through a small cutout (visible in Fig. 7-4) in the copper can to a removable card called a side coax. The DIBs also include infrared LEDs, which are powered at regular intervals to illuminate and neutralize trapped space charge in the detectors. The side coax card provides both electrical and thermal connections between the tower housing and the individual detectors and also contains a number of the passive elements of the ionization amplification system, which reduces Johnson noise contributions compared to housing these elements at higher temperature stages.

Towers

The tower assembly is another hexagonal, low-radioactivity copper structure used to support the detectors, provide a thermal link between the fridge and the detectors, and connect the readout wiring to signal amplification cards. The tower has four stages, which are connected to the 4 K, 600 mK, 50 mK, and 10 mK refrigeration stages (see Fig. 7-5). Detectors are mounted below the tower structure shown in Fig. 7-5, and each detector is connected to the tower wiring through the side coax card. Each detector's readout wiring runs along one of the 6 tower faces to SQUET (SQUID and FET) cards located on the top of the towers. Individual CDMS detectors are referred to by their tower and their position within the tower (i.e., T2Z3 is the third detector from the top in Tower 2).



Figure 7-4: CDMSII detector holders, with placeholders occupying the detector space. The DIB cutout is visible on the empty placeholder standing on its side. Courtesy: Dennis Seitz.




Figure 7-5: Picture of a CDMSII tower, with temperature stages, readout wiring, and side coax card connections labeled. Courtesy: Dennis Seitz.




Figure 7-6: Picture and schematic layout of a SQUET card. Courtesy: Dennis Seitz.

SQUET Cards and striplines

The SQUET cards (shown in Fig. 7-6) are composed of 2 PCBs each: one that houses a pair of FET amplification circuits to read out the 2 ionization channels per detector, and a second that contains the SQUIDs, shunt resistors, and input and feedback coils used to read out and amplify the phonon signal. The FETs are operated at ~ 140 K (self-heating), and are thus placed inside of a copper gusset and suspended on a Kapton membrane to minimize the effects of heat radiation/conduction to the rest of the card, which is connected to the 4 K stage. Signals are taken from the SQUET card to the warm electronics by a set of striplines, which are flat ribbon cables with copper wires sandwiched between two copper ground planes, all enclosed in Kapton. The striplines are heatsunk twice along the path to the warm electronics to ameliorate thermal loading on the 4 K stage. One additional stripline is employed to carry signals from all of the internal thermometry in the fridge.

Figure 7-7 shows a CAD rendition of an assembled CDMSII detector stack.




Figure 7-7: CAD rendering of a fully assembled detector stack: detectors, towers, and SQUETS. At right is a cross-sectional view, with one ZIP removed. Courtesy: Patrick Wikus.

7.2.4 Warm Electronics

Front-end boards

Signals from the tower run through the “E-stem” pictured in 7-2 to the warm electronics box (“E-box”), a vacuum-tight chamber with an array of vacuum-sealed connectors. The signals are distributed from the E-box to a set of front-end boards (FEBs), one FEB per detector. The FEBs house the circuitry to bias the ionization and phonon readout lines and to control the LEDs used to neutralize the detectors. The FEBs also contain further amplification stages for the signals from the detectors. The FEBs are the last piece of the signal chain that is actually contained in the RF room.

Triggering

The signals from the FEBs are next sent to the Receiver-Trigger-Filter boards, which are located in a separate room (the electronics room) one floor up from the experiment within the Soudan Lab. The RTF boards serve two primary purposes. The

first is to prepare signals for digitization by adjusting the baseline signal from the sensors and applying a two-pole, 366-kHz, anti-aliasing Butterworth filter. The second function of the RTF boards is to issue triggers based on the filtered signal. The phonon signals are used to generate three triggers, known as P-whisper, P-low, and P-high. Each of these triggers can be issued when the sum of the four phonon pulses exceeds a pre-defined threshold, and the threshold for each trigger increases from P-whisper to P-high. P-low is the primary trigger utilized in data taking. The charge signal is also used to generate two similarly-defined triggers, Q-low and Q-high. The trigger thresholds are not set to correspond to a particular event energy, but rather to prevent excessive noise triggers; the measured efficiency of the primary trigger P-low tends to be approximately 100% until recoil energies in the 1-3 keV range. The scintillating veto panel signals are also fed to the electronics room, where the PMT signals are also used for triggering. Each PMT signal is compared against a single threshold value for trigger generation. Because the PMT signals are extremely fast ($\mathcal{O}(1\text{-}10\text{ ns})$), the signals are processed through a filter network that stretches the signal prior to digitization.

Data acquisition (DAQ)

The signals from the RTF boards and the veto triggers are fed to a trigger logic board, which can issue a global trigger that causes the DAQ to digitize and record the ZIP and veto panel signals. The conditions for a global trigger are:

1. The RTF boards issue a P-low trigger for any detector.
2. Two or more of the forty veto panels issue a trigger.
3. Random triggers are issued intermittently to measure the baseline noise levels of all signals.

During low-background and ^{252}Cf calibration running, a global trigger causes the signals from all ZIP detectors and all veto panels to be digitized and recorded. During ^{133}Ba calibration running, a selective readout mode is employed wherein only the ZIP that issued the P-low trigger is recorded, and veto traces are ignored. Signals from the detectors are digitized by Struck SIS 3301 ADCs. These digitizers sample at 100 MHz, but an averaging window of 80 samples is applied to reduce the sampling rate to 1.25 MHz to reduce the size of the raw data. Upon receiving a trigger, the ADC records 2048 samples ($1640\mu\text{s}$) in a window of $[-512, 1595]$ samples surrounding the trigger. Veto triggers are digitized at 5 MHz, and 1024 samples are recorded. When an event is written to disk, a trigger history buffer is also written out to record the trigger history surrounding the event of interest.

7.2.5 DAQ, Data Storage, and Data Processing

A small cluster of computers in the electronics room runs custom-written Java and C++ software to control the CDMSII DAQ. The details of the software are presented in [127]. The primary program for operator oversight is a GUI called RunControl. RunControl allows remote monitoring of the experiment from anywhere, but the ability to issue commands through RunControl is limited to the current shift crew in the mine or the surface facility and a small set of super-users who can access RunControl regardless of location. In Runs 125-128, the DAQ deadtime while acquiring and recording traces after receiving a global trigger is ~ 50 ms (20 Hz), an acceptable level given the ~ 0.3 Hz event rate in low-background (WIMP search) running mode. During barium calibration runs, the experiment is put into the selective readout mode described above, and the dead time falls to ~ 15 ms.

The raw event data is stored on local disks in the mine, backed up to tape,

transferred to another set of storage devices located at the Soudan lab surface facility, and transferred to computers at Fermilab. A cluster in the surface building performs an initial processing of the raw data to provide preliminary reduced datasets (event rate, noise spectra, roughly calibrated phonon and charge energy distributions, etc.) that the operators and a set of automated scripts use to verify the performance of the experiment. Primary data processing to produce the datasets used for WIMP search analysis takes place at Fermilab. Details of the data processing are presented in the next section.

7.3 Data processing and calibration

This section describes the data processing pipeline for calibration and WIMP-search data acquired by CDMSII during Runs 125-128, which ran from July 2007 to September 2008 and represented the final WIMP search runs of the CDMSII experiment. I give an overview of the various stages of data processing and analysis and highlight key algorithms related to pulse reconstruction, energy determination, etc. During the processing of this data, the CDMSII data processing software was undergoing a transition from primarily MATLAB-based software to C++-based software linked to the ROOT data analysis framework [128], and different stages of the processing were performed using different packages. Analyzers are free to analyze the data using ROOT or the CDMS Analysis Package (CAP) in MATLAB. The data processing and analysis pipeline is the result of work by many members of the collaboration, and I make an effort to cite CDMSII internal notes both to credit the extensive contributions of other collaborators and to allow future analyzers to readily follow up any section in further detail.

7.3.1 Data processing

First-level processing

The first-level processing converts the raw trace information, veto and trigger histories, etc. into a set of reduced quantities (RQs) for analysis. Many of the RQs at this stage are represented in uncalibrated units until the next stage of processing. In Runs 125-128, the first-tier processing was performed using a C++ package known as BatRoot [129] linked to ROOT. The RQs are stored in ROOT NTuple format but can also be accessed by CAP.

Ionization Pulse Reconstruction The ionization signal amplitude is determined using optimal filtering. The shape of an ionization signal pulse is essentially fixed by the electronics and can be empirically determined for each detector by averaging a large number of well-measured ionization pulses. The noise is characterized separately for each data series utilizing randomly triggered ‘events’ containing no pulses. Reviews of optimal filter theory and the algorithm employed in CDMSII can be found in [130, 90]. The algorithm simultaneously fits templates to the inner and outer ionization electrode pulses to estimate the signal heights (proportional to energy) and account for cross-talk between the electrodes. The signal height estimates are expressed in arbitrary units until energy calibration in the next stage of processing. The optimal filter algorithm also estimates the start time of the charge pulses.

Phonon Pulse Reconstruction The phonon pulses, on the other hand, do not exhibit a single self-similar shape from event-to-event. Rather, the exact shape of the four phonon pulses depend on both the event position and energy. Optimal fil-

tering using a single template therefore incurs significant systematic errors related to the position and energy dependence. Nevertheless, optimal filtering does provide the first-pass energy estimation and allows us to make estimates of the initial event position through the energy partitioning. The template pulses follow the functional form $A(t) = A_0(1 - e^{-t/\tau_1})e^{-t/\tau_2}$. τ_1 and τ_2 represent a rise time and a fall time and are estimated empirically from a set of well-measured phonon pulses. The integral of the raw trace is also saved as a separate energy estimator; the optimal filter estimator performs better than the integral because it considers the measured noise characteristics, but the integral improves at higher signal-to-noise, while the optimal filter degrades at higher energies due to energy dependence systematics. Similar to the charge estimators, the phonon optimal filter and integral energy estimators are stored in arbitrary units until calibration. Pulse shape characteristics (rise times and fall times) are determined using the ‘walking’ algorithm described in Sec. 4.3.2. The relative delay between these characteristic times in the four TES channels also provide useful estimators for the event position. Figure 7-8 shows an example of this algorithm applied to the four phonon pulses from a single event, and highlights the problems faced by the walking algorithm at low energies, where noise fluctuations degrade performance. Unlike the energy estimators, the pulse shape characteristics are expressed in physical units after first-stage processing. Two pulse shape parameters are of particular importance due to their use in discriminating low-yield surface ERs from bulk NRs in the WIMP search analysis: `pde1`, which is the time between the beginning of the ionization pulse and the 20% risetime of the phonon pulse in the highest-energy phonon channel, and `pminrt`, the 10-40% phonon pulse risetime in the highest-energy phonon channel.




Figure 7-8: Illustration of the walking algorithm applied to four phonon pulses from a single event. The first row shows the charge trace, and the derived start time of the charge signal is shown as a solid red line across the lower four rows. These four rows show the four phonon traces; the full trace is on the left and a zoomed image around the onset of the pulse is shown at right. Unfiltered traces are shown in black, and filtered traces are shown in blue. The 20% risetime of each phonon pulse is indicated by a red X. The channel A trace highlights the poor resolution of this algorithm at low energy. Image courtesy of Scott Hertel [131]

7.3.2 Energy Calibration

Ionization calibration

Energy calibration of the CDMSII detectors begins with the ionization signal and is described in detail by Kyle Sundqvist in [132]. Gammas from ^{133}Ba are used to calibrate the ionization signal. Recall from the detector physics discussions in Sections 3.3 and 3.6 that electron recoils from gammas or betas free $n = \frac{E_R}{E_{eh}}$ charge carrier pairs, where E_R is the recoil energy and E_{eh} is the mean electron-hole pair creation energy, equal to 3.96 eV in Si. Because the charges drift to the ionization electrodes, Ramo's theorem indicates that the height of the ionization signal should scale linearly with n and thus with the recoil energy of an electron recoil. The first step involves residual cross-talk subtraction, performed by diagonalizing a 2x2 correlation matrix; the cross-talk is of order 1-10%, depending on the detector. A preliminary calibration is also applied to the ionization signals to transform from the arbitrary optimal filter output to energy. This step is illustrated in Fig. 7-9.




Figure 7-9: Inner and outer charge signals: uncalibrated (left) and cross-talk corrected with preliminary calibration (right). Images courtesy of Kyle Sundqvist [132]

After the cross-talk correction and preliminary calibration, the germanium detectors receive a secondary calibration. The 356 keV line of ^{133}Ba can be identified in the Ge detectors, and the inner charge signal can be calibrated based on this line. Insufficient 356 keV events are contained entirely in the outer ionization ring to calibrate it directly, but events that share signal between the inner and outer electrodes can be used to calibrate the energy scale of the outer electrode. When the calibrated charge signal is plotted against the event position estimators derived from the phonon delay parameters (described in Sec. 7.3.3), some position dependence is evident. The origins of this position dependence are not well-understood, but it is calibrated away using a 5th-order polynomial in the delay parameters. Calibrating the silicon detectors is a somewhat more difficult problem; as can be seen in Fig. 3-3, the Compton cross-section in Si at 356 keV dominates the photoelectric cross-section by orders of magnitude, unlike in Ge. Therefore, the 356 keV line is effectively not present in the Si detectors. However, the Si detectors can be calibrated using gamma events that Compton scatter in the Si detector and are then fully absorbed in an adjacent detector (or vice versa), by constraining the total summed energy to be equal to 356 keV. This process is illustrated in Fig. 7-10. The ionization energies of the inner and outer channels are called q_i and q_o in CAP, and the total energy is q_{sum}




Figure 7-10: The 356 keV line of ^{133}Ba is not evident in Si detectors alone. Therefore, Si detectors are calibrated using events that also scatter in an adjacent detector (either Ge or Si). Image courtesy of Kyle Sundqvist [132]

Phonon Energy Calibration

The phonon energy calibration is a two-step process, and the same two steps are followed for integral and optimal filter quantity energy calibration. I performed the phonon energy calibration for R125-128 [133, 134, 135] using scripts written by other collaborators during previous analyses. The first step is to perform a relative calibration of the four phonon channels, which is accomplished by aligning the histograms of the individual channel signal divided by the sum of the signals in all four channels; the results of this relative calibration are presented in Fig. 7-11(a)


After the relative calibration is performed, the absolute energy scale of the

summed phonon signals is calibrated. Because the TES sensors exhibit strong non-linearities at the 356 keV line, the phonon signal is calibrated to the ionization signal at lower energies. Events within an ionization energy window from 65-100 keV are selected, and the calibration proceeds as follows: We define the ionization yield as $y = \frac{\mathbf{qsum}}{\mathbf{pr}}$, where $\mathbf{pr}$ is the recoil energy. The recoil energy can be estimated from the total phonon signal $\mathbf{pt}$ by subtracting the Luke phonon contribution, which is equal to $\mathbf{qsum} \frac{eV_b}{E_{eh}}$, where e is the magnitude of the electron charge and V_b is the bias voltage. The ionization yield can thus be written as


$$y = \frac{E_Q}{E_R} = \frac{E_Q}{E_P - E_Q \frac{eV_b}{E_{eh}}} \quad (7.1)$$

Because we define the ionization yield of electron recoils to be equal to 1, we calibrate the total phonon signal $\mathbf{pt}$ such that $\mathbf{pt} = \mathbf{qsum} \left(1 + \frac{eV_b}{E_{eh}}\right)$. Fig. 7-11(b) presents the results of this calibration in ionization yield versus recoil energy. Red and blue points in the figure represent calibrations performed using only R125 data and using combined R125-128 data, respectively. The curved shape of the distribution arises naturally from phonon and ionization signal noise due to the definition of the axes ($x = \mathbf{pt} - \mathbf{qsum} \frac{eV_b}{E_{eh}}$ and $y = \frac{\mathbf{qsum}}{x}$).

At the end of these procedures, the ionization calibration is complete, and a preliminary phonon calibration is available, pending correction of the position and energy dependences of the optimal filter response (and the phonon timing parameters) described in 7.3.3. A MATLAB package called PipeCleaner is used to apply the calibrations to the data and make calibrated energy quantities available to be read from disk, allowing analyzers easy access without having to carry around hard-coded calibration constants or calibration filenames. This package has since been deprecated for a C++/ROOT package known as BatCalib.



(a) Calibrated phonon energy partitioning. The title detector 414 indicates detector T2Z4, and the four colors indicate the four channels. From [133].



(b) Ionization yield vs. recoil energy after absolute phonon calibration, Detector T2Z4, Run 125 Ba calibration. Red and blue represent calibrations performed using only R125 data and using combined R125-128 data, respectively. From [133].

7.3.3 Position and Energy Dependent calibration

The final step of data processing/calibration is known among the CDMS collaborators as position correction, but a more apt description is position and energy dependence correction. In this step of calibration, the observed ionization yields and phonon timing parameters are investigated on a 5-dimensional manifold (x and y energy partitioning, x and y phonon delay, recoil energy), and lookup tables to correct local variations (where local is defined on the 5-D manifold) in yield/timing to corresponding global means across detectors and runs. Though I generated the position correction tables for the c58 silicon detectors, I am indebted to Zeeshan Ahmed, who developed a MATLAB package called CorrTools for this purpose, and to Bruno Serfass, who applied the generated position correction tables to the c58 Si data. For a detailed description of the development of the position correction scheme, see [52].

X-Y Position determination


The energy partitioning and relative timing of phonon pulses among the four channels provide information about the X-Y recoil position within the detector (I put off discussing Z position discrimination until Sec. 7.7). When an event happens under a particular channel, that channel is likely to receive a larger fraction of the total energy than the other three channels, and the pulse in that channel is also likely to be ‘faster’ than the pulses in other three channels. We label the four phonon channels A-D, with corresponding energies p_a - p_d , and construct the following position estimators from the energies:

$$x_{\text{part}} = \frac{p_c + p_d - p_a - p_b}{p_a + p_b + p_c + p_d} \quad (7.2)$$


$$y_{\text{part}} = \frac{p_a + p_d - p_b - p_c}{p_a + p_b + p_c + p_d} \quad (7.3)$$

With these definitions, we have defined the quadrants as follows: A = upper left, B = lower left, C = lower right, D = upper right. These estimators take on a square shape when plotted in a scatter plot (shown in Fig. 7-11(c) for barium calibration events), and is known within CDMS as a box plot (the corresponding plots for the iZIP4 and mZIP presented in Chapter 4 are termed triangle plots). The ‘boxiness’ of this plot, as opposed to the cylindrical shape of the detector, is explained in [136] as a consequence of simply having only four ‘pixels’ measuring these integrated energy quantities. A ‘radial partitioning’ parameter can also be constructed: $r_{\text{part}} = \sqrt{x_{\text{part}}^2 + y_{\text{part}}^2}$.

A separate position estimator can be constructed from the relative pulse delays, though its definition is slightly different from the partition estimators. The 20% risetime of the primary pulse, i.e., the pulse with the largest amplitude, is compared to the 20% risetimes in the two adjacent channels. For example, if channel D (upper



(c) Box plot from silicon detector T2Z4. Events passing the Q-inner cut are shown in blue, events with >10 keV of energy in the outer ionization rail are shown in red.



(d) Box plot from silicon detector T2Z4. Events passing the Q-inner cut are shown in blue, events with >10 keV of energy in the outer ionization rail are shown in red.

right) contains the primary pulse, then the delay parameters are defined as:

$$\mathbf{xdel} = \text{PAr}t20 - \text{PD}rt20 \quad (7.4)$$

$$\mathbf{ydel} = \text{PC}rt20 - \text{PD}rt20 \quad (7.5)$$

where ‘rt20’ indicates the 20% risetime in the corresponding channel. When $\mathbf{xdel}$ and $\mathbf{ydel}$ are plotted in a scatter plot, the result is known as a ‘delay plot’. The delay plot forms a more circular shape than the box plot, as seen in Fig. 7-11(d). Similar to the radial partitioning, a radial delay quantity can be formed: $\mathbf{rdel} = \sqrt{\mathbf{xdel}^2 + \mathbf{ydel}^2}$.

It is crucial to note that these position estimators do exhibit distinct degeneracies. In particular, events that deposit energy in the outer electrode and are thus known to have occurred at high radius appear at smaller radius of the box and delay plots than the outermost events that only deposit energy into the inner electrode. Figures 7-11(c) and 7-11(d) illustrate this issue; events that deposit no energy in the outer channel are shown in blue, and events with >10 keV deposited in the outer electrode

are shown in red, but the red points occupy regions that also contain many blue points. This degeneracy is termed ‘foldback’ and was a significant driver of the ZIP layout redesign into the mZIP and iZIP configurations, which all include outer radial phonon channels to break these degeneracies. However, this degeneracy can also be broken fairly well in the CDMSII-style detectors by utilizing both partition and delay information. Figure 7-11 shows a plot of `rdel` vs. `rppart` in Detector T2Z4. This plot is known as a ‘shrimp’ plot within the collaboration, and nearly all outer (though not all) events appear at the ‘head’ of the shrimp in this space.





Figure 7-11: ‘Shrimp’ plot (`rdel` vs. `rppart`) from Detector T2Z4, Channel A. In this space, the foldback degeneracies evident in partition and delay alone are broken, and almost all outer events appear at the head of the shrimp.

Position and Energy Dependence of Phonon Pulse Shapes


A variety of factors contribute to position-dependent variations in the pulse shape [137, 138]. Fig. 7-3(a) shows that QET coverage is compromised at high radius, due to packing the square QETs onto a circular detector face. Additionally, because phonons are generally not absorbed on their first contact with the surface, reflections from the detector sidewalls also contribute to position-dependent variations. Energy dependence arises as TESs begin to saturate, i.e., the input phonon energy becomes sufficiently high to cause TESs to leave the linear regime or heat to above the superconducting transition. TES transition temperatures also vary across the surface, causing position and energy-dependent pulse-shape variations. Figures 7-12(a) and 7-12(b) illustrate these effects. Figure 7-12(a) shows a shrimp plot from Detector T5Z3, similar to that presented in Fig. 7-11 but colored by `pminrt`. This timing parameter clearly varies with event position, but it varies smoothly along the shrimp, which allows the position-dependence to be empirically corrected (Sec. 7.3.3). The energy dependence of this pulse-shape parameter is presented in Fig. 7-12(b); it is not as obvious as the position dependence, but is present. The changes in pulse shape also affect the measured phonon energy and therefore the ionization yield (Fig. 7-12(c) and 7-12(d)), as the fixed pulse shape employed in the optimal filtering routine is not a perfect description of the true signal pulses.

Correction of Position and Energy Dependence

If the variations in phonon pulse shape with event position and energy are not accounted for, attempting to use pulse shape parameters to discriminate electron recoils and nuclear recoils is effectively hopeless, and yield-based discrimination is significantly degraded as well. However, within localized regions of the detector, electron




(a) Shrimp plot colored by measured p_{minrt} from Run 126, Detector T5Z3, Channel A. [139]




(b) Measured p_{minrt} vs. pr , Run 126, Detector T5Z3, with the 4 phonon channels distinguished by marker color.[139]

recoils and nuclear recoils of similar energies exhibit sufficiently distinct pulse shapes and yields to allow these discrimination techniques. Defining discrimination cuts that vary versus event position and energy is an unappealing option, given the likely high complexity of the functional forms that would be required. Therefore, the energy and position dependences of the pulse shape parameters and the ionization yield are calibrated out prior to setting the final WIMP search discrimination cuts.

The correction method proceeds as follows. A large sample of electron recoils from ^{133}Ba calibration data, distributed throughout the entire detector, is defined after basic data quality cuts are set (see Sec. 7.5). For each event in this dataset, a set of nearest neighbors is found on a 5-D manifold (energy, x/y partition, x/y delay). The local averages of the yield and the pulse shape parameters of interest are determined from this set of nearest neighbors and entered into a lookup table. The value of a particular parameter is then corrected by comparing the local average among that event's nearest neighbors with the global average of that parameter (except in the case of yield, where the local averages are corrected to a value of 1). I review the basics of applying the position correction method here, and interested readers can



(c) Shrimp plot colored by measured yield from Run 126, Detector T5Z3, Channel A.[139]



(d) Measured yield vs. event energy, Run 126, Detector T5Z3, with the 4 phonon channels distinguished by marker color.[139]

find a wealth of information regarding the history, development, and tests of the position correction scheme is available in Chapter 4 of the Caltech doctoral thesis of Zeeshan Ahmed [52].

First-pass table generation

A set of “good” gammas is defined from the ^{133}Ba calibration data using basic data quality cuts described in Sec. 7.5. The gammas are split into separate tables according to run, detector, and detector quadrant (determined using x/y partition and delay). Runs 127 and 128 were combined for most detectors (not T3Z3, which was not analyzed in R127) due to poor statistics in R128, the shortest of the four runs. For each event, a set of N nearest neighbors is determined using the following distance parameter:

$$d = \sqrt{\delta x_{\text{ppart}}^2 + \delta y_{\text{ppart}}^2 + \frac{\delta x_{\text{del}}^2}{L_d} + \frac{\delta y_{\text{del}}^2}{L_d} + \frac{\text{prg}}{L_E}} \quad (7.6)$$

where prg represents the gamma recoil energy and L_d and L_E are scaling factors for the delay and energy parameters in the metric (34 and 600, respectively, for the silicon detector tables). The optimization of these two weights and the number of nearest neighbors N is discussed in [52]. N is defined relative to the number of total events in the detector quadrant, and is optimized in a tradeoff between competing effects: if N is too small, statistical errors overwhelm the determination of the local mean, and if N is too large, then there will be excessive position dependence within the nearest neighborhood itself. For the silicon position correction, N is set to 64 events per 100k, with a minimum of 15 events in detectors with low statistics. After the tables are generated, they are used to perform the position correction on the events used to generate the tables; this step allows us to remove events that are poorly corrected from the tables.

Lookup table cleanup


After generation of the first pass tables, four quality cuts are used to ‘clean’ the lookup tables prior to applying the correction to the full calibration and WIMP search dataset [139]. The cuts are defined with fairly high passage efficiency, as they are meant only to remove outliers.

- Manifold χ^2 cut: this cut removes events that are far from their nearest neighbors in partition and delay space. The nearest neighborhood of most ‘good’ events is distributed relatively tightly around the event. However, events that are located off of the primary manifold may be quite distant from their nearest neighbors, meaning that the measured local average is not necessarily a reflection of the true local average. These events are rejected by defining a 4-D χ^2 using the x/y partition and delay parameters, and the cut is placed at the 99th


percentile point of the 4-D χ^2 distribution. An example event rejected by this cut is presented in Fig. 7-13(a)

- Timing χ^2 cut: similarly, events whose timing parameters (**pdel** and **pminrt**) are drastically different from those of the events comprising its nearest neighbor set are cut. The events are rejected by defining a 2-D timing parameter χ^2 and cutting all events above the 99th percentile point of the 2-D χ^2 distribution.
- Neighborhood size cut: the nearest neighborhood of an event should be relatively tightly distributed, and any events with an excessively wide nearest neighborhood (on the 4-D partition and delay manifold) are rejected. An example event rejected by this cut is presented in Fig. 7-13(b)
- First-pass corrected yield cut: the events used to generate the first-pass lookup tables are position corrected, and an electron recoil ionization yield band is fit to the corrected yield quantities (the procedure for fitting the bands is described in Sec. 7.6.1). Any events whose corrected yield lies outside of the 3σ electron recoil band are removed from the tables. This cut is illustrated in Fig. 7-14

After the application of these cuts, the lookup tables are regenerated, as the nearest neighbor selection of other events may have changed with the removal of these events. These second-pass correction tables are then applied to the entire calibration and WIMP search datasets. Example distributions of yield and **pdel** before and after position correction are compared in Fig. 7-15



(a) Example event rejected by the manifold χ^2 cut is shown in blue, with nearest neighborhood shown in red. All events in first-pass table from Run 125, Detector T2Z4, Quadrant A are shown in gray.



(b) Example event rejected by the neighborhood size cut is shown in blue, with nearest neighborhood shown in red. All events in first-pass table from Run 125, Detector T2Z4, Quadrant A are shown in gray.




Figure 7-14: Electron recoil band after first-pass correction in Run 126, Detector T2Z4. Events in red pass the first-pass yield cut, while those in blue fail it.[139]




Figure 7-15: Comparison of uncorrected (left) and position-energy corrected variables (right) in Detector T2Z4. Top: Yield vs. energy of gammas from WIMP search data. Bottom: 20% phonon risetime (pdel) vs. energy of gamma events from WIMP search data

7.4 Blinding

The WIMP search data is ‘blinded’ to prevent analyzers from using events in the presumed signal region from tuning any WIMP search analysis cuts to achieve some desired result [140, 141, 142, 143]. Potential WIMP candidates in the WIMP search data are selected using a fairly widely-defined blinding cut and removed until all data analysis cuts are finalized. With potential WIMP candidates removed from the data, all WIMP search cuts must be placed using calibration data and events outside of the signal region in WIMP search data, thereby mitigating potential biases from the analyzers. The goal of the blinding cut is to aim for ‘over-blindness’, that is, many of the criteria in the blinding cut mask a wider region than the true eventual signal region. The blinding cut removes any events meeting ALL of the following criteria:

- Yield in the NR band: The nuclear recoil yield bands are defined using nuclear recoils from the ^{252}Cf calibration data. This criterion passes events within 3σ of the mean NR yield. The NR band in the final signal region definition is substantially tighter than 3σ .
- Recoil Energy: Events with recoil energies above the analysis threshold (the greater of 2 keV or 6σ above the noise for a particular detector) and less than 130 keV are selected by this criterion. The eventual analysis range is 7-100 keV in most detectors.
- Inner charge cut: This cut selects events with less than 5 keV of energy observed in the outer charge electrode. The eventual fiducial volume cut is substantially tighter than this definition.
- Single-detector scatter cut: Events must not deposit energy above threshold (again, the greater of 2 keV or 6σ above the noise) in more than one detector

to be blinded.

- Veto anti-coincidence: Events must not have a veto trigger issued within a window of $50\ \mu\text{s}$ prior to the global detector trigger to be blinded.

Blinding is performed immediately after the initial energy scale calibration, prior to position correction, and the blinded sample is redefined after position correction. The position correction of the yield and energy could potentially (and in fact does) cause events to move into and out of the nuclear recoil band, so the blinding cut must be updated after position correction to ensure continued blindness of the data.

Most readers will (rightly!) be skeptical of any blinding scheme that blinds at one stage of the analysis and then re-blinds after subsequent calibration and data quality cut application. However, this decision is defensible, as the properties of events in the low background data are not studied in detail prior to position correction. Unblinded events near the blinded region, which may potentially move into the blinded region post-position-correction, are absolutely not separated and studied. Data quality cuts in this analysis can be broadly categorized as event reconstruction quality cuts or detector performance quality cuts. Event reconstruction quality cuts that remove poorly reconstructed events are set using calibration data, and employ quantities that are independent of the ER vs. NR nature of the recoil, e.g., pre-pulse baseline noise, trigger glitches, etc. Data quality cuts on the WIMP search data that are defined prior to position correction and reblinding are generally set to remove periods of high noise, determined using randomly triggered data; obviously degraded detector performance, e.g., large-scale shifts in multiple RQ distributions; or other known issues, e.g., neutron activation after ^{252}Cf exposure or superfluid He leakage into the detector space. Though skepticism of this blinding scheme is justified, the essential purpose of a blind analysis, to prevent analyzer bias in setting a signal region, is

maintained.

7.5 Data Quality

Data quality cuts are used to ensure that the detector operating conditions are stable, detector performance is up to par, and event reconstruction is of sufficient quality. I separate the data quality cuts into detector performance cuts and event-by-event cuts. Most detector performance cuts (though not all) remove entire periods of data for some or all detectors rather than cutting individual events from data series. The causes of degraded detector performance or operating conditions during these time periods are sometimes well-understood (e.g., detectors undergoing tuning at the beginning of a run, superfluid He leakage into the detector space), though not always. The livetime (and thus the total exposure) of the c58 silicon analysis is computed after application of these cuts. Event-by-event cuts, on the other hand, remove individual events from the data due to a variety of pathologies (e.g., increased pre-pulse baseline noise, trigger glitches). The effects of these cuts are quantified using efficiencies rather than being included in the total livetime determination.

7.5.1 Detector Performance Cuts

Good Detectors Cut

This cut simply removes specific detectors from entire runs due to known (though not necessarily understood) problems with readout or neutralization [144]. Of the eleven silicon detectors, three are entirely unused in c58:

1. T3Z1: Channel C of T3Z1 intermittently enters a state in which the noise spectrum is substantially altered and the pulse amplitudes are drastically reduced.

Experimental operators have attempted to retune the TESs and SQUIDs to prevent this behavior, but unsuccessfully, and this detector is thus unused in this analysis.

2. T1Z6: Neither of the LEDs used to neutralize the detector work for T1Z6, and the detector is thus poorly neutralized much of the time, leading to poorly defined ER and NR yield bands. The detector status wiki maintained by CDMSII [144] lists this detector as usable in R125, but not R123-4, R126-8, and I exclude it from this analysis due to its known problems.
3. T2Z6: Neither of the LEDs on T2Z6 work, and this detector is thus poorly neutralized much of the time. The detector status wiki maintained by CDMSII [144] lists this detector as usable in R125 and R127, but not R123-4, R126, or R128, and I exclude it from this analysis due to the known neutralization problems.

Of the remaining eight detectors, two have problems that cause them to be removed from certain runs in this analysis, though not all:

1. T1Z4: LEDs on T1Z4 appear to have stopped working after R125, and the detector is poorly neutralized and not included in R126-128.
2. T3Z3: T3Z3 was very difficult to neutralize in R127 for unknown reasons, and is thus removed from this run.

It should be noted that all detectors are still utilized in the multiple-detector scatters cut. In CAP, the Good Detectors cut is named `cGoodDet_c58`.

Stable Detector Bias cuts

A pair of simple cuts removes periods of time during which either the TES sensors or the ionization sensors on a particular detector were not in a stable, finalized state. One cut focuses on the TES sensors and primarily flags the periods near the beginning of runs, when the TES sensors and SQUID readout chain undergo tuning for optimal response, though a few data series post-tuning are also removed due to improper or non-standard experimental configurations. The second of these cuts places similar cuts on the ionization channel biases, removing all time periods when the biases are not -3 V on Ge detectors or -4 V on Si detectors. These values are not tuned, but are set to closely match the electron-hole pair creation energies in the two materials (3 eV in Ge, 3.96 eV in Si). This matching ensures that the Luke phonon signal is equal to the recoil phonon signal for ERs, i.e., the phonon signal is amplified by a factor of ~ 2 by drifting the charge carriers. This cut thus removes no ‘tuning’ periods but does remove series when the experiment was improperly configured or during a variety of non-standard runs, e.g., high-field CDMSlite runs. In CAP, the stable TES bias cut is named `cStabTuning_c58`, and the stable ionization bias cut is named `cQNormBias_c58`.

Trigger Rate Cut

The normal trigger rate during WIMP search data series is ~ 0.1 - 0.3 Hz. Series in which trigger rates exceed 0.7 Hz are flagged by a cut known as the trigger burst cut [145]. The high trigger rates can arise due to electronics problems or due to increased phonon noise. In Runs 125 and 127, it is believed that superfluid helium leaked into the detector space and caused a large number of phonon-only events (ionization signal consistent with zero). The time period affected by the helium films is cut by a

separate cut, described below in Sec. 7.5.3, but most of the series cut by the trigger burst cut were also taken during this time period. In CAP, the trigger burst cut is called `cTrigBurst_c58` [145].

The KS Test Cut

Kolmogorov-Smirnov (KS) tests were used to verify the stability of detector performance from series to series. The KS test compares two measured cumulative distributions and computes a reference statistic that indicates the likelihood that the two measured CDFs were drawn from the same underlying distribution. For each run, thirty data series were selected at random for comparison against all series in the run. ^{133}Ba calibration series, ^{252}Cf calibration series, and WIMP search data series were treated separately [146, 147, 148]. KS tests were performed on seven distributions: the yield, `pminrt`, `rdel`, `rppart`, 80-40% faltime of the summed phonon pulse, goodness-of-fit parameter for the charge optimal filter, and the partitioning of energy between qinner - qouter. For each data series, the 30 KS statistics generated by comparison to each of the thirty randomly selected series are averaged, leaving seven KS statistics for each series (one for each distribution).

A data series is automatically cut based on three criteria:

- Any of the seven averaged KS statistics is less than 0.001.
- Any of the seven averaged KS statistics is less than 2σ below the corresponding run-averaged KS statistic.
- The detector contains fewer than ten events in the series.

The automatic cut was then adjusted to account for two specific issues. The first is that the goodness-of-fit parameter for the charge optimal filter template varies much

more strongly from series to series than the other six parameters. The second is to manually adjust the 2σ threshold if a large number of outlier series have driven the threshold to be too permissive. The results of the KS test are also plotted in an array for inspection. This plot is called a ‘kilt plot’, and an example is presented in Fig. 7-16 (Det. T2Z1, Run 125). The columns in the array represent the thirty comparison series, and the rows represent all of the series taken during the run. KS statistics indicating a good match are shown in blue, and those indicating a poor match are in red. In the example, it is clear that detector performance degraded severely toward the end of the run (this was later hypothesized to be due to superfluid helium leaking into the detector space). The KS test cuts are called `cBadDet_Ba_c58` and `cBadDet_bg_c58` in CAP.

The Neutralization Cut

Ensuring that the detectors are properly neutralized is crucial in the WIMP search data. When space charge accumulates, charge collection becomes less efficient, and gamma yields can begin to ‘droop’ down toward the nuclear recoil band. In Sec. 7.2.3, I mentioned that infrared LEDs are present on the detector interface boards; at regular intervals, the detectors are grounded, and these LEDs are powered to illuminate the detector and maintain detector neutralization. The interval is every 11 hours during WIMP-search running, and additional neutralization is performed after any ^{133}Ba calibration series over 30 minutes long. The time between detector neutralizations is set to be shorter than the timescale of neutralization loss in the CDMSII detectors.

The neutralization cut is implemented to remove data series in which, for whatever reason, a detector appears to have lost neutralization. The cut is defined on a

detector-by-detector basis, i.e., if only one detector appears to have lost neutralization, then the series is flagged as bad for that detector but still usable for all others. The primary indicator of poor detector neutralization is an increased fraction of low-yield events. Separate cuts are developed for ^{133}Ba data and WIMP search data. The high event rate during ^{133}Ba calibration provides sufficient statistics to measure the low-yield fraction over time during the data series and therefore salvage periods of a data series during which the detectors appear well-neutralized. Thus, ^{133}Ba series are split up into ‘chunks’ of 100,000 events each when defining the neutralization cut. In contrast, WIMP-search data contains too few events to reliably do this, and so each data series must be evaluated in its entirety.

To define the cut, the mean ratio of low-yield events (yield < 0.8) to total events for the run as a whole is computed (separately for ^{133}Ba and WIMP search data). The low-yield fractions of individual series or ‘chunks’ are then compared to the run-averaged low-yield fraction based on three criteria. In the criteria below, ‘series’ can refer to whole series or to 100,000-event chunks.

1. Compute the 95% binomial confidence interval on the individual series low-yield fraction. If the run mean low-yield fraction is below the lower 95% confidence limit on the series low-yield fraction, cut the series.
2. Compute the Poisson probability of observing the number of low-yield events (or more) observed in the series, given the total number of events in the series and the run-averaged low yield fraction. Cut any series for which this probability is less than 10^{-10} .
3. (WIMP search only) Compute the 5σ confidence interval on the run-averaged low-yield fraction. Cut any series with a low-yield fraction above this interval.

Many series fail more than one of the criteria, but any series failing any of the criteria are cut from the final data set. Fig. 7-17 shows an example plot of neutralization vs. time in Detector T2Z4 in Run 126. Black series pass the neutralization cut, while those in any color are cut. The neutralization cuts are stored in CAP under the names `cBadNeut_bg_12[5/6/7/8]_Si [149]`.




Figure 7-16: KS test ‘kilt plots’ for Detector T2Z1 in Run 125. The columns represent the thirty randomly chosen comparison series, and the rows represent the >300 WIMP search series taken during Run 125. Blue indicates a good match, and red indicates a poor match. The onset of poor detector performance is obvious towards the end of the run. Courtesy: Scott Hertel [148]




Figure 7-17: Rate of low-yield events ($y < 0.8$) vs. time in Detector T2Z4, Run 126. Series shown in black pass the neutralization cut, red series fail criterion 1, blue series fail criterion 3, and cyan series failed more than one criterion. Error bars hidden to improve the legibility of the figure.[149]

Helium Films Cut

Toward the end of Runs 125 and 127, a large number of phonon pulses with no corresponding ionization pulses appeared in a number of detectors. The phonon trigger rates (nominally ~ 0.3 Hz) on many detectors simultaneously rose to unrealistically high values, which severely affected the livetime of the detectors that were unaffected by the no-ionization events. The problem persisted until the run was ended, the fridge was warmed above 4.2 K, and the vacuum spaces (including the detector space) were pumped out; upon re-cooling, the pathological detectors behaved normally again. The problematic behavior is believed to be due to helium that leaked into the detector space. At the tens of mK temperatures of the base detector space, this helium would condense into a superfluid state and form films on the detectors, which are believed to have caused the increase in phonon-only events. The rate of events with ionization yields below 0.1 is used as a means of checking for the presence of He films. Because these films should remain in place until the fridge is warmed, once the rate of low-yield events is observed to spike, all subsequent data series are removed for the affected detector. The shift in the low-yield event rate is rather drastic, and thus no specific criteria were used to flag these series; a simple scan of the low-yield rate vs. time by eye is sufficient to identify problematic series. Figure 7-18 shows an illustrative example plot of the low-yield event rate vs. time in Detector T2Z1, Run 127.

In the Si analysis, we conservatively cut all detectors in a tower if signs of He films appear on any of the detectors in the tower, to avoid any possible systematic effects that might arise due to different detectors being cut at different times, on-site operators changing the trigger conditions for affected detectors, etc. Future re-analyzers of this data set may wish to test less conservative cuts to maximize




Figure 7-18: Rate of zero-ionization events ($y < 0.1$) vs. time in Detector T2Z1, Run 411. The onset of Helium films towards the end of the run is clearly visible. The cut is placed at the red dashed line. X's indicate data series already cut by other data quality cuts, while squares indicate data series considered good before application of `cHeFilm.c58`. Error bars hidden to improve the legibility of the figure. [150]

the livetime, but the difference between setting a strict He film cut and a loose He film cut is $\sim 2\%$ of the total livetime of the c58 dataset. In CAP, this cut is named `cHeFilm.c58` [151, 150].

Noise Performance Cut

This cut removes series that exhibit elevated noise levels in the phonon or charge channels. The amplitude and start time resolution of the optimal filtering routine can be derived using optimal filtering theory from the phonon and charge noise spectra measured in the randomly triggered events. A series is cut if any of the amplitude or timing resolutions of any of the phonon or ionization channels is greater than

25% above the mean resolution for the corresponding channel over the course of the run. This cut was initially set prior to some of the other data quality cuts described above; for the Si analysis, this cut was recomputed after implementing other data quality cuts, as the mean resolution over the course of the run may have been strongly affected by data series known to be bad for other reasons. This cut is called `cBadResTight_c58` in CAP [152, 153]. The setting of this cut is illustrated in Fig. 7-19, which presents the amplitude resolution of all 6 channels of detector T1Z4 vs. time in Run 125. All resolutions are normalized to the run mean for that channel, and any series in which one of the channel resolutions increases above 1.25 (red dashed line) is cut.




Figure 7-19: Setting of the `cBadResTight_c58` cut. Points indicate the amplitude resolution of all 6 channels of detector T1Z4 vs. time in Run 125. All resolutions are normalized to the run mean for that channel, and any series in which one of the channel resolutions increases above 1.25 (red dashed line) is cut [153].

A second cut is placed based on another method of estimating the ionization

noise. Events with phonon energy less than 4σ above the noise are selected, and the summed ionization energy (qsum) of these events is computed. The qsum of these events is used as an estimator of the charge noise. qsum is averaged over a 20 event window, and a cut is placed to remove periods where the moving average of the charge noise is more than 6σ above the mean charge noise (mean taken over the series). qsum is also averaged over a 20-series window, and any series whose moving average increases above 4σ above the mean charge noise for the run as a whole. This cut is named `cHighQNoise_c58` in CAP [154].

MINOS beam cut

A GeV-scale muon neutrino beam (NuMI) created at Fermilab is directed towards the MINOS far detector, located in a separate cavern within the Soudan Underground Lab. These neutrinos could cause showers that produce neutrons in the rock or materials surrounding the CDMS experiment. The background from the NuMI beam is expected to be insignificant, but all events within a $60\ \mu\text{s}$ window around beam firings are cut from the final dataset out of caution. The name of this cut in CAP is `cNuMI_c58` [155].

7.5.2 Total Exposure

After all data quality cuts are defined, the total exposure of the c58 Si dataset can be computed by simply multiplying the detector mass and the detector livetime. The collaboration has traditionally quoted the masses of the Si detectors at $\sim 100\ \text{g}$, but more detailed estimates have indicated masses of $\sim 106\ \text{g}$, which I will use. The time that the DAQ waits between events (in ms) is recorded in the CAP RQ `LiveTime`, and the total livetime is obtained by summing this over all events not rejected by

the data quality cuts above.

	T1	T2	T3	T4	T5
Z1		167.83		208.19	
Z2		115.50			
Z3			160.90	217.28	203.42
Z4	62.55	187.21			
Z5					
Z6					

Table 7.1: Total livetime (in days) for each detector in the c58 Si analysis

7.5.3 Event-by-Event Cuts

Triggering

In addition to the trigger burst cut described in 7.5.1, a number of cuts are in place to ensure that all events in the WIMP search analysis have normal, well-recorded trigger information. Any events whose trigger records appear problematic (i.e., missing or invalid values recorded) are removed by a cut called `cErrMask_c58` [156]. Electronics ‘glitches’ can cause a nominally normal-looking global trigger to be issued. These glitches are believed to be the result of coherent noise on the QET bias lines, which can briefly heat the QETs and cause the phonon triggers to fire. The glitches are identifiable by comparing the number of phonon triggers and the number of charge triggers that were issued. If the number of phonon triggers minus the number of charge triggers is greater than 4, the event is cut by `cGlitch_c58`[157].

Charge Optimal Filter χ^2 Cut

The goodness-of-fit χ^2 (CAP RQ: `QSOFchisq`) between the optimal filter template and the ionization pulse is used to evaluate the quality of event reconstruction on an event-by-event basis. This cut is particularly sensitive to ‘pileup’ events, that is, events in which more than one recoil occurs within the trace window recorded after a trigger. These events can occur during high-rate calibration data runs, but are clearly unusable for calibrating the energy scale, yield, or timing of ERs and NRs. Because the ionization pulse is significantly shorter than the total trace length, a second recoil causes significant deviations from the template pulse shape, allowing pileup events to be easily cut using the optimal filter χ^2 . Figure 7-20 shows a scatter plot of the optimal filter χ^2 vs ionization energy. It seems odd that the goodness-of-fit value decreases at high energies, as the signal-to-noise increases with energy. However, the template itself is defined by averaging a set of measured pulses, and the noise in the template degrades the goodness-of-fit at high energies. The cut is computed by first binning the data into four energy bins: [20 40], [45 85], [90 130] and [135 175] keV. The χ^2 distribution is fit to Gaussians within each bin, and a quadratic is then fit to the 3.5σ points in the four energy bins to define the cut [158]. This cut is known in CAP as `cChiSq_c58`. No similar cut is computed for the phonons due to the inherent variation in pulse shape vs. energy and position.




Figure 7-20: Charge optimal filter goodness-of-fit vs. ionization energy in Detector T4Z3, Run 125. The black solid line indicates the cut used in this analysis, while the red dashed line indicates a previous definition of the cut, defined prior to some updates to the charge optimal filtering routine. Courtesy: Jianjie Zhang [158]

Baseline Noise Cuts

`cBadResTight_c58` and `cHighQNoise_c58` are used to reject periods or series of unusually high noise. Analogous cuts can also be generated to cut individual noisy events. When a trigger is issued, the previous 512 samples are recorded from a continuous buffer. The $\sim 400 \mu\text{s}$ of samples that are recorded prior to the trigger can be used to set cuts on the baseline noise and as a second pileup-rejection cut. The standard deviations of the first 512 samples in each channel are stored in the RQs `P*std` and `Q*std`, where ‘*’ can be [A, B, C, D] for phonon channels or [I, O] for the

inner and outer charge channels, respectively. Two cuts are then placed using these parameters, with CAP names `cPstd_c58` and `cQstd_c58` [159, 160].

`cPstd_c58` cuts events based on the phonon baseline standard deviations, and functions as a cut on both noise and pileup. Because the charge trace falltimes are quite short, the charge χ^2 cut is sensitive to a second recoil that occurs after a trigger is issued. In contrast, the phonon falltimes are significantly longer, and the pre-pulse baseline of the phonon traces can sometimes contain persistent signal from a prior event in the detector. To define the cut, P^*std is first computed for all traces in the series. The values for a series are then fit to a Gaussian, and any events with pre-pulse standard deviations more than 5σ from the series mean are rejected. The efficiency of this cut is high, generally greater than 99%. The cut and its effect are illustrated in Fig. 7-21.

`cQstd_c58` is not applicable as a pileup-rejection cut, but does reject events that have significantly elevated ionization noise, primarily due to microphonic pickup of the vibrations induced by the cryocooler. The cryocooler performs a mechanical compression cycle at a frequency of 0.8 Hz, but the compression cycle can induce vibrations throughout the experiment at all frequencies. Most detectors are insensitive to these vibrations, but a subset of detectors intermittently picks up additional noise in the charge channels due to the cryocooler. Of the silicon detectors, T1Z4 and T2Z1 are sensitive to the cryocooler vibration noise. This cut is defined similarly to `cPstd_c58`, but the threshold is set at 4σ rather than 5σ . The cut and its effect on Detector T2Z1 are presented in Fig. 7-22




Figure 7-21: Baseline noise vs. time in the four phonon channels of Detector T2Z1, Run 125. Blue events pass $cPstd_c58$, while red events fail. The cyan lines indicate the limits of the cut. Courtesy: David Moore [160]




Figure 7-22: Baseline noise vs. time in the two charge channels of Detector T2Z1, Run 125. Blue events pass `cQstd.c58`, while red events fail. The cyan lines indicate the limits of the cut. The effect of cryocooler microphonic pickup in the outer channel of this detector is clear. Courtesy: David Moore [160]

Phonon Start Time

This cut ensures that the start time of the pulse, determined using the walking algorithm previously described, is within a reasonable window around the global trigger (at the 513th sample). Events in which the pulses are not actually near the trigger occur when a detector issues a global trigger, and then an unrelated recoil occurs in a different detector within the trace window. These types of events can occur during calibration data, when event rates are high, but are extremely rare in WIMP search data, when event rates are low (0.3 Hz event rate and ~ 1 ms trace window). The ionization reconstruction algorithms search within a window of $[-100, +10]\mu\text{s}$ around the trigger, while the phonon reconstruction algorithms search within a $[-50, +200]\mu\text{s}$ window. These events may be well-reconstructed in phonons, but the ionization reconstruction can be poor. The cut `cGoodPStartTime` [161, 162] ensures that the event lies within the overlap of the two reconstruction windows ($[-50, +10]\mu\text{s}$). This issue could also be handled by widening the search windows, but this is computationally expensive, and so events whose reconstructed start times do not fall within the search window around the trigger are rejected. Figure 7-23 illustrates the definition of the cut and shows an example event that fails the cut. Again, this problem occurs in calibration data, and is not an issue in low-rate WIMP search mode.




Figure 7-23: Example ‘cross-detector pileup’ event failing the `cGoodPStartTime` cut. The charge search window is shown in magenta, and the phonon search window is shown in green. Courtesy: Jeff Filippini [161].

7.6 Physics Cuts

After setting the data quality cuts, the next step is to define a set of physically motivated cuts that delineate the WIMP signal region and separate WIMPs from the background ERs. This section describes nearly all of these cuts – the timing cut used to distinguish low-yield surface ERs from true NRs is described in its own section, as setting this cut requires more detailed considerations of the tradeoff between background leakage and signal efficiency than any of the other cuts in this section.

7.6.1 The ionization yield bands

As discussed in Sec. 3.3, nuclear recoils in the CDMS energy range produce far fewer electron-hole pairs than electron recoils at the same energy, and WIMPs are expected to recoil from nuclei rather than electrons. Thus, the great majority of the ER background can be rejected by a cut on the ionization yield. The nuclear recoil signal region for this analysis is initially defined by the cut `cNR_c58` as a $\pm 2\sigma$ band around the mean NR ionization yield at the measured event energy (the actual signal band is -1.8σ to $+1.2\sigma$, optimized when considering surface event rejection). The nuclear recoil band is defined using ^{252}Cf calibration data. To set the cut, the nuclear recoils are first binned in recoil energy. The NR yield within each bin is then fit to a Gaussian distribution, and the fitted means and sigmas are then fit to energy-dependent polynomials of the forms:

$$\mu_y(\text{pr}) = A * \text{pr}^B \sigma_y(\text{pr}) = C * A * \text{pr}^{B-1} + \frac{D}{\text{pr}} \quad (7.7)$$

where μ_y is the fitted Gaussian mean, σ_y is the fitted Gaussian width, pr is the measured recoil energy, and A-D are the fit parameters [163]. The electron recoil band is determined using exactly the same method, and the fitted bands are presented in Fig. 7-24.

The efficiency of the NR band cut is computed by comparing the number of ^{252}Cf events in a $\pm 2\sigma$ NR band to the number of events in the $\pm 4\sigma$ band. The NR band was later constrained to -1.8σ to $+1.2\sigma$ during the surface event rejection analysis, and the NR band efficiency obviously must be recomputed. The $\pm 4\sigma$ band is used for comparison to prevent excessive gamma contamination in the NR efficiency measurement.




Figure 7-24: The yield bands for electron recoils (red) and nuclear recoils (green) vs. event energy. The mean yields of the ER and NR bands are presented as dashed lines, and the $\pm 2\sigma$ limits are indicated by solid lines.

7.6.2 Charge Fiducial Volume Cut

The cut `cQin.c58R` rejects events with any signal above noise in the outer ionization electrode. Events near the outer edge of the detector may exhibit reduced ionization yield due to fringing electric fields. Similar to the yield band definition, the energy in the outer electrode (q_o) is binned according to the energy in the inner electrode (q_i), the noise band around $q_o = 0$ is fit to Gaussians, and the mean and width of the Gaussian fits are fit to polynomials in q_i . The cut requires q_o to be within $\pm 2\sigma$ around the mean [164]. The cut is expanded when both q_i and q_o approach the noise, and the cut criterion changes from requiring q_o to be within 2σ to the criterion $q_i^2 + q_o^2 < R^2$, where R is fit from the noise for each detector. An illustration of the cut is shown in Fig. 7-25.




Figure 7-25: Plot of q_o vs. q_i showing the effect of the charge fiducial volume cut, Detector T4Z3, Run 125. The limits of the cut are presented as green solid lines. The polynomial fits extend to $q_i \sim 5$ keV, below which the cut is widened

Based purely on the areal coverage of the inner and outer electrodes, one would expect the fiducial volume cut efficiency to be $\sim 85\%$. However, the measured efficiency is somewhat different. At higher energies, the likelihood that the event was composed of multiple scatters within the detector increases, and correspondingly, the efficiency of the cut decreases. Because a WIMP event is expected to be a single scatter within the detectors, measuring the cut efficiency using neutrons results in a systematic underestimation of the true WIMP fiducial volume. The DMC would provide a natural candidate to test the true WIMP fiducial volume, but unfortunately, simulation of silicon detectors has not yet been implemented. Instead, GEANT4 simulations of ^{252}Cf calibration data were performed during the c34 analysis, and geometric considerations were used to assign corresponding ionization energies into the inner and

outer electrodes. This study estimated that the efficiency of `cQin_c58` measured using neutrons in germanium detectors is expected to be $94.75\% \pm 0.89\%(\text{stat})^{+0.0084\%}_{-0.0050\%}$ (syst) of the actual WIMP efficiency [165].

7.6.3 Scintillator Veto Cut

The plastic scintillator panels that surround the experiment were described in Sec. 7.2.1. These panels protect against cosmogenic neutrons by tagging muons or secondary showers that could interact within the experiment to produce neutrons. When a global trigger is issued during WIMP search running, veto traces are read out in a window of $[-185, +25]\mu\text{s}$ around the trigger. In addition, the time of any scintillator panel hit that exceeds a hardware threshold of ~ 2 MeV is recorded into a history buffer. The cut is defined by two criteria: an event is veto-coincident if 1) the history buffer records any veto hit above the hardware threshold within a window of $[-50, 0]\mu\text{s}$ around the global trigger, or 2) the trace from any panel exceeds a software threshold at any time within the full $[-185, +25]\mu\text{s}$ window around the global trigger.

Though the expected muon rate at Soudan is quite low ($\sim$ one per minute), the gamma background causes the veto hit rate to be roughly 400 Hz, as the veto panels are located outside of the passive shielding. To ensure that the livetime of the experiment does not suffer due to the veto anti-coincidence requirement, the software threshold for the veto cut is set at a level (~ 3.8 MeV) that ignores the gamma background while still being sensitive to muons or showers. Figure 7-26 presents the spectrum of trace amplitudes from a single panel, with the hardware threshold indicated by a green dashed line and the software threshold indicated by a red dashed line.

The efficiency of the veto cut is $97.876\% \pm 0.013\%$ [167], measured by determining




Figure 7-26: Spectrum of trace amplitudes from a single veto scintillator panel, with the hardware threshold indicated by a green dashed line and the software threshold indicated by a red dashed line. Courtesy: Matt Fritts, [166]

how many randomly triggered events fail the veto cut.

7.6.4 Single-Detector Scatter

The very low WIMP-nucleon cross-section implies that a WIMP candidate should only interact a single time in the experiment. In contrast, many background particles are expected to interact numerous times, and thus tagging multiple scatters is an effective way to remove backgrounds while maintaining high signal efficiency. The ZIP detectors do not have the necessary temporal or spatial resolution to identify events that scatter more than once in a single detector, but events that recoil in 2 or more ZIPs are identified and removed by the cut `cSingle_c58R`. To determine the energy thresholds for `cSingle_c58R`, all randomly triggered traces taken in a data series are processed using the same optimal filtering algorithms applied to real events.

The derived energy distributions of the random triggers are then fit to Gaussians. An event is classified as a single scatter if a phonon energy of more than 6σ above the noise mean is observed in one detector, and the phonon energy is less than 4σ above the mean noise in all other detectors. Measuring the efficiency of the single-scatter cut is not straightforward, as neither of the calibration sources provide a known population of true single scatters. Instead, the efficiency is measured by determining how many randomly triggered events, which should not have real energy depositions in any detector, fail the singles cut.

A few poorly performing detectors required different treatment, as excessive noise caused the singles cut to lose efficiency. Some of the observed poor performance may have been due to the onset of helium films, but the singles cut was defined prior to the helium film cut, and I chose to continue using the same definition of a “single scatter” as was used in previous analyses. The poorly performing detectors were:

- T1Z1, T5Z5 and T5Z6 in all runs use both the charge signal and the phonon signal for vetoing multiple scatters. The charge energy threshold for multiples tagging is set at 4σ above the mean energy of random triggers, just like the phonon energy.
- T5Z1 in Run 125: High phonon noise. The phonon signal is ignored, and the charge signal is used as a veto instead.
- T2Z2 in Run 125: High phonon noise. The 4σ threshold for multiples tagging was raised to 5σ
- T5Z2 in Run 127: Unused for multiples tagging due to high phonon noise and loss of charge signals.

- T2Z1, T5Z3, T5Z4 in Run 127: High phonon noise at the end of the run (likely due to Helium film). If the phonon noise in a single channel is high, add the charge signal as a multiple scatters veto. If two or more phonon channels exhibit high noise, use only the charge signal as a veto.

7.6.5 Energy thresholds

Separate energy thresholds are defined for the ionization and phonon signals. The cut `cQThresh_c58R` requires that the charge signal in the inner electrode be above the 4.5σ point of the charge noise distribution for the series. The series-dependent noise distributions are used to ensure that series with higher observed noise levels do not either 1) cause the threshold to be set excessively high or 2) leak noise events into the signal region. The overall charge noise distribution for each entire run is also computed, and the charge signal is also required to be above the 4.5σ point of the overall run noise distribution. That is, the noise distribution of the overall run sets an effective minimum threshold, and the thresholds can be adjusted upward for particularly noisy series.

A separate pair of thresholds is set on the recoil energy of the event. Unlike the ionization threshold, the lower recoil energy threshold is not set based on the noise distribution of randomly triggered events. Rather, the recoil energy cut ensures that the phonon pulse shape parameters are well-measured, that the position correction can thus be performed well (as the timing parameters `xdel` and `ydel` are in the metric), that the bulk electron recoil leakage into the nuclear recoil band is negligible, and that there are sufficient statistics of low-yield surface electron recoils to set the surface event rejection cut. For this analysis, the lower recoil energy threshold is set to 7 keV for most detectors, with two exceptions. T4Z1 exhibits a wider electron

recoil band than most detectors, and the threshold is set at 10 keV to prevent gamma leakage (this poor performance is why it was chosen to be at the end of the tower rather than the interior). The readout hardware attached to T4Z3 in Run 125 was much noisier than average, and so the threshold in this run is set at 15 keV; the hardware was replaced between Run 125 and 126, and the threshold is set to 7 keV in these runs. Finally, an upper energy threshold is also set on the recoil energy of the event at 100 keV. This threshold exists because the expected WIMP-nucleon interaction rate decreases exponentially with energy, and energies above 100 keV contribute very little to the analysis. Future analyzers may decrease the lower recoil-energy threshold to perform a low-energy analysis similar to that in [168] or increase the upper recoil energy threshold to investigate inelastic dark matter models as in [169].

7.7 Surface Event Rejection Cut

Surface events in the CDMSII detectors exhibit reduced yield and represent a significant possible background. Fortunately, surface events can be identified using phonon pulse shapes, in particular the shape of the rising edge of the phonon pulse. In this section, I describe the definition and performance of the surface event rejection cut (or timing cut) used in the c58 Si analysis. A note regarding terminology - surface events are also referred to as betas or β s throughout the remainder of this document, but surface events can come from actual electrons or from photons that recoil within the $\mathcal{O}(100\mu\text{m})$ nearest the surface, termed the ‘dead layer’ where low yields are common.

7.7.1 Calibration Sample

The WIMP efficiency and surface event rejection of the timing cut are defined using nuclear recoils from ^{252}Cf data and low-yield events selected from ^{133}Ba data (the energy distribution of these events is re-weighted to reflect the low-yield ER energy distribution in WIMP-search data). All events in both samples must pass all of the data quality and physics cuts described above (except for the veto cut, as veto information is not recorded during calibration runs). Events in the surface event sample must be at least 5σ below the electron recoil band and exhibit yields less than 0.7 and greater than 0.1. Many of the surface events in a given detector can be tagged as phonon-side or charge-side events using coincident recoils in the detectors above or below, respectively. However, this information is only available on the charge side of the two endcap detectors (T2Z1 and T4Z1). In these cases, a phonon-side beta fraction is estimated using the number of untagged and charge-side betas in the endcaps and the ratio of phonon-side to all betas (tagged and untagged) in the interior detectors.

7.7.2 The χ^2 cut

Following a method developed by Joseph Manungu Kiveni [170], we utilize the timing parameters `pdelc` (the position-corrected time between the 20% risetime of the charge pulse and the 20% risetime of the phonon pulse) and `pminrtc` (the position-corrected time between the 10% and 40% risetimes of the phonon pulse) to differentiate nuclear recoils from surface electron recoils. To determine the similarity between an individual low-background event and the neutron and surface event calibration

datasets, a χ^2 distance is defined:

$$\chi^2 = \sum_{jk}^N (x_j - \mu_j(E)) C_{jk}^{-1}(E) (x_k - \mu_k(E)) \quad (7.8)$$

where the j, k indices represent the two variables `pdelc` and `pminrtc`, μ_j is the mean of x_j , C is the covariance matrix, and E is the event energy (`pric`, the position-corrected recoil energy determined using the inner ionization signal). The means and the covariance matrix elements are determined in bins and fitted to the functions

$$\mu_j(E) = A_{1,j} + \frac{A_{2,j}}{E} + A_{3,j} E^2 \quad (7.9)$$

$$C_{jk}(E) = B_{1,jk} + \frac{B_{2,jk}^2}{E^2} \quad (7.10)$$

Figures 7-27 and 7-28 show example fits of the means and covariance matrix elements, respectively, in Detector T4Z3. Care must be taken to ensure that the determinant of the fitted covariance matrix remains positive at all energies, as it must by definition. Fits of the means and covariance matrix elements are performed separately for the neutron and surface event calibration sets to define two χ^2 variables, called χ_n^2 (χ^2 distance from the nuclear recoil distribution) and χ_b^2 (χ^2 distance from the beta distribution). Figure 7-29 shows $\chi_b^2 - \chi_n^2$ vs. recoil energy for neutrons (green) and face-tagged betas (blue) in Detector T4Z3.

The surface events are then rejected using two cuts. The first is called the neutron consistency cut, which demands that χ_n^2 is less than a particular value. This cut is not optimized in any methodical way, but is instead set to allow 90% neutron passage. The second cut is called the rejection cut, and requires that $\chi_b^2 - \chi_n^2$ be greater than a particular value. The position of the rejection cut will be referred to as η . The

rejection cut is defined in three energy bins for each detector, and the optimization of this cut is described in the upcoming sections.




Figure 7-27: Fits of the mean $pdele$ and $pminrt$ in Detector T2Z4. Neutrons are shown in green, phonon-side surface events in blue, and charge-side surface events in red. [171]




Figure 7-28: Fits of the elements of the covariance matrix between p_{del} and p_{minrt} in Detector T2Z4. Neutrons are shown in green, phonon-side surface events in blue, and charge-side surface events in red. [171]




Figure 7-29: $\chi_b^2 - \chi_n^2$ vs. recoil energy of neutrons (green) and surface events (blue) in Detector T2Z4. [171]

7.7.3 Nuclear Recoil Consistency Cut

The NE consistency cut is defined immediately after defining the timing χ^2 variables. Unlike the surface event rejection cut, the NR consistency cut is un-optimized and defined in an energy-independent fashion. The nuclear recoil consistency cut requires the event to be similar to the set of ^{252}Cf nuclear recoils. This cut is implemented as an upper limit on χ_n^2 and is defined to accept 90% of the neutron calibration events, without accounting for energy dependence in χ_n^2 .

7.7.4 Wimp Search/Calibration Systematics

It is important to consider the systematic differences between the calibration beta sample and the distribution of surface events in the WIMP search data. Figure 7-30 shows the timing distributions of side-tagged betas on interior detectors in barium calibration data in R125-128. The detectors clearly respond differently to phonon-side and charge side betas, and this must be accounted for when estimating leakage; because the timing also depends on the energy (7-27), systematic differences in the energy distributions of surface events in barium calibration data and WIMP search data must also be considered. The method used to account for these systematic effects in the leakage estimation are described in the next section.

7.7.5 Leakage Estimation

To estimate the expected leakage for a given cut defined on calibration betas in the presence of the systematics above, the calibration surface events are binned according to energy and face, the expected leakage in each bin is measured, and the expected leakages are scaled using scaling factors measured in a previously unblinded data




Figure 7-30: Timing distributions of side-tagged betas on interior detectors in barium calibration data in R125-128. [172]

set. As in [52], the expected leakage on detector i is estimated using the following equation:

$$n_i = N_i \sum_{e,f} b_{e,f}^{(i)} s_{e,f}^{(i)} \quad (7.11)$$

where e, f represent energy and face bins, respectively, N_i is the expected number of nuclear recoil single scatters (NRSSs), $b_{e,f}^{(i)}$ is the surface event passage fraction of the timing cut on detector i in bin e, f measured in calibration data, and $s_{e,f}^{(i)}$ is the fraction of face-tagged low-yield events in bin e, f in WIMP search data. The energy bins are 7-20 keV, 20-30 keV, and 30-100 keV. $s_{e,f}^{(i)}$ is a set of weighting factors that effectively scales the expected leakage measured in calibration data to account for the differences in energy and face distribution between WIMP search and calibration data. N_i is measured for each detector from the unblinded Run 123-124

data and scaled up to account for the increased livetime of the c58 analysis. $s_{e,f}^{(i)}$ is also measured using multiple scatters from the c34 data but requires no livetime scaling.

Table 7.2 shows the number of NRSSs in c34, the livetimes of c34 and c58, and the derived expected number of NRSSs in c58. Note that T4Z3 is a special case - in c34 and Run 125, it has a low-energy threshold of 15 keV due to excess noise, but after a hardware replacement between Runs 125 and 126, it is capable of reaching a 7 keV threshold in Runs 126-128. To estimate the expected number of NRSSs in c58 for T4Z3, Run 125 is treated the same as other detectors - scale the number of c34 betas by the ratio of c34 livetime to Run 125 livetime. In Runs 126-128, the number of NRSSs observed in c34 is also scaled by the ratio of NRSSs in 7-100 keV to the number of NRSSs in 15-100 keV observed in other interior detectors to account for the change in threshold from c34 to Runs 126-8.

Figure 7-31 shows the fractions of face-tagged low-yield events in the energy and face bins in c34 WIMP search data (red) and in c58 WIMP search data (black). The phonon-side fractions on the endcap detectors have been estimated as described in Sec. 7.7.1. In general, the variation across detectors is smaller than the statistical errors on individual detectors, and so I simply use the average energy/face fractions across the various detectors. In a few detector/energy/face bins, $s_{e,f}^{(i)}$ is 0, and using the average $s_{e,f}^{(i)}$ ensures that these bins are still accounted for in the expected leakage calculation. It is not possible to directly measure the scaling factors for the endcaps, given that one detector face has no adjacent detector to tag multiples. It is possible to determine weighting factors for the endcaps using the singles/multiples and p-side/q-side ratios in the interior detectors along with the measured numbers of singles and tagged multiples on the endcaps. It was found that the resulting weighting factors for the endcaps were in statistical agreement with those of the interior detectors,

Detector	c34 NRSSs	c34 LT (days)	c58 LT (days)	Exp. NRSSs in c58
T1Z4	7	48.95	62.55	8.945 ± 3.38
T2Z1	23	87.30	167.83	44.216 ± 5.39
T2Z2	6	81.01	115.50	8.554 ± 2.18
T2Z4	14	119.37	187.21	21.957 ± 3.23
T3Z3	9	52.72	160.90	27.4669 ± 5.92
T4Z1	13	148.91	208.19	18.175 ± 2.81
T4Z3	2	146.87	217.28	6.666 ± 2.75
T5Z3	10	85.78	203.42	23.714 ± 4.31

Table 7.2: Factors entering the computation of the expected number of NRSSs in the c58 silicon WIMP search data. The observed number of NRSSs in c34 (col. 2) is multiplied by the ratio of c58 livetime to c34 livetime (in kilogram-days, columns 4 and 3, respectively) to produce the expected number of NRSSs in c58 (column 5). Errors in column 5 are $\sqrt{N}$. The expected number of NRSSs in T4Z3 is also corrected to account for the change in low-energy threshold between Runs 125 and 126.

and so the mean weighting factors over all interior detectors are also applied to the endcaps.




Figure 7-31: Weighting factors $s_{e,f}^{(i)}$ for the 6 energy/face bins. Interior detectors are shown in blue, and endcap detectors are shown in red after correction for the lack of phonon-side tagging. The mean weighting factors across the 6 interior detectors and error on the mean are shown as a red band. [172]

With these definitions, the expected leakage L can be computed versus the cut position η in each detector and energy bin. The leakage curve is then fit to the functional form:

$$L(\eta) = A \times \ln(1 + e^{-B\eta^C}) \quad (7.12)$$

Figure 7-32 presents fits to the estimated leakage in the three energy bins for Detector T4Z3. Note that even with the increased statistics of the calibration dataset (relative to the low background dataset), the tails of the calibration β distributions are not particularly well-constrained. Therefore, during the fitting process, I aim to guarantee that the fits produce conservative leakage estimates by fitting to the "right edge" of the stairs shown in the figure and adjusting fit ranges to more heavily weight the beta tails on the right. Therefore, in some examples, the fit does not perform particularly well at high leakages (low values of the rejection parameter), but it is far more important to have accurate and conservative fits towards the low-leakage end of the distribution.

7.7.6 Spectrum-Averaged Exposure

The raw livetime presented in Table 7.1 does not directly determine the WIMP sensitivity of a detector. The efficiency of all event-by-event background rejection cuts vs. energy and the expected WIMP recoil energy spectrum must be included. These three parameters can be combined to determine a spectrum-averaged exposure (SAE) as follows:

$$SAE = LT \frac{\int_{E_{min}}^{E_{max}} \frac{dR}{dE} \epsilon(E) dE}{\int_{E_{min}}^{E_{max}} \frac{dR}{dE} dE} \quad (7.13)$$

where LT is the total livetime, $\epsilon(E)$ is the combined efficiency of all event-by-event cuts at the recoil energy E , and $\frac{dR}{dE}$ is the expected WIMP recoil energy spectrum (the differential rate per unit energy) at a given WIMP mass. LT is measured directly, $\epsilon(E)$ can be computed using calibration data, and $\frac{dR}{dE}$ can be computed using Eq. 2.11. The efficiencies of all cuts except the timing cut are computed in energy bins and fit vs. energy, and then the total SAE can be computed as a function of the timing cut position and used in concert with the expected leakage vs. cut position to define an optimized cut. Similar to the leakage, the SAE is computed within 3 different energy bins, and the timing cut is optimized in each of these bins. Figure 7-33 shows the SAE for a 60 GeV/c² WIMP vs. the rejection cut position (η) in the three recoil energy bins for all detectors, along with fits to the function:

$$SAE(\eta) = (A\eta + B) \times e^{-C\eta^D} \quad (7.14)$$




Figure 7-32: Estimated and fit leakage vs. rejection cut position in all Detectors in recoil energy bins: 7-20 keV (red), 20-30 keV (blue), 30-100 keV (green). The cut position is offset by 10 to allow negative cut values to appear on the log scale x-axis. Because the cut will be placed at low expected leakage, the fit is only required to be good at leakages below ~ 0.5 events.




Figure 7-33: Measured and fit Spectrum-Averaged Exposure to a $60 \text{ GeV}/c^2$ WIMP vs. rejection cut position in all detectors in recoil energy bins: 7-20 keV (red), 20-30 keV (blue), 30-100 keV (green). The cut position is offset by 10 to allow negative cut values to appear on the log scale x-axis. The fit is only required to be good at values of the rejection parameter >10 .

7.7.7 Optimization of the Rejection Cut

With the leakage (L) and exposure (SAE) vs. cut position (henceforth labeled η) curves computed in the previous step, we can compute SAE vs. L curves, which are shown in Figure 7-34. The next task is to determine how to determine the optimal set of cut values from these curves that provides maximal exposure at a given leakage. To do this, I compute the derivative of the exposure with respect to the leakage ($\frac{dSAE}{dL}(L)$) for each detector and energy bin. The slope $\frac{dSAE}{dL}(L)$ vs. the leakage L is presented in Figure 7-35 for a WIMP mass of $60 \text{ GeV}/c^2$. Requiring the slope to be constant across all detectors and energy bins ensures that no exposure can be gained by reappportioning a differential amount of leakage from one detector/bin to another. By scanning a range of slopes and recording the corresponding cut positions in each bin, we can thus determine an optimal exposure vs. leakage curve, and choose our set of cuts to lie somewhere on the curve.

Because the expected WIMP spectrum depends on the WIMP mass, the routine to optimize the cut actually uses 3 WIMP masses to select cut positions in the 3 energy bins: $15 \text{ GeV}/c^2$ in the 7-20 keV bin, $30 \text{ GeV}/c^2$ in the 20-30 keV bin, and $60 \text{ GeV}/c^2$ in the 30-100 keV bin. The optimization routine proceeds as follows:

1. Begin with the SAE fits for a $15 \text{ GeV}/c^2$ WIMP.
2. Select a value of $\frac{dSAE}{dL}(L)$, and determine SAE and L and η in all detectors and energy bins at that value of $\frac{dSAE}{dL}(L)$. Because $\frac{dSAE}{dL}(L)$ is constant across all detectors and energy bins, no exposure can be gained by changing the cut to move leakage from one bin to another.
3. Repeat Step 2 over a range of values for $\frac{dSAE}{dL}(L)$. Scanning over the full range of slopes produces an optimal SAE vs. leakage curve, and a unique set of η

values in each bin corresponds to each point on this SAE vs. leakage curve.

4. Choose the set of η values that provide the best sensitivity. I choose $\frac{N_{68}(L)}{SAE}$ as the metric for sensitivity, where $N_{68}(L)$ is the 68% upper Poisson confidence limit on the number of expected leakage events. To find the final cut, choose the values of η that provide the optimal (lowest value) of the sensitivity.
5. Fix the cuts for the 7-20 keV bin at these values of η .
6. Repeat steps 1-5, using a 30 GeV/c² WIMP for the 20-30 keV bin and a 60 GeV/c² WIMP for the 30-100 keV bin. In each re-optimization, the cuts set in the lower energy bins in prior iterations are held fixed. That is, the cuts in the 20-30 keV bin are optimized to provide the best sensitivity to a 30 GeV/c² WIMP accounting for the additional exposure and leakage already provided by the fixed cuts in the 7-20 keV bin.

7.7.8 Simultaneous Optimization of the Nuclear Recoil Bands

Because the surface event background is composed of electron recoils with reduced ionization yields, they predominantly leak into the nuclear recoil region from above. The yields of nuclear recoil events, on the other hand, are roughly normally distributed around the NR band mean, which implies that the optimal signal region is likely not the $\pm 2\sigma$ nuclear recoil band. Rather, a more constrained NR band may remove a substantial amount of leakage while preserving a favorable proportion of the nuclear recoils. Figure 7-36 presents the cumulative yield distributions of calibration neutrons and low-yield multiple-scatter events from WIMP search data, showing that the distribution of low-yield multiples is indeed heavily weighted towards the upper end of the nuclear recoil band.




Figure 7-34: SAE vs. L in all energy bins and detectors for a $60 \text{ GeV}/c^2$ WIMP.

The NR band can easily be optimized simultaneously with the rejection cut assuming that the two are independent. Rather than simply optimizing the sensitivity along the SAE vs. leakage curve, the upper and lower limits of the NR band can be added as additional dimensions to the optimization space. The range of potential NR band limits is constrained to be within $\pm 2\sigma$ and the neutron and surface event efficiencies are computed in steps of 0.1σ , so it is not computationally expensive to simply directly scan the entire space of potential NR band widths at each potential position on the SAE vs. leakage curve to find the absolute optimum sensitivity. The final optimal NR band ranges from -1.8σ to $+1.2\sigma$ of the NR band mean.




Figure 7-35: $\frac{dSAE}{dL}$ vs. L in all energy bins and detectors for a 60 GeV/ c^2 WIMP. Curves for individual bins are cut off to only show regions where $\frac{d^2SAE}{dL^2} < 0$.

7.7.9 Final rejection cut

The final rejection cut is illustrated in Figures 7-37. The rejection cut is illustrated by the magenta line; if no line is visible in a particular detector/energy bin, then the rejection cut was set to infinity in that bin. This occurs when the slope of the exposure vs. leakage curve in a particular bin never reaches the slope that corresponds the optimal timing cut; i.e., more exposure can be gained by apportioning additional leakage into other bins than by including the bin in question. The calibration betas are shown in blue; those events that pass the neutron consistency cut are circled in red.




Figure 7-36: Cumulative distribution of yields from -2σ to $+2\sigma$ of the NR band for neutrons (blue solid) and surface events (black dashed). The surface event distribution is heavily weighted towards higher yields.

7.8 Conclusion

After the timing cuts are set, we compute the final efficiency of all WIMP search cuts. Figure 7-38 presents the nuclear recoil efficiency (left axis) and total exposure (right axis) vs. recoil energy in this analysis. The overall efficiency is computed by averaging the individual detector/run efficiency curves weighted by the corresponding livetime in that detector/run. The sudden jumps in the efficiency curve arises from the different detector thresholds and from the energy-binned nature of the timing cut.

Finally, Figure 7-39 presents the sensitivity of this analysis, which I define as the

upper 90% one-sided confidence limit on the WIMP-nucleon cross-section that would be set if no events are observed upon unblinding. Results from other experiments are included for comparison.




Figure 7-37: The rejection cut is illustrated by the magenta line; if no line is visible in a particular detector/energy bin, then the rejection cut was set to infinity in that bin. The calibration betas are shown in blue; those events that pass the neutron consistency cut are circled in red.




Figure 7-38: Nuclear recoil efficiency (right axis) and total exposure (left axis) vs. recoil energy of the c58R Si analysis.




Figure 7-39: Sensitivity of the c58R Si analysis (the limit that would result if no events are seen in the signal region). Results from other direct dark matter searches are included for comparison.

Chapter 8

WIMP Search Results from the CDMSII silicon detectors

8.1 Introduction

With all data quality, physics, and background rejection cuts in hand, it is time to unblind the signal region! Unblinding took place on December 25th, 2012, and thus the work in this chapter is all extremely recent at the time of writing. The collaboration is preparing to publish the results, but studies related to much of the work detailed below are ongoing, and some of the numbers presented below may be subject to change after this thesis is finished and submitted. Every effort has been made to ensure consistency between this chapter and the forthcoming publication, but the ongoing nature of the investigations means that certain details may change as new information comes to light. Of course, I am referring only to numbers such as the total background estimate, exact positions of limit curves/contours, and goodness-of-fit numbers; the results of the unblinding itself are set in stone and will not change.

For the official publication of these results, I refer readers to the pre-print at [173].

In this chapter, I first review a full computation of the expected leakage from neutrons, photons, and surface events. I then present the results of unblinding. Three events within the signal NR band and passing the timing cut are observed. Throughout this chapter, I will refer to these three events as ‘candidates’. I proceed with some basic checks to ensure that these events do not exhibit any obvious pathologies that went unnoticed in the analysis of calibration data. Finally, a limit curve on the WIMP-nucleon scattering cross-section vs. WIMP mass is computed under standard assumptions about the astrophysical dark matter distribution and without background subtraction, under the hypothesis that all three candidates are true WIMPs. This assumption produces the most conservative set of limits. I also present a contour outlining the region of interest indicated by the energies of the three events; this contour is calculated using a profile likelihood method. These results do not rise to the level of WIMP discovery.

8.2 Pre-unblinding

Before fully unblinding and revealing the events in the signal region *passing* the timing cut, I first look at events in the signal region *failing* the timing cut. Comparing this against the estimate of nuclear recoil single scatters (NRSSs) failing the timing cut shown in Table 7.2 provides a final check before fully unblinding the signal region. The counts of observed events in the signal region failing the timing cut are shown in Table 8.1. The expected number of events are computed from the observed number of c34 signal region events, though the estimates do not exactly match those of Table 7.2 because the optimization routine restricted the NR band from the $\pm 2\sigma$ used in c34 to -1.8σ - $+1.2\sigma$. P-values for the observed vs. expected numbers are

computed through Poisson likelihood estimates, and the total P-value is computed using Fisher’s method for combining binned p-values.

Detector	Observed NRSSs failing	Expected NRSSs failing	p-value
T1Z4	4	5.11 ± 2.56	0.7291
T2Z1	21	36.52 ± 8.38	0.0837
T2Z2	9	7.13 ± 3.19	0.6729
T2Z4	8	14.12 ± 4.71	0.2431
T3Z3	15	12.21 ± 6.10	0.7090
T4Z1	15	13.98 ± 4.43	0.8629
T4Z3	5	2.26 ± 2.26	0.4319
T5Z3	11	18.97 ± 6.71	0.2503
Total	88	110.30 ± 14.70	0.5507

Table 8.1: Observed and expected counts of NRSSs failing the timing cut

Though the observed number of NRSSs failing the timing cut is smaller than expected, the disagreement is not particularly worrisome given the large statistical uncertainties on the expected number of events. Subsequent analysis has revealed that the yield distributions after position correction in c34 and c58 are different, and accounting for this systematic difference does bring the numbers into better agreement (90.177 ± 13.107 expected, [174]). The difference in yield distributions could be explained by both the implementation of a new charge trace optimal filter in c58 and by substantial updates to the position-correction scheme compared to that used in c34. Note that the expected numbers of events shown in Tables 7.2 and 8.1 are different; these differences arise due to the constriction of the NR band from $\pm 2\sigma$ to -1.8σ to $+1.2\sigma$.

8.3 Final Leakage Estimates

8.3.1 Surface Event Leakage

The pre-unblinding surface event leakage estimates presented in Section 7.7.5 relied on expected numbers of NRSSs failing the timing cut, which were scaled from previously unblinded silicon data. Then, the passage rate of the timing cut for surface events was measured on barium calibration data, weighted according to energy and face bin using Equation 7.11, with weighting factors also measured in c34. The expected leakage was then fit vs. timing cut position. Statistical and systematic errors on the fits were not accounted for; the fits were simply generated to be overly conservative to guard against these errors.

The above process is a satisfactory method of estimating the leakage to set the timing cut. After ‘pre-unblinding’, the true number of NRSSs failing the timing cut and the face/energy distributions of multiples in c58 can be used to generate improved leakage estimates, and statistical and systematic errors become more important to accurately estimate. The leakage is estimated using equations of the form

$$n_i = N_i r_i \tag{8.1}$$

where n_i is the number of expected NRSSs *passing* the timing cut on detector i , N_i is the known number of NRSSs *failing* the timing cut, and r_i is an estimator of the timing cut’s pass-fail ratio. r_i is measured on three datasets: multiple scatter events inside the signal NR band in WIMP search data, multiple scatters outside of the signal NR band but below the ER band in WIMP search, and Ba calibration multiples below the ER band. Systematics arise from potential differences in behavior between the set of NRSS surface events and these three datasets.

Statistical uncertainties pose somewhat of a challenge. The WIMP search datasets used to estimate r_i have fairly small numbers of total events, and because the timing cut is set to fail the vast majority of surface events, r_i is often estimated to be 0 for a particular detector. Statistical error estimators that use asymptotic assumptions will fail, and simple propagation of Poisson errors will result in mostly 0 ± 0 estimates. To produce reliable estimates of the error on the leakage, Jeff Filippini developed a Bayesian Monte Carlo technique for Runs 123-124. A full description of the method is available in Appendix C of [90].

A simple, and probably not fully satisfying description of the technique follows. The basic idea is to throw a large number of Monte Carlo trials for the total expected leakage n to produce the posterior distribution of the total expected leakage. The actual posterior distribution for n (the sum over $n_i = N_i r_i$) is somewhat complicated, but it is equivalent to draw the posterior distributions of N_i and r_i individually, given the observed counts and a set of assumed priors, and then compute the posterior distribution for n . What remains is the choice of the priors. Choosing the appropriate conjugate prior (please see [90]) allows us to bypass difficult numerical integrations and sample directly from the posterior distributions for N_i and r_i . N_i is Poisson-distributed, and the appropriate conjugate prior is the gamma distribution. Various factors that go into computing r_i in the three methods can be Poisson or binomially-distributed; the appropriate conjugate prior to the multinomial distribution is the Dirichlet distribution. It is computationally cheap to throw hundreds of thousands of trials from the gamma and Dirichlet distributions to produce the posterior distribution of n . The final leakage estimate is then taken as the median of this distribution, and statistical errors are obtained by computing the [15.87, 84.13] percentiles. Certain parameters of the conjugate priors must be chosen; some details on this choice are included in the systematics section.

Method 1: WIMP search multiples in the NR band

Measuring the surface event pass/fail ratio using multiple scatters inside of the signal NR band in WIMP search data should provide the least systematic error when applied to single-scatter surface events inside the NR band. Unfortunately, there are relatively few WIMP search NR multiples, and only 1 of them (over all 8 detectors) passes the timing cut. The leakage is derived from the equation

$$n_i = \sum_i N_i \frac{b_i}{B_i} \quad (8.2)$$

where N_i is the observed count of NRSSs failing the timing cut on Detector i , b_i is the observed count of NRMS passing the timing cut on Detector i , and B_i is the observed number of NRMS failing the timing cut on Detector i . The Bayesian leakage estimate for Method 1 treats N_i , b_i and B_i each as Poisson parameters, and I throw 100k trials to compute the posterior distribution. The resulting expected leakage is

$$n = 0.90^{+1.51}_{-0.64}(\text{stat.})^{+0.40}_{-0.32}(\text{syst.}) \quad (8.3)$$

Method 2: WIMP search multiples outside of the NR band

Method 2 utilizes WIMP search multiples outside of the NR band and more than 3σ below the ER band. Similar to what was done using the Ba calibration data in Section 7.7.5, the multiples outside of the NR band are binned according to energy (7-20 keV, 20-30 keV, 30-100 keV) and face (phonon-side or charge-side), and the factor r_i is derived from the passage fraction in each bin weighted according to the

observed energy/face distribution of WIMP search multiples inside of the NR band.

$$n_i = N_i \sum_{e,f} \frac{m_{e,f}^{(i)}}{M_{e,f}^{(i)}} s_{e,f}^{(i)} \quad (8.4)$$

In this equation n_i and N_i are defined as in Method 1. $m_{e,f}^{(i)}$ is the number of multiples outside of the NR band that pass the timing cut, and $M_{e,f}^{(i)}$ is the number of multiples that fail the timing cut. $s_{e,f}^{(i)}$ is the weighting factor used to account for systematic differences in energy and face distribution between the events outside and inside of the signal NR band. Unlike in Section 7.7.5, we can now use the NR band multiples observed in Runs 125-128 rather than those in Runs 123-124. Similarly to the previous estimates of the weighting factors, the statistical errors on the factors are larger than the variation across detectors, and correcting the weighting factors for the endcaps (which can have face-tagging information on only one face) provides scaling factors statistically similar to those on the interior. Thus, the mean weighting factors from the interior detectors are used for all detectors. To extract the Bayesian leakage estimates in this method, N_i , $m_{e,f}^{(i)}$, and $M_{e,f}^{(i)}$ are treated as Poisson-distributed variables with posteriors drawn from the gamma distribution, and $s_{e,f}^{(i)}$ are treated as multinomially-distributed variables and drawn from the Dirichlet distribution. The final leakage estimate using this method is:

$$n = 0.68_{-0.25}^{+0.44}(\text{stat.})_{-0.25}^{+0.40}(\text{syst.}) \quad (8.5)$$

Method 3: Ba calibration multiples below the ER band

Method 3 utilizes the same form as Method 2, but $m_{e,f}^{(i)}$ ($M_{e,f}^{(i)}$) are taken from the face-tagged multiples in barium data that pass (fail) the timing cut, providing higher

statistics than either Method 1 or Method 2. The leakage estimate using this method is:

$$n = 0.31^{+0.13}_{-0.09}(stat.)^{+0.21}_{-0.11}(syst.) \quad (8.6)$$

Combined Surface Event Leakage Estimate

To combine the surface event leakage estimate, all three methods are simulated at once, with the correlating parameters (the $s_{e,f}^{(i)}$) drawn once in each trial and used in both Methods 2 and 3. The resulting posterior distributions from all 3 methods are combined to provide the combined leakage estimate. The result is somewhat dominated by Method 3, which has higher statistics and thus a tighter posterior distribution. Figure 8-1 presents the resulting posterior distributions for the surface event leakage. Systematic errors are excluded from the simulations to allow an estimate of the statistical errors on the final leakage estimate. The systematic errors described above are added in quadrature and applied to the final leakage estimate. The resulting expected leakage is:

$$n = 0.41^{+0.20}_{-0.08}(stat.)^{+0.28}_{-0.24}(syst.) \quad (8.7)$$

Systematics

Multiple sources of systematics are present in these leakage estimates.

1. The choice of the prior distribution. To choose the parameters of the prior distributions for each method, I start by generating thousands of mock datasets and computing the ‘true’ leakage n for each mock dataset. Then I run 5000 trials of the Bayesian leakage estimation scheme for each mock dataset and




Figure 8-1: Posterior distributions of the expected leakage from surface events, with statistics only (blue) and statistics + systematics (red).

scanning over a set of potential prior parameters. The parameter that minimizes the mean bias (the difference between the ‘true’ leakage of the mock dataset and the median of the derived prior distribution) is chosen, but there is still some spread in the bias over the different mock datasets. The systematics assigned to Methods 1, 2, and 3 below from the choice of prior are ± 0.02 , $^{+0.01}_{-0.08}$, and ± 0.01 , respectively.

2. Using the mean weighting factors across all detectors (Methods 2 and 3 only). I utilize the mean weighting factors over all interior detectors in Eq. 8.4, as the individual detectors’ weighting factors have large statistical errors relative to the variation from detector-to-detector. I utilize the standard deviation of the individual detector weighting factors over the mean weighting factors to

estimate this systematic, which gives a $\pm 12\%$ systematic error on the leakage estimates of Methods 2 and 3.

3. The assumption that singles track multiples, i.e., the assumption that the passage fraction of multiples is the same as that of the singles. This assumption can be tested on the WIMP search multiples and singles outside of the signal region and below the ER band (cut at -3σ). The best estimate for the singles passage fraction over the multiples passage fraction is 1.7, with a 68% confidence interval of [1.13, 2.54], and this result is in agreement with the singles/multiples passage fractions for the timing cut used in the Run 123-124 silicon analysis. The correct way to handle this systematic offset is *not* to just assign a 70% upper systematic error on the estimated leakages, but rather to scale the estimated leakage by 1.7, and add systematic errors that account for the confidence interval on the singles/multiples passage fractions, [-34%, +44%]. This scaling and systematic error are already contained in all of the estimates quoted above.
4. The assumption that the variation in surface event timing cut passage rate depends primarily on recoil energy in Methods 2 and 3. In fact, variation of this passage rate vs. the charge energy may be equally or more important, particularly at low nuclear recoil energies where the phonon signal is still well above the noise but the charge pulse signal-to-noise ratio can be as low as 5, affecting the measurement of phonon delay. This systematic is tested using a binning and scaling scheme similar to Methods 2 and 3, but binning in charge energy rather than recoil energy, and the result is a [-46%, +46%] systematic error. This systematic is not quoted in the leakage estimates above but is included in the final leakage estimate.

8.3.2 Bulk Gamma Leakage

With the lack of any low-energy activation lines in silicon, it is difficult to directly measure the energy resolution of bulk gammas at energies near the threshold of this analysis. However, after position correction, the bulk gamma yield distribution is very well-described by a Gaussian, so it is easy to determine what fraction of this Gaussian overlaps with the signal nuclear recoil band. Figure 8-2 presents this gamma leakage fraction averaged over all detectors on the left. The expected bulk gamma leakage can then be estimated by binning and counting the number of bulk gammas in the $\pm 2\sigma$ electron recoil band that pass all analysis cuts, which is shown in the middle panel of Figure 8-2. This study was performed prior to the surface event rejection cut definition, so I only counted all gammas passing all other cuts. Finally, the expected number of bulk gammas vs. energy in the nuclear recoil band can be computed from the gamma energy distribution and the overlap of the gamma yield Gaussian with the NR band. The right panel of Figure 8-2 shows the integrated expected gamma leakage vs. energy, where the integrated leakage is defined as the sum of the expected leakage from an energy of 'x' up to 100 keV.

Because the final expected gamma leakage is less than .01 events, I do not include bulk gammas in the final background estimate.

8.3.3 Cosmogenic neutrons

Multiple simulations of the cosmogenic neutron production in the Soudan facility with the c58 veto and detector setup have been performed using GEANT4 and FLUKA. These simulations produce overall rates of neutron production in muon spallation showers, as well as event ratios (veto-coincident to unvetoes nuclear recoils, singles vs. multiples, etc.). To estimate the cosmogenic neutron leakage in this




Figure 8-2: Estimating the bulk gamma leakage. Left: Fraction of bulk gammas expected to leak into the signal NR band, averaged over all detectors. Middle: Histogram of all WIMP search gammas that pass all cuts except the timing cut (Q-in, single, veto-anticoincident, data quality, thresholds, etc.). Right: Integrated expected leakage vs. energy.

analysis, I utilize the measured number of veto-coincident nuclear recoils and the event ratios from simulations; the event ratios should be less sensitive to assumptions regarding the incident muon flux than the absolute rates from simulation. There are 25 total (single and multiple scatters) veto-coincident nuclear recoils ($\pm 2\sigma$ NR bands) in the silicon detectors in the c58 WIMP search data. In GEANT4 simulations presented by Angela Reisetter [175], 3 unvetoes were observed in the Si detectors, and 245 vetoed total NRs were observed. I use the equation

$$N_{\text{veto-anticosingle,data,exp.}} = N_{\text{veto-co,data,obs.}} \frac{N_{\text{veto-antico,single,MC}}}{N_{\text{veto-co,MC}}} \quad (8.8)$$

to determine an expected raw rate of unvetoes NR singles in the $\pm 2\sigma$ band (0.31 events). This number is then scaled by the relative acceptance of the -1.2σ to $+1.8\sigma$ signal NR band (.9075) and the average NR timing cut efficiency weighted by the simulated veto-coincident NR spectrum (.1678) to get a final leakage estimate of

$$N_{veto-antico,single,data,exp.} = 0.047^{+0.045}_{-0.026}(stat.)^{+0.104}_{-0.039}(syst.) \quad (8.9)$$

Statistical errors are easily derived from Poisson and binomial errors on Eq. 8.8. The quoted upper systematic error is computed by considering other binnings of the data/MC, for example, considering only veto-coincident singles from data (7) and MC (18) in Eq. 8.8 produces an estimate of 0.151 events. The lower systematic error is computed using a second simulation [176], which observed only 4 veto-antico single NRs in the operating Si detectors compared to 1428 total veto-coincident NRs.

8.3.4 Radiogenic neutrons

Neutrons can be produced in radioactive processes in the materials surrounding the detectors, through fission or (α,n) reactions. The primary sources are trace uranium and thorium contents in the copper hardware, polyethylene shielding, and lead shielding. The uranium and thorium contents of the experimental materials can be determined both with gamma-counting measurements or with global Monte Carlo fits to the detectors' observed gamma spectrum during WIMP search data taking. Simulations of the neutron flux with the known U/Th concentrations can thus be used to determine the expected rate of neutron interactions in the detectors [177]. The results of these simulations are summarized in Table 8.2.

With 140 kg-days of exposure, we expect <0.062 radiogenic neutrons in the low-background data. Statistical errors from the MC are negligible. However, estimating the neutron rates from measured γ rates relies on the assumption of secular equilibrium in the radioactive chain. This assumption introduces a systematic error that has not yet been quantified, and so I assign a 100% systematic error to this estimate in the upper direction. In the lower direction, note that in Table 8.2, approximately

Component	U (ppb)	Th (ppb)	Mass (kg)	Si n rate (#/kg-year)
Copper cold hardware	0.16	0.25	14.95	1.1×10^{-2}
Copper icebox cans	0.18	0.56	260	6.0×10^{-2}
Inner poly	<0.12	<0.12	120	$<1.6 \times 10^{-2}$
Inner lead	<0.05	<0.2	1917	$<2.4 \times 10^{-2}$
Outer lead	<0.05	<0.4	12190	$<5.3 \times 10^{-2}$
Total				< 0.164

Table 8.2: Contamination levels in the experimental materials and corresponding expected Si neutron rates.

55% of the total expected neutron leakage comes from upper limits rather than well-defined rates, so I arbitrarily assign a systematic in the lower direction of 55% of the total leakage.

$$N_{radiogenic} = 0.062^{+0.062}_{-0.027}(syst.) \quad (8.10)$$

Please note that new simulations of both the cosmogenic and radiogenic neutrons are underway, and these numbers may be updated prior to publication of these results.

8.4 Unblinding

With the timing cut and leakage estimate finalized, and the pre-unblinding revealing no excesses in the observed number of events failing the surface event rejection cut in the signal region, it is time to fully unblind the signal region and look for events passing the rejection cut.

When the signal region was unblinded, three events were observed. Two of these

events were in Detector T4Z3 (433, in CAP nomenclature), one at 9.51 keV recoil energy during Run 127 and another at 12.29 keV in Run 128. The third event is found in Detector T5Z3 (443) at a recoil energy of 8.20 keV during Run 126. There are no events just outside the signal band in yield that pass the timing cut. However, there is a single event in T5Z3 that passes the surface event rejection cut but fails the neutron consistency cut. Table 8.3 records some basic information about the events; I also number the candidates for reference throughout this chapter. Figure 8-3 shows all good single-scatters above thresholds from the WIMP search data in each of the 8 Si detectors used in this analysis. Events in red pass the neutron consistency cut, and events with black crosses pass the rejection cut. The signal NR band is outlined in green, the -3σ edge of the ER band is in blue, the charge threshold cut is shown in magenta, and the recoil energy threshold is the black dashed line. The width of these shaded lines indicate variations in the band and threshold definitions across runs and series. Figure 8-4 presents the same events with the same coloring scheme from Figure 8-3 in a space of normalized yield vs. the surface event rejection parameter. The normalized yield is the yield defined in terms of the number of σ s from the NR band mean. The signal region is outlined by a black box, with the timing cuts in the three energy bins represented by the dashed lines at the left edge of the box. The signal events in T4Z3 and T5Z3 are fairly well separated from other low background singles and from the rejection cut boundaries.

Candidate #	Detector	Series #	Event #	Recoil Energy (keV)
1	T4Z3	180701_1326	70474	9.51
2	T4Z3	180906_1139	80120	12.29
3	T5Z3	180314_1555	60048	8.20
“Near-miss event”	T4Z3	180626_1715	120442	9.61

Table 8.3: Information about the events in the signal region in this analysis.




Figure 8-3: Yield (yic) vs. recoil energy (pric, keV) of all good low-background events from Runs 125-128 in the 8 Si detectors of this analysis. Events in red pass the neutron consistency cut, and events with black crosses pass the rejection cut. The signal NR band is outlined in green, the -3σ edge of the ER band is in blue, the charge threshold cut is shown in magenta, and the recoil energy threshold is the black dashed line. The width of these lines indicate variations in the band/threshold definitions across runs and series.




Figure 8-4: All good low-background events from Runs 125-128 in the 8 Si detectors of this analysis. Events in red pass the neutron consistency cut, and events with black crosses pass the rejection cut. The y-axis is ‘normalized yield’, the yield expressed in terms of number of σ s from the NR band mean. The x-axis is the rejection parameter $\chi_b^2 - \chi_n^2$. The black box outlines the signal region, with the three dashed lines at the left edge indicating the rejection cut in the three energy bins.

Another interesting space to look at these events in is the space of timing variables, rather than in the timing χ^2 space. The top row in Figure 8-5 shows the candidate events in phonon delay vs. phonon risetime along with the neutron and barium surface event calibration event distributions used to set the rejection cut. The bottom row removes the barium surface event distributions and shows the neutrons, the candidate events, and all other signal region singles in these two detectors. These figures give a clearer picture of where the events lie relative to the surface event background. The three WIMP candidates are clearly marked in red. Finally, Figures 8-6, 8-7, and 8-8 present the raw phonon and charge traces from the candidate events.




Figure 8-5: Phonon delay vs. phonon risetime in the two detectors containing the three WIMP candidates. The top row shows the neutrons and barium surface event calibration sets, along with the candidate events. The bottom row shows the neutrons, the candidate events, and all other signal region singles in these two detectors.




Figure 8-6: Phonon and charge traces of Candidate 1




Figure 8-7: Phonon and charge traces of Candidate 2




Figure 8-8: Phonon and charge traces of Candidate 3

8.5 Checking the events

An array of studies were performed to understand the 3 events that were found in the signal region. I summarize a subset of the checks that turned up no obvious concerns, and go into a bit more detail on the studies that turned up potential concerns. Of course, if I look at a large number of event parameters, I'm bound to turn up something that seems concerning. These studies have so far revealed no major pathologies that would cause these events to be non-viable WIMP candidates.

8.5.1 Basic checks

1. At first glance, it seems odd to find two events out of three in a single detector. However, considering the distributions of exposure among the detectors, I find that given three events, 2 would be found in a single detector $\sim 48\%$ of the time.
2. The events have normal baseline noise levels.
3. Charge optimal filter χ^2 values indicate that the events match the OF template about as well as other events with similar ionization energies.
4. The overall gamma and surface event rates in the series containing the candidates are normal.
5. The noise levels of the series containing the candidates are normal.
6. There is no unusual veto activity in the time preceding the candidate events.
7. The KS test results for the series in question do not indicate any issues.

8. The delay and risetime (the two quantities that go into the computation of χ_n^2 and χ_b^2) are well-separated from the surface event distributions and consistent with the neutron distributions. The yields of the events are also deep in the tails of the surface event distribution.
9. Two of the three candidates are near the edges of the position correction manifold; however, the neutrons tend to form a different manifold than the γ s used to create the correction manifold, and the candidates are within the main body of the neutron distributions. All events easily pass the manifold χ^2 cut that was placed to ensure the quality of position correction.
10. The events are well-separated in time. The probability of getting no events out of 3 during Run 125, which was the longest run, is computed under simplifying assumptions to be 17%. This is a conservative estimate, as there was somewhat reduced exposure (per unit livetime) in Run 125 relative to other runs due to triggering issues in T2Z2 and noisy hardware in T4Z3 (which was replaced between Run 125 and 126).
11. The events are relatively close to the ionization threshold, which is set 4.5σ above the mean ionization energy measured on noise pulses. The thresholds in these series were not abnormal relative to surrounding series.
12. The only check that revealed something potentially problematic is that during one of the events, there was energy in an adjacent detector that was nearly high enough to trigger the multiple scatters veto. The other two events are good singles candidates.

8.6 Constraints on WIMP parameter space

The results of this analysis are now used to set constraints on the WIMP-nucleon scattering cross-section and mass. The optimal interval technique [178] is used to set an upper limit. A standard Maxwellian dark matter halo is assumed, with $v_0 = 220$ km/s, $v_{escape} = 544$ km/s, and ρ_0 , the local WIMP density, equal to $0.3 \text{ GeV}/c^2/\text{cm}^3$. The Helm nuclear form factor is used to compute the spin-independent recoil spectrum, and the spin-structure functions in [179] are used to compute spin-dependent scattering limits. The overall analysis efficiency presented in Figure 7-38 is of course taken into account when computing the limits. In all limits below, the three events are assumed to be WIMPs and no background subtraction is performed, so these are conservative upper limits on the WIMP-nucleon cross-section vs. mass.

8.6.1 Spin-Independent constraints

The 90% upper confidence limit on the spin-independent WIMP-nucleon cross-section is presented in Figure 8-9, along with a set of limits from other experiments for comparison. The limit does not exclude any new parameter space under these assumptions. However, it is worth noting that the competing limits at low WIMP mass, which come from CDMSII and EDELWEISS germanium detectors and the XENON collaboration, are more sensitive to astrophysical uncertainties than the results from Si detectors due to the differing masses of the target nuclei. For example, a $10 \text{ GeV}/c^2$ WIMP would have to travel at ~ 400 km/s, ~ 530 km/s, and ~ 670 km/s to induce a 7 keV recoil in Si, Ge, and Xe, respectively. Some authors [180] have developed means of integrating out astrophysical uncertainties to compare results from different target nuclei independently of astrophysical assumptions, but these methods have not yet been applied to this result.

Given the observation of three events and the low expected background, it is more interesting to consider the parameter space consistent with the hypothesis that these three events are true WIMP recoils. Julien Billard has led a likelihood-based analysis of these results. For the moment, the likelihood function considers the PDFs of the expected leakage and energy distributions from surface events, neutrons, and ^{206}Pb recoils from ^{210}Po decay:

$$\begin{aligned} \mathcal{L}_{i,j}(m_\chi, \sigma_{\chi-n}, \vec{\nu}_{i,j}) &= \frac{\exp\{-(\mu_{i,j} + \nu_{i,j,\beta} + \nu_{i,j,Pb} + \nu_{i,j,n})\}}{N_{i,j}!} \\ &\times \prod_{k=1}^{N_{i,j}} \left(\mu_{i,j} f_s^{i,j}(E_k) + \sum_{l=\beta,n,Pb} \nu_{i,j,l} f_l^{i,j}(E_k) \right) \times \prod_{l=\beta,n,Pb} \mathcal{L}_{i,j,l}(\nu_{i,j,l}) \end{aligned} \quad (8.11)$$

The indices i and j refer to specific detectors and runs, respectively; the index k refers to the k -th observed event in the i -th detector during the j -th run; and the index l loops over background contaminations from surface events (β), neutrons (n), and lead recoils (Pb). $\mu_{i,j}$ refers to the expected rate of WIMP detections, and $f_s^{i,j}(E_k)$ refers to the expected energy distribution of detected WIMPs; both terms are determined from the WIMP mass, cross-section, detector livetime, and detection efficiency for Detector i during Run j . $\nu_{i,j,l}$ refers to the expected background rate for each type of background l in Detector i during Run j , and the corresponding energy distributions $f_l^{i,j}(E_k)$ and expected leakage PDFs $\mathcal{L}_{i,j,l}(\nu_{i,j,l})$ are derived from data-driven models. The minimization of $\ln(\mathcal{L}(m_\chi, \sigma_n))$, with the energy and detector distribution of the three observed events, provides best fit values of $\hat{m}_\chi = 8.6 \text{ GeV}/c^2$ and $\hat{\sigma}_n = 1.9 \times 10^{-41} \text{ cm}^2$. The region of parameter space around this best fit has received considerable attention in recent years due to the potential signals observed by CoGent, DAMA/LIBRA, and CRESST-II. All of these results are in tension with

exclusion limits from the CDMSII germanium data and from the XENON collaboration.

Figure 8-9 presents the results of this likelihood analysis. A goodness-of-fit test to the WIMP+background hypothesis provides a p-value of 68%, while the background-only hypothesis fits the data with a p-value of 4.5%. A profile likelihood ratio test favors the WIMP+background hypothesis over the background-only hypothesis with a p-value of 0.19%. Though this result favors the WIMP interpretation over the best estimates of the known backgrounds, this result does not rise to the level of a discovery.

8.6.2 Spin-dependent constraints

Figure 8-10 presents constraints on spin-dependent WIMP-neutron interactions with the ^{29}Si isotope, which makes up $\sim 4.67\%$ of the isotopes in the detectors. a_p is assumed to be equal to zero when setting this limits. Limits from XENON100 and ZEPLIN-III [192, 193] are presented for comparison. These xenon-based experiment measure the WIMP-neutron cross-section on the isotopes ^{129}Xe and ^{131}Xe and produce better limits at most WIMP masses than the Si analysis presented here. However, the available spin-dependent limits from these experiments do not extend as low in WIMP mass as the results of this Si analysis. Notably, the results of this analysis, even with the observed three events, rules out much of the parameter space corresponding to an interpretation of the DAMA/LIBRA signal as arising from annual modulation and a spin-dependent WIMP-neutron coupling.

Limits on the WIMP-proton cross-section are not presented here, as direct detection experiments that use target nuclei with odd numbers of protons (e.g., Cs, F, Cl, I) have achieved WIMP-proton SD scattering cross-section limits that are orders of

magnitude below the sensitivity of this silicon-based analysis. Interestingly, indirect searches also produce some of the best limits on WIMP-proton SD scattering by searching for neutrinos resulting from WIMP annihilation in the sun, assuming that WIMP capture in the sun, which depends on the WIMP-hydrogen interaction rate, is in equilibrium with WIMP annihilation in the sun. The expected neutrino spectrum is strongly model-dependent on the branching ratios of WIMP annihilation, but the results still lead direct dark matter searches at some masses by a fairly wide margin.




Figure 8-9: Experimental upper limits (90% confidence level) for the WIMP-nucleon spin-independent cross section as a function of WIMP mass. We show the limit obtained from the exposure analyzed in this work alone (*black dots*), and combined with the CDMS II Si data set reported in [181] (*blue solid line*). Also shown are limits from the CDMS II Ge standard [182] and low-threshold [183] analysis (*dark and light dashed red*), EDELWEISS low-threshold (*orange diamonds*) [184], XENON10 S2-only [185] (*light dash-dotted green*), and XENON100 [186] (*dark dash-dotted green*). The filled regions identify possible signal regions associated with data from CoGeNT [187] (*magenta*, 90% C.L., as interpreted by Kelso *et al.* including the effect of a residual surface event contamination described in [188]), DAMA/LIBRA [189, 190] (*yellow*, 99.7% C.L.), and CRESST [191] (*brown*, 95.45% C.L.) experiments. 68% and 90% C.L. contours for a possible signal from these data are shown in blue and cyan, respectively. The asterisk shows the maximum likelihood point at (8.6 GeV/c², 1.9×10^{-41} cm²).




Figure 8-10: Limit on the spin-dependent WIMP-neutron cross-section from this analysis. Results from other direct dark matter searches are included for comparison.

Bibliography

- [1] E. Komatsu, K. M. Smith, J. Dunkley et al. *The Astrophysical Journal Supplement Series*, 192(2):18, 2011.
- [2] R. A. Alpher and R. C. Herman. *Phys. Rev.*, 74:1737–1742, Dec 1948. doi:10.1103/PhysRev.74.1737.
- [3] G. Gamow. *Phys. Rev.*, 74:505–506, Aug 1948. doi:10.1103/PhysRev.74.505.2.
- [4] G. Gamow. *Nature*, 162:680–682, Oct 1948. doi:10.1038/162680a0.
- [5] A. A. Penzias and R. W. Wilson. *ApJ*, 142:419–421, July 1965. doi:10.1086/148307.
- [6] A. McKellar. *Publications of the Dominion Astrophysical Observatory Victoria*, 7:251, 1941.
- [7] T. Shmaonov. *Pribory i Tekhnika Experimenta*, 1957.
- [8] J. F. Denisse, J. Lequeux, and E. Le Roux. *Comptes Rendus*, 244B, 1957.
- [9] D. Larson, J. Dunkley, G. Hinshaw et al. *The Astrophysical Journal Supplement Series*, 192(2):16, 2011.

- [10] N. Jarosik, C. L. Bennett, J. Dunkley et al. *The Astrophysical Journal Supplement Series*, 192(2):14, 2011.
- [11] D. J. Eisenstein, W. Hu, and M. Tegmark. *The Astrophysical Journal Letters*, 504(2):L57, 1998.
- [12] D. J. Eisenstein, I. Zehavi, D. W. Hogg et al. *ApJ*, 633:560–574, November 2005. doi:10.1086/466512.
- [13] S. Perlmutter, G. Aldering, G. Goldhaber et al. *The Astrophysical Journal*, 517(2):565, 1999.
- [14] M. Hamuy, M. M. Phillips, N. B. Suntzeff et al. *The Astrophysical Journal*, 112:2398, December 1996. doi:10.1086/118191.
- [15] R. A. Alpher, H. Bethe, and G. Gamow. *Phys. Rev.*, 73:803–804, Apr 1948. doi:10.1103/PhysRev.73.803.
- [16] D. Kirkman, D. Tytler, N. Suzuki et al. *Astrophysical Journal Supplements*, 149:1–28, November 2003. doi:10.1086/378152.
- [17] F. Zwicky. *Astrophysical Journal*, 86:217, October 1937. doi:10.1086/143864.
- [18] A. Vikhlinin, A. Kravtsov, W. Forman et al. *Astrophysical Journal*, 640:691–709, April 2006. doi:10.1086/500288.
- [19] L. M. J. S. Volders. *Bulletin of the Astronomical Institutes of the Netherlands*, 14:323, September 1959.
- [20] V. C. Rubin and W. K. Ford, Jr. *Astrophysical Journal*, 159:379, February 1970. doi:10.1086/150317.

- [21] Y. Sofue, Y. Tutui, M. Honma et al. *The Astrophysical Journal*, 523:136–146, September 1999. doi:10.1086/307731.
- [22] T. S. van Albada, J. N. Bahcall, K. Begeman et al. *The Astrophysical Journal*, 295:305–313, August 1985. doi:10.1086/163375.
- [23] D. Clowe, M. Bradač, A. H. Gonzalez et al. *The Astrophysical Journal*, 648:L109–L113, September 2006. doi:10.1086/508162.
- [24] R. D. Peccei and H. R. Quinn. *Phys. Rev. Lett.*, 38:1440–1443, Jun 1977. doi:10.1103/PhysRevLett.38.1440.
- [25] L. Roszkowski. *Pramana*, 62:389, February 2004. doi:10.1007/BF02705097.
- [26] J. L. Feng. *Annual Review of Astronomy and Astrophysics*, 48(1):495–545, 2010. doi:10.1146/annurev-astro-082708-101659.
- [27] G. Jungman, M. Kamionkowski, and K. Griest. *Physics Reports*, 267:195–373, March 1996. doi:10.1016/0370-1573(95)00058-5.
- [28] G. Bertone, D. Hooper, and J. Silk. *Physics Reports*, 405:279–390, January 2005. doi:10.1016/j.physrep.2004.08.031.
- [29] W. J. G. de Blok. *Advances in Astronomy*, 2010:789293, 2010. doi:10.1155/2010/789293.
- [30] S. S. McGaugh, W. J. G. de Blok, J. M. Schombert et al. *The Astrophysical Journal*, 659:149–161, April 2007. doi:10.1086/511807.
- [31] S. D. M. White and C. S. Frenk. *The Astrophysical Journal*, 379:52–79, September 1991. doi:10.1086/170483.

- [32] J. F. Navarro, C. S. Frenk, and S. D. M. White. *The Astrophysical Journal*, 462:563, May 1996. doi:10.1086/177173.
- [33] L. Hui. *Phys. Rev. Lett.*, 86:3467–3470, Apr 2001. doi:10.1103/PhysRevLett.86.3467.
- [34] P. Bode, J. P. Ostriker, and N. Turok. *The Astrophysical Journal*, 556:93–107, July 2001. doi:10.1086/321541.
- [35] D. N. Spergel and P. J. Steinhardt. *Phys. Rev. Lett.*, 84:3760–3763, Apr 2000. doi:10.1103/PhysRevLett.84.3760.
- [36] A. Klypin, A. V. Kravtsov, O. Valenzuela et al. *The Astrophysical Journal*, 522:82–92, September 1999. doi:10.1086/307643.
- [37] B. Moore, S. Ghigna, F. Governato et al. *The Astrophysical Journal Letters*, 524(1):L19, 1999.
- [38] G. Kauffmann, S. D. M. White, and B. Guideroni. *MNRAS*, 264:201, 1993.
- [39] B. Willman. *Advances in Astronomy*, 2010, 2010. doi:10.1155/2010/285454.
- [40] J. S. Bullock. *ArXiv e-prints*, September 2010.
- [41] G. D. Martinez, Q. E. Minor, J. Bullock et al. *The Astrophysical Journal*, 738:55, September 2011. doi:10.1088/0004-637X/738/1/55.
- [42] J. Wolf, G. D. Martinez, J. S. Bullock et al. *Monthly Notices of the Royal Astronomical Society*, 406(2):1220–1237, 2010. ISSN 1365-2966. doi:10.1111/j.1365-2966.2010.16753.x.
- [43] G. Aad et al. 2012.

- [44] S. Chatrchyan et al. *JHEP*, 1209:094, 2012. doi:10.1007/JHEP09(2012)094.
- [45] Y. Bai, P. J. Fox, and R. Harnik. *Journal of High Energy Physics*, 12:48, December 2010. doi:10.1007/JHEP12(2010)048.
- [46] J. Goodman, M. Ibe, A. Rajaraman et al. *Physical Review D*, 82(11):116010, December 2010. doi:10.1103/PhysRevD.82.116010.
- [47] E. A. Baltz, M. Battaglia, M. E. Peskin et al. *Physical Review D*, 74(10):103521, November 2006. doi:10.1103/PhysRevD.74.103521.
- [48] P. Konar, K. Kong, K. T. Matchev et al. *New Journal of Physics*, 11(10):105004, October 2009. doi:10.1088/1367-2630/11/10/105004.
- [49] M. Beltrán, D. Hooper, E. W. Kolb et al. *Journal of High Energy Physics*, 9:37, September 2010. doi:10.1007/JHEP09(2010)037.
- [50] J. Carr, G. Lamanna, and J. Lavalle. *Reports on Progress in Physics*, 69(8):2475, 2006.
- [51] J. Lewin and P. Smith. *Astroparticle Physics*, 6(1):87 – 112, 1996. ISSN 0927-6505. doi:10.1016/S0927-6505(96)00047-3.
- [52] Z. Ahmed. *A Dark Matter Search Using the Final CDMSII Dataset and a Novel Detector of Surface Radiocontamination*. Ph.D. thesis, California Institute of Technology, 2012.
- [53] A. Kurylov and M. Kamionkowski. *Phys. Rev. D*, 69:063503, Mar 2004. doi:10.1103/PhysRevD.69.063503.
- [54] J. L. Feng, J. Kumar, D. Marfatia et al. *Physics Letters B*, 703:124–127, September 2011. doi:10.1016/j.physletb.2011.07.083.

- [55] M. T. Ressell, M. B. Aufderheide, S. D. Bloom et al. *Phys. Rev. D*, 48:5519–5535, Dec 1993. doi:10.1103/PhysRevD.48.5519.
- [56] B. A. Dolgoshein, V. N. Lebedenko, and B. U. Rodionov. *Journal of Experimental and Theoretical Physics Letters*, 11:513–516, June 1970.
- [57] M. Schumann. *ArXiv e-prints*, jun 2012.
- [58] R. Bernabei, P. Belli, A. Bussolotti et al. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 592(3):297 – 315, 2008. ISSN 0168-9002. doi:10.1016/j.nima.2008.04.082.
- [59] S. C. Kim, H. Bhang, J. H. Choi et al. *Physical Review Letters*, 108(18):181301, May 2012. doi:10.1103/PhysRevLett.108.181301.
- [60] C. Aalseth et al. 2012.
- [61] I. Phillips, D.G., E. Aguayo, I. Avignone, F.T. et al. *J.Phys.Conf.Ser.*, 381:012044, 2012. doi:10.1088/1742-6596/381/1/012044.
- [62] E. Behnke, J. Behnke, S. J. Brice et al. *Phys. Rev. Lett.*, 106:021303, Jan 2011. doi:10.1103/PhysRevLett.106.021303.
- [63] S. Archambault, F. Aubin, M. Auger et al. *Phys.Lett.*, B682:185–192, 2009. doi:10.1016/j.physletb.2009.11.019.
- [64] M. Felizardo, T. Morlat, A. Fernandes et al. *Phys.Rev.Lett.*, 105:211301, 2010. doi:10.1103/PhysRevLett.105.211301.
- [65] S. Ahlen, J. Battat, T. Caldwell et al. *Phys.Lett.*, B695:124–129, 2011. doi:10.1016/j.physletb.2010.11.041.

- [66] R. Bernabei et al. *Eur.Phys.J.*, C67:39–49, 2010. doi:10.1140/epjc/s10052-010-1303-9.
- [67] C. Aalseth, P. Barbeau, J. Colaresi et al. *Phys.Rev.Lett.*, 107:141301, 2011. doi:10.1103/PhysRevLett.107.141301.
- [68] Z. Ahmed et al. *Phys.Rev.Lett.*, 106:131302, 2011. doi:10.1103/PhysRevLett.106.131302.
- [69] E. Aprile et al. *Phys.Rev.Lett.*, 109:181301, 2012. doi:10.1103/PhysRevLett.109.181301.
- [70] E. Aprile, M. Alfonsi, K. Arisaka et al. *Physical Review Letters*, 109(18):181301, November 2012. doi:10.1103/PhysRevLett.109.181301.
- [71] Bertolini. *Semiconductor Detectors*. North-Holland Publishing Company, 1968.
- [72] J. Lindhard, V. Nielsen, M. Scharff et al. *Matematisk-fysiske Meddelelser*, 33(10):1–42, 1963.
- [73] E. Bonderup. *Penetration of Charged Particles though Matter, Lecture Notes*. Institute of Physics, University of Aarhus, Aarhus, Denmark,, 1979.
- [74] S. Agostinelli, J. Allison, K. Amako et al. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250 – 303, 2003. ISSN 0168-9002. doi:10.1016/S0168-9002(03)01368-8.
- [75] P. Luke. *Appl. Phys. Lett.*, 65:2884, 1994.
- [76] N. W. Ashcroft and D. N. Mermin. *Solid State Physics*. Thomson Learning, 1976. ISBN 0030839939.

- [77] J. P. Wolfe. *Imaging Phonons*. Cambridge University Press, 1998. ISBN 0521620619.
- [78] S. W. Leman. *Development of Phonon-Mediated Transition-Edge-Sensor X-Ray Detectors for use in Astronomy*. Ph.D. thesis, Stanford University, 2006.
- [79] A. Bakalyarov, A. Y. Balysh, S. Belyaev et al. *Phys.Part.Nucl.Lett.*, 2:77–81, 2005.
- [80] S.-I. Tamura. *Physical Review B*, 27:858–866, 1983.
- [81] M. e. a. Msall. *Physical Review B*, 56:9557, 1997.
- [82] S.-I. Tamura. *Physical Review B*, 31:2574, 1985.
- [83] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. *Phys. Rev.*, 108:1175–1204, Dec 1957. doi:10.1103/PhysRev.108.1175.
- [84] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. *Phys. Rev.*, 106:162–164, Apr 1957. doi:10.1103/PhysRev.106.162.
- [85] S. Kaplan, C. Chi, D. Langenber et al. *Physical Review B*, 14:4854, 1976.
- [86] M. Kurakado. *Nuclear Instruments and Methods*, 196:275–277, 1982.
- [87] P. L. Brink. *Non-Equilibrium Superconductivity induced by X-ray Photons*. Ph.D. thesis, Oxford University, 1995.
- [88] S. W. Leman. *ArXiv e-prints*, sep 2011.
- [89] K. D. Irwin and G. C. Hilton. In C. Enss, ed., *Cryogenic Particle Detection*, volume 99 of *Topics in Applied Physics*, page 63149. Springer- Verlag, 2005.

- [90] 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.
- [91] A. Anderson, S. Leman, M. Pyle et al. *Journal of Low Temperature Physics*, 167:135–140, 2012. ISSN 0022-2291. doi:10.1007/s10909-012-0555-1.
- [92] A. J. Anderson, S. W. Leman, M. Pyle et al. *Journal of Low Temperature Physics*, 167:135–140, May 2012. doi:10.1007/s10909-012-0555-1.
- [93] K. D. Irwin, G. C. Hilton, D. A. Wollman et al. *Journal of Applied Physics*, 83(8):3978–3985, 1998. doi:10.1063/1.367153.
- [94] C. Kittel. *American Journal of Physics*, 22:250, 1954.
- [95] C. Kittel. *Introduction to Solid State Physics*. John Wiley & Sons, Inc., New York, 6th edition, 1986.
- [96] B. Cabrera, M. Pyle, R. Moffatt et al. *ArXiv e-prints*, apr 2010.
- [97] J. Singleton. *Band Theory and Electronic Properties of Solids*. Oxford University Press, 2001. ISBN 019850606449.
- [98] B. Cabrera, S. Leman, K. McCarthy et al. CDMS Detector Monte Carlo Documentation. http://titus.stanford.edu/cgi-test/cvsweb.cgi/CDMS_DetectorMC/CDMS_DMC_Manual.pdf, 2012.
- [99] E. Herring, C. and Vogt. *Physical Review*, 101, 1956.
- [100] C. Jacoboni and L. Reggiani. *Rev. Mod. Phys.*, 55:645–705, Jul 1983. doi:10.1103/RevModPhys.55.645.

- [101] W. Shockley. *Journal of Applied Physics*, 9(10):635–636, 1938.
- [102] S. Ramo. *Proceedings of the IRE*, 27(9):584 – 585, sept. 1939. ISSN 0096-8390. doi:10.1109/JRPROC.1939.228757.
- [103] T. Shutt, N. Wang, B. Ellman et al. *Phys. Rev. Lett.*, 69:3531–3534, Dec 1992. doi:10.1103/PhysRevLett.69.3531.
- [104] T. S. et al. *Physical Review Letters*, 69:3531, 1992.
- [105] T. Shutt, J. Emes, E. Haller et al. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 444:340 – 344, 2000.
- [106] B. Serfass. Run 362 UCB ELog. http://cdms.berkeley.edu/cdms_restricted/elog/RunDiary/548, 2008.
- [107] A. Phipps, K. Sundqvist, A. Lam et al. *Journal of Low Temperature Physics*, 167:1112–1118, 2012. ISSN 0022-2291. doi:10.1007/s10909-012-0472-3.
- [108] K. Sundqvist and B. Sadoulet. *Journal of Low Temperature Physics*, 151:443–447, 2008. ISSN 0022-2291. doi:10.1007/s10909-007-9666-5.
- [109] B. Cabrera, M. Pyle, R. Moffatt et al. *ArXiv e-prints*, April 2010.
- [110] V. Aubry-Fortuna and P. Dollfus. *JOURNAL OF APPLIED PHYSICS*, 108(12), DEC 15 2010. ISSN 0021-8979. doi:{10.1063/1.3520656}.
- [111] V. Aubry-Fortuna, A. Broniatowski, and P. Dollfus. *AIP Conf. Proc.*, 1185:639, 2009. doi:10.1063/1.3292422.

- [112] M. Pepin. Global Gamma Monte Carlo c58R Analysis. http://titus.stanford.edu/cdms_restricted/pepin/GlobalGammaNote/index.html, 2013.
- [113] J. Cooley. SuperCDMS Proposal sensitivity table. http://cdms.berkeley.edu/wiki/lib/exe/fetch.php?media=backgrounds:act_doe_v2.pdf, 2012.
- [114] S. W. Leman et al. *Chinese Journal of Physics*, 49:349–358, 2011.
- [115] S.W. Leman et al. *J. Appl. Pys.*, 110:094515, 2011.
- [116] S.W. Leman. *Review of Modern Physics*, submitted 2011.
- [117] A.J. Anderson and S.W. Leman and M. Pyle and E. Figueroa-Feliciano and K. McCarthy and T. Doughty for the SuperCDMS Collaboration. *Journal of Low Temperature Physics*, submitted 2011.
- [118] A. Anderson, S. Leman, T. Doughty et al. *IEEE Transactions on Applied Superconductivity*, submitted 2011.
- [119] M. Pyle and B. Serfass. iZIP: Initial Attempts to Create High Quality Pt estimators without position correction. http://titus.stanford.edu/cdms_restricted/mpyle1/ebooks/IZIP2/YPt_poscor_free/Ypt_pcfree.html, 2010.
- [120] S. Smith. *Development of Transition Edge Sensor Distributed Read-Out Imaging Devices for Applications in X-ray Astronomy*. Ph.D. thesis, University of Leicester, 2006.
- [121] D. S. Akerib, M. S. Armel-Funkhouser, M. J. Attisha et al. *Phys. Rev. D*, 72:052009, Sep 2005. doi:10.1103/PhysRevD.72.052009.

- [122] D. S. Akerib, J. Alvaro-Dean, M. S. Armel-Funkhouser et al. *Phys. Rev. Lett.*, 93:211301, Nov 2004. doi:10.1103/PhysRevLett.93.211301.
- [123] D. S. Akerib, M. J. Attisha, C. N. Bailey et al. *Phys. Rev. Lett.*, 96:011302, Jan 2006. doi:10.1103/PhysRevLett.96.011302.
- [124] D. S. Akerib, M. S. Armel-Funkhouser, M. J. Attisha et al. *Phys. Rev. D*, 73:011102, Jan 2006. doi:10.1103/PhysRevD.73.011102.
- [125] C. I. Collaboration, Z. Ahmed, D. S. Akerib et al. *Science*, 327:1619, mar 2010. doi:10.1126/science.1186112.
- [126] Z. Ahmed, D. S. Akerib, S. Arrenberg et al. *Phys. Rev. Lett.*, 102:011301, Jan 2009. doi:10.1103/PhysRevLett.102.011301.
- [127] J. Sander. *Results from the Cryogenic Dark Matter Search Using a Chi Squared Analysis*. Ph.D. thesis, University of California, Santa Barbara, 2007.
- [128] ROOT - A data analysis framework. <http://root.cern.ch/drupal/>.
- [129] L. Hsu and B. e. a. Serfass. SuperCDMS Software Central. <http://cdms.berkeley.edu/wiki/doku.php?id=software>.
- [130] 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.
- [131] S. Hertel. Improving the Walked Timing Parameters at Small Energy. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/081119_1/webpage1/pulses-50k.png, 2008.

- [132] K. Sundqvist. Run 125-128 Charge Crosstalk & Position Correction. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090225/postprocessv3/R125_Charge_Calib.html, 2009.
- [133] K. McCarthy. Phonon Calibration for R125-128. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090226/phononcal/R58%20phonon%20calibration.html, 2009.
- [134] K. McCarthy. Comparison of individual vs overall calibration of runs 125-128. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090327_2/R58_phonon_calibration.html, 2009.
- [135] K. McCarthy. Phonon Calibration Checks for R125-128. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090310/check_phononcal/R58_phonon_calibration_checks.html, 2009.
- [136] M. Pyle and W. Ogburn. Subpixel Position Systematics: Why box plots are boxy? http://titus.stanford.edu/cdms_restricted/detector_physics/iZIP/ebook/101031/BoxPlts.html, 2010.
- [137] M. Pyle and B. Serfass. Phonon Position Correction. http://cdms.berkeley.edu/wiki/doku.php?id=analysis:r123:phonon_pos_corr#position_and_energy_dependences_of_the_phonon_parameters, 2007.
- [138] M. Pyle. Possible explanation of Energy Dependence of Timing Parameters. http://titus.stanford.edu/cdms_restricted/Soudan/r118r119/ebook/060314/index.html, 2006.

- [139] K. McCarthy. Updates to Si position correction. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/120717_2/20120717_poscorr.html, 2012.
- [140] K. McCarthy. c58 Blinding Cut. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090521_2/blind_cut.html, 2009.
- [141] K. McCarthy. R126 Blinding Cut. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/081104/index.html, 2008.
- [142] K. McCarthy. R127 Blinding Cut. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/081110_2/index.html, 2008.
- [143] K. McCarthy. R128 Blinding Cut. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/081110_3/index.html, 2008.
- [144] C. Bailey. Detector Status Wiki. <http://cdms.berkeley.edu/wiki/doku.php?id=detectors:soudan:home>, 2009.
- [145] J. Manungu. Trigger Rate Analysis for Background Data, Run 125 to 128. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090701/cTrigBurst_c58/cTrigBurst.html, 2009.
- [146] S. e. a. Hertel. c58 Post-Run Cf KS tests. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090610/c58CfKS.html, 2009.
- [147] S. e. a. Hertel. c58 Post-run Ba KS tests. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090521_2/blind_cut.html, 2009.

- [148] S. e. a. Hertel. c58 Post-Run WIMP Search KS tests. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090520/c58BgKS.html, 2009.
- [149] K. McCarthy. Periods of Poor Neutralization in c58 Low Background Data. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/120828_3/20120731_neutralization.html, 2012.
- [150] K. McCarthy. Removal of Helium film events from c58 Si low background data. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/120815/20120806_Hefilms.html, 2012.
- [151] M. Kos. He Film Studies. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090721/heFilmStudies.html, 2009.
- [152] S. Arrenberg. OF power stability checks and definition of the bad resolution cut for the combined analysis of run 125-128. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090623/cBadRes_OFStability/cBadRes_OFStability.html, 2009.
- [153] K. McCarthy. Check of cBadRes cut for c58 Si detectors. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/120821/20120821_BadRes.html, 2012.
- [154] D. Moore. Time intervals with high charge noise in R125-128. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090618/, 2009.
- [155] K. McCarthy. NuMi Beam Coincidence Cut. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090625/cNuMi.html, 2009.

- [156] T. Bruch. Documentation of ErrorMask values in c58 WS data. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090622/errmaskc58.html, 2009.
- [157] J. Hall. cGlitch_c58. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090812/cGlitch_c58.html, 2009.
- [158] J. Zhang. Charge χ^2 Cut and Efficiencies for the Reprocessed Run 125-128 Data. http://www.hep.umn.edu/cdms/cdms_restricted/misc/110524/cChiSq.html, 2011.
- [159] D. Moore. cPstd and cQstd cuts for c58. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090317_2/index.html, 2009.
- [160] D. Moore. Updates to cPstd_c58 and cQstd_c58. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090622_4/, 2009.
- [161] J. Filippini. Charge OF Pathologies. http://titus.stanford.edu/cdms_restricted/Soudan/r118r119/ebook/070512b/crossPileup.shtml, 2007.
- [162] K. McCarthy. Check of cGoodPStartTime for Run 125-128. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090612/start_time.html, 2009.
- [163] K. McCarthy. Si yield bands post position correction. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/120828_2/20120730_yield_bands.html, 2012.
- [164] D. Speller. Adjustments to the Inner Charge Cut for the c58 Re-analysis. http://titus.stanford.edu/cdms_restricted/speller/c58R/072312/Adj_to_cQin_c58R_July2012.html, 2012.

- [165] O. Kamaev. Efficiency correction for cQin and cTiming cut for R123-4 Si detectors. http://titus.stanford.edu/cdms_restricted/Soudan/R123/ebook/100517/cQinTeffic_run34_Si.html, 2010.
- [166] M. Fritts. *Background Characterization and Discrimination in the Final Analysis of the CDMS II Phase of the Cryogenic Dark Matter Search*. Ph.D. thesis, University of Minnesota, 2011.
- [167] S. Fallows and M. Fritts. The Veto Cut for r125-128. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/090622_3/, 2009.
- [168] Z. Ahmed, D. S. Akerib, S. Arrenberg et al. *Phys. Rev. Lett.*, 106:131302, Mar 2011. doi:10.1103/PhysRevLett.106.131302.
- [169] Z. Ahmed, D. S. Akerib, S. Arrenberg et al. *Phys. Rev. D*, 83:112002, Jun 2011. doi:10.1103/PhysRevD.83.112002.
- [170] J. Kiveni. *A Search for WIMP Dark Matter using an Optimized Chi-square Technique on the Final Data from the Cryogenic Dark Matter Search Experiment (CDMS II)*. Ph.D. thesis, Syracuse University, 2012.
- [171] K. McCarthy. Intro to the c58R Si timing cut analysis. http://titus.stanford.edu/cdms_restricted/kevmc/20121017_timingchi2/20121016_timing_intro.html, 2012.
- [172] K. McCarthy. Beta Face/Energy Systematics for the c58R Si analysis. http://titus.stanford.edu/cdms_restricted/kevmc/20121023_betas/20121023_beta_syst.html, 2012.
- [173] R. Agnese et al. *Phys.Rev.Lett.*, 2013.

- [174] K. McCarthy. c58R Si analysis - Systematics checks, pre-unblinding checks. http://titus.stanford.edu/cdms_restricted/Soudan/R125-128/ebook/121219/20121211_timingchi2.html, 2012.
- [175] A. Reisetter. Summary of Cosmogenic NR simulations. http://titus.stanford.edu/cdms_restricted/analysis/091007/cosmoSummary.pdf, 2009.
- [176] A. Reisetter. 5-Tower Cosmogenic Simulation Summary. http://www.hep.umn.edu/cdms/cdms_restricted/mc/080404b/NRsummary.pdf, 2008.
- [177] X. Qiu. *Advanced Analysis and Background Techniques for the Cryogenic Dark Matter Search*. Ph.D. thesis, University of Minnesota, 2009.
- [178] S. Yellin. *Phys. Rev. D*, 66:032005, Aug 2002. doi:10.1103/PhysRevD.66.032005.
- [179] V. I. Dimitrov, J. Engel, and S. Pittel. *Phys. Rev. D*, 51:R291–R295, Jan 1995. doi:10.1103/PhysRevD.51.R291.
- [180] P. J. Fox, J. Liu, and N. Weiner. *Phys.Rev.*, D83:103514, 2011. doi:10.1103/PhysRevD.83.103514.
- [181] J. P. Filippini. Ph.D. thesis, University of California, Berkeley, 2008.
- [182] Z. Ahmed et al. *Science*, 327(5973):1619–1621, 2010.
- [183] Z. Ahmed et al. *Physical Review Letters*, 106(13):131302, March 2011.
- [184] E. Armengaud et al. *Phys.Rev.*, D86:051701, 2012. doi:10.1103/PhysRevD.86.051701.

- [185] J. Angle et al. *Physical Review Letters*, 107(5):051301, July 2011. doi:10.1103/PhysRevLett.107.051301.
- [186] J. Angle et al. *Phys. Rev. Lett.*, 100:021303, 2008.
- [187] C. Kelso, D. Hooper, and M. R. Buckley. *Phys. Rev. D*, 85(4):043515, February 2012. doi:10.1103/PhysRevD.85.043515.
- [188] C. E. Aalseth et al. *Physical Review Letters*, 107(14):141301, September 2011. doi:10.1103/PhysRevLett.107.141301.
- [189] R. Bernabei et al. *Eur. Phys. J.*, C67:39–49, 2010. doi:10.1140/epjc/s10052-010-1303-9.
- [190] C. Savage, G. Gelmini, P. Gondolo et al. *JCAP*, 0904:010, 2009. doi:10.1088/1475-7516/2009/04/010.
- [191] G. Angloher et al. *European Physical Journal C*, 72:1971, April 2012. doi:10.1140/epjc/s10052-012-1971-8.
- [192] M. Garny, A. Ibarra, M. Pato et al. 2012.
- [193] D. Y. Akimov, H. Araujo, E. Barnes et al. *Phys.Lett.*, B709:14–20, 2012. doi:10.1016/j.physletb.2012.01.064.