

SEARCH FOR LOW-MASS DARK MATTER WITH SUPERCDMS
SOUDAN AND STUDY OF SHORTED ELECTRIC FIELD
CONFIGURATIONS IN CDMS DETECTORS

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Kristiana E. Schneck

October 2015

Preface

The area of dark matter is one of the most interesting and exciting topics in physics today. Existing at the intersection of particle physics and astrophysics, the existence of a new dark matter particle can be used to explain many astrophysical and cosmological observations, as well as to reconcile outstanding issues in the standard model of particle physics. Experiments such as SuperCDMS are built to detect dark matter in the lab by looking for low-energy nuclear recoils produced by collisions between dark matter particles and atoms in terrestrial detectors. SuperCDMS Soudan is particularly well-suited to follow up on possible hints of low-mass dark matter seen by other recent experiments because of its low thresholds and excellent background discrimination. Analyzing SuperCDMS Soudan data to look for low-mass dark matter comes with particular challenges because of the low signal-to-noise very near threshold. However, with a detailed background model developed by scaling high-energy events down into the low-energy signal region, SuperCDMS Soudan produced world-leading limits on the existence of low-mass dark matter.

In addition, a few SuperCDMS Soudan detectors experienced cold hardware problems that can affect the data collected. Of particular interest is one detector considered for the low-mass WIMP search that has one of its charge electrodes shorted to chassis ground. Three events were observed in this detector upon unblinding the SuperCDMS Soudan low-energy data, even though <1 event was expected based on pre-unblinding calculations. However, the data collected by the shorted detector may have been compromised since an electrode shorted to ground will modify the electric field in the detector. The SuperCDMS Detector Monte Carlo (DMC) provides an excellent way to model the effects of the modified electric field, so a new model of the

expected backgrounds in the low-mass WIMP search is developed using the DMC to try to explain how the short may have affected the data collected.

This thesis is organized as follows: Chapter 1 gives a broad introduction to dark matter, discussing the astrophysical and cosmological evidence for its existence, listing several possible particle physics candidates, and outlining several experimental strategies to look for dark matter. Chapter 2 is an overview of CDMS detector technology and the experimental setup at the Soudan Underground Laboratory, with a focus on how data coming out of Soudan is analyzed. Chapter 3 presents results from a search for low-mass dark matter at SuperCDMS Soudan and discusses the interpretation of the results. Chapter 4 contains follow-up work that uses the CDMS Detector Monte Carlo (DMC) to understand the possible systematics associated with a detector that had one of its charge electrodes shorted to ground. This chapter represents the first time the DMC has been used to inform ongoing CDMS analysis. Chapter 5 takes a brief detour into the world of effective field theory (EFT), examining the consequences of an expanded set of possible WIMP-nucleon interactions in the EFT framework. Finally, Chapter 6 wraps up the material of the previous chapters and discusses how the research presented in this thesis can be applied as CDMS moves toward SuperCDMS SNOLAB.

Acknowledgements

One of the wonderful things about working on CDMS here is that I was able to be involved with two groups: the SLAC group, led by Richard Partridge, and the Stanford group, led by Blas Cabrera. The wide range of expertise in the two groups meant that I was able to work on a wide range of projects, from scanning newly-fabricated Soudan and SNOLAB detectors, to testing high electron-mobility transistors for the SNOLAB cold electronics, to the detector simulations and data analysis that make up the bulk of my thesis. Rich and Blas have always asked challenging questions about my work and provided helpful suggestions, and I am grateful for their guidance.

I would also like to thank my other committee members, Risa Wechsler and Dan Akerib, as well as my defense committee chair, Steve Sano. Risa and Dan asked really interesting questions during my defense and provided helpful feedback on my thesis. Steve was my choral director for five years in the Stanford Chamber Chorale, and I'm grateful that he was willing to take a step outside of his area of expertise to serve as my chair.

This work would not have been possible without the huge amount of expertise in the CDMS collaboration. In particular, I would like to thank the members of the low-threshold analysis team: Adam Anderson, Julien Billard, Ray Bunker, Lauren Hsu, Elias Lopez Asamar, and Brad Welliver; the detector monte carlo group: Peter Redl, Rob Agnese, Katrina Coletti, Jorge Morales, and Daniel Brandt; and the effective field theory group: David Cerdeño, Vuk Mandic, and Hannah Rogers. I also had many helpful discussions with Liam Fitzpatrick about his work on an effective field theory for direct detection, which was enormously helpful for the work documented in Chapter 5.

A big thank you also to everyone who shared onsite shift duties with me at Soudan: Rito Basu Thakur, Jon Wilson, Katrina Coletti, Brad Welliver, Rob Agnese, Jorge Morales, Adam Anderson, Hannah Rogers, D'Ann Barker, Prisca Cushman, Silvia Scorza, Peter Redl, Rich Partridge, Dennis Wright, Daniel Brandt, and others I may have forgotten, as well as the full-time Soudan crew. We had many adventures, from debugging FEDUP checks and commissioning the first CDMSlite run, to the time the DAQ crashed at 4:30 pm the Friday before Christmas. Working at Soudan was always an interesting experience!

Last, but certainly not least, I would like to thank my family and friends who have supported me throughout this endeavor. Thank you for listening, for taking me out for ice cream, and for traveling with me across California, the United States, and the world. I could not have done this without you.

Contents

Preface	iv
Acknowledgements	vi
1 Introduction	1
1.1 Astrophysical and cosmological evidence for dark matter	2
1.1.1 Gravity: galaxy clusters and rotation curves	2
1.1.2 Evidence from gravitational lensing	3
1.1.3 Cosmological evidence for dark matter	5
1.2 Dark matter candidates	7
1.3 Experimental signatures of dark matter	11
1.3.1 Indirect detection	12
1.3.2 Collider searches	14
1.3.3 Direct detection	14
1.4 Direct detection signals	18
1.4.1 Standard halo model and spin-independent interaction	18
1.4.2 Alternate astrophysical models	22
1.4.3 Additional signatures of dark matter	24
2 SuperCDMS at Soudan	27
2.1 The Soudan Underground Laboratory	28
2.2 SuperCDMS at Soudan	29
2.2.1 CDMS detector physics	31
2.2.2 iZIP detectors	33

2.2.3	Readout Electronics	35
2.2.4	Data processing	37
2.2.5	Calibration	41
2.3	Recent CDMS results	45
2.3.1	CDMS II Silicon	46
2.3.2	CDMS II low-energy likelihood	46
2.3.3	CDMSlite—a low ionization threshold experiment	48
3	Search for Low-Mass WIMPS with SuperCDMS	50
3.1	Event selection criteria and background model development	50
3.1.1	Selecting high-quality physics events	50
3.1.2	Background model development	53
3.1.3	Final event selection methods	54
3.2	Livetime estimates and final analysis efficiency	58
3.3	Expected sensitivity for final event selection methods	60
3.3.1	Energy scale for exclusion limits	60
3.3.2	Sensitivity calculation	62
3.4	Candidate events and agreement with expected backgrounds	63
3.5	Exclusion limits	64
3.5.1	Limits on spin-independent and spin-dependent cross sections	64
3.5.2	Limits assuming an alternate halo model	67
4	Simulating detectors with shorted channels	69
4.1	Detector Monte Carlo-based background model	70
4.1.1	DMC input event types and energies	73
4.1.2	Formatting DMC output for processing	75
4.1.3	Pulse simulation method and data processing	79
4.2	Validation	80
4.2.1	Rescaling phonon partitions	80
4.2.2	The bulldozer algorithm—modeling local saturation in the TES simulation	82
4.2.3	Assessing agreement between Soudan and DMC barium data .	84

4.3	Comparing shorted detector T5Z3 to fully-functional detectors	88
4.3.1	Qualitative effects of the short	88
4.3.2	Comparing the DMC to the SuperCDMS Soudan background model	91
4.3.3	Hypothesis for source of T5Z3 events in the low-threshold analysis	94
4.4	Other insights from the DMC	97
5	Effective field theory scattering	100
5.1	Bias created by standard dark matter assumptions	105
5.2	Exclusion limits on a set of EFT operators	108
5.3	Principal component analysis of EFT interference	111
5.4	G2 event rate comparisons	115
5.5	Prospects for the future	121
6	Conclusions	123
6.1	Low-mass dark matter searches at SNOLAB	123
6.2	Moving forward with the detector monte carlo	125
6.3	Effective field theory beyond SuperCDMS Soudan	128
A	Low-frequency noise discrimination	132
A.1	Low-frequency noise in SuperCDMS Soudan data	132
A.2	Selection criteria for low-frequency noise events	133
A.3	Low frequency noise template development	134
A.4	Determining selection criteria	135
A.5	Estimating the efficiency of the low-frequency noise selection criteria .	138
B	EFT exclusion limits	141
C	HEMT testing in the Stanford ^3He fridge	144
	Bibliography	148

List of Tables

3.1	Detector masses and livetimes for this analysis, yielding a total exposure of 577 kg-days.	59
3.2	Eleven candidate events observed in this analysis.	64
C.1	Summary of results from several Stanford ^3He tests	147

List of Figures

1.1	Rotation curve (rotational velocity vs. distance to center) for the Andromeda galaxy [3]. The solid line represents a least-squares fit to a fourth-order polynomial that is required to remain nearly flat at a radius of 120 arc-minutes.	3
1.2	Measured rotation curve of the Milky Way (circles and triangles), with models of the contribution to the rotation curve from the galactic bulge and disk (luminous matter) and halo (dark matter) [8].	4
1.3	Mass density vs. radius for cluster 0024+1654 [10]. The thick line shows the total mass as measured using strong gravitational lensing, while the thin line shows the mass of the galaxy-only component. Galaxies represent only 17% of the mass of the cluster inside $107 h^{-1}$ kpc.	5
1.4	Observations of the Bullet Cluster, a merging cluster of galaxies [14]. Left: Magellan images of the cluster, with weak lensing mass contours overlaid in green. Right: Chandra observations of the X-ray emission from hot gas, showing clear evidence of shocks from the merger. . . .	6
1.5	Left: All-sky map of CMB temperature fluctuations produced by Planck [20]. Right: Power spectrum of CMB temperature as measured by Planck, with $\pm 1\sigma$ uncertainties. The best-fit Λ CDM theoretical spectrum is plotted in red, and the bottom panel shows the residuals with respect to this model.	7
1.6	Content of the universe, as measured by the Planck cosmic microwave background satellite [20].	8

1.7	Dark matter parameter space, with ranges of mass and cross section for possible candidates indicated [33]. Red, pink, and blue boxes represent hot, warm, and cold dark matter candidates respectively.	9
1.8	WIMP density as a function of inverse temperature ($\sim$ time) in the early universe for different choices of the velocity-averaged annihilation cross section $\langle\sigma v\rangle$ [37]. The black curve shows the WIMP density if the expansion of the universe was slow enough to maintain thermal equilibrium.	10
1.9	95% C.L. upper limits on the dark matter annihilation cross section for annihilation to $b\bar{b}$ (left) and $\tau^+\tau^-$ (right) from Fermi observations of DES dwarf-spheroidal galaxy candidates [56]. Combined with earlier results, these limits exclude the thermal relic cross section (thin dashed grey line) below masses of ~ 100 GeV/ c^2	13
1.10	Limits on spin-independent dark matter interactions from monophoton searches by ATLAS (left, figure: [66]) and CMS (right, figure: [67]). . .	15
1.11	WIMP discovery limit (“neutrino floor”) compared with recent direct detection results [71, 73].	18
1.12	Nuclear form factor (Eqn. 1.6) for several common direct detection targets.	20
1.13	Differential recoil spectra for a 10 GeV/ c^2 (left) and 100 GeV/ c^2 WIMP (right) with spin-independent cross section of 10^{-44} cm ² , assuming a local dark matter density ρ_0 of 0.3 GeV/ c^2 /cm ³ , a galactic escape velocity of 544 km/s, $v_E = 232$ km/s, and $v_0 = 220$ km/s.	23
1.14	Expected dark matter recoil rate for standard halo model and alternate model for an 8.6 GeV/ c^2 WIMP with a cross section of 1.9×10^{-41} cm ² (the CDMS II Si best-fit result).	24
1.15	Event rates for minimum (December), maximum (June), and mean earth velocity for a 10 GeV (left) and 100 GeV (right) mass WIMP. The annual modulation is more significant for the lower mass, where experiments are only sensitive to the tail of the distribution.	25

1.16	Integrated count rate as a function of time for the CDMS II best-fit Si result in several energy bins.	26
2.1	The sign that greets visitors to the level of the physics halls, complete with one of the lab's ubiquitous bats.	27
2.2	Muon flux versus depth for various underground laboratories.	28
2.3	SuperCDMS shielding and muon veto.	30
2.4	Left: SuperCDMS Soudan tower layout. Right: channel layout for a SuperCDMS Soudan detector.	31
2.5	Cartoon of the physics processes that occur in a CDMS detector. An interaction produces ionization and prompt phonons. As the charges drift across the crystal, additional Luke phonons are produced.	32
2.6	Yield vs. recoil energy in a CDMS detector [96]. Blue points show bulk electron recoils, while red points show events from a source that produces surface lead recoils and other lead decay chain events. The nuclear recoil band is shown in green.	33
2.7	Left: Electric field geometry near the surface of an iZIP detector. Right: Charge collection on side 2 vs. side 1 of a CDMS, determined by the iZIP electric field geometry. A source of surface events was placed near side 1, producing the population of red points with Side 2 Charge Collection consistent with zero. Bulk recoil events, as determined by the dashed-line boundaries, are shown in blue.	34
2.8	QET on the face of a SuperCDMS detector. The dark region is the detector substrate, the light region is the aluminum fins, and the thin red lines show the tungsten TES.	36
2.9	Left: Sample of phonon pulses for T1Z1 (colored traces) and phonon pulse template (black) constructed by averaging. Right: Position-dependent residual traces for T1Z1, constructed by subtracting the phonon template from individual phonon pulses. The covariance of these residuals is used in the non-stationary optimal filter. [Figure: Yohan Ricci]	41

2.10	Charge calibration procedure for inner-channel events. [figure: Adam Anderson]	42
2.11	Charge calibration procedure for outer-channel and shared events. [figure: Adam Anderson]	43
2.12	Example nuclear recoil energy scale for detector T2Z2, measured using ^{252}Cf neutron calibration data. The cyan curves show MCMC samplings of the charge model parameterization. [figure: Julien Billiard] .	45
2.13	Left: Signal region and three candidate events from the CDMS II silicon result [47]. Right: Allowed region for the CDMS II silicon result (blue enclosed region), with other contemporary results including CDMS II Ge (dashed maroon), CRESST (enclosed magenta), and XENON100 (dash-dot green).	47
2.14	Limit set by a likelihood analysis of CDMS II low-energy germanium data (black) [101]. This limit is ~ 5 times stronger than the corresponding limit (magenta) with no background subtraction. The deviation between the limit and the expected sensitivity (green bands) is also seen in the multiples data, indicating that it is due to an unaccounted-for systematic, not a WIMP component.	48
2.15	Left: Spectrum from CDMSlite Run 1, clearly displaying the K-shell (~ 10.4 keV) and L-shell (~ 1.3 keV) activation lines. Right: Exclusion limits from CDMSlite Run 1. [figure from [103]]	49
3.1	Fiducial volume selection for detector T2Z1. The surface event criteria is shown on the left, and the radial criteria is shown on the right. Red events pass the selection criteria, while black events fail the criteria. [Figure: Todd Doughty]	52
3.2	Nuclear recoil bands for detector T1Z1, defined using neutron calibration data. [Figure: Julien Billard]	53
3.3	Background model (lead = yellow, gammas = red, 1.3 keV line = blue) [Figure: Adam Anderson]	55

3.4	Serially-tuned event selection cuts. Notched nuclear recoil band (left), phonon radial cut for a $10 \text{ GeV}/c^2$ WIMP (middle) [figure: Adam Anderson], and phonon symmetry cut for a $10 \text{ GeV}/c^2$ WIMP (right) [figure: Brad Welliver]	56
3.5	Simultaneously-tuned event selection cuts on normalized charge energy (left), phonon radial partition (middle), and phonon z-partition (right) as a function of total phonon energy. [figure: Adam Anderson]	57
3.6	Distribution of BDT scores for events in background model. Black points show SuperCDMS data [104]. The corresponding uncertainty on the residuals is shown in tan in the lower subplot.	58
3.7	Analysis efficiency for 4 levels of cuts in this analysis [104].	60
3.8	Final efficiency of all selection criteria for each of the seven detectors included in this analysis.	61
3.9	Expected sensitivity of three analysis methods (from left to right: individually-tuned cuts, simultaneously-tuned cuts, and boosted decision tree), with 68% and 95% statistical uncertainties shown by the blue bands. CDMS-II Si closed contour in green, CDMSlite limit in yellow, LUX limit in magenta.	62
3.10	Left: Candidate events in this analysis. Events were consistent with expected background, with the exception of the three high-energy events in detector T5Z3 (yellow stars). Right: Expected backgrounds for each detector computed using the pulse simulation background model, with observed candidate events shown in black. [figure: Adam Anderson]	63
3.11	Limits from this analysis (black) with 68% and 95% CL systematic uncertainties shown in gray. The pre-unblinding expected sensitivity in the absence of a signal is shown as 68% (dark green) and 95% (light green) C.L. bands.	65
3.12	Limits from this analysis (black) with 95% CL uncertainties shown in gray for WIMP-neutron (left) and WIMP-proton (right) spin-dependent scattering.	66

3.13	Comparison of WIMP-proton and WIMP-neutron form factors from Dimitrov et al [119] and Klos et al [120] (left), and the resulting limits on WIMP-neutron spin-dependent scattering (right).	67
3.14	Left: Spin-independent limits, comparing the standard Maxwellian velocity distribution to the alternate halo model discussed in Chapter 5. Right: Differential event rates for CDMS II Si best-fit result for the 3 halo models used in left-hand plot.	68
4.1	Transport of electrons (blue) and holes (red) in fully-functional CDMS detector (left) and detector with QOS1 short (right). The bias voltage applied to each channel is indicated in black.	70
4.2	Electric field lines (red) and equipotentials (blue) for fully-functional electric field (left) and T5Z3 (shorted QOS1) field	71
4.3	Total phonon pulse fall times by Soudan detector. The separation between wide-electrode (blue) and narrow-electrode (green) is clear, and needs to be accounted for in the DMC.	72
4.4	Comparison between Soudan phonon template (orange-red) and DMC phonon template (blue) for T1Z1.	73
4.5	Lead decay chain energies and conversion probabilities.	75
4.6	Uncalibrated FET energy versus input recoil energy for side 1 (left) and side 2 (right). Blue points show all simulations run while green points show data after removing events which we <i>a priori</i> expect to have poor charge collection.	76
4.7	Uncalibrated FET energy versus input recoil energy with calibration fits for side 1 (left) and side 2 (right). Data is binned into 1 keV bins and each bin is fit to a gaussian, resulting in the red points. The error bars show the standard deviation for each bin The means of all the gaussians are fit to a line, producing the black result.	77
4.8	DC subtraction and relative calibration for simulated TES pulses. . .	78

4.9	Histogram of total phonon energy for T1Z1 L-shell activation line data. The cyan histogram shows the total phonon energy estimated using the Soudan template, while the blue histogram shows the total phonon energy estimated using the DMC template. The Soudan template energy is biased low because it does not properly fit the peakiness of the DMC pulses.	80
4.10	Raw DMC data (blue) and rescaled DMC data (red) for detector T2Z1, compared to the expected Soudan distribution (black) for phonon r-partition (left) and phonon z-partition (right).	81
4.11	Schematic of the “bulldozer” algorithm for modeling local saturation. Any phonons collected in a TES that push the histogram of phonon arrival times above the saturation threshold (grey line) in a given time bin are bulldozed to the next time bin. [figure: Peter Redl]	83
4.12	Total phonon energy templates generated on standard DMC data (magenta) and DMC data following the bulldozer algorithm (red). The bulldozer algorithm produces pulses that more closely match the Soudan phonon template (black). [figure: Peter Redl]	84
4.13	Raw DMC data (blue) and rescaled DMC data (red) for detector T1Z1, compared to the expected Soudan distribution (black) for phonon r-partition for standard DMC data with no local saturation (left) and DMC data after application of the bulldozer algorithm (right).	85
4.14	Likelihoods comparing barium data from all Soudan low-threshold detectors to T2Z1 DMC barium data.	86
4.15	Likelihoods comparing barium data from all Soudan low-threshold detectors as well as unscaled and scaled DMC data for T2Z1 to T2Z1 Soudan barium data.	87
4.16	Results for a BDT trained on Soudan data and tested on both Soudan (red) and DMC data (blue) for detector T2Z1. The total phonon energy spectrum for the californium data was re-weighted to match the expected spectrum of a $5 \text{ GeV}/c^2$ WIMP (left) or $15 \text{ GeV}/c^2$ WIMP (right)	88

4.17	Truth yield (mean inner charge channel energy divided by truth recoil energy) versus position for fully functional detector T2Z1 (left) and shorted detector T5Z3 (right) for simulated barium calibration data between 2 and 13.1 keV total phonon energy.	89
4.18	Charge radial partition for side 2 (holes) versus position (top) and charge z-partition versus position (bottom) for fully functional detector T2Z1 (left) and shorted detector T5Z3 (right) for simulated barium calibration data between 2 and 13.1 keV total phonon energy. The radial partition for side 1 is ill-defined for T5Z3 due to the outer charge channel short.	90
4.19	Fraction of simulated barium events that pass the charge fiducial volume and nuclear recoil band cuts (the “preselection” cuts discussed in Chapter 3) for fully-functional T2Z1 (left) and shorted T5Z3 (right).	91
4.20	Phonon radial partition (rescaled and normalized) versus truth position for barium data from fully-functional detector T2Z1 (left) and shorted detector T5Z3 (right).	92
4.21	Histograms of phonon radial partition (left), phonon z-partition (center), and charge energy (right) for low-energy gammas for the DMC (unscaled partitions: blue, rescaled partitions: cyan) and data-driven (magenta) models for detector T1Z1.	93
4.22	Charge energy versus total phonon energy for ^{210}Pb surface events in detector T2Z1. There is a qualitative disagreement between DMC and data-driven simulation in the yield for this event type.	94
4.23	Histograms of boosted decision tree scores trained against a $15\text{ GeV}/c^2$ WIMP for ^{206}Pb sidewall events for detector T1Z1 (left) and T5Z3 (right). The data-driven background model is shown in magenta and the DMC is shown in blue.	95

4.24	Histograms of boosted decision tree scores trained against a $5 \text{ GeV}/c^2$ WIMP (left) or $7 \text{ GeV}/c^2$ WIMP (right) for low-background gamma events for detector T5Z2. The data-driven background model is shown in magenta and the DMC is shown in blue. The disagreement is large between DMC and data-driven models for this detector and several others, indicating that further validation work may be needed.	96
4.25	Mean BDT score versus truth z-position for recycled DMC events for T1Z1 (left) and T5Z3 (right) for ^{206}Pb sidewall events. T5Z3 has more high-score outliers near the top surface near the shorted region; an upward fluctuation in the noise may push one of these events above the BDT threshold and into the signal region.	97
4.26	BDT score standard deviation versus truth z-position for recycled DMC events for T1Z1 (left) and T5Z3 (right) for ^{206}Pb sidewall events. There is a possible increase in the standard deviation of the scores very near the top surface of the detector ($z \sim 2.5 \text{ cm}$); a larger standard deviation might lead to statistical fluctuations pushing a small number of events above the boosted decision tree threshold.	98
4.27	Comparison between DMC 1 keV line events (blue) and data-driven 1 keV line events (magenta, cyan) with two different weighting schemes for detector T2Z2. The updated weighting scheme (magenta), which includes surface events, matches the DMC better and can partially account for an excess of very low energy events seen in a likelihood analysis of the data from the low-mass WIMP search of Chapter 3. .	99
5.1	Differential event rate for a 50 GeV WIMP interacting in germanium as a function of energy for several EFT interactions. Although the interactions have vastly different strengths for a given set of coefficients, here the coefficients are chosen so that the integrated event rate is equal for the four operators shown.	104

5.2	Co-added energy spectrum from 100 simulated experiments (blue histogram) assuming the dark matter interaction proceeds according to the isoscalar $\mathcal{O}_3$ operator for a $10 \text{ GeV}/c^2$ (left) and a $300 \text{ GeV}/c^2$ WIMP (right). The detection efficiency is assumed to be independent of energy. The smooth cyan, magenta, and black curves show the expected spectrum for the standard spin-independent rate for several WIMP masses, while the dashed dark blue curve shows the $\mathcal{O}_3$ spectrum from which the simulated experiments were sampled.	106
5.3	Distribution of 90% confidence level upper limits calculated using the optimum interval method for the simulated experiments discussed above and shown in Fig. 5.2, sampled from the event rate for isoscalar $\mathcal{O}_3$. Shaded blue bands show the 68% and 95% confidence level uncertainty on the distribution. The zero-background Poisson limit is shown in magenta.	107
5.4	Upper limits on the dimensionless isoscalar coefficients c_3^0 (left) and c_8^0 (right) as a function of WIMP mass for SuperCDMS Soudan (light blue) [104], CDMS II Ge reanalysis (dark blue) [133], and CDMS II Si (red) [134], and estimated limits for LUX (black) [85], for the Maxwellian halo (solid) and an alternate halo model (dashed). Limits on the remaining isoscalar operators can be found in Appendix B.	109
5.5	Polar limits on $\mathcal{O}_1$ isospin for SuperCDMS Soudan (blue) [104], LUX [85] (black), and CDMS II Si (red) [134] at a WIMP mass of $6 \text{ GeV}/c^2$	111
5.6	Constructive interference eigenvectors for 2D $\mathcal{O}_4$ isospin interference. Proton-dominated interactions occur along the $x = y$ diagonal, while neutron-dominated interactions occur along the $x = -y$ diagonal.	113
5.7	Eigenvectors for 4D $\mathcal{O}_4/\mathcal{O}_5$ interference in germanium for a 10 GeV WIMP (left), and relative event rates for parameters consistent with expectations for SuperCDMS SNOLAB Ge iZIP.	114
5.8	Eigenvectors for 4D $\mathcal{O}_1/\mathcal{O}_3$ interference in germanium for a 10 GeV WIMP (left), and relative event rates for parameters consistent with expectations for SuperCDMS SNOLAB Ge iZIP.	114

5.9	Relative event rates for LZ (black), SuperCDMS SNOLAB Ge iZIP (blue), and SuperCDMS SNOLAB Si (red), normalized to 1 observed event in SuperCDMS Ge ($3 \text{ GeV}/c^2$) or LZ ($10, 300 \text{ GeV}/c^2$). From left to right are shown the rates for a 3, 10, and $300 \text{ GeV}/c^2$ WIMP, assuming isoscalar interactions and the standard Maxwellian halo model. The $3 \text{ GeV}/c^2$ case also shows the rates from SuperCDMS SNOLAB Ge high-voltage (light blue), which has similar parameters to SuperCDMS Si high-voltage, but a target mass of 6 kg. (The rate for LZ at $3 \text{ GeV}/c^2$ is essentially zero for kinematic reasons.) The top row shows cumulative event rates, while the bottom row shows events per time per target mass. True interaction strengths may differ from this calculation since the interaction may proceed via a linear combination of operators.	116
5.10	Event rate suppression relative to $\mathcal{O}_1$ scattering in LZ (black), SuperCDMS SNOLAB Ge iZIP (blue), and SuperCDMS SNOLAB Si (red) for interference in germanium, with interference ranging from constructive (left) to maximally destructive (right), as determined by the magnitude of the corresponding eigenvalue. The rate for SuperCDMS SNOLAB Ge high-voltage (light blue) is shown for the $3 \text{ GeV}/c^2$ case where LZ sees no events above threshold. The seven operator-operator interference cases are shown, as well as pure isoscalar $\mathcal{O}_1$, which is used as a reference point.	118
6.1	Current limits from direct detection experiments (solid lines) and expected limits from the planned G2 experiments (dashed lines). SuperCDMS SNOLAB will have world-leading sensitivity to low-mass dark matter.	125
6.2	Projected background-free $\mathcal{O}_1$ isoscalar-isovector limits for a $6 \text{ GeV}/c^2$ WIMP in the G2 experiments (solid), compared with current limits (dashed).	130

A.1	Examples of low-frequency noise seen in two different detectors in low-energy Soudan data. The left-hand event (from detector T3Z1) is reconstructed to have an energy of 7 keV, and the right-hand event (from detector T4Z1) has an energy of 4 keV.	132
A.2	Representative low-frequency noise trace for detector T4Z2.	134
A.3	Low frequency noise template and phonon pulse template for T4Z2.	135
A.4	$\Delta\chi^2$ versus ptOF using the low frequency noise template for detector T4Z1.	136
A.5	Example event with positive $\Delta\chi^2$ and an energy of 2.3 keV, shown in frequency domain (left) and time domain (right).	137
A.6	Example event with negative $\Delta\chi^2$ and an energy of 5.2 keV, shown in frequency domain (left) and time domain (right).	137
A.7	Histogram of $\Delta\chi^2$ for random triggers. The true distribution of randoms is shown in blue, while the expected distribution of randoms with no low-frequency noise component is shown in magenta. Flat cut thresholds were set such that 95% of the distribution with no low-frequency noise (magenta) passed the threshold.	138
A.8	Histogram of $\Delta\chi^2$ for californium calibration data. The fits to the 2σ and 3σ bands of good pulses are shown.	139
A.9	Histogram of $\Delta\chi^2$ for unblind low-background data.	140
A.10	Efficiency of low-frequency noise selection criteria for detector T4Z2. [Figure: Adam Anderson]	140
B.1	Upper limits on the dimensionless isoscalar coefficients for all effective field theory operators as a function of WIMP mass for SuperCDMS Soudan (light blue) [104], CDMSII Ge reanalysis (dark blue) [133], and CDMSII Si (red) [134], and estimated limits for LUX (black) [85], for the Maxwellian halo (solid) and an alternate halo model (dashed).	143
C.1	Block diagram of fridge setup for HEMT IV measurement.	145

C.2	HEMT board as installed in Stanford ^3He fridge. The HEMTs are located on the underside of the top (green) board, and the bottom (yellow) board acts as a guard for the HEMTs and wire bonds. . . .	146
C.3	Drain current vs drain-source voltage for two of the three high-electron mobility transistors tested on October 16, 2012, measured at 4K. . . .	146
C.4	Fit to equation C.2 for the HEMTs 1 and 2 at $V_{DS} = 175$ mV. . . .	147

Chapter 1

Introduction

Physicists have long endeavored to understand the workings of the universe at the most fundamental level. The triumph of particle physics in the 20th century was the standard model, which describes the interactions of fundamental particles such as electrons and quarks. However, around the same time, astrophysical and cosmological evidence began to suggest that the standard model did not provide the full picture of all interactions and types of particles. A new type of particle was proposed to explain these observations. Known as dark matter, this particle (or family of particles) does not interact with standard model particles through electromagnetism or the strong force, and, as such, it is not observable using most standard techniques. This chapter discusses the evidence for dark matter, from the earliest calculations of galactic dynamics, to state-of-the-art observations of the cosmic microwave background. It also discusses several possible models for dark matter as well as current means of detecting this new particle.

1.1 Astrophysical and cosmological evidence for dark matter

1.1.1 Gravity: galaxy clusters and rotation curves

The earliest evidence for dark matter comes from observations of galactic dynamics. In the 1930s, Fritz Zwicky studied the movement of galaxies within the Coma Cluster. He measured the velocities of the galaxies using their Doppler shifts and determined the total mass of the cluster with the virial theorem [1, 2]. After comparing this result to the mass estimated from the number of galaxies and the luminosity of the cluster, his calculations suggested that there was significantly more mass than could be accounted for with visual observations. He called this extra mass “dark matter.”

Individual galaxies also show similar behavior: more mass than is visually observed is needed to account for the large rotational velocities at high radii. Vera Rubin and collaborators determined the rotational velocity as a function of distance to the center of the galaxy using the Doppler shift of spectra from numerous HII regions, which contain partially-ionized hydrogen gas, in the Andromeda galaxy [3]. The measured velocities from this study are shown in Fig. 1.1.

The rotational velocity was found to be nearly flat at high radius, indicating that the mass of the galaxy increases linearly with radius [4, 5]. Further studies found similar results across a range of galaxies [6]. The flat portion of the rotation curves extends past the light-emitting regions of the galaxy. This differs from the expectation that, if the mass of the galaxy tracks with the light, the rotation curve should fall off as $1/\sqrt{r}$ at large radii. These observations provide further evidence for the existence of a significant massive and non-luminous component in galaxies [7]. More recent models of the contributions to the rotation curve from the various galactic components support this hypothesis, as shown in Fig. 1.2 [8]. For the Milky Way, a significant dark matter halo component (dash-dot black curve) is required to fit the observations.




Figure 1.1: Rotation curve (rotational velocity vs. distance to center) for the Andromeda galaxy [3]. The solid line represents a least-squares fit to a fourth-order polynomial that is required to remain nearly flat at a radius of 120 arc-minutes.

1.1.2 Evidence from gravitational lensing

General relativity predicts that mass distributions located between an observer and a distant light source can “bend” the light as it propagates, producing an effect known as gravitational lensing. This phenomena can be used to determine the magnitude and distribution of dark matter in two ways: strong lensing, which creates visible distortions of the lensed source, and weak lensing, which considers small distortions of the lensed sources that are detected by analyzing large numbers of sources in aggregate.

Strong lensing of background sources produces multiple images of a single source, magnification of the lensed source, and distortion effects known as Einstein rings. Large luminous arcs (partial Einstein rings) were observed in galaxy clusters as early as the 1970s; however, these structures were not connected to lensing phenomena until the late 1980s [9]. Further studies of clusters, including Abell 1689 and 0024+1654, demonstrate that strong lensing can be used to accurately determine the mass of clusters [10, 11]. In addition, comparisons of the mass of the galaxies in the cluster to the total mass of the cluster, as in Fig. 1.3, show that galaxies comprise only a




Figure 1.2: Measured rotation curve of the Milky Way (circles and triangles), with models of the contribution to the rotation curve from the galactic bulge and disk (luminous matter) and halo (dark matter) [8].

fraction of the total mass of galaxy clusters.

Weak gravitational lensing, on the other hand, looks for small distortions (of the order of a few percent) in the shapes of large populations of galaxies. Early studies detected systematic alignments between background galaxies and attributed to lensing by a massive foreground galaxy [12]. Modern weak lensing work often focuses on cosmic shear, looking for distortion of galaxies in one preferential direction as photons are lensed by the mass distribution along the line of sight. Since cosmic shear measurements often focus on wide surveys of the sky, they can be used to map out the large-scale structure of the mass in the universe [13].

One famous example of lensing evidence for dark matter is from the somewhat audaciously-titled paper “A Direct Empirical Proof of the Existence of Dark Matter” [14]. This study examined weak lensing observations of a merging galaxy cluster known as the Bullet Cluster and compared the results to X-ray emission from the hot gas. The distributions of mass and hot gas in the Bullet Cluster are shown in




Figure 1.3: Mass density vs. radius for cluster 0024+1654 [10]. The thick line shows the total mass as measured using strong gravitational lensing, while the thin line shows the mass of the galaxy-only component. Galaxies represent only 17% of the mass of the cluster inside $107 h^{-1} \text{ kpc}$.

Fig. 1.4. The weak lensing maps approximately trace the distribution of galaxies, even though the majority of the baryonic matter in the cluster is found in the hot gas. The separation between the bulk of the mass as measured by weak lensing and the dominant baryonic component cannot be easily explained by theories which propose modifications to gravity. Such a result can be caused by the existence of a non-luminous particle that comprises the majority of the mass of the cluster.

1.1.3 Cosmological evidence for dark matter

The structure and evolution of the universe also provides evidence for the existence of dark matter. The Λ CDM cosmological model [15, 16] (named for cold dark matter and Λ , or dark energy), supported by evidence from high-redshift supernovae [17, 18], the cosmic microwave background [19, 20], big bang nucleosynthesis [21, 22], baryon acoustic oscillations [23], and the large-scale structure of the universe [24], describes




Figure 1.4: Observations of the Bullet Cluster, a merging cluster of galaxies [14]. Left: Magellan images of the cluster, with weak lensing mass contours overlaid in green. Right: Chandra observations of the X-ray emission from hot gas, showing clear evidence of shocks from the merger.

the evolution of the universe from the Big Bang through the current epoch.

The cosmic microwave background (CMB), discovered in the mid-1960s by Arno Penzias and Robert Wilson [25], has provided a wealth of information about the evolution of the early universe. Produced when the early universe cooled enough to allow the ionized plasma to recombine into hydrogen atoms (an age of approximately 400,000 years, or a temperature of 3000 K), this radiation is isotropic and uniform in temperature. As the universe cooled and expanded, the CMB photons were redshifted, so today the CMB is a near-perfect black body spectrum with a temperature of ~ 2.73 K.

Subsequent observations from satellites such as COBE and WMAP discovered small anisotropies in the CMB at the level of one part in 10^5 [26, 27]. These anisotropies are the result of quantum fluctuations in the early universe, imprinted in the density of the universe as it underwent exponential inflation a tiny fraction of a second after the big bang [28, 29]. Today, this appears as small fluctuations in the temperature map of the CMB, as shown in Fig. 1.5, left [20].

From the temperature fluctuations map, CMB experiments can construct the power spectrum as a function of the multipole moment ℓ , which is inversely proportional to the angular scale. Figure 1.5, right, shows the CMB power spectrum as

measured by Planck. The theoretical power spectrum predicted by the best-fit parameters for the Λ CDM model is shown in red, and is in excellent agreement with the data.




Figure 1.5: Left: All-sky map of CMB temperature fluctuations produced by Planck [20]. Right: Power spectrum of CMB temperature as measured by Planck, with $\pm 1\sigma$ uncertainties. The best-fit Λ CDM theoretical spectrum is plotted in red, and the bottom panel shows the residuals with respect to this model.

The peaks in the power spectrum above ℓ of ~ 100 are caused by oscillations produced by two competing forces in the photon-baryon plasma in the early universe: radiation pressure, which tends to smooth out anisotropies, and gravitational attraction, which produces clumping. At the time of recombination, these oscillations were frozen in and appear as the series of peaks seen in the power spectrum [30, 31]. The locations and amplitudes of these peaks contain information about the curvature of the universe and the density of baryons and dark matter. This yields values for the parameters of the Λ CDM model, including the fractional density of dark energy, dark matter, and ordinary matter (baryons) shown in Fig. 1.6. These results show that the total amount of matter in the universe is ~ 5 times larger than the baryonic matter, further emphasizing that the majority of the matter in the universe is dark.

1.2 Dark matter candidates

The astrophysical and cosmological evidence for non-luminous matter suggests the existence of a new type of particle that makes up the majority of the mass density of




Figure 1.6: Content of the universe, as measured by the Planck cosmic microwave background satellite [20].

the universe and is not found in the standard model. Candidates for such a particle are often well-motivated by outstanding problems in the standard model [32]. Several of the most commonly-discussed particles are shown in Fig. 1.7 [33]. The red boxes, which depict hot or warm dark matter candidates, are ruled out by large scale structure constraints: neutrinos or other hot dark matter candidates tend to wash out large-scale structure since the particles will freely stream away instead of clumping together [34, 35].

One particularly attractive dark matter candidate is the Weakly Interacting Massive Particle, or WIMP [36]. Such a particle could be produced thermally in the early universe and freeze out to the correct relic density (this is known as the “WIMP miracle”). Figure 1.8 depicts the freeze-out process [37]. At early times (high temperatures), WIMPs are in thermal equilibrium: the rate of annihilation and production is approximately equal. As the universe expands and cools, WIMP production is suppressed and the density begins to drop off. Eventually, the universe expands to a




Figure 1.7: Dark matter parameter space, with ranges of mass and cross section for possible candidates indicated [33]. Red, pink, and blue boxes represent hot, warm, and cold dark matter candidates respectively.

sufficiently large size that particles cannot find partners to annihilate with, so annihilation also ceases. This results in a relic density that is frozen in. Thermally-averaged annihilation cross sections at the weak scale ($\langle\sigma v\rangle \sim 10^{-26} \text{ cm}^3/\text{s}$) can produce the observed relic density.

WIMPs can be found in numerous extensions to the standard model. Supersymmetry (SUSY) provides a particularly attractive way to generate stable, massive, neutral particles [38]. To solve the gauge hierarchy problem, which asks why the scale of the Higgs field is $\sim 100 \text{ GeV}/c^2$ instead of at the Planck scale, SUSY posits the existence of a new partner particle with opposite spin statistics for each particle in the standard model. Without such particles, the scale of the Higgs field must be fine-tuned to this low value; however, the addition of superparticles naturally brings




Figure 1.8: WIMP density as a function of inverse temperature ($\sim$ time) in the early universe for different choices of the velocity-averaged annihilation cross section $\langle\sigma v\rangle$ [37]. The black curve shows the WIMP density if the expansion of the universe was slow enough to maintain thermal equilibrium.

its value down to $\sim 100 \text{ GeV}/c^2$.

In SUSY models that contain a new conserved quantity known as R-parity, the lightest supersymmetric particle (LSP) provides a possible dark matter candidate because it is stable [39, 40]. In the minimal supersymmetric standard model (MSSM), the LSP is usually a linear combination of the superpartners of the photon, Z^0 , and Higgs bosons called the neutralino. The simplest and most highly-constrained version of supersymmetry, the CMSSM, is in tension with experimental data including null results from direct detection experiments and the discovery of a $126 \text{ GeV}/c^2$ Higgs [41, 42]. However, in the more general pMSSM, many dark matter candidates remain that can be accessed by experiments that are currently operating or will come online in the near future [43].

Recent results from DAMA [44], CoGeNT [45], CRESST-II [46], and CDMS II Si

[47] are consistent with dark matter particles with masses of approximately $10 \text{ GeV}/c^2$, igniting interest in models that can produce low-mass dark matter particles. SUSY models in the Next-to-MSSM (NMSSM) can produce particles with masses of $10 \text{ GeV}/c^2$ or less that have the proper relic density while satisfying all constraints from the LHC [48, 49]. However, there are additional models that can also produce dark matter with masses of a few GeV/c^2 in the appropriate amounts. One such model is asymmetric dark matter, which claims that dark matter in the early universe is related to the observed matter/antimatter asymmetry [50, 51]. This effect naturally produces low-mass dark matter particles, since in this model, $m_\chi \sim \frac{\rho_{DM}}{\rho_b} m_p \sim 5 \text{ GeV}/c^2$.

Another theoretically well-motivated candidate for dark matter is the axion, motivated by a solution to the strong CP problem [52]. Such a particle is expected to have a mass that is a fraction of an eV, and can be detected by conversion to photons in the presence of a magnetic field. The Axion Dark Matter Experiment (ADMX) is designed to detect axions using RF cavities that can scan over a range of frequencies [53]. The conversion between axions and photons is expected to be resonant, so an axion would show up as a small bump in the power as a function of frequency. Upgrades to the experiment are expected to probe a wide range of the possible axion dark matter parameter space [54]. In addition, experiments designed to look for WIMPs are also sensitive to other types of axions that are not dark matter candidates: CDMS II produced limits on the axion-photon coupling for Solar axions [55], and SuperCDMS SNOLAB, with its larger exposure, is expected to improve upon those limits. Though these other particles are also viable dark matter candidates, this thesis is primarily concerned with Weakly Interacting Massive Particles (WIMPs).

1.3 Experimental signatures of dark matter

The evidence detailed in section 1.1 demonstrates the existence of dark matter via gravitational effects on astrophysical scales. Numerous experiments are currently running that are designed to detect dark matter through non-gravitational interactions with standard model particles. These experiments are interested in one of three types

of signals: indirect detection experiments such as Fermi and AMS look for the products of dark matter self-annihilation, collider experiments at the LHC look for dark matter produced in proton-proton collisions, and direct detection experiments such as CDMS and LUX look for ambient dark matter particles scattering off of atoms in terrestrial detectors.

1.3.1 Indirect detection

Indirect detection experiments look for the signature of dark matter particles annihilating to produce standard model particles. Satellites such as Fermi look for gamma rays produced directly by annihilation or in a shower of secondary particles. Many such searches examine data from dwarf spheroidal satellite galaxies (dSphs) in the Milky Way, which are excellent places to look for annihilation products due to their proximity to Earth, high dark matter densities, and lack of other gamma-ray producing processes.

One recent result examined the gamma ray emission from eight Milky Way dwarf spheroidal galaxy candidates recently discovered by DES [56]. This result found no significant excess of gamma rays in the dSphs considered, and sets strong limits on the velocity-averaged annihilation cross section. These limits are shown in Figure 1.9 for annihilation to $b\bar{b}$ (left) and $\tau^+\tau^-$ (right). The result from the eight DES candidate dSphs excludes the thermal relic cross section below masses of $\sim 20 \text{ GeV}/c^2$ for annihilation to $b\bar{b}$ and $\tau^+\tau^-$, and combined with results from other dSphs, excludes the thermal relic cross section up to $\sim 100 \text{ GeV}/c^2$ [57].

In addition, the cosmic ray experiment AMS is sensitive to dark matter annihilating to electrons and positrons. AMS and other cosmic ray experiments have observed an excess of positrons above energies of $10 \text{ GeV}/c^2$ [58]. One possible explanation for the rise in the positron fraction is the annihilation or decay of TeV-scale dark matter particles. Such particles could annihilate to hadronic final states or leptons; positrons would be produced by the decay of π^+ (hadronic final states), τ^+ , or μ^+ , or by direct annihilation to e^+/e^- .

By modeling the expected astrophysical sources of e^+/e^- , as well as any potential




Figure 1.9: 95% C.L. upper limits on the dark matter annihilation cross section for annihilation to $b\bar{b}$ (left) and $\tau^+\tau^-$ (right) from Fermi observations of DES dwarf-spheroidal galaxy candidates [56]. Combined with earlier results, these limits exclude the thermal relic cross section (thin dashed grey line) below masses of $\sim 100 \text{ GeV}/c^2$.

dark matter annihilation signal, physicists can look for possible annihilation signals in the AMS positron data. Limits derived from this data place strong limits on the existence of dark matter with masses below $\sim 300 \text{ GeV}/c^2$ and exclude the thermal relic cross section by one to a few orders of magnitude at low mass, depending on the annihilation final state [59]. These limits, as well as the limits from Fermi discussed above, constrain thermal relic dark matter that annihilates to $b\bar{b}$ or leptons, but they do not limit alternate dark matter models such as asymmetric dark matter.

Other indirect searches have examined Fermi data from the galactic center, another astrophysical region with high dark matter density and therefore high probability of dark matter self-annihilation. These searches have found evidence for a possible excess of gamma rays at low energies that may arise from dark matter particles with masses of ten to a few tens of GeV/c^2 [60, 61, 62, 63, 64, 65]. However, the backgrounds for the galactic center are more complicated than those of the dwarf spheroidal galaxies, so such signals should be followed up on by analyses of many dSph galaxies.

1.3.2 Collider searches

Another method of searching for dark matter involves looking for evidence of its production at colliders. Given enough energy, two quarks colliding at the LHC could produce a pair of dark matter particles. Dark matter particles alone are expected to produce no signal in detectors, so another particle or jet is required to tag such an event. These searches are known as mono-X searches, since they look for events with a single photon, jet, lepton, or other particle in coincidence with “missing” energy—energy necessary to satisfy conservation of energy and momentum, but which does not produce a signal in any of an experiment’s trackers or calorimeters.

Since dark matter interacts so weakly, a dark matter particle that is produced in a particle collision will show up as missing energy and momentum in a detector. These experiments look for dark matter by comparing the distribution of missing energy to the expected backgrounds, which include interactions that produce neutrinos, or W-boson decays where the lepton escapes the detector without producing a signal. An excess above the estimated backgrounds could be interpreted as a dark matter signal. From the observed number of events, limits can be calculated on the interaction strength between dark matter and quarks.

Figure 1.10 shows the limits on spin-independent dark matter interactions from monophoton searches by ATLAS (left) and CMS (right) [67, 66], compared to recent direct detection results. Mono-lepton searches yield results comparable within a few orders of magnitude [68, 69]. These limits are more constraining than direct detection limits at low masses, since direct detection is limited by thresholds for very low mass particles. However, collider limits can be model-dependent, since they require assumptions about the particle physics to translate a production rate in colliders to a scattering rate in direct detection experiments.

1.3.3 Direct detection

Direct detection experiments look for the signatures of the rare interactions between dark matter and the atoms in terrestrial detectors. A typical dark matter particle is expected to produce a nuclear recoil with an energy of a few to a few tens of keV [70].




Figure 1.10: Limits on spin-independent dark matter interactions from monophoton searches by ATLAS (left, figure: [66]) and CMS (right, figure: [67]).

Because many types of WIMP-standard model interactions can produce a nuclear recoil, direct detection provides a less model-dependent way of searching for dark matter. The goal of direct detection experiments is to measure the energy deposited by interactions in a detector and reduce the rate of non-dark matter (background) events to near zero. When the expected background rates are near-zero, any observed events in the detector can be interpreted as possible dark matter candidate events.

The observed number of candidate events and their associated energies can be used to infer or constrain the dark matter scattering rate. Because this rate is dependent on the dark matter mass and cross section, constraints on the scattering rate can be translated to constraints on the dark matter mass and cross section. Section 1.4 contains a more detailed discussion of the differential scattering rate and other characteristics of the expected signal of dark matter interactions in direct detection experiments.

Backgrounds for direct searches

The challenge in direct detection experiments is to distinguish between the nuclear recoil signature of dark matter and any other background events. Typical sources of background include Compton-scattering photons, radioactive contamination of the experimental setup, and cosmic-ray-induced particle showers. Experiments are designed to maximize the background-rejection capabilities using a combination of detector properties and external shielding.

Dark matter direct detection experiments are typically located deep underground to shield against the high levels of cosmic radiation present at the surface, and often feature additional shielding such as lead bricks or a water tank to further protect the experiment. Neutrons produced by cosmic ray showers or radioactive processes are particularly insidious because they also produce a nuclear recoil signal. As a result, experiments typically employ additional shielding composed of materials such as polyethylene, which moderate the energy of neutrons.

Experiments employ many other techniques to reduce background with the intent to reduce the background levels to $\ll 1$ event. A common strategy is to use two independent signals, such as ionization and scintillation or ionization and heat, to distinguish between electron recoils and nuclear recoils. Liquid noble experiments such as XENON and LUX use the ratio of ionization to scintillation to select nuclear recoils, while semiconductor solid-state experiments such as CDMS use the ratio of ionization to heat (phonons).

The location of an event within the detector is also a good indicator of the type of interaction. Scatters that occur near the outer surfaces of a detector are more likely to result from radioactive contamination on the surfaces of the detectors. In addition, the energy collected from events near the surface may be reduced, biasing the ratios used to discriminate between electron recoils and nuclear recoils. Determining the position of every event allows experiments to define a “fiducial volume,” or the region in the bulk of the detector where energy collection is good and background radioactive decay events are rare.

Chapter 2 discusses in more detail how SuperCDMS uses shielding and data analysis techniques to reduce the background rates to near-zero levels.

The neutrino floor

Another source of background that will become important for future direct detection experiments is coherent neutrino scattering [71]. Neutrinos have many of the same properties as dark matter particles and are abundantly produced in the Sun and by supernovae. Because neutrinos, like dark matter, are weakly-interacting particles, it is essentially impossible to shield against them. In addition, neutrinos are predicted to scatter coherently off nuclei, just as a hypothetical dark matter particle is expected to do.

As a result, coherent neutrino scattering produces an irreducible background for direct dark matter searches. Given a large enough detector volume and integration time, a direct detection will see events from coherent neutrino scattering that cannot be distinguished from dark matter on an event-by-event basis. The flux and energy spectra for neutrino interactions is well-understood, so the rate of neutrino coherent scattering in direct detection experiments can be straightforwardly computed [72]. The rate of neutrino interactions can be compared to the expected dark matter interaction rate (see Section 1.4 below) to find the approximate dark matter cross sections below which it becomes difficult to improve the sensitivity of the experiment. This is known as the “neutrino floor.”

Figure 1.11 shows the neutrino floor as a function of WIMP mass, compared to recent direct detection results. The thick orange line shows the mass and cross section at which the rate of neutrino and dark matter interactions is approximately equal. An experiment that is sensitive to dark matter with properties along the line will see ~ 1 neutrino event. The next generation (G2) dark matter experiments are likely to observe ^8B coherent neutrino scattering, and the G2 experiments will come within approximately an order of magnitude of the diffuse supernova background (DSNB) and atmospheric neutrino floor above $> 1\text{GeV}/c^2$.




Figure 1.11: WIMP discovery limit (“neutrino floor”) compared with recent direct detection results [71, 73].

1.4 Direct detection signals

The ultimate goal of dark matter direct detection experiments is to use the observed number of events and their energies to infer an event rate and energy spectrum. From this spectrum, the dark matter parameters of interest (mass and scattering cross section) can be extracted.

1.4.1 Standard halo model and spin-independent interaction

The expected event rate for a dark matter particle with a given mass and cross section interacting with a terrestrial detector can be calculated with a few basic assumptions [74]. The differential particle density for dark matter is

$$dn = n_0 f(\vec{v}, \vec{v}_E) d^3v \quad (1.1)$$

which is normalized so $n_0 = \int_0^{v_{esc}} dn$. Here, n_0 is the mean dark matter number density ($= \rho_0/m_\chi$), $\vec{v}$ is the velocity of the dark matter particle incident upon the target nucleus, $\vec{v}_E$ is the velocity of Earth relative to the dark matter distribution, and v_{esc} is the local galactic escape velocity.

Typically experiments assume the simplest Maxwellian velocity distribution

$$f(\vec{v}, \vec{v}_E) = \frac{1}{k} \exp \left[-\frac{(\vec{v} + \vec{v}_E)^2}{v_0^2} \right], \quad (1.2)$$

where the constant k ensures that the distribution integrates to unity.

Using the differential number density, the event rate per unit mass is

$$dR = \frac{N_A}{A} \sigma v dn \quad (1.3)$$

where N_A is Avogadro's number, A is the atomic mass of the target, and σ is the WIMP-nucleus cross section.

The differential event rate (number of nuclear recoils per kg of detector per day per keV of recoil energy) is [75]

$$\frac{dR}{dE_R} = \frac{N_A \rho_0}{A k m_\chi} \int \frac{d\sigma}{dE_R} v f(\vec{v}, \vec{v}_E) d^3v \quad (1.4)$$

For the coherent spin-independent interaction, the differential cross section $d\sigma/dE_R$ can be simplified by considering quantities arising from two different physical effects: σ_0 , the point-like WIMP-nucleus cross section, and $F(q)$, an energy-dependent nuclear form factor for spin-independent interactions.

$$\frac{d\sigma}{dE_R} = \frac{\sigma_0}{E_R^{max}} F^2(q) \quad (1.5)$$

Here, the momentum transfer is $q = \sqrt{2m_N E_R}$, and m_N is the mass of the target nucleus. It is common to choose the Helm parameterization of the form factor, which

is related to the a Fourier transform of a hard sphere, with a skin depth:

$$F(q) = 3 \frac{j_1(qr_0)}{qr_0} \exp \left[-\frac{1}{2} s^2 q^2 \right] \quad (1.6)$$

where r_0 is the nuclear radius, s is the skin depth and $j_1(qr_0)$ is the spherical Bessel function of index 1. Additional form factors, which relate to nuclear properties such as spin and angular momentum, are discussed in more detail in Chapter 5. A plot of the form factor versus nuclear recoil energy for several common direct detection targets is shown in Fig 1.12.




Figure 1.12: Nuclear form factor (Eqn. 1.6) for several common direct detection targets.

From kinematic considerations, the nuclear recoil energy of a WIMP-nucleus collision is

$$E_R = \frac{\mu^2}{m_N} v^2 (1 - \cos \theta)$$

where $\mu = m_\chi m_N / (m_\chi + m_N)$ is the WIMP-nucleus reduced mass and θ is the scattering angle in the center of mass frame. From this expression, the maximum

recoil energy is $E_R^{max} = 2(\mu^2/m_N)v^2$. The differential scattering rate then becomes

$$\frac{dR}{dE_R} = \frac{N_A \rho_0 \sigma_0 m_N}{2 A k m_\chi \mu^2} F^2(q) \int \frac{f(\vec{v}, \vec{v}_E)}{v} d^3v \quad (1.7)$$

The point-like WIMP-nucleus cross section σ_0 is related to the spin-independent WIMP-nucleon cross section σ_{SI} (the quantity traditionally quoted by direct detection experiments) by

$$\sigma_0 = A^2 \frac{\mu^2}{\mu_{nucleon}^2} \sigma_{SI} \quad (1.8)$$

where $\mu_{nucleon}$ is the WIMP-nucleon reduced mass. Note that this expression assumes the WIMP interacts identically with neutrons and protons (that is, $f_n = f_p$). If this assumption is violated, the expression becomes

$$\sigma_0 = \left(Z + (A - Z) \left[\frac{f_n}{f_p} \right] \right)^2 \frac{\mu^2}{\mu_{nucleon}^2} \sigma_{SI} \quad (1.9)$$

By tuning the ratio f_n/f_p , it is possible to significantly decrease the signal in certain target elements, while leaving the signal in other targets relatively unchanged. In particular, it has been noted that choosing $f_n/f_p = -0.7$ can partially reconcile the tension between the CDMS II Si closed contour and the LUX and SuperCDMS exclusion limits [76].

For the simplest Maxwellian velocity distribution, the integral over velocity can be evaluated for $v_{min} = q/(2\mu) = \sqrt{m_N E_R / (2\mu^2)}$ and v_{max} constrained by $|\vec{v} + \vec{v}_E| \leq v_{esc}$

$$\int \frac{f(\vec{v}, \vec{v}_E)}{v} d^3v = 2\pi \int_{-1}^{+1} d(\cos \theta) \int_{v_{min}}^{v_{max}} v e^{-\frac{|v - v_E|^2}{v_0^2}} dv. \quad (1.10)$$

As discussed in [77], the bound $|\vec{v} + \vec{v}_E| \leq v_{esc}$ leads to the maximum WIMP velocity as a function of scattering angle in the galactic rest frame:

$$v_{max}(\theta_{gal}) = \sqrt{v_{esc}^2 - v_E^2(1 - \cos \theta_{gal})} - v_E \cos \theta_{gal} \quad (1.11)$$

The integral in Eqn. 1.10 must be evaluated for three separate cases:

$$\begin{aligned} 0 &\leq v_{min} \leq v_{esc} - v_E \\ v_{esc} - v_E &\leq v_{min} \leq v_{esc} + v_E \\ v_{esc} + v_E &\leq v_{min} \leq \infty \end{aligned}$$

For the first case ($0 \leq v_{min} \leq v_{esc} - v_E$), the integral evaluates to

$$\frac{\sqrt{\pi}}{4} \frac{v_0}{v_{earth}} \left[\operatorname{erf} \left(\frac{v_{min} + v_{earth}}{v_0} \right) - \operatorname{erf} \left(\frac{v_{min} - v_{earth}}{v_0} \right) \right] - \exp \left(-\frac{v_{esc}^2}{v_0^2} \right).$$

For $v_{esc} - v_E \leq v_{min} \leq v_{esc} + v_E$, the integral evaluates to

$$\frac{\sqrt{\pi}}{4} \frac{v_0}{v_{earth}} \left[\operatorname{erf} \left(\frac{v_{esc}}{v_0} \right) - \operatorname{erf} \left(\frac{v_{min} - v_{earth}}{v_0} \right) \right] - \frac{v_{esc} + v_E - v_{min}}{2v_E} \exp \left(-\frac{v_{esc}^2}{v_0^2} \right),$$

and it evaluates to zero for $v_{esc} + v_E \leq v_{min} \leq \infty$.

A plot of the differential recoil spectrum versus recoil energy for several common direct detection targets and assuming standard astrophysical parameters is shown in Fig 1.13 for a 10 and 100 GeV/ c^2 WIMP, evaluated at parameters typical of our galaxy ($v_0 = 220$ km/s, $v_{earth} = 232$ km/s and $v_{esc} = 544$ km/s) [78]. For the 100 GeV/ c^2 case, the effects of the A^2 enhancement of the rate are clear: the rate is larger at energies of tens of keV for the heavier elements, germanium and xenon. For the 10 GeV/ c^2 case, kinematic considerations come into play: the maximum recoil energy for a given WIMP mass and velocity becomes inversely proportional to the nucleus mass for $m_\chi \ll m_N$. As a result, the tail of the recoil spectrum extends to higher recoil energies for the lighter elements, argon and silicon.

1.4.2 Alternate astrophysical models

As shown in [79, 80], N-body simulations are not always well described by the Maxwellian velocity distribution. Consequently, alternate halo models have been proposed. Several papers [81, 82] have also pointed out that it may be possible to reconcile the tension between the CDMS-II Si result and XENON100's exclusion limit




Figure 1.13: Differential recoil spectra for a 10 GeV/ c^2 (left) and 100 GeV/ c^2 WIMP (right) with spin-independent cross section of 10^{-44} cm², assuming a local dark matter density ρ_0 of 0.3 GeV/ c^2 /cm³, a galactic escape velocity of 544 km/s, $v_E = 232$ km/s, and $v_0 = 220$ km/s.

by choosing a WIMP velocity distribution other than the standard Maxwellian distribution. The proposed function, which is shown to fit the velocity distributions of dark matter haloes in N-body simulations, is given in Eqn. 1.12.

$$f(v) = \exp \left[-\frac{v}{v_0} \right] (v_{esc}^2 - v^2)^p \quad (1.12)$$

This function goes smoothly to zero as the velocity goes to the escape velocity of the galaxy, unlike the Maxwellian distribution, which requires an artificial cutoff at $v = v_{esc}$. This has the effect of decreasing the rate in the tail of the distribution, as shown in Fig. 1.14. This figure uses $p = 2.7$ and $v_0/v_{esc} = 0.6875$ that correspond to results from a simulation of a Milky Way-like galaxy that includes baryons [83]. For this choice of v_0/v_{esc} and p , this function falls off faster than the standard Maxwellian distribution. This difference can significantly affect the expected dark matter event rate, especially for low-mass WIMPs for which experiments are only sensitive to the high-velocity tail of the distribution. The values obtained from the Eris simulations for the parameters of this alternate halo model are in the region of parameter space that [82] demonstrates can possibly reconcile the tension between CDMS II Si [47] and XENON100 [84]. However, this effect cannot also account for the tension with LUX [85] because of that experiment's lower energy threshold.




Figure 1.14: Expected dark matter recoil rate for standard halo model and alternate model for an $8.6 \text{ GeV}/c^2$ WIMP with a cross section of $1.9 \times 10^{-41} \text{ cm}^2$ (the CDMS II Si best-fit result).

1.4.3 Additional signatures of dark matter

In addition to the energy spectrum, dark matter has several additional signatures that experiments can try to detect. Collaborations such as CoGeNT [45], DAMA [44], and CDMS II [86] look for the annual modulation of the dark matter signal, caused by the yearly change in velocity of the earth with respect to the dark matter halo. The velocity of the earth as a function of time can be approximated

$$v_E \simeq 232 + 15 \cos \left[2\pi \frac{t - 152.5}{365} \right] \text{ km/s} \quad (1.13)$$

where t is the time in days since January 1. As the velocity of the earth varies with time, so does the integral over the dark matter velocity distribution in Eqn. 1.7, producing slight variations in the differential recoil spectrum. The effect of the annual modulation of the earth's velocity on the resulting spectrum is shown in Fig 1.15 for WIMPs of two different masses scattering in germanium. Around the month of June, larger velocity of the earth has the effect of increasing the relative velocity between

WIMPs and the terrestrial detector, pushing larger numbers of particles above the minimum velocity needed to produce a recoil of a given energy. The inverse is true for months near December, when the velocity of the earth moving through the dark matter halo is at its smallest.




Figure 1.15: Event rates for minimum (December), maximum (June), and mean earth velocity for a 10 GeV (left) and 100 GeV (right) mass WIMP. The annual modulation is more significant for the lower mass, where experiments are only sensitive to the tail of the distribution.

For the CDMS II Si best-fit result, this leads to a small modulation in the event rate in germanium over the course of a year, as shown in Fig 1.16. The total modulation rate above an assumed 2 keV threshold is shown in dark blue, and corresponds to a modulation amplitude of ~ 0.15 counts/kg/day. When separated into 1-keV bins similar to what was done for the CDMS II annual modulation analysis, the modulation rate becomes less than 0.05 counts/keV/day. A signal of this magnitude is difficult to detect with current experiments, but it may become clearer with the larger exposures of the G2 experiments.

Dark matter is also expected to exhibit a directional signature due to the movement of the solar system through the dark matter halo. Detecting this signal requires measuring the energy of an interaction as well as reconstructing its three-dimensional track [87]. Collaborations such as DRIFT [88, 89], DMTPC [90], and MIMAC [91] use time-projection chambers to reconstruct the initial direction of a recoil. Directional signals of dark matter also provide a way to discriminate against coherent neutrino




Figure 1.16: Integrated count rate as a function of time for the CDMS II best-fit Si result in several energy bins.

scattering [92], since dark matter appears to come from the constellation Cygnus, while the highest flux of coherently-scattering neutrinos comes from the sun. However, current directional detection technologies are difficult to scale to the ton-years of exposure needed to reach the neutrino floor.

Chapter 2

SuperCDMS at the Soudan Underground Laboratory

SuperCDMS Soudan, an upgrade of the Cryogenic Dark Matter Search (CDMS II), is operating 15 germanium detectors in the Soudan Underground Laboratory to search for rare interactions between dark matter and normal matter. SuperCDMS Soudan uses the same shielding, electronics, and cryogenic infrastructure as CDMS II [93].




Figure 2.1: The sign that greets visitors to the level of the physics halls, complete with one of the lab’s ubiquitous bats.

2.1 The Soudan Underground Laboratory

The Soudan Underground Laboratory is situated in northern Minnesota, in a former iron mine in the community of Soudan. The mine's large rock overburden and stable infrastructure provided an excellent location for experiments that require very low background rates, beginning with the Soudan proton decay experiments in the 1980s and continuing to MINOS, CoGeNT, and CDMS today. The current physics halls are located on level 27, nearly half a mile below the surface. Visitors to the lab are reminded of the depth when stepping off the hoist cage by the sign in Fig. 2.1.

Soudan is one of a number of underground physics facilities operating worldwide. These labs, used for experiments looking for rare events such as dark matter and neutrino interactions, are often built in mines or under large mountains to reduce the rate of cosmic ray interactions that could mimic the desired signals. The large rock overburdens of these labs significantly reduces the flux of cosmic ray muons, as shown in Fig. 2.2. For the Soudan Underground Lab, at a depth of 780 meters (or 2090 meters water equivalent), the surface muon flux is reduced by a factor of 5×10^4 .




Figure 2.2: Muon flux versus depth for various underground laboratories.

The physics lab at Soudan consists of two large caverns: one to house the large MINOS far detector, and one to house CDMS and several smaller experiments. To support the full-time staff and rotating CDMS shift crew, the halls also contain most of the amenities needed to support 10-hour days underground, including office space, a fully-functional kitchen, and a ping-pong table. While on shift, a CDMS scientist is typically responsible for controlling the data acquisition system and performing calibrations, in addition to monitoring the cryogenics and trouble-shooting problems as necessary.

2.2 SuperCDMS at Soudan

The CDMS experiment is located within an RF-shielded clean room (the RF room) to protect the detectors from radiogenic contamination. To provide additional protection beyond the clean environment of the RF room, shielding was constructed around the vacuum cans that contain the detectors to reduce the rate of background interactions, as shown in Fig. 2.3. This shielding consists of several layers: an active muon veto, which can be used to reject events that are coincident with a signal in the veto panels, a layer of polyethylene shielding, two layers of lead shielding, followed by an inner layer of polyethylene. The polyethylene moderates low-energy neutrons produced in radioactive decays to energies below detection threshold. Inside the first layer of polyethylene is a layer of lead shielding, the inner 4.5 cm of which is ancient lead recovered from a sunken ship near Nantes, France. A second layer of polyethylene inside the lead shielding provides further neutron moderation.

Within the multiple layers of shielding lies the CDMS “icebox,” the volume kept cold by the Oxford Instruments dilution refrigerator. The icebox consists of six nested copper cans, which provide a small amount of additional shielding. The copper used in the icebox was screened to ensure low levels of radioactivity. The copper cans and contents of the icebox are thermally connected to the fridge via the cold stem, but are isolated from any cryogenics or radioactive materials in the fridge itself.

The 15 germanium detectors contained within the icebox are kept at a temperature of ~ 50 mK by the dilution refrigerator during normal data-taking. CDMS detectors




Figure 2.3: SuperCDMS shielding and muon veto.

consist of 1-inch thick, 3-inch diameter cylindrical germanium crystals with sensors patterned on the flat faces using photolithography. These detectors are stacked in groups of 3 to form 5 towers, as shown in Fig. 2.4, left. Detectors are labeled by tower number and position within the tower, e.g. T4Z1 is the top detector in tower 4. The sensors on the top and bottom faces of the crystals are separated into four phonon sensors and two charge sensors per side, as shown in Fig. 2.4, right. The phonon sensors are labeled A (outer guard ring sensor) and B through D (three inner channels) for each side, while the charge sensors are labeled by inner (QI) and outer (QO) for each side.

During the course of commissioning the detectors and subsequent operations, shorts were discovered on the phonon or charge channels for several detectors. These shorts are monitored throughout the course of the run by periodic continuity checks known as FEDUP checks (possibly because the shifters responsible for performing the checks quickly become fed up with the process of moving readout cables to and from the electronics that perform the tests). Once a short is identified, it is noted within the data processing software, and the subsequent calibration and data analysis also attempts to account for the effects of the short. Chapter 4 of this thesis discusses simulations used to study how one particular detector with a charge bias channel




Figure 2.4: Left: SuperCDMS Soudan tower layout. Right: channel layout for a SuperCDMS Soudan detector.

shorted to chassis ground behaved in the low-mass dark matter search discussed in Chapter 3.

2.2.1 CDMS detector physics

When an interaction occurs in a CDMS detector, phonons and ionization are produced as shown in Fig. 2.5. The charges are drifted towards readout sensors by applying a bias of $\pm 2V$ to electrodes on the top and bottom faces of the crystals. The electrons propagate towards the top surface, biased at $+2V$, at an angle relative to the electric field vector. This phenomena, due to the band structure of germanium, is well-understood both experimentally and theoretically [94]. The holes propagate along the electric field lines to the bottom side of the detector, biased at $-2V$.

As the charges propagate through the electric field, they produce additional Luke phonons in an amount proportional to the number of electron-hole pairs and the bias voltage of the crystal [95]. The total phonon energy measured by the detectors is then the sum of the recoil energy (the prompt phonons) and the Luke contribution for an event:

$$E_{tot} = E_{recoil} + E_{Luke} = E_{recoil} + n_{eh}eV_b = E_{recoil} \left[1 + Y \frac{eV_b}{\epsilon} \right] \quad (2.1)$$

where $E_Q = \epsilon n_{eh}$, $E_{recoil} = YE_Q$, and $\epsilon = 3\text{eV}$ is the energy to produce a single




Figure 2.5: Cartoon of the physics processes that occur in a CDMS detector. An interaction produces ionization and prompt phonons. As the charges drift across the crystal, additional Luke phonons are produced.

electron-hole pair. The measurement of both the total phonon energy and the charge energy E_Q determines the recoil energy of each event independent of recoil type.

The primary means of discriminating between signal (bulk nuclear recoils) and background (electron recoils) is the ratio of charges to recoil energy, a quantity known as yield (Y in the equation above). Electron recoils produce proportionally more charge per keV of recoil energy than nuclear recoils. The charge signal is calibrated to an electron-equivalent energy scale (keVee) so electron recoil events have yield centered around unity. Nuclear recoil events, on the other hand, have yield typically around 0.3. The separation between the nuclear recoil and electron recoil bands can be seen in Fig. 2.6. The dashed green line shows the mean of the nuclear recoil band (measured with neutron calibration data), while the solid green lines show the $\pm 2\sigma$ contours. The many- σ separation between nuclear recoils and electron recoils gives CDMS excellent discrimination against electron recoil backgrounds.

There are numerous processes that can reduce the amount of charge collected in the detectors. These processes are especially problematic for background electron recoils, since reducing the amount of charge collected could move an event from the electron recoil band to the nuclear recoil band, mimicking a signal. One cause of reduced charge collection is edge effects: charges can become trapped near the edges of the detectors, preventing them from reaching the charge electrodes. Other




Figure 2.6: Yield vs. recoil energy in a CDMS detector [96]. Blue points show bulk electron recoils, while red points show events from a source that produces surface lead recoils and other lead decay chain events. The nuclear recoil band is shown in green.

effects near the detector surface, such as carriers diffusing into the wrong electrodes, also contribute to reduced charge collection. SuperCDMS detectors are designed to account for these effects, as discussed in subsection 2.2.2 below. In addition, poor space-charge neutralization within the crystals, where impurities in the crystals maintain a net charge, can act as traps for drifting charges at cryogenic temperatures and low electric fields. This effect is mitigated by routinely neutralizing the detectors by flashing LEDs. The LED photons produce electron-hole pairs in the crystal, which can neutralize the charged impurity sites, reducing the likelihood that a drifting charge will become trapped. The LED flashes are included as an automated process in the Soudan data-acquisition software: after the detectors have been in active data-taking mode for several hours, data acquisition is paused for an LED flash.

2.2.2 iZIP detectors

The detectors currently being operated at Soudan, known as interleaved Z-sensitive Ionization and Phonon detectors (iZIPs), are an improvement to the original ZIP technology used in CDMS II. The limiting background for CDMS II was events occurring very near the surface of the detector, which suffer from reduced charge collection. To

mitigate these effects, charge sensors (biased at $\pm 2\text{V}$), and phonon sensors (grounded) were interleaved on both flat faces of the detector. (Contrast this with the CDMS II detectors, which had phonon sensors on one side and charge sensors on the other.) This interleaved design creates the electric field pattern seen in Fig. 2.7, left.




Figure 2.7: Left: Electric field geometry near the surface of an iZIP detector. Right: Charge collection on side 2 vs. side 1 of a CDMS, determined by the iZIP electric field geometry. A source of surface events was placed near side 1, producing the population of red points with Side 2 Charge Collection consistent with zero. Bulk recoil events, as determined by the dashed-line boundaries, are shown in blue.

For events that occur in the bulk of the detector ($\gtrsim 1\text{ mm}$ away from detector face), the electrons and holes will propagate to opposite faces of the detector, producing approximately equal signals on both sides of the detector. In contrast, for events that occur within $\sim 1\text{ mm}$ of the surface, one sign of charge will propagate to the biased electrode, while the opposite sign will be drawn to the grounded phonon electrode. This will produce an asymmetric charge signal where charge is only recorded on one face, as shown in Fig. 2.7, right. This asymmetry allows such surface events to be clearly identified in the data. The iZIP design was shown to provide discrimination against surface events at better than 1 part in 10^5 , corresponding to less than 1 event in the planned SuperCDMS SNOLAB exposure [96].

In addition to using the electric field geometry to reject events that occur near the flat surfaces of the detectors, the distribution of energy between sensors is used to reject events that occur near the sidewalls of the detectors. Such events may be

caused by decays from radiogenic contamination of the housings, and they may also suffer from reduced charge collection due to edge effects in the electric field. The “radial partition,” or ratio of energy collected in the outer channel to total energy collected, is calculated independently for the two charge measurements (side 1 and side 2) and the phonon signal. The majority of the energy for events that occur near the sidewalls is typically collected in the outermost channel (QOS1/S2 or PAS1/S2 in Fig. 2.4, right). Therefore, the charge and phonon radial partitions can serve as a second means of rejecting backgrounds.

2.2.3 Readout Electronics

The detector towers inside the icebox also contain the cold hardware and electronics required to operate the detectors, as well as the heat-sinking and wiring to bring the signals to room temperature. SuperCDMS uses the same readout electronics as CDMS II: field effect transistors (FETs) to read out the charge signals and superconducting quantum interference devices (SQUIDs) to read out the phonon signals, one for each charge and phonon channel. The electronics cards containing the FETs are weakly heat-sunk to the 4K temperature stage, though they self-heat to a temperature of $\sim 130\text{K}$. The electronic noise in the FETs is a significant factor in determining the detection threshold for SuperCDMS Soudan, so SuperCDMS SNOLAB plans to use High Electron-Mobility Transistors (HEMTs), which can be operated at significantly lower temperatures. A brief characterization study of some of the SNOLAB HEMTs is discussed in Appendix C.

The phonon signals are read out by SQUIDs operated at 600 mK. Each SQUID reads out the current from one phonon channel, which consists of arrays of quasiparticle-assisted electrothermal-feedback transition-edge sensors (QETs) in parallel. A single QET from a Soudan detector is shown in Fig. 2.8. Each QET consists of a micron-wide tungsten transition edge sensor (TES) connected to aluminum collection fins. A phonon absorbed in the superconducting aluminum fins will break a Cooper pair, creating quasiparticles, some of which absorb into the TES. This heats up the TES, which is held stable in its superconducting-to-normal transition by an electrothermal

feedback process [97]. The TES is held at a constant voltage, so as it heats up and its resistance changes, the resulting change in the current is read out by the SQUIDS.




Figure 2.8: QET on the face of a SuperCDMS detector. The dark region is the detector substrate, the light region is the aluminum fins, and the thin red lines show the tungsten TES.

The signals are carried to the room temperature electronics by striplines that pass through the electronics stem to connect with the front-end boards (FEBs), which are located inside the RF room. These boards contain electronics to control the LED flash mode, the cold electronics, and detector settings such as bias voltages for the channels, as well as to amplify the signals. The signals are then brought outside the RF room to the electronics room on the level immediately above the RF room.

The electronics room contains the DAQ control computer, as well as the remaining readout electronics, including the trigger electronics. The signals from the FEB boards are received by the receiver-trigger-filter (RTF) boards, which generate the trigger input signals. Since the RTF boards were originally designed for CDMS II but SuperCDMS Soudan detectors have twice as many channels, each RTF board monitors the incoming signals from half the channels of one detector. The RTF board applies a band-pass filter to the analog signal, creates total phonon and charge pulses by summing the 4 individual phonon or 2 charge channels, and determines if a trigger

was issued by comparing the signals with externally-defined thresholds. For the low-mass WIMP search discussed in Chapter 3, the RTF boards for the best-performing detectors underwent CROW (CDMS RTF optimization work) modifications to optimize the band-pass filters in order to improve the signal to noise of the filtered pulse, with the goal of lowering trigger thresholds. The pulses that are filtered by the RTF boards are then passed to a Struck digitizer. The detector signals are digitized by taking 2048 samples (one every $0.8\ \mu\text{s}$) for each charge channel and 4096 samples (one every $1.6\ \mu\text{s}$) for each phonon channel (the phonon pulses have significantly longer fall times than the charge pulses).

The global trigger condition for events in the iZIP detectors is determined by a trigger logic boards (TLB), which determines whether and when to issue a global trigger. For most fully-functional detectors, a global trigger is issued if one of a detector's two trigger boards has a phonon signal that crosses threshold (the logical OR of the two boards). However, for a few detectors with particularly noisy channels, a global trigger is only issued when both of a detector's trigger boards have a phonon signal that crosses threshold (the logical AND of the two boards).

The RTF boards have convenient lights that blink whenever a trigger is issued, aiding CDMS onsite shifters in tuning the trigger thresholds if necessary. Throughout data-taking at Soudan, occasional periods of bursty trigger behavior occurred because of changes to the noise environment. Much of this problematic data could be removed at the software level after data processing, but these trigger bursts can also be controlled by raising trigger thresholds. This is usually accomplished by a shifter going into the electronics room, adjusting the trigger thresholds using the DAQ control computer, and monitoring the rate reported by the blinking lights and the DAQ monitor computer until they subside to the desired trigger rate.

2.2.4 Data processing

After the raw digitized pulses are brought to the surface, they are processed through the CDMS reconstruction software known as *cdmsbats*. This software package contains all the algorithms necessary to produce reduced CDMS data, which consists

of a large number of events, each with series of features including energy estimators, yield, and energy partitions.

At the heart of *cdmsbats* is an algorithm known as the optimal filter, which fits each pulse to a pre-determined template in the frequency domain to determine the pulse height and time offset. This algorithm is covered in detail in the appendices of Sunil Golwala and Jeff Filippini's theses [98, 99], so I will only discuss the basics here.

We would like to find the best-fit amplitude a for the signal

$$s(t) = n(t) + aT(t) \quad (2.2)$$

where $s(t)$ is the measured signal in the time-domain, $n(t)$ is the noise, and $T(t)$ is a fitting template normalized to 1. These templates are determined by averaging together a large number of pulses from bulk electron recoil calibration events and then normalizing the averaged pulse to unity. This procedure is performed for each detector separately to account for detector-to-detector variations in the pulse shape. Several additional pulse shapes were defined to match other abnormal pulse shapes in order to remove non-physics events caused by glitches and low-frequency noise from the data. Appendix A discusses the development of the low-frequency noise optimal filter template and subsequent selection criteria.

Fits are often performed by minimizing the χ^2 , which is a measure of the signal's deviation from the expected shape. The CDMS optimal filter performs this fit in the frequency domain, since the noise spectrum $n(t)$ in the denominator of the time-domain χ^2 is typically correlated in time. The optimal filter minimizes the discrete frequency-domain χ^2

$$\chi^2 = \sum_n \frac{|\tilde{S}_n - a\tilde{T}_n|^2}{J(f_n)} \quad (2.3)$$

where $\tilde{S}_n$ is the frequency-domain signal (consisting of a pulse plus noise), a is the amplitude of the pulse, $\tilde{T}_n$ is the frequency-domain pulse fitting template (assumed to be normalized to unity), and J_n is the noise power spectrum. Minimizing the χ^2

to find the best-fit amplitude $\hat{a}$ gives

$$\hat{a} = \frac{\sum_n \frac{\tilde{T}_n^* \tilde{S}_n}{J_n}}{\sum_n \frac{|\tilde{T}_n|^2}{J_n}} \quad (2.4)$$

This can be thought of as the application of a filter $\phi \equiv \tilde{T}_n^*/J_n$ to the signal, hence the name “optimal filter.” This filter de-weights the bins that have high noise, leaving the other bins relatively unchanged.

In addition to the amplitude, the time shift of the pulse must be estimated since the exact start time of the pulse varies between traces. In this case, the template is shifted by a time offset t_0 in the χ^2 expression of Eqn. 2.3:

$$\chi^2 = \sum_n \frac{\left| \tilde{S}_n - a e^{-2\pi i t_0 f_n} \tilde{T}_n \right|^2}{J(f_n)}, \quad (2.5)$$

which yields a best-fit amplitude of

$$\hat{a} = \frac{\sum_n \frac{e^{2\pi i t_0 f_n} \tilde{T}_n^* \tilde{S}_n}{J_n}}{\sum_n \frac{|\tilde{T}_n|^2}{J_n}} \quad (2.6)$$

The time-shift can also be determined by minimizing the χ^2 , this time with respect to t_0 . The expression $\partial\chi^2/\partial t_0$ is a nonlinear equation with no analytic solution; however, by taking the derivative of $\hat{a}$ with respect to t_0 , we find that $\partial\chi^2/\partial t_0 \propto \partial\hat{a}/\partial t_0$. Thus, the best-fit value of t_0 is that which maximizes the amplitude $\hat{a}(t_0)$.

Modifications to the optimal filter algorithm were developed specifically to account for cross-talk between charge channels on the detectors, caused by the mutual capacitance of the electrodes. Additional templates T_{Ix} and T_{Ox} are constructed that match the expected crosstalk shape in the inner and outer channels. The expression for the signal in a given channel then can be formulated as a 2×2 matrix equation:

$$\begin{bmatrix} s_I \\ s_O \end{bmatrix} = \begin{bmatrix} T_I & T_{Ix} \\ T_{Ox} & T_O \end{bmatrix} \begin{bmatrix} a_I \\ a_O \end{bmatrix} + \begin{bmatrix} n_I \\ n_O \end{bmatrix} \quad (2.7)$$

The χ^2 is then a function of two parameters a_I and a_O , and the best-fit values for these amplitudes can be found via a matrix generalization of Eqn. 2.6. In addition, the optimal filter delays for the charge pulses on the two sides of the detector were constrained to be within $2\mu\text{s}$ of each other. The cross-talk optimal filter and the time-delay constraint could be turned off in the case of detectors with known charge problems.

Because the shape of the phonon pulses is dependent upon event location, another type of optimal filter was developed that accounts for this additional dependence. This algorithm, called the non-stationary optimal filter, treats the position-dependent part of the pulse (primarily the first $\sim 100\mu\text{s}$) as a source of non-stationary noise. In this case, the power spectrum $J(f_n)$ becomes a covariance matrix because of the frequency-frequency correlations in the non-stationary position dependent part of the pulse. This covariance matrix is calculated by subtracting the phonon pulse template from a number of phonon pulses (Fig. 2.9, left), producing a number of residual spiky pulses (Fig. 2.9, right). The variance and covariance of these spiky residual pulses is used in the power spectrum calculation. For the short timescales of the position-dependent spike in the pulse, the amplitude in the corresponding frequency bin is large, de-weighting those bins in the χ^2 , therefore mitigating the position-dependence in the pulse and improving the resolution. This algorithm is slow because it requires inverting a large matrix, so it is applied only to the total phonon pulse and only above an energy threshold.

CDMSbats also contains other data-processing algorithms. These include the RTFT walk, which measures the rise times and fall times of the pulses, as well as a time-domain phonon pulse fitting algorithm and a multiple-template optimal filter. After running the optimal filter and other data-processing algorithms, the data is calibrated as described in subsection 2.2.5 below. From the calibrated data, the software calculates additional calibrated reduced quantities, such as the yield and energy partitions, for each event.




Figure 2.9: Left: Sample of phonon pulses for T1Z1 (colored traces) and phonon pulse template (black) constructed by averaging. Right: Position-dependent residual traces for T1Z1, constructed by subtracting the phonon template from individual phonon pulses. The covariance of these residuals is used in the non-stationary optimal filter. [Figure: Yohan Ricci]

2.2.5 Calibration

The detector response is calibrated periodically by means of external radioactive sources. Two tubes were constructed to allow sources to pass through the shielding, one near the electronics stem and one near the cold stem. It is the job of onsite scientists to insert and remove the sources, which are on the end of long flexible wires, from the experiment (occasionally in strange circumstances, as in summer 2012, when the AC power to the RF room was temporarily cut following the discovery that the AC filter was producing high noise in the detectors, leaving the room in complete darkness).

Electron recoil calibration was performed using a ^{133}Ba γ source, and nuclear recoil calibration was performed using a ^{252}Cf neutron source. Since the ^{133}Ba data is also used to assess the trigger thresholds and stability of the detectors, approximately 15 hours of barium calibration data are taken per week. On the other hand, ^{252}Cf calibrations are performed for several days straight once every few months since neutrons can activate the germanium crystals, which leads to higher event rates as the activated isotopes decay by electron capture.

Calibration of CDMS detectors is a multi-step process, beginning with the overall charge calibration. ^{133}Ba produces several distinct gamma-ray lines as it decays, including 276 keV, 303 keV, 356 keV, and 384 keV. Many of these lines are detectable in CDMS detectors and can be used for calibration. For the purposes of calibrating the charge channels, the voltage bias was set so that holes were collected on the side being calibrated. This was done to avoid the inherent smearing effects from the oblique propagation of electrons in germanium: due to the angle of propagation, many more events with energy shared between the inner and outer channel are found for the side that collects electrons.




Figure 2.10: Charge calibration procedure for inner-channel events. [figure: Adam Anderson]

Figure 2.10 demonstrates the algorithm for calibrating the inner charge channel. Once the height of the uncalibrated pulses are calculated for barium calibration data using the optimum filter, a histogram of the uncalibrated charge energy is constructed for the inner charge channel of the side being calibrated, and the peaks are identified by eye. (Note that the 356 keV line is typically visible in all detectors, but the other lines are not always identifiable above the broad continuum.) The data in the region around the peak is binned and fit to a functional form that is the sum of a linear

background and skew Gaussian. The position of the peak is determined by finding the maximum of the best-fit function, and the calibration constant for the inner channel is determined by fitting the peak positions to a linear function.




Figure 2.11: Charge calibration procedure for outer-channel and shared events. [figure: Adam Anderson]

To determine the outer channel calibration, the band of 356 keV events with energy shared between the inner and outer channels is identified by eye, as shown in the upper left quadrant of Fig. 2.11. A number of points are selected in this band, and fit to a line. Then, the band of events shared between inner and outer channels (blue box in Fig. 2.11, upper left) is re-fit by computing the distance to this line and placing data into 5 bins in inner-channel energy. The baseline of the distribution is subtracted for each bin, and the mean of the distribution is found. The band of shared events is then re-fit using the new estimated means. Finally, the shared events are fit to a linear function constrained to pass through 356 keV, and the result is extrapolated to the outer channel axis to yield the outer-channel calibration.

Prior to performing the absolute phonon calibration, a relative calibration must be applied to the eight phonon channels. We expect the falling exponential behavior of the pulses for each of the eight channels to be identical since the distribution of

athermal phonons is uniform in the crystal after several hundred μs , so the power absorbed in each channel is identical. The relative calibration is accomplished by fitting the late-time portion of the trace (800 μs after the trigger to the end of the trace) to an exponential decay plus a constant. The integral of the exponential fit function is computed for each channel, and it is subsequently normalized to one reference channel via a least-squares fit.

Once the relative phonon calibration is applied, the total phonon pulse is constructed by summing the eight individual-channel pulses. The uncalibrated amplitude is computed for the total phonon pulse using the non-stationary optimal filter, which treats position-dependence in the rising edge of the pulse as an extra non-stationary noise term. The expected total phonon energy for each barium calibration event is computed according to Eqn. 2.1 using the calibrated charge quantities to estimate the recoil energy and Luke phonon energy. This estimation of the total phonon energy is then fit versus the uncalibrated pulse height to determine the overall phonon calibration. Since the phonon calibration was discovered to be sensitive to the base temperature of the fridge on the order of a few percent, an additional correction was applied to correct this dependence.

After the overall phonon and charge calibrations are applied, the charge energy as a function of phonon energy is determined for nuclear and electron recoils (what is referred to as the electron recoil and nuclear recoil “bands”). Electron recoil or nuclear recoil calibration data is broken up into energy bins, and the median charge energy in a bin is determined by fitting the binned data to a gaussian. The means and standard deviations of the gaussians are then fit to an empirical functional of Eqn. 2.8 form to determine the bands.

$$f(p_t) = \alpha_1 + \alpha_2 p_t + 10^{\alpha_3} \times \text{erf} \left[-\frac{p_t}{10^{\alpha_4}} \right] \quad (2.8)$$

Figure 2.12 shows the nuclear recoil band resulting from the above procedure, computed in units of mean charge energy versus total phonon energy (ptNF). The measurement agrees with the theoretical Lindhard prediction (black) at high energies

[100, 74], but differs below total phonon energies of approximately 10 keV. The magenta points show the binned measurement, while the red curve shows the best fit to the data using the charge model of Eqn. 2.8. A Markov-chain Monte Carlo, which samples the parameter space of Eqn. 2.8, was used to understand the possible distributions of the parameters in the charge model, and the results from this sampling are shown in shades of cyan.




Figure 2.12: Example nuclear recoil energy scale for detector T2Z2, measured using ^{252}Cf neutron calibration data. The cyan curves show MCMC samplings of the charge model parameterization. [figure: Julien Billiard]

2.3 Recent CDMS results

CDMS II and SuperCDMS Soudan have produced a number of recent results relevant to this thesis. The low thresholds of CDMS detectors yield excellent sensitivity to low-mass dark matter, and CDMS has developed numerous techniques to improve its sensitivity, including lowering the thresholds via Luke amplification (CDMSlite) and modeling known backgrounds to perform a likelihood analysis. These results serve as motivation for the low-mass WIMP search discussed in Chapter 3, since SuperCDMS Soudan, with its low thresholds and excellent background discrimination, is well-poised to confirm or extend exclusion limits and follow up on possible low-mass dark

matter signals.

2.3.1 CDMS II Silicon

CDMS II used silicon detectors in addition to germanium detectors to search for dark matter. Since silicon is a much lighter element than germanium, these detectors provided an excellent means to follow up on potential low-mass dark matter signals seen by DAMA [44], CoGeNT [45], and CRESST [46]. This result considered data taken using the silicon detectors of CDMS II between July 2007 and September 2008, for a total exposure of 23.4 kg-days [47].

Three candidate events were found in this dataset upon unblinding. These events are shown in the discrimination space of yield versus timing parameter (used to identify surface events) in Fig. 2.13, left. The estimated backgrounds from surface events, neutrons, and ^{206}Pb recoils for this dataset was less than one event, and the probability that these known backgrounds would produce 3 or more events was 5.4%. A profile likelihood test indicates that the probability for the hypothesis that only includes the expected backgrounds is 0.19% when tested against a WIMP+background hypothesis. The best-fit contours from this profile likelihood test are shown in blue in Fig. 2.13, right.

2.3.2 CDMS II low-energy likelihood

A second result from CDMS II is a reanalysis of the original low-energy germanium result [101] that uses likelihood analysis to effectively subtract the known backgrounds [102]. This analysis used simulation and calibration data to model the expected backgrounds from Compton scatters and ^{210}Pb decay-chain events.

The original low-energy germanium analysis from CDMS II did not attempt a background subtraction, instead opting to compare to the number of expected background events estimated from a portion of the dataset left unblinded. As a result, the analysis included all observed candidate events when setting a limit, even though these events could largely be explained by known backgrounds. A likelihood analysis,




Figure 2.13: Left: Signal region and three candidate events from the CDMSII silicon result [47]. Right: Allowed region for the CDMSII silicon result (blue enclosed region), with other contemporary results including CDMSII Ge (dashed maroon), CRESST (enclosed magenta), and XENON100 (dash-dot green).

on the other hand, is expected to produce a stronger limit in the presence of well-modeled backgrounds, since those backgrounds are effectively subtracted out when setting the limit.

The expected electron recoil and lead decay chain backgrounds were modeled using barium calibration data and Geant4 simulations, respectively. The fits showed no significant evidence for a WIMP component, so an exclusion limit was calculated. This limit is shown in black in Fig. 2.14. The result using likelihood techniques represents a significant improvement over the limit produced using the optimum interval method with no background subtraction (thin magenta line).

This analysis demonstrates the efficacy of accurately modeling backgrounds. The limit calculated from this likelihood analysis, which subtracts the known backgrounds, produces approximately a factor of five improvement over the limit with no background subtraction. Future analyses will need robust background models to achieve similar improvements in sensitivity, demonstrating the need for accurate detector modeling like the work discussed in Chapter 4.




Figure 2.14: Limit set by a likelihood analysis of CDMS II low-energy germanium data (black) [101]. This limit is ~ 5 times stronger than the corresponding limit (magenta) with no background subtraction. The deviation between the limit and the expected sensitivity (green bands) is also seen in the multiples data, indicating that it is due to an unaccounted-for systematic, not a WIMP component.

2.3.3 CDMSlite—a low ionization threshold experiment

A novel readout mode introduced for SuperCDMS Soudan is known as CDMSlite—a low ionization threshold experiment [103]. This mode takes advantage of the Luke effect (outlined in Eqn. 2.1) to amplify very low-energy recoils above the total phonon energy threshold by increasing the bias voltage on the detector from 4V to 70V. However, this amplification comes at the expense of electron recoil/nuclear recoil rejection, since at high voltages the phonon signal is simply proportional to the charge signal.

To operate a Soudan detector in CDMSlite mode, a special adaptor electronics board was created to hold one side at high voltage and allow the standard Soudan electronics to measure total phonon energy using the other side. The operating bias voltage was 69 V, leading to a gain of 24 for electron recoils with yield of 1. This gain

gave a baseline resolution of $14 \text{ eV}_{\text{ee}}$ (electron-equivalent) and a trigger threshold of $170 \text{ eV}_{\text{ee}}$.

A single 0.6-kg detector was operated for a total of 15.7 live-days (10.3 live-days after removing periods of unstable gain or abnormal operation). The resulting spectra is seen in Fig. 2.15, left. The K-shell ($\sim 10.4 \text{ keV}$) and L-shell ($\sim 1.3 \text{ keV}$) lines are clearly identifiable in the spectra, with a relatively flat background in between.

The electron-equivalent energies were converted to nuclear recoil energies in order to set a limit on WIMPs. Because the ionization yield is not measured in this readout mode, the theoretical prediction from Lindhard was used to perform this calculation. No background subtraction was performed, and the optimal interval method was used to calculate the 90% C.L. upper limit, shown in black in Fig. 2.15, right. The extremely low threshold achieved by Luke amplification leads to increased sensitivity to low-mass WIMPs: this analysis produced limits that excluded previously-untested regions of parameter space below WIMP masses of $6 \text{ GeV}/c^2$.




Figure 2.15: Left: Spectrum from CDMSlite Run 1, clearly displaying the K-shell ($\sim 10.4 \text{ keV}$) and L-shell ($\sim 1.3 \text{ keV}$) activation lines. Right: Exclusion limits from CDMSlite Run 1. [figure from [103]]

Chapter 3

Search for Low-Mass Dark Matter with SuperCDMS Soudan

SuperCDMS is well-suited to searching for low-mass dark matter due to its low thresholds and excellent background discrimination. Most scatters from low-mass dark matter will have recoil energies of less than 10 keV, but CDMS detectors with thresholds of a few keV will be sensitive to such small energies. This chapter discusses the search for low-mass WIMPs performed using the data collected at Soudan between October 2012 and July 2013 with total phonon energy between 2 and 13.1 keV ($\sim 1.6 - 10$ nuclear recoil equivalent energy (keV_{nr})). This data was taken using the seven detectors with the lowest thresholds, for a total exposure of 577 kg-days [104].

3.1 Event selection criteria and background model development

3.1.1 Selecting high-quality physics events

To perform a search for dark matter, we must first remove all events that would be inconsistent with a dark matter particle scattering in the detectors. We remove poorly-reconstructed events and events where multiple pulses occur during the read-out window by placing a cut on the optimal filter χ^2 of charge and phonon pulses.

As a second means of removing events with pileup or high noise, we require that the standard deviation of the pre-pulse portion of each trace fall within 4σ of the mean of the distribution. In addition, we remove events with abnormal pulse shapes, such as electronics glitches and low-frequency noise triggers, by considering the differences between the fit to the pulse template and specially-developed abnormal pulse templates. The development of the low-frequency noise templates and selection criteria is discussed in further detail in the Appendix A. We also require that the phonon pulse start time as determined by the optimal filter fitting algorithm did not fall near the edges of the optimal filter time window, which is defined as the time range from $200\ \mu\text{s}$ before the global trigger to $100\ \mu\text{s}$ after the global trigger.

In addition to defining event-by-event selection criteria, we remove periods of livetime associated with poor detector performance. We removed periods of time when the base temperature of the fridge was too high or too low, which may have affected the detector calibration. The 20 days following neutron calibration were not considered for WIMP-search analysis due to the elevated event rates from internal activation lines. Since this search focuses on the low-energy region (near trigger threshold), we only consider periods of time where the thresholds were stable. This allows for sufficient calibration statistics that the threshold can be reliably estimated. We also remove periods of abnormally high trigger rates or high charge noise. Finally, we use two-sample Kolmogorov-Smirnov (KS) tests on a variety of phonon and charge quantities to assess data quality and remove periods of poor quality.

We also remove from consideration events that deposit energy in multiple detectors within the trigger window, since the probability for a WIMP to scatter in multiple detectors is negligible. We use timing information from the muon veto to remove events coincident with the veto, which may have been due to spurious neutrons produced by cosmic rays. During the periods when the NuMI neutrino beam was on, events within $200\ \mu\text{s}$ of a beam dump are also removed from consideration, though the probability of such an event originating from a neutrino scatter is small. We also require that, for each event in a given detector, that detector must be the source of the global trigger, i.e. for an event that occurred in T1Z1, T1Z1 must have issued a trigger.

Events that occur near the surface of the detectors tend to have reduced yield due to charge trapping. Since low-yield events can leak into the signal region, especially at the low energies considered in this analysis, we define a fiducial volume to remove events that occur near the surfaces of the detectors. Separate criteria were defined to remove events occurring near the flat surfaces and events occurring at high radius. Figure 3.1, left, shows the definition of the surface event rejection criteria. Due to the iZIP electric field geometry, events near the surface will record a charge signal on only one side of the crystal, so we require that events fall in the symmetric $x=y$ band in this plot. Figure 3.1, right, shows the definition of the high-radius rejection criteria. Events which deposit charge energy primarily in the outer guard electrode are rejected because they suffer from reduced charge collection. Both cuts were designed to pass the low-energy random triggers taken to characterize the noise environment.




Figure 3.1: Fiducial volume selection for detector T2Z1. The surface event criteria is shown on the left, and the radial criteria is shown on the right. Red events pass the selection criteria, while black events fail the criteria. [Figure: Todd Doughty]

Finally, we require that events deposit charge energy consistent with a nuclear recoil. The nuclear recoil bands were defined in the plane of charge energy (defined as the mean charge collected on the two inner electrodes) versus total phonon energy, as shown in Fig. 3.2. Californium neutron calibration data was binned in phonon energy, and the data in each bin was fit to a gaussian. The mean, $\pm 1\sigma$, and $\pm 2\sigma$

contours were fit to the function

$$f(p_t) = \alpha_1 + \alpha_2 p_t + 10^{\alpha_3} \times \text{erf} \left[-\frac{p_t}{10^{\alpha_4}} \right] \quad (3.1)$$

This function was chosen since it gives a linear relationship between charge and phonons at high energies, as predicted by Lindhard. The resulting fits are shown in Fig. 3.2. A 2σ nuclear recoil band was used for the individually-tuned final event selection method discussed below, while a 3σ band was used for the other two methods since they included tighter discrimination cuts on charge and phonon energy.




Figure 3.2: Nuclear recoil bands for detector T1Z1, defined using neutron calibration data. [Figure: Julien Billard]

3.1.2 Background model development

We blind all data that may be consistent with a WIMP scatter in order to avoid bias when setting discrimination quantities. All single-scatter events above threshold and below 10 keVnr (13.1 keV total phonon energy) which were within the 3σ nuclear recoil band, were not coincident with the muon veto, and did not have abnormal pulse shapes were blinded. Twenty days following the neutron calibration were also left unblinded due to the high activation rates. As a result, it was necessary to create

an independent model of the expected backgrounds to tune the final level of event selection.

This model was created by selecting high-energy template events from barium calibration or WIMP-search data and scaling them to the energies relevant to this analysis. Events were selected for each detector to naturally account for the small variation in energy collection and partitioning between detectors. Three types of backgrounds were considered: electron recoils from low-energy gammas, high-radius and surface events from the lead recoil chain, and L-shell germanium activation lines. High-energy barium calibration events were used as templates for the low-energy gammas, WIMP-search events which had low yield and deposited most of their energy on one surface or in the outer electrode were used as templates for the lead 210 recoil chain, and K-shell germanium activation lines were used as templates for the L-shell activation lines. Radiogenic and cosmogenic neutrons were assumed to be sub-dominant and were not considered in the model.

The pulses from the selected template events were scaled to low energies and summed with noise traces from Soudan. The resulting pulses were then run through the standard CDMS reconstruction algorithms. Events were then weighted according to the expected rate and exposure for this analysis. The electron recoil rate was scaled to the rate of unblinded low-energy electron recoils. L-shell activation line events were scaled to the rate of K-shell events since the two rates are directly proportional. The lead decay chain was scaled to the expected number of events in the unblinded post-Cf data. Further detail on the background model and the event selection methods can be found in Adam Anderson's thesis [105].

3.1.3 Final event selection methods

Several methods were developed for the final level of selection criteria. Each one used four discrimination variables to remove events that are more likely to come from background, leaving only WIMP candidate events behind. These four variables were the mean charge energy collected on the two inner electrodes, the total phonon energy, the phonon z-partition (difference between phonon energy collected on the top side




Figure 3.3: Background model (lead = yellow, gammas = red, 1.3 keV line = blue) [Figure: Adam Anderson]

and phonon energy collected on the bottom side, divided by the total phonon energy), and the phonon r-partition (ratio of phonons collected in the outer phonon channel to total phonon energy). The selection criteria were defined to specifically remove three types of events: 1.3 keV activation line events, which can leak into the nuclear recoil band, high-radius events, which tend to come from the lead decay chain, and asymmetric surface events.

The first method for event selection featured three selection criteria that were tuned serially against separate backgrounds. First, the nuclear recoil band was tightened to the 2σ level, and a notch in the nuclear recoil band was defined to remove the 1.3 keV line at the 3σ level as defined by fits to the activation line in data taken immediately following neutron calibration (Figure 3.4, left). The second selection criteria was a cut on the phonon radial partition. This cut preferentially removes events coming from lead decay chain events, which typically deposit energy near the sidewall of the crystal. The cut was defined using the background model events from the lead decay chain and neutron calibration data re-weighted to match the expected energy spectrum of a 5, 7, 10, and 15 GeV/c^2 WIMP. The cut as a function of energy can be seen in Fig. 3.4, center, for the 10 GeV/c^2 case. Events were required to pass the logical OR of the criteria for the four WIMP masses. Finally, selection criteria were defined to remove events near the flat surfaces of the detectors, which tend to suffer from reduced charge collection. This cut removes events with large positive or negative phonon z-partition. A similar method to the radial cut was used to define

these criteria, and events were required to pass the logical OR of the 5, 7, 10, and 15 GeV/c^2 WIMP mass criteria. The cut as a function of energy can be seen in Fig. 3.4, right, for the 10 GeV/c^2 case. Events which passed the AND of the notched nuclear recoil band, phonon radial cut, and phonon symmetry cut were considered candidate events.




Figure 3.4: Serially-tuned event selection cuts. Notched nuclear recoil band (left), phonon radial cut for a 10 GeV/c^2 WIMP (middle) [figure: Adam Anderson], and phonon symmetry cut for a 10 GeV/c^2 WIMP (right) [figure: Brad Welliver]

A second method for final event selection used the same discrimination variables as the serially-tuned cuts, but the selection criteria were optimized simultaneously to maximize the expected sensitivity of the full exposure. Similar to the serially-tuned radial and symmetry cuts, selection criteria were defined using events from the background model and neutron calibration data re-weighted to match the expected WIMP spectrum for several WIMP masses. The cut boundaries in normalized charge energy (distance from the mean of the nuclear recoil band in standard deviations), phonon radial partition, and phonon z-partition are shown in Fig. 3.5.

The third method for final event selection was a boosted decision tree. This method was trained using the events from the background model, as well as neutron calibration data rescaled to the expected dark matter recoil energy spectra to model the expected signal. Four discrimination variables were used: the mean charge energy collected on the two inner electrodes, the total phonon energy, the phonon z-partition (difference between phonon energy collected on the top side and phonon energy collected on the bottom side), and the phonon r-partition (ratio of phonons




Figure 3.5: Simultaneously-tuned event selection cuts on normalized charge energy (left), phonon radial partition (middle), and phonon z-partition (right) as a function of total phonon energy. [figure: Adam Anderson]

collected in the outer phonon channel to total phonon energy). Four separate BDTs were trained against the expected signal for a 5, 7, 10, and 15 GeV/c^2 WIMP. Events which passed any of the four BDTs were considered candidate events.

Figure 3.6 shows the output BDT score for the 10 GeV/c^2 WIMP BDT. The boosted decision tree maps multivariate output to a single number between -1 and +1, where -1 is completely background-like and +1 is completely signal-like. The stacked colored histograms show the various background model components, the grey histogram shows the expected signal from a 10 GeV/c^2 WIMP, and the black points show the SuperCDMS data after applying the WIMP selection criteria discussed in Section 3.1.1. As shown by the residuals in the figure, the background model is in excellent agreement with the observed data after unblinding, with a p-value of 0.14 for this BDT. P-values for the other three BDTs range from 0.08 to 0.26.

Once the BDT scores were calculated, the boundary between signal events and background must be determined. Two methods were used to determine this value: optimizing the sensitivity of each detector individually, or optimizing the total sensitivity of all seven detectors simultaneously. The latter was chosen as the preferred method since it is expected to provide better sensitivity. BDT score thresholds for individual detectors fell between 0.24-0.37.




Figure 3.6: Distribution of BDT scores for events in background model. Black points show SuperCDMS data [104]. The corresponding uncertainty on the residuals is shown in tan in the lower subplot.

3.2 Livetime estimates and final analysis efficiency

The expected number of WIMP events is proportional to the detector mass multiplied by the live time spent waiting for an event to occur (a quantity known as exposure). Since events which occurred during periods of poor detector performance were not included in the analysis, the livetime associated with those bad periods was removed from consideration using a series of cuts. Periods of bad operation include the 20 days following californium calibration, when the event rate is elevated due to activation of the detectors; periods of unstable fridge base temperature, which may affect the detector calibrations; periods of high noise as determined by KS tests; and periods when the detector settings were not finalized. In addition, the livetime associated with individual events with certain known issues was also removed. In particular, if the pre-pulse baseline of an event was noisy, it is reasonable to expect that the livetime before the event was also noisy, so the livetime from such events was not included in the calculation. The total livetime was calculated on a detector-by-detector basis by summing the time spent waiting for each event that passed all livetime selection criteria. These values are shown in Table 3.1.

Detector	Detector mass [g]	Livetime [days]
T1Z1	609.5	129.4
T2Z1	597.4	136.4
T2Z2	591.3	134.4
T4Z2	597.4	143.8
T4Z3	594.3	138.5
T5Z2	606.5	133.8
T5Z3	591.3	131.7

Table 3.1: Detector masses and livetimes for this analysis, yielding a total exposure of 577 kg-days.

The total analysis efficiency is shown in Fig. 3.7, weighted by the exposure of each detector. The black curve shows the efficiency of the quality cuts such as pulse-shape and χ^2 cuts which remove noisy or poorly reconstructed events. This efficiency was measured using a simulation where experimental noise was added to the pulse templates used in the fitting algorithms. This provided a selection of events with known good pulse shape to determine the efficiency, which is simply the fraction of good events which pass the quality cuts. The grey curve shows the quality cuts plus trigger and threshold efficiencies, which were measured using barium calibration events. The discrete jumps show where the thresholds for different detectors turn on. The blue curve shows the efficiency of the preselection cuts, which include the 3σ nuclear recoil band and charge fiducial volume cuts. This efficiency was measured using high-quality neutron calibration data.

The final efficiency including the boosted decision tree cut is shown in red. Similar to the preselection cut efficiency, this was measured using neutron calibration data from fission of ^{252}Cf . This data contains contamination from photons produced during the fission process, and from neutrons which scatter multiple times in a single detector. Multiple internal scatters can bias the efficiency measurement since WIMPs are expected to only produce single-scatter nuclear recoils. As a result, both the preselection and BDT efficiencies needed to be corrected to account for these effects. The electron recoil subtraction was performed by normalizing barium calibration data to the expected californium gamma rate and subtracting this scaled rate from the overall number of ^{252}Cf events, while the multiple scatter correction was performed

using Geant4 [106] simulations of californium interactions in SuperCDMS detectors. 68% confidence level uncertainties on the measurement and the gamma and multiple scatter corrections are shown in the red band.

Due to limited statistics, the efficiencies of the preselection and BDT criteria were measured in 1 keV bins, which produced non-physical discrete jumps at the bin edges. Since the true efficiency is expected to be a continuous function of energy, a spline was used to interpolate between the discrete bins, producing the smooth curves of Fig. 3.7.




Figure 3.7: Analysis efficiency for 4 levels of cuts in this analysis [104].

3.3 Expected sensitivity for final event selection methods

3.3.1 Energy scale for exclusion limits

Dark matter event rates are typically quoted as a function of nuclear recoil energy. However, the quantity measured in CDMS detectors is the total phonon energy, which is equal to the recoil energy plus the Luke phonon energy. Since there are uncertainties

on the energy scale for nuclear recoils and therefore the number of Luke phonons produced, we consider the total phonon energy of the candidate event, not the nuclear recoil energy, when determining the exclusion limit. As a result, it is necessary to compute the expected WIMP event rate as a function of total phonon energy. This quantity is determined by a simple change of variables:

$$\frac{dR}{dp_T} = \frac{dR}{dE_R} \frac{dE_R}{dp_T} \quad (3.2)$$

The quantity $\frac{dE_R}{dp_T}$ is determined for each detector individually using the charge model of Eqn. 3.1. As discussed in Chapter ??, the uncertainty in the parameters of this charge model was characterized using an MCMC. Samples of the parameter space from the MCMC procedure were used to compute $\frac{dE_R}{dp_T}$, which, combined with the theoretical $\frac{dR}{dE_R}$ discussed in Section 1.4, give the spectrum as a function of total phonon energy. The total rate is then the sum of the rate for each of the seven individual detectors, multiplied by that detector's detection efficiency shown in Fig 3.8.




Figure 3.8: Final efficiency of all selection criteria for each of the seven detectors included in this analysis.

3.3.2 Sensitivity calculation

The optimum interval method [107, 108] was used to determine both the expected sensitivity of the experiment and the final exclusion limit. The sensitivity and limit calculations use the standard halo assumptions [74], with a WIMP mass density of $0.3 \text{ GeV}/c^2/\text{cm}^3$, most probable WIMP velocity of 220 km/s, mean circular velocity of the Earth with respect to the galactic center of 232 km/s, galactic escape velocity of 544 km/s, the Helm form factor, and a velocity distribution which correctly takes into account the effect of the Earth's velocity on the escape-velocity cutoff [109].




Figure 3.9: Expected sensitivity of three analysis methods (from left to right: individually-tuned cuts, simultaneously-tuned cuts, and boosted decision tree), with 68% and 95% statistical uncertainties shown by the blue bands. CDMS-II Si closed contour in green, CDMSlite limit in yellow, LUX limit in magenta.

The expected sensitivity for each method was calculated prior to unblinding using the background model developed to train the event selection criteria. Many independent realizations of the expected background were performed, and the resulting sensitivity of each pseudo-experiment was determined in the method described above. The 68% and 95% confidence level uncertainty bands on the expected sensitivity are shown in shades of blue in Figure 3.9. Since the boosted decision tree provided the best sensitivity, it was chosen as the primary analysis method.

3.4 Candidate events and agreement with expected backgrounds

Figure 3.10, left, shows the candidate events observed upon unblinding the data. Events occurred in four of the seven detectors considered for this analysis. The eleven candidate events occurred during normal periods of operation and were generally of high quality, with the exception of the lowest-energy event, which had an abnormal pulse shape. As shown in Fig. 3.10, right, the observed events were consistent with the background expectation, with the exception of the three events observed in T5Z3. The probability to observe at least 3 in T5Z3 events given the background expectation of $0.13^{+0.06}_{-0.04}$ is 4×10^{-4} , while p-values for the other six detectors range from 0.03 to 1. However, this detector has its outer ionization channel on side 1 shorted to chassis ground. We believed that we had properly accounted for the effects of the short in the background model, but it is possible that the modified electric field near the short may have affected the analysis. Further studies of this effect are discussed in chapter 4.




Figure 3.10: Left: Candidate events in this analysis. Events were consistent with expected background, with the exception of the three high-energy events in detector T5Z3 (yellow stars). Right: Expected backgrounds for each detector computed using the pulse simulation background model, with observed candidate events shown in black. [figure: Adam Anderson]

The dates, event energies, and which BDT (5, 7, 10, or 15 GeV/c^2) the event

passed are listed in Table 3.2 for the eleven candidates in this analysis. As expected, the low-energy events are more likely to pass the boosted decision tree trained to select a 5 or 7 GeV/ c^2 WIMP, since the lower the mass of the particle, the lower the expected energy of an interaction. In addition, the three events on T5Z3 pass only the 15 GeV/ c^2 selection criteria. This is because they are relatively high-energy events, but it also indicates the 15 GeV/ c^2 BDT should be the first line of investigation to try to understand the source of these events.

Detector	Event date	Total phonon energy [keV]	Passed which BDTs?
T2Z1	May 5, 2013	2.30	5 GeV/ c^2
T2Z1	June 19, 2013	2.41	5 GeV/ c^2
T2Z2	November 6, 2012	2.09	5 GeV/ c^2
T2Z2	February 23, 2013	3.24	7 GeV/ c^2
T5Z2	October 5, 2012	7.18	7, 10, and 15 GeV/ c^2
T5Z2	November 21, 2012	2.15	5 GeV/ c^2
T5Z2	June 15, 2013	3.56	7 GeV/ c^2
T5Z2	June 27, 2013	2.76	5 and 7 GeV/ c^2
T5Z3	October 12, 2012	9.79	15 GeV/ c^2
T5Z3	February 23, 2013	11.99	15 GeV/ c^2
T5Z3	April 24, 2013	8.74	15 GeV/ c^2

Table 3.2: Eleven candidate events observed in this analysis.

3.5 Exclusion limits

3.5.1 Limits on spin-independent and spin-dependent cross sections

The result exclusion limit on the spin-independent cross section is shown in black in Figure 3.11. The grey bands show 68% and 95% confidence level systematic uncertainties on the limit, resulting from propagating uncertainties in the fiducial volume efficiency, nuclear recoil energy scale, and trigger efficiency. All uncertainties are uncorrelated across detectors except the systematic uncertainty of the fiducial-volume

efficiency. The limit is consistent with the expected sensitivity of the analysis below WIMP masses of 10 GeV; above 10 GeV, the discrepancy is due to the three high-energy events on T5Z3.




Figure 3.11: Limits from this analysis (black) with 68% and 95% CL systematic uncertainties shown in gray. The pre-unblinding expected sensitivity in the absence of a signal is shown as 68% (dark green) and 95% (light green) C.L. bands.

Aside from the SuperCDMS result in black, closed contours shown are CDMS II Si [47] (*blue*, 90% C.L.), CoGeNT [45] (*yellow*, 90% C.L.), CRESST-II [46] (*pink*, 95% C.L.), and DAMA/LIBRA [110] (*tan*, 90% C.L.). 90% C.L. exclusion limits shown are CDMS II Ge [111] (*dotted dark red*), CDMS II Ge low-threshold [101] (*dashed-dotted red*), CDMSlite [103] (*solid dark red*), CDMS II low-energy likelihood [102] (*dashed dark red*), LUX [85] (*solid green*), XENON10 S2-only [112, 113] (*dashed dark green*), EDELWEISS low-threshold [114] (*dashed orange*), and CRESST low-threshold (*solid pink*) [115]. Many of the closed contours are excluded at high confidence by the SuperCDMS result. In particular, it is difficult to reconcile the SuperCDMS result with the CoGeNT closed contour, since the two experiments both use germanium as a target, so astrophysical and isospin violation arguments cannot account for the discrepancy. This limit also excludes new regions of WIMP parameter space between masses of 4 and 6 GeV.

In addition, we compute exclusion limits on spin-dependent scattering for pure




Figure 3.12: Limits from this analysis (black) with 95% CL uncertainties shown in gray for WIMP-neutron (left) and WIMP-proton (right) spin-dependent scattering.

WIMP-neutron and WIMP-proton interactions. The resulting limits are shown in Figure 3.12 in black, with 68% and 95% CL systematic uncertainties shown in gray. The SuperCDMS Soudan WIMP-neutron spin-dependent limits are competitive with earlier limits from CDMS II Ge [111] (dashed dark red) and XENON100 [84] (dashed dark green), and with the CDMSlite limits (solid dark red), and lead the field at WIMP masses below $\sim 6 \text{ GeV}/c^2$. Since germanium has one isotope, ^{73}Ge , with an unpaired neutron and a natural abundance of $\sim 7\%$ and no isotopes with an unpaired proton, the limits on WIMP-proton scattering from this analysis are less stringent than limits from experiments such as COUPP [116] (cyan band), SIMPLE [117] (dark magenta), and PICASSO [118] (dark cyan), which use target elements that have unpaired protons.

The limits in Fig. 3.12 use spin-dependent nuclear form factors for ^{73}Ge calculated by Dimitrov et al [119]. To understand the systematics associated with this calculation, we also compute limits using a more recent calculation of the nuclear form factors by Klos et al [120]. A comparison of the two calculations is shown in Figure 3.13. At the momentum transfers relevant to this analysis ($\mathcal{O}(10\text{MeV})$), the neutron form factors for the two calculations are nearly identical, and therefore the corresponding exclusion limits differ by only a few percent, as seen in Fig. 3.13, right. However, the Klos et al calculation of the proton form factor includes two-body currents, so protons

never fully decouple from neutrons. As a result, the form factor at low momentum transfer is significantly larger than the earlier Dimitrov et al calculation.




Figure 3.13: Comparison of WIMP-proton and WIMP-neutron form factors from Dimitrov et al [119] and Klos et al [120] (left), and the resulting limits on WIMP-neutron spin-dependent scattering (right).

3.5.2 Limits assuming an alternate halo model

To further understand the systematics associated with the dark matter assumptions that go into calculating the limit, we recompute the exclusion limit on spin-independent scattering using the alternate halo model discussed in chapter 5. A comparison of limits calculated using the standard halo model (blue) to limits calculated the alternate distribution with two sets of model parameters (green, red) is shown in Figure 3.14.

The green curves use values for p and v_0/v_{esc} taken from a simulation of a Milky Way-like galaxy that includes baryons. These parameters yield a differential recoil spectrum that is similar to the standard halo model (blue), though it falls off faster, especially in the tail of the distribution. This is because the alternate functional form goes smoothly to zero at the escape velocity, while the standard halo model does not. The resulting limit for these parameters agrees with the standard halo model within the systematic uncertainties, except for at the very lowest masses, where only the tail of the distribution is above threshold. For comparison, the red curve uses




Figure 3.14: Left: Spin-independent limits, comparing the standard Maxwellian velocity distribution to the alternate halo model discussed in Chapter 5. Right: Differential event rates for CDMS II Si best-fit result for the 3 halo models used in left-hand plot.

values for p and v_0/v_{esc} from a set of simulated halos that only include dark matter particles. This choice of parameters produces a differential recoil spectrum that falls off significantly faster, so the resulting limit is weaker over most of the mass range. However, at the highest masses, where most of the recoil spectrum falls within the energy range considered, the limits for all three sets of parameters agree to within the uncertainties.

Chapter 4

Simulating the behavior of detectors with shorted charge channels

SuperCDMS has developed a detailed simulation [121] which includes all relevant detector and readout physics for the detectors currently running at Soudan. The CDMS Detector Monte Carlo (DMC) can be used to study the behavior of SuperCDMS detectors, in particular the several detectors at Soudan that have cold hardware problems that affect their operation. Of particular interest are detectors have shorted charge bias channels, which can modify the electric field in the crystal. One of the detectors with a shorted bias channel is T5Z3, which saw unexpected events in the low-mass WIMP search in Chapter 3, so we hope to understand the origin of these events using the detector monte carlo.

Figure 4.1 shows a qualitative demonstration of the effects of T5Z3's modified electric field. In a fully-functional detector, the holes (red) propagate directly along the electric field lines, while the electrons propagate obliquely due to the band structure of germanium. However, in T5Z3, which has its outer charge electrode on the top side shorted to chassis ground, there is a region at high radius that has very low electric field. As a result, the electrons which would have been collected in the outer electrode are directed to the inner electrode. In addition, fewer holes propagate to




Figure 4.1: Transport of electrons (blue) and holes (red) in fully-functional CDMS detector (left) and detector with QOS1 short (right). The bias voltage applied to each channel is indicated in black.

the opposite surface due to increased trapping on the sidewall near the region of low electric field.

Small-scale simulations like the one in Fig. 4.1 demonstrate that the shorted electric field in detector T5Z3 may have significantly affected the analysis. To understand how the short may have affected the data in aggregate, simulations of large numbers of DMC events must be performed. This chapter details the development of a Detector Monte Carlo based background model, the steps necessary to turn DMC output into analyzable data, validation of the DMC against Soudan data, and several insights from the DMC, including the effects of the shorted charge channel on T5Z3.

4.1 Detector Monte Carlo-based background model

The primary goal of this simulation work is to reproduce the background model from Chapter 3 using the detector monte carlo simulation. Each of the 7 detectors was modeled separately in the DMC since several properties vary between detectors, including FET pulse shape, phonon pulse fall time, and electric field configuration.

The two electric field configurations considered for this study are shown in Fig. 4.2: the fully-functional field (left) and the T5Z3 field with its top outer electrode shorted

to ground. The electric field files needed for the simulation were produced using COMSOL. The charge electrodes were held at the proper bias voltages (± 2 V for fully-functional electrodes, or 0 V for the shorted T5Z3 electrode), and the phonon sensors were grounded, producing the scalloped pattern seen in the figure. The electric field models used in this analysis also included the detector housings and DIBs (detector interface boards), which were held at ground. Although they were not considered for this study, additional electric field files featuring other arrangements of shorted channels were produced for use with ongoing analyses of SuperCDMS Soudan data.




Figure 4.2: Electric field lines (red) and equipotentials (blue) for fully-functional electric field (left) and T5Z3 (shorted QOS1) field

Several parameters in the phonon propagation simulation were tuned so the fall times of the simulated TES pulses match the fall times of Soudan phonon pulses: the aggregate parameter corresponding to the probability that a phonon is absorbed in the aluminum fins, and the percentage of the detector surface covered by aluminum. Since two iZIP designs with slightly different aluminum coverages were installed at Soudan, the aluminum coverage was set to match the known coverage from the fabrication mask design for each detector. The major difference in the masks was the width of the charge electrodes: iZIP v4.4 had 8-micron-wide electrodes, corresponding to an aluminum coverage of 6.2%, while iZIP v5.0 had 40-micron-wide electrodes, corresponding to an aluminum coverage of 7.5%. Since the rate of phonon absorption is proportional to the aluminum coverage (more aluminum to absorb phonons means

a faster absorption rate), we expect that the fall time of the phonon pulses should be faster for the wide-electrode detectors operating at Soudan. This is clearly seen in Fig. 4.3: the detectors with wide electrodes have total phonon pulse fall times that are $\sim 100 \mu\text{s}$ faster than the detectors with narrow electrodes.




Figure 4.3: Total phonon pulse fall times by Soudan detector. The separation between wide-electrode (blue) and narrow-electrode (green) is clear, and needs to be accounted for in the DMC.

In addition, the parameters of the TES simulations were set so that each of the eight phonon channels had identical parameters including the transition temperature and transition width. This was to ensure that the saturation point for each channel was identical. Because the TES simulation assumes that all TESs in parallel in a single channel heat up and cool together, local saturation of individual TESs is not modeled in this simulation. As a result, the simulated TES pulses tend to be peakier in the initial tens of microseconds than true pulses from Soudan. To account for this, new phonon pulse templates were developed for DMC data using a similar algorithm as the Soudan template generation. Simulated barium events with good signal-to-noise and that deposited energy at least 5 mm away from any detector surface were summed and the peak normalized to unity. The resulting template for T1Z1 is shown in Fig. 4.4 alongside the Soudan template for that detector. These templates were

used to process the data to produce energy estimators and partitions, as discussed in more detail in subsection 4.1.3.




Figure 4.4: Comparison between Soudan phonon template (orange-red) and DMC phonon template (blue) for T1Z1.

The shape of the FET pulses is mainly governed by the readout electronics. The DMC FET simulation does not attempt to *a priori* reproduce the FET pulse shapes, instead calculating the expected amplitude of a pulse using the Ramo theorem [122] and scaling pre-determined pulse shape templates to the desired amplitude. Detector-to-detector variations in the FET pulse shapes were taken into account by using the detector-specific Soudan optimal filter fitting templates as pulse shape templates for the FET simulations.

4.1.1 DMC input event types and energies

Several types of expected background events were simulated to correspond with the pulse simulation used for the low-mass WIMP search discussed in Chapter 3. These are:

- L-shell lines at 1.1, 1.2, and 1.3 keV for activated isotopes of Zn, Ga, and Ge, respectively

- Barium calibration data, used for simulation validation and to model the low-energy gamma background
- ^{210}Pb housing contamination data, including sidewall events from ^{210}Pb , ^{210}Bi , and ^{206}Pb decays
- ^{210}Pb detector face contamination data, including surface events from ^{210}Pb , ^{210}Bi , and ^{206}Pb decays

The input files containing the event information (energy deposited, event type, event location) were generated using Geant4 simulations of the Soudan barium sources or contaminated housings or detectors. The two barium sources (inserted through tubes near the cryo-stem and e-stem) were simulated separately, but the resulting events were combined into a single input file. At the low energies considered in this work, the spectrum is nearly flat since the gamma interactions in the crystals are typically Compton scatters.

Figure 4.5 shows the lead decay chain energies and conversion probabilities considered in this work. The source of these types of interactions is plate-out of ^{222}Rn onto the surfaces of detectors and housings during fabrication [123]. This decay chain produces relatively low energy events that cannot penetrate deep into a detector. As a result, the lead decay chain primarily produces surface events and zero-charge events (events occurring near the sidewall of the detector). Since surface events and sidewall events behave differently in iZIP detectors, the lead decay chain events were separated into events from contamination of the detector faces (which produce surface events) and events from contamination of the detector housings (which produce zero charge/sidewall events).

This work is concerned with effects at low energy, so only events with total recoil energy between 0.5 and 15 keV were simulated. For the barium simulations, each detector simulated had a unique input file because the position of each detector relative to the sources is different. The lead decay chain input is only dependent upon the position of the detector within the stack, so separate files were used for the top, middle, and bottom detectors in a tower. In addition, the ratio of 1.1, 1.2, and 1.3 keV events was fixed for all detectors because a single input file was used for




Figure 4.5: Lead decay chain energies and conversion probabilities.

all activation line simulations. During post-simulation processing, the relative rates of the three L-shell lines were weighted to correspond to the measured rates in each Soudan detector.

4.1.2 Formatting DMC output for processing

After running the full DMC simulation, which includes the phonon and charge propagation simulation, the FET simulation, and the TES simulation, the readout pulses needed to be formatted for processing through the standard CDMS analysis algorithms. The DMC produces TES and FET pulses with somewhat arbitrary units. In order to use DMC pulses in the pulse simulation framework implemented for the low-mass WIMP search, we must calibrate them to the same energy scale as Soudan data. The pulses must also be resampled at the Soudan digitization rate. In addition, the TES pulses require further formatting: the DC component of the current must be subtracted off, and a relative calibration must be applied to the individual channels. The calibration of the DMC pulses was performed similar to the manner used for Soudan data as described in Section 2.2.5, though the DMC calibration is aided by the *a priori* knowledge of the true recoil energy and interaction type for each event.

FET calibration for electron recoils

Electron recoils which occur in the bulk of the crystal have a yield of 1, that is $E_{\text{charge}} = E_{\text{recoil}}$. Since the recoil energy is known *a priori*, we can fit a line to the

uncalibrated FET energy vs. the recoil energy for electron recoils to determine the calibration constants. This is performed in a manner outlined by the following figures.

First, a least-squares fit to the optimal filter template is performed for each charge channel (QIS1, QOS1, QIS2, QOS2) to determine the uncalibrated FET energy for an event. We then plot the side-summed charge energy ($QIS\# + QOS\#$) versus the true recoil energy, as shown in Figure 4.6. There is a highly-populated band of bulk events, but also a “haze” of events which have reduced charge collection. Before fitting this data to determine the calibration constants, we remove all events with input position under the outer channel or within 2 mm of the faces of the detector. We also remove all events with uncalibrated FET energy near zero. This results in the green points seen in Figure 4.6.




Figure 4.6: Uncalibrated FET energy versus input recoil energy for side 1 (left) and side 2 (right). Blue points show all simulations run while green points show data after removing events which we *a priori* expect to have poor charge collection.

After removing events which we *a priori* expect to have poor charge collection, the remaining points are binned into 1 keV bins. The data in each bin was then fit to a gaussian to determine the mean and standard deviation, shown by the red points in Figure 4.7. The mean of each bin was then fit to a line which passes through zero, resulting in the black line, which is used as the final calibration constant for that side.

The calibration constants for side 1 (electrons) and side 2 (holes) as determined by the binned fit procedure are then used to determine the charge energy of the data points for each side on an event-by-event basis. This preserves the charge trapping




Figure 4.7: Uncalibrated FET energy versus input recoil energy with calibration fits for side 1 (left) and side 2 (right). Data is binned into 1 keV bins and each bin is fit to a gaussian, resulting in the red points. The error bars show the standard deviation for each bin. The means of all the gaussians are fit to a line, producing the black result.

inherent to the DMC simulation. The resulting calibrated energy of the events is then used in determining the total phonon energy.

FET calibration for other recoil types

Since the FET calibration procedure for electron recoils preserves the calibration constants for later use, the FET calibration procedure for events other than electron recoils uses the pre-determined constants. This preserves the charge trapping and nuclear recoil energy scale as implemented in the DMC. As with the electron recoil case, the calibrated energy of the events is used in determining the total phonon energy.

For this work, the barium simulations were used to determine the calibration constants on a detector-by-detector basis. These calibration constants were then applied to the L-shell activation line data and the lead decay chain data to maintain a consistent energy scale across all recoil types.

TES pulse formatting and calibration

The TES pulses must also be formatted and calibrated. Again, the DMC calibration procedure follows the basic logic of the Soudan calibration outlined in Section 2.2.5. Figure 4.8 shows the steps of the formatting procedure used to produce pulses ready to be calibrated. Raw TES simulation output is shown in Fig. 4.8. The raw pulses are downward-going, with a non-zero DC baseline. CDMS fitting algorithms assume that there is no DC component to the pulse, so it must be subtracted off. Subtracting off the DC component and inverting the pulse produces the figure in the center. As in Soudan data, we expect the late-time tails of the pulses to have equal power, so a relative calibration is applied. This is performed by integrating the tail of the pulse and scaling all eight channels to the reference channel, again on an event-by-event basis.




Figure 4.8: DC subtraction and relative calibration for simulated TES pulses.

After the DC component is subtracted from the individual TES pulses and the relative calibration is applied, we construct the total phonon pulse by summing all eight channels. This pulse is then calibrated on an event-by-event basis by determining the height of the uncalibrated pulse and multiplying by the scale factor necessary to produce the expected energy. We scale to the total phonon energy, which is the sum of the recoil energy and the energy of the Luke phonons produced as charges propagate across the crystal:

$$P_{total} = E_{recoil} + E_{Luke} = E_{recoil} + \frac{V_b}{\epsilon} E_{charge}$$

The recoil energy is given by the DMC input file, and the bias voltage is 4V. The charge energy is calculated as above, though there are two choices for the total energy: the maximum of the two sides, or the mean of the two sides. The mean charge energy for the two sides provides a better estimate of the Luke energy for surface events, where the effective voltage seen by the charges is half of the total bias across the crystal. Therefore, the mean charge energy is used to calculate the total phonon energy.

4.1.3 Pulse simulation method and data processing

After formatting the pulses output by the simulation, they were processed through `cdmsbats` using the infrastructure developed for the background model discussed in Chapter 3. The calibrated DMC pulses were summed with random traces used to characterize the Soudan noise environment to model the expected resolution effects on a detector-by-detector basis. Then, the DMC pulse with added noise is run through the standard CDMS reconstruction algorithms.

For the purposes of this study, the noise used for the pulse simulation was sampled from a short range of dates during the first half of June 2013. A short range of dates was chosen to minimize systematics caused by changes to the noise environment. However, future work should use noise sampled from the full data-taking period to understand the systematic variation with time and to provide a better comparison between DMC and the data-driven pulse simulation discussed in Chapter 3.

Because of the differences in phonon pulse shape between the detector monte carlo and Soudan, new DMC-specific phonon energy and partition quantities that use the DMC phonon template in Fig. 4.4 were implemented in the CDMS data-processing code. Since the DMC phonon templates are a better fit to the DMC pulse shapes, the new phonon energy and partition quantities provide better estimates of the true energy. This effect is demonstrated in Fig. 4.9, which shows a histogram of the total phonon energy for T1Z1 as determined by fits to the Soudan phonon template (cyan) and DMC phonon template (blue). Using Eqn. 2.1 with yield of 1 and voltage bias of 4V, we expect the peak of the distribution to occur around $(1 + 4/3) \times 1.3 \simeq 3$ keV for

electron recoils with good energy collection. The blue histogram is centered around this value, while the cyan histogram is centered at a lower value because the peakiness of the DMC pulses is underestimated by the Soudan phonon template. As a result, the phonon energy and partition estimators determined by fitting to the DMC phonon template were used throughout this study.




Figure 4.9: Histogram of total phonon energy for T1Z1 L-shell activation line data. The cyan histogram shows the total phonon energy estimated using the Soudan template, while the blue histogram shows the total phonon energy estimated using the DMC template. The Soudan template energy is biased low because it does not properly fit the peakiness of the DMC pulses.

4.2 Validation

4.2.1 Rescaling phonon partitions

Because the pulse shapes produced by the DMC simulation do not exactly match the pulse shapes in Soudan detectors, the distribution of phonon radial and z-partitions in the simulation did not exactly match the expected Soudan distribution. Rather than fine-tune the TES simulation to exactly match Soudan behavior, the processed

data was rescaled to match the Soudan distribution using the function

$$f(\text{data}, \vec{\alpha}) = (\text{data} + \alpha_0) \alpha_1 + \alpha_2. \quad (4.1)$$

Figure 4.10 shows the effects of this rescaling for detector T2Z1. Maximum likelihood estimation is used to determine the best fit rescaling parameters $\vec{\alpha}$. The best fit parameters are those that scale the raw DMC data (blue) to most closely match the expected Soudan distribution (black). The resulting rescaled data (red) more closely matches the expected Soudan distribution. Comparing the relative p-values from two-sample KS tests between Soudan data and the unscaled or scaled DMC distribution (see title of figure) confirms that the rescaled distributions better match the expected distribution.




Figure 4.10: Raw DMC data (blue) and rescaled DMC data (red) for detector T2Z1, compared to the expected Soudan distribution (black) for phonon r-partition (left) and phonon z-partition (right).

However, some residual differences remain between the simulated and Soudan data. These are clearest in the tails of the z-partition distribution and the valley between the two peaks in the radial partition distribution. These differences may be an effect of some of the approximations made in the DMC, as well as small differences in the event selection criteria applied to the DMC and Soudan data. In addition, the two sharp peaks in the radial partition distribution are clearly separated in the truth radius of the event: the first peak is from events that occur under the inner three

channels, while the second (larger r-partition) peak is from events that occur under the outer channel.

4.2.2 The bulldozer algorithm—modeling local saturation in the TES simulation

As shown in Fig. 4.10, even after rescaling the phonon radial partitions, some residual differences remain between the DMC and Soudan distributions. We believe that these differences are caused by the lack of local saturation in the TES simulation: in the simulation, the temperatures of all the TESs in a channel rise and fall together, while in the actual detectors, some TESs very near the interaction site may have higher temperatures than those farther away, causing the TESs near the interaction to saturate. Without local saturation, the phonon pulses tend to be peakier, especially near the site of the interaction. As a result, the partition quantities become sharper when local saturation is not modeled.

A new algorithm called the bulldozer was developed to model this effect. This algorithm works by dividing the face of the detector into bins, looking at the number of phonons that arrive in a spatial bin during a time step, and if the number of phonons exceeds some threshold, moving the phonons above the threshold to the next time bin. This has the effect of “bulldozing” off the top of the phonon arrival time histogram, as shown in Fig. 4.11. The new phonon arrival times are then saved into a file that can be read into the TES simulation.

This algorithm was run on a small subset of simulated barium calibration data for detector T1Z1. Phonon templates were generated using the resulting simulated TES pulses using the method depicted in Fig. 4.4. The resulting template is shown in Fig. 4.12 alongside the un-bulldozed DMC template and the Soudan phonon template for detector T1Z1. The bulldozed template is in much better agreement with the Soudan template than the un-bulldozed DMC template, and any remaining differences will likely be washed out by the poor signal-to-noise at the low energies considered in this work.




Figure 4.11: Schematic of the “bulldozer” algorithm for modeling local saturation. Any phonons collected in a TES that push the histogram of phonon arrival times above the saturation threshold (grey line) in a given time bin are bulldozed to the next time bin. [figure: Peter Redl]

Figure 4.13 shows the resulting phonon r-partition distribution (the ratio of outer-channel energy to total energy) from applying the bulldozer algorithm to barium data from T1Z1. The bulldozed data (right) does not show the extreme double-peaked structure of the un-bulldozed data (left). However, it appears that the local saturation is over-estimated with the current thresholds in the bulldozer algorithm, as there is even less double-peaked structure in the bulldozed DMC data than the true Soudan data. As a result, the rescaling code discussed in section 4.2.1 does not produce an optimal fit, because it is attempting to fit a single-peaked distribution to a double-peaked distribution.

Though there are problems with the agreement between Soudan data and DMC simulation that require further tuning, the bulldozer algorithm does correctly model local TES saturation on a qualitative level. This algorithm can be compared against the standard DMC data to understand the remaining issues with the simulation. However, the bulldozer is somewhat cumbersome because it is I/O limited, so it was only run on a subset of the seven low-threshold detectors. Two fully-functional detectors, T1Z1, and T2Z1, intended as controls, as well as T5Z3 were run through the




Figure 4.12: Total phonon energy templates generated on standard DMC data (magenta) and DMC data following the bulldozer algorithm (red). The bulldozer algorithm produces pulses that more closely match the Soudan phonon template (black). [figure: Peter Redl]

bulldozer. Analysis of these events shows that the bulldozer algorithm qualitatively improves the phonon partitions, but more tuning is needed to bring the DMC and Soudan data into full agreement.

4.2.3 Assessing agreement between Soudan and DMC barium data

Likelihood tests

Likelihood tests with no free parameters were used to assess the agreement between Soudan data and DMC data for each detector. Since the bulldozer local saturation algorithm was only run on a subset of the detectors, these tests used the simulations that do not include local saturation of the TESs. The likelihood was constructed by placing the data for each phonon partition quantity into bins and assuming the differences between the test data and the comparison data were Gaussian-distributed.




Figure 4.13: Raw DMC data (blue) and rescaled DMC data (red) for detector T1Z1, compared to the expected Soudan distribution (black) for phonon r-partition for standard DMC data with no local saturation (left) and DMC data after application of the bulldozer algorithm (right).

Thus, the likelihood that data from a Soudan detector agrees with the DMC data is

$$\mathcal{L} \propto \prod_{i=1}^{i=n_{bins}} \exp \left[-\frac{(X_i^{Soudan} - X_i^{DMC})^2}{2} \right]. \quad (4.2)$$

For each detector modeled in the DMC, the likelihood was calculated for all seven Soudan low-threshold detectors. If the DMC is a good model of a given Soudan detector, the likelihood should be larger when the DMC detector modeled matches the Soudan detector.

Figure 4.14 shows the results for comparisons between DMC data for detector T2Z1 and Soudan data from all low-threshold Soudan detectors. As expected, the likelihood is largest when the Soudan detector on the x-axis matches the DMC data. This indicates that the DMC data for T2Z1 is a better fit to Soudan detector T2Z1 than any of the other Soudan detectors.

One estimate of the systematic uncertainties at Soudan is the variation in behavior between detectors. To show that the DMC data is nearing this systematic limit, additional likelihood tests were performed to show that the agreement between DMC and Soudan data for a given detector is at least as good as the agreement between




Figure 4.14: Likelihoods comparing barium data from all Soudan low-threshold detectors to T2Z1 DMC barium data.

Soudan data for two different detectors. In this case, the likelihood is

$$\mathcal{L} \propto \prod_{i=1}^{i=n_{bins}} \exp \left[-\frac{(X_i^{test} - X_i^{Soudan})^2}{2} \right], \quad (4.3)$$

where the test data runs over all Soudan detectors, as well as unscaled and scaled DMC data. The results for this test for detector T2Z1 are shown in Fig. 4.15. Here, X^{Soudan} is the barium data for T2Z1, and X^{test} runs over all detectors as well as the DMC data. As expected, the test statistic should be best when X^{test} is T2Z2 Soudan data, that is, that $X^{test} = X^{Soudan}$. The effects of the rescaling discussed above are clear: the likelihood is significantly larger for the scaled DMC data than the unscaled DMC data. In addition, the fact that the agreement between DMC and Soudan data for a given detector is generally better than the agreement between two Soudan detectors indicates that the DMC will soon be limited by the systematic variation between detectors.

Tests using machine learning algorithms

As a second test of the agreement between Soudan and DMC data, a boosted decision tree was trained to recognize the differences between Soudan barium data and Soudan californium data using four variables: mean inner charge channel energy, total phonon




Figure 4.15: Likelihoods comparing barium data from all Soudan low-threshold detectors as well as unscaled and scaled DMC data for T2Z1 to T2Z1 Soudan barium data.

energy, phonon r-partition, and phonon z-partition. These are the same four variables used in the boosted decision tree that provided the final level of event selection in the low-mass WIMP search in Chapter 3. The californium data was re-weighted to match the expected total phonon energy spectrum for 5, 7, 10, and 15 GeV/c^2 WIMPs. The boosted decision trees (one for each detector and WIMP mass) were then tested on Soudan data as well as DMC data, and the distributions of the scores were compared.

Figure 4.16 shows the BDT scores on the test data for T2Z1, with the Cf signal data re-weighted to match the spectrum of a 5 GeV/c^2 WIMP (left) or 15 GeV/c^2 WIMP (right). The signal data (re-weighted californium data from Soudan, black) is reasonably well-separated from the background data. In addition, the distribution of BDT scores for the two types of background data (Soudan barium, red, and DMC barium, blue) are in close agreement, indicating that the machine learning algorithm is picking up the same structures in the DMC data as the Soudan data. This serves as further evidence that the DMC is overall a good model for Soudan data.




Figure 4.16: Results for a BDT trained on Soudan data and tested on both Soudan (red) and DMC data (blue) for detector T2Z1. The total phonon energy spectrum for the californium data was re-weighted to match the expected spectrum of a $5 \text{ GeV}/c^2$ WIMP (left) or $15 \text{ GeV}/c^2$ WIMP (right)

4.3 Comparing shorted detector T5Z3 to fully-functional detectors

4.3.1 Qualitative effects of the short

The large-scale simulations show many of the behaviors expected from the example event shown in Fig. 4.1. Studying the charge propagation as a function of the position of the event within a crystal demonstrates that the short has a clear effect on charge propagation, especially underneath the channel with the short.

Figures 4.17 and 4.18 show the truth yield (mean inner charge channel energy divided by truth recoil energy) and charge partitions for a fully-functional detector and T5Z3. These figures were produced using simulated barium data. Since a typical barium calibration event scatters multiple times within a crystal, the average position, weighted by energy of each scatter, is used as the truth position of the event. In addition, events in these plots were required to have a total phonon energy after event reconstruction between 2 and 13.1 keV total phonon energy, to coincide with the energy range considered for the low-threshold analysis in Chapter 3.

Comparing the fully functional detector T2Z1 to the shorted detector T5Z3, two




Figure 4.17: Truth yield (mean inner charge channel energy divided by truth recoil energy) versus position for fully functional detector T2Z1 (left) and shorted detector T5Z3 (right) for simulated barium calibration data between 2 and 13.1 keV total phonon energy.

effects can be seen in the shorted region (upper right-hand corner of the T5Z3 plot): in the yield plot, there is a population of events with yield ~ 0.5 , and in the partition plots, the topology of the event as indicated by the partitions is affected. These low-yield, improperly partitioned events may leak past the preselection cuts used for the low-threshold analysis in several ways. First, the reduced yield may move these events into the nuclear recoil band. Second, the effect on the radial partitions may cause the selection criteria to fail to reject events that occur at high radius (including the problematic low-yield events). Third, many events in the shorted region appear as side 1 surface events (charge z partition ~ 1) since the holes do not see a bias voltage to pull them to the opposite face.

To quantify these effects, Fig 4.19 shows the fraction of simulated barium events that pass the “preselection” cuts discussed in Chapter 3 for fully-functional detector T2Z1 and shorted detector T5Z3. The pass fraction is defined as the number of high-quality physics events that pass a given cut, divided by the total number of high-quality physics events. As the above figures suggest, more events pass the radial fiducial volume criteria (blue) and the 3σ nuclear recoil band cut (green) for T5Z3 than for T2Z1 because of the effects of the short. However, many of the high-radius events that fail to be rejected by the high-radius selection criteria for T5Z3




Figure 4.18: Charge radial partition for side 2 (holes) versus position (top) and charge z-partition versus position (bottom) for fully functional detector T2Z1 (left) and shorted detector T5Z3 (right) for simulated barium calibration data between 2 and 13.1 keV total phonon energy. The radial partition for side 1 is ill-defined for T5Z3 due to the outer charge channel short.

are removed by the charge symmetry requirement (cyan), as expected from Fig. 4.18, which shows that many of the events affected by the short have charge partition consistent with a side 1 surface event. As a result, the overall pass fraction for all preselection cuts (black) for T2Z1 and T5Z3 are similar.

The phonon signal contains contributions from Luke phonons, which are related to charge propagation, so the effect of the short can also be seen in the phonon partitions. Because the electric field lines near the short pull charges that would typically propagate to the outer channel into the inner channel, the Luke phonons generated by the charge propagation are not as localized under the outer channel. The effects are particularly strong in the DMC data considered here, which does




Figure 4.19: Fraction of simulated barium events that pass the charge fiducial volume and nuclear recoil band cuts (the “preselection” cuts discussed in Chapter 3) for fully-functional T2Z1 (left) and shorted T5Z3 (right).

not include local saturation in the TESs. As a result, the separation between the two peaks in the phonon radial partition for the unscaled DMC data is less well-defined for T5Z3. This can be seen clearly in Fig. 4.20, which shows the rescaled phonon radial partition versus truth position for a fully-functional detector and T5Z3. Because the range of the phonon radial partition varies between detectors, here the values are normalized by subtracting the mean of the distribution and dividing by the standard deviation. The demarcation in phonon partition between the outer and inner channels is less pronounced in T5Z3, indicating that the phonon radial partition may have poorer discrimination power for T5Z3 than for a fully-functional detector.

4.3.2 Comparing the DMC to the SuperCDMS Soudan background model

The agreement between barium calibration data simulated in the detector monte carlo and real barium data from Soudan is generally good. However, we would also like to test the agreement between the DMC and the SuperCDMS Soudan low-threshold background model, since the algorithms used to select candidate events for the low-mass WIMP search were trained on the data-driven model, not real Soudan data.




Figure 4.20: Phonon radial partition (rescaled and normalized) versus truth position for barium data from fully-functional detector T2Z1 (left) and shorted detector T5Z3 (right).

Figure 4.21 shows the phonon partitions (without the bulldozer for the DMC) and charge energy for T2Z1, comparing the data-driven low-energy gammas to detector monte carlo barium data. (The DMC barium data is a good proxy for low-energy gammas since the primary interaction process at those energies is Compton scattering, so the recoil energy spectrum is flat.) These plots demonstrate that the rescaling parameters computed on Soudan barium data can also bring the DMC into close agreement with the data-driven simulation. There are some small disagreements in the tails of the phonon partition distributions, as well as the double-peaked structure of the radial partition histogram, but these are to be expected given the differences in the DMC and data-driven models, as well as the lack of local saturation in the DMC.

In addition, the charge energy (Fig. 4.21, right) shows the expected behavior. There is a large peak at zero keV in both distributions that corresponds to low-energy events and events that occur at high radius (zero-charge events). The spectrum of bulk electron recoils above a few keV is also in good agreement. There is a slight discrepancy between the two histograms (which is also present in comparisons between the DMC and Soudan barium calibration data) that may be caused by small differences between low-energy charge trapping in the DMC and in real data.

The agreement between DMC and data-driven simulations is generally good for




Figure 4.21: Histograms of phonon radial partition (left), phonon z-partition (center), and charge energy (right) for low-energy gammas for the DMC (unscaled partitions: blue, rescaled partitions: cyan) and data-driven (magenta) models for detector T1Z1.

the energy estimators and phonon partitions, especially for low-energy gammas, activation lines, and ^{210}Pb decay chain sidewall events. Any small differences between the two simulations for these event types are due to differences in the assumptions about the input events for the simulation. However, as shown in Fig. 4.22, there are larger differences between the two simulations for ^{210}Pb decay chain surface events, especially the ^{210}Bi and ^{210}Pb decays, which contain a population of electron recoils. For this event type, the DMC produces events with very small charge energies; however, the data-driven simulation also contains a population of events with yield closer to the value expected for bulk electron recoils. Because these comparisons only point out discrepancies and do not explain their underlying causes, further work comparing simulated events to Soudan data is needed to understand the underlying issue. It may be that the DMC does not properly model the charge physics for certain types of surface events or that there is a problem with the electric field model very near the surface. Otherwise, if the problem is in the data-driven simulation, it is possible that the selection of template events is biased if the behavior of high-energy ^{210}Bi and ^{210}Pb events used as templates is very different from true low-energy ^{210}Bi and ^{210}Pb decays.

As another test of the agreement between the low-threshold data-driven background model and the detector monte carlo, we apply the boosted decision tree used as the final level of event selection to the DMC and compare with the scores the algorithm assigns to the data-driven model. Figure 4.23 shows this comparison for ^{206}Pb




Figure 4.22: Charge energy versus total phonon energy for ^{210}Pb surface events in detector T2Z1. There is a qualitative disagreement between DMC and data-driven simulation in the yield for this event type.

sidewall events for detector T1Z1 (left) and T5Z3 (right). For these two detectors, the agreement between data-driven and DMC models is quite good.

However, the agreement between data-driven and DMC background models is less good for other detectors, such as T5Z2 shown in Fig. 4.24. For this detector, and several others including the tower 2 detectors, the histogram of BDT scores is generally peaked at a larger value for DMC than for the data-driven background model. This indicates that there is a disagreement between the underlying data for the DMC and the data-driven model. However, it is not clear if the discrepancy is caused by mis-modeling in the DMC or the data-driven model. It may be possible to better understand these issues by re-training the boosted decision tree and more closely examine the behavior of the test and training samples.

4.3.3 Hypothesis for source of T5Z3 events in the low-threshold analysis

The observed systematic effect on T5Z3 is on the order of a few percent: three events passed the boosted decision tree cut out of approximately 100 in the preselection




Figure 4.23: Histograms of boosted decision tree scores trained against a 15 GeV/ c^2 WIMP for ^{206}Pb sidewall events for detector T1Z1 (left) and T5Z3 (right). The data-driven background model is shown in magenta and the DMC is shown in blue.

region. The detector monte carlo provides an excellent way to study this small effect, since the statistics that we are able to generate grow with CPU time, not integration time at Soudan. We can also study the effect of the noise environment on a given event type by processing the same DMC event multiple times using different samples of the noise.

The current implementation of the DMC pulse simulation works by recycling DMC events: a given DMC event is reused multiple times in the simulation with different noise samples added onto the simulated pulse. As a result, a single DMC event will have a range of reconstructed energies, partitions, and BDT scores after processing. The range of values is caused by the differences in the noise sample added to each event.

To understand how the noise environment affects the reconstruction parameters for an event, we compute the mean and standard deviation of the boosted decision tree scores for each unique DMC event. The mean scores for T1Z1 and T5Z3 ^{206}Pb sidewall events are shown versus truth z-position in Fig. 4.25. (These two detectors were considered because the overall distribution of scores agrees well with the data-driven model, as shown in Fig. 4.23.)

Since the three events on T5Z3 passed only the 15 GeV/ c^2 BDT selection criteria,




Figure 4.24: Histograms of boosted decision tree scores trained against a $5 \text{ GeV}/c^2$ WIMP (left) or $7 \text{ GeV}/c^2$ WIMP (right) for low-background gamma events for detector T5Z2. The data-driven background model is shown in magenta and the DMC is shown in blue. The disagreement is large between DMC and data-driven models for this detector and several others, indicating that further validation work may be needed.

we are most interested in examining the BDT scores for that WIMP mass (black points). There are a few outliers for T1Z1, as expected from the histograms in Fig. 4.23. Many more outliers are present for T5Z3, especially at z position greater than $\sim 2 \text{ cm}$. The events in these figures all occur at high radius because they come from contamination of the detector housings, so these outlier events primarily occur in the shorted region for T5Z3. In addition, they occur at relatively high recoil energies ($\gtrsim 7 \text{ keVnr}$), so such events may be candidates for the three observed events on T5Z3 in the SuperCDMS Soudan low-mass WIMP search.

In addition, we can examine the standard deviations of the BDT scores for ^{206}Pb sidewall events. The standard deviation for each unique event for detectors T1Z1 and T5Z3 is shown versus z position in Fig. 4.26. There is a possible uptick in the standard deviation near the top surface of T5Z3, indicating that the noise environment may more strongly affect events in the shorted region, possibly causing a small percentage to fluctuate above the threshold.

Though this effect is suggestive, it is not sufficient to explain the three events observed on T5Z3. Currently the DMC statistics are not sufficient to confirm that




Figure 4.25: Mean BDT score versus truth z-position for recycled DMC events for T1Z1 (left) and T5Z3 (right) for ^{206}Pb sidewall events. T5Z3 has more high-score outliers near the top surface near the shorted region; an upward fluctuation in the noise may push one of these events above the BDT threshold and into the signal region.

the outliers mentioned here are significant; in addition, we would like to use more noise samples for each unique DMC event to confirm that the standard deviation also increases near the shorted region. Running more unique DMC events, as well as processing these DMC events with the full range of noise samples (October 2012–July 2013) will provide more statistics to try to confirm or reject this hypothesis.

4.4 Other insights from the DMC

Detector Monte Carlo simulations have also informed ongoing analysis issues aside from the shorted charge channels.

The Detector Monte Carlo is expected to play a large role in analyses of CDMSlite data moving forward. During CDMSlite Run 1, it was discovered that the activation line peaks were broader than expected. After modeling the electric field of the CDMSlite detector, it was realized that the electric field at high radius is smaller, and some electric field lines terminate on the sidewall of the crystal instead of the opposite face. As a result, events that occur at high radius may not see the full 70V bias. Such events would produce fewer Luke phonons due to the smaller bias voltage,




Figure 4.26: BDT score standard deviation versus truth z -position for recycled DMC events for T1Z1 (left) and T5Z3 (right) for ^{206}Pb sidewall events. There is a possible increase in the standard deviation of the scores very near the top surface of the detector ($z \sim 2.5$ cm); a larger standard deviation might lead to statistical fluctuations pushing a small number of events above the boosted decision tree threshold.

therefore appearing at a lower energy.

To test this effect, activation line events were simulated for a detector in CDMSlite configuration. As expected, the activation lines had tails at low energy. Upon further investigation, the events in the low-energy tail were found to occur at high radius. In addition, studies of the charge propagation in a CDMSlite-style detector showed that some events that occur at high radius do not propagate through the full 70V and therefore produce fewer Luke phonons. These studies showed that using position information in a CDMSlite crystal to develop a fiducial volume and remove high-radius events can improve the end result.

Another insight was gleaned from the DMC in the course of validating the L-shell activation lines against the data-driven pulse simulation used for the low-mass WIMP search in Chapter 3. While performing preliminary fits for a likelihood analysis of all data in the preselection region (before application of the boosted decision tree), an excess of very low-energy events was discovered in T2Z2. However, it was noticed that the distribution of total phonon energies for the DMC L-shell simulated events extended to lower total phonon energies, which may partially account for this excess as described below.

The template events for the data-driven background model were selected to be a delta function in total phonon energy centered at 3 keV. Selecting events in this way effectively only chooses bulk events: the reduced yield for surface events produces a smaller value of total phonon energy. As a result, the total phonon energy histogram is symmetric about the mean as shown in black in figure 4.27, since the only broadening is due to noise resolution. However, the DMC data (blue), which includes both surface and bulk events, is asymmetric, with a tail extending to lower energies due to reduced charge collection. When the selection of L-shell events for the data-driven model was modified to include surface events, the resulting histogram (pink) is in much better agreement with the DMC, and also yields a better fit to the low-energy excess seen in the likelihood analysis of T2Z2.




Figure 4.27: Comparison between DMC 1 keV line events (blue) and data-driven 1 keV line events (magenta, cyan) with two different weighting schemes for detector T2Z2. The updated weighting scheme (magenta), which includes surface events, matches the DMC better and can partially account for an excess of very low energy events seen in a likelihood analysis of the data from the low-mass WIMP search of Chapter 3.

Chapter 5

Effective field theory scattering in direct detection experiments

Astrophysical and cosmological evidence indicates that the majority of the matter in the universe takes the form of non-luminous particles called dark matter, though the exact nature of the dark matter particle remains unknown [32]. A generic weakly-interacting massive particle (WIMP) is a very attractive dark matter candidate [124]. Numerous experiments are engaged in efforts to detect rare collisions between WIMPs and target nuclei in terrestrial detectors. Results from DAMA [44], CoGeNT [45], CRESST-II [46], and CDMS II Si [47] can be interpreted in terms of interactions of WIMPs with masses of 6-30 GeV/c^2 . A similar range of masses could also account for a possible excess in the gamma-ray flux near the galactic center in Fermi-LAT data [125, 126]. Under standard assumptions for spin-independent WIMP-nucleon interactions, however, such interpretations are difficult to reconcile with the limits set by CDMSlite [103], SuperCDMS [104], LUX [85], and PICO [127].

Standard WIMP scattering calculations make simplifying assumptions about the type of interaction between the nucleon and the dark matter particle: typically only isospin-conserving spin-independent couplings, or spin-dependent couplings to either the proton or neutron are considered. This results in constraints on the three corresponding WIMP-nucleon cross sections. Relaxing such assumptions can suppress the interaction for some target elements by orders of magnitude relative to others [128].

However, such solutions often require a high degree of fine-tuning.

Recently, an effective field theory (EFT) approach for WIMP scattering has been developed that considers all leading-order and next-to-leading order operators that can occur in the effective Lagrangian that describes the WIMP-nucleus interaction [129, 130, 131]. This formalism introduces new operators that rely on a range of nuclear properties in addition to the standard spin-independent and spin-dependent cases. It also explicitly includes isospin interference and interference between operators, creating a rich parameter space of possible dark matter interactions that are very sensitive to the specific choice of detector material.

The EFT framework parametrizes the WIMP-nucleus interaction in terms of fourteen operators, $\mathcal{O}_i$, which are listed in Eq. 5.1 and include the standard spin-independent and spin-dependent interactions. These operators feature explicit dependence on $\vec{v}^\perp$ (the relative velocity between the incoming WIMP and the nucleon) and the momentum transfer $\vec{q}$, in addition to the WIMP and nucleon spins, $\vec{S}_\chi$ and $\vec{S}_N$. Note that $\mathcal{O}_2$ is not considered since it cannot arise from the non-relativistic limit of a relativistic operator at leading order. In addition, each operator can independently couple to protons or neutrons. We formulate this isospin dependence in terms of isoscalar and isovector interactions, following the conventions of [130].

$$\begin{aligned}
\mathcal{O}_1 &= 1_\chi 1_N \\
\mathcal{O}_3 &= i\vec{S}_N \cdot \left[\frac{\vec{q}}{m_N} \times \vec{v}^\perp \right] \\
\mathcal{O}_4 &= \vec{S}_\chi \cdot \vec{S}_N \\
\mathcal{O}_5 &= i\vec{S}_\chi \cdot \left[\frac{\vec{q}}{m_N} \times \vec{v}^\perp \right] \\
\mathcal{O}_6 &= \left[\vec{S}_\chi \cdot \frac{\vec{q}}{m_N} \right] \left[\vec{S}_N \cdot \frac{\vec{q}}{m_N} \right] \\
\mathcal{O}_7 &= \vec{S}_N \cdot \vec{v}^\perp \\
\mathcal{O}_8 &= \vec{S}_\chi \cdot \vec{v}^\perp \\
\mathcal{O}_9 &= i\vec{S}_\chi \cdot \left[\vec{S}_N \times \frac{\vec{q}}{m_N} \right] \\
\mathcal{O}_{10} &= i\vec{S}_N \cdot \frac{\vec{q}}{m_N} \\
\mathcal{O}_{11} &= i\vec{S}_\chi \cdot \frac{\vec{q}}{m_N} \\
\mathcal{O}_{12} &= \vec{S}_\chi \cdot \left[\vec{S}_N \times \vec{v}^\perp \right] \\
\mathcal{O}_{13} &= i \left[\vec{S}_\chi \cdot \vec{v}^\perp \right] \left[\vec{S}_N \cdot \frac{\vec{q}}{m_N} \right] \\
\mathcal{O}_{14} &= i \left[\vec{S}_\chi \cdot \frac{\vec{q}}{m_N} \right] \left[\vec{S}_N \cdot \vec{v}^\perp \right] \\
\mathcal{O}_{15} &= - \left[\vec{S}_\chi \cdot \frac{\vec{q}}{m_N} \right] \left[\left(\vec{S}_N \times \vec{v}^\perp \right) \cdot \frac{\vec{q}}{m_N} \right]
\end{aligned} \tag{5.1}$$

These operators contribute to six types of nuclear response functions. The spin-independent response is denoted by M and is typically the strongest of the six functions since it is related to the number of nucleons in the target nucleus. The main contribution to this response comes from the standard spin-independent operator $\mathcal{O}_1$, but it also contains higher-order contributions from operators $\mathcal{O}_5$, $\mathcal{O}_8$, and $\mathcal{O}_{11}$. There are two spin-dependent responses, Σ' and Σ'' , that correspond to projections of spin parallel and perpendicular to the momentum transfer. A linear combination of these two responses yields $\mathcal{O}_4$, which is related to the standard spin-dependent response.

Many of the other operators also appear in one of these two responses. A novel type of response introduced in the EFT, Δ , is related to the net angular momentum of an unpaired nucleon and contains contributions from operators $\mathcal{O}_5$ and $\mathcal{O}_8$. A second novel response is Φ'' , which is sensitive to the product of angular momentum and spin. This response tends to favor heavier elements, and the most dominant contribution to this response is from $\mathcal{O}_3$. The last response considered in the EFT, $\tilde{\Phi}'$, contains contributions from operators $\mathcal{O}_3$, $\mathcal{O}_{12}$, and $\mathcal{O}_{15}$. $\tilde{\Phi}'$ is discussed less frequently in the literature since it is difficult to find a model that produces this response, but we consider it here for completeness.

The EFT also includes two operator-operator interference terms: $\Sigma'\Delta$ and $M\Phi''$. Σ' interferes with Δ because velocity-dependent responses are sensitive to properties such as angular momentum that depend on the motion of the nucleon within the nucleus. This interference term is particularly significant for germanium, which has large responses to both Σ' and Δ . The $\Sigma'\Delta$ response contains interference between $\mathcal{O}_4$ and $\mathcal{O}_5$, as well as between $\mathcal{O}_8$ and $\mathcal{O}_9$. In addition, since both M and Φ'' are scalar responses, interference between the two can be significant, especially for elements like xenon that have large responses to both. The $M\Phi''$ response contains interference between operators $\mathcal{O}_1$ and $\mathcal{O}_3$, operators $\mathcal{O}_{11}$ and $\mathcal{O}_{12}$, and operators $\mathcal{O}_{11}$ and $\mathcal{O}_{15}$.

Since the various responses are related to different nuclear properties, the strength of the resulting interaction can vary by many orders of magnitude. The expectation values of these properties are listed in [129]. For instance, the spin-dependent responses Σ' and Σ'' depend on the square of the spin of an unpaired nucleon, which ranges from 5×10^{-6} for protons in germanium (which has one isotope with an unpaired nucleon, which is a neutron) to 0.2 for protons in fluorine (which has an unpaired proton). The angular momentum of a nucleon, which governs the strength of the Δ response, ranges from $\mathcal{O}(1 \times 10^{-3})$ to $\mathcal{O}(1)$, while $(L \cdot S)^2$, which governs the strength of the Φ'' response, ranges from 0.1 for light nuclei to several hundred for heavier nuclei. The strongest response is M , which is related to the square of the number of nucleons.

The strength of an EFT interaction is parametrized by numerical coefficients, c_i^τ , associated with each operator $\mathcal{O}_i$, where $\tau = 0$ or 1 denotes the isoscalar ($c_i^0 = 1/2(c_i^p +$

c_i^n)) and isovector ($c_i^1 = 1/2(c_i^p - c_i^n)$) combinations, respectively. The coefficients have dimensions of $1/\text{energy}^2$, so we multiply by the weak mass scale ($m_{weak} = 246.2$ GeV) to produce dimensionless quantities. The c_i^τ are related by a change of basis to generalized versions of f_n and f_p and can take on any value, positive or negative. The coefficients appear as $c_i^\tau c_j^{\tau'}$ in the interaction, indicating that operators interfere pair-wise, at most.

The event rate for an arbitrary EFT interaction follows from Eqn. 1.4. The differential cross section $\frac{d\sigma}{dE_R}$ can be computed by replacing the Helm form factor with one of the nuclear responses discussed above. The numerical coefficients c_i^τ are incorporated in the WIMP response, which is the effective field theory analogue of Equation 1.9. Unlike the spin-independent case, where the form factor and WIMP response are nearly constant over the energy range of interest (approximately 1-100 keVnr), the differential cross section for a general EFT interaction can have strong dependence on momentum transfer, especially for operators such as $\mathcal{O}_3$ and $\mathcal{O}_{11}$, as shown in Figure 5.1.




Figure 5.1: Differential event rate for a 50 GeV WIMP interacting in germanium as a function of energy for several EFT interactions. Although the interactions have vastly different strengths for a given set of coefficients, here the coefficients are chosen so that the integrated event rate is equal for the four operators shown.

5.1 Bias created by standard dark matter assumptions

Because of the additional momentum dependence of several of the EFT operators, the differential event rate for an arbitrary dark matter interaction could be very different than for the standard calculation. Consequently, it is possible that a limit-setting algorithm that expects the (approximately) exponential event rate of the standard spin-independent interaction could misinterpret a potential signal from a more general EFT interaction as background.

To demonstrate the possible bias that could arise from assuming the standard spin-independent event rate when setting limits, we perform simulated experiments assuming that the dark matter scattering is purely due to a single isoscalar EFT operator. Figure 5.2 shows the co-added results of 100 simulated experiments sampled from the energy spectrum of isoscalar $\mathcal{O}_3$ scattering in germanium for two different dark matter masses, assuming an energy-independent (or “flat”) detection efficiency. The operator coefficients were set to give each simulated experiment an expectation value of 10 events. This expectation was then convolved with a Poisson distribution to select the number of events for a given simulated experiment.

Unlike the standard spin-independent event rate, the event rate for $\mathcal{O}_3$ depends on the square of the momentum transfer, so the event rate is suppressed at low recoil energies. This effect is illustrated in Fig. 5.2, where the black, cyan, and magenta curves show the standard spin-independent scattering rate for a range of WIMP masses and the blue histogram corresponds to the simulated spectrum expected from $\mathcal{O}_3$ interactions. For the case of a $10 \text{ GeV}/c^2$ WIMP mass, the distribution of events is more closely matched by the spin-independent rate for a higher-mass WIMP. For the $300 \text{ GeV}/c^2$ case, no spin-independent rate calculation matches the observed spectrum of events; if experimenters only consider the spin-independent WIMP rate, they may erroneously conclude that they have observed an unexpected background. Alternatively, such an observation could lead to the conclusion that detection efficiency as a function of energy was measured incorrectly: an observation of few events at low energies, where the spin-independent rate exponentially increases, could also be




Figure 5.2: Co-added energy spectrum from 100 simulated experiments (blue histogram) assuming the dark matter interaction proceeds according to the isoscalar $\mathcal{O}_3$ operator for a $10 \text{ GeV}/c^2$ (left) and a $300 \text{ GeV}/c^2$ WIMP (right). The detection efficiency is assumed to be independent of energy. The smooth cyan, magenta, and black curves show the expected spectrum for the standard spin-independent rate for several WIMP masses, while the dashed dark blue curve shows the $\mathcal{O}_3$ spectrum from which the simulated experiments were sampled.

attributed to the true detection efficiency being lower than the measured efficiency.

To further demonstrate this issue, we calculate the 90% confidence level upper limit on the spin-independent cross section for each simulated experiment sampled from the $\mathcal{O}_3$ spectrum using the optimum interval method [107, 108] and the standard Maxwellian halo model with halo parameters as above with no background subtraction. Each simulated experiment was assumed to have an exposure of 1000 kg days and a flat efficiency of 60% between 1 and $100 \text{ keV}_{\text{nr}}$. The distribution of limits on the spin-independent cross section for simulated experiments sampled from the $\mathcal{O}_3$ spectrum is shown in blue in Fig. 5.3. Figure 5.3 also shows the resulting median limit on the spin-independent cross section from simulated experiments sampled from the spin-independent distribution in black.

The distribution of limits on the spin-independent cross section for the simulated experiments sampled from the $\mathcal{O}_3$ energy spectrum deviates from the zero-background limit shown in magenta as well as from the mean limit derived from similar simulated experiments sampled from the spin-independent distribution. As expected, the simulated-experiment limits are weaker than the zero-background limits for both $\mathcal{O}_3$




Figure 5.3: Distribution of 90% confidence level upper limits calculated using the optimum interval method for the simulated experiments discussed above and shown in Fig. 5.2, sampled from the event rate for isoscalar $\mathcal{O}_3$. Shaded blue bands show the 68% and 95% confidence level uncertainty on the distribution. The zero-background Poisson limit is shown in magenta.

and spin-independent cases due to the presence of candidate events. However, because the energy distribution of the candidate events sampled from $\mathcal{O}_3$ is different than the expected spin-independent rate assumed in the limit-setting algorithm, the limits also deviate from the expected shape for the true spin-independent experiment. Here, this is an effect of the optimum interval method, which tries to match the observed event energies to the expected (spin-independent) energy spectrum and ignores energy ranges where the observed event rate does not match the expected spectrum. Other methods, such as likelihood techniques, which also assume a particular energy spectrum for the WIMP component, will display similar effects.

In the $10 \text{ GeV}/c^2$ case, we expect the limit to be weakest around a mass of $10 \text{ GeV}/c^2$, where the rate expected by the limit algorithm matches the observed event rate. However, because the observed events for $\mathcal{O}_3$ scattering are skewed towards higher recoil energies, the limit tends to be weaker at larger WIMP masses where the tail of the spin-independent event rate extends to higher recoil energies. In other words, a low-mass WIMP that scatters via $\mathcal{O}_3$ may look like a higher-mass WIMP because the energy spectrum skews towards higher values. For the $300 \text{ GeV}/c^2$ case, the distribution of limits agrees with the Poisson zero-background limit at low masses;

the observed events occur at recoil energies that cannot be produced by a low-mass WIMP. At higher masses, the distribution of limits is still close to the zero-background limit because the shape of the observed spectrum is very different from the expected spin-independent WIMP rate. For this high-mass case, experimenters may conclude that some or all of the 10 candidate events in these simulated experiments are backgrounds, since the spin-independent cross section that would produce 10 events is excluded in the majority of these simulated experiments.

The difference in the limits between the spin-independent and EFT cases demonstrates the importance of correctly modeling the expected WIMP signal. Algorithms that assume the standard spin-independent rate when calculating limits may interpret events from EFT interactions with different spectral shapes as background because of the mis-match in spectral shapes, and thus, this assumption could lead to a bias in the exclusion limits reported by experiments, especially in the case where events are observed.

5.2 Exclusion limits on a set of EFT operators

Because of the potential for bias demonstrated above, we would like to compute exclusion limits that take the differences in spectral shape for EFT interactions into account. The strength of the interaction in the EFT framework is governed by a set of 28 numerical coefficients corresponding to the 14 operators, one for each isospin. Others have attempted to find global fits in this multi-dimensional parameter space, combining data from many direct detection experiments [132]. Since the parameter space is large and relatively unconstrained by current experiments, we choose instead to calculate exclusion limits on the coefficients for individual EFT operators for three different target elements: germanium (SuperCDMS Soudan and CDMS II), silicon (CDMS II), and xenon (LUX).

We use the optimum interval method to calculate 90% upper confidence limits on the numerical coefficients of EFT operators. The optimum interval method incorporates information about the candidate event energies and energy-dependent detection efficiencies, which can yield stronger exclusion limits in the presence of unknown




Figure 5.4: Upper limits on the dimensionless isoscalar coefficients c_3^0 (left) and c_8^0 (right) as a function of WIMP mass for SuperCDMS Soudan (light blue) [104], CDMS II Ge reanalysis (dark blue) [133], and CDMS II Si (red) [134], and estimated limits for LUX (black) [85], for the Maxwellian halo (solid) and an alternate halo model (dashed). Limits on the remaining isoscalar operators can be found in Appendix B.

backgrounds than likelihood methods that consider only a single energy bin in the presence of backgrounds. This is particularly important here because of the spectral differences that can arise from different EFT interactions. We consider a single operator at a time and present the exclusion limit on the square of the EFT coefficient, which is proportional to the total interaction cross section. We compare the effects of two halo models on the limits. The first uses standard halo assumptions as in [74], with a WIMP mass density $\rho_0 = 0.3 \text{ GeV}/c^2/\text{cm}^3$, most probable WIMP velocity of 220 km/s, mean circular velocity of the Earth with respect to the galactic center of 232 km/s, galactic escape velocity of 544 km/s, and a velocity distribution that correctly takes into account the effect of the Earth's velocity on the escape-velocity cutoff [77]. The second halo model uses the functional form of Eq. 1.12 with $p = 2.7$ and $v_0/v_{esc} = 0.6875$, determined by fits to the Eris simulation of a Milky-Way-like galaxy [83], and other halo parameters as above.

Figure 5.4 shows the upper limits for two example operators, isoscalar $\mathcal{O}_3$ (left) and $\mathcal{O}_8$ (right), as a function of WIMP mass. Solid lines correspond to the Maxwellian

halo, whereas dashed lines show the limit calculated assuming the alternate velocity distribution function discussed above. The SuperCDMS Soudan, CDMSII Ge (re-analysis), and CDMSII Si limits use the candidate events, thresholds, and detection efficiencies discussed in [104], [133], and [134] respectively, while the estimated LUX limit assumes zero observed events and functional form for the detection efficiency that follows a hyperbolic tangent versus energy centered at $2.5 \text{ keV}_{\text{nr}}$ but with a step function cutoff that goes to zero below $3 \text{ keV}_{\text{nr}}$.

Because of the different nuclear responses for the three target elements considered, the relative strength of the limits varies from operator to operator. In particular, $\mathcal{O}_8$ (Fig. 5.4, right) includes contributions from the Δ response, which is greater in germanium than in silicon or xenon. This contribution strengthens the SuperCDMS Soudan constraint relative to LUX and CDMSII Si. In addition, the shape of the curve for a single target element changes from operator to operator. For example, $\mathcal{O}_3$ depends on the square of the momentum transfer, naturally suppressing the event rate at low energies. As a result, the limits at low WIMP mass for $\mathcal{O}_3$ are weaker than for other operators.

The difference between the two WIMP velocity distributions becomes apparent when the only events expected above the detection thresholds are due to WIMPs in the high-velocity tails. Since both CDMS and LUX have thresholds of a few keV, this disparity appears only at the lowest WIMP masses. The difference is also more pronounced for LUX, since its target nucleus, xenon, is heavier than silicon or germanium. A dark matter particle must have a higher velocity to deposit a given recoil energy in xenon than in germanium or silicon; higher-energy recoils become comparatively rarer. For the SuperCDMS Soudan result, the difference in velocity distributions leads to a factor of two difference in the limit around $4 \text{ GeV}/c^2$, whereas for LUX, the difference in velocity distribution leads to a factor of two difference around $7 \text{ GeV}/c^2$.

Since the EFT explicitly includes isospin dependence, we can also use the optimum interval method to set polar limits on isospin. For a given WIMP mass and a given angle between the isoscalar and isovector components of an operator, we set a 90% upper confidence limit on the isoscalar-isovector radius. Varying the polar




Figure 5.5: Polar limits on $\mathcal{O}_1$ isospin for SuperCDMS Soudan (blue) [104], LUX [85] (black), and CDMS II Si (red) [134] at a WIMP mass of $6 \text{ GeV}/c^2$.

angle produces exclusion ellipses in the isoscalar-isovector plane, as in Fig. 5.5, which shows limits for operator $\mathcal{O}_1$ and a $6 \text{ GeV}/c^2$ WIMP. The major axis of each ellipse corresponds to the value of c_1^0/c_1^1 that yields maximum suppression of the scattering rate. Note that although the exposures for CDMS and SuperCDMS are significantly lower than for LUX, there are regions of parameter space allowed by LUX but excluded by SuperCDMS and CDMS at 90% confidence. This example demonstrates that a combination of experiments using several target nuclei can constrain the EFT parameter space better than any single experiment.

5.3 Principal component analysis of EFT interference

The EFT framework also provides a more general description of interference among operators such as the “xenophobic” isospin violation case discussed in the literature [128]. It not only allows for interference between the isospin components of individual operators, but also among different operators. The generalized interference can be

written as a matrix equation in the large EFT parameter space, but because operators interfere in pairs, and only certain pairs interfere, this large matrix can be decomposed into block-diagonal form. We consider the 2×2 case of isospin interference and the 4×4 case of isospin and operator-operator interference.

The generalized amplitude for the 4×4 case can be written as the product of the vector of operator coefficients c_i^τ with the amplitude matrix, where superscript 0 and 1 indicate isoscalar and isovector, respectively, and the subscripts indicate the operator being considered: [135]

$$\begin{bmatrix} c_i^0 & c_i^1 & c_j^0 & c_j^1 \end{bmatrix} \begin{bmatrix} A_{ii}^{00} & A_{ii}^{01} & A_{ij}^{00} & A_{ij}^{01} \\ A_{ii}^{10} & A_{ii}^{11} & A_{ij}^{10} & A_{ij}^{11} \\ A_{ji}^{00} & A_{ji}^{01} & A_{jj}^{00} & A_{jj}^{01} \\ A_{ji}^{10} & A_{ji}^{11} & A_{jj}^{10} & A_{jj}^{11} \end{bmatrix} \begin{bmatrix} c_i^0 \\ c_i^1 \\ c_j^0 \\ c_j^1 \end{bmatrix} \quad (5.2)$$

The amplitudes $A_{ij}^{\tau\tau'}$ are the product of the WIMP and nuclear response functions for the interaction specified by c_i^τ and $c_j^{\tau'}$ and depend on properties such as target element, WIMP mass, WIMP spin, WIMP velocity, and nuclear recoil energy. We evaluate the $A_{ij}^{\tau\tau'}$ without integrating over the dark matter velocity distribution to avoid introducing more variables. Amplitudes are summed over the isotopes for a given element according to their natural abundances.

Finding the eigenvectors of this matrix will give the “principal components” of the interaction space. We expect that three of the four eigenvalues should be small, since the matrix for a single isotope is an outer product and therefore should have a single nonzero eigenvalue. The vector with the largest eigenvalue corresponds to the maximal amplitude for scattering in the interference space under consideration, while the three small eigenvalues correspond to local extrema in the scattering amplitude which tend to suppress the event rate. To be maximally sensitive to the parameter space for a given interference case, we would like to choose target elements whose constructive interference eigenvectors span the space of interactions.

As an example, we first consider isospin interference for a single operator in an already well-understood case. Figure 5.6 shows the constructive isospin interference eigenvectors for scattering via operator $\mathcal{O}_4$ (the spin of the WIMP interacting with the

spin of an unpaired nucleon) for several elemental targets, evaluated at a WIMP mass of $100 \text{ GeV}/c^2$ and nuclear recoil energy of 100 keV . The vectors are plotted in the space of the isoscalar coefficient versus the isovector coefficient. The proton-neutron space can be recovered from this basis via a 45-degree rotation. The amplitude in a given direction indicates the target's response to that operator and illustrates the sensitivity of each material to the corresponding operator. In addition, if we were to plot polar limits as in Fig. 5.5 for $\mathcal{O}_4$, we would see that the direction of the constructive interference vector corresponds to the minor axis of the ellipse. In the two-dimensional case, the destructive interference vector is perpendicular to the constructive vector and corresponds to the major axis of the ellipse in a polar limit plot.




Figure 5.6: Constructive interference eigenvectors for 2D $\mathcal{O}_4$ isospin interference. Proton-dominated interactions occur along the $x = y$ diagonal, while neutron-dominated interactions occur along the $x = -y$ diagonal.

We can apply this same procedure to the more general 4D case to demonstrate the complementarity of the different target elements. Figures 5.7 and 5.8, left, show all 2D projections of the four-dimensional eigenvectors in the interference space for $\mathcal{O}_4/\mathcal{O}_5$ (Fig. 5.7) and $\mathcal{O}_1/\mathcal{O}_3$ (Fig. 5.8), in addition to the two pure isoscalar operators for each interference case. All eigenvectors are normalized to unit length. The eigenvalues λ for the four principal components are shown in the legend, ordered from largest to smallest magnitude.

The total event rate is evaluated in germanium assuming parameters consistent with the expectation for SuperCDMS SNOLAB Ge iZIP: target mass of 50 kg, trigger threshold of 0.5 keVnr, 60% combined nuclear recoil and fiducial volume efficiency, and exposure of 1000 days. The event rate is then normalized relative to the stronger of the two pure isoscalar operators and plotted in Figures 5.7 and 5.8, right. Note that these plots show the event rate for a 10 GeV/ c^2 WIMP; the dependence of the relative rates on WIMP mass is sub-dominant.




Figure 5.7: Eigenvectors for 4D $\mathcal{O}_4/\mathcal{O}_5$ interference in germanium for a 10 GeV WIMP (left), and relative event rates for parameters consistent with expectations for SuperCDMS SNOLAB Ge iZIP.




Figure 5.8: Eigenvectors for 4D $\mathcal{O}_1/\mathcal{O}_3$ interference in germanium for a 10 GeV WIMP (left), and relative event rates for parameters consistent with expectations for SuperCDMS SNOLAB Ge iZIP.

Many interesting effects can be seen in these plots. First, we note that the maximal constructive interference case (the largest eigenvalue) represents typically only a modest enhancement relative to the dominant pure isoscalar operator. However, the event rates for the other cases can vary by many orders of magnitude. For the case of maximal destructive interference (smallest eigenvalue), the event rate is an order of magnitude or more smaller than the sub-dominant pure isoscalar operator, even though the sum of the squares of the operator coefficients is the same for all cases. The variation in rates is purely due to the suppression that occurs when two operators interfere.

In addition, the direction of the various interference vectors in the left-hand plots indicates which interactions germanium is most or least sensitive to. For the $\mathcal{O}_4/\mathcal{O}_5$ case, the blue eigenvector, which has the largest eigenvalue and therefore indicates the maximal event rate in germanium for these operators, points in the direction of $c_4^0 = -c_4^1$, which corresponds to WIMP-neutron spin-dependent scattering. Since germanium has one isotope, ^{73}Ge , with an unpaired neutron (and no isotopes with unpaired protons), choosing WIMP-neutron scattering should maximize the event rate. For the $\mathcal{O}_1/\mathcal{O}_3$ case, the green eigenvector points approximately in the direction that corresponds to $f_n/f_p = -0.8$ for spin-independent scattering, which is the location of the local minimum in the 2D $\mathcal{O}_1$ -only interference case.

5.4 G2 event rate comparisons

Three target elements will be used in the upcoming G2 experiments: germanium, silicon, and xenon. Under the standard spin-independent scattering framework, where the rate scales as $\sim A^2$, experiments that use xenon as a target element have the greatest sensitivity for WIMP masses above a few GeV/c^2 . However, in order to probe operators dependent on other nuclear properties, the complementarity of the three G2 target elements merits further investigation.

When considering the possible observations the G2 experiments may make, the difference in experimental parameters such as detector mass and trigger threshold must also be taken into account. The proposed LZ detector will have a 5600 kg




Figure 5.9: Relative event rates for LZ (black), SuperCDMS SNOLAB Ge iZIP (blue), and SuperCDMS SNOLAB Si (red), normalized to 1 observed event in SuperCDMS Ge ($3 \text{ GeV}/c^2$) or LZ ($10, 300 \text{ GeV}/c^2$). From left to right are shown the rates for a 3, 10, and $300 \text{ GeV}/c^2$ WIMP, assuming isoscalar interactions and the standard Maxwellian halo model. The $3 \text{ GeV}/c^2$ case also shows the rates from SuperCDMS SNOLAB Ge high-voltage (light blue), which has similar parameters to SuperCDMS Si high-voltage, but a target mass of 6 kg. (The rate for LZ at $3 \text{ GeV}/c^2$ is essentially zero for kinematic reasons.) The top row shows cumulative event rates, while the bottom row shows events per time per target mass. True interaction strengths may differ from this calculation since the interaction may proceed via a linear combination of operators.

fiducial mass of xenon, while SuperCDMS will be operating 57 kg of germanium and silicon. Figure 5.9 shows the relative event rates for the three G2 target elements assuming scattering proceeds via a single isoscalar EFT operator. This figure only shows the relative WIMP rates for the G2 experiments; background rates are not taken into consideration. Note that the true interaction, which may come from a linear combination of operators, could enhance or suppress these rates.

We normalize the event rate so that SuperCDMS Ge observes one event for a given operator in the $3 \text{ GeV}/c^2$ case and LZ observes one event for a given operator in the 10 and $300 \text{ GeV}/c^2$ cases. The LZ rate (black) assumes a 5600 kg fiducial mass, an exposure of 1000 days, a 100% trigger efficiency between 1 and 30 keV_{nr} , and a flat 50% nuclear-recoil selection efficiency. The SuperCDMS Ge rate (blue) assumes 50 kg of germanium operating in standard iZIP mode [96], an exposure of 1000 days, and a 100% trigger efficiency between 0.5 and 100 keV_{nr} , and a flat 60% combined fiducial-volume and nuclear-recoil selection efficiency. The SuperCDMS Si rate (red) assumes 1 kg of silicon and an exposure of 1000 days. Since the silicon detectors will be operated in high-voltage mode [103], the trigger threshold will be much lower, so we assume a 60% combined trigger and fiducial-volume efficiency up to 50 keV_{nr} , with a trigger threshold of 70 eV. We also plot the event rate for SuperCDMS Ge high-voltage (light blue) for the $3 \text{ GeV}/c^2$ WIMP case. For the SuperCDMS Ge high-voltage detectors, we assume a target mass of 6 kg, trigger threshold of 80 eV, and all other parameters identical to SuperCDMS Si high-voltage.

Though silicon, germanium, and xenon have similar nuclear properties (e.g., all three have isotopes with unpaired neutrons), the variation in the event rate across operators and target elements is large. For the $3 \text{ GeV}/c^2$ case, the strength of the silicon signal relative to the germanium signal varies by three orders of magnitude, depending on the operator assumed. The signal in LZ is very close to zero for such a low-mass dark matter particle because the velocity required for a $3 \text{ GeV}/c^2$ WIMP to deposit energy above the assumed 1 keV threshold is greater than the galactic escape velocity. However, for WIMP masses above a few GeV/c^2 , LZ's exposure, which is approximately 100 times larger, leads to event rates that are enhanced by approximately the same factor. In addition, the relative rate for SuperCDMS Si HV




Figure 5.10: Event rate suppression relative to $\mathcal{O}_1$ scattering in LZ (black), SuperCDMS SNOLAB Ge iZIP (blue), and SuperCDMS SNOLAB Si (red) for interference in germanium, with interference ranging from constructive (left) to maximally destructive (right), as determined by the magnitude of the corresponding eigenvalue. The rate for SuperCDMS SNOLAB Ge high-voltage (light blue) is shown for the $3 \text{ GeV}/c^2$ case where LZ sees no events above threshold. The seven operator-operator interference cases are shown, as well as pure isoscalar $\mathcal{O}_1$, which is used as a reference point.

becomes smaller at higher masses, since, by design, it is mainly sensitive to the small energy depositions produced by low-mass WIMPs.

To examine the effects of the different possible interactions for experiments with similar fiducial masses, we also plot the event rate per time per target mass (Fig. 5.9,

bottom). Here, we see that both Ge and Si SuperCDMS detectors operating in high-voltage mode are more sensitive to low-mass WIMPs because of their lower thresholds. In particular, the germanium high-voltage rate per kg day (light blue) is nearly an order of magnitude larger than the standard germanium iZIP rate (blue) for certain operators. For higher masses, the rates for xenon (black) and germanium are comparable within an order of magnitude, but the nuclear properties of silicon (red) make it less sensitive to these interactions. In addition, SuperCDMS Ge sees a modest enhancement to the overall event rate at high WIMP masses where the distribution of events extends beyond the assumed $30 \text{ keV}_{\text{nr}}$ upper limit for LZ. This effect is most prominent for operators such as $\mathcal{O}_3$ and $\mathcal{O}_{15}$, which have a q^2 dependence that suppresses the rate at low energies, though it is not enough to overcome the effects of LZ's larger target mass in the total number of events.

The variation in signal strength across target elements in this effective field theory solidifies the case for using multiple targets to detect dark matter. Previous work has shown that complementary target elements can break the degeneracy between the standard spin-dependent and spin-independent operators [136, 137], and others have shown that this concept can also be applied to the larger EFT parameter space [132]. Such considerations are particularly important when incorporating the effect of interference on the event rate. Because of the presence of both isospin interference and operator-operator interference, there are many combinations of interactions that may greatly suppress the event rate for one particular element. Even if a single experiment sees no signal due to interference effects, a complementary target with different nuclear properties may still observe events.

To demonstrate the effect of interference on the relative event rate, we determine regions of extremal interference in germanium using the principal component analysis method detailed above. The event rate suppression relative to $\mathcal{O}_1$ for the three G2 experiments for germanium constructive interference and destructive interference are shown in Fig. 5.10 for WIMPs with masses of 3, 10, or $300 \text{ GeV}/c^2$, assuming the standard Maxwellian halo model and the same experimental parameters as in Fig. 5.9. Again, this figure does not consider the relative background rates for the

three experiments. We consider all seven possible cases of four-dimensional operator-operator interference. The sum of the squares of the EFT coefficients is equal for all cases presented; however, because of the relative strength of various operators and the presence of interference, the rate can be suppressed by many orders of magnitude. We characterize the interference using the magnitude of the eigenvalue: the largest eigenvalue corresponds to the maximally-enhanced event rate, while small eigenvalues correspond to varying levels of destructive interference.

The relative event rates in Fig. 5.10 indicate that constructive interference can only modestly enhance the event rate. In the case of $\mathcal{O}_1/\mathcal{O}_3$ interference, the maximal rate is only $\sim 1.5\%$ larger than the pure $\mathcal{O}_1$ rate. For operators such as $\mathcal{O}_4$ that depend on the spin of a nucleon in the nucleus, the enhancement relative to the respective isoscalar operator tends to be slightly larger. In particular, the constructive interference eigenvector for $\mathcal{O}_4/\mathcal{O}_5$ and $\mathcal{O}_4/\mathcal{O}_6$ interference corresponds to WIMP-neutron spin-dependent scattering and is approximately a factor of 2 larger than the isoscalar $\mathcal{O}_4$ rate.

Since germanium, silicon, and xenon have similar properties, the event rate in SuperCDMS and LZ is suppressed equally for most interference cases. However, there are a few notable exceptions. From Fig. 5.9, we see that for a $3 \text{ GeV}/c^2$ WIMP interacting via a pure isoscalar operator, the event rate in SuperCDMS Si high-voltage tends to be at least an order of magnitude smaller than the rate in SuperCDMS Ge. When interference is considered, the rate in silicon may become equal to or larger than that in germanium. As an example, the $\mathcal{O}_1/\mathcal{O}_3$ right-most destructive interference case in Fig. 5.10 corresponds to maximal $\mathcal{O}_1$ isospin violation in germanium ($f_n/f_p \sim -0.8$) as discussed in [128]. For this choice of coefficients, the rate in xenon and germanium is suppressed relative to pure isoscalar $\mathcal{O}_1$ scattering in that target by a factor of ~ 500 and ~ 2000 , respectively, while the rate in silicon is suppressed by a factor of ~ 100 . A second instance of this suppression is seen for $\mathcal{O}_4/\mathcal{O}_6$ interference at $3 \text{ GeV}/c^2$ in the second plot from the left: the rate in both silicon and germanium is suppressed, but the suppression in germanium is much larger, leading to a greater number of events observed in silicon.

In addition, there exist several cases for higher WIMP masses where the rate in LZ

is smaller than that in SuperCDMS Ge, despite LZ's $100\times$ larger exposure. Maximal destructive interference (right-most plot) for $\mathcal{O}_4/\mathcal{O}_5$ and $\mathcal{O}_8/\mathcal{O}_9$ suppresses the event rate in xenon enough that SuperCDMS will see orders of magnitude more events than LZ, even for larger WIMP masses where LZ typically has an advantage. For additional interference cases the rate in LZ is less than an order of magnitude larger than that in SuperCDMS Ge. Although the cases presented here are arguably fine-tuned, the existence of regions of parameter space where interference suppresses the rate in one experiment by orders of magnitude relative to another further supports the need for multiple experiments which use a variety of target elements.

5.5 Prospects for the future

The interaction between dark matter particles and nuclei might be much more complicated than direct detection experiments have typically assumed. The inclusion of new operators within the framework of an EFT might have profound consequences for current and proposed experiments. As a result, in this richer parameter space, data from multiple experiments with different targets is essential in order to determine the precise nature of the interaction. In addition, when modeling dark matter signals, experiments must consider how an interaction due to an arbitrary EFT operator can affect the energy distribution of dark matter events.

The importance of using multiple target elements to constrain dark matter interactions can already be seen when plotting limits from current experiments. As we have shown, the differences in target element properties lead to variations in the shape of the interaction strength versus mass limit curve. In addition, a combination of target elements can produce better constraints on dark matter, especially when considering multiple dark matter interactions and the possibility of interference. This complementarity of different target elements will become increasingly important in the case of a statistically significant detection.

The additional interactions introduced by the EFT formalism become especially significant when experiments use statistical techniques which rely on assumptions

about the shape of the dark matter recoil spectrum to distinguish between background and a potential dark matter signal. Machine learning techniques, such as the boosted decision tree used in the SuperCDMS Soudan result [104], and likelihood analyses, such as the one performed on CDMS II low-energy data [102], require accurate models of both the signal and the expected background. So far, direct detection experiments have focused primarily on building accurate models of their expected backgrounds, while assuming a simple signal model. However, mis-modeling the signal could also have significant consequences. If a WIMP signal that does not conform to the standard spin-independent assumptions is present in the data, it could produce unknown effects on the final result because it may not match either the signal or the background model. In the case of algorithms such as the optimum interval method that compare the observed events to the expected WIMP spectrum but do not attempt to subtract background, WIMP signal events may be interpreted as background, leading to limits that are too strict (i.e. parameter space is incorrectly excluded) when using the optimum interval method, which effectively ignores events that do not match the expected spin-independent rate.

These considerations become especially important as the community moves forward with the proposed G2 experiments. SuperCDMS SNOLAB and LZ will have unprecedented sensitivity to dark matter scattering for a wide range of WIMP masses, and the combination of target elements allows one experiment to verify a potential signal seen by the other. However, the variation in signal strengths across EFT operators and experimental target elements could lead to experimental results that appear to be in conflict under the standard dark matter assumptions. In particular, interference between operators can suppress the relative event rates by several orders of magnitude for germanium, silicon, and xenon. If the true dark matter interaction includes such interference, it is possible that one experiment will observe a statistically significant signal while the other does not. The effective field theory framework can account for such apparent inconsistencies, and, in the event of a statistically significant signal, it will pave the way for future likelihood analyses to determine the nature of the dark matter interactions.

Chapter 6

Conclusions

This work has demonstrated several new data analysis techniques that will become increasingly useful as SuperCDMS Soudan moves towards SuperCDMS SNOLAB. The SuperCDMS Soudan low-mass WIMP search, which featured a detailed background model, showed that machine learning algorithms can be applied to direct detection data to great effect, increasing the sensitivity of a given dataset by a factor of a few. In addition, this work showed that the CDMS detector monte carlo can be used to model the expected backgrounds in fully-functional and shorted detectors and *a posteriori* understand the systematic effects of T5Z3's shorted outer charge channel on the SuperCDMS Soudan low-mass WIMP search. Finally, this work introduced tools for considering WIMP-nucleon scattering via effective field theory operators. The wide range of possible responses and the possibility for interference that is inherent in the EFT framework demonstrates that SuperCDMS SNOLAB, with its germanium and silicon detectors, will be complementary to LZ in the next generation of direct detection experiments.

6.1 Low-mass dark matter searches at SNOLAB

Because of its low thresholds, excellent background discrimination, and relatively light target elements, SuperCDMS Soudan has excellent sensitivity to low-mass WIMPs. The improved iZIP detectors used for SuperCDMS Soudan reject the surface event

background at greater than 1 in 10^5 , and the dual-sided design of the phonon sensors allows for fiducialization at very low energies, where the charge channel resolution is poor, using the phonon information.

The SuperCDMS Soudan low-mass WIMP search demonstrated the power of the phonon sensor fiducialization and set leading limits at the time on the existence of low-mass dark matter. A model of the expected backgrounds from activation lines, Compton scatters, and radiogenic contamination was developed and used to determine background rejection criteria while keeping the WIMP-search dataset blinded. The framework developed for this background simulation, where noise is added to pulses of known energy and then processed through the standard CDMS reconstruction algorithms, was also instrumental in building the detector monte carlo-based background simulations discussed in Chapter 4.

This background model showed that machine learning algorithms, long used in the realm of accelerator experiments, can also be applied to direct detection data to develop useful criteria to discriminate between signal and background. In particular, as shown in Fig. 3.9, the boosted decision tree used in the low-mass WIMP search increased the sensitivity over more traditional discrimination methods by a factor of a few. Experiments such as EDELWEISS have since begun using machine learning techniques in their own low-mass WIMP searches. In addition, with this background model, we demonstrated that it is possible to accurately estimate the expected sensitivity and propagate systematic uncertainties to the final limit. These uncertainty-propagating methods are already being used to understand how the systematic uncertainty on the nuclear recoil scale for CDMSlite affects the limits from a second, extended run in high voltage mode at Soudan.

Such techniques will be particularly useful for SuperCDMS SNOLAB, which is expected to focus on the low-mass region of dark matter parameter space. As shown in Fig. 6.1, SuperCDMS SNOLAB dominates the WIMP landscape at masses below $\sim 10 \text{ GeV}/c^2$. Although SuperCDMS SNOLAB is expected to have even lower thresholds than SuperCDMS Soudan, challenges are always present when discriminating between signal and background very near threshold. The background modeling and discrimination techniques used for SuperCDMS Soudan should translate well to

SuperCDMS SNOLAB.




Figure 6.1: Current limits from direct detection experiments (solid lines) and expected limits from the planned G2 experiments (dashed lines). SuperCDMS SNOLAB will have world-leading sensitivity to low-mass dark matter.

6.2 Moving forward with the detector monte carlo

The CDMS detector monte carlo (DMC) is a powerful tool for modeling the expected performance of CDMS detectors. This work represents the first attempt to use the detector monte carlo to inform ongoing CDMS data analysis. Processing and calibration algorithms inspired by the algorithms used on Soudan data were developed to turn DMC output into data that can be analyzed using the same tools as data taken at Soudan. The output energy and partition quantities were then studied to understand the systematic issues observed for detector T5Z3 in the SuperCDMS Soudan low-mass WIMP search.

The DMC is especially powerful for modeling detectors with known issues, such as T5Z3 in the SuperCDMS Soudan low-mass WIMP search. Using the DMC, we demonstrated that the modified electric field created by the shorted charge channel can affect the energy collection and partitioning, especially for events at high radius. We also hypothesize that events occurring near the sidewall of the shorted detector may be more susceptible to noise fluctuations that would push them into the signal region. However, the observed effect in the SuperCDMS Soudan low-mass WIMP search is of the order of a few percent (three events passed the final selection criteria out of over a hundred in the preselection region), so it will be necessary to simulate a wider range of events to fully understand the effect. In addition, using noise samples from a larger time period than the two weeks in June 2013 considered for this analysis will allow for better understanding of the possible systematics associated with time-varying noise environments.

This work demonstrates that it is possible to use the DMC to model detectors with cold hardware problems that may otherwise limit their utility for dark matter searches. Since one of the goals of the ongoing high-threshold (zero-background) analysis is to maximize exposure while maintaining very low backgrounds, this work lays the foundation for further modeling to understand the backgrounds at higher energies. In particular, the detector monte carlo group is already undertaking large-scale simulations of barium calibration data (millions of events per detector, instead of the tens of thousands considered here) to understand if there are any outlier events that systematically leak through the WIMP selection criteria.

Although the ongoing high-threshold SuperCDMS Soudan analysis will not be considering events in detectors with charge shorts in its first-pass analysis, this work also lays the foundation for using the detector monte carlo to better understand the behavior of detectors with modified electric fields. Since the overarching goal of a zero-background analysis is to maximize the exposure (and therefore the number of detectors in use) while maintaining very low background rates, if the DMC provides insight into how charge shorts affect the detector performance, it may be possible to use detectors such as T5Z3 and T1Z3 (which has both of its side 1 charge channels shorted to ground) in future analysis of this dataset.

However, work remains to be done to fully understand the behavior of the detector monte carlo. Currently the limiting factor is the phonon partitions, which do not agree with the observed distributions from Soudan data. This disagreement is largely driven by the lack of local saturation in the DMC. The bulldozer algorithm discussed in subsection 4.2.2 qualitatively models local saturation, which is not present in the standard DMC. However, further tuning of the algorithm is needed to bring the DMC phonon partitions into agreement with Soudan data.

In addition, comparisons of surface events simulated in the DMC and the low-mass WIMP search data-driven background model have uncovered discrepancies in the yield between the two. In the future, validating both these simulations against low-energy lead decay chain events observed at Soudan will inform the physics models in the DMC, especially as the collaboration moves from the current MATLAB-based detector monte carlo to the Geant4 DMC currently under development [138]. One source of validation data that will be available in the very near future is calibration data from a detector whose housing was exposed to radon gas. Radioactive decays from the radon that plated out onto the housing will produce sidewall events exactly analogous to the sidewall data simulated in this work. Validating simulated DMC sidewall events against the sidewall source data will produce even more confidence in the detector monte carlo.

In addition, we still need to understand some of the discrepancies in the distribution of boosted decision tree scores between the two simulations considered here. Retraining the algorithm on both the data-driven and the DMC simulations and examining the output in more detail should provide information about why the disagreement in the BDT score is larger than the disagreement in any underlying variable. Including the bulldozer algorithm in the DMC data that the BDT is tested on may also improve the agreement between the DMC and data-driven models.

Once the few remaining issues with the detector monte carlo are ironed out, many more projects are possible with the DMC. Currently, simulations of low-mass WIMPs in a high-voltage detector are being used to design the sensor layout to optimize the fiducial volume for CDMSlite-style detectors for SuperCDMS SNOLAB. The DMC could also provide the basis for a future likelihood analysis of Soudan data, similar

to the analysis performed on CDMSII data. This technique was shown to improve the sensitivity of CDMSII by a factor of a few, and with the ability to use the DMC to straightforwardly model both signal and background for SuperCDMS Soudan, a likelihood analysis is a clear way to both demonstrate the power of the detector monte carlo and improve the reach of SuperCDMS Soudan data. In addition, since the high voltage detectors planned for SuperCDMS SNOLAB are expected to be background-limited, the detector monte carlo provides an excellent candidate for modeling those backgrounds so they can be subtracted, again increasing the reach of the experiment.

6.3 Effective field theory beyond SuperCDMS Soudan

This work has considered numerous effective field theory operators beyond the standard spin-independent and spin-dependent interactions while remaining agnostic to the underlying physics models that produce the interactions. Within this framework, a very large range of relative event rates are possible for different target elements because the range of nuclear properties on which the responses depend is large. The large range of possible event rates across targets can help account for the current tension between exclusion limits and possible WIMP signals currently seen in the low-mass region of WIMP parameter space, and it also emphasizes the need to include several different target elements in future direct detection experiments.

In addition, the inclusion of new operators with dependence on momentum transfer and velocity can drastically alter the shape of the theoretical WIMP energy spectrum. Such effects will need to be considered in the event of an excess of events over expected backgrounds: a spectrum of candidate events that does not have the expected exponential shape derived in Chapter 1 may still be caused by WIMP scattering if the interaction proceeds via an operator with this additional momentum or velocity dependence.

As work on effective field theories for direct detection becomes more mature, new areas of inquiry will open up for future analyzers. For example, when the nonrelativistic effective field theory considered in this work is connected to simplified models, which extend the standard model by adding a single dark matter particle and a single

particle to mediate the dark matter-standard model interaction, it becomes clear that some operators considered in the nonrelativistic EFT cannot arise in relativistic simplified models [139]. In the framework of simplified models, it also becomes possible to distinguish between different models because the coupling constants depend on the parameters of the model.

The simplified model analysis also yielded two additional operators that can describe the scattering of spin-1 WIMPs off of nuclei:

$$\begin{aligned}\mathcal{O}_{17} &= i \frac{\vec{q}}{m_N} \cdot \mathcal{S} \cdot \vec{v}^\perp \\ \mathcal{O}_{18} &= i \frac{\vec{q}}{m_N} \cdot \mathcal{S} \cdot \vec{S}_N,\end{aligned}\tag{6.1}$$

where $\mathcal{S}$ is the symmetric combination of polarization vectors as defined in [139]. In the simplified model framework, these two operators are unique to spin-1 WIMPs, while certain other operators are unique to spin-0 or spin-1/2. The presence or absence of some of these operators in a potential signal yields additional information about the underlying dark matter physics.

The effective field theory framework is rich with possibilities for understanding the underlying dark matter physics by using more sophisticated algorithms than the optimal interval method considered in this work to set limits or define closed contours. Work is already being done within the CDMS collaboration as well as by Catena and Gondolo [132] to combine results from multiple direct detection experiments and perform likelihood fits in the full EFT parameter space. Several other authors have considered a list of plausible models within the simplified model framework and used simulations to show that these models can be reconstructed [140]. Extending a likelihood analysis to consider the full EFT parameter space or the models in the simplified model framework will be an important step in understanding the underlying nature of dark matter, especially in the event of a possible detection.

In addition to the simplified model framework, other theoretical work, which connects the non-relativistic EFT to the relativistic “chiral EFT” of WIMP-quark interactions, shows that not all of the operators in the non-relativistic EFT are independent at relativistic energies and also concludes that some non-relativistic operators do not

appear in the relativistic theory [141]. The effective field theory work considered here could also be extended by constraining certain linear combinations of operators as specified by the chiral effective field theory. This could be done using the optimal interval method by fixing the angle for a polar limit plot (such as Fig. 5.5), or by a more sophisticated likelihood analysis.

This work has demonstrated that the effective field theory framework also makes the case for testing a possible dark matter signal with many different target elements. In particular, extending polar plots such as Fig. 5.5 to the expected exposures of the G2 experiments, as shown in Fig. 6.2, shows that including silicon detectors in SuperCDMS SNOLAB can significantly better constrain spin-independent scattering in the isoscalar-isovector plane. These projections are limited to the energy range above ~ 2 keVnr where zero background events are expected; if low-energy backgrounds are well-understood, the true limits computed from G2 experimental data may even more strongly constrain the parameter space because of the additional data below 2 keVnr.




Figure 6.2: Projected background-free $\mathcal{O}_1$ isoscalar-isovector limits for a $6 \text{ GeV}/c^2$ WIMP in the G2 experiments (solid), compared with current limits (dashed).

SuperCDMS Soudan is the culmination of years of hard work by the scientists and

engineers of the SuperCDMS collaboration. This thesis has built upon the foundations of CDMS scientists past and present to demonstrate that new analysis approaches and techniques, including applying machine learning algorithms for event classification, using insights from the CDMS detector monte carlo, and considering new theories of dark matter interaction, can increase the sensitivity of the experiment and provide new avenues of inquiry for CDMS data analyzers. As SuperCDMS moves from Soudan to SNOLAB, the techniques developed in this thesis will become increasingly important for modeling the expected backgrounds and for understanding the nature of dark matter in the event that a dark matter signal is observed.

Appendix A

Low-frequency noise discrimination

A.1 Low-frequency noise in SuperCDMS Soudan data

Triggers on low-frequency ($\mathcal{O}(100)$ Hz) noise are often seen in the low-energy Soudan data. These events have phonon traces with shapes similar to those shown in Fig A.1 and no information in the charge traces. The CDMS optimal filter often reconstructs the amplitude of these events to be above the 2 keV energy threshold, and for some detectors, these events dominate the low energy data.




Figure A.1: Examples of low-frequency noise seen in two different detectors in low-energy Soudan data. The left-hand event (from detector T3Z1) is reconstructed to have an energy of 7 keV, and the right-hand event (from detector T4Z1) has an energy of 4 keV.

There appears to be a single dominant low-frequency shape for each detector, so it should be possible to define a low-frequency noise optimal filter template and reject the majority of these events using the difference between the χ^2 for the fit to the low-frequency noise template and the standard pulse template.

A.2 Selection criteria for low-frequency noise events

Since these pathological events are fairly common in the low-energy data at Soudan, a set of cuts that preferentially selects these events is defined, and these events are summed to create an average template that can be run through the optimal filter. There are several characteristics shared by all low-frequency noise events:

- The upward fluctuation of the low-frequency noise causes the DAQ to trigger. This guarantees that the noise will be more or less in phase, eliminating much of the destructive interference that could potentially occur in averaging traces.
- The event has a slow rise time compared to real phonon pulses, which typically have rise times of a few hundred μs .
- The event has an energy at most a few keV above the trigger threshold.

For each of the 15 Soudan detectors, I examined raw phonon traces with these characteristics and picked a single trace that I thought was representative of the low-frequency noise for that detector. An example representative trace for detector T4Z1 is shown in Figure A.2.

After passing each representative trace through a 50 kHz low-pass filter, a selection of unblinded data from the periods following californium calibrations was processed using the optimal filter in *cdmsbats* to calculate the χ^2 of the fit to the representative trace. The difference between the χ^2 of the fit to the standard phonon pulse template and this χ^2 provided another selection criteria for low-frequency noise traces.




Figure A.2: Representative low-frequency noise trace for detector T4Z2.

A.3 Low frequency noise template development

Low frequency noise templates were generated using the above selection criteria in a method similar to the generation of standard phonon pulse templates. The traces used to generate the template were selected from a set of low-energy events which issued a trigger, had a 90% rise time greater than 1ms, did not contain pileup, and were not caused by an electronics glitch. These traces were then summed with no time shifting performed, run through a 10 kHz low-pass filter, and normed to 1 to create the template for each detector. A comparison of the low frequency noise template and the phonon pulse template for T4Z1 is shown in Fig A.3.

Test data processing using these templates

All unblind WIMP search data from the days following californium calibration was processed using these templates to test how well a $\Delta\chi^2$ selection criteria picks up low frequency noise events. I define

$$\Delta\chi^2 = \chi^2(\text{phonon pulse template}) - \chi^2(\text{low frequency noise template})$$




Figure A.3: Low frequency noise template and phonon pulse template for T4Z2.

so events that are true phonon pulses should have negative $\Delta\chi^2$, while events that are triggers on low frequency noise should have positive $\Delta\chi^2$. A plot of $\Delta\chi^2$ versus total phonon energy for detector T4Z1 is shown in Fig A.4 below.

The distribution in Fig A.4 makes intuitive sense: at low phonon energy (ptOF), most events have positive $\Delta\chi^2$ since events below the 2 keV trigger threshold are mainly noise. At higher ptOF , the events appear proportional to $-\text{ptOF}^2$, as expected from the quadratic dependency of χ^2 on pulse amplitude. Example triggered events with positive and negative $\Delta\chi^2$ are shown in Figs A.5 and A.6, respectively. As indicated by the value of $\Delta\chi^2$, the LF noise template is a much better fit for the event in Fig A.5 and the pulse template is a much better fit for the event in Fig A.6.

A.4 Determining selection criteria

For the low-mass WIMP search discussed in chapter 4, the selection criteria were defined to remove low-frequency noise with high confidence. Because low-frequency noise is a significant component of the event rate at the low energies relevant to that analysis, the selection criteria were defined to remove all low-frequency noise from the




Figure A.4: $\Delta\chi^2$ versus ptOF using the low frequency noise template for detector T4Z1.

signal region. Thus, a two-part threshold was devised: at low energies, an energy-independent threshold was determined to pass the majority of random triggers (a good proxy for very low-energy nuclear recoils). At higher energies, the threshold was fit to follow the parabolic band of good events.

To determine the flat portion of the threshold, all random triggers from low-background data that pass basic reconstruction quality cuts were examined, as shown in Figure A.7. The true histogram (blue), has more power in the positive- $\Delta\chi^2$ tail of the histogram due to the presence of low-frequency noise in the noise environment. Since we do not expect low-frequency noise to show up at negative $\Delta\chi^2$, the noise traces with negative $\Delta\chi^2$ were reflected about $\Delta\chi^2=0$ to obtain a “symmetrized” distribution with no low-frequency noise component. The flat threshold was set so that 95% of this symmetrized distribution passes the threshold.

The parabolic portion of the thresholds was determined by fitting the band of good pulses obtained in californium data. This data, which is from high-statistics calibration data (where the rate from the source significantly exceeds the low-frequency noise rate) is expected to be a fairly pure sample of good pulses. Events were binned in




Figure A.5: Example event with positive $\Delta\chi^2$ and an energy of 2.3 keV, shown in frequency domain (left) and time domain (right).




Figure A.6: Example event with negative $\Delta\chi^2$ and an energy of 5.2 keV, shown in frequency domain (left) and time domain (right).

1keV bins, and each bin was fit to a gaussian to determine the mean and standard deviation. The 2σ and 3σ bands were fit to parabolas, as shown in Figure A.8. The 2σ band was used as the cut threshold at energies above ~ 10 keV, but near threshold, the σ level was tuned by eye to remove as much low-frequency noise as possible.

The resulting discrimination criteria is shown in Figure A.9. All events which fall below the flat line and the green parabola are considered “good events.” Note the presence of red points that fall below the flat threshold but are removed by the parabolic threshold. These events are the “scariest” for the low-mass WIMP search, since the goal was to push the threshold as low as possible while minimizing




Figure A.7: Histogram of $\Delta\chi^2$ for random triggers. The true distribution of randoms is shown in blue, while the expected distribution of randoms with no low-frequency noise component is shown in magenta. Flat cut thresholds were set such that 95% of the distribution with no low-frequency noise (magenta) passed the threshold.

backgrounds. For the ongoing high-threshold WIMP search, which is designed to be background-free, the thresholds were raised so this region was excluded; as a result, the cut was reverted to the flat threshold.

A.5 Estimating the efficiency of the low-frequency noise selection criteria

A pulse simulation was used to measure the efficiency of the selection criteria. The phonon fitting template was scaled to a random energy, added to a noise trace, and processed through the standard reconstruction algorithms. This produced a dataset of pure good pulses; the only variance was caused by the variance in the noise environment. The efficiency was defined as the fraction of events in this dataset which passed the selection criteria, divided by the total number of events. It was measured as a function of energy, as shown in Figure A.10. (This same method was used to






Figure A.8: Histogram of $\Delta\chi^2$ for californium calibration data. The fits to the 2σ and 3σ bands of good pulses are shown.


estimate the efficiency of other pulse-shape criteria, including the χ^2 cut that rejects pileup events and the $\Delta\chi^2$ cut that rejects electronics glitches.)


At high energies, the efficiency of the low-frequency noise selection criteria is close to unity, since the band of noise events and the band of good pulses are well-separated. As the energy approaches threshold (2 keV analysis threshold for this detector; the trigger threshold was >2 keV), the efficiency also decreases. This is because the noise events and good pulses are less well-separated at low energies where the signal-to-noise is poor. As a result, a larger number of good pulses fail the selection criteria.

Figure A.9: Histogram of $\Delta\chi^2$ for unblind low-background data.Figure A.10: Efficiency of low-frequency noise selection criteria for detector T4Z2.
[Figure: Adam Anderson]

Appendix B

EFT exclusion limits








Figure B.1: Upper limits on the dimensionless isoscalar coefficients for all effective field theory operators as a function of WIMP mass for SuperCDMS Soudan (light blue) [104], CDMSII Ge reanalysis (dark blue) [133], and CDMSII Si (red) [134], and estimated limits for LUX (black) [85], for the Maxwellian halo (solid) and an alternate halo model (dashed).

Appendix C

Testing high electron mobility transistors in the Stanford ^3He fridge

The charge amplifiers currently operating at Soudan use Field Effect Transistors (FETs), which operate at ~ 130 K. To decrease the noise in the charge signal, the charge amplifiers designed for SNOLAB plan to use High Electron Mobility Transistors (HEMTs), which can be operated at cryogenic temperatures. To understand the behavior of the HEMTs and to better model their behavior in the SPICE circuit simulation package, the properties of the HEMTs were tested at 4K and the SPICE parameters were extracted by fitting the IV curves to the expected behavior.

Measurement method

The IV curves presented in Figure C.3 were acquired using a Kiethley 2400 source meter in four-wire mode to provide the drain-source bias and a Kiethley 6517B electrometer to provide the gate-source bias. The Kiethley source meter interfaced with a MATLAB program that stepped the drain-source bias voltage and measured the drain current. A block diagram of the measurement setup is shown in Fig C.1. Three separate HEMTs were wire-bonded to a board and installed in the Stanford ^3He

fridge, as shown in Figure C.2, and cooled to 4K.




Figure C.1: Block diagram of fridge setup for HEMT IV measurement.

We want to use these curves to determine the SPICE parameters for the HEMT by fitting the data to the equation

$$I_D = \begin{cases} 0 & V_{GS} \leq V_T \\ \beta(V_{GS} - V_T)^2 & V_{GS} > V_T \end{cases} \quad (\text{C.1})$$

where the drain-source voltage is held constant and the HEMT is operating in the saturation region. In the linear region, the drain current is for $V_{GS} > V_T$

$$I_D = \beta[2(V_{GS} - V_T)V_{DS} - V_{DS}^2] \quad (\text{C.2})$$

Fig C.4 shows the results of fitting the data for HEMTs 1 and 2 to Eqn C.2 for $V_{DS} = 175$ mV. (Fits to other values of V_{DS} in the saturation region yield nearly identical results.) The best-fit value of β (the transconductance parameter) is 3.2×10^{-4} mA/mV², and the best-fit value of V_T (the threshold voltage at which drain current begins to flow) is -81 mV. The assumed error of 0.001 mA on I_D was chosen somewhat arbitrarily, so the error values are not representative of the true uncertainty in the measurement.




Figure C.2: HEMT board as installed in Stanford ^3He fridge. The HEMTs are located on the underside of the top (green) board, and the bottom (yellow) board acts as a guard for the HEMTs and wire bonds.




Figure C.3: Drain current vs drain-source voltage for two of the three high-electron mobility transistors tested on October 16, 2012, measured at 4K.

The transconductance can be found by taking the derivative of Eqn C.2 with respect to V_{GS} .

$$g_m = 2\sqrt{\beta I_D} \quad (\text{C.3})$$

For β as calculated above, this value is

$$g_m = 36\sqrt{I_D} \text{ mS}$$

where I_D is in mA.




Figure C.4: Fit to equation C.2 for the HEMTs 1 and 2 at $V_{DS} = 175$ mV.

A summary of the results from the data in Figure C.3 and results from two earlier runs is given in table C.1. There is a high amount of variability in the results, which may have been caused by space charge in the HEMTs in the earlier runs. Note that g_m is calculated for I_D in mA.

Table C.1: Summary of results from several Stanford ^3He tests

Date of Run	Batch (A/B) & HEMT #	Geometry	V_T [mV]	β [mA/mV ²]	g_m [mS]
May 31	A-1	1x64	-40	1.3×10^{-5}	$7.1 \sqrt{I_D}$
July 10	A-1	1x64	-207	6.2×10^{-5}	$16 \sqrt{I_D}$
July 10	A-2	2x32	-220	1.7×10^{-4}	$26 \sqrt{I_D}$
October 16	B-1	2x32	-81	3.2×10^{-4}	$36 \sqrt{I_D}$
October 16	B-2	2x32	-75	3.5×10^{-4}	$37 \sqrt{I_D}$
October 16	B-3	2x32	-73	3.4×10^{-4}	$37 \sqrt{I_D}$

Bibliography

- [1] F. Zwicky, **Die Rotverschiebung von extragalaktischen Nebeln.** *Helvetica Physica Acta* **6**, 110–127 1933.
- [2] F. Zwicky, **On the Masses of Nebulae and of Clusters of Nebulae.** *Astrophys.J.* **86**, 217 October 1937.
- [3] V. C. Rubin and W. K. Ford, Jr., **Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions.** *Astrophys.J.* **159**, 379 February 1970.
- [4] V. C. Rubin, N. Thonnard, and W. K. Ford, Jr., **Extended rotation curves of high-luminosity spiral galaxies. IV - Systematic dynamical properties, SA through SC.** *Astrophys.J.* **225**, L107–L111 November 1978.
- [5] A. Bosma, **The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types.** Ph.D. thesis, Groningen Univ., 1978.
- [6] V. C. Rubin, W. K. J. Ford, and N. . Thonnard, **Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/.** *Astrophys.J.* **238**, 471–487 June 1980.
- [7] V. Trimble, **Existence and nature of dark matter in the universe.** *Ann. Rev. Astron. Astrophys.* **25**, 425–472 1987.

- [8] A. Klypin, H. Zhao, and R. S. Somerville, **Λ CDM-based Models for the Milky Way and M31. I. Dynamical Models.** *Astrophys.J.* **573**, 597–613 July 2002.
- [9] G. Soucail *et al.*, **Further data on the blue ring-like structure in A 370.** *Astronomy and Astrophysics* **184**, L7–L9 October 1987.
- [10] J. A. Tyson, G. P. Kochanski, and I. P. Dell’Antonio, **Detailed Mass Map of CL 0024+1654 from Strong Lensing.** *Astrophys.J.* **498**, L107–L110 May 1998.
- [11] A. N. Taylor *et al.*, **Gravitational Lens Magnification and the Mass of Abell 1689.** *Astrophys.J.* **501**, 539–553 July 1998.
- [12] J. A. Tyson, R. A. Wenk, and F. Valdes, **Detection of systematic gravitational lens galaxy image alignments - Mapping dark matter in galaxy clusters.** *Astrophys.J.* **349**, L1–L4 January 1990.
- [13] A. Refregier, **Weak Gravitational Lensing by Large-Scale Structure.** *Ann. Rev. Astron. Astrophys.* **41**, 645–668 2003.
- [14] D. Clowe *et al.*, **A Direct Empirical Proof of the Existence of Dark Matter.** *Astrophys.J.* **648**, L109–L113 September 2006.
- [15] V. Springel, C. S. Frenk, and S. D. M. White, **The large-scale structure of the Universe.** *Nature* **440**, 1137–1144 April 2006.
- [16] J. A. Frieman, M. S. Turner, and D. Huterer, **Dark Energy and the Accelerating Universe.** *Ann.Rev.Astron.Astrophys.* **46**, 385–432 September 2008.
- [17] A. G. Riess *et al.*, **Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant.** *Astron.J.* **116**, 1009–1038 September 1998.
- [18] S. Perlmutter *et al.*, **Measurements of Ω and Λ from 42 High-Redshift Supernovae.** *Astrophys.J.* **517**, 565–586 June 1999.

- [19] G. Hinshaw *et al.*, **Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Parameter Results.** *Astrophys.J.Supplement* **208**, 19 October 2013.
- [20] Planck Collaboration *et al.*, **Planck 2015 results. XIII. Cosmological parameters.** *ArXiv e-prints* 1502.01589
- [21] J. Beringer *et al.*, **Review of Particle Physics.** *Phys.Rev.D* **86**, 010001 July 2012.
- [22] K. A. Olive and Particle Data Group, **Review of Particle Physics.** *Chinese Physics C* **38**, 090001 August 2014.
- [23] W. J. Percival *et al.*, **Baryon acoustic oscillations in the Sloan Digital Sky Survey Data Release 7 galaxy sample.** *Monthly Notices of the Royal Astronomical Society* **401**, 2148–2168 February 2010.
- [24] A. A. Klypin, S. Trujillo-Gomez, and J. Primack, **Dark Matter Halos in the Standard Cosmological Model: Results from the Bolshoi Simulation.** *The Astrophysical Journal* **740**, 102 October 2011.
- [25] A. A. Penzias and R. W. Wilson, **A Measurement of Excess Antenna Temperature at 4080 Mc/s..** *Astrophys.J.* **142**, 419–421 July 1965.
- [26] G. F. Smoot *et al.*, **Structure in the COBE differential microwave radiometer first-year maps.** *Astrophys.J.* **396**, L1–L5 September 1992.
- [27] C. L. Bennett *et al.*, **Nine-year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Final Maps and Results.** *Astrophys.J.Supplement* **208**, 20 October 2013.
- [28] A. Linde, **Inflationary Cosmology.** in *Inflationary Cosmology*, Vol. 738 of *Lecture Notes in Physics*, Berlin Springer Verlag, edited by M. Lemoine, J. Martin, and P. Peter (, , 2008), p. 1.

- [29] A. Linde, **Inflationary Cosmology after Planck 2013**. *ArXiv e-prints* 1402.0526
- [30] P. J. E. Peebles and J. T. Yu, **Primeval Adiabatic Perturbation in an Expanding Universe**. *Astrophys.J.* **162**, 815 December 1970.
- [31] R. A. Sunyaev and Y. B. Zeldovich, **Small scale entropy and adiabatic density perturbations — Antimatter in the Universe**. *Astrophys. & Space Sci.* **9**, 368–382 December 1970.
- [32] J. L. Feng, **Dark matter candidates from particle physics and methods of detection**. *Ann. Rev. Astro. Astrophys.* **48**, 495–545 2010.
- [33] K.-Y. Choi, J. E. Kim, and L. Roszkowski, **Review of axino dark matter**. *Journal of Korean Physical Society* **63**, 1685–1695 November 2013.
- [34] S. D. M. White, C. S. Frenk, and M. Davis, **Clustering in a neutrino-dominated universe**. *Astrophys.J.* **274**, L1–L5 November 1983.
- [35] G. R. Blumenthal, S. M. Faber, J. R. Primack, and M. J. Rees, **Formation of galaxies and large-scale structure with cold dark matter**. *Nature* **311**, 517–525 October 1984.
- [36] B. W. Lee and S. Weinberg, **Cosmological Lower Bound on Heavy-Neutrino Masses**. *Physical Review Letters* **39**, 165–168 1977.
- [37] G. Steigman, B. Dasgupta, and J. F. Beacom, **Precise relic WIMP abundance and its impact on searches for dark matter annihilation**. *Phys.Rev.D* **86**, 023506 July 2012.
- [38] S. P. Martin, **A Supersymmetry Primer**. *Perspectives On Supersymmetry. Series: Advanced Series on Directions in High Energy Physics, ISBN: 978-981-02-3553-6/ISBN 978-981-02-3553-6. WORLD SCIENTIFIC, Edited by Gordon L Kane, vol. 18, pp. 1-98* **18**, 1–98 July 1998.

- [39] G. Jungman, M. Kamionkowski, and K. Griest, **Supersymmetric dark matter**. *Physics Reports* **267**, 195–373 1996.
- [40] K. Griest and M. Kamionkowski, **Supersymmetric dark matter**. *Physics Reports* **333334**, 167 – 182 2000.
- [41] G. Aad *et al.*, **Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC**. *Physics Letters B* **716**, 1–29 September 2012.
- [42] S. Chatrchyan *et al.*, **Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC**. *Physics Letters B* **716**, 30–61 September 2012.
- [43] M. Cahill-Rowley *et al.*, **Complementarity and Searches for Dark Matter in the pMSSM**. *ArXiv e-prints* 1305.6921
- [44] R. Bernabei *et al.*, **New results from dama/libra**. *Eur. Phys. J. C* **67**, 39–49 2010.
- [45] C. E. Aalseth *et al.*, **Cogent: A search for low-mass dark matter using p -type point contact germanium detectors**. *Phys. Rev. D* **88**, 012002 Jul 2013.
- [46] G. Angloher *et al.*, **Results from 730 kgâ€days of the cresst-ii dark matter search**. *Eur. Phys. J. C* **72**, 1971 2012.
- [47] R. Agnese *et al.*, **Dark Matter Search Results Using the Silicon Detectors of CDMS II**. *Phys. Rev. Lett.* **111**, 251301 2013.
- [48] J. Kozaczuk and S. Profumo, **Light NMSSM neutralino dark matter in the wake of CDMS II and a 126 GeV Higgs boson**. *Physical Review D* **89**, 095012 May 2014.
- [49] D. G. Cerdeño, J.-H. Huh, M. Peiró, and O. Seto, **Very light right-handed sneutrino dark matter in the NMSSM**. *Journal of Cosmology and Astroparticle Physics* **11**, 027 November 2011.

- [50] D. B. Kaplan, **Single explanation for both baryon and dark matter densities.** *Physical Review Letters* **68**, 741–743 February 1992.
- [51] S. M. Barr, **Baryogenesis, sphalerons, and the cogeneration of dark matter.** *Physical Review D* **44**, 3062–3066 November 1991.
- [52] R. Peccei and H. R. Quinn, **CP Conservation in the Presence of Instantons.** *Physical Review Letters* **38**, 1440–1443 1977.
- [53] S. Asztalos *et al.*, **Large-scale microwave cavity search for dark-matter axions.** *Phys. Rev. D* **64**, 092003 November 2001.
- [54] S. J. Asztalos *et al.*, **SQUID-Based Microwave Cavity Search for Dark-Matter Axions.** *Physical Review Letters* **104**, 041301 January 2010.
- [55] Z. Ahmed *et al.*, **Search for Axions with the CDMS Experiment.** *Physical Review Letters* **103**, 141802–5 October 2009.
- [56] A. Drlica-Wagner *et al.*, **Search for Gamma-Ray Emission from DES Dwarf Spheroidal Galaxy Candidates with Fermi-LAT Data.** *ArXiv e-prints* 1503.02632
- [57] Fermi-LAT Collaboration, **Searching for Dark Matter Annihilation from Milky Way Dwarf Spheroidal Galaxies with Six Years of Fermi-LAT Data.** *ArXiv e-prints* 1503.02641
- [58] L. Accardo *et al.*, **High statistics measurement of the positron fraction in primary cosmic rays of 0.5–500 GeV with the alpha magnetic spectrometer on the international space station.** *Phys. Rev. Lett.* **113**, 121101 Sep 2014.
- [59] L. Bergstrom *et al.*, **New Limits on Dark Matter Annihilation from Alpha Magnetic Spectrometer Cosmic Ray Positron Data.** *Physical Review Letters* **111**, 171101–6 October 2013.

- [60] A. Morselli, B. Cañadas, and V. Vitale, **The Indirect Search for Dark Matter from the centre of the Galaxy with the Fermi LAT.** *Il Nuovo Cimento* **C034N3**, 311–316 2011.
- [61] D. Hooper and L. Goodenough, **Dark matter annihilation in the Galactic Center as seen by the Fermi Gamma Ray Space Telescope.** *Physics Letters B* **697**, 412–428 March 2011.
- [62] D. Hooper and T. Linden, **Origin of the gamma rays from the Galactic Center.** *Physical Review D* **84**, 123005 December 2011.
- [63] K. N. Abazajian and M. Kaplinghat, **Detection of a Gamma-Ray Source in the Galactic Center Consistent with Extended Emission from Dark Matter Annihilation and Concentrated Astrophysical Emission.** *Phys.Rev.* **D86**, 083511 2012.
- [64] T. Daylan *et al.*, **The Characterization of the Gamma-Ray Signal from the Central Milky Way: A Compelling Case for Annihilating Dark Matter.** *ArXiv e-prints* 1402.6703
- [65] F. Calore, I. Cholis, and C. Weniger, **Background model systematics for the Fermi GeV excess.** *ArXiv e-prints* 1409.0042
- [66] G. Aad *et al.*, **Search for new phenomena in events with a photon and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector.** *Phys.Rev.* **D91**, 012008 2015.
- [67] V. Khachatryan *et al.*, **Search for new phenomena in monophoton final states in proton-proton collisions at $\sqrt{s} = 8$ TeV.** *ArXiv e-prints* 1410.8812
- [68] V. Khachatryan *et al.*, **Search for physics beyond the standard model in final states with a lepton and missing transverse energy in proton-proton collisions at $\sqrt{s} = 8$ TeV.** *ArXiv e-prints* 1408.2745

- [69] G. Aad *et al.*, **Search for new particles in events with one lepton and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector.** *Journal of High Energy Physics* **9**, 37 September 2014.
- [70] M. W. Goodman and E. Witten, **Detectability of certain dark-matter candidates.** *Phys. Rev. D* **31**, 3059–3063 Jun 1985.
- [71] J. Billard, E. Figueroa-Feliciano, and L. Strigari, **Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments.** *Phys. Rev. D* **89**, 023524 Jan 2014.
- [72] L. E. Strigari, **Neutrino coherent scattering rates at direct dark matter detectors.** *New Journal of Physics* **11**, 105011 October 2009.
- [73] P. Cushman *et al.*, **Snowmass CF1 Summary: WIMP Dark Matter Direct Detection.** *ArXiv e-prints* 1310.8327
- [74] J. Lewin and P. Smith, **Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil.** *Astropart. Phys.* **6**, 87 – 112 1996.
- [75] F. Donato, N. Fornengo, and S. Scopel, **Effects of galactic dark halo rotation on WIMP direct detection.** *Astroparticle Physics* **9**, 247–260 October 1998.
- [76] K. Hamaguchi, S. P. Liew, T. Moroi, and Y. Yamamoto, **Isospin-violating dark matter with colored mediators.** *Journal of High Energy Physics* **5**, 86 May 2014.
- [77] S. Arrenberg, **Searching for dark matter with the cryogenic dark matter search experiment.** Ph.D. thesis, University of Zurich, 2011.
- [78] M. C. Smith *et al.*, **The RAVE survey: constraining the local Galactic escape speed.** *Monthly Notices of the Royal Astronomical Society* **379**, 755–772 August 2007.

- [79] M. Kuhlen *et al.*, **Dark matter direct detection with non-Maxwellian velocity structure.** *Journal of Cosmology and Astroparticle Physics* **2**, 30 February 2010.
- [80] C. W. Purcell, A. R. Zentner, and M.-Y. Wang, **Dark matter direct search rates in simulations of the Milky Way and Sagittarius stream.** *Journal of Cosmology and Astroparticle Physics* **8**, 27 August 2012.
- [81] Y.-Y. Mao *et al.*, **Halo-to-halo Similarity and Scatter in the Velocity Distribution of Dark Matter.** *Astrophys. J.* **764**, 35 February 2013.
- [82] Y.-Y. Mao, L. E. Strigari, and R. H. Wechsler, **Connecting direct dark matter detection experiments to cosmologically motivated halo models.** *Phys. Rev. D* **89**, 063513 March 2014.
- [83] A. Pillepich, M. Kuhlen, J. Guedes, and P. Madau, **The Distribution of Dark Matter in the Milky Way’s Disk.** *Astrophys. J.* **784**, 161 April 2014.
- [84] E. Aprile *et al.*, **Dark Matter Results from 225 Live Days of XENON100 Data.** *Physical Review Letters* **109**, 181301 November 2012.
- [85] D. Akerib *et al.*, **First results from the lux dark matter experiment at the sanford underground research facility.** *Phys. Rev. Lett.* **112**, 091303 Mar 2014.
- [86] CDMS Collaboration *et al.*, **Search for annual modulation in low-energy CDMS-II data.** *ArXiv e-prints 1203.1309*
- [87] F. Mayet, J. Billard, and D. Santos, **Directional Detection of Dark Matter.** in *EAS Publications Series*, Vol. 53 of *EAS Publications Series*, edited by F. Mayet and D. Santos (, , 2012), pp. 3–10.
- [88] J. B. R. Battat *et al.*, **First background-free limit from a directional dark matter experiment: results from a fully fiducialised DRIFT detector.** *ArXiv e-prints 1410.7821*

- [89] E. Daw *et al.*, **Spin-dependent limits from the DRIFT-II directional dark matter detector**. *Astroparticle Physics* **35**, 397–401 February 2012.
- [90] J. Monroe, **Status and Prospects of the DMTPC Directional Dark Matter Experiment**. *EAS Publ.Ser.* **53**, 19–24 2012.
- [91] D. Santos *et al.*, **MIMAC: Micro-tpc Matrix of Chambers for dark matter directional detection**. *Journal of Physics Conference Series* **469**, 012002 December 2013.
- [92] P. Grothaus, M. Fairbairn, and J. Monroe, **Directional dark matter detection beyond the neutrino bound**. *Phys. Rev. D* **90**, 055018 Sep 2014.
- [93] D. S. Akerib *et al.*, **Exclusion limits on the WIMP-nucleon cross section from the first run of the Cryogenic Dark Matter Search in the Soudan Underground Laboratory**. *Phys.Rev.D* **72**, 052009 September 2005.
- [94] R. A. Moffatt *et al.*, **Imaging the Oblique Propagation of Electrons in Germanium Crystals at Low Temperature and Low Electric Field**. *ArXiv e-prints* 1505.00052
- [95] P. Luke, **Voltage-assisted calorimetric ionization detector**. *J. Appl. Phys.* **64**, 6858 1988.
- [96] R. Agnese *et al.*, **Demonstration of Surface Electron Rejection with Interleaved Germanium Detectors for Dark Matter Searches**. *Appl. Phys. Lett.* **103**, 164105 2013.
- [97] K. D. Irwin, **An application of electrothermal feedback for high resolution cryogenic particle detection**. *Applied Physics Letters* **66**, 1998–2000 1995.
- [98] S. Golwala, **Exclusion limits on wimp-nucleon elastic scattering cross-section from the cryogenic dark matter search**. Ph.D. thesis, University of California, Berkeley, 2000.

- [99] J. 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.
- [100] J. Lindhard, V. Nielsen, M. Scharff, and P. Thomsen, *Integral Equations Governing Radiation Effects* (Munksgaard i komm., Copenhagen, Denmark, , 1963).
- [101] Z. Ahmed *et al.*, **Results from a low-energy analysis of the cdms ii germanium data**. *Phys. Rev. Lett.* **106**, 131302 Mar 2011.
- [102] R. Agnese *et al.*, **Maximum likelihood analysis of low energy CDMS II germanium data**. *Phys. Rev. D* **91**, 052021 March 2015.
- [103] R. Agnese *et al.*, **Search for low-mass weakly interacting massive particles using voltage-assisted calorimetric ionization detection in the supercdms experiment**. *Phys. Rev. Lett.* **112**, 041302 Jan 2014.
- [104] R. Agnese *et al.*, **Search for Low-Mass Weakly Interacting Massive Particles with SuperCDMS**. *Physical Review Letters* **112**, 241302 June 2014.
- [105] A. Anderson, **A search for light weakly-interacting massive particles and applications to neutrino physics**. Ph.D. thesis, Massachusetts Institute of Technology, 2015.
- [106] S. Agostinelli *et al.*, **Geant4: A Simulation toolkit**. *Nucl. Instrum. Methods A* **506**, 250–303 2003.
- [107] S. Yellin, **Extending the optimum interval method**. *ArXiv e-prints* 0709.2701
- [108] S. Yellin, **Finding an upper limit in the presence of an unknown background**. *Phys. Rev. D* **66**, 032005 Aug 2002.
- [109] C. Savage, K. Freese, and P. Gondolo, **Annual modulation of dark matter in the presence of streams**. *Phys. Rev. D* **74**, 043531 August 2006.

- [110] C. Savage, G. Gelmini, P. Gondolo, and K. Freese, **Compatibility of DAMA/LIBRA dark matter detection with other searches**. *JCAP* **0904**, 010 2009.
- [111] Z. Ahmed *et al.*, **Dark matter search results from the cdms ii experiment**. *Science* **327**, 1619–1621 2010.
- [112] J. Angle *et al.*, **Search for light dark matter in xenon10 data**. *Phys. Rev. Lett.* **107**, 051301 Jul 2011.
- [113] J. Angle *et al.*, **Erratum: Search for light dark matter in xenon10 data [phys. rev. lett. 107, 051301 (2011)]**. *Phys. Rev. Lett.* **110**, 249901(E) Jun 2013.
- [114] E. Armengaud *et al.*, **Search for low-mass wimps with edelweiss-ii heat-and-ionization detectors**. *Phys. Rev. D* **86**, 051701 Sep 2012.
- [115] G. Angloher *et al.*, **Results on low mass WIMPs using an upgraded CRESST-II detector**. *European Physical Journal C* **74**, 3184 December 2014.
- [116] E. Behnke *et al.*, **First dark matter search results from a 4-kg CF₃I bubble chamber operated in a deep underground site**. *Phys. Rev. D* **86**, 052001 September 2012.
- [117] M. Felizardo *et al.*, **Final Analysis and Results of the Phase II SIMPLE Dark Matter Search**. *Physical Review Letters* **108**, 201302 May 2012.
- [118] S. Archambault *et al.*, **Constraints on low-mass WIMP interactions on ¹⁹F from PICASSO**. *Physics Letters B* **711**, 153–161 May 2012.
- [119] V. I. Dimitrov, J. Engel, and S. Pittel, **Scattering of weakly interacting massive particles from ⁷³Ge**. *Phys. Rev. D* **51**, R291–R295 Jan 1995.
- [120] P. Klos, J. Menéndez, D. Gazit, and A. Schwenk, **Large-scale nuclear structure calculations for spin-dependent wimp scattering with chiral effective field theory currents**. *Phys. Rev. D* **88**, 083516 Oct 2013.

- [121] S. W. Leman, **Invited Review Article: Physics and Monte Carlo techniques as relevant to cryogenic, phonon, and ionization readout of Cryogenic Dark Matter Search radiation detectors.** *Review of Scientific Instruments* **83**, 091101 September 2012.
- [122] Z. He, **Review of the Shockley-Ramo theorem and its application in semiconductor gamma-ray detectors.** *Nuclear Instruments and Methods in Physics Research A* **463**, 250–267 May 2001.
- [123] P. Redl, **Accurate Simulations of ^{206}Pb Recoils in SuperCDMS.** *Journal of Low Temperature Physics* 1–6 2014.
- [124] G. Steigman and M. S. Turner, **Cosmological constraints on the properties of weakly interacting massive particles.** *Nuclear Physics B* **253**, 375–386 1985.
- [125] D. Hooper and T. Linden, **Origin of the gamma rays from the galactic center.** *Phys. Rev. D* **84**, 123005 Dec 2011.
- [126] F. Calore, I. Cholis, and C. Weniger, **Background model systematics for the Fermi GeV excess.** *ArXiv e-prints 1409.0042*
- [127] C. Amole *et al.*, **Dark Matter Search Results from the PICO-2L C_3F_8 Bubble Chamber.** *ArXiv e-prints 1503.00008*
- [128] J. L. Feng, J. Kumar, and D. Sanford, **Xenophobic dark matter.** *Physical Review D* **88**, 015021 July 2013.
- [129] A. L. Fitzpatrick *et al.*, **Model Independent Direct Detection Analyses.** *ArXiv e-prints 1211.2818*
- [130] N. Anand, A. L. Fitzpatrick, and W. C. Haxton, **Model-independent WIMP Scattering Responses and Event Rates: A Mathematica Package for Experimental Analysis.** *ArXiv e-prints 1308.6288*

- [131] A. L. Fitzpatrick *et al.*, **The effective field theory of dark matter direct detection.** *Journal of Cosmology and Astroparticle Physics* **2**, 4 February 2013.
- [132] R. Catena and P. Gondolo, **Global fits of the dark matter-nucleon effective interactions.** *Journal of Cosmology and Astroparticle Physics* **9**, 45 September 2014.
- [133] R. Agnese *et al.*, **Improved WIMP-search reach of the CDMS II germanium data.** *ArXiv e-prints* 1504.05871
- [134] R. Agnese *et al.*, **Silicon detector results from the first five-tower run of CDMS II.** *Physical Review D* **88**, 031104 August 2013.
- [135] K. Schneck *et al.*, **Dark matter effective field theory scattering in direct detection experiments.** *Phys. Rev. D* **91**, 092004 May 2015.
- [136] G. Bertone, D. G. Cerdeño, J. I. Collar, and B. Odom, **Identification of Weakly Interacting Massive Particles Through a Combined Measurement of Axial and Scalar Couplings.** *Physical Review Letters* **99**, 151301 October 2007.
- [137] D. G. Cerdeño *et al.*, **Complementarity of dark matter direct detection: the role of bolometric targets.** *Journal of Cosmology and Astroparticle Physics* **7**, 28 July 2013.
- [138] D. Brandt *et al.*, **Semiconductor phonon and charge transport Monte Carlo simulation using Geant4.** *ArXiv e-prints* 1403.4984
- [139] J. B. Dent, L. M. Krauss, J. L. Newstead, and S. Sabharwal, **A General Analysis of Direct Dark Matter Detection: From Microphysics to Observational Signatures.** *ArXiv e-prints* 1505.03117
- [140] V. Gluscevic *et al.*, **Identifying the Theory of Dark Matter with Direct Detection.** *ArXiv e-prints* 1506.04454

- [141] M. Hoferichter, P. Klos, and A. Schwenk, **Chiral power counting of one- and two-body currents in direct detection of dark matter**. *ArXiv e-prints 1503.04811*