

**A Search for Light Weakly-Interacting Massive
Particles with SuperCDMS and Applications to
Neutrino Physics**

by

Adam J. Anderson

A.B., University of Chicago (2010)

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


Doctor of Philosophy in Physics

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2015

© Massachusetts Institute of Technology 2015. All rights reserved.



Signature redacted

Author

Department of Physics

15 April 2015

Signature redacted

Certified by

Enectalí Figueroa-Feliciano

Associate Professor of Physics

Thesis Supervisor

Signature redacted

Accepted by ...

Nergis Mavalvala

Professor of Physics

Associate Department Head for Education

A Search for Light Weakly-Interacting Massive Particles with SuperCDMS and Applications to Neutrino Physics

by

Adam J. Anderson

Submitted to the Department of Physics
on 15 April 2015, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Physics

Abstract

Cosmological and astrophysical evidence indicates that 85% of the matter content of the universe is in the form of non-baryonic dark matter. A large number of experiments are currently undertaking searches for weakly-interacting massive particles (WIMPs), the leading class of particle candidates for dark matter. This thesis describes the results of such a search with the SuperCDMS experiment, which uses Ge detectors cooled to 50 mK to detect ionization and phonons produced by particle interactions. We perform a blind analysis of 577 kg d of exposure on 7 detectors targeting WIMPs with masses $< 30 \text{ GeV}/c^2$, where anomalous results have been reported by previous experiments. No significant excess is observed and we set an upper limit on the spin-independent WIMP-nucleon cross section of $1.2 \times 10^{-42} \text{ cm}^2$ at $8 \text{ GeV}/c^2$. We also set constraints on dark matter interactions independent of the dark matter halo physics, as well as on annual modulation of a dark matter signal.

Cryogenic detectors similar to SuperCDMS also have potential applications in neutrino physics. We study several configurations in which dark matter detectors could be used with an intense neutrino source to detect an unmeasured Standard Model process called coherent neutrino scattering. This process may be useful, for example, as a calibration for next-generation dark matter detectors, and for constraining eV-scale sterile neutrinos. In addition, small cryogenic X-ray detectors on sounding rockets with large fields-of-view have the unique ability to constrain sterile neutrino dark matter. We set limits on sterile neutrino dark matter using an observation by the XQC instrument, and discuss prospects for a future observation of the galactic center using the Micro-X instrument.

Thesis Supervisor: Enectalí Figueroa-Feliciano
Title: Associate Professor of Physics

Acknowledgments

Much of the work described in this thesis was made possible by the combined efforts of the superb faculty, students, and technical staff of the SuperCDMS collaboration, whose expertise has made the experiment a success. SuperCDMS is the continuation of a long series of experiments, from the early days at Stanford in the 1990s, through CDMS II in the 2000s, and we all benefit from the collective knowledge produced by that tradition. While I have had the pleasure of working very closely with only a subset of the 100 or so members of SuperCDMS, many people were essential for building the detectors, operating the experiment, and performing data analysis that made my work possible.

I am also deeply indebted to my advisor Enectali Figueroa-Feliciano. Since the beginning of my PhD, he provided me with substantial trust and freedom, yet he constantly pushed for my involvement in interesting projects, including the light WIMP search and neutrino physics studies. His approach of nudging my direction without micromanaging my work has been ideal for me. He is also a first-class experimentalist with a knack for immediately asking the right question in data analysis and figuring out the most elegant solutions to problems in the lab.

When I arrived in the CDMS group at MIT, Scott Hertel, Steve Leman, and Kevin McCarthy introduced me to dark matter searches, patiently answered my endless questions—no matter how inane—and introduced me to the physics of our detectors. During and after their departure from MIT, I had the immense privilege of working with Julien Billard. After two years of daily French-length lunches and coffees, endless discussions, and late nights of working together, I can safely say that I have never learned more from someone in so short a time. His untrammelled rigor and work ethic always encouraged me to do better and more precise work. He is a great colleague, a superb physicist, but an even better friend: he also taught me how to drive, how to sail, and most crucially, how to say, “Je voudrais avoir une bière, s’il vous plaît.” During the same time, it has been great fun to work with Alex Leder and his infectious “charm offensive”, now poised for success in the CUORE experiment.

At MIT, outside of CDMS, I have enjoyed occasional collaborations with Joe Formaggio and Janet Conrad on neutrino physics topics, which form a part of this thesis. They both have a unique ability to think outside of the box and have been extremely supportive and encouraging in my work with them. The members of the Micro-X project in Tali’s group, Meredith Danowski, David Goldfinger, Phil Oakley, Sarah Trowbridge, John Rutherford, and Patrick Wikus, deserve credit for tolerating our rampant theft of their vacuum parts in the lab and helping us with cryogenic problems. David Goldfinger and Daniel Castro patiently tolerated my ignorance of X-ray astrophysics during our studies of sterile neutrino dark matter. And Yoni Kahn succeeded in interesting me in the halo-independent methods and phenomenological studies that he has undertaken with Patrick Fox and Matthew McCullough.

Within the CDMS collaboration, the core members of the low-threshold analysis team, Julien Billard, Ray Bunker, Lauren Hsu, Elías Lopez-Asamar, Kristi Schneck, and Brad Welliver, were critical for producing the CDMS result described in this thesis. Lauren Hsu’s relentlessness and meticulous attention to detail as analysis co-

ordinator were responsible for the project finishing successfully and before the G2 downselect. Blas Cabrera, Wolfgang Rau, Bernard Sadoulet, Richard Schnee, Steve Yellin, and many others consistently provided valuable constructive advice. The efforts of Daniel Brandt, Steve Leman, Peter Redl, Kristi Schneck, and the rest of the detector Monte Carlo group have become increasingly important as our simulations mature and are belatedly integrated into the analysis efforts. No thesis is complete without thanking Matt Pyle for having his trenchant advice and generalized ridiculousness always just a phone call away. Many other CDMS grad students and postdocs too numerous to list have provided useful advice on everything from cryogenics to detectors to data.

I first became interested in neutrinos, dark matter, and things difficult to detect while working for Ed Blucher, who is responsible for getting me started with physics research. And in a weirdly fateful sort of way, my entire academic trajectory to experimental physics, through Chicago and MIT, probably ultimately derives from the inimitable influence of Ron Zaraza, like many others before me.

On the personal front, I have been fortunate to have had the company of many bright friends and flatmates. During first year, sailing with the crew of the S.S. Erie—Carl Brozek, Tim Boutelle, and Jonathan Sellon—from the southern shores of Lake Michigan was a great pleasure until her untimely demise in the Charles River basin during tropical storm Mimi. Since then, the staff of the Pearl of Cambridgeport—Carlo Bozo, Brohannes Brahms, and more recently Emanuele Leoncini—have been hilarious and fun friends. The world will never know such a good B&B. Yoni Kahn again deserves credit for getting me into the GRE project, being understanding of my research commitments, and repeatedly showing me that ambitious ideas do succeed if you believe in them. All of the “astrograds” deserve thanks, for their friendliness and for many great lunchtime discussions. And most of all, the tolerant companionship of Ailsa Keating, through all of my manic moods and schemes, has been a sustained source of happiness. *J’espère qu’il continue.*

Finally, my family has been unwaveringly supportive in all of my pursuits. Erin has always nudged me in the more reasonable direction, and she never hesitated to give me a weekend of her time and a couch to crash on in New York City if I was feeling burnt out in Cambridge. My parents, each in their own very different ways, always encouraged me to chart my own course and not worry about convention or failure. More valuable than these habits of mind, however, has been their consistent love and support through the years.

Contents

1	Dark Matter: Observations and Evidence	17
1.1	Λ CDM cosmology	17
1.2	Measuring Ω_m and Ω_b	19
1.2.1	Galaxies	19
1.2.2	Type Ia Supernovae	22
1.2.3	Cosmic microwave background	22
1.2.4	Big-bang Nucleosynthesis	25
1.3	Summary	27
2	Dark Matter: Candidates and Searches	29
2.1	Weakly-Interacting Massive Particles	29
2.1.1	Freeze-out Mechanism	29
2.1.2	Theories	32
2.2	Direct Detection	34
2.2.1	Elastic Scattering Rate	35
3	iZIP Detector Physics	43
3.1	Detector Basics	43
3.2	Charge	47
3.2.1	Physics and Readout	47
3.2.2	Ionization-based Surface Event Discrimination	51
3.3	Phonons	54
3.3.1	Phonon Propagation	54
3.3.2	Phonon Readout	56
3.3.3	Phonon-Based Background Discrimination	58
3.4	Data Reconstruction	62
4	SuperCDMS Experiment	65
4.1	Soudan Underground Laboratory	65
4.2	Shielding and Layout	66
4.3	Cryogenics	67
4.4	Cold Hardware	69
4.5	Warm Electronics and Triggering	72
4.6	Operation	73
4.6.1	Experimental Running Modes	73

5	Data Analysis for a Light-WIMP Search	75
5.1	Philosophy	75
5.1.1	Experimental Context	75
5.1.2	Analysis Strategy	76
5.2	Data Set and Blinding	79
5.3	Charge Fiducial Volume Cuts	80
5.4	Quality Cuts	81
5.4.1	Data Period Cuts	81
5.4.2	Pulse Quality Cuts	87
5.4.3	Quality Cut Efficiencies	93
5.5	Trigger and Analysis Thresholds	97
5.6	Nuclear Recoil Energy Scale	101
5.7	Background and Detector Simulations	106
5.7.1	Physics Simulations	106
5.7.2	Detector Pulse Simulation	117
5.8	WIMP Consistency Requirements	122
5.9	Discrimination Algorithms	123
5.9.1	Signal Model	125
5.9.2	Cut-Based Algorithm	125
5.9.3	Boosted Decision Tree	130
5.10	Background Estimates and Systematics	136
5.10.1	Statistical Uncertainty from Background Normalizations	136
5.10.2	^{210}Pb Normalization	136
5.10.3	Yield Shape Normalization	139
5.10.4	Finite Monte Carlo Statistics	142
5.10.5	Combined Estimates	143
5.11	Selection Criteria Efficiencies	145
6	Constraints on Light WIMPs	149
6.1	Candidate Events	149
6.2	Comparison with Background Expectations	149
6.3	Constraints on Spin-Independent Dark Matter	154
6.4	Constraints on Annual Modulation	156
6.4.1	Maximum Modulation Method	158
6.4.2	Unbinned Maximum Likelihood	159
6.5	Halo-Independent Constraints	164
6.5.1	Unbinned Halo-Independent Limits	165
6.5.2	Comparing XENON10, LUX, and CDMS II Si with SuperCDMS168	
6.6	Future Directions	170
7	Coherent Neutrino Scattering	173
7.1	Coherent Neutrino Scattering	174
7.2	Coherent Neutrino Scattering in Dark Matter Detectors	175
7.2.1	Detection	175
7.2.2	Rates and Time to Discovery	176

7.2.3	Physics with a Detector Close to the Neutrino Source	183
7.2.4	Conclusions	184
7.3	Oscillation of eV-Scale Sterile Neutrinos	185
7.4	Sterile Searches at DAR Neutrino Source	185
7.4.1	Motivation for an Active-to-Sterile Oscillation Search	185
7.4.2	The Neutrino Source	186
7.4.3	Experimental Setup	187
7.4.4	Measurement Strategy and Sensitivity	193
7.4.5	Conclusions	196
7.5	Sterile Searches with Electron-Capture Neutrino Sources	197
7.5.1	The ^{37}Ar Source	199
7.5.2	The Detector	200
7.5.3	Sensitivity and Outlook	209
7.5.4	Summary	214
8	Constraints on Sterile Neutrino Dark Matter	217
8.1	Introduction	217
8.2	keV Dark Matter with Rockets	219
8.2.1	Dark Matter Interpretations of X-ray Lines	219
8.2.2	Dark Matter Signal for Large FOV Observations	220
8.3	Analysis of XQC Data	222
8.4	Estimates for Future Observations	227
8.4.1	Backgrounds	228
8.4.2	Signal and Sensitivity Estimates	232
8.5	Sterile Neutrino Interpretation	234
8.6	Conclusion	234
8.A	Statistical Model for XQC	236

List of Figures

1-1	Example of galactic rotational velocity curves.	20
1-2	Hubble image of Bullet Cluster.	21
1-3	Hubble diagram for Type Ia supernovae.	23
1-4	CMB temperature angular power spectrum measured by Planck. . . .	24
1-5	Relative abundances of light nuclei produced during BBN.	26
2-1	Comoving number density of dark matter as a function of time in early universe.	30
2-2	Best-fit WIMP mass and cross section from global fits of cMSSM to LUX and LHC data.	33
2-3	Asymmetric DM parameter space in which neutron stars collapse to form black holes.	34
2-4	Schematic of the secluded DM scenario.	35
2-5	Rates for elastic scattering of 5, 20, and 100 GeV WIMPs on Ge. . . .	41
2-6	Rates for elastic scattering of WIMPs on Si, Ar, Ge, and Xe for 5 GeV and 20 GeV WIMPs.	42
3-1	3-inch iZIP detector.	44
3-2	Phonon and ionization channel layout on the mask of the iZIP.	45
3-3	Schematic cross section of iZIP electric field.	48
3-4	Electronic band structure of Ge.	50
3-5	Ionization yield of Lindhard for Ge ($k = 0.157$).	51
3-6	Simulation of oblique electron propagation in Ge.	52
3-7	Circuit diagram of ionization readout electronics.	52
3-8	Data from ^{210}Pb surface event source on T3Z1.	53
3-9	Distributions of frequency and angle with respect to the electric field direction of Luke phonons in Ge and Si.	55
3-10	Two images of a QET sensor with TES highlighted.	57
3-11	Schematic of superconducting gap in the Al / W interface of the QET. .	57
3-12	Simplified diagram of TES circuit and readout.	59
3-13	Example phonon pulse showing position-dependent and ballistic components.	61
3-14	Data from ^{210}Pb surface event source on T3Z1.	61
4-1	Muon fluxes at various underground labs.	66
4-2	Cross-sectional views of SuperCDMS shielding at Soudan.	68

4-3	Arrangement of detectors in towers.	69
4-4	Cross-sectional view of tower assembly used to mount iZIP detectors.	70
4-5	Pictures of various cold hardware components.	71
4-6	Signal-to-noise as a function of Butterworth filter bandpass.	73
5-1	Selection of spin-independent dark matter limits circa 2013.	77
5-2	Basic flow of the analysis strategy used for the light WIMP search.	78
5-3	Schematic of charge fiducial volume definitions.	82
5-4	Variation in tower 3 temperature during Run 133.	83
5-5	Phonon channel T_c by detector, compared with phonon gain as a function of temperature.	84
5-6	Phonon Optimal Filter χ^2 for Ba calibration data on T2Z1	87
5-7	Histogram of number of charge and phonon triggers in commissioning data	89
5-8	Traces from a glitch event compared with a physics event	90
5-9	Typical low-frequency noise total phonon pulse compared with typical physics pulse.	91
5-10	Schematic of low-energy events that rail in the optimal filter search window	92
5-11	Phonon optimal filter start time as a function of energy.	93
5-12	Quality cut efficiency as a function of total phonon energy, estimated by Monte Carlo pulse simulation.	96
5-13	Effect of trigger requirement in Ba data.	98
5-14	Example fits of trigger efficiency to Ba calibration during early 2013.	100
5-15	Analysis thresholds as a function of time during Run 133.	102
5-16	Live time-averaged trigger and analysis threshold efficiencies.	103
5-17	Example of ^{252}Cf calibration data near NR band, for T4Z3.	105
5-18	Example of best-fit charge model, uncertainty, and comparison with Lindhard, for T4Z3.	105
5-19	L-shell activation lines in post-Cf data.	110
5-20	K-shell activation lines in all Run 133 data.	111
5-21	^{210}Pb decay chain.	112
5-22	^{210}Pb simulation geometry.	113
5-23	Recoil energy spectra of components of ^{210}Pb simulation for T2Z2.	114
5-24	Alphas in Run 133 WIMP-search data.	116
5-25	Typical 100 keV phonon and charge traces from Ba data used for gamma pulse simulation.	119
5-26	Events selected as templates for pulse simulation of ^{210}Pb background on T1Z1.	121
5-27	Reweightings of Cf data for construction of WIMP signal model used in cut tuning.	126
5-28	Projections of signal model KDE in the 3-4 keV bin on T2Z1.	128
5-29	Projections of background model KDE in the 3-4 keV bin on T2Z1.	129
5-30	Example of optimized cut boundaries as a function of total phonon energy on T4Z3 from cut-based discrimination algorithm.	131

5-31	Example of optimized cut boundaries in bin of 4-6 keV total phonon energy on T4Z3 from cut-based discrimination algorithm.	131
5-32	Input variables to the BDT on T2Z2.	134
5-33	Pdfs of BDT output trained on 10 GeV WIMP.	135
5-34	Cdfs of BDT output trained on a 10 GeV WIMP.	137
5-35	Slices of sensitivity parameter as a function of BDT cut position on each detector.	138
5-36	Pdf of scale factor parameterizing uncertainty for ^{210}Pb background. .	140
5-37	Half gaussian distribution for sidewall event yield systematic uncertainty.	143
5-38	PDF of systematic uncertainty in total background, coadded over detectors.	144
5-39	Efficiency of BDT and WIMP consistency cuts.	147
5-40	Combined acceptance of event selection criteria.	148
6-1	Unblinded WIMP-search data.	150
6-2	Background rates by detector compared with observed event rates. . .	152
6-3	Cross-sectional view of T5Z3 electric field.	153
6-4	Comparison of inner charge quantities on T5Z3 compared to fully-functional detector.	154
6-5	Comparison of data in preselection region to background model for 10 GeV/ c^2 BDT.	155
6-6	Limits on spin-independent WIMP-nucleon cross section.	157
6-7	Upper limit on counts in maximum modulation method for annual modulation constraints	160
6-8	Maximum likelihood fits to annual modulation.	162
6-9	Maximum likelihood constraints on modulation rate and phase. . . .	163
6-10	Comparison of annual modulation constraints with CDMS II and Co-GeNT.	164
6-11	Halo-independent comparison between XENON10, LUX, SuperCDMS, and CDMS II Si.	169
6-12	DM halo-independent constraints from XENON10, CDMSII, and SuperCDMS data.	170
7-1	Energy distribution of neutrinos in a DAR source	177
7-2	Recoil energy distributions for coherent scattering 1.5 km from a DAR neutrino source for Ne, Ar, Ge, and Xe. The rates reported assume 100% detection efficiency.	178
7-3	CNS recoil energy distribution for a Ge experiment 2.3 km from a DAR neutrino source.	180
7-4	Expected sensitivity of dark matter experiment with neutrino calibration source.	182
7-5	Expected ^{39}Ar background energy spectrum using measured performance and theoretically-achievable performance.	191
7-6	Expected signal and background rates for sterile neutrino search. . . .	192

7-7	Sensitivity to LSND allowed parameter space using a Ge-based detector in the baseline scenario.	195
7-8	Sensitivity to LSND allowed parameter space using a Ge-based detector in the dedicated-running scenario.	196
7-9	Sensitivity to LSND allowed parameter space using a Ge-based detector in the baseline scenario.	197
7-10	Sensitivity to disappearance amplitudes accessible with charged current searches.	198
7-11	Schematic of bolometer thermal model.	202
7-12	Calculation of threshold for Si and Ge targets from 20–100 g at different operating temperatures.	204
7-13	Simulated current readout for 10–50 eV recoils using the model parameters in Table 7.5.2.	206
7-14	Conceptual schematic of the experimental setup for a bolometric measurement of coherent scattering from a high-intensity ^{37}Ar neutrino source.	207
7-15	Plot of the relative signal error versus signal-to-noise ratio.	213
7-16	Distribution of events on a 500 kg Si array as a function of time of source deployment.	215
7-17	Ratio of data and Monte Carlo for a simulated neutrino oscillation signal, and likelihood curves for a 300 d exposure.	215
7-18	Likelihood contours for a 300 d run with Si and Ge detector and ^{37}Ar and ^{51}Cr neutrino sources.	216
8-1	Common dark matter density profiles.	220
8-2	Comparison of sterile neutrino decay rate in large-FOV observation with observations using convention X-ray telescopes.	221
8-3	Data and fit of data from the fifth flight of XQC.	223
8-4	XQC detector efficiency.	225
8-5	Upper limit on the flux of an unidentified line in the XQC data.	226
8-6	Total efficiency of Micro-X for X-ray detection as a function of energy.	228
8-7	Expected X-ray background from a 20° radius region centered on the galactic center.	229
8-8	Background from a 20° radius region centered on the galactic center.	231
8-9	Mock data in the energy range of interest for the 3.5 keV line.	233
8-10	Expected limit from an observation of a 20° field around the galactic center by Micro-X.	233
8-11	Constraints on decaying sterile neutrino dark matter.	235
8-12	Simulated energy spectrum in the XQC (<i>blue</i>) and Micro-X (<i>green</i>) X-ray absorbers.	239

List of Tables

1.1	Time-dependence of the Robertson-Walker scale factor for universes dominated by different types of matter.	18
3.1	Variables relevant for iZIP physics analysis.	63
5.1	^{252}Cf calibrations used during Run 133 and the subsequent periods of data that were left unblinded because of high rates of activation. . . .	81
5.2	Live time efficiency of cQhighnoise_v53.	85
5.3	Live time efficiency of cBadOFRes_v53.	86
5.4	Dates with data series excluded by cStableTrigThresh_133_LT due to rapid changes in trigger threshold that prevented the trigger efficiency from being reliably estimated.	86
5.5	Live time efficiency of quality cuts.	94
5.6	Gamma rates.	107
5.7	Electron capture isotopes detectable in SuperCDMS.	108
5.8	Activation line rates.	109
5.9	Sideband regions used to normalize background simulation.	117
5.10	Correspondence between physics backgrounds and pulse simulation template categories.	118
5.11	Statistical uncertainty of background component normalizations. . . .	139
5.12	Uncertainty of ^{210}Pb background normalization.	141
5.13	Systematic uncertainty of sidewall beta background due to ionization yield.	142
5.14	Background estimates by detector.	145
6.1	Candidate events and energies by detector.	151
7.1	CNS rates for various detector configurations and thresholds.	178
7.2	Detector parameters assumed for GEODM.	179
7.3	Experimental parameters assumed for GEODM.	181
7.4	Experimental parameters for GEODM with a source close to the detector. .	184
7.5	Experimental parameters for a short-baseline sterile neutrino search using CNS.	188
7.6	Detector parameters for a short-baseline sterile neutrino search using CNS.	192
7.7	List of properties of selected electron capture neutrino sources.	199

7.8	Model parameters for a 50 g Si detector coupled to a Mo/Au TES operated at 15 mK.	205
7.9	List of relevant source and detector parameters used for sensitivity analysis.	211
7.10	List of systematic uncertainties expected for a low-threshold germanium detector array.	214
8.1	Basic signal and rates expected for an observation of the galactic center.	222
8.2	Flux limits on a line at 3.53 keV from XQC data.	227
8.3	LMXBs included in the background model for the Micro-X observation.	231
8.4	Components of the XQC background PDF.	237
8.5	Parameters of the XQC background PDF.	238

Chapter 1

Dark Matter: Observations and Evidence

The dark matter problem is a classic example of a “known unknown”. It is “known” in the sense that numerous lines of evidence across all time scales, across vast length scales, and using very diverse physical probes, suggest that the majority of matter in the universe is a feebly-interacting, non-baryonic type that we call “dark matter”. Yet it is “unknown” in the sense that we have not detected dark matter directly in any terrestrial laboratory experiment. Without study in a controlled laboratory setting, we do not know anything precise about the particle that is the dark matter. We do not know its mass, its interactions with normal matter, and if the dark matter is a single particle or a “dark sector” with complex dynamics like the matter we see in the universe around us. This section describes the basic matter-energy budget of the universe and a sampling of the pieces of evidence that underlie our current understanding of the properties of dark matter.

1.1 Λ CDM cosmology

Nearly all of modern cosmology, including our understanding of the matter content of the universe, is based on the Robertson-Walker spacetime metric or approximations to it. Observations of the large-scale structure of the universe indicate that it is both approximately *isotropic*—it looks the same in all directions—and *homogeneous*—it looks the same at each point. The universe is obviously not exactly homogeneous or isotropic because of the existence of galaxies and other structure, but this is a useful approximation because of a theorem in general relativity [10] that states that the only possible metric for a homogeneous and isotropic universe is given by

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

where K is a constant, $a(t)$ is a function called the *scale factor* that sets the overall length scale at each time, and where we have used spherical coordinates in which t is time, (r, θ, ϕ) are the usual spherical polar spatial coordinates, and $d\Omega = d\theta^2 +$

non-relativistic matter	$a(t) \propto t^{2/3}$
relativistic matter (radiation)	$a(t) \propto t^{1/2}$
vacuum energy	$a(t) \propto \exp Ht$

Table 1.1: Time-dependence of the Robertson-Walker scale factor for universes dominated by different types of matter.

$\sin^2 \theta d\phi^2$ is the solid angle. There are three values of K that produce unique spacetime geometries: $K = 0$ (Euclidean), $K = -1$ (hyperbolic), and $K = +1$ (spherical) [11].

Given a metric, the dynamical behavior of the universe is given by solving Einstein's equation

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R + \Lambda g_{\mu\nu} = 8\pi G T_{\mu\nu}, \quad (1.2)$$

where $R_{\mu\nu}$ and R are the Ricci tensor and scalar respectively, G is Newton's constant, $T_{\mu\nu}$ is the stress-energy tensor, and Λ is the cosmological constant. To calculate the equations of motion for the Robertson-Walker metric, we therefore need to specify a stress-energy tensor for the matter content of the universe. On large scales, galaxies, dark matter, and vacuum energy all behave as a *perfect fluid*, having no viscosity or heat flow, with a stress-energy tensor given by

$$T_{\mu\nu} = \text{diag}(-\rho, p, p, p). \quad (1.3)$$

Plugging in numbers, we arrive at the Friedmann equation

$$H^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{K}{a^2}, \quad (1.4)$$

for the dynamics of the expansion of spacetime. The dimensionless ratio, the Hubble constant, $H = \dot{a}/a$ measures the rate of expansion. Observations indicate that we live in flat space that contains vacuum energy, nonrelativistic matter, and relativistic matter.

The Friedmann equation is quite useful because it provides a means by which to calculate the time-evolution the scale factor $a(t)$, telling us under which conditions the universe is expanding or contracting as a function of time, provided we know how the energy density ρ of the universe varies with the scale factor. For non-relativistic matter, we have $\rho = \rho_0(a/a_0)^{-3}$ because if lengths grow by some factor, volumes grow by the cube of that factor, causing the energy density to also drop by the cube. For relativistic matter, such as a photon gas, we know that $\rho \propto T^4$ and the entropy $S \propto VT^3$. Since S is conserved if space expands adiabatically, we must have $T \propto 1/a$, and therefore $\rho = \rho_0(a/a_0)^{-4}$. Finally, vacuum energy is characterized by a constant energy density $\rho = \rho_V$. For each of these three energy densities, we can easily solve the Friedmann equation, giving the time-dependence shown in Table 1.1. In all three cases, we remarkably find that the universe expands with time.

If Ω_Λ , Ω_m , and Ω_R are the fractions of each component, respectively, then we can

rewrite the total energy density as

$$\rho = \frac{3H_0^2}{8\pi G} \left[\Omega_\Lambda + \Omega_m \left(\frac{a_0}{a} \right)^3 + \Omega_R \left(\frac{a_0}{a} \right)^4 \right] \quad (1.5)$$

We can further split the matter component of the energy density into two parts $\Omega_m = \Omega_b + \Omega_c$, where Ω_b represents the component due to baryons, and Ω_c represents everything else, which we will see next is nonbaryonic dark matter.

1.2 Measuring Ω_m and Ω_b

A large number of observables, sensitive to very different systematic uncertainties, provide measurements of both Ω_m and Ω_b . While there is minor tension between some of these probes, all observations agree with the basic conclusion that $\Omega_m h^2 \sim 0.15$ while $\Omega_b h^2 \sim 0.02$; that is, the overwhelming majority of the matter in the universe is nonbaryonic. Astrophysical observations furthermore suggest that this matter must be “dark”, in the sense that it has only gravitational and weak-scale interactions.

1.2.1 Galaxies

Galactic rotational velocities

Some of the earliest compelling evidence for dark matter was the observation that a variety of luminous astrophysical bodies move faster than one would expect if they only experienced the gravitational attraction of other luminous matter. The most dramatic evidence of this type is the rotational velocity of galaxies [12]. If we idealize a galaxy as an axisymmetric rotating distribution of matter, then a test particle of mass m at a radius r from the center of the galaxy ought to experience a tangential acceleration due to gravity, equal to its centripetal acceleration

$$\frac{mv^2}{r} = \frac{GM(r)}{r^2}, \quad (1.6)$$

where $M(r)$ is the total enclosed mass at radii less than r . If a galaxy consists solely of luminous matter, then we would expect that at large radii, the outermost luminous matter would have rotational velocities that scale as

$$v \propto \sqrt{\frac{M(r)}{r}}. \quad (1.7)$$

Although this is an idealized picture, it is a very robust prediction that the rotational velocities should drop at radii beyond the bulk of the luminous mass in galaxies.

Detailed measurements of rotational velocity curves can be performed by measuring the redshift of known atomic lines as a function of radius from the center of a galaxy. The emission lines $H\alpha$ ($n = 3 \rightarrow 2$ transition of H), HI (21-cm transition of H), and CO are particularly useful for this measurement because they tend to be




Figure 1-1: Compilation of galactic rotational velocity curves from [14] measured using a combination of CO, H α , and HI line data from spiral galaxies (including Sb, Sc, SBb, and SBc types).

produced by interstellar gas, which has a particularly low velocity dispersion. HI is further useful because it often extends to radii 3 to 4 times greater than the stars [13].

Extensive measurements of many galaxies using a variety of tracers of the rotational velocity have convincingly shown that rotational velocities are very nearly constant at the highest measurable radii. Figure 1-1 shows the velocity curves for an ensemble of spiral galaxies, clearly showing the constant velocity at high radii. The accepted interpretation of these data is that there exists a halo of non-luminous dark matter surrounding the visible part of the galaxies. Two ingredients are required to prop up the rotational velocities at high radius, even at the edge of the luminous matter distribution: a.) the mass of dark matter must be significantly greater than the mass of luminous matter, and b.) the scale radius of the dark matter halo must be significantly larger than for the luminous galaxies.

Galaxy Clusters

The kinematics of galaxy clusters is a probe of the cluster mass, which can therefore provide additional evidence for the existence of dark matter. Clusters are considered to be more useful than galaxies for measuring the total dark matter density because they sample a larger fraction of the observed universe. Clusters masses can be measured using the virial theorem, x-ray observations, and gravitational lensing.

The virial theorem method [15] relies on the fact that a system in a state of equilibrium (“virialized”), obeys the virial theorem: $2T + V = 0$. Substituting expressions




Figure 1-2: Hubble image of the Bullet Cluster, overlaid with the distribution of the total matter density (*blue*) inferred from gravitational lensing, and the distribution of baryonic matter (*pink*) inferred from X-ray observations by Chandra [16].

for the kinetic energy and the potential, which depend on the total mass of the cluster, we can solve for the mass

$$M = \frac{2 \langle v^2 \rangle}{G \langle 1/r \rangle}, \quad (1.8)$$

where $\langle v^2 \rangle$ is the mean mass-weighted square velocity relative to the center of mass, and $\langle 1/r \rangle$ is the mean inverse separation of galaxies from the center of mass. To measure Ω_m , this estimate of the mass and the cluster luminosity is used to form a mass-to-light ratio M/L . The total matter density of the universe is estimated by multiplying this ratio by an estimate of the total luminosity density of the universe $\rho_m = (M/L)\mathcal{L}$. This gives a value of

$$\Omega_m = 0.15 \pm 0.02 \pm 0.04, \quad (1.9)$$

where the first uncertainty is the systematic in the total luminosity density and the second is the uncertainty from the mass-to-light ratio inferred from clusters [11]. These values of total matter density can be compared with independent probes of the baryonic matter density from BBN, described below, to infer that the majority of matter density in the universe is non-baryonic.

X-ray observations and gravitational lensing of galaxy clusters also provide evidence that the majority of matter is nonbaryonic in nature. The most striking examples of these techniques involve collisions of galaxy clusters, such as the “Bullet Cluster” shown in Figure 1-2. The image consists of two galaxy clusters that are believed to have collided with each other. During mergers of galaxy clusters, indi-

vidual galaxies act like collisionless particles, while the X-ray-emitting intracluster plasma acts like a fluid, experiencing ram pressure. The Hubble space telescope image in Figure 1-2 is overlaid with the distribution of total matter in blue, inferred from gravitational lensing. Chandra observations are used to estimate the distribution of the intracluster plasma, shown in pink. The total matter distribution traces the galaxies, rather than the X-ray-emitting gas that is known to comprise the majority of the baryonic mass of galaxy clusters. This provides direct evidence that the majority the total matter in the cluster is weakly interacting—because its kinematics are similar to the collisionless galaxies—and not comprised of hot X-ray-emitting gas. The interpretation is that during the collision, the hot gas interacted, slowed down, and formed the distinctive shock front visible in Figure 1-2, while the nonbaryonic dark matter and galaxies interacted only minimally in the collision.

1.2.2 Type Ia Supernovae

We observed in Table 1.1 that the expansion rate of the universe depends on its composition. It is possible to measure the expansion history directly using observations of high-redshift Type Ia supernovae. Type Ia supernovae are believed to be produced by a white dwarf in a binary system. White dwarfs are dense stars in which nuclear reactions have stopped occurring, and which are supported only by electron degeneracy pressure. If the white dwarf accretes enough matter from its companion, it can approach the Chandrasekhar limit: the maximum amount of mass that can be supported by electron degeneracy pressure. The ensuing instability leads to high temperatures and pressures that convert light elements into heavy ones, triggering a thermonuclear explosion. Since supernovae always occur near the Chandrasekhar limit, they a fairly uniform amount of fissile material and little variation in peak luminosity [11].

The very high and uniform luminosity of Type Ia supernovae have enabled their use for measuring cosmological distances. This technique was famously used by [17, 18] in the discovery of the accelerating universe. Figure 1-3 shows how different values of matter and dark energy density affect the apparent magnitude of the supernovae as a function of redshift. For flat cosmologies, the preferred matter density from supernovae data is

$$\Omega_m = 0.28^{+0.09}_{-0.08} {}^{+0.05}_{-0.04}, \quad (1.10)$$

where the first error is statistical and the second is systematic [11]. This value of Ω_m is larger than inferred by galaxy clusters, but not unreasonably so given statistical and systematic uncertainties in both measurements.

1.2.3 Cosmic microwave background

The cosmic microwave background (CMB) provides the most precise measurements of the energy densities. After the temperature of the universe drops below the ionization energy of hydrogen, and protons and electrons recombine to form neutral hydrogen, the mean free path of photons in the primordial plasma suddenly rises, exceeding




Figure 1-3: Hubble diagram for Type Ia supernovae, from [17]. Upper panel shows the data for two supernovae samples, compared with three cosmologies with different matter densities. Bottom panel shows the residual of the data with a matter-only cosmology




Figure 1-4: CMB temperature angular power spectrum measured in Planck 2013 data release. The green shaded region shows the magnitude of the cosmic variance, which is also included in the error bars, and the green line shows the best fit model. From [19].

the Hubble length, so that the photons are essentially free-streaming. This radiation is imprinted with the density fluctuations that existed at recombination, although interactions between us and this surface of last scattering can slightly modify the pattern of anisotropy on the sky. Figure 1-4 shows the angular power spectrum of temperature fluctuations of the CMB on the sky. The peaks in the power spectrum are the result of acoustic oscillations in the photon-baryon plasma of the early universe. The gravitational attraction of baryons resulted in overdensities that were supported against collapse by radiation pressure from photons. This produced outgoing pressure waves of radiation and baryonic matter from the overdensities. These oscillations are imprinted in the CMB temperature by the series of acoustic peaks in the angular power spectrum.

Fitting the Λ CDM prediction for the angular power spectrum provides strong constraints on many of the parameters in the model. In particular, when the baryon density $\Omega_b h^2$ is large compared to the total matter density, odd peaks become larger than even ones, due to the effect on the acoustic oscillations. This also reduces the diffusion length of photons, which pushes the power to smaller angular scales.

The strongest constraints on cosmological parameters using CMB data now come from the Planck satellite. Because of degeneracy with the Hubble parameter, the CMB measurements are primarily sensitivity to the reduced baryon and cold dark

matter densities $\Omega_b h^2$ and $\Omega_c h^2$. The Planck 2013 data release reports values of

$$\Omega_b h^2 = 0.02207 \pm 0.00033 \quad (1.11)$$

$$\Omega_c h^2 = 0.120 \pm 0.003, \quad (1.12)$$

corresponding to a universe that is approximately 85% non-baryonic dark matter. Planck also reports a measurement of $\Omega_m = 0.314 \pm 0.020$, in broad agreement with the value from supernovae.

1.2.4 Big-bang Nucleosynthesis

Models of nuclei production in the universe are also sensitive to the baryon density, and provide constraints on $\Omega_b h^2$ that are in agreement with the measurements of the CMB, using completely different observables. At early times, protons and neutrons were kept in thermal equilibrium via weak interactions with an interconversion rate $\Gamma_{np} \propto T^5$, which fell faster than the Hubble expansion rate $H \propto T^2$. This resulted in the ratio of protons to neutron “freezing out” of thermal equilibrium around a temperature of 1 MeV, when Hubble expansion overtook the proton-neutron interconversion rate. As temperature dropped below the binding energy for various nuclei, nuclear reactions could begin to occur to form heavier nuclei. In reality, production of deuterium was delayed slightly because of photodissociation until the temperature dropped well below its binding energy of 2.23 MeV. Because of the low energy particle density at this stage of the universe, only 2-body interactions were possible. This fact, the absence of stable nuclei with $A = 5$ or 8 , and the high Coulomb barrier for the reactions ${}^3\text{He}({}^4\text{He}, \gamma){}^7\text{Li}$ and ${}^3\text{He}({}^4\text{He}, \gamma){}^7\text{Be}$ blocked nucleosynthesis of heavier elements. Only D, ${}^4\text{He}$, and ${}^3\text{He}$ were synthesized in significant quantities, with very small amounts of ${}^7\text{Li}$. Models of big-bang nucleosynthesis (BBN) yield quantitative predictions for the abundances of these light elements. The abundances are importantly parameterized by a single number: the baryon density, normalized to the relic blackbody photon density, η . This simplicity is not surprising, because the relative densities of all particle species during the evolution of nucleosynthesis are fixed by precisely-measured weak and nuclear cross sections.

Abundances of elements can only be observed at times much after nucleosynthesis. Since stars produce heavier elements by nuclear fusion, abundances that are measured too late will not match nucleosynthesis predictions. Light element abundances have historically been measured by a variety of techniques, but more recent measurements of the D/H ratio have been using absorption lines in light from high-redshift quasars. Measurements of the other light element abundances have been made with a wide variety of techniques. Figure 1-5 shows that the abundance measurements are in general agreement—over 10 orders of magnitude—with only a modest discrepancy in the ${}^7\text{Li}$ abundance. The inferred baryon-photon ratio is directly related to the cosmological baryon density, and provides the constraint

$$0.021 \leq \Omega_b h^2 \leq 0.025, \quad (1.13)$$




Figure 1-5: Relative abundances of light nuclei produced during BBN (D , ${}^3\text{He}$, and ${}^7\text{Li}$, in decreasing order from top). Yellow boxes indicate the observed light element abundances, while curved bands show the predictions of the abundances from BBN as a function the baryon-to-photon ratio. The width of the predicted bands incorporates the uncertainty from the nuclear cross sections. The blue vertical band indicates the baryon density inferred from CMB measurements, while the pink vertical band indicates the range inferred from the BBN measurements (95% C.L.). From [20].

at 95% C.L. [20], which again is in good agreement with the CMB, and well below the total matter density.

1.3 Summary

In the previous sections, we have seen that Λ CDM generally provides a good fit to data from the Type Ia supernovae, CMB, and BBN. Supernovae and CMB data consistently point to a matter density of $\Omega_m \sim 0.3$, while CMB and BBN data imply that only 85% of the matter density can be in the form of normal baryonic matter (e.g. protons and neutrons). Furthermore, the presence of a small baryonic and large non-baryonic component of matter seems to be reflected in the kinematics and X-ray emission of both galaxies and galaxy clusters. In other words, the evidence for some kind of dark matter is overwhelming and paints a consistent picture. In the next section, we discuss possible particle models for this non-baryonic component and strategies for detecting it directly.

Chapter 2

Dark Matter: Candidates and Searches

Given the evidence described in the previous chapter, we can enumerate properties that any theoretical dark matter model must satisfy. The dark matter must be massive, so that we observe its gravitational interactions. It must have very weak self-interactions and electromagnetic interactions, to avoid constraints on self-interaction from the bullet cluster and the simple fact that we do not see significant emission in any part of the electromagnetic spectrum. It must be nonbaryonic, in order to evade constraints from BBN and the CMB. It must be stable on cosmological time scales, so that it remains in galaxies at the present epoch. And it must have a viable production mechanism early in the history of the universe. There are a wide variety of candidates that satisfy these properties, such as axions [21], sterile neutrinos [22], primordial black holes [23], and weakly-interacting massive particles (WIMPs) [24]. In this chapter, we will focus on the WIMPs because they are the candidate sought by SuperCDMS. Sterile neutrino dark matter will be discussed in Chapter 8.

2.1 Weakly-Interacting Massive Particles

WIMPs are class of stable particles that possess gravitational and weak interactions and are defined by their production via thermal freeze out in the early universe. None of the particles in the Standard Model (SM) seem to be a good candidate for dark matter. The neutrino would naively appear to be a natural WIMP candidate, but cosmological constraints on the neutrino mass imply that it cannot comprise more than a small fraction of the dark matter density. The existence of dark matter therefore points to new particle physics requiring modification of the SM.

2.1.1 Freeze-out Mechanism

The original motivation for WIMPs came from the observation that thermal freeze out in the early universe provides a mechanism for producing DM with the relic density that we observe today, using weak-scale masses and annihilation cross sections. The




Figure 2-1: Comoving DM number density (left axis) and thermal relic density (right axis) as a function of temperature (bottom axis) and time (top axis), illustrating the freeze-out mechanism in the early universe. The solid gray line is the trajectory for an annihilation cross section that produces the correct relic density, while the shaded regions are for cross sections that differ by 1, 2, and 3 orders of magnitude. From [25].

freeze-out story is as follows. In the hot early universe, the DM was in thermal equilibrium with SM particles, being produced and annihilated at equal rates. As the universe expanded, it also cooled. Once the universe cooled below the DM mass m_χ , its production was energetically blocked, so the DM particles could only annihilate, producing exponential Boltzmann suppression of the number density. Eventually the universe expanded so much, however, that particles could not “find” each other to annihilate, and the comoving number density reached the constant *relic density* that we observe today. This qualitative picture is illustrated in Figure 2-1.

It is useful to repeat this argument in slightly more quantitative terms along the lines of [25] and [24] to gain some additional insight into the physics of freeze out. In thermal equilibrium, the number density of DM particles just given by the integral over the usual Fermi-Dirac or Bose-Einstein distributions

$$n_\chi^{eq} = \frac{g}{(2\pi)^3} \int f(\mathbf{p}) d^3\mathbf{p}, \quad (2.1)$$

where g is the number of internal degrees of freedom of the DM. The evolution of the

number density over time is described by the Boltzmann equation

$$\frac{dn_\chi}{dt} = -3Hn_\chi - \langle\sigma_A v\rangle ((n_\chi)^2 - (n_\chi^{eq})^2). \quad (2.2)$$

The first term on the right accounts for the expansion of the universe. The factor $\langle\sigma_A v\rangle$ is the thermally-averaged annihilation cross section for the DM. This term and the factors after it encode the net rate of annihilation and production of DM.

The Boltzmann equation has no closed form solution, but some rough analytic approximations give insight into the physics. Freeze out occurs roughly when the Hubble expansion rate is comparable to the annihilation rate $H = \langle\sigma_A v\rangle n_f$. At these times, the universe is radiation-dominated, so the Hubble parameter is $H \sim T^2/M_{pl}$. At freeze out, the DM number density n_f is given by the Maxwell-Boltzmann distribution

$$n_f \sim (m_\chi T_f)^{3/2} \exp(-m_\chi/T_f). \quad (2.3)$$

The thermal relic density is given by

$$\Omega_\chi = \frac{m_\chi n_0}{\rho_c} = \frac{m_\chi T_0^3}{\rho_c} \frac{n_0}{T_0^3}, \quad (2.4)$$

where subscript 0 denotes present quantities, and ρ_c is the critical density. Taking all these pieces together, we can solve for

$$\Omega_\chi \sim \frac{m_\chi}{T_f} \frac{T_0^3}{\rho_c M_{pl} \langle\sigma_A v\rangle}. \quad (2.5)$$

Formally speaking, this expression has two moving parts, the DM mass and the cross section. It is well-known, however, that the ratio m_χ/T_f is effectively constant and equal to about 20 in natural units, so the density is inversely related to the cross section. In particle physics theories, the cross section is typically a function of the mass scale of the particle. Using dimensional analysis, we can write the cross section as

$$\sigma_A = k \frac{g_{weak}^4}{16\pi^2 m_\chi^2}, \quad (2.6)$$

where $g_{weak} \simeq 0.65$ is the weak gauge coupling, and k is a model-dependent fudge factor. We immediately see that the preferred mass of the WIMP scales as the 1/2 power of the relic density, so there are a wide range of plausible WIMP masses. Plugging in numbers, we find preferred masses in the range of 100 GeV - 1 TeV. This fact is often called the “WIMP miracle” because particles with roughly weak-scale masses produce the observed DM relic abundance. The argument is obviously not quantitatively precise, as order unity factors can modify equation (2.6) in specific theories, but it does provide a very strong rationale for looking for weak-scale dark matter.

2.1.2 Theories

WIMP candidates exist in many plausible extensions of the Standard Model, including the neutralino of supersymmetric models. We discuss supersymmetry and several models of light dark matter that could explain the anomalies that motivate the low-mass search in Chapter 5.

Supersymmetric Models

Supersymmetry is an additional spacetime symmetry that relates each SM particle to a new *superpartner* particle that differs in spin by $1/2$. Supersymmetry has a number of theoretical virtues. It provides a natural way of solving the hierarchy problem: the problem that radiative corrections to the Higgs boson mass should cause it to be much larger than its observed value unless there is an incredibly fine-tuned cancellation. Since the Feynman rules for fermion and scalar loops have opposite sign, radiative corrections for each particle cancel with that for its superpartner, solving the fine-tuning problem. The additional particles of supersymmetry also cause the gauge couplings of the SM to unify near the GUT scale. And finally, the supersymmetry provides a natural dark matter candidate. For technical reasons outlined in [26], the neutralinos (neutral mixtures of the bino, neutral wino, and neutrino higgsinos) are the most plausible WIMP candidates of the superpartners. The gravitino is also a perfectly viable dark matter candidate, but it is not a WIMP, and it would have interaction cross sections too low to be realistically detectable.

Supersymmetric theories introduce large numbers of free parameters associated with the new particles, so phenomenological studies are typically performed by fixing the majority of the parameters. Examples of the constrained models include constrained MSSM (cMSSM; 4 free parameters), and the phenomenological MSSM (pMSSM; 19 free parameters). LHC constraints have significantly improved constraints on the cMSSM parameter space. Recent global fits of the cMSSM to a large number of searches from LHC and well as the LUX experiment can be used to infer preferred regions of neutralino mass and WIMP-nucleon cross section relevant to DM searches, as shown in Figure 2-4. The preferred masses in this global fit is near 1 TeV, but lighter WIMPs are allowed in the less restrictive parameter space of the pMSSM. Ref. [28] finds available parameter space, albeit finely-tuned, for WIMPs as light as 10 GeV.

Light WIMPs

Neutralino dark matter tends to favor masses in the 100 GeV to 1 TeV range, but lighter WIMPs are natural in other plausible models. The absence of any widely-accepted experimental evidence for supersymmetry at the LHC or in dark matter experiments furthermore justifies considering alternative scenarios.

Asymmetric dark matter (ADM) was originally proposed as a way to relate the dark matter and baryon densities [29], but it has received renewed attention in light of low-mass experimental anomalies in direct detection experiments [30]. The basic scenario is motivated by the observation that the baryon and dark matter densities




Figure 2-2: Best-fit WIMP mass and spin-independent WIMP-proton cross section from global likelihood fits of cMSSM to LUX and LHC data. Red contours are at 68% C.L., while red contours are at 95% C.L.. Dashed lines indicate fits using alternative parameter constraints. From [27].

are quite similar $\rho_{DM}/\rho_B \sim 5$. While this could be a coincidence, the similarity could also be dynamically generated. In asymmetric dark matter, the asymmetry between matter and anti-matter in the SM sector is transferred to the DM sector by some mechanism (or vice-versa). A variety of mechanisms have been proposed, but electroweak sphalerons and higher-dimension operators and renormalizable interactions have received the most attention [31]. After the transfer of the asymmetry, both the SM and DM number densities are comparable

$$n_\chi - n_{\bar{\chi}} \sim n_b - n_{\bar{b}}. \quad (2.7)$$

This, combined with $\rho_{DM}/\rho_B \sim 5$, sets the natural mass scale for asymmetric DM candidates to be $m_\chi \sim 5m_p$.

Like most WIMP candidates, ADM may be constrained by direct detection experiments, described in more detail in the following section. Neutron stars also provide unique constraints on ADM. Because there is an asymmetry of matter over anti-matter, the relic DM does not self-annihilate and can accumulate to very high densities in neutron stars. In the case of scalar ADM, the lack of degeneracy pressure means that densities can be quite high. The accretion of DM can push the neutron star over the Chandrasekhar limit, and causing collapse into a black hole [32].

Another class of DM models that accommodates a wide range of WIMP masses is the secluded DM scenario [33]. In typical WIMP models, there is typically a close link between the annihilation and elastic scattering cross sections. Secluded DM




Figure 2-3: Asymmetric DM parameter space in which neutron stars collapse to form black holes, for different ADM relic densities (*gray regions*). Red line corresponds to the local density in the solar neighborhood. Right panel assumes that the ADM forms a Bose-Einstein condensate and DM in the BEC collapse to form a black hole. In the diagonally-hatched region, the black hole will evaporate due to Hawking radiation. In the square-hatched region Hawking radiation interferes with DM accretion by heating the DM. From [32].

decouples these two cross sections by connecting the DM sector to the visible sector by a metastable mediator whose mass is lighter than the DM. This suppresses elastic scattering rates probed by direct detection, but has no effect on annihilation cross sections. The scenario therefore provides a way of evading the increasingly stringent bounds from the direct detection, while not affecting annihilation rates. Concrete examples of this class of models include $U(1)'$ models, in which an additional gauge group is added to the SM. The $U(1)'$ vector boson does not couple directly to the SM, but instead undergoes kinetic mixing with the SM hypercharge gauge bosons. Small mixing angles then suppress elastic scattering rates.

2.2 Direct Detection

Because WIMPs have weak interactions, they can scatter off nuclei in terrestrial experiments. This enables *direct detection* of astrophysical DM from the galactic halo in a terrestrial experiment. Direct detection has several attractive features. Unlike searches at colliders, whose interpretation depends on the assumed model of dark matter production, direct detection experiments directly probe elastic scattering cross section in a manner that is fairly independent of the underlying particle physics. Direct detection also detects the astrophysical dark matter directly. In contrast, even if new weakly-interacting particles were produced at the LHC, it would be difficult to know whether the observed particle constitutes the dark matter without confirmation from the direct detection.




Figure 2-4: Schematic of the secluded DM scenario. Separate DM and SM sectors are coupled by a mediator. If the mediator is light compared to the DM mass, then elastic scattering rates in direct detection experiments can be strongly suppressed, even with weak-scale WIMP annihilation cross sections that could be detectable by astrophysical observations. From [33].

2.2.1 Elastic Scattering Rate

Calculating the direct detection rate in a terrestrial experiment gives important insight into the variables that are relevant for designing and optimizing an experiment, as well as the major systematic uncertainties that affect the interpretation of experimental results.

Consider the case of a WIMP of speed v and number density n scattering off a target nucleus N . By definition, the differential rate as a function of the target nucleus recoil energy E_R is given by

$$\frac{dR}{dE_R} = nv \frac{d\sigma}{dE_R}(v, E_R). \quad (2.8)$$

Since the WIMP mass is unknown, it is typically more useful to write the number density as $n = \rho_0/m_\chi$, where ρ_0 is the approximately-known local DM density. Galactic WIMPs are characterized by a continuous distribution of speeds $f(v)$, often called the WIMP *velocity distribution function* (VDF), not just a single speed. And we are typically interested in the rate of events per unit mass of target material, rather than the rate for a single target nucleus, so we must divide equation (2.8) by m_N . With these modifications, the WIMP-nucleus scattering rate takes its usual form [34]

$$\frac{dR}{dE_R}(E_R) = \frac{\rho_0}{m_\chi m_N} \int_{v_{min}(E_R)}^{\infty} dv \, v f(v) \frac{d\sigma}{dE_R}(v, E_R), \quad (2.9)$$

where we have integrated the velocity distribution from the minimum velocity $v_{min}(E_R)$ capable of producing a recoil of energy E_R .

Cross Section

The differential cross section $d\sigma/dE_R$ combines kinematic factors with factors related to the underlying particle physics model. From non-relativistic elastic scattering

kinematics, the recoil energy of the nucleus is given by

$$E_R = \frac{\mu_N^2 v^2 (1 - \cos \theta^*)}{m_N}, \quad (2.10)$$

where $\mu_N = m_N m_\chi / (m_N + m_\chi)$ is the reduced mass of the WIMP-nucleon system and θ^* is the scattering angle in the center-of-mass frame. Since the WIMP-nucleus system is a non-relativistic scattering problem, we can treat it in the usual way with the incident WIMP as a plane wave scattered by the spherically symmetric potential of the nucleus, whose solutions can be obtained by the partial wave expansion. One can easily show in the partial wave expansion that the scattering is isotropic in the low-energy limit applicable to WIMP-nucleus scattering [35]. In this case, we have

$$\frac{d\sigma}{d\theta^*} = \int \frac{d\sigma}{d\Omega^*} \sin \theta^* d\phi^* = \frac{\sigma_0}{2} \sin \theta^*, \quad (2.11)$$

and therefore

$$\frac{d\sigma}{dE_R} = \frac{d\sigma}{d\theta^*} \frac{d\theta^*}{dE_R} = \frac{m_N \sigma_0}{2\mu_N^2 v^2}, \quad (2.12)$$

where σ_0 is the cross section. The differential cross section is then typically written as

$$\frac{d\sigma}{dE_R} = \frac{m_N \sigma_0 F^2(E_R)}{2\mu_N^2 v^2}, \quad (2.13)$$

where $F^2(E_R)$ is the *nuclear form factor*. The form factor encodes the fact that as the momentum transfer increases, the scattering becomes able to probe more of the structure of the nucleus, and the coherence of the interaction is lost.

One calculates the form factor by solving the scattering problem using the Born approximation. To first order, the scattering amplitude is given by

$$f(q) = -\frac{1}{4\pi} \frac{2m}{\hbar^2} \int V(\mathbf{r}) e^{i\mathbf{q}\cdot\mathbf{r}} d^3r, \quad (2.14)$$

where q is the momentum transfer, and $V(\mathbf{r})$ is the nuclear potential. The differential cross section is then given by

$$\frac{d\sigma}{d\Omega} = |f(q)|^2. \quad (2.15)$$

For a sufficiently short range force, the potential can be approximated by the nuclear density function; in other words, $V(\mathbf{r}) = k\rho(\mathbf{r})$. If we insert this ansatz into equation (2.14) and collect all the q -dependence, we can split the cross-section into a factor for $q = 0$, and the form factor that encodes all of the q -dependence

$$\frac{d\sigma}{d\Omega} = \sigma_0 F^2(q), \quad (2.16)$$

where

$$F(q) = \int \rho(r) e^{i\mathbf{Q}\cdot\mathbf{r}} d^3r = \frac{4\pi}{q} \int_0^\infty \rho(r) r \sin qr \, dr. \quad (2.17)$$

We see that the form factor is simply the Fourier transform of the nuclear density, and the problem is thus reduced to specifying a nuclear density function. One common choice is to choose $\rho(r)$ to be the convolution of two functions

$$\rho(r) = \int \rho_0(r') \rho_1(r - r') d^3 r'. \quad (2.18)$$

In this expression $\rho_0(r)$ is a kind of idealized density distribution, while $\rho_1(r)$ is a regularization function that smooths out the edges of the idealized density distribution. The classic Helm form factor is obtained by choosing a uniform sphere for the idealized density distribution

$$\rho_0(r) = \frac{3Ze}{4\pi r_n^3} \quad (2.19)$$

for $r < r_n$. The smearing function is just taken to be Gaussian

$$\rho_1(r) = \frac{1}{(2\pi s^2)^{3/2}} e^{-r^2/2s^2}. \quad (2.20)$$

It turns out to be quite difficult to evaluate the integrals explicitly, so we can just use the convolution theorem. Recall that $\mathcal{F}\{f * g\} = \mathcal{F}\{f\}\mathcal{F}\{g\}$. Since the form factor $F(q)$ is just the Fourier transform of the density function, we just have

$$F(q) = \mathcal{F}\{\rho_0\}\mathcal{F}\{\rho_1\} \quad (2.21)$$

$$= \mathcal{F}\{\rho_0\} e^{-q^2 s^2/2} \quad (2.22)$$

$$= \frac{3Ze}{4\pi r_n^3} e^{-q^2 s^2/2} \frac{4\pi}{q} \int_0^{r_n} r \sin qr dr \quad (2.23)$$

$$= \frac{3Ze}{q^3 r_n^3} e^{-q^2 s^2/2} (\sin qr_n - qr_n \cos qr_n). \quad (2.24)$$

Lewin and Smith [36] favor the following choice of parameters:

$$r_n = \sqrt{R^2 - 5s^2} \quad (2.25)$$

$$R = [0.89A^{1/3} + 0.3] \text{ fm} \quad (2.26)$$

$$s = 1 \text{ fm}. \quad (2.27)$$

The final ingredient in the rate calculation is to reparameterize the cross-section so that experiments using different target nuclei can be compared easily. The total cross section for WIMP scattering σ_0 depends on the number of nucleons, so this dependence needs to be factored out in order for different experiments to measure the same underlying quantity. For spin-independent interactions, the cross-section can generally be written as

$$\sigma_0 = \frac{4\mu_N^2}{\pi} [Zf^p + (A - Z)f^n]^2, \quad (2.28)$$

where Z is the number of protons in the target nucleus, A is the number of total nucleons, and f^p and f^n are the coupling constants to protons and neutrons respectively. WIMP scattering is conventionally assumed to be isospin-conserving, in which case $f^p = f^n = f$, and the cross section reduces to

$$\sigma_0 = \frac{4\mu_N^2}{\pi} A^2 f^2. \quad (2.29)$$

The total elastic scattering cross section depends on the number of nucleons and therefore varies for different experiments. Since f is the only unknown and is the same for all experiments, it is common to express the cross section in terms of the WIMP-nucleon cross section, normalized to the value for a single proton

$$\sigma_p = \frac{4\mu_n^2}{\pi} f^2. \quad (2.30)$$

Combining the previous two equations, we then have the relation between the WIMP-nucleus and WIMP-nucleon cross sections

$$\sigma_0 = \frac{\mu_N^2}{\mu_n^2} A^2 \sigma_p. \quad (2.31)$$

There are several important points to note about the previous few equations. First, isospin-conserving dark matter is a *coherent* interaction, in the sense that it scales with the square of the number of nucleons A . Heavier nuclei, such as Ge will tend to produce a higher event rate than lighter nuclei such as Si. Second, the total cross section has been parameterized by one unknown quantity σ_n —essentially the total cross section normalized per nucleon—which crucially can be compared across experiments using different targets. And finally, the assumption of isospin conservation is completely arbitrary. Interest in isospin-violating DM was triggered in recent years by [37] which pointed out that $f^p \neq f^n$ can reduce the tension between null and positive results in two experiments using different target nuclei. In particular, [38] proposed that the ratio f^p/f^n could even be tuned specifically so that Xenon-based experiments specifically would have a significantly reduced rate compared to other targets, known as the “Xenophobic” dark matter scenario.¹ Although the xenophobic scenario is an exotic possibility, exotic physics is not needed to produce isospin violation: dark matter comprised of Dirac neutrinos has $f^p/f^n \sim 0.08$ [36]!

Halo Physics

Equation (2.9) depends on the WIMP velocity distribution function, which we need to specify explicitly. To a rough approximation, the DM halo should be non-relativistic (otherwise the DM would not be bound), in thermal equilibrium (because it is old), and sufficiently weakly interacting so that the effect of baryons can be neglected (otherwise we would have already observed the DM through non-gravitational interac-

¹The rate in Xe cannot be tuned exactly to zero using isospin violation, however, because of the presence of different isotopes in naturally occurring Xe.

tions). The phase space density for a DM particle in the halo under these assumptions is given by Maxwell-Boltzmann statistics

$$f(\mathbf{x}, \mathbf{v}) = C \exp(-E(\mathbf{x}, \mathbf{v})/k_B T), \quad (2.32)$$

with an energy

$$E(\mathbf{x}, \mathbf{v}) = \frac{1}{2} m_\chi |\mathbf{v}|^2 + m_\chi \Phi(\mathbf{x}), \quad (2.33)$$

where $\Phi(\mathbf{x})$ is the gravitational potential, and C is a normalization constant. In this case, the velocity distribution function (VDF) is given by marginalizing over the spatial coordinates

$$f(\mathbf{v}) = \int d^3 \mathbf{x} f(\mathbf{x}, \mathbf{v}) \quad (2.34)$$

$$= 4\pi C \int d\mathbf{x} \exp(-m_\chi |\mathbf{v}|^2 / 2k_B T) \exp(-m_\chi \Phi(\mathbf{x}) / k_B T) \quad (2.35)$$

$$\propto \exp(-m_\chi |\mathbf{v}|^2 / 2k_B T) \quad (2.36)$$

$$\propto \exp(-|\mathbf{v}|^2 / 2\sigma_v^2). \quad (2.37)$$

The VDF is therefore just a Maxwell-Boltzmann distribution with dispersion σ_v , which is related to the circular velocity of the DM by $\sigma_v = \sqrt{3/2} v_c$. Since the Milky Way as a finite escape velocity, the velocity distribution is also typically truncated at this v_{esc} . There are a variety of important halo parameters that highlight the sensitivity of DM experiments to systematic uncertainty in the halo physics.

- **Density:** The local density of dark matter is canonically taken to be $\rho_0 = 0.3 \text{ GeV}/c^2 \text{ cm}^{-3}$. Although this value appears to be roughly similar to a variety of measurements, Ref. [39] points out that the origin of this exact value is unclear. Most methods of determining ρ_0 involve constructing a model of the matter distribution in both the dark matter and luminous components of the Milky Way, and then estimating its parameters from observational data such as rotation curve measurements, velocity dispersions of halo stars, and local surface mass density. The chosen parameterizations of the mass densities can have a significant effect on the extracted value of ρ_0 , resulting in significant unmodeled systematic uncertainties. Ref. [39] points out, for example, that [40] and [41] use power law and Einasto DM profiles, respectively, with modest differences in their modeling of the stellar component. They arrive at densities of $0.30 \pm 0.05 \text{ GeV}/c^2 \text{ cm}^{-3}$ and $0.39 \pm 0.03 \text{ GeV}/c^2 \text{ cm}^{-3}$, which suggests the presence of unquantified systematic errors. High-resolution simulations of MW-like galaxies, such as [42], furthermore indicate that assumption of spherical symmetry for the DM density is not accurate, and there maybe overdensities of DM in the stellar disk as high as 40%. Their analysis argues for adopting a local density of $0.466 \pm 0.033(\text{stat}) \pm 0.077(\text{syst}) \text{ GeV}/c^2 \text{ cm}^{-3}$.

Setting aside detailed criticism of the various studies, it is clear that the local density has a large uncertainty. Fortunately, this uncertainty impacts all ex-

periments equally, because it simply scales the overall rate observed in a direct detection experiment. It does not affect the comparison of results between different experiments.

- **Circular speed:** The canonical value of circular speed used by DM experimentalists is $v_c = 220 \pm 20 \text{ km s}^{-1}$, which comes from a review of galactic constants dating from the 1980s [43]. This review surveyed a large number of determinations using a variety of methods and took a simple average of the values and uncertainties. More recent determinations using galactic masers have found slightly higher speeds in the range of 235-255 km s^{-1} , as reviewed in [39], highlighting the systematic uncertainties of different measurements.
- **Escape speed:** The integral over the WIMP VDF is cut off at a maximum speed given by the escape speed of the Milky Way. In analyses of experimental data, this is typically taken to be a hard cutoff. The most recent widely cited value of escape speed is $498 < v_{esc} < 608 \text{ km s}^{-1}$ (90% C.L.) estimated high-velocity stars in the RAVE survey [44]. The escape velocity plays an important role in the interpretation of direct detection experiments at low WIMP mass. At low WIMP mass, experiments are only sensitive to the high-speed tail of the Maxwell-Boltzmann distribution, and a large escape velocity will enable an experiment to be sensitive to lighter WIMPs.
- **VDF shape:** The assumption that the WIMP VDF is of Maxwell-Boltzmann form is a convenient simplifying assumption, but it is likely inaccurate. Modern N-body simulations indicate that there is significant halo-to-halo variation in the VDF among galaxies with similar mass and structure to the Milky Way. The VDF is moreover strongly affected by the merger history of the Milky Way [45].
- **Earth's motion:** The Maxwell-Boltzmann distribution for the WIMP speed is valid in the rest frame of the galaxy, but detectors are located on Earth, which is orbiting both the sun and the Milky Way. This is accounted for by performing a time-dependent, Galilean transformation $\mathbf{v} \rightarrow \mathbf{v} + \mathbf{v}_e(t)$. The earth's motion relative to the galactic rest frame $\mathbf{v}_e(t)$ must include the effects of the sun's motion relative to the local standard of rest (LSR)², the sun's peculiar motion relative to the LSR, and the Earth's orbit around the sun. The circular motion of the LSR is typically taken to be $v_c = 220 \text{ km s}^{-1}$, with compilations of measurements giving a mean value of $v_c = 220 \pm 20 \text{ km s}^{-1}$ [43]. The peculiar motion of the sun in galactic coordinates is given by $\mathbf{v} = (10.0 \pm 0.4, 5.2 \pm 0.6, 7.2 \pm 0.4) \text{ km s}^{-1}$ [34]. And the rotational speed of the earth around the sun is 29.8 km s^{-1} [46]. This component of the earth's annual motion gives rise to an important annual modulation of the rate expected in direct detection experiments.

²The mean motion of material in the neighborhood relative to the galactic frame.




Figure 2-5: Rates for elastic scattering of 5, 20, and 100 GeV WIMPs on Ge with a WIMP-nucleon cross section of 10^{-42} cm^2 . Width of band is maximum range of annual modulation due to the earth's orbit around the sun.

Taking all of these pieces together again, we can rewrite the elastic scattering rate again and explore how it scales with experimental parameters such as mass and nuclear target. The total differential rate is given by

$$\frac{dR}{dE_R}(E_R) = \frac{\rho_0 \sigma_p}{2\mu_n^2 m_\chi} A^2 \int_{v_{\min}(E_R)}^{v_{\text{esc}}} d^3\mathbf{v} \frac{f(\mathbf{v} + \mathbf{v}_e)}{v} F^2(E_R). \quad (2.38)$$

The differential rate for several WIMP masses and a Ge target is shown in Figure 2-5. Several features are worth mentioning. Due to kinematics, the recoil energy spectrum extends to higher energies as the WIMP mass increases. This means that experiments with a nonzero energy threshold have a corresponding mass threshold, below which they are not sensitive to WIMPs.

The recoil spectrum is also approximately exponential, with an unknown slope because the mass is unknown. This renders it very difficult to distinguish from many background sources, except lines. But bump-like features could even be produced by non-minimal dark matter interactions, such as inelastic dark matter [47]. This motivates the use of zero-background experiments, which have expected backgrounds much less than one event, so that any excess of events can be interpreted as significant.

Also important is the scaling of the signal with nuclear target. As we discussed above the total nuclear cross section is proportional to A^2 , so heavier nuclei tend to be preferable. Because experiments have nonzero energy thresholds, however, lighter nuclei may be preferable at low WIMP mass, as shown in Figure 2-6. The Xe-based LUX experiment used a nuclear recoil energy threshold of 3 keVnr [48], and the Ge and Si detectors of CDMS II and SuperCDMS have generally achieved thresholds of about 1-2 keVnr. Given these thresholds, it is possible for Ge and Si to be competitive or better than Xe at low masses, even with smaller exposures.




Figure 2-6: Rates for elastic scattering of WIMPs on Si, Ar, Ge, and Xe for 5 GeV and 20 GeV WIMPs with a WIMP-nucleon cross section of 10^{-42} cm^2 .

Regardless of which single target has the highest rate under the standard assumptions, the use of multiple nuclear target elements in multiple experiments is generally helpful. We have seen in this section that many factors, such as the isospin coupling and halo physics can significantly affect the WIMP interaction rates differently for each target. The use of multiple targets and experiments helps eliminate these “blind spots” in our searches.

Chapter 3

iZIP Detector Physics

The ultimate goal of any direct DM experiment is to achieve as low a background as possible while maintaining as large a fiducial volume and signal acceptance as possible. This goal can either be achieved by designing the detector and readout to have excellent *discrimination* between signal and background or by using *clean materials* that contain less background in the first place. Because SuperCDMS at Soudan uses the same low-background setup as CDMS II with only a slightly larger detector mass, the primary improvements are in the detector design which enables significantly better background discrimination with higher acceptance compared with CDMS II detectors. SuperCDMS at Soudan uses new detectors known as interleaved z -sensitive ionization and phonon (iZIP) detectors, which have much better rejection of the surface events that were the limiting background of CDMS II. In this section, we describe the rudiments of the detector physics and then how the iZIP design combines with the phonon and charge physics to enable background discrimination. A more detailed discussion of the design principles of the iZIP can be found in the theses of Matt Pyle [49] and Scott Hertel [50].

3.1 Detector Basics

To set the stage, let's first consider at a cartoon level what an iZIP is, and what occurs inside of it in the instant just after a particle interacts in the detector. Once we're oriented, we'll spend the rest of the chapter unpacking each of step of the cartoon-level picture. If you are comfortable just knowing *what stuff happens* rather than *why stuff happens*, this sketch should be all you need for understanding the analysis in the subsequent chapters, and you can skip the rest of this chapter. But the details of the detector physics are interesting and useful for understanding some of the design drivers of iZIP detectors.

iZIP detectors consist of 0.6 kg Ge crystals with ionization and phonon sensors photolithographically patterned on the top and bottom sides, as shown in Figure 3-1. The crystals are 1 inch thick and 3 inches in diameter. Ionization and charge sensors are interleaved so that the ionization and phonon signals are read out on both sides. Ionization channels are partitioned into an inner electrode and an outer guard ring,




Figure 3-1: 3-inch iZIPs detector, in housing, used in SuperCDMS at Soudan. The spiral pattern on the top is the parallel array of phonon absorbers and TESs. The ionization electrodes are interleaved between the phonon channels, but they are too small to be visible in the image.




Figure 3-2: Layout of phonon and ionization channels on the iZIP. Each side has interleaved phonon and ionization sensors. Charge readout is partitioned into two channels per side, while phonon readout is partitioned into four channels on each side. Both sides are identical, but clocked by 60° .

while the phonon channels are partitioned into three inner channels and one outer guard ring, as shown in Figure 3-2.

When a nuclear recoil occurs in the detector, a chain of events is set into motion:

- *Prompt phonons* are produced by the recoiling nucleus and electron-hole pairs are excited.
- Because charge electrodes are held at ± 2 V on each side of the crystal, the charge carriers drift to opposite sides of the detector in a few μs .
- While they drift, charge carriers reach a terminal velocity in the Ge, yet they are still moving through a potential. This additional energy is converted into a type of phonons called *Neganov-Luke phonons*¹ [51, 52].
- Charge carriers recombine at the electrodes on the detector surface, and the gap energy is converted into phonons called *recombination phonons*.

¹With no disrespect intended to Dr. Neganov, we'll hereafter refer to these phonons simply as "Luke phonons", following colloquial SuperCDMS usage.

- High-energy prompt phonons have a very short mean free path, which rapidly rises as they downconvert in energy, in a process called *quasi-diffusive propagation*. At this stage, phonons are not homogeneous in the detector, and their distribution contains information about the position of the event (or of the charge carriers that produced Luke phonons).
- Within approximately 100 μs , the mean free path of phonons has increased to become comparable to the size of the detector, and propagation becomes *ballistic*. In the ballistic stage, the distribution of phonons is obviously homogeneous, and position information is completely lost.
- During this phonon propagation process, if a phonon hits a surface that is instrumented by a phonon sensor, it has some chance of being absorbed and measured.
- The phonon sensors, called quasiparticle-assisted electrothermal-feedback transition-edge sensors (QET),² absorb the phonons when they break Cooper pairs in superconducting Al fins, producing Bogoliubov quasiparticles. The quasiparticles diffuse and eventually are trapped in W transition-edge sensors where they are read out. This process occurs continuously over the course of several ms in the iZIP detectors used in Soudan.

The division of the energy between the ionization and phonon systems can be understood as evolving through three separate stages. In the first stage, there is just the recoil energy of the event, which is partitioned into an ionization component and a prompt phonon component

$$E_0 = E_{\text{recoil}} = (E_{\text{eh}})^Q + (E_{\text{prompt}})^P, \quad (3.1)$$

where we have used the superscripts Q and P to denote the portions of the energy in the charge and phonon systems, respectively. Near the end of the charge drift, additional energy has been added to the system in the form of Luke phonons. The energy contained in these phonons arises from the work performed by the electric field during the drift of the charge carriers. The total energy at this stage is

$$E_1 = E_{\text{recoil}} + E_{\text{Luke}} = (E_{\text{eh}})^Q + (E_{\text{prompt}} + E_{\text{Luke}})^P. \quad (3.2)$$

Finally after the charge carriers recombine at the electrodes, the energy of the charge carriers is returned to the crystal in the form of recombination phonons

$$E_2 = (E_{\text{eh, recombination}} + E_{\text{prompt}} + E_{\text{Luke}})^P. \quad (3.3)$$

This seemingly trivial accounting is actually important because it tells us that the full recoil energy of the event is eventually returned to the crystal, and we need only subtract off the Luke component to estimate the recoil energy.

²Quite a mouthful.

In the end, each event leaves us with both a charge energy and a phonon energy. There are two obvious ways that this information can be used to gain insight into the nature of the event. First, electron recoils induced by low-energy gammas and betas produce deposit a larger fraction of their recoil energy into ionization than do nuclear recoils. This fraction, known as the *ionization yield*, is 1 for electron recoils, while it is around 0.1-0.3 for low-energy nuclear recoils in Ge. The discrimination power of ionization yield depends considerably on the efficiency of charge collection and energy resolution, but it is excellent for bulk electron recoils: one internal estimate of the iZIP background fraction from gammas was $< 4.7 \times 10^{-6}$ (90% C.L.) above 8 keVnr (using a charge fiducial volume with 30% efficiency). Electron recoils very near the detector surfaces (e.g. very low-energy gammas or betas), however, typically have reduced ionization collection, and yield-based discrimination may be considerably less effective for these species.

Position information is the other way to separate signal and background because WIMPs interact uniformly in the detector, but betas and low-energy gammas will stop in the outer layer. The trajectories of charges, on average, “point back” to the event location, so segmenting the charge readout into multiple channels provides crucial position information. Events with a large signal in the outer charge channel tend to interact near the detector sidewall, for example. By interleaving ionization and phonon electrodes, the electric field is furthermore manipulated to provide exquisite rejection of surface events on the top and bottom of the detector. This is depicted in Figure 3-3. Interactions in the top or bottom ~ 1 mm of the detector produce a charge signal only on one side, while events in the bulk produce a charge signal on both sides.

Phonons are also sensitive to the position of events. As noted in the bullets above, it takes 100-200 μ s for the phonon population to become fully ballistic and homogeneous in the detector. Before this time, charge channels near the quasi-diffusively expanding phonon “ball” will tend to absorb significantly more phonon energy. This early or prompt component of the phonon pulse in each channel is therefore roughly sensitive to the distance of the event from the channel.

3.2 Charge

3.2.1 Physics and Readout

SuperCDMS detectors are made of Ge, and the behavior of the charge signal depends crucially on several aspects of the semiconductor properties. Ge has a band gap (0.74 eV), meaning that electrons must gain a fixed threshold of energy before they can leave the localized valence band and move freely in the conduction band. At finite temperature, the probability that an energy state is occupied is given by the Fermi-Dirac distribution, and some charge carriers in the tail of the distribution will populate the conduction band. The 50 mK at which SuperCDMS detectors are operated is sufficiently low, however, that charge carriers are essentially “frozen out”, meaning that there are no electrons in the conduction band and the crystal behaves




Figure 3-3: Schematic cross sectional view of iZIP electric field. Grey lines depict the approximate orientation of electric field lines. Charge carriers produced in the center of the detector experience a relatively uniform field, drifting along the red (e^-) and blue (h^+) dotted lines. At the top and bottom surfaces, the scalloped pattern of the electric field causes charge carriers to drift to the same side. The 1-sided charge signal of surface events provides excellent discrimination from bulk nuclear recoils. (Courtesy of Scott Hertel)

as an insulator.

When energy is deposited in the detector, electrons and holes are excited into the conduction band. How many charge carriers are produced? Charge carriers are randomly populated into conduction band states by ionizing radiation, with a mean energy that has been measured to be 2.96 eV / eh pair [54]—greater than the size of the band gap.

As discussed in the previous section the fraction of the recoil energy that produces charge carriers is highly dependent on the incident radiation. The theory of Lindhard is typically used to model the ionization yield of solid state detectors [55]. The ionization yield as a function of recoil energy is summarized concisely in Ref. [36] and given by

$$y(E_R) = \frac{kg(\epsilon(E_R))}{1 + kg(\epsilon(E_R))}, \quad (3.4)$$

with

$$\epsilon = 11.5E_R Z^{-7/3} \quad (3.5)$$

$$k = 0.133Z^{2/3}A^{1/2} \quad (3.6)$$

$$g(\epsilon) = 3\epsilon^{0.15} + 0.7\epsilon^{0.6} + \epsilon. \quad (3.7)$$

In these expressions, the recoil energy E_R has units of keV, Z is the atomic number, and A is the number of nucleons. Figure 3-5 shows the ionization yield using standard values from the Lindhard theory, which is in the range of 0.2-0.3 over the range of recoil energies relevant to most Ge-based DM experiments. This comparatively low yield is one major motivation for using phonon-based readout: the raw signal available is significantly higher. Over several decades, a large number of measurements have confirmed that Lindhard is a reasonable approximation to the true ionization yield. Of interest to DM searches are the SuperCDMS ^{252}Cf calibrations discussed in Chapter 5, as well as measurements using the CoGeNT [56] and TEXONO [57] detectors.

Once charges are produced, they propagate under the influence of the electric field, with momentum described by the equation of motion

$$\hbar \frac{d\mathbf{k}}{dt} = e\mathbf{E}. \quad (3.8)$$

Near a minimum of the conduction band, we can approximate the dispersion relation of the semiconductor to leading order by a quadratic equation

$$E = E_0 + \frac{\hbar^2}{2} [\mathbf{M}^{-1}]_{\alpha\beta} \mathbf{k}_\alpha \mathbf{k}_\beta, \quad (3.9)$$

and the particle acceleration by

$$\mathbf{a} = \frac{e}{\hbar^2} [\mathbf{M}^{-1}]_{\alpha\beta} \mathbf{E}_\beta, \quad (3.10)$$

where $[\mathbf{M}^{-1}]$ is the inverse effective mass tensor. In the special case that the band




Figure 3-4: Electronic band structure of Ge. The four valence bands are shown shaded in gray. The energy gap to each of the electron valleys are indicated on the diagram. From [53].




Figure 3-5: Ionization yield of Lindhard model using standard parameters for Ge ($k = 0.157$).

structure is isotropic (i.e. the energy levels have the same shape in any direction of crystal momentum), then $[M^{-1}]$ reduces to simply a scalar $1/m^*$ [58]. From Figure 3-4 it is clear that the holes have a roughly isotropic band structure because the maximum of the valence band is centered at $[000]$ and approximately parabolic in this neighborhood. Electrons on the other hand have a conduction band minimum in the $[111]$ direction. Since the electric field in SuperCDMS detectors is applied in the $[100]$ direction, this results in electron propagation along four axes oriented at 33 degrees from the electric field axis. This is qualitatively visible in a simulation from [59], shown in Figure 3-6, which first discussed the phenomenon in the context of SuperCDMS detectors.

Once charges propagate through the detector, they are read out by a charge amplifier whose circuit diagram is shown in Figure 3-7. The charge propagation in the detector induces an image charge on the coupling capacitor. The amplifier in turn charges the feedback capacitor, which slowly discharges through the feedback resistor. The typical rise time of an ionization pulse is about $1 \mu\text{s}$ (comparable to the charge drift time and digitization rate, so the pulse shapes are not sensitive to the drift velocity), while the typical fall time is about $50 \mu\text{s}$.

3.2.2 Ionization-based Surface Event Discrimination

The interleaving of ionization and phonon electrodes at alternating potentials produces one-sided ionization readout for surface events on the top and bottom surfaces, as in Figure 3-3. CDMS II detectors had a surface event discrimination of about 1:200, which was not adequate for scaling to much larger detector masses, but the




Figure 3-6: Simulation of oblique propagation of electrons in the $[111]$ crystal direction in Ge. Taken from [59].




Figure 3-7: Circuit diagram of ionization readout electronics. The section labeled “Gate” is a JFET amplifier. C_D is the detector capacitance, C_p is a parasitic capacitance, C_c is a coupling capacitor, and C_{FB} is a feedback capacitor. Taken from [60].

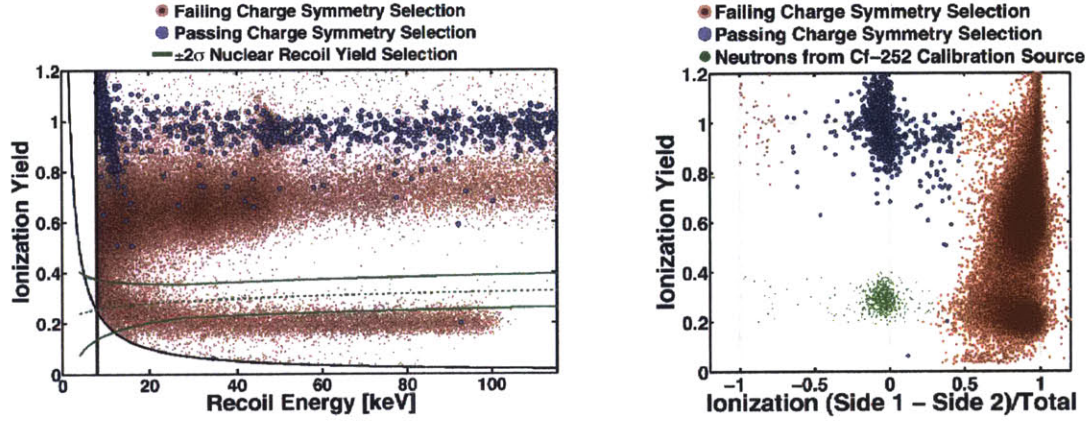


Figure 3-8: Calibration data from ^{210}Pb surface event source above side 1 of detector T3Z1 in Soudan. In both panels, red events fail a requirement that charge be collected on both sides, while blue events pass this selection. **Left:** Events in the plane of ionization yield and recoil energy. The green band is a $\pm 2\sigma$ band where nuclear recoils are expected. The vertical black line is a recoil energy threshold, while the curved black line is an ionization energy threshold. Events with symmetric charge collection are primarily bulk gammas from radioactive decays external to the detector. **Right:** Events in the plane of ionization yield and the normalized difference in charge energy between the two sides. Green events are nuclear recoils from a ^{252}Cf calibration source. Note that surface events appear at a value of +1 in the normalized charge symmetry quantity because the calibration source was above side 1 of the detector. From [61].

interleaved design of the iZIP essentially solves this problem. To characterize the surface event rejection ability of the iZIP, two detectors operating in Soudan had ^{210}Pb surface events sources placed above one of their sides. The source plates were Si wafers that were temporarily exposed to Radon gas.

After several months of data-taking in WIMP-search mode, the surface event rejection power of the ionization readout was measured. The data, after some basic quality cuts and a cut on the outer charge channel energy to remove events near the sidewall, are shown in Figure 3-8. During 906 live hours of exposure, 7.2×10^4 electrons and 1.6×10^4 ^{206}Pb recoils were observed with no events in the signal region, allowing a limit to be put on the surface event rejection factor of the iZIP at 1.7×10^{-5} at 90% C.L. with 50% acceptance in the 8-115 keVnr energy range for a $60\text{GeV}/c^2$ WIMP [61]. This surface event rejection is sufficient to maintain no background from surface events in an experiment with hundreds of tons of target mass.

3.3 Phonons

Phonons are collective vibrational excitations of the crystal lattice. The collective motion of the crystal lattice is described by eigenfunctions of the hamiltonian

$$H = \sum_j \frac{P_j^2}{2M_j} + \sum_{j \neq j'} \frac{1}{2} \frac{Z_j Z_{j'} e^2}{4\pi\epsilon_0 |\mathbf{R}_j - \mathbf{R}_{j'}|} - \sum_{ij} \frac{Z_j e^2}{4\pi\epsilon_0 |\mathbf{r}_i - \mathbf{R}_j|}, \quad (3.11)$$

where the first sum is the kinetic energy, the second describes the interaction of ions in the lattice, and the third describes the interaction of electrons with ions in the lattice. This is analytically intractable, so we make some simplifying approximations. Assume that the ions only see a fixed, time-averaged adiabatic potential from the electrons, Taylor expand the potential about the equilibrium position of the ions, and throw away terms above the leading quadratic ones. The hamiltonian for N particles is therefore rewritten as a collection of $3N$ independent harmonic oscillators, an approach known as the harmonic approximation (discussed very clearly in Chapter 3 of [62] and Chapter 23 of [63]). The frequencies of the harmonic oscillators correspond to the normal modes of the system, one $\omega_s(\mathbf{k})$ per mode (the subscript s denotes the polarization of the displacement, relative to the wave vector \mathbf{k}), and the energy of each mode is $(n_{s\mathbf{k}} + 1/2)\hbar\omega_s(\mathbf{k})$, with $n_{s\mathbf{k}}$ an integer. In analogy with photons, the energy of the normal modes are quantized in units of $\hbar\omega_s(\mathbf{k})$, so these excitation quanta are called *phonons*, each characterized by their wave number \mathbf{k} and their polarization s either along the direction of propagation (*longitudinal*), or perpendicular (*transverse*).

3.3.1 Phonon Propagation

When an event occurs in the iZIP, three types of phonons are produced:

- **Prompt phonons:** These are high-frequency phonons ($\nu > 1$ THz) produced by the recoil itself in its immediate vicinity. Their initial frequency distribution is not well-understood, but this is irrelevant because they rapidly downconvert to lower energies, destroying the initial distribution [64]. During the initial phonon propagation, two scattering processes are relevant:

1. *Anharmonic decay:* This inelastic process results from the deviation of the crystal hamiltonian from the perfect harmonic approximation that we used to derive the stationary states of the crystal. In other words, the higher order terms that we threw away after Taylor expanding equation (3.11) are not completely negligible. Because the crystal is not perfectly harmonic, the stationary states are not perfect normal modes [65]. This nonideal behavior manifests itself in a process by which a single phonon can decay into two lower-energy phonons. The rate for this process in Ge is given by [64]

$$\Gamma = (6.43 \times 10^{-55} \text{ s}^4) \nu^5. \quad (3.12)$$




Figure 3-9: Frequency (*left*) and angular distribution of Luke phonons produced in Ge at 300 V m^{-1} (*black*) and Si at 400 V m^{-1} (*gray*) for electrons (*solid*) and holes (*dashed*). Note that the Luke phonons are strongly focused in the direction of the charge propagation. Figure taken from [66].

Since the scattering rate is quintic in frequency, the phonon downconversion process occurs very rapidly after the energy deposition.

2. *Isotope scattering*: This elastic process occurs due the presence of ions of different isotopes in the crystal lattice. This process is very important in Ge because of the significant presence of five different naturally occurring isotopes. The rate for this process in Ge is given by

$$\Gamma = (3.67 \times 10^{-41} \text{ s}^3) \nu^4. \quad (3.13)$$

At short times after the event, the high scattering rates mean that the mean free path of the phonons is quite short, and the phonon ball expands slowly with increasing mean free path over time, until the phonon energy ultimately becomes homogeneous in the detector.

- **Luke phonons**: When charge carriers propagate throughout the crystal, they produce phonons along their trajectory [52]. Charge carriers rapidly reach a constant drift velocity in the electric field due to the effect of scattering off impurities in the crystal. As a result of this process, phonons are produced in proportion to the electric potential traversed by the charge carriers. The amount of Luke phonon energy produced by n pairs of charge carriers traversing a potential ΔV is given by

$$E_{\text{Luke}} = ne\Delta V. \quad (3.14)$$

These Luke phonons have frequency and angular distributions shown in Figure 3-9. The angular distribution is strongly focused along the direction of the charge propagation, which implies that the Luke phonon signal measured in the detector should have a strong spatial correlation with the charge information.

This is one of the important mechanisms by which phonons can provide redundant position information in addition to the charge signal: if the charge signal of a surface event is asymmetric, then the phonon signal will share some of this asymmetry.

- **Recombination phonons:** When charge carriers recombine at the electrodes on the surfaces, they release their gap energy in the form of ballistic phonons. The presence of these phonons ensures that the full recoil energy of the event is returned to the crystal in phonons.

3.3.2 Phonon Readout

Phonons can be read out in two distinct configurations: *thermal* and *athermal*. In a thermal readout design, the timescale for phonons to downconvert to a thermal distribution is much shorter than the sensor bandwidth. In this mode, the sensor is essentially acting like a thermometer. The athermal design used by the iZIP involves having a large sensor bandwidth coupled with fairly fast phonon collection. In this way, high-energy phonons are absorbed and sensed before decaying to a thermal distribution. Thermal detectors are limited in size because of their heat capacity: if the heat capacity is large, then the temperature increase due to an event will be small and energy resolution poor. Athermal detectors are not subject to this limitation, and can be made much larger without substantial degradation in resolution. This is a key motivation for the iZIP’s athermal design.

In each phonon channel of the iZIP, there are 458 QET phonon sensors photolithographically patterned on the detector surface. The layout of each QET is shown in Figure 3-10. Superconducting Al fins are coupled to a W transition-edge sensor (TES). Electrons inside the superconductor form Cooper pairs, which are characterized by a gap energy $2\Delta_{Al} = 340 \mu\text{eV}$ (at 50 mK) needed to break the pairs. This gap is significantly above the energy of thermal phonons ($k_B T(50 \text{ mK}) = 4.3 \mu\text{eV}$), so the QET is only sensitive to high-energy phonons. When phonon reaches the Al interface, it has some probability of breaking Cooper pairs to form Bogoliubov quasiparticles [67]. The challenge is then to concentrate the quasiparticles into some sensor element that can read out their energy.

The W TES coupled to the Al fins serves to trap and measure the energy of the quasiparticles. Since the W is a superconductor like the Al, it has a gap energy, but the gap is the much smaller value of $2\Delta_W \sim 20 \mu\text{eV}$ (at 50 mK). Because the W is actually deposited on top of the Al, as shown in Figure 3-11, there is also an overlap region where the superconducting gap transitions between $340 \mu\text{eV}$ and $20 \mu\text{eV}$. The portion of the W TES with the highest current density is maintained at the critical temperature of the superconductor, and therefore actually has no energy gap. If quasiparticles encounter the W / Al interface during their diffusion, they can relax due to the decreasing energy gap, releasing their energy in the form of phonons. These phonons, in turn, can heat the electron system of the W TES.

The energy that is now concentrated in the W can then be easily read out by operating it as a standard TES. The theory and operation of TESs is described in




Figure 3-10: Two images of an individual QET. The long light regions are phonon-absorbing Al fins, which are coupled to a central W TES that reads out the phonon energy deposited into the fins. The TES is highlighted in red in the right panel.




Figure 3-11: From [50].

many excellent reviews and theses (e.g. [68–70]), so we will be brief and qualitative. TESs are maintained in the middle of the superconducting transition, where the resistance curve is very steep, and therefore very small inputs of energy or temperature fluctuations will produce a large change in resistance. SuperCDMS uses *voltage-biased* TESs operating in a circuit similar to the simplified diagram in Figure 3-12. A constant current is applied to the QET bias resistor, which then flows through a parallel circuit of a shunt resistor $R_{sh} \simeq 15 - 20$ mOhm and the parallel array of 458 QETs with normal resistance $R_n \simeq 700$ mOhm. Because the shunt resistance is so much smaller than the TES resistance of the QETs, the current flowing through the shunt resistor is nearly constant, and the voltage across the TESs is approximately constant; hence, the TESs are voltage-biased.

Voltage bias is important because the TES is stable when operated in a voltage-biased configuration. Let V be the equilibrium voltage across the TES, and R be the equilibrium resistance of the TES when biased somewhere in its superconducting transition. The Joule power dissipated is simply

$$P = \frac{V^2}{R}. \quad (3.15)$$

If the temperature suddenly rises, due to an event in the detector, then the resistance will rise. In the limit of perfect voltage-bias, V is constant, so the Joule power decreases, and the TES returns to equilibrium. This feedback mechanism between the TES temperature and resistance is known as electrothermal feedback. The feedback is negative in the limit of perfect voltage bias. In the limit of perfect current bias, $P = I^2 R$, so the electrothermal feedback is positive. Other interesting complexities arise in a fully quantitative treatment, such as [68]. If the negative feedback is too strong, for example, then electrothermal oscillations can produce instability in voltage-biased operation.

The TESs on the iZIP are read out using an amplifier circuit sketched in Figure 3-12. An inductor, called the *input coil* is placed in series with the TES and coupled to a superconducting quantum interference device (SQUID). When the SQUID is biased, it will produce a small voltage in response to a small magnetic field. This sensitivity allows it to be used as a low-noise current amplifier. The SQUID is also coupled to another inductor, called the *feedback coil*. The amplifier is a feedback loop that reacts to changes in the SQUID voltage by applying a current to the feedback coil to null out changes in the total magnetic flux through the SQUID. The voltage at the output of amplifier is therefore proportional to the current through the input coil and TES. The current through the TES is proportional to the resistance, and therefore the phonon power being absorbed by the TES. We have therefore achieved our goal of measuring the phonon energy.

3.3.3 Phonon-Based Background Discrimination

Much like the charge, the phonon signal in the iZIP can provide discrimination between both electron and nuclear recoils, as well position information to discriminate

Typical phonon channel connections




Figure 3-12: Simplified diagram of TES circuit and readout. The TESs in the QETs are voltage biased by the presence of a shunt resistor in parallel. Readout is performed by a SQUID in a flux-locked feedback loop.

surface events from bulk events. There are two mechanisms responsible for this ability:

1. **Correlation between Luke phonons and charge:** Since charge carriers produce Luke phonons focused along their direction of drift, the phonon signal is larger on the phonon channels near electrodes that collect the charges. When charge collection is asymmetric, phonon energy is also biased to be asymmetric. Since phonons become ballistic and homogeneous before all the energy is absorbed, this asymmetry is much less pronounced than in the charge signal. While this redundant information might seem unequivocally worse than the ionization signal alone, the phonons have two key advantages over the charge. First, as discussed in Sections 3.1 and 3.2.1 the full recoil energy is available in the phonon signal, while only 10-20% of the energy is available in charge. Phonon collection is also much more reliable than the charge, which varies considerably throughout the detector and depending on the neutralization state of the detector. And second, the phonon baseline noise is actually better than the charge ($\sigma_{pt} \sim 200$ eV vs. $\sigma_Q \sim 400$ eV). This means that the signal-to-noise for nuclear recoils is about 10 times better in phonons than in charge, for example, at 5 keVnr. Even modest discrimination from the Luke phonon signal will be better than ionization information at low energies.
2. **Greater absorption of quasi-diffusive phonons at early times:** As discussed in Section 3.3.1, the high-energy phonons present a few μ s immediately after the event have very short mean free path, and are therefore highly localized. If an event occurs near a phonon channel, that channel will have a high rate of absorption of the prompt phonons, while other channels will have much lower absorption until most phonons become ballistic after a couple hundred of μ s.

Figure 3-13 shows a visual example of the discussion above. The first 300 μ s are characterized by significant pulse shape variation across different channels, which reflects the location of the event and the charge propagation. After this, phonon propagation is ballistic and the spatial distribution is homogeneous, resulting in uniform phonon absorption.

As an example of the surface-event discrimination using the phonon signal, we can use the surface event calibration data on T3Z1 to perform an analysis similar to the one in Section 3.2.1. Figure 3-14 shows a plot analogous to Figure 3-8, but using the phonon information instead of the charge information to tag surface events. The discrimination is not as effective in the charge. Although the phonon discrimination is more powerful at low energies, excessive low-frequency noise on T3Z1 precluded extending the analysis to low energies. In the recoil energy range of 8-115 keVnr, the surface event rejection factor is $4.5 \pm 0.9 \times 10^{-4}$ for an analysis with a 46% acceptance above 8 keVnr for a 60 GeV/ c^2 WIMP.




Figure 3-13: Example phonon pulse from calibration data showing the position-dependent part of the pulse at early times, followed by the ballistic tails during which the phonon powers absorbed by each channel are equal.




Figure 3-14: Calibration data from ^{210}Pb surface event source above side 1 of detector T3Z1 in Soudan. In both panels, red events fail a requirement that charge be collected on both sides, while blue events pass this selection. **Left:** Events in the plane of ionization yield and recoil energy. The green band is a $\pm 2\sigma$ band where nuclear recoils are expected. The vertical black line is a recoil energy threshold, while the curved black line is an ionization energy threshold. Events with symmetric phonon signal are primarily bulk gammas from radioactive decays external to the detector. **Right:** Events in the plane of ionization yield and the normalized difference in phonon energy between the two sides. Green events are nuclear recoils from a ^{252}Cf calibration source. Note that surface events appear at a positive value of the phonon symmetry quantity because the calibration source was above side 1 of the detector. From [61].

3.4 Data Reconstruction

The raw data that is produced by one event in each detector consists of 8 phonon traces (4096 samples per trace) and 4 charge traces (2048 samples per trace) for each channel. The raw data is analyzed using data processing software called `cdmsbats`. While `cdmsbats` uses many algorithms to reduce the information about the trace shapes to a smaller set of quantities, the only algorithm that is very relevant to this thesis is the optimal filter, which is used to estimate event energies. The basic theory of optimal filtering, relevant to CDMS, is described with great clarity in Appendix B of [71] and Appendix A of [72]. We simply review a basic few facts about optimal filters in order to provide some context for the analysis in Chapter 5.

1. **Basic optimal filter:** Denote an event trace by S , a template for that trace by A , and the PSD of the noise by J . Assuming that the noise at each frequency is independent, the minimum-variance estimator for the amplitude a of the event trace \hat{a} , is given by minimizing the quantity

$$\chi^2(a) = \sum_n \frac{|\tilde{S}_n - a\tilde{A}_n|^2}{J_n}, \quad (3.16)$$

where the sum is over the n frequency components, and the tilde denotes the FFT. The motivation for working in the frequency domain is that a χ^2 of this form only provides optimal results if each degree of freedom is independent. In the time domain, this is only true for white noise (flat PSD in frequency). The phonon noise, however, does not have a flat PSD. Nevertheless, the amplitudes of individual *frequency* bins are still expected to be uncorrelated if the noise is a gaussian random process, which motivates performing the fit in the frequency domain, rather than the time domain. Alternatively, we can think of J_n as the diagonal elements of a covariance matrix. Since we have included no off-diagonal elements in the covariance matrix, there must be no correlations between frequencies in order for equation (3.16) to remain a proper χ^2 quantity.

2. **Optimal filter with time shift:** In reality, events start at different times relative to the trigger time, so we could do better by adding a time-shifting degree of freedom to the fits. Since we are working in the frequency domain, this amounts to multiplying the template by a complex phase factor

$$\chi^2(a, t_0) = \sum_n \frac{|\tilde{S}_n - ae^{-2\pi i t_0 f_n} \tilde{A}_n|^2}{J_n}, \quad (3.17)$$

where f_n is the frequency. In the CDMS analysis, the value of t_0 is restricted to a window around the trigger time called the *optimal filter search window*.

3. **Non-stationary optimal filter:** The astute reader will now realize that even the optimal filter with the time-shift is actually pretty suboptimal. We spent a

Variable	Description
pa10F, pb10F, pc10F, pd10F	side 1 OF phonon energies
pa20F, pb20F, pc20F, pd20F	side 2 OF phonon energies
psumOF	sum of all phonon energies ($\sum_* p*OF$)
ptOF	OF energy of sum of all phonon traces
ptNF	non-stationary OF energy of sum of all phonon traces
PTOFdelay	best-fit OF time shift for fit to summed phonon traces
prpartOF	$(pa10F + pa20F)/ptOF$
pzpartOF	$(\sum_* p*10F - \sum_* p*20F)/ptOF$
qi10F, qo10F	side 1 OF charge energies (inner, outer)
qi20F, qo20F	side 2 OF charge energies (inner, outer)
qimeanOF	$qi10F + qi20F/2$
qsummaxOF	$\max\{qi10F + qo10F, qi20F + qo20F\}$

Table 3.1: Variables relevant for iZIP physics analysis.

lot of time in this chapter arguing that phonon pulses have complex position-dependent pulse shapes in the first 300 μs , yet we are fitting them with a fixed template! In other words, the residuals in the numerator of equation (3.17) are not distributed according to the PSD in the denominator. The easiest way to understand this problem is to think about the time-domain covariance of the residual $S(t_n) - aA(t_n)$. Because of the position-dependence, the covariance matrix contains large off-diagonal elements around the start of the event. Matt Pyle developed an approximate way of dealing this problem, called the *non-stationary optimal filter* in CDMS jargon, which essentially just estimates the off-diagonal elements of the equivalent frequency-domain covariance matrix \mathbf{J}_{mn} , and then minimizes the χ^2 [49]

$$\chi^2(a, t_0) = \sum_{m,n} (\tilde{S}_n - ae^{-2\pi i t_0 f_n} \tilde{A}_n)^\dagger \mathbf{J}_{mn}^{-1} (\tilde{S}_m - ae^{-2\pi i t_0 f_m} \tilde{A}_m). \quad (3.18)$$

The non-stationary optimal filter is really just a standard χ^2 fit, which accounts for the fact that position-dependence introduces off-diagonal correlations between certain frequencies. This has the function of deweighting the position-dependent part of the pulse, and using the position-independent tails as the basis for the energy estimate. There are many subtleties and practical difficulties involved in implementing this algorithm, but that is beyond this scope of our treatment.

The optimal filter gives us an energy estimate for each trace, and we combine these into quantities that are useful for the analysis. All analysis variables used in this thesis are summarized in Table 3.1.

Chapter 4

SuperCDMS Experiment

SuperCDMS at Soudan operates in a low-background cryogenic setup originally constructed for the CDMS II experiment, which took data during 2003-2009. Since the lab and hardware for SuperCDMS is quite old, more detailed descriptions can be found in previous theses (e.g. [60, 72]). This chapter gives a brief review of the setup to orient the reader. This is also a good place to clarify the unfortunate and oft-confused nomenclature of the CDMS series of experiments. *CDMS* was an early experiment that took data with Ge detectors in the late 1990s [71]. *CDMS II* took data in Soudan during the 2000s. *SuperCDMS at Soudan* is the experiment described here, using 15 Ge iZIP detectors in the cryostat at Soudan. And *SuperCDMS at SNOLAB* will be a much larger experiment installed at SNOLAB using more numerous and larger iZIP detectors in a completely rebuilt cryogenic facility. The term *CDMS* is also used to refer to the entire series of experiments.

4.1 Soudan Underground Laboratory

SuperCDMS at Soudan is located in the Soudan Underground Laboratory in northern Minnesota in a former underground iron mine. The underground portion of the mine at Soudan was operational from 1892 to 1962 and was profitable because of the high-grade of its ore. The mine is currently operated as a state park, with two large caverns at the 27th level used for particle physics experiments. In addition to SuperCDMS, the far detector of the MINOS neutrino oscillation experiment fills one of the caverns.

The lab is located 713 m beneath the surface with an overburden of 2090 meters water equivalent (m.w.e.), which reduces the cosmic ray muon flux by a factor of approximately 50,000. While the muons and muon-induced secondaries that pass close by a detector can easily be rejected by an active muon veto, muons interacting in the rock far from the experiment produce energetic showers of particles, including neutrons. Neutrons produce nuclear recoils like those expected from WIMPs, so experiments must be performed deep underground where the rate of cosmogenic neutrons is low. The overburden at Soudan is sufficient to reduce the neutron background in the analysis described in Section 5 to 0.1 event. This level of background is too high for significantly larger exposures, which is one of the motivations for building a




Figure 4-1: Muon flux as a function overburden at various underground labs. The muon flux at Soudan is suppressed by a factor of about 5×10^4 relative to the flux at the surface. From [73].

larger experiment at SNOLAB.

4.2 Shielding and Layout

The shielding in SuperCDMS is optimized to suppress neutrons, gammas, and muon secondaries. Various parts of the experimental apparatus generate significant gamma backgrounds. The dilution refrigerator outside the shield has significant amounts of stainless steel, which contains ^{60}Co . Indium and silver used in cryogenic seals and solder joints inside the dilution refrigerator have a variety of radioisotopes. And the PMTs used in the muon veto also have contaminants [60].

As shown in Figure 4-2, the passive shielding consists of alternating layers of polyethylene and lead. Polyethylene is useful as a neutron moderator because it has a high density of free protons (H), whose mass is well-matched to neutrons. Neutrons therefore lose a maximal amount of energy in each collision. Lead is useful for shielding gammas because of its high density and Z , giving it an attenuation length of about 1.2 cm for 1 MeV gammas, versus 15 cm for polyethylene. The alternating layers protect against photoproduction of neutrons by gammas in the lead, which would be inside the outer layer of shielding. The inner layer of lead is “ancient” lead, which was recovered from the ballast of a sunken ship near Nantes, France. Lead smelted at the surface has significant concentrations of the radioisotope ^{210}Pb ($t_{1/2} \sim 22$ y). The decay of ^{210}Bi , a daughter of ^{210}Pb , produces a high-energy beta that can generate significant bremsstrahlung in the lead. The ancient lead has a much smaller fraction of these radioisotopes. The copper cans of the cryostat provide an additional average shielding thickness of about 3 cm. Surrounding the passive shielding is an active

muon veto, with 99.98% efficiency for through-going muons. An average of 1 muons per minute is incident on the veto. The entire shield is 99% hermetic, with two holes for the cold stem and the electrical stem [74].

Inside of the shielding, there are a series of concentric copper cans, known as the “icebox”, which contain the detectors. Each can is attached via the cold stem to a stage of the dilution refrigerator (described below). There are 15 iZIP detectors arranged in 5 “towers” each of 3 detectors inside in the innermost can, as shown in Figure 4-3.

4.3 Cryogenics

The critical temperatures of the TESs vary across detectors between about 70 mK to 110 mK. The operating point of the experiment must be somewhat below this so that cooling to the substrate is strong enough to rapidly return the TESs to their quiescent temperature after an event occurs. The experiment is cooled using a commercial Oxford $^3\text{He}/^4\text{He}$ dilution refrigerator with a cooling power of $400\ \mu\text{W}$ at 100 mK. During stable operation of SuperCDMS, the base temperature measured at tower 3 of the experiment was typically in the range of 50-60 mK. Since the dilution unit is a commercial product, it is not built of low-radioactivity materials. Because of this, it is located outside of the cryostat and shielding that contain the detectors.

The dilution refrigerator achieves its low and stable operating temperature through a closed cycle circulation of the $^3\text{He}/^4\text{He}$ mixture. The mixture is precooled by a combination of liquid N_2 and liquid He before it condenses in the mixing chamber of the refrigerator. ^3He atoms are more strongly bound to ^4He atoms than to each other because the ^3He has a larger zero-point energy due to its smaller mass. Because ^3He obeys Fermi statistics, its kinetic energy increases with its density in the $^3\text{He}/^4\text{He}$ mixture. When the concentration of ^3He exceeds 6.5%, it is no longer bound in the mixture, so a second phase rich in ^3He spontaneously forms. The rich phase and the “dilute” phase are immiscible, with the lighter rich phase floating on top. The ^3He atoms in the dilute phase have a lower entropy, so by pumping ^3He out of the dilute phase, ^3He atoms are moved from the rich into the dilute phase. The corresponding heat of solution for this transition is $Q = T\Delta S = (-84\ \text{J K}^{-2})T^2$ [75]. Closed-cycle circulation of the gas mixture through the mixing chamber therefore cools the fridge, and this circulation can continue stably for months or even years. Commissioning and WIMP-search data analyzed in Chapter 5, for example, was taken during a period from April 2012 to July 2013 during which the fridge remained stably at its base temperature. As the mix exits the dilution unit, it is cleaned in external liquid N_2 and liquid He cold traps, which prevent impurities due to leaks or outgassing from freezing out and blocking the circulation line.

Each stage of the dilution refrigerator is connected to a copper can by cold fingers that pass through the shielding along the cold stem. An outer shield is maintained at 77 K by a liquid N_2 bath, while a liquid He bath cools another stage to 4K. A stage with pumped liquid He is about 1 K, the still is about 800 mK, the cold plate is about 180 mK, and the mixing chamber is around 50 mK. A cryocooler located near




Figure 4-2: **Top:** Cross-sectional view of SuperCDMS shielding from the side. the cryostat in the center is surrounded by alternating layers of polyethylene and lead shielding to suppress background neutrons and gammas. The shielding is enclosed in an active muon veto. One hole in the shielding, labeled “E-stem” contains readout cabling, while the hole labeled “C-stem” contains a cold finger that extends from the dilution refrigerator to cool the cryostat. **Bottom:** Same as top panel but view shown from above. Figure courtesy of Joel Sander.




Figure 4-3: Arrangement of detectors in towers.

the electronics stem also helps cool the 77K and 4K stages. The cryocooler contains a moving piston which generates significant vibrational noise that impacts the analysis, as discussed in Section 5.4.2.

4.4 Cold Hardware

The iZIP detectors used for SuperCDMS at Soudan were designed with the intent of reusing the cold hardware and electronics used in CDMS II. The iZIP is thicker (1 inch) than the ZIPs (1 cm) used in CDMS II, and it has twice as many charge and phonon channels. Each iZIP therefore takes the place of two ZIPs in the towers and uses the readout electronics for two ZIP detectors. Figure 4-4 shows the stack of detector housings mounted in a tower inside the innermost copper can of the experiment. Detectors in the stack are identified by their tower number and position in the tower. So the bottom detector in tower 2 would be designated “T2Z3”. Attached to the inside of each housing are two detector interface boards (DIBs), which contain pins that protrude from the side of the detector housing and pads for wire-bonding connections onto the detector faces. The DIB also has LEDs used to flash the detectors, heating them up and removing any buildup of space charge.

The stack of detectors is connected to four stages, each of which is heatsunk to a different temperature in the fridge. Running along the sides of the hexagonal detector stack are six connectors called the “side coax” (two per detector), which route conductors from the phonon and charge channels to the SQUIDS and FETs at the warmer stages. The side coax contains the bias and feedback resistors of the charge amplifiers, heatsunk at the base temperature in order to reduce Johnson noise.

Stack Assembly Cross Section

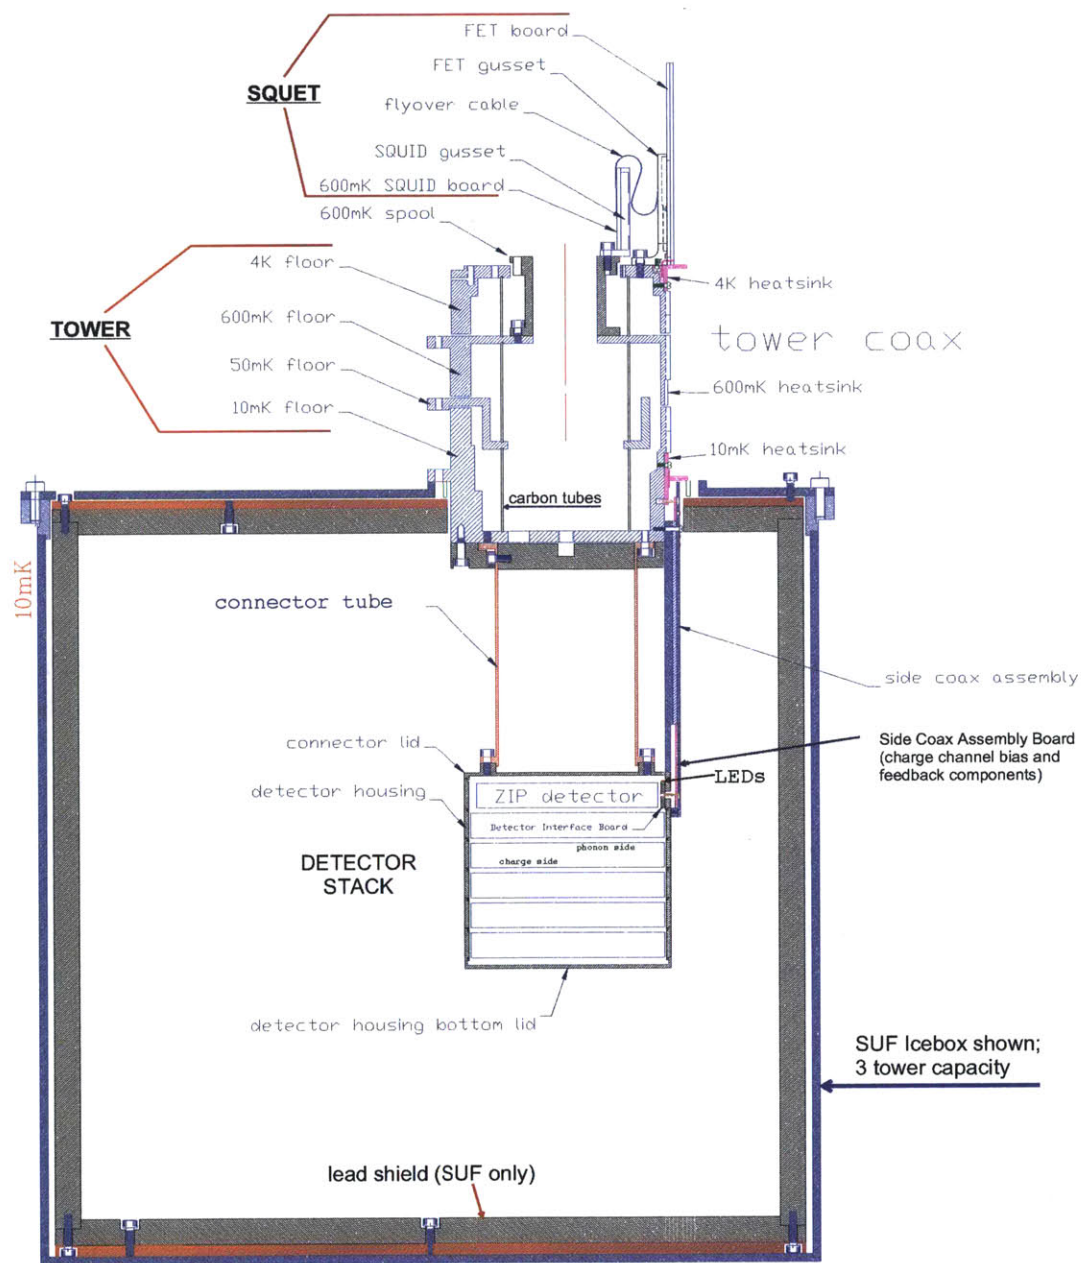



Figure 4-4: Cross-sectional view of the tower assembly used to mount the ZIP detectors used in CDMS II. SuperCDMS at Soudan used the same cold hardware and electronics, but installed only 3 iZIPs per tower instead of the 6 ZIPs per tower used in CDMS II. Drawing prepared by Dennis Seitz [76].



(a) SQUET



(b) side coax



(c) tower assembly

Figure 4-5: Various components of cold hardware. In (a) the larger PCB contains the FET, while the smaller PCB contains chips with SQUIDS and input and feedback coils; (b) shows the side coax with PCB at bottom where capacitors and feedback resistors for the charge amplifiers are located; (c) shows the tower assembly (upside down) with one CDMS II detector mounted. From Dennis Seitz [76].

Near the top of the tower are two connected PCBs containing the SQUIDs and FETs (one per side coax, two per iZIP), collectively called by the unfortunate portmanteau “SQUET”. The FET portion of the card is heatsunk to the 4K stage. The FET consists a doped semiconductor whose charges are largely frozen out at 4K, and it normally operates around 120 K. There is a heater on the card which warms the FET when initially turned on; the FET subsequently self-heats. The SQUID portion of the card is heat sunk to the still at about 800 mK. The SQUID card contains the shunt resistor and chips that contain the SQUID and input and feedback coils. Heat sinking at this stage reduces Johnson noise from the shunt resistor. The top end of the FET card has pins that mate with a flexible PCB going from 4K to 300K, while the bottom end of the FET card connects to the tower and side coax. There are two SQUET cards per detector. Each SQUET reads out two charge channels on the same side of the detector, plus four phonon channels (two from each side).

4.5 Warm Electronics and Triggering

Cables leaving the electronics stem are fed into the room-temperature electronics that control parts of the amplification, signal conditioning, triggering, and digitization. The first stage of this chain is the front-end boards (FEB). The FEBs contain the electronics that bias the phonon and charge channels, as well as parts of the amplifier chain like the SQUID feedback circuit. Again, there are two FEBs per detector, or one per SQUET card.

The output of each FEB connects to a receiver-trigger-filter (RTF) board. The RTF board first applies a Butterworth filter to the signal from the detectors, before summing the phonon signals and the charge signals into a total phonon and a total charge signal. The RTF board then issues a trigger based on the whether the filtered and summed charge or phonon signal is above a given threshold. Because the charge has poor signal-to-noise for low-energy nuclear recoils, triggering is always performed on the summed phonon channel.

Because there are two RTF boards A and B per detector, the decision to trigger the experiment and acquire an event can be based whether triggers are recorded in (A and B), A alone, B alone, or (A or B). The ‘or’ condition (called P1oOR) is typically used because it results in a slightly lower effective energy threshold. Some detectors not analyzed in Chapter 5 have very bad noise on certain channels from problems like shorts to ground, vibrational noise, or failing SQUIDs. In such cases, triggering was switched to the ‘and’ condition or using just one RTF board to prevent excessive triggering of the experiment.

The Butterworth filter applied to the signals before triggering affects the energy threshold of the trigger. A narrower filter reduces both signal and noise, so the bandpass needs to be optimized based on the expected iZIP pulse shape. The bandpass of [900 Hz, 18 kHz] used during the first part of Run 133 was inherited from CDMS II, which used detectors with significantly shorter pulses. The bandpass was reoptimized detector-by-detector to improve the post-filter signal-to-noise, based on the expected pulse shapes of iZIP detectors. Brad Welliver and Ray Bunker changed the resistors

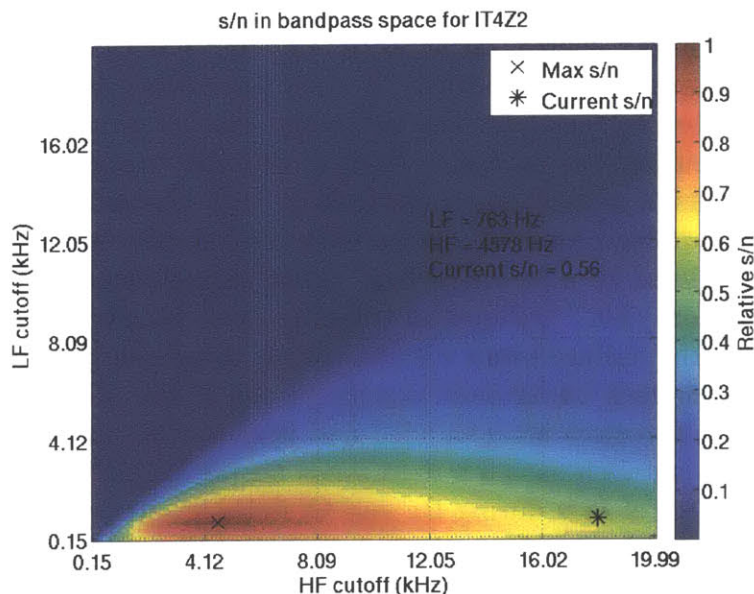


Figure 4-6: Signal-to-noise as a function of Butterworth filter bandpass in the RTF board for a representative detector. Courtesy of Brad Welliver.

in the RTF boards during the first week of January 2013 to implement the new bandpasses. The bandpass of most detectors shifted to lower frequencies. A representative example of the signal-to-noise as a function of the filter window is shown in Figure 4-6. The tuning reduced slightly reduced the energy thresholds on several detectors, and it produced sharper trigger efficiency curves. These two effects slightly increased the low-energy exposure.

4.6 Operation

4.6.1 Experimental Running Modes

Data is taken in three primary running modes: WIMP search, ^{133}Ba calibration, and ^{252}Cf calibration. WIMP search data is taken without any calibration source and is used in the DM analysis. Data is taken with a ^{133}Ba calibration source inserted partially inside the shielding for about 15 h each week. ^{133}Ba produces strong gamma lines at 303 keV (18%), 356 keV (62%), 384 keV (9%), and many other lower energies that are harder to resolve in the Compton continuum. The positions of these lines are used for the charge calibration of the experiment, as well to monitor the overall stability of the calibration. In addition, the regular Ba data is useful for measuring the trigger efficiency on the experiment on a regular basis. Finally, calibration data is taken with ^{252}Cf source every several months. ^{252}Cf produces neutrons over a broad range of MeV-scale energies, which are used to calibrate the response of the detector to nuclear recoils.

During data-acquisition, the detectors must be flashed periodically by the LEDs in order to restore proper neutralization of the crystal. The Ge crystals contain impurities that may capture charge carriers. At very low electric fields ($<0.4 \text{ V cm}^{-1}$), trapping is dominated by “cascade capture”, where charge carriers trap on charged centers to form a neutral site. At higher fields of the iZIP, trapping is dominated by the capture of carriers by neutral impurities to form “over-charged” D^-/A^+ states [77]. Because of these trapping processes, even a fully neutral detector will develop space charge after being biased for a period of time. Space charge reduces charge collection, suppressing ionization yield. Differential trapping rates for electrons and holes will also interfere with the surface event discrimination using charge symmetry. These effects significantly worsen background discrimination in the iZIP. To prevent space charge buildup, the detectors are grounded and flashed with an LED at regular intervals. This occurs between each 3 h data series during WIMP search running, and after every 25 minutes of bias during Ba and Cf data. Photons from the LED heat the crystal, liberating charges from shallow trapping sites, and producing a large population of charge carriers that neutralize ionized impurities.

Chapter 5

Data Analysis for a Light-WIMP Search

After the initial phase of installing the iZIP detectors, understanding the commissioning data, and characterizing the surface event rejection power of the charge fiducial volume in the detectors, SuperCDMS decided to pursue an analysis targeting light WIMPs in the 3-30 GeV mass range. Because of the kinematic mismatch between the WIMP and a Ge nucleus, searches for WIMPs in this mass range improve in sensitivity by lowering the threshold of the analysis. The CDMS II experiment historically used recoil energy thresholds of 8-10 keV. This choice ensured net backgrounds of ~ 1 event, but discarded sensitivity to WIMPs with masses around a few GeV. An alternative approach is to analyze data down the trigger threshold of the experiment, which was successfully used by Ray Bunker and Dave Moore in their analyses of CDMS II data from Run 21 at SUF [78, 79] and Runs 123-128 at Soudan [80, 81]. Searches with low-energy data quickly become background-limited, and only the detectors with the lowest threshold contribute significantly to the sensitivity. A light WIMP search therefore does not gain by analyzing data from all detectors. Because it is sufficient to analyze only a subset of detectors, this analysis was a logical choice for the first science result using the normal operating mode of iZIP detectors.

5.1 Philosophy

5.1.1 Experimental Context

During 2012 and 2013, during the bulk of this work, there was extensive interest in light dark matter due to excess events above background reported by the silicon detectors of CDMS II [82], CoGeNT (Ge) [83], DAMA/LIBRA (NaI) [84], and CRESST II (CaWO_4) [85]. Under the standard assumptions of a Maxwellian dark matter halo and isospin-conserving spin-independent elastic WIMP interactions, these signal claims were mutually incompatible and in strong tension with exclusion limits of the xenon-based experiments XENON10 [86] and XENON100 [87] as well as with the exclusion limit of the germanium detectors of CDMS II [80]. Many authors ob-


served that the tension could be significantly reduced by isospin-violating dark matter that would couple differently to each target nucleus [88], such as the “xenophobic” dark matter scenario [38], or by assuming a non-standard dark matter halo, as is well motivated by N -body simulations [45]. Figure 5-1 shows limits from mid-2013, highlighting the discrepancies between the various signal claims and exclusion limits, as well as the focal point for the claims near a WIMP mass of 10 GeV and a WIMP-nucleon cross-section of 10^{-41} cm².

Improving the sensitivity of this analysis over the low-energy analysis of CDMS II germanium data from Soudan is only possible by exploiting the background discrimination of the iZIP detectors. Because the SuperCDMS at Soudan uses the same cryostat, shielding, and copper as CDMS II, the overall background rate is quite similar between SuperCDMS and CDMS II. Moreover, the ionization and charge resolution of iZIPs in Soudan are comparable to CDMS II detectors, so the yield discrimination is not significantly better in the iZIP. This means that the only handles available for improving the sensitivity over CDMS II are improvements to background rejection algorithms and background modeling. Because of its interleaved phonon/charge sensor design and sensitivity to athermal phonons, the iZIP detector allows us to make major improvements to fiducialization that reduce the effective background. And more careful simulations of the backgrounds and detectors can easily improve on the simplistic background modeling historically performed for CDMS II analyses.


5.1.2 Analysis Strategy

The overall strategy for this search is fairly standard for particle physics analyses and is summarized in Figure 5-2. The analysis was performed blindly, in the sense that the data was first split into two portions: an “open” portion containing high-energy data outside of the nuclear recoil band and data following neutron calibrations with high rates of activation, and a “blinded” portion containing data used to calculate the final limit. The blinding is essential to prevent bias in the event selection, and is discussed in more detail in Section 5.2. We constructed a background model and detector simulation, purpose-built for this analysis, which we then parameterized using calibration data and the open sidebands. The background model provides a complete simulation of the discriminator variables used to separate signal and background, allowing us to use it to train multivariate discrimination algorithms. Following the training, the signal region was unblinded and discriminators were applied to the data.

While this strategy seems conventional from the perspective of, say, collider physics, several factors make it challenging in the context of SuperCDMS. First, before starting this analysis, SuperCDMS did not have a fully validated detector simulation. While extensive development had gone into the ambitious SuperCDMS detector Monte Carlo (DMC) [64, 65], which simulates the physics of iZIPs from first principles, it was not integrated into the analysis pipeline and validated on the kinds of backgrounds most important to this analysis. Given the aggressive timescale required to release this analysis before the G2 funding downselect, it was believed that using the full DMC was too ambitious, requiring us to develop simpler but still robust methods. Second, the data-driven background modeling performed in CDMS II analyses was far



(a) Upper limits from XENON100 (solid green), XENON10 (dashed green), CDMS II (Ge) (solid red), low-energy CDMS II (Ge) (dashed red), low-energy EDELWEISS (solid orange), CDMSlite (solid blue) and COUPP (solid purple). Closed contours from CDMS II (Si) (blue), CoGeNT (orange), DAMA/LIBRA (pink), and CRESST II (green).



(b) Same as Figure 5-1(a), but zoomed into low-mass region.

Figure 5-1: Selection of spin-independent dark matter limits circa 2013. Upper limits are at 90% CL, while closed contours are at 95% CL.




Figure 5-2: Basic flow of the analysis strategy used for the light WIMP search. WIMP-search data was split into two portions: one event selection containing the signal, and another containing all other data ("sidebands"). The signal region was blinded until the discrimination algorithms had been trained using simulations parameterized by the sideband data and Cf calibration data taken periodically during the experiment. After unblinding, an upper limit on the WIMP-nucleon cross section was calculated using data passing the discriminators.

too crude to be useful for low-energy SuperCDMS. Low-energy analyses in CDMS II did not use any background model, while analyses of higher-energy data simply normalized ^{133}Ba calibration data and multiple-scatter events in WIMP-search data to the sideband data above the nuclear recoil band, extrapolating the event rate into the signal region. Although a variety of careful attempts were made to correct for the systematic bias between single- and multiple-scatter surface events, it difficult to validate this modeling approach, which leaves much room for criticism [65, 72]. Criticism aside, we will see that low-energy backgrounds in SuperCDMS contain a large contribution from ^{210}Pb daughter surface events that are completely absent in the multiple-scatter and ^{133}Ba calibrations, rendering the calibration and WIMP-search multiple-scatter samples useless.

Because of quenching, low energy recoils have very little ionization energy (e.g. about 1.5 keVee at 10 keVnr, in a charge channel with a baseline noise of $\sim 300 - 400$ eV). The yield-based discrimination and ionization fiducial volumes that are the key strengths of the iZIP at high energies, are only weak discriminators at low energies. It is therefore necessary to use the phonon position information to gain better discrimination. We combine a variety of these weak discriminators with a boosted decision tree (BDT) to obtain a single strong discriminator variable.

Once the discriminators are set, we optimize a “cut-and-count” analysis that we then apply to the unblinded data. As has been historically common in dark matter searches, data are used for exclusion only, to set an upper limit on the WIMP nucleon cross section under the conservative assumption that all events passing the cut are WIMPs. While it would clearly be preferable to perform a discovery analysis with full background subtraction, doing so requires full confidence in the model of the background to subtract. Because we do not have calibration data for all backgrounds in this analysis, it is difficult to reliably validate our background modeling to this level. In the cut-and-count approach without background-subtraction, background mismodeling can only cause the cuts to be set suboptimally. Provided the WIMP efficiency is correctly measured, it can never produce an overly aggressive limit.

5.2 Data Set and Blinding

We analyze WIMP-search data taken between 1 October 2012 and 1 June 2013, during Run 133 of SuperCDMS at Soudan. During this period, 15 iZIP detectors were in run continuously in the cryostat. We analyze data from seven of these detectors, designated as T1Z1, T2Z1, T2Z2, T4Z2, T4Z3, T5Z2, and T5Z3. An eighth detector, T1Z3, was originally included in the analysis but subsequently discarded before unblinding because its shorted charge channels QIS1 and QOS1 complicated the analysis and appeared to significantly degrade its background discrimination. The seven detectors were selected qualitatively on the basis of their low trigger thresholds, stable noise and operational performance, and comparative lack of shorted ionization and phonon channels. Before any data quality cuts, this totals 182.6 live days of data, or 767 kg d of exposure.

Calibration with a ^{252}Cf neutron source was performed during Run 133 at the

times listed in Table 5.1. During these calibrations, neutrons capture on the naturally occurring isotope ^{70}Ge , producing the isotope ^{71}Ge that decays by electron capture to ^{71}Ga with a half-life of 11.4 d. If a K-shell electron is captured, the energy released by the atomic de-excitation is 10.36 keV, while 1.30 keV is released for L-shell electron capture [89]. Because of the higher rate of background at low energies near the region of interest for light WIMPs, periods of data within 10 or 20 d following the Cf calibrations are not used for setting limits in this analysis (hereafter called “post-Cf data”). After removing the post-Cf data from the WIMP-search dataset, 169.4 live days of data, or 711.5 kg d of exposure remain.

All single-scatter events, passing some very basic quality requirements, and in the energy range 0-13.1 keV in total phonon energy (`psum0F`) were blinded at the start of the analysis (i.e. hidden from the analyzers). Blinding the data set prevents the selection criteria, analysis optimization, and especially the background modeling from being biased by the observed data. While it demands a higher level of understanding of the detector and background physics, a blind analysis in a low signal-to-background setting such as ours is essential to prevent the analyzer from unfairly drawing his conclusion-of-choice from the data. Our highly conservative blinding strategy was initially chosen for two reasons. First, one can always unblind more data if the analysis requires it, but one can never “re-blind” data, so it is best to start with an overly conservative prescription. And second, the poor resolution of the experiment means that the signal and background distributions have significant overlap at low energy. Data above 13.1 keV was set aside for a later high-threshold analysis and blinded according to a more complex but less strict scheme that hides single-scatter nuclear recoils that interact in the bulk of each detector. Events above 13.1 keV and outside of the nuclear recoil band are not blinded.

Although the post-Cf data is not useful for the final WIMP-search analysis, it is very useful for characterizing backgrounds present at low energies. The periods of post-Cf data listed in Table 5.1 remained open for use in the analysis, mainly in order to estimate the background from L-shell EC decay of ^{68}Ge and ^{71}Ge , and as a crude cross-check of the background modeling. Leaving this data unblinded obviously creates some modest potential for bias in the analysis. Section 5.7.1 discusses how this risk was mitigated when estimating the L-shell backgrounds. In addition, discrimination algorithms are never trained directly on the data itself, so even if a WIMP signal is present in the data, the discrimination algorithms will not be biased to exclude it. And finally, the total exposure in the open dataset, before quality cuts, is only 130.2 kg d, or about 18% of the total exposure. Random triggers taken throughout the entire run were also not blinded.

5.3 Charge Fiducial Volume Cuts

As described in Chapter 3, the iZIP has the ability to fiducialize using both the ionization and phonon signals of the detector. Although the background discrimination in the phonon signal is more effective at lower energies than in the charge because of better signal-to-noise in the electronics, the charge signal is easier to interpret because

date of ^{252}Cf calibration	subsequent “open” data period
11 Apr 2012	11 Apr - 21 Apr 2012
31 Jul 2012	not unblinded (unstable)
22 Aug 2012	22 Aug - 9 Sep 2012
30 Aug 2012	22 Aug - 9 Sep 2012
15 Jan 2013	15 Jan - 4 Feb 2013
8 Jul - 9 Jul 2013	not unblinded (end of Run 133)

Table 5.1: ^{252}Cf calibrations used during Run 133 and the subsequent periods of data that were left unblinded because of high rates of activation.

it has better *intrinsic* discrimination. To a very good approximation, sidewall events with sufficient ionization yield tend to produce a signal in the outer charge channel, while bulk events produce no signal in the outer charge channel. And surface events on the top and bottom surfaces produce charge collection only on one side, while bulk events are symmetric (assuming no charge trapping). A reasonable strategy for reducing the surface event background is therefore to require that the charge signal in the outer channels is consistent with noise and that the charge signal on the two sides are consistent with being equal. Although further optimization could be possible, the low signal-to-noise already means that this “charge fiducial volume” has fairly high efficiency at low energies.

5.4 Quality Cuts

SuperCDMS analyses are very sensitive to data stability and quality. Periods of poor charge collection can result in reduced ionization yield, causing electron recoils or sidewall surface events to appear as bulk nuclear recoils. Elevated phonon noise can increase the number of triggers above threshold, appearing as a population events with low ionization yield that could be misinterpreted as low-mass WIMPs. Many other pathologies are possible. As a result, extensive effort is devoted to developing a large number of quality cuts. The large number of cuts needed also illustrates why blinding is so crucial for SuperCDMS. Without it, there would be an overwhelming tendency to adjust cuts to remove signal-like events in the real data. The subsections below summarize the quality cuts developed for and used in this analysis.

5.4.1 Data Period Cuts

These selection criteria remove entire periods of time with undesirable data. When a period of data on only a subset the detectors is removed from the WIMP-search analysis, that detector is still used as a veto to reject multiple-scatter events.

- **Trigger Rate Stability:** Although the overall rate of physics events is fairly constant near ~ 1 Hz, there are short periods with significantly higher trig-

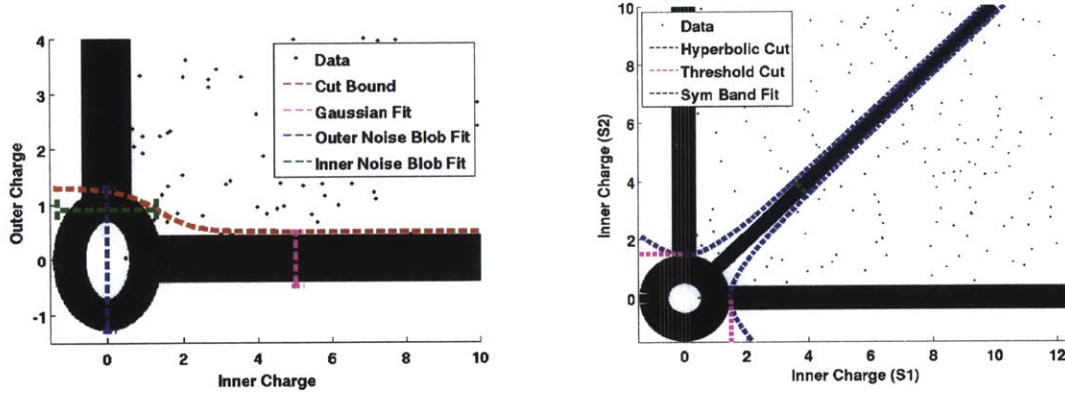


Figure 5-3: Schematic showing the location of the main population of events in the space of charge channel signals, and the approximate location of the charge fiducial volume cuts. The left panel shows the inner charge signal versus the outer charge signal on one side (a similar plot exists for the opposite side). Events tend to form a bands with charge collection on strictly on either the inner or outer charge channel, with a low-density population of events having their charge shared between the inner and outer electrodes. The charge fiducial volume cuts `cQin1_v53_LT` and `cQin2_v53_LT` require that events have a charge signal consistent with noise on the outer charge channel (below the red line). The right panel shows the distribution of events in the plane of inner charge signal on side 1 versus inner charge signal on side 2. There are bands of events corresponding to charge collection strictly on side 1 (S1 surface events), strictly on side 2 (S2 surface events), or equal on the two sides. The charge fiducial volume cut `cQisym_v53_LT` requires that events be in the band with equal collection on both sides. (Courtesy of Todd Doughty)




Figure 5-4: Variation in tower 3 temperature during the portion of Run 133 used for gain correction.

ger rate. The origin of these periods is not thoroughly understood, but low-frequency phonon noise of environmental origin is a significant contributor. The cut `cTrigBurst_133` removes these periods by two criteria:

- Entire series with average trigger rates above 1.3 Hz are removed. The threshold was set arbitrarily to balance acceptance against removing outliers.
- Subsets of series with elevated trigger rates are removed. Three criteria were used to identify periods to remove:
 1. Individual periods of 200 s were removed if they had an average trigger rate greater than 3 Hz.
 2. Periods with rates more than 8σ above the median rate for 200 s intervals in the series were discarded. This removes periods in which the rate spikes significantly above the median rate for the period, even if the absolute rate of the offending series is still low.

The periods within a series that satisfy either of the above criteria must also either have a combined live time rate greater than 3 Hz or the total event rate in the bursting period must be more than twice the rate in the non-bursting period.

Both criteria together remove 19.1% of events with a total live time efficiency of 96.7% in WIMP-search data taken between the beginning of Run 133 and 1 June 2013.

- **Base Temperature:** One new, but important observation first made in Run 133 was that the gain of the phonon channels has a small dependence on the base temperature of the fridge. Over the course of normal WIMP-search running in Run 133, the fridge temperature (measured only on tower 3) typically varied in the range of 52-63 mK (see Figure 5-4). The mechanism for this is not well understood, but several plausible causes include:

1. *TES dynamics:* Since the electron-phonon coupling scales as $P_{ep} \propto T_{\text{TES}}^5 - T_{\text{bath}}^5$, a larger difference between the TES and bath temperatures will alter the electron-phonon coupling and therefore the bias point in the superconducting transition and gain. Because of the quintic dependence on




Figure 5-5: Phonon channel T_c by detector, compared with phonon gain as a function of temperature.

temperature, this should only be relevant when the bath temperature is close to the TES temperature. Figure 5-6 shows, however, that there is not a particularly strong correlation between detectors with low T_c and strong temperature dependence, casting doubt on this hypothesis.

2. *Temperature variation across towers:* The “base” temperature was only measured on tower 3 due to failures with thermometers on other towers. If the temperature is different on each tower, then the correlation between T_c and temperature-dependence could be broken, as we see empirically in Figure 5-6.
3. *TES self-heating:* Noise down the QET lines could produce self-heating the TESs, increasing the temperature closer to T_c and similarly breaking the T_c and temperature-dependence.

Regardless of the cause, gain fluctuations need to be corrected in calibration. There is insufficient data at very low and very high temperatures to perform a reliable correction, and some periods have no data because of readout problems. These periods are removed by the cut `cGoodBaseTemp_v53`. On detectors T2Z1, T2Z2, T2Z3, T4Z3, and T5Z1 data below 52 mK or above 64 mK is removed. The other detectors have smaller temperature dependence, so only data above 64 mK is removed. An algorithm is applied to recover some data with no base temperature measurement. The total live time loss during the analyzed dataset is 4.9%.

- **DC Offsets:** The phonon readout relies on the SQUID feedback loop being locked and in stable operation. SQUIDS can occasionally lock in part of the $V - \Phi$ curve with reduced sensitivity due to trapped flux in the SQUID, or they can lose lock altogether. In such cases, phonon pulses may have reduced amplitude or may not appear in the problematic channels. These periods can be identified by abnormal DC offsets of the phonon channels. To remove these periods, we therefore require that the DC offsets for every phonon channel on

	T1	T2	T3	T4	T5
Z1	99.5%	100%	99.8%	99.9%	99.5%
Z2	99.8%	99.8%	99.9%	100%	96.5%
Z3	99.9%	92.3%	100%	99.9%	99.7%

Table 5.2: Live time efficiency of `cQhighnoise_v53`.

every detector be within 250 mV of its mean value in the `cGoodDCOffset_v53` cut. This requirement removes $< 3\%$ of live time on every detector, except T4Z3, on which it removes 12% of live time.

- **High Charge Noise:** The baseline charge noise fluctuates over time for reasons that are not necessarily well-understood. The noise baseline can be monitored on a continuous basis throughout each series using the random triggers taken in the middle of the run. In addition to secular trends in the charge noise environment, there are brief periods with extremely high charge noise. These periods are identified using a moving-average algorithm developed by Brad Welliver. In qualitative terms, the algorithm for the `cQhighnoise_v53` cut removes periods whose moving-average OF charge amplitudes are above a time-dependent threshold. These events are potentially problematic: if low-frequency phonon noise were to trigger the experiment at the same time as elevated charge noise, such an event could reconstruct with a charge energy above the typical charge noise distribution, faking a low-energy nuclear recoil. As shown by Table 5.2, `cQhighnoise_v53` has excellent live time efficiency.
- **Phonon Resolution:** An estimate of the optimal filter resolution for every phonon channel in the experiment can be monitored on a series-by-series basis using the random triggers. To ensure that the phonon noise environment is controlled, Run 133 was divided into time blocks with similar noise characteristics. Ideally, we would remove data in which any channel deviated significantly from the baseline of the time block. Unfortunately the resolution of individual phonon channels fluctuates enough that this strategy would incur prohibitive live time losses. Instead, all phonon OF resolutions for each detector are summed in quadrature to estimate the OF resolution of the total phonon energy

$$\sigma_t = \sqrt{\sum_i \sigma_i^2}. \quad (5.1)$$

This quantity is both more stable than the individual channel resolutions, and more relevant for the analysis than the individual phonon channel resolutions. Periods are removed by the `cBadOFRes_v53` cut when σ_t is more than 25% above the mean value for the time block. The live time efficiency of this requirement is shown in Table 5.3.

- **Trigger Errors:** In a very small number of events, the trigger information

	T1	T2	T3	T4	T5
Z1	90.6%	97.5%	96.1%	99.8%	97.4%
Z2	98.1%	99.7%	96.5%	99.9%	99.5%
Z3	98.5%	96.6%	99.3%	99.0%	96.8%

Table 5.3: Live time efficiency of cBadOFRes_v53.

Detector	Dates Excluded
T1Z1	2 Jan - 9 Jan 2013
T2Z1	2 Jan - 3 Jan 2013
T2Z2	2 Jan - 7 Jan 2013
T4Z2	2 Jan - 4 Jan 2013
T4Z3	2 Jan - 3 Jan 2013 1 Feb - 4 Feb 2013
T5Z2	28 Dec 2012 - 3 Jan 2013
T5Z3	12 Dec 2012 - 14 Dec 2012 31 Dec 2012 - 4 Jan 2013

Table 5.4: Dates with data series excluded by cStableTrigThresh_133_LT due to rapid changes in trigger threshold that prevented the trigger efficiency from being reliably estimated.

recorded by the the DAQ does not correspond to a physically valid situation. A variety of failure modes are classified, but most of these involve the lack of a valid trigger bit (indicates which detector issued the global trigger) recorded at the nominal trigger time or inconsistent veto triggering information. These events are removed by the cut cErrMask_v53, with a live time efficiency of 99.75%.

- **Trigger Threshold Stability:** Because of variations in the baseline noise environment during Run 133, the trigger threshold was adjusted multiple times on all detectors. Because this analysis uses data at the trigger threshold of each detector, it is critical that the trigger-threshold be well-measured at all times. The trigger efficiency was monitored using ^{133}Ba calibration data, typically taken over three days per week for a total of about 15 h per week. There were several periods during Run 133 in which the trigger threshold was adjusted frequently, and there was insufficient Ba calibration data to reliably measure the trigger efficiency. Such periods are listed in Table 5.4 and excluded from the analysis by the cut cStableTrigThresh_133_LT.
- **Bias Time:** WIMP-search events taken while the detector has been biased for more than 3 hours are removed by the cGoodBiasTime_133 cut. Data biased for longer than this is not guaranteed to have neutralization and charge collection.
- **Bias Polarity:** During Run 133, calibration and WIMP-search data were typ-




Figure 5-6: Phonon optimal filter χ^2 quantity as a function of total phonon energy for ^{133}Ba calibration data, showing cut definition. Though the cut strength is somewhat arbitrary, most events above the boundary are pileup, while most below the boundary are high-quality events. T2Z1 is shown as a representative example.

ically taken with a bias polarity of +2V/-2V (S1/S2), although data was occasionally taken with the reverse bias polarity. The cut `cBiasPM_133` removes any WIMP-search data that is not in the standard bias configuration.

- **Miscellaneous bad series:** 18 of the 1643 series taken for this analysis were marked as bad by the onsite shift for miscellaneous reasons. They were removed from the analysis sample by the cut `cBadSeries_133`. The cut `cFinalPhononSettings_133` removed series during which SQUID and QET settings were not in their finalized state, though no series were removed from the data analyzed here.

5.4.2 Pulse Quality Cuts

The selection criteria below remove individual events from the analysis.

- **Phonon optimal filter χ^2** The χ^2 obtained from the optimal filter provides a measure of the goodness of fit between the phonon pulse in data and the optimal filter template that is used to estimate the energy. Events with high values of χ^2 may have poorly estimated energies and should therefore be discarded. This is most important as a means to remove pileup events from calibration data. This also provides a crude “catch-all” method of discriminating pathological events caused by various kinds of electronics noise, such as “glitch” events, low-frequency noise triggers, and square pulses (all discussed in more detail below).

At low energies, the phonon optimal filter χ^2 for the total phonon trace has a value centered near 4095 (4096 digitizer bins minus one d.o.f. for the fitted amplitude of the pulse), but this rises at energies above ~ 200 keV on most detectors, as shown in Figure 5-6. A 3σ upper cut called `cPChiSq_v53` is set in energy bins and smoothed with a quartic polynomial on each detector. The 3σ cut threshold is essentially arbitrary, but it does a good job separating pileup events from traces containing a single event.

- **Charge Optimal Filter χ^2 :** In the spirit of the charge χ^2 , `cQChiSq_v53` is similar to the phonon version. Cuts were placed on the optimal filter χ^2 for the charge channels on each side (`QS10Fchisq` and `QS10Fchisq`) as a function of the total charge energy on each side (`qsum10F`, `qsum20F`). The cut threshold was placed slightly looser than 3σ in order to balance acceptance and rejection, albeit in a somewhat arbitrary manner.
- **Prepulse Noise Baseline:** A buffer of 512 digitizer bins of pre-trigger data are recorded with each event. This portion of the trace provides a snapshot of the noise environment on an event-by-event basis. If it is abnormally noisy, then the event is suspect. The standard deviation of this prepulse region is calculated for each charge and phonon channel of each event. Events with baseline standard deviations more than 4σ away from the typical value for the series are removed by the cuts `cQstd_v53` and `cPstd_v53`.
- **Glitches:** Since the early days of CDMS II at Soudan, a population of events known as “glitch events” has been observed. There are two classes of events both called glitches. The first class is characterized by very short phonon fall times (typically $\tau_{\text{glitch}} \sim 100\mu\text{s}$ instead of the $\tau_{\text{pulse}} \sim 750\mu\text{s}$ typical of iZIP phonon pulses), with little or no charge signal. Such events typically trigger on a large number of detectors in phonons, but very few in charge. Phonon glitches are believed to be produced by voltage spikes on the QET bias line that directly heat the TES. The resulting pulse shape therefore should be reasonably close to the impulse response of the TES, which has a much shorter fall time than the phonon collection time that sets the pulse fall time of τ_{pulse} ([49] has an excellent discussion of the physical origin of glitches and the weaknesses of this argument). The second class of events is characterized large spikes in many charge channels, but few phonon triggers. These distinct characteristics provide at least two obvious methods for discrimination:

1. *Trigger multiplicity:* Single-scatter WIMPs or calibration data should have a relatively similar number of charge and phonon triggers. At low energies, one would expect somewhat fewer charge triggers relative to phonon triggers because the signal-to-noise is worse in the charge, but the difference should not be too great. The first selection criteria to veto glitches, `cGlitch_133` removes events with at least six more phonon triggers than charge (`ntrigp - ntrigq > 6`) and events with at least one more charge trigger than phonon (`ntrigq - ntrigp > 1`). As shown in Figure 5-7,

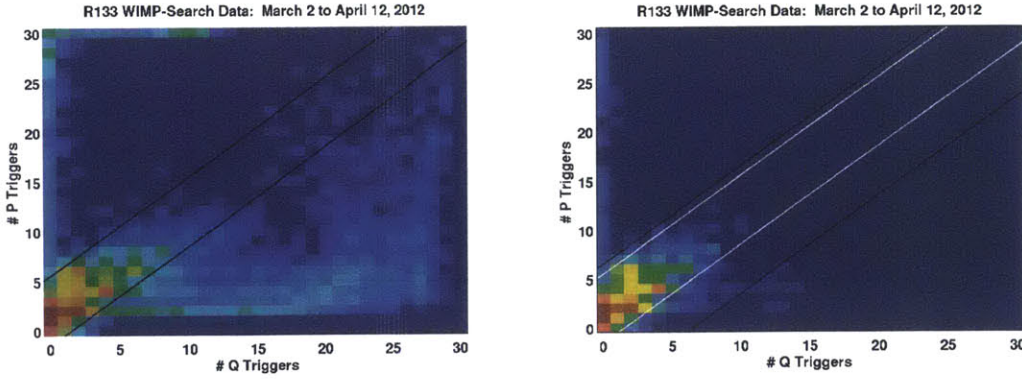


Figure 5-7: Histogram of number of DIBs with charge triggers vs. number of DIBs with phonon triggers in a commissioning dataset. Left panel shows all data, while right panel shows a subset of data passing the charge χ^2 requirement. Note the significant population of events with an unequal number of charge and phonon triggers. Black lines on left panel indicate acceptance region, which was set in order to ensure high-efficiency on a subset of events that also pass the phonon and charge χ^2 requirements. (Courtesy of Ray Bunker)

this removes the large number of events with a mismatch in the number of charge and phonon triggers.

2. *Pulse shape*: Figure 5-8 shows the significant difference in pulse shape between a typical glitch event and physics event. This difference is significant enough that the phonon goodness-of-fit cut `cPChiSq_v53` removes the majority of the glitches. Somewhat better performance can be achieved, however, via a likelihood ratio test statistic. Each total phonon trace is fit to a physics pulse template and to a glitch pulse template, and the χ^2 for each is computed. The log-likelihood ratio test statistic $\Delta\chi^2 \equiv \chi_{\text{NF}}^2 - \chi_{\text{glitch}}^2$ is used as a discriminator between glitch events and physics events. Because of pulse shape variations in physics pulses, the two classes of events form parabolic bands, so this test statistic is not truly distributed as a χ^2 distribution. Nevertheless, it still has the strong discrimination power that one would expect from the Neyman-Pearson lemma. The cut `cGlitch1_v53` is an energy-dependent parabolic band, defined in a somewhat arbitrary fashion, to separate the two event types.

- **Low-frequency noise**: SuperCDMS data taken at Soudan was plagued by low-frequency phonon noise that caused excess triggers in the experiment. Figure 5-13 shows a typical low-frequency noise event that triggered. While the origins of all low-frequency noise triggers are not understood, a major source of the noise is known to be caused by mechanical vibrations that couple to the TESs. The pulse-tube cryocooler on the e-stem of the fridge is known to produce vibrations that are correlated in phase with bursts of low-frequency noise. Indeed, if

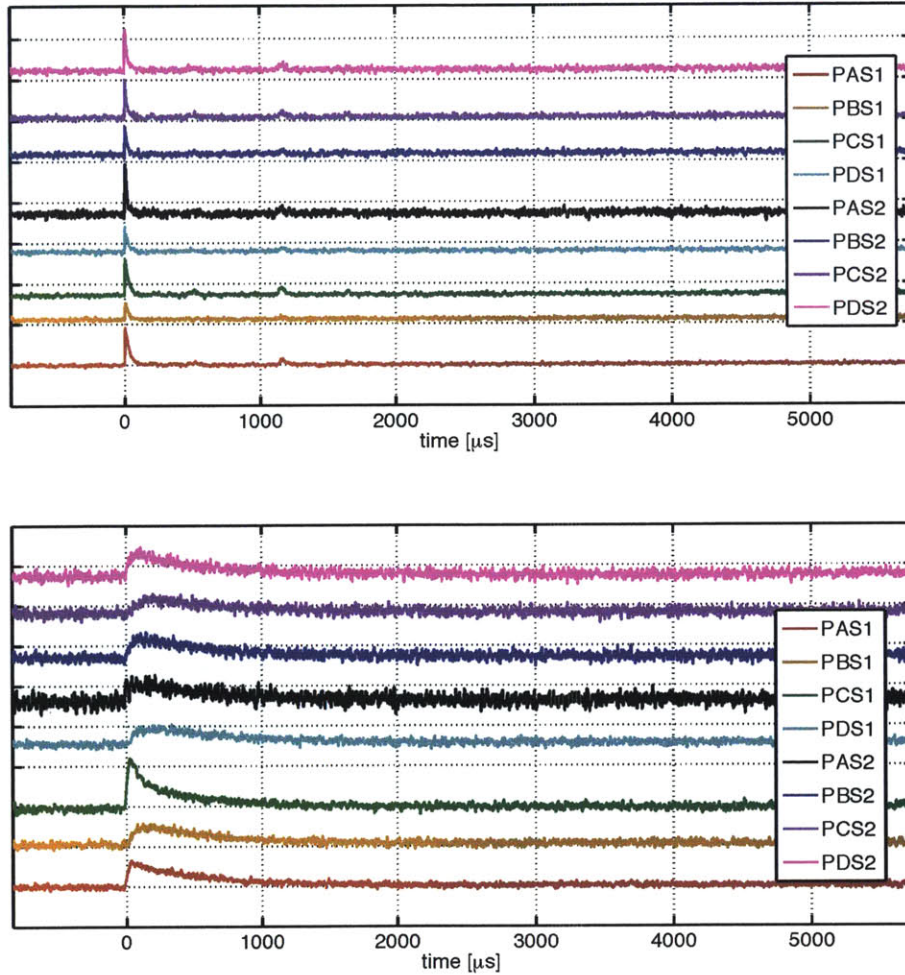


Figure 5-8: Typical phonon traces from a glitch event (*top*) compared with a good physics event (*bottom*). Note the significantly shorter fall time in the glitch event.




Figure 5-9: Typical total phonon trace for a low-frequency noise event (*left*), compared with a typical physics event. Note the faster rise time in the physics pulse, and the second and third “bumps” in the low-frequency noise trace.

one looks at a phonon channel on an oscilloscope and listens to the chirp of the cryocooler, regular bursts of low-frequency noise are synchronized with the cryocooler cycle. Since Run 133, important progress has been made reducing the effects of the cryocooler noise. Replacing the cryocooler cold head significantly reduced the amount of vibrational noise, while Ritoban Basu Thakur developed a method of using an accelerometer to monitor vibrations and veto particularly noise periods.

In spite of these later efforts, the low-energy data analyzed here is heavily contaminated with low-frequency noise triggers. Because physics event pulse shapes are also dominated by low-frequency information in the 100 Hz range, the general-purpose phonon quality cut `cPChiSq_v53` does not provide very good discrimination. To compound problems, low-frequency noise does not have a unique pulse shape. An attempt was made to discriminate low-frequency noise from physics events using an optimal filter χ^2 quantity computed using low-frequency information only, but the best discrimination was achieved using a likelihood ratio test statistic similar to the one used for `cGlitch1_v53`, developed by Kristi Schneck. A large number of low-frequency noise events were averaged together to create a low-frequency noise “template” for each detector. Although there is significant pulse shape variation between low-frequency noise events, there is enough similarity that the log-likelihood ratio between the physics template and the low-frequency noise template provides good discrimination between the two classes of events—even down to threshold.

- **Phonon Start Time:** As described in Section 3.4, the start time of the optimal filter template is varied in a search window to achieve the best fit. The search window was set to be wide enough so that most events should be reconstructed near the middle of the window. Because the phonon pulses have a non-zero rise time, however, it is conceivably possible that very low-energy pulses will




Figure 5-10: Schematic showing how low-energy events can have a start time outside the optimal filter search window. Because low-energy events trigger near the peak of the pulse, if the pulse rise time is larger than the pre-trigger portion of the window, then the optimal filter will rail. These events will have biased energy estimates and are therefore removed by `cGoodPStartTime_v53`.

rail on the leading edge of the window, as shown schematically in Figure 5-10. The search window was optimized to avoid this situation, but any events that do rail may have biased energy estimators and should therefore be removed. Similarly, calibration data can have “cross-detector” pileup events, which rail at the end of the optimal filter search window. This occurs when detector A issues a global trigger, and then another event triggers detector B just outside of the OF search window. The phonon χ^2 on detector B will be reasonable because start time is near the edge of the search window, but the energy may be slightly biased, and so these events are removed by the cut `cGoodPStartTime_v53`.

Events railed at the edges of the search window are removed simply by enforcing that

$$(E_{pt} > E_{min}) \cap (-190\mu\text{s} < t_{OF} < t_{min}),$$

where E_{min} is a detector-dependent low-energy threshold for the cut (typically around 3 keV in total phonon energy), and t_{min} is a detector-dependent time threshold (typically around $35\mu\text{s}$). The cut is not applied to random triggers or other very low-energy events because the time-resolution is poor at very low energies, which would compromise the cut efficiency.

- **Square Pulses:** Several detectors experienced periods during Run 133 in which a significant number of “square” pulses were recorded. These were vetoed by a requirement that the integral of the phonon trace not be excessively negative, and that the fall time be loosely consistent with the value expected for physics events. The combination of these requirements forms the cut

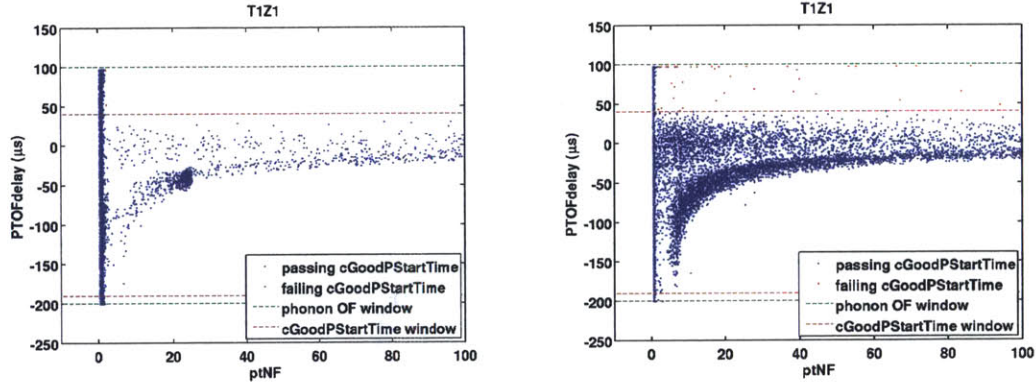


Figure 5-11: Phonon optimal filter start time as a function of energy for WIMP-search data (*left*) and Cf calibration data (*right*), showing events that pass and fail `cGoodPStartTime_v53`.

`cSquarePulse_v53`.

5.4.3 Quality Cut Efficiencies

There are three methods that are appropriate for computing the efficiency of quality cuts, depending on the cut definition. The first method is to compute the live time loss incurred by the cut. This is appropriate for cuts that remove entire periods of data. Somewhat less obviously, it is also appropriate for cuts such as `cPstd_v53` and `cQstd_v53`, which decide whether to pass or fail an event based on the background noise *environment*, rather than based on an intrinsic property of the event. The efficiency of these cuts is given by the fraction of the time in which if an event *were* to occur, then it *would* have been rejected. The simplest estimate for this fraction is just

$$\hat{f} = \frac{\sum_{\text{fail}} t_i}{\sum_{\text{all}} t_i},$$

where t_i is the total live time immediately preceding each event. So the total live time loss is therefore

$$\Delta T = \sum_{\text{fail}} t_i$$

the sum of the live times preceding each event that fails the cuts. The live time losses are summarized in Table 5.5.

The second “method” is to realize that some cuts incur no live time loss. This assumption is appropriate for the trigger-based glitch cut `cGlitch_133`. A single-scatter WIMP candidate should always have 2 phonon triggers and 0, 1, or 2 charge triggers, so a WIMP will never be accidentally classified as a glitch. However, the time after the event when the DAQ is not live obviously should not count toward the live time of the experiment.

The third method is to estimate the fraction of single-scatter WIMPs that pass a

Cut	T1Z1	T2Z1	T2Z2	T4Z2	T4Z3	T5Z2	T5Z3
cGoodBiasTime_133	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%
!cBadSeries_133	99.2%	99.2%	99.2%	99.2%	99.2%	99.2%	99.2%
!cTrigBurst_133	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%	96.3%
cFinalPhononSetting_133	100%	100%	100%	100%	100%	100%	100%
cGoodDCOffset_v53	99.6%	99.7%	99.7%	99.2%	98.5%	99.7%	99.3%
cGoodBaseTemp_v53	99.9%	95.1%	95.1%	99.9%	95.1%	99.9%	99.9%
!cQhighnoise_v53	99.8%	100%	99.8%	100%	99.9%	96.3%	99.9%
!cBad0FRes_v53	94.4%	96.1%	99.9%	99.8%	98.7%	100%	97.6%

Table 5.5: Live time efficiency of main quality cuts.

cut. This is appropriate for cuts that have a fixed mis-identification rate on events that trigger. Estimating these efficiencies from nuclear recoils in ^{252}Cf calibration data is confounded by the fact that we must first apply our quality cuts to identify a sample of good nuclear recoils, free of pileup, to use in the efficiency estimation. Instead of using data, we rely on a Monte Carlo pulse simulation to construct fake pulses and then use it to estimate the cut efficiencies. The use of a simulation is justified by two observations. First, the efficiency of all quality cuts is expected to be quite high by construction, so even large systematics on the rate of failing events should be small in absolute terms. Second, variation in the quality cut efficiencies between detectors and over time are expected to be driven by variation in the baseline noise of each detector. A pulse simulation that incorporates experimental noise from throughout all of Run 133 should therefore produce more reliable results than calibration data that is taken at fixed intervals.

To generate fake pulses, the phonon templates used for the optimal filter are first scaled to a random energy picked from a uniform distribution between 0 and 30 keV in total phonon energy. Next, the charge templates used for the optimal filter were scaled to an energy calculated from the true phonon energy, assuming a Lindhard ionization yield with $k = 0.157$ and full, symmetric charge collection on the inner two electrodes of the iZIP, with no charge collection in the outer channel. These simplifying assumptions are justified by the fact that the ionization yield was measured to be similar to Lindhard on most detectors and WIMPs with non-zero charge collection would be removed from the analysis anyway by the charge fiducial volumes. This method also implicitly assumes that pulse-shape variations do not significantly affect the acceptance of quality cuts. This is a fairly reasonable assumption because the quality cuts whose efficiencies are estimated from simulation are all based on the total phonon trace (summed over channels), which shows less position dependence than individual channels. The scaled charge and phonon templates are then summed with a random trigger taken from the Run 133 WIMP-search dataset, and processed through the usual analysis algorithms. The templates are summed with the randoms assuming no delay between the trigger and the pulse start time. While the true data has a delay due to the nonzero rise time of the phonon pulse, this assumption should

not affect the quality of the energy estimators because the optimal filter should always reconstruct events properly when the pulse start time is contained in the optimal filter search window. This pulse simulation procedure is repeated for every random trigger event taken in the analysis, to build up a fake data set.

The simulated data contain events constructed from randoms taken at the beginning of the run (827,000 events), in the middle of the run interspersed with real triggers (1,084,017 events), and at the end of the run (827,000 events). The efficiency of the quality cuts was estimated by taking the fraction of simulated events constructed from in-run randoms passing `!cGlitch1_v53`, `!cLFnoise1_v53`, `cPChiSq_v53`, and `cQChiSq_v53`. In-run randoms only were used because they were taken at the same time as physics data, and were believed to be more representative of the noise environment of real data. Figure 5-13 shows the nuclear recoil efficiencies as a function of energy for detector.

The phonon start time requirement `cGoodPStartTime_v53` cannot be determined using the method above because the simulation explicitly assumes that there is no delay between the trigger and the start of the phonon pulse. It also cannot be determined from calibration data because calibration data contains a significant amount of pileup events that fail `cGoodPStartTime_v53`. Because the optimal filter search window is quite wide and was optimized to select all non-pileup events, this cut should have very high acceptance, and we simply assume that it is 100% efficient with negligible uncertainty. As a cross-check, we confirm that 0 events from the open post-Cf WIMP-search data fail this cut on any detector, as expected.

While the pulse simulation provides robust estimate of the quality cut efficiency, we can also check the results using calibration and WIMP search data by performing a likelihood ratio test on the binned efficiency measurements.¹ To check the efficiencies of the `!cLFnoise1_v53` and `!cGlitch1_v53` cuts, the χ^2 cuts were first applied to Cf

¹Consider a two binned measurements of an efficiency, one from the simulation of Figure 5-13, and another from a validation sample of Cf or WIMP-search data. Each bin represents a binomial random variable, and we want to test whether the data in each bin have a common efficiency or two different efficiencies. We do this using a likelihood ratio test, with the null hypothesis being that each measurement is drawn from the same efficiency function, and the alternative hypothesis being that the two measurements are drawn from different efficiency functions. The likelihood ratio test statistic is given by

$$\lambda = \frac{\max_{\bar{p}} \prod_i \mathcal{L}(p_i; N_i^{Cf}, n_i^{Cf}) \mathcal{L}(p_i; N_i^{sim}, n_i^{sim})}{\max_{\bar{p}^{Cf}, \bar{p}^{sim}} \prod_i \mathcal{L}(p_i^{Cf}; N_i^{Cf}, n_i^{Cf}) \mathcal{L}(p_i^{sim}; N_i^{sim}, n_i^{sim})} \quad (5.2)$$

where the product in the numerator and denominator is taken over the energy bins. The likelihood function for each bin is given by the binomial distribution for n successes in N trials given a probability p

$$\mathcal{L}(p; N, n) \propto p^n (1 - p)^{N-n}. \quad (5.3)$$

We obtain the p -value using the fact that twice the negative log is asymptotically chi-square with M degrees of freedom for an efficiency measurement with M energy bins

$$-2 \ln \lambda \sim \chi^2(M). \quad (5.4)$$




Figure 5-12: Combined quality cut efficiency of !cGlitch1_v53, !cLFnoise1_v53, cPChiSq_v53, and cQChiSq_v53, as a function of total phonon energy, estimated from the Monte Carlo pulse simulation. The slight differences in energy-dependence between different detectors is due to relative differences in the position of the cut boundary. These differences can be driven by time-variation of the baseline noise, which affects the distributions of $\Delta\chi^2$ used to define !cGlitch1_v53 and !cLFnoise1_v53. Statistical error bars are present on each bin, but too small to see.

data to remove pileup. Then the fraction of events also passing one of the two cuts of interest was calculated as a function of energy. These fractional passage rates in Cf data were compared with the same quantities computed on simulation data, above the trigger threshold of each detector. For all detectors, the average passage rates are quite similar, and most detectors have p -values that are > 0.05 .

5.5 Trigger and Analysis Thresholds

When SuperCDMS triggers during WIMP-search data-taking, data is recorded on all detectors. In other words, if an event triggers detector A, data will be recorded on detector B even if there is no trigger. This has two consequences. First, strong noise fluctuations could reconstruct to energies at which the trigger efficiency is non-zero, thereby polluting the signal sample. This is an additional source of background. And second, the experiment is actually sensitive to sub-trigger threshold events, provided that they occur in coincidence with a trigger on another detector. For example, a low-frequency noise fluctuation could issue a global trigger on detector A in coincidence with a sub-threshold WIMP scatter on detector B. Although this is extremely improbable, it formally means that the net trigger efficiency is the actual efficiency of the trigger, plus the efficiency for measuring sub-threshold nuclear recoils. The latter is difficult to assess quantitatively. These two complications motivate our requirement that WIMP candidate events must cause a trigger the detector of interest. This requirement, called `cTriggeredEvent_133_LT`, solves both problems described above. The efficiency of this requirement is simply the efficiency of the trigger on each detector, which can be measured with routine Ba calibration data.

The trigger threshold of the experiment was adjusted periodically in response to changing overall baseline noise in order to maintain a fairly low and constant trigger rate of about 1 Hz. Since data at threshold is used in the analysis, the trigger efficiency must be well-measured at these energies. Ben Loer performed painstaking archaeology to determine all periods of constant trigger efficiency with useable data and then fit the trigger efficiency for each period. The functional form assumed for the trigger efficiency is

$$\epsilon(E_{pt}) = \frac{1}{2} \operatorname{erfc} \left(\frac{\mu - E_{pt}}{\sigma_0 + \sigma_m \sqrt{E_{pt}}} \right). \quad (5.5)$$

This is a typical error function trigger efficiency with an energy-dependent resolution term of $\sigma(E) = \sigma_0 + \sigma_m \sqrt{E_{pt}}$, rather than the usual constant σ_0 . This functional form was chosen because it fit the data better than a model with a constant resolution term, but it is physically motivated by the fact that the phonon resolution degrades with energy due to position-dependence.² Figure 5-14 shows fits to the trigger efficiencies

²The fit quality was essentially identical using either a linear or square-root energy-dependence of the resolution, but the square-root form was chosen because its convergence was slightly better. Matt Pyle has argued in his internal notes estimating the sensitivity of SuperCDMS at SNOLAB that the total phonon resolution scales as $\sigma_{pt}^2 = \sigma_0^2 + (f_p E_{pt})^2$, so the linear term should be most




Figure 5-13: Example of Ba calibration data on T4Z3, showing the effect of `cTriggeredEvent_133_LT`, which enforces that events on this detector issue a trigger. Red events are removed by `cTriggeredEvent_133_LT`, while blue events pass. Note the large number of sub-threshold events that are recorded on this detector with reasonable charge and phonon energies, illustrating the experimental sensitivity below the trigger threshold. This effect is much more noticeable in Ba calibration data than it would be single-scatter WIMP-search data due to the large number of multiple scatters in Ba data.

during a period in early 2013. Fits were performed for every period of distinct trigger efficiency. The trigger efficiency is assumed to remain constant within each period of constant threshold, which appears to be a reasonable assumption.

In addition to the trigger requirements, three types of time-dependent hard analysis thresholds were used:

1. **Global analysis threshold:** During periods of Run 133, the trigger thresholds on certain detectors were slightly below 2 keV in total phonon energy. The performance of event-by-event data quality cuts, however, degrades below 2 keV because of the poor signal-to-noise. This is particularly true for the low-frequency noise cuts. We therefore choose, somewhat arbitrarily, to only analyze data above 2 keV on all detectors except T1Z1, which used a 3 keV threshold. The higher threshold on T1Z1 was motivated by a noisy SQUID in the phonon channel PAS1 readout electronics, which could periodically cause noise fluctuations that reconstruct to higher energies.
2. **Trigger-dependent analysis threshold:** T1Z1, T4Z2, and T5Z3 have significant amounts of exposure during which their trigger thresholds are 0% between 2 keV and 3-5 keV in ptNF. Events that trigger the detectors and reconstruct at, say, 3 keV would be extremely unlikely to be WIMPs during these periods because the trigger is insensitive to these energies. To mitigate against spurious triggers at energies with little trigger sensitivity, we introduce another analysis threshold set effectively at 1σ below the 50% midpoint of the fitted trigger efficiency curve

$$E_{1\sigma} = \mu - \frac{\sigma_0 + \sigma_m \sqrt{\mu}}{\sqrt{2}}, \quad (5.6)$$

where μ , σ_0 , and σ_m are the parameters from the fits of the trigger efficiency. The choice of setting the threshold specifically at 1σ below the mean of the trigger efficiency is arbitrary, but balances the desire to use most of the data near threshold with the desire to avoid triggers at energies with little trigger sensitivity.

3. **Blinding cut thresholds:** The high-threshold blinding requirements described in Section 5.2 require that data be above a time-dependent threshold in order to be hidden from view. During certain periods of time, this blinding threshold was above both the global analysis threshold and the trigger-dependent analysis threshold, implying that some very low-energy events could potentially be visible to the analyzer and used in the analysis. For consistency with the blinding requirement, we also required that all data used in the analysis be above this time-dependent blinding threshold.³

relevant in the Taylor expansion at low energies. Using the square-root form over the linear one is therefore slightly awkward, but still incorporates the qualitative behavior that the phonon resolution degrades due to position-dependence at higher energies.

³As a technical aside, the high-threshold blinding cut threshold was set using the `psumOF` energy estimator rather than the `ptNF` energy estimator used in this analysis. In the cut construction, we also use `psumOF` for consistency, however we estimate the efficiency as a function of `ptNF` as an error

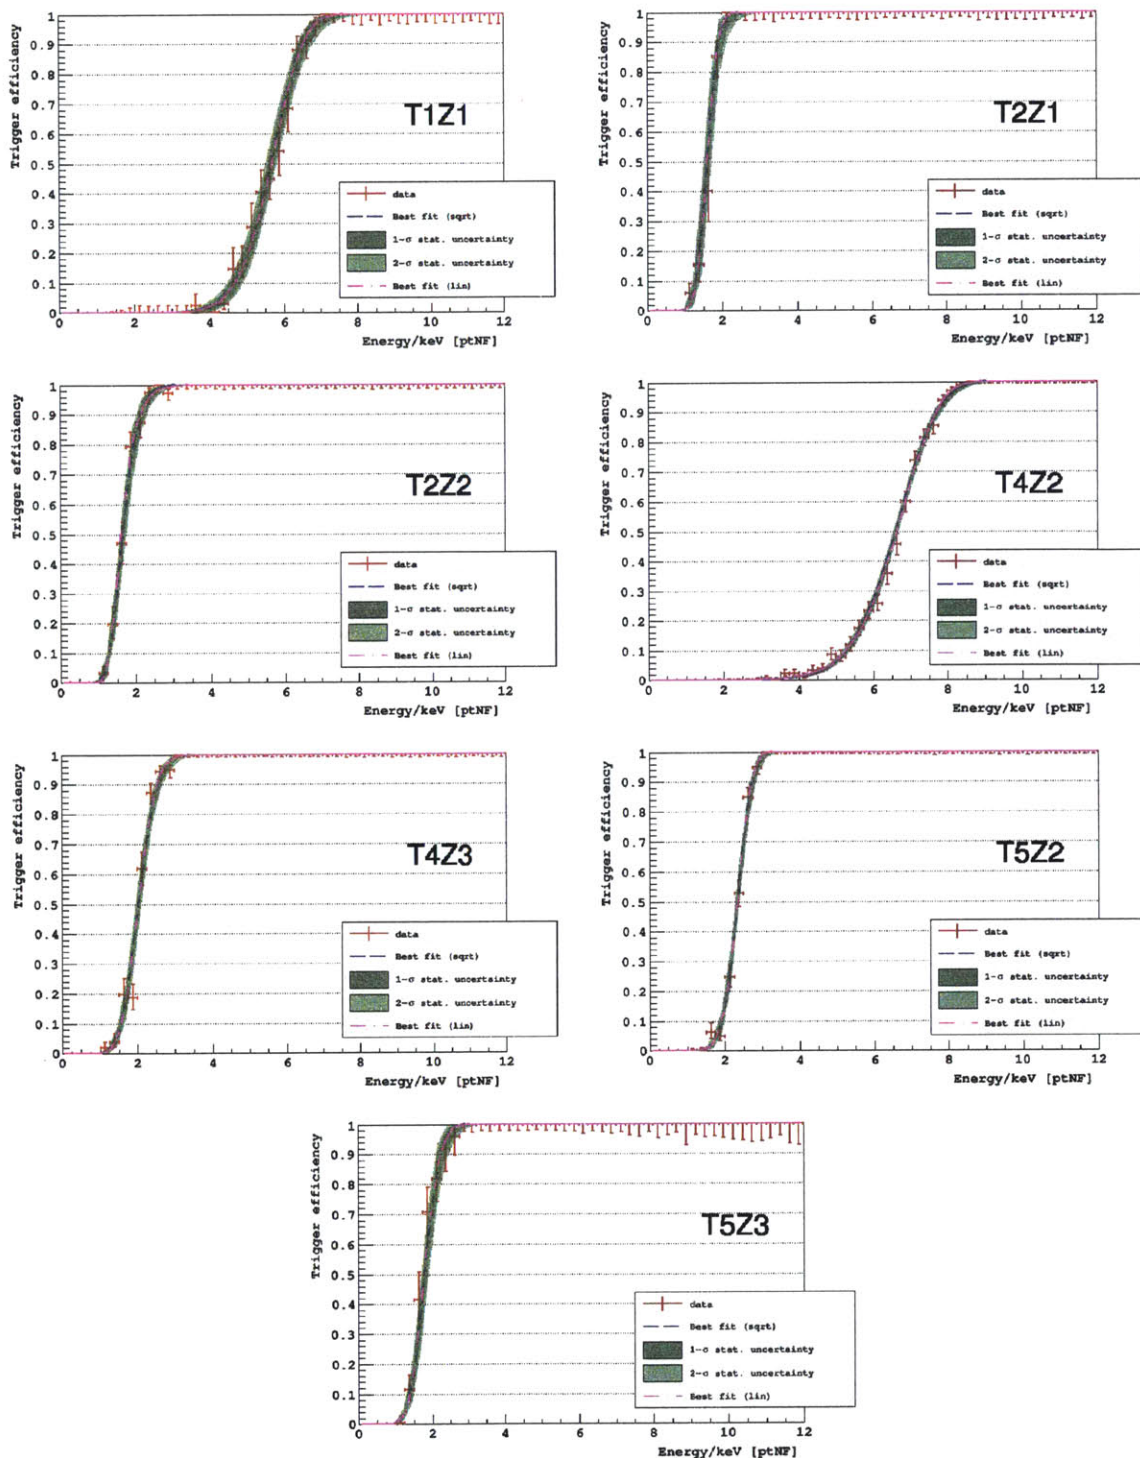


Figure 5-14: Fits of trigger efficiency functional form to Ba calibration during early 2013. Blue dashed line shows the best fit of the square-root model, while the red dashed line shows the best fit of the linear model. The green bands shows the extreme range of trigger efficiencies obtained by varying fit parameters within $\pm 1\sigma$ and $\pm 2\sigma$ of the best fit, according to the covariance matrix. (Courtesy of Ben Loer)

The three analysis thresholds are applied simultaneously so that events must be above all three in order to contribute to the analysis. Figure 5-15 shows the variation in the three analysis thresholds as a function of time. In particular, note that T1Z1 and T4Z2 have much higher analysis thresholds relative to the other detectors during much of Run 133. These detectors therefore do not contribute greatly to the sensitivity at the lowest WIMP masses of this analysis, but they do contribute some sensitivity at WIMP masses $\gtrsim 10$ GeV

Since we do not include time as a variable in the final analysis, we must integrate over the time-dependence of the trigger and analysis thresholds. This is equivalent to performing a live time-weighted average of the trigger and threshold efficiencies over all periods of unique trigger and analysis efficiency. The live time-averaged efficiency is given by

$$\bar{\epsilon}(E) = \sum_i \frac{T_i}{T} \epsilon_i^{trig}(E) \epsilon_i^{thresh}(E), \quad (5.7)$$

with the sum taken over each period of constant trigger efficiency and

$$\begin{aligned} T_j &= \text{livetime in the } j\text{th period of constant trigger and threshold} \\ T &= \text{total livetime} \\ E &= \text{energy} \\ \epsilon_i^{trig}(E) &= \text{trigger efficiency functional form for } i\text{th period} \\ \epsilon_i^{thresh}(E) &= \text{efficiency for analysis threshold of } i\text{th period} \end{aligned}$$

Figure 5-16 shows the average efficiencies as a function of energy for the trigger alone, analysis thresholds alone, and total.

Uncertainties on the trigger threshold efficiency from the fits are propagated by Monte Carlo simulation. The analysis threshold efficiencies are assumed to have no uncertainty. For each time period of the trigger efficiency fits, fit parameters are sampled from a multivariate gaussian given by the fit covariance matrix and used to generate efficiency curve samples. Samples from different periods are live time-averaged together with the analysis thresholds, as described above. The resulting uncertainty is generally quite small due to the good fit quality and the averaging together of different measurements.

5.6 Nuclear Recoil Energy Scale

As discussed in Section 3.1 the total phonon energy measured by a detector is the sum of the recoil energy and the Luke phonon energy, which is proportional to the charge energy. Computing limits and comparing to Monte Carlo simulations requires knowing the recoil energy of the event, which involves an assumption about the charge

function that accounts for the small (approximately gaussian) random difference between `psum0F` and `ptNF`. In contrast, since the global analysis threshold and trigger-dependent analysis thresholds are defined in terms of the primary analysis energy estimator `ptNF`, their efficiencies are simply heaviside functions.




Figure 5-15: Global analysis threshold (*purple*), trigger-dependent analysis threshold (*green*), and blinding analysis threshold (*blue*) as a function of time during Run 133 for each detector. The global analysis threshold is constant at 2 keV on all detectors except T1Z1 (3 keV). The trigger-dependent analysis threshold requires that events have an energy above 1σ -equivalent below the mean of the trigger efficiency curve. The blinding analysis threshold requires that WIMP-search candidates events be above the threshold used for the high-threshold blinding cut. The analysis threshold used for a given detector at a given time is the highest of the three thresholds.




Figure 5-16: Live time-averaged efficiency for all analysis thresholds and the trigger as a function of total phonon energy. The red curve shows the live time-averaged efficiency only, the green curve shows the live time-averaged trigger efficiency only, while the blue curve shows the combined efficiency. Note that the combined efficiency is *not* the simple product of the live time-averaged trigger and analysis threshold efficiencies: the multiplication of the trigger and analysis threshold efficiencies occurs before, not after, the live time averaging.

energy of that event. This defines three different energy scales

- **keV:** measured total phonon energy
- **keVnr:** recoil energy assuming the ionization yield of a nuclear recoil (“nuclear-recoil-equivalent energy”)
- **keVee:** recoil energy assuming the ionization yield of an electron recoil (“electron-equivalent energy”)

If we measure a total phonon energy, there are two ways to convert to the recoil energy

1. **Use measured ionization energy:** Since we measure both ionization and phonons, it is possible to use the measured ionization energy for each event to subtract the Luke phonon energy. Because of the poor charge resolution at low energy, this would produce very poor recoil energy resolution.
2. **Use a model for the ionization yield:** If we have a model of the ionization energy as function of total phonon energy, then we can use the model prediction for the charge to estimate and subtract the Luke phonon energy. The recoil energy resolution of this method is comparable to the total phonon resolution, but the assumption of a charge model introduces potential systematics.

We use the second method to determine the nuclear recoil energy scale. We parameterize the ionization energy as a function of total phonon energy by

$$E_Q = f(E_p) = \alpha_1 + \alpha_2 E_p + 10^{\alpha_3} \operatorname{erf} \left(-\frac{E_p}{10^{\alpha_4}} \right). \quad (5.8)$$

The parameters α_1 , α_2 , α_3 , and α_4 are determined for each detector individually by fitting the band of nuclear recoils in ^{252}Cf calibration data. The recoil energy of each event is then obtained by simply subtracting off the Luke phonon contribution from the total measured phonon energy

$$E_r = E_p - \frac{E_Q}{3 \text{ eV}} e\Delta V,$$

where we have divided the charge energy by 3 eV because this is the mean eh pair creation energy in Ge. This functional form was chosen because it fit the data reasonably well with good convergence.

An example of calibration data is shown in Figure 5-17. Nuclear recoils form a band in the plane of total phonon and ionization energies, which we refer to as the *nuclear recoil band*. Data in the nuclear recoil band are binned in total phonon energy and the mean ionization energies in each bin are fit to the model as a function of total phonon energy. An example fit is shown in Figure 5-18. Note that the agreement with the Lindhard ionization model introduced in Section 3.2.1 is generally good. When the fitted charge models deviate from Lindhard, they tend to fit for a lower ionization yield. This results in more conservative limits than if the Lindhard model were used.




Figure 5-17: Example of ^{252}Cf calibration data with NR band fit, for T4Z3. Blue points are within a $\pm 2\sigma$ gaussian band (blue lines) around the band mean, predicted by Lindhard (black dashed line). Red points are the mean of the band in energy bins. Light blue points are outside the nuclear recoil band. Note the band of events with high ionization yield, which are due to high-energy gammas, also produced by the ^{252}Cf source. (Courtesy of Julien Billard.)

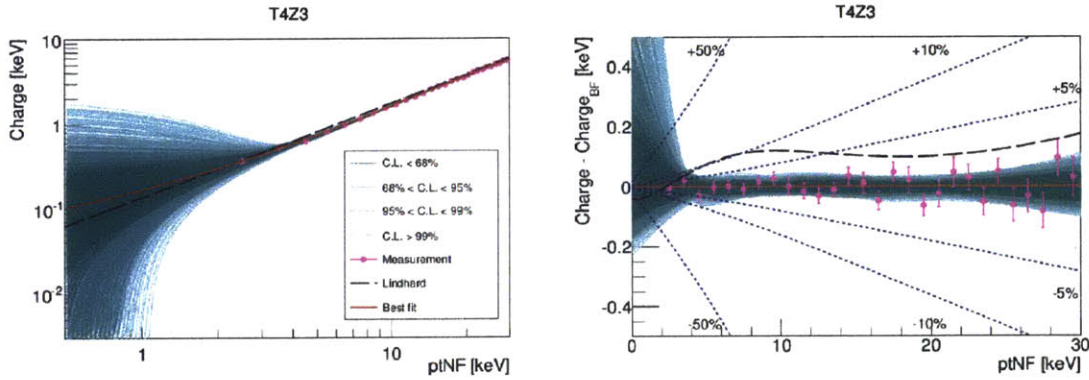


Figure 5-18: **Left:** Best fit charge model (red line), charge model of Lindhard (black dashed line), and charge models sampled from the uncertainty of the fit at various confidence levels (teal). **Right:** Residual of left panel. Blue dashed lines show lines of fractional deviation from the best-fit charge model. Note that Lindhard is generally within 10% of the best-fit charge model above our 2 keV total phonon energy threshold. (Courtesy of Julien Billard.)

5.7 Background and Detector Simulations

The background model for this analysis necessarily consists of two components: a *physics simulation* that describes the rates of the event types and their deposited energy in the detector, and a *detector simulation* that describes how that energy is manifested in the observable quantities of the detector. The physics backgrounds at low energy consist primarily of gammas and electrons produced by high-energy gammas from radioactive decays in the passive material surrounding the detector (“gamma-sourced” events), lines from L-shell electron capture of the isotopes ^{68}Ge , ^{71}Ge , ^{67}Ga , and ^{68}Zn (“activation line” events), and daughters of the decay of ^{210}Pb (“surface” events). While extensive development has occurred on a first-principles detector simulation by Daniel Brandt, Steve Leman, Kevin McCarthy, Peter Redl, and others, the aggressive timeline required for this analysis and the lack of validation with Soudan detector data available at the time motivated a simpler pulse-simulation approach to the detector simulation.

5.7.1 Physics Simulations

Gamma background

While gammas may be produced by decays within the detector, the majority of gammas that interact in detectors are due to decays that occur within the passive material surrounding the detectors. The primary gamma lines are due to the decay of U, Th, and K and their daughters. These gamma lines are mostly in the 100 keV to few MeV energy range, and they are all above the 0-10 keV energy range of interest. Events at low energies from these high-energy gammas are due to Compton scattering, which produces an approximately flat low-energy spectrum. We therefore take the low-energy spectrum of gamma-induced events to be flat in energy, normalized using the event rate in the electron recoil band where contamination from other event categories is negligible.

The gamma background is normalized using an approximate $\pm 2\sigma$ bulk electron recoil band between 5-12 keV total phonon energy. Events are counted in the sidebands data from the full Run 133 low-threshold exposure. Note that we do not apply any charge fiducial volume requirement to the normalization region in order to maximize the number of counts in the region and reduce the statistical uncertainty. Because this region contains a pure sample of gamma-sourced backgrounds, we can directly normalize the same region of parameter space in the trigger-efficiency-corrected simulation to this count rate.

One might question whether the gamma background in the range of 5-12 keV total phonon energy is actually only due to Compton scattering of high-energy gammas. The main reason to doubt this interpretation comes from measurements by the CoGeNT detector. Like the SuperCDMS iZIPs, CoGeNT is made of Ge, and is therefore subject to the same cosmogenic backgrounds as SuperCDMS. CoGeNT measures weak lines from ^{49}V (4.97 keVee, $t_{1/2} = 330$ d), ^{51}Cr (5.46 keVee, $t_{1/2} = 28$ d), ^{54}Mn (5.99 keVee, $t_{1/2} = 312$ d), ^{55}Fe (6.54 keVee, $t_{1/2} = 2.7$ y), $^{56,57,58}\text{Co}$ (7.11 keVee,

Detector	events (norm. region)	events (2-13.1 keV pre-cuts)
T1Z1	146 ± 12.1	552 ± 46
T2Z1	131 ± 11.4	693 ± 61
T2Z2	69 ± 8.3	383 ± 46
T4Z2	54 ± 7.3	219 ± 30
T4Z3	162 ± 12.7	828 ± 65
T5Z2	125 ± 11.2	754 ± 67
T5Z3 (BS)	32 ± 5.6	140 ± 25
T5Z3 (AS)	154 ± 12.4	803 ± 65

Table 5.6: Rates of gamma-sourced events in the normalization region and in the full energy range after correcting for the trigger efficiency, but before the effect of fiducial volume and nuclear recoil band requirements. Normalization region is a roughly $\pm 2\sigma$ portion of the electron recoil band in 5-12 keV total phonon energy.

$t_{1/2} = 77$ d, 272 d, 71 d), and ^{56}Ni (7.71 keVee, $t_{1/2} = 6$ d). These lines are unfortunately not resolved in the open post-Cf data because of the fairly low exposure and the fact that SuperCDMS has worse electron-recoil energy resolution than CoGeNT, so we are unable to include them in the model in a well-motivated way. See Chapter 6 for the unblinded low-background recoil energy spectrum.

Activation lines

Several electron capture radioisotopes are produced in Ge by cosmogenic activation (^{68}Ge , ^{68}Ga , ^{65}Zn) and activation from neutron calibration (^{71}Ge). Cosmogenic production of the long-lived isotopes ^{68}Ge and ^{65}Zn is well-known from simulations [90], while the short-lived isotope ^{68}Ga is a product of ^{68}Ge decay. Because these isotopes are produced by neutrons and cosmic rays, which have a long range in Ge, the activation in the detectors is spatially uniform to an excellent approximation. Following the electron capture, the electrons in the daughter atom are in an excited state which quickly relaxes by releasing X-rays and Auger electrons. This low-energy radiation has a short path length in Ge and therefore is fully contained in each detector, so that the energy of each electron capture decay is the full binding energy of the hole left vacant by the capture. Table 5.7 summarizes the key isotopes that contribute to SuperCDMS data.

Because the ratio of K-shell electron captures to L-shell electron captures is constant, we can use the strong K-shell lines in the unblinded high-energy data to estimate the rate of L-shell events that will lie near the signal region. To estimate the rate from the activation lines, we scale the K-shell rate in the full dataset by the ratio of L-shell to K-shell events measured in the post-Cf data

$$N_L^{\text{tot}} = N_L^{\text{pCf}} \frac{N_K^{\text{tot}}}{N_K^{\text{pCf}}}, \quad (5.9)$$

Isotope	K-shell energy	L-shell energy	Half life
^{68}Ge	10.36 keV	1.30 keV	270.95 d
^{71}Ge	10.36 keV	1.30 keV	11.43 d
^{68}Ga	9.66 keV	1.20 keV	67.71 m
^{65}Zn	8.98 keV	1.10 keV	243.93 d

Table 5.7: Binding energies and lifetimes for electron capture isotopes that are detectable in low-energy SuperCDMS data. The atomic de-excitation following the electron capture produces multiple x-rays and other radiation, but because events are fully contained, the event is reconstructed as a single event with energy equal to the binding energy of the hole left by the captured electron.

where the subscripts ‘L’, ‘K’ denote the shell of the electron capture, the superscript ‘tot’ denotes the number of events in the total exposure, and the superscript ‘pCf’ denotes the number of events in the post-Cf exposure. The event rates are defined to be *before application of the trigger efficiency*.

We use post-Cf data from April 2012 and January 2013 only to measure the L/K ratio. It was not possible to reliably measure the trigger efficiencies for the August 2012 post-Cf data due to frequent adjustment of the trigger thresholds. The selection criteria used in the measurement of N_L^{pCf} are the trigger condition `cTriggeredEvent_133_LT`, quality cuts described in Section 5.4.2, and the single-scatter condition. The selection criteria used in the measurement of N_K^{pCf} are the same, in addition to the charge fiducial volume. The application of the charge fiducial volume improves the resolution of the phonon lines, reducing the uncertainty on their relative sizes, at the expense of a modest loss in statistics. We use all Run 133 data to measure N_K^{tot} , applying the same selection criteria. Note that no correction is needed for the fact that the charge fiducial volume cut is applied to the estimate of the K-shell rates, but not the L-shell rate; the K-shell estimates appear in both the numerator and denominator in equation (5.9), so this efficiency term cancels.

The L-shell rates are extracted by fitting data in total phonon energy to a functional form

$$f_L(E) = N_s \mathcal{N}(E; \mu_L, \sigma_L) + N_b [f \exp(\lambda E) + (1 - f)B], \quad (5.10)$$

which consists of a gaussian component for the L-shell line with rate N_s , and a background component consisting of the sum of an exponential and flat distribution. The fit parameters are the overall signal N_s and background N_b rates, the fraction f of background in the exponential component, the mean of the line μ , and the resolution σ_L . This L-shell peak has contributions from the three isotopes in Table 5.7, but the three lines cannot be individually resolved, so they are modeled by a single “effective” peak with a floating resolution and mean. The inclusion of the exponential is included to avoid bias in case a small WIMP contribution exists in the post-Cf data. More realistically, it fits for the ^{210}Pb component that rises somewhat at low energies. The

Detector	L/K ratio	L-shell estimate (pre-trigger)	L-shell estimate (post-trigger)
T1Z1	0.105 ± 0.035	137.3 ± 46.1	10.6 ± 3.6
T2Z1	0.107 ± 0.029	87.4 ± 24.7	84.7 ± 24.0
T2Z2	0.092 ± 0.017	223.2 ± 43.1	214.1 ± 41.4
T4Z2	0.174 ± 0.026	289.8 ± 47.8	23.0 ± 3.8
T4Z3	0.092 ± 0.025	81.8 ± 22.8	78.0 ± 21.8
T5Z2	0.147 ± 0.030	110.5 ± 23.5	105.6 ± 22.5
T5Z3 (BS)	0.099 ± 0.028	16.0 ± 5.5	1.3 ± 0.4
T5Z3 (AS)	0.099 ± 0.028	35.5 ± 11.1	25.9 ± 8.1
mean L/K:	0.117 ± 0.010		
Ge (theory) [91]:	0.114		
Zn (theory) [91]:	0.108		

Table 5.8: Summary of results from study of electron capture isotopes. Ratios of L/K capture are in good agreement with theory. Third and fourth columns show estimates for the number of L-shell background events expected in Run 133 before and after accounting for the effect of the trigger efficiency.

K-shell rates are extracted by fitting data to a functional form

$$f_K(E) = N_b B + N_s [f_1 \mathcal{N}(E; \mu_{Zn} + \Delta, \sigma_K) + f_2 \mathcal{N}(E; \mu_{Ga} + \Delta, \sigma_K) + (1 - f_1 - f_2) \mathcal{N}(E; \mu_{Ge} + \Delta, \sigma_K)], \quad (5.11)$$

which consists of a flat background and three gaussians to model the K-shell lines. The fit parameters are the overall signal N_s and background N_b rates, the fractions f_i in each signal component, a resolution σ_K common to all lines, and an overall energy shift Δ to account for any calibration error. Fits are performed using unbinned extended maximum likelihood. Errors on fit parameters are taken from the covariance matrix for ease of propagation, but were checked to be in good agreement with profile likelihood errors. Figure 5-19 shows fits to the L-shell lines in post-Cf data, while Figure 5-20 shows the fits to the K-shell lines for all data. Similar fits are performed on the K-shell lines in post-Cf data. These are used to compute the L/K electron capture ratio for each detector, shown in Table 5.8 to be in excellent agreement with values expected for the Ge and Zn isotopes present in the detectors. Table 5.8 also lists the final background estimate before and after correcting for the trigger efficiency.

²¹⁰Pb daughters

The final component of the low-energy backgrounds in the iZIP is daughters from ²¹⁰Pb decay. During the assembly and fabrication of the experiment, Rn daughters plate out on the surfaces of detectors and their housings. These products quickly decay down to the long-lived isotope ²¹⁰Pb, which has a half-life of 22.2 years. The decay of ²¹⁰Pb is somewhat complex and is summarized in Figure 5-21. There is a




Figure 5-19: Fits to L-shell activation lines in post-Cf data. Error bars are much larger than Poisson because events have been weighted by the reciprocal of the trigger efficiency, in order to extract the pre-trigger L-shell rates.




Figure 5-20: Fits to K-shell activation lines in all Run 133 data. Detector T2Z2 is notable for being the only detector with a clearly resolved ^{68}Ga line at 9.66 keV. For reasons that are not well-understood, this detector has a higher level of the cosmogenic isotope ^{68}Ge , which decays to ^{68}Ga . Differences in the ratio of the ^{68}Ge and ^{65}Zn lines could be due to different levels of activation during neutron calibrations between the detectors.




Figure 5-21: Primary radiation from the decay of ^{210}Pb .

low-energy beta continuum from the decay of the ^{210}Pb . There is an 84% branching ratio into an excited state of ^{210}Bi with a total energy of 46.5 keV that may be released in a combination of gammas and conversion electrons. Since the ^{210}Bi remains on the surface of the detector or housing and only some of the conversion electrons and gammas enter the detector, the full 46.5 keV need not be measured from each decay. ^{210}Bi then decays by beta emission with a 1.1 MeV endpoint, producing little background below 10 keV. Finally, the resulting ^{210}Po undergoes an alpha decay to $\alpha + ^{206}\text{Pb}$. The ^{206}Pb nucleus ends up with 103 keV of kinetic energy. When the decay occurs in the housing surrounding the detector, however, the ^{206}Pb nucleus may lose energy in the housing before entering the detector resulting in a continuous energy spectrum between zero energy and 103 keV. The CDMS II low-threshold analysis found a significant background component of events at low or zero ionization energy, which were attributed in part to ^{210}Pb daughters [80]. Although the iZIP has considerably better ionization collection near the surface than CDMS II detectors, events at high radius are still expected to have low ionization yield, causing these events to form a major background. The estimation of the background for the ^{210}Pb chain consists first of a Geant4 simulation of the various components of the decay chain, followed by the overall normalization of those spectra to sidebands in the data.

The Geant4 [8] simulation of the ^{210}Pb spectral shape, developed by Peter Redl, is performed using the SuperCDMS simulation package called SuperSim, with the simulation geometry outlined in Figure 5-22. The simulation strategy is similar to that used in the iZIP surface event rejection analysis [61] and the CDMS II maximum likelihood analysis [92] that followed the result described here. Three different components of the decay products are simulated: decays in the housing, decays in the endplates above the top detector in the tower or below the bottom detector in the tower, and decays on the surfaces between two detectors, as shown in Figure 5-22. Since the geometry of each tower is identical, only a single tower is simulated and results are scaled to the rate expected for each individual detector in the normalization




Figure 5-22: Geometry of the ^{210}Pb in SuperSim. Contamination from three sources is simulated separately: contaminants on the endplates of the tower may decay and interact in the top and surface of the top detector or the bottom surface of the bottom detector, decays in the housing produce events in the detector sidewall, and decays on the surfaces adjacent to a neighboring detector produce surface events on the top or bottom surfaces. (Courtesy of Peter Redl)

procedure. The surfaces of the detectors and the Cu housings are contaminated with ^{214}Po , which decays by alpha emission to ^{210}Pb . As a result of the decay, the ^{210}Pb becomes implanted in a detector or housing surface. This implantation is crucial because it affects the energy loss of subsequent products in the decay chain before they hit the detector. For example, ^{206}Pb nuclei implanted in the housing may lose a significant fraction of energy before hitting the detector. Correct modeling of the implantation and propagation of the heavy ions involved in this decay chain requires using the “Screened Nuclear Recoil Physics List” (SNRPL) [93, 94] available with Geant4, which produces results compatible with the widely-used SRIM simulation package [95]. The energy deposition in the simulation is split into three groups, with spectra shown in Figure 5-23. The simulation shows the characteristic edge of the ^{206}Pb spectrum at 103 keV, as well as the conversion electron and gamma lines in the ^{210}Pb spectrum.

The overall normalization of the simulated spectra is performed in two ways with different systematics. The first method uses the rate of alpha decays of ^{210}Po to




Figure 5-23: Recoil energy spectra of components of Geant4 ^{210}Pb simulation for T2Z2. Top panel shows rate of sidewall events from housing simulation, while bottom panel shows rate of events from iZIP-iZIP simulation on top and bottom surfaces. Insets show low-energy region, with events below threshold overlaid in grey. The drop in events below threshold in the ^{210}Pb component of the simulation is due energy cuts in the simulation, but this does not affect rates above threshold. Component labeled ^{210}Pb is the combined energy deposition from all decay products that hit the detector.

normalize the rate of the entire chain, assuming secular equilibrium. This method has the advantage that the alphas have 5.3 MeV of energy and are therefore relatively easy to tag in the detectors. In the plane of total phonon energy and charge energy, alphas appear as a broad peak around 5 MeV in phonon energy at low ionization energy, below the main band of events from gammas. Events at these phonon energies cause strong TES saturation, so the peak is smeared out. The ionization yield is low because the alphas produce nuclear recoils near the detector surfaces. There is also a “tail” of events visible extending below the main alpha peak which are likely also due to alphas that lose a significant fraction of energy in the detector housing or end plate before hitting the iZIP. This tail is reproduced in simulations, but at a significantly lower rate than observed in the data, for reasons that are not understood. To count alphas, we define an acceptance region of these low-energy events, shown in Figure 5-24. The region is defined somewhat arbitrarily, but clearly captures the main alpha line and stragglers.

The second method for estimating the normalization of the ^{210}Pb chain is to count the number of events in a low-yield band where ^{206}Pb events should appear. This method may be biased systematically high because of additional non-Pb events that fall into the low-yield region, such as gammas with poor charge collection. It nevertheless provides a second method that can be used to bracket the range of systematic uncertainty of this background. The data region used to count ^{206}Pb sidewall events is defined by

- $-0.05 < \text{yield (ytNF)} < 0.2$
- S1, S2 inner charge consistent with noise (cQouter_LT_1, cQouter_LT_2)
- $15 \text{ keV} < \text{recoil energy (precoiltnf)} < 90 \text{ keV}$

For surface events on the top and bottom surfaces, events are counted in the region

- $-0.05 < \text{yield (ytNF)} < 0.2$
- NOT S1, S2 inner charge consistent with noise (cQouter_LT_1, cQouter_LT_2)
- $15 \text{ keV} < \text{recoil energy (precoiltnf)} < 90 \text{ keV}$

The alpha method and the low-yield counting have a systematic discrepancy, which we absorb as a systematic error described in Section 5.9.

Neutrons

The neutron background consists of two components: radiogenic and cosmogenic. Both were reestimated using Geant4-based simulations and a methodology used for the CDMS II experiment. The radiogenic component was estimated by fitting for the contamination levels of various passive materials in the experiment using the energy spectrum of gammas at above 10 keV. Fits were performed for the contamination levels of the ^{232}Th and ^{238}U chains, as well as ^{60}Co and ^{40}K . The resulting radiogenic neutron background from these contaminants, summed across detectors and




Figure 5-24: Outer alphas (sidewall events) in high-energy Run 133 WIMP-search data. Red events are in the alpha selection region, while blue events are events outside the selection region. Note that the nominal alpha energy of 5.3 MeV recoil energy is reduced to ~ 5 MeV total phonon energy because of TES saturation.

Background type	Sideband region for normalization
gammas	$\pm 2\sigma$ electron recoil band above upper edge of 2σ nuclear recoil band 5-12 keV pt
L-shell activation	K-shell lines with pt in 18-28 keV and -cQin1_v53_LT -cQin2_v53_LT -cQisym_v53_LT L-shell lines in post-Cf data with pt in 2.5-5 keV
^{210}Pb chain from housing	pt $\gtrsim 2000$ keV (see Figure 5-24) total charge $\lesssim 1000$ keV OR $-0.05 < \text{yield} < 0.2$ S1, S2 charge consistent with noise recoil energy in 15-90 keV
^{210}Pb chain on top and bottom	pt $\gtrsim 2000$ keV (see Figure 5-24) total charge $\lesssim 1000$ keV OR $-0.05 < \text{yield} < 0.2$ one of S1, S2 charge above noise recoil energy in 15-90 keV

Table 5.9: Sideband regions used to normalize background simulations, described in detail in Section 5.7.1

correcting for WIMP acceptance, is $(1.417 \pm 0.002(\text{stat}) \pm 0.053(\text{syst})) \times 10^{-2}$. The cosmogenic background estimate is based on simulations of cosmic rays through the Soudan cavern. The estimate for this background is slightly larger, but still small compared to the other backgrounds 0.08 ± 0.13 .

5.7.2 Detector Pulse Simulation

One of the important novelties of this analysis compared with the CDMS II analyses is the use of a pulse simulation as a kind of crude detector model that allows us to map Geant4 Monte Carlo simulations of the background physics directly into quantities that we actually measure in the analysis. We identify regions at high energies where background events can be sampled relatively cleanly. These high-energy events have excellent signal-to-noise, and so can be treated approximately as noiseless templates, which can be scaled down to the 0-10 keVnr energy range of interest and summed with random noise traces to form simulated background events. These simulated traces are then processed through the same data pipeline as standard CDMS data.

Because the high-energy template events do contain noise, it is important that they are chosen from energies high enough that the resulting template-plus-random does not have significantly worse resolution than real data. In practice, this requirement is

Background type	Template source in sideband
gammas	single-scatter ^{133}Ba calibration
L-shell activation	K-shell activation in WIMP-search
^{210}Pb from housing ^{210}Bi from housing	Qouter mid-yield in WIMP-search
^{206}Pb from housing	Outer low-yield in WIMP-search
^{210}Pb on top and bottom ^{210}Bi on top and bottom	Qinner asymmetric mid-yield in WIMP-search
^{206}Pb on top and bottom	Outer asymmetric low-yield in WIMP-search

Table 5.10: Source of templates used to simulate each of the background event types included in the physics simulations described in Section 5.7.2

not very restrictive. If typical data has a baseline noise n , and we scale our template by a factor f , then the noise of the template-plus-random simply adds in quadrature

$$n_{sum} = \sqrt{(fn)^2 + n^2} = n\sqrt{f^2 + 1}. \quad (5.12)$$

If we scale pulses by a factor of 1/3, the resolution degrades by only 5%.

There are many reasons to be skeptical of this simplistic approach, but the key reason that it works well for this analysis is that the signal-to-noise in the signal region is relatively poor. The distribution of the phonon partition variables is therefore mostly driven by resolution effects instead of detector physics effects. Even if the detector physics effects are only approximately captured by the pulse simulation, the difference would only be detectable with very high statistics.⁴ Here we discuss the construction of the pulse simulation used in this analysis, deferring further study and refinements of the model to the post-unblinding discussion in Chapter 6.

Each of the backgrounds described in Section 5.7.1 produce interactions that look different in the detector. We model this in the pulse simulation by lumping the backgrounds into six different categories each using a different template selection. The template categories are summarized in Table 5.10, and described in detail in the sections below.

Gammas

Templates for the gamma background are taken from Ba calibration. A random sample of events of 1000 events near 100 keV total phonon energy (43 keVee) are selected to be used as templates. Figure 5-25 shows the typical phonon pulses and charge pulses from an event of this energy, with generally excellent signal-to-noise (recall $\sigma_{pt} \sim 0.2 - 0.4$ eV). Only basic quality cuts are applied, so gammas as well as a small number of ejected electron that populate the yield region below the ER band are included. The ratio of yield \sim 1-events to low-yield surface events may differ between

⁴The high-threshold analysis of SuperCDMS data would therefore have a much harder time using this approach with success.

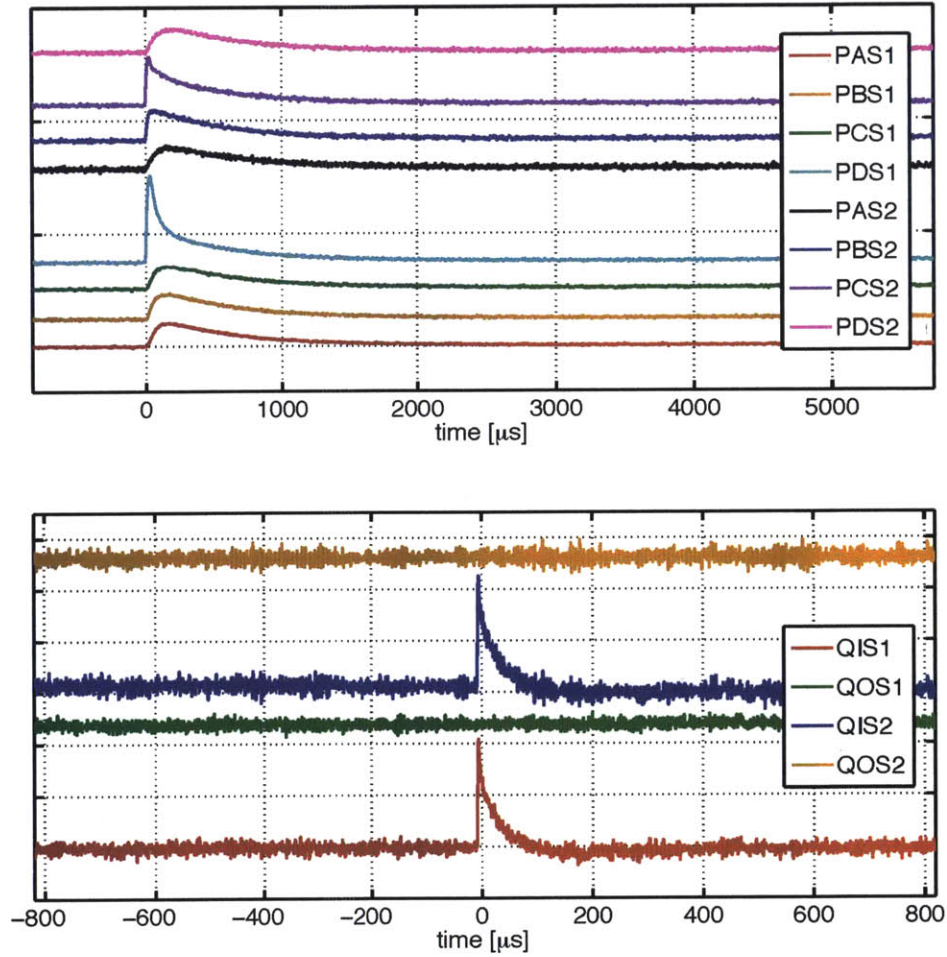


Figure 5-25: Typical 100 keV phonon (*top*) and charge (*bottom*) traces from Ba data used for gamma pulse simulation.

low energies and the energy region of interest. We do not correct for this potential difference. There also may be a difference in the spatial distribution of events between Ba data and low-background gammas. The main ^{133}Ba line is at 356 keV where the attenuation length in Ge is about 2 cm, while a significant fraction of gammas from the radiogenic backgrounds have few-MeV energies where the attenuation length is 4-5 cm [96].

L-shell Lines

The templates for the L-shell lines are taken from the K-shell events near 10 keV. These events have total phonon energy near 25 keV, which still represents a total signal-to-noise of 100. Although the activation line events are also electron recoils due to photons, like the gamma events above, the spatial distribution of the two event categories is significantly different. The activation line events are highly uniform throughout the detector because the electron capture isotopes are produced by neutrons and cosmic rays that have long path length in the detectors. In contrast, radiogenic gammas are more likely to interact in the outer surfaces of the detector. Events are selected in a narrow window around 10.4 keV in the recoil energy estimator `precoiltnf`.

^{210}Pb Backgrounds

Features characteristic of the ^{210}Pb backgrounds are visible in WIMP-search data above 20 keV, particularly when the charge channels are used to select events interacting near the detector surfaces. Figure 5-26 shows WIMP-search data from near the detector surfaces, clearly showing the ^{206}Pb band at low-yield as well as a broad band of electron and photon surface events from the other stages in the decay chain. The decay products of ^{210}Pb are simulated using four different types of templates to account for spatial variation and differences in ionization yield. The first two categories consist of events from decays that occur in the copper housings of the detector, while the second two categories consist of events that interact with the top and bottom faces of the detector. Because the charge information provides highly efficient tagging above 5 keV of events near the sidewalls and events on the top and bottom surface, we use this charge information to divide template candidates into disjoint “Qouter” and “Qinner” samples depending on whether they have charge signal consistent with a sidewall event or a face surface event. Within each charge classification, potential templates are split into a high-yield sample and a low-yield sample. The betas and gammas that result from the decays of ^{210}Pb and ^{210}Bi are electron recoils and therefore are expected to have higher ionization than the nuclear recoils induced by the ^{206}Pb nucleus.

- **Qouter mid-yield (decay products of ^{210}Pb and ^{210}Bi):** These events consist primarily of electrons (betas and conversion electrons) and low-energy gammas which do not penetrate deeply into the detector and therefore have reduced ionization yield. To select these templates, we first require that the charge energy on the inner ionization electrodes be consistent with no signal




Figure 5-26: Events selected as templates for the pulse simulation on ^{210}Pb background on T1Z1. Top panel shows “Qouter” events that have a charge signal consistent with an interaction in the outer radial surface of the detector, while bottom panel shows “Qinner” events that have charge consistent with an interaction in the top or bottom surface of the detector. Green events are selected as templates for betas, while purple events are selected as templates for ^{206}Pb nuclei.

(cQout1_v53_LT and cQout2_v53_LT). We then require that the events lie in the ionization yield (ytNF) range of 0.2 to 0.8. While this range is chosen somewhat arbitrarily, we can identify fairly clearly the beta-decay spectrum from ^{210}Pb decay in the top panel of Figure 5-26 below about 40 keV, due to the high density of points. Our ionization yield acceptance range roughly covers the yields where the betas are found, suggesting that it should do a reasonable job selecting electrons.

- **Qouter low-yield (^{206}Pb nuclei):** These events have the same Qouter selection criteria as the betas, but have a lower ionization yield requirement of between -0.05 and 0.2. This is chosen because it very clearly brackets the low-yield band of ^{206}Pb nuclei visible in Figure 5-26. This band is immediately identifiable by its 103 keV endpoint.
- **Qinner mid-yield (decay products of ^{210}Pb and ^{210}Bi):** The Qinner surface events are selected by first requiring that they fail the Qouter selection (NOT cQout1_v53_LT and cQout2_v53_LT) to enforce that the Qinner and Qouter templates are drawn from disjoint populations. Next, we require that the templates be surface events by looking for asymmetric charge collection. The energy measured on one of the inner charge electrodes, side 1 (qi10F) or side 2 (qi20F), must be between -1 keV and +1 keV. In addition, the event must have charge signal consistent with noise on the two outer charge electrodes (cQin1_v53_LT and cQin2_v53_LT). The mid-yield templates are chosen from a yield range of 0.3 to 0.85 and a recoil energy range of 40 keV to 150 keV. A higher range is used for the Qinner template selection than for the Qouter templates because the charge collection on the top and bottom surfaces near the electrodes is better than near the detector sidewall.
- **Qinner low-yield (^{206}Pb nuclei):** There are very few ^{206}Pb nuclei visible in the data. On several detectors, there are zero events above 40 keV that can reasonably be attributed to ^{206}Pb . In addition, one can see from Figure 5-23 that the total ^{206}Pb background on the detector faces is expected to be $\sim 1\%$ of the total ^{210}Pb background. Because of the difficulties in doing a data-driven pulse simulation with a rare background, we neglect this background source.

5.8 WIMP Consistency Requirements

Before constructing selection criteria to optimize the sensitivity of the analysis, we first define criteria to enforce that the events that we do accept into the signal region at least have characteristics that are *consistent* with nuclear recoils. These WIMP consistency requirements are intended to have high efficiency, but remove events that are far away from the main nuclear recoil distribution. We define consistency cuts on all of the variables that are used by the final discrimination cuts described in Section 5.9: ionization (qimean0F) and total phonon energy (ptNF), radial phonon partition (prpart0F), and z phonon partition (pzpart0F). We also

require that WIMP candidates be *single-scatters*, meaning that only one detector is above a phonon threshold.

- **$\pm 3\sigma$ nuclear recoil band:** (cNR_qimean_3sig_v53_LT) The nuclear recoil band is shown in Figure 5-18. The band is defined by dividing the ^{252}Cf calibration data into bins of total phonon energy and then fitting the charge energy distribution in each phonon energy bin to a gaussian functional form. The $\pm 3\sigma$ quantiles are fit, as a function of total phonon energy to a functional form, defining the edges of the nuclear recoil band.
- **Phonon partition consistency:** (cConsistency_v53_LT) This requirement simply removes extreme outliers in both prpartOF and pzpartOF. It is defined to remove the most extreme 1% of events in both the positive and negative directions for both variables, for ^{252}Cf data that has been reweighted to match the energy spectrum of a $5 \text{ GeV}/c^2$ WIMP.
- **Single-scatter:** (cPmultTight_v53_LT) This requirement removes events in which multiple detectors are above their respective phonon thresholds. Phonon thresholds are defined as function of time so that the WIMP acceptance is 99%.

5.9 Discrimination Algorithms

We lastly define a set of *discrimination cuts* that are tuned to optimize the expected sensitivity of the analysis by separating signal from background. Three different discrimination strategies, in increasing level of complexity, were designed: a partially-optimized cut-based analysis, a fully-optimized cut-based analysis, and a boosted-decision tree (BDT). For the sake of brevity and because of its significantly worse sensitivity, the partially-optimized analysis is not described here.

Both the fully-optimized cut based analysis (hereafter called “cut-based analysis”) and the BDT use four discriminator variables to separate signal and background

- ionization energy (qimeanOF)
- total phonon energy (ptNF)
- radial phonon partition (prpartOF)
- z -phonon partition (pzpartOF)

The first two variables together provide an effective way of separating the WIMP signal from the gamma background in the electron recoil band. The third variable can discriminate sidewall events, and the fourth variable can discriminate top/bottom surface events at low energies. Additional variables were not included because of concerns about the robustness of the background modeling in high-dimensional spaces. Significant background mismodeling could cause the discrimination algorithms to become overly sensitive to regions of parameter space in which the background model was unreliable.

Both the cut-based algorithms and the BDT are optimized in order to provide the best expected sensitivity of the final analysis. Before unblinding, the optimal interval method [97] was chosen as the default limit-setting method. The optimal interval method sets an upper limit on the WIMP-nucleon cross section using data from the energy interval that provides the best upper limit, including the proper statistical penalty for the freedom to choose the most constraining interval. Estimating the sensitivity of a proposed cut ensemble with the optimal interval method is numerically expensive: Monte Carlo realizations of the energy spectrum must first be generated, and then the optimal interval limit must be calculated for each in order to get the mean or median upper limit. When numerically optimizing a cut location, this procedure must furthermore be repeated many times. Instead of finding the cut ensemble that minimizes the expected optimal interval upper limit, we minimize the expected 90% CL upper limit for a Poisson counting experiment.

If a particular realization of the experiment observes N total events, then the 90% CL Poisson upper limit on the mean number of events is $\mu_{90}(N)$ such that

$$0.1 = \sum_{n=0}^N P(\mu_{90}; n) = \sum_{n=0}^N \frac{\mu_{90}^n e^{-\mu_{90}}}{n!}. \quad (5.13)$$

If we expect a rate μ_{bg} of background events passing a trial cut ensemble, we can compute the corresponding *expected* 90% CL upper limit $\bar{\mu}_{90}(\mu_{bg})$ with the relation

$$\bar{\mu}_{90}(\mu_{bg}) = \sum_{n=0}^{\infty} \mu_{90}(n) P(\mu_{bg}; n). \quad (5.14)$$

The corresponding expected upper limit on the cross section is then just

$$\sigma_{90} = \frac{\bar{\mu}_{90}}{\mu_{-42}(\{K_i\}, \epsilon(E))} 10^{-42} \text{ cm}^2 \quad (5.15)$$

where $\mu_{-42}(K_i, \epsilon(E))$ is the number of expected events from a WIMP with cross section 10^{-42} cm^2 , given a particular choice of cuts K_i with an efficiency $\epsilon(E)$. Because the Poisson distribution is numerically cheap to evaluate and the values of $\mu_{90}(n)$ in equation 5.13 can be tabulated for reasonably small values of n , calculating the expected upper limit from the Poisson distribution is much faster than using the optimal interval method.

At a single WIMP mass, the Poisson optimization method should produce a similar optimal cut ensemble as a optimal interval method. Cuts optimized with the optimal interval method will result in a background distribution that, on average, will set limits using a preferred energy interval with maximal sensitivity. If data in this energy interval tends to produce optimal sensitivity, then it will also tend to give the best Poisson upper limits. Even though the numerical values of the Poisson and optimal interval methods may be different, the energy range with the best sensitivity is the same in both cases.

5.9.1 Signal Model

We construct a model of the signal from Cf calibration data for use in optimizing the discrimination selection criteria. We first apply some standard data quality cuts to the calibration data and then apply the nuclear recoil band to isolate a high-efficiency (unbiased) but quite pure sample of nuclear recoils. The Cf nuclear recoils are then reweighted in energy bins in order to match the spectrum expected for WIMPs at each optimization mass, correcting for the efficiency of the selection criteria.

The Cf data used passed the same data quality requirements as typical WIMP-search data and were required to trigger the experiment (`cTriggeredEvent_v53_LT`). Events were also required to pass the ionization fiducial volume (`cQin1_v53_LT`, `cQin1_v53_LT` and `cQisym_v53_LT`), 3σ nuclear recoil band (`cNR_qimean_3sig_v53_LT`), and phonon partition consistency requirements (`cConsistency_v53_LT`). Both single and multiple scatters were used, since the distinction does not affect the efficiency of requirements on a single detector.

For a given optimization WIMP mass M , each event i in the Cf data is given a weight

$$w^i = f^{ij} \times \frac{dR}{dE_{pt}}(E_{pt}^i; M) \times \epsilon(E_{pt}^i). \quad (5.16)$$

The data are binned in total phonon energy E_{pt} so that an event in energy bin j is weighted by the reciprocal of the number of events in this bin, $f^{ij} = 1/N_j$ in order to transform the Cf energy distribution into a uniform distribution. The weights are then multiplied by the WIMP rate as a function of total phonon energy $dR/dE_{pt}(E_{pt}^i; M)$ for a reference WIMP-nucleon cross section, in order to reproduce the WIMP energy spectrum. Finally, the weights are multiplied by the nuclear recoil acceptance as a function of energy for all cuts that are applied to the real data before the final discrimination cuts. Figure 5-27 shows the application of successive stages of the reweighting procedure.

5.9.2 Cut-Based Algorithm

The cut-based discrimination algorithm uses linear cuts in three discriminator variables, set in bins of total phonon energy (ptNF). The discriminators used are the distance in charge energy (qimeanOF) from the mean of the nuclear recoil band, the radial phonon partition, and the z phonon partition. The cuts in all energy bins and all discriminator variables are numerically optimized simultaneously to produce the minimum expected upper limit based on the background model. Cuts are optimized for 5, 7, 10, and 15 GeV WIMPs separately for each detector.

While data would ideally be binned in energy as finely as possible, this is limited in practice by the difficulty of numerically minimizing the expected upper limit over a large number cuts values. After testing a range of different binning configurations and numerical minimizers for this specific problem, we determined that we could reliably minimize the expected upper limit over approximately 25 cut positions using the “interior-point” algorithm implemented by the `fmincon` routine in MATLAB [98]. The Nelder-Mead simplex algorithm [99] and stochastic optimization techniques (e.g.




Figure 5-27: Left panel shows typical energy spectrum of Cf data (T2Z1) after quality and trigger selection criteria before reweighing (*blue*) and after reweighing to a uniform distribution (*green*). Right panel shows the data after applying weights to match a 10 GeV WIMP spectrum (*red*), correct for quality cut efficiency (*blue*), trigger and analysis threshold efficiencies (*green*), charge fiducial volume efficiency (*purple*), 3σ nuclear recoil band efficiency (*orange*), and phonon partition consistency requirements (*yellow*). Both panels are pdfs, so units on vertical axis are arbitrary.

simulated annealing [100]) were also tested, but they tended to be either slower or less robust. Different binning configurations were chosen for each WIMP mass in order to match the endpoint and steepness of the energy spectrum.

The Cf data tends to have a small enough number of events in each energy bin that statistical fluctuations become important. In some regions of the background parameter space, the statistics are also poor. These problems are especially severe for the optimization of the 5 GeV cut since the majority of events above threshold would be below 5 keV, reducing the number of statistics available in the signal and background model. Statistical fluctuations in the signal and background models produce kinks in the sensitivity function that causes the numerical minimizer to easily become trapped in local minima. To alleviate this problem, we use a weighted 3-dimensional kernel density estimate (KDE) to model the discrimination variables in each energy bin [101, 102]. The KDE uses a 3-dimensional gaussian kernel function

$$\mathcal{K}(\mathbf{x}) = \frac{1}{(2\pi)^{3/2}} \exp(-\mathbf{x}^T \mathbf{x}/2). \quad (5.17)$$

Because each of the events used in the KDE has a different weight w^i applied in order to correctly reproduce the signal and background distribution shapes, we also must apply these weights in the KDE sums. Finally, we must specify the bandwidth of the kernel in our KDE. The covariance matrix for a 3-dimensional gaussian has 6 degrees of freedom that must be specified. To simplify the bandwidth estimation, we assume

a diagonal covariance matrix. After these modifications, the KDE of the signal or background distributions takes the form

$$\hat{f}(\mathbf{x}; \{\mathbf{d}^i\}) = \sum_{i=1}^{N_{data}} w^i \frac{1}{\sqrt{\det \Sigma}} \mathcal{K}(\Sigma^{-1/2}(\mathbf{x} - \mathbf{d}^i)) \quad (5.18)$$

$$= \sum_{i=1}^{N_{data}} w^i \frac{1}{\sigma_1 \sigma_2 \sigma_3} k\left(\frac{x_1 - d_1^i}{\sigma_1}\right) k\left(\frac{x_2 - d_2^i}{\sigma_2}\right) k\left(\frac{x_3 - d_3^i}{\sigma_3}\right), \quad (5.19)$$

where the $\{\mathbf{d}^i\}$ are the data coordinates, and the σ_i are the bandwidths for the three observables of the KDE, and the function

$$k(x) = \frac{1}{\sqrt{2\pi}} \exp(-x^2/2) \quad (5.20)$$

is the unit gaussian.

Optimization of the bandwidths σ_1 , σ_2 , and σ_3 is the main challenge of KDE techniques, particularly in dimensions higher than two, where the full distribution cannot be visualized. We test two common bandwidth choices: Silverman's rule of thumb [103] and maximum-likelihood cross-validation (MLCV) [101].⁵ Figures 5-28 and 5-29 show the 1- and 2-dimensional projections of the full KDE, overlaid with a (weighted) random sample of the underlying data. There is good agreement between the data and KDE, and it is apparent that the two bandwidth optimization methods give very similar results. We use the Silverman method because it is much faster to compute.

In each energy bin of each observable, we define an acceptance region $[a, b]$. The endpoints a, b are the free cut parameters that we vary to optimize the sensitivity. For computational efficiency, the acceptance region for `prpartOF` has its lower endpoint fixed to zero ($[0, b]$) because we expect all backgrounds to have higher values of `prpartOF` than the Cf calibration data. The total signal acceptance in energy bin j

⁵Silverman's rule of thumb is the following prescription for the bandwidth

$$\sigma = \left(\frac{4}{3n}\right)^{1/5} \sqrt{\text{var } x}, \quad (5.21)$$

where n is the number of points, and $\text{var } x$ is the variance of the variable x . This choice is optimal in the case that the data is drawn from a normal pdf, and it tends to perform well for distributions that are approximately gaussian. MLCV treats bandwidth selection as a parameter-estimation problem, by maximizing the likelihood that each data point is drawn from a KDE constructed of all the other points. In the case of univariate data, the MLCV bandwidth is found by maximizing the likelihood function

$$\mathcal{L}(\sigma) = \frac{1}{(n-1)\sigma} \prod_{j=1}^n \sum_{\substack{i=1 \\ i \neq j}}^n k\left(\frac{x_i - x_j}{\sigma}\right). \quad (5.22)$$

Because we have assumed a diagonal form for the covariance matrix of the kernel function, the bandwidth can be estimated for each 1-dimensional distribution independently.

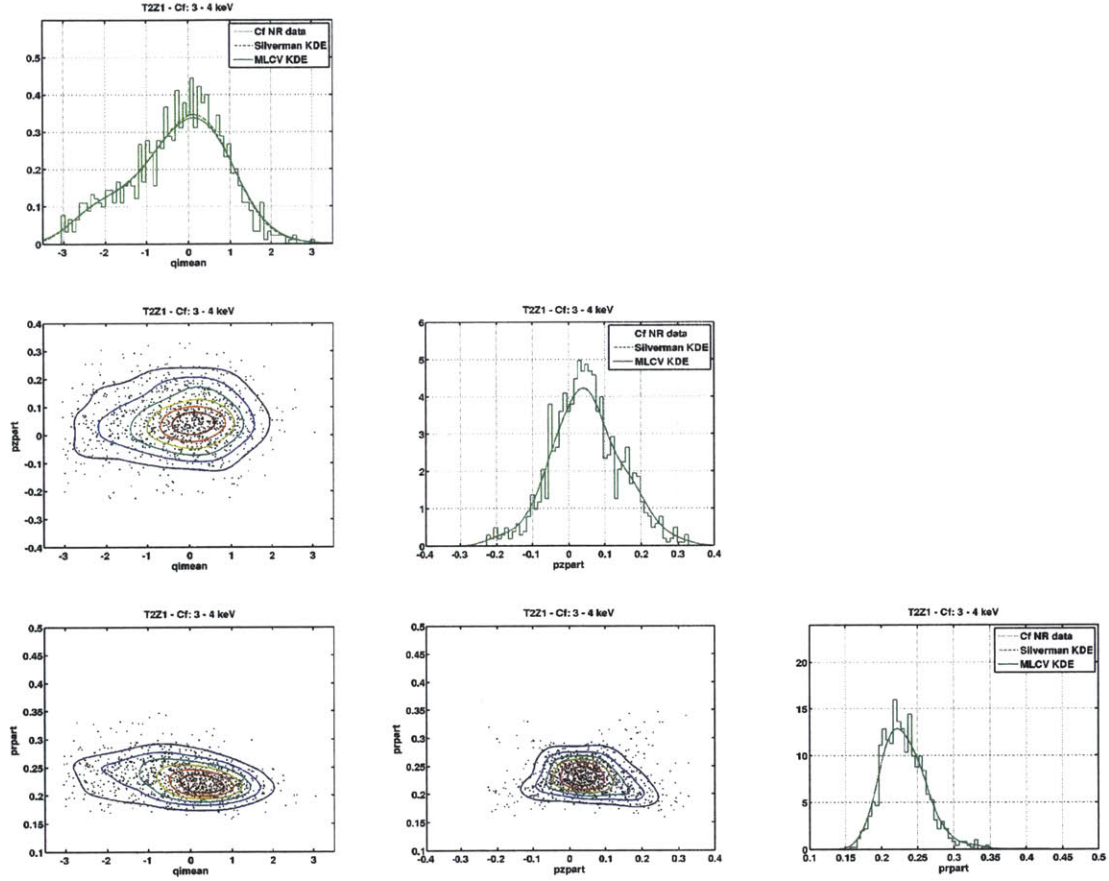


Figure 5-28: Projections of signal model KDE in the bin of 3-4 keV total phonon energy on T2Z1. Colored lines on 2-dimensional projections are lines of constant density of the KDE using the Silverman bandwidth. Both the Silverman and MLCV bandwidth estimates are shown to provide very similar results in the 1-dimensional projections.

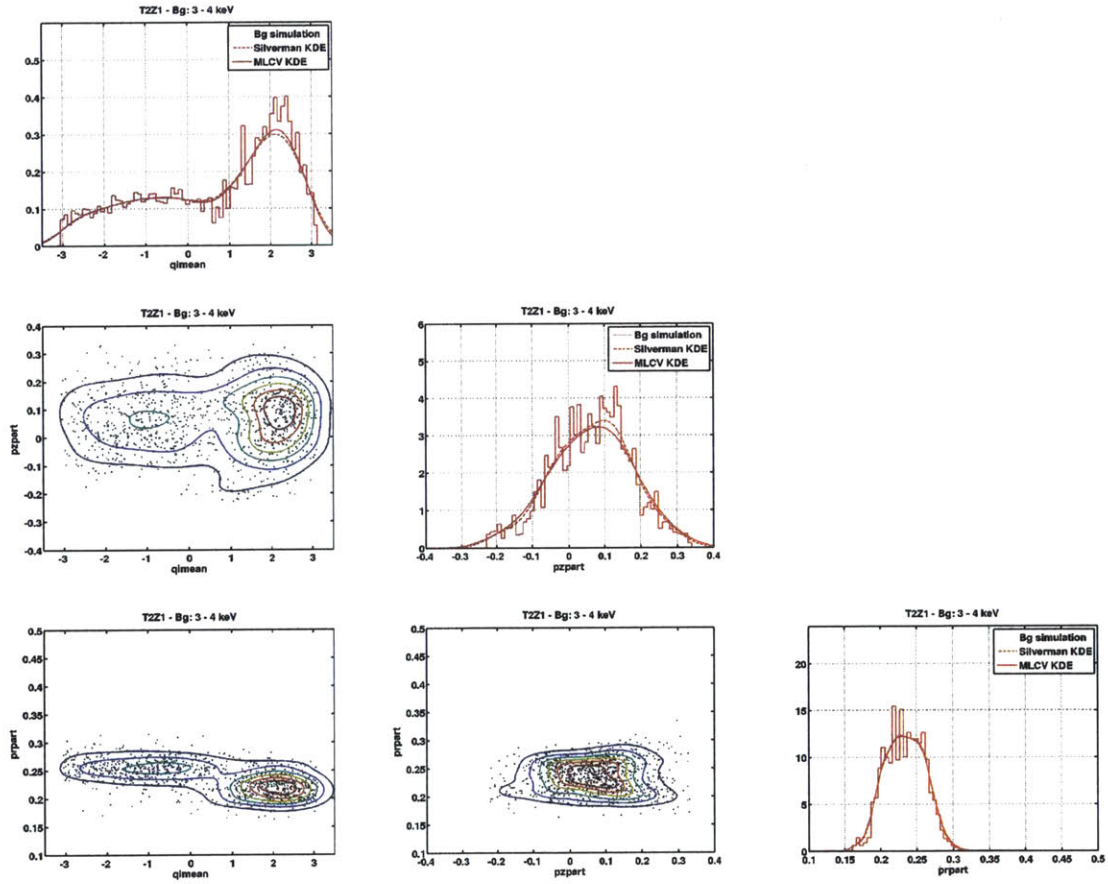


Figure 5-29: Same as Figure 5-28, except for the total background model of T2Z1 in the bin of 3-4 keV total phonon energy.

is therefore just the integral of the Cf KDE

$$\epsilon^j = \int_{a_1}^{b_1} \int_{a_2}^{b_2} \int_{a_3}^{b_3} dx_1 dx_2 dx_3 \hat{f}(\mathbf{x}; \{\mathbf{d}^i\}) \quad (5.23)$$

$$= \sum_{i=1}^{N_{\text{bin } j}} w^i \frac{1}{\sigma_1 \sigma_2 \sigma_3} \int_{a_1}^{b_1} dx_1 k\left(\frac{x_1 - d_1^i}{\sigma_1}\right) \quad (5.24)$$

$$\int_{a_2}^{b_2} dx_2 k\left(\frac{x_2 - d_2^i}{\sigma_2}\right) \int_{a_3}^{b_3} dx_3 k\left(\frac{x_3 - d_3^i}{\sigma_3}\right) \quad (5.25)$$

The background in each bin is obtained by exactly the same method applied to the background KDE

$$B^j = \sum_{i=1}^{N_{\text{bin } j}} \beta_i \frac{1}{\sigma_1 \sigma_2 \sigma_3} \int_{a_1}^{b_1} dx_1 k\left(\frac{x_1 - d_1^i}{\sigma_1}\right) \quad (5.26)$$

$$\int_{a_2}^{b_2} dx_2 k\left(\frac{x_2 - d_2^i}{\sigma_2}\right) \int_{a_3}^{b_3} dx_3 k\left(\frac{x_3 - d_3^i}{\sigma_3}\right). \quad (5.27)$$

After running the optimization, we obtain optimized cut boundaries such as in Figure 5-30 and 5-31. These figures clearly show that the cut boundary avoids regions with particularly high background density, such as near the L-shell activation lines and ^{210}Pb at low energies, as in Figure 5-30. In the 2-dimensional plane, the acceptance rectangles selected by the algorithm show similarly expected behavior, avoiding the main background regions.

The use of rectangular acceptance regions, however, is not necessarily completely optimal. Higher acceptance could probably be achieved in Figure 5-31 left panel, for example, but removing curved regions around the background. This is difficult to implement in practice because of the highly non-gaussian shape of the background in the input variables. Finally, it is only practically possible to optimize detectors on an individual basis in this cut-based algorithm. Since each detector has 25 free cut parameters, simultaneously optimizing detectors is computationally intractable. Since some detectors have significantly better average discrimination than others (e.g. because of differences in threshold), the combination of cuts across detectors may not be optimal, without further modifications of this method. These drawbacks motivate the use of more sophisticated multivariate classifiers discussed below.

5.9.3 Boosted Decision Tree

The multivariate analysis method based on boosted decision trees (BDTs) is a powerful algorithm for classifying events into signal and background categories that is widely used in high-energy physics (see [104] for a pedagogical review). BDTs have several features that are attractive for physics analyses. They can combine a large number of weak classifiers into a single strong discriminator that determines the extent to which an event is signal-like or background-like. They are non-parametric, in




Figure 5-30: Example of optimized cut boundaries as a function of total phonon energy on T4Z3 from cut-based discrimination algorithm. Red points are samples from the weighted background simulation, green points are samples from the Cf-based signal model, and blue line is the cut boundary as a function of the total phonon energy. Panels show (left to right) charge energy distance from NR band mean, radial phonon partition, and z phonon partition.




Figure 5-31: Example of optimized cut boundaries in bin of 4-6 keV total phonon energy on T4Z3 from cut-based discrimination algorithm. Red points are samples from the weighted background simulation, green points are samples from the Cf-based signal model, and black dashed box is the cut acceptance region. Panels show (left to right) charge energy distance from NR band mean vs. radial phonon partition, z phonon partition vs. radial phonon partition, and z phonon partition vs. charge energy distance from NR band mean.

the sense that the BDT algorithm does not assume an underlying distribution of the data or require tedious preconditioning of the training data. And they are generally easy to train with good performance, given sufficient training data.

A *decision tree* is a series of binary decision criteria, which determine whether to classify an event as either signal or background. Let the training data for the tree be $\{\mathbf{S}_i\}$ with weights $\{s_i\}$ for signal and $\{\mathbf{B}_i\}$ with weights $\{b_i\}$ for background, each containing some input variables (x^1, \dots, x^n) . The training of the tree proceeds as follows. For each input variable, find the value of the variable that provides the best separation between the signal and background training samples. Then find the variable whose splitting value gives the best separation. Assign the chosen splitting value to that *node* of the tree. Split the training samples according to the splitting criteria, and repeat the procedure on the subsamples. If the separation cannot be improved by splitting at a given node, or if the node satisfies some stopping condition, then do not perform any split at that node, and call it a *terminal node* or *leaf*. This process continues until all open nodes are declared leaves. Various stopping conditions are possible, but decision trees commonly use *depth* as a stopping condition; that is, the maximum number of nodes that can be traced from top to bottom in the tree.

To apply a decision tree to an event \mathbf{x} , we simply following the splitting values in each node to determine to which leaf to assign the event. Once the leaf is determined, there are a variety of choices for what numerical value to assign to the event. Most often, the value assigned is the purity $p = \sum_i s_{n(i)} / (\sum_i s_{n(i)} + \sum_i b_{n(i)})$ —the sum of the signal weights in the leaf divided by the total sum of the weights—or simply a value of +1 if $p > 1/2$ (signal-like) or -1 if $p < 1/2$ (background-like).

Not surprisingly, simple decision trees have a variety of problems. The most significant drawback of decision trees is that they are *unstable*: small changes in training data can produce large changes in the tree structure and output. For this reason, they tend to have poor predictive performance [105]. *Boosting* is a way of combining multiple weak decision trees in order to improve the stability and predictive power of the output [106]. The process of boosting always involves training an ensemble of trees and taking the output to be a weighted average of the outputs of the individual trees

$$F(\mathbf{x}) = \sum_{n=1}^{N_{trees}} \alpha_n T_n(\mathbf{x}). \quad (5.28)$$

The most common algorithm for generating the weights is known as AdaBoost [107]. A formal description of the algorithm can be found in the references, but training proceeds in the following qualitative manner. After the first decision tree is trained, the weights $\{s_i\}$ and $\{b_i\}$ of the training events that are misclassified by the first tree are increased. Next, another tree is trained on the reweighed training sample. Additional trees are trained in this manner until the desired number of trees is reached. The weight α_n that is assigned to each tree is related to its total misclassification rate, so that trees with the lowest misclassification rate receive the highest weight in the sum. The power of this algorithm is that, for a wide class of problems, even if each tree has a relatively high misclassification rate, the ensemble can provide excellent separation between signal and background.

BDT Construction

Once weighted signal and background Monte Carlo samples are available, it is straightforward to train multivariate classifiers such as BDTs using the TMVA package within ROOT [108].⁶ Similar to the cut-based analysis, we use four input variables: charge energy (`qimean0F`), total phonon energy (`ptNF`), radial phonon partition (`prpart0F`), and z phonon partition (`pzpart0F`). Figure 5-32 shows an example of the input distributions to the BDT.

The signal and background distributions are split into two equally-sized “test” and “training” samples. The BDTs are trained using the training samples and the background estimates are made on the test samples. The split is made in order to mitigate against bias from *overtraining*. A classifier is said to be overtrained when it performs very well on the specific realization of signal and background to which it was trained, but fails to perform well on other realizations of the background. Even if the BDT is only moderately overtrained, background estimates could be biased too low if they are measured with the sample used to train the data.

To limit the amount of overtraining, the number of trees trained and the depth of each tree is adjusted. The misclassification rate of additional trees generally levels off, so there is little reason to add an arbitrary number of trees. Deeper trees are able to partition the space into a finer arrangement of hypercubes, providing better discrimination, at the cost of potential overtraining. The Kolmogorov-Smirnov (KS) test statistic between the BDT response of the test and training samples is used as a diagnostic of the degree of overtraining [109]. To achieve KS p -values in the range of 0.01 to 0.2, the number of trees needed for each detector is generally between 400 and 1000 and the maximum depth of each tree is between 3 and 5.

The trained trees can be evaluated on the test sample and give rise to distributions such as those in Figure 5-33, showing separation between signal and background that is much better than the separation in any one of the individual input variables in Figure 5-32. The discrimination is clearly not perfect, and much of this remaining degeneracy is a consequence of the fact that the signal and background distributions are truly degenerate in some regions of the input variable parameter space. One might wonder why only four variables are included in the BDT, given that BDTs can easily be trained on a much larger number of input variables without much computational overhead. While it is true that significantly better discrimination may be achievable with more input variables, it is unclear whether the background model would be able to accurately describe many of the other pulse shape analysis parameters, such as rise-time quantities, that can be affected by local saturation. The decision was therefore made to focus on a small set of variables that the background model was likely to model accurately.

Although we do not have calibration data for the crucial sidewall ^{210}Pb backgrounds, we can get a crude sense of the reliability of the background modeling by comparing the BDT output distribution of the post-Cf data to the BDT output of

⁶Julien Billard deserves tremendous credit here for actually performing the tedious training of the BDTs, cross-checking the output, introducing me to TMVA, and most of all for being the one to advocate using BDTs once it became clear that we had a suitably advanced background model.

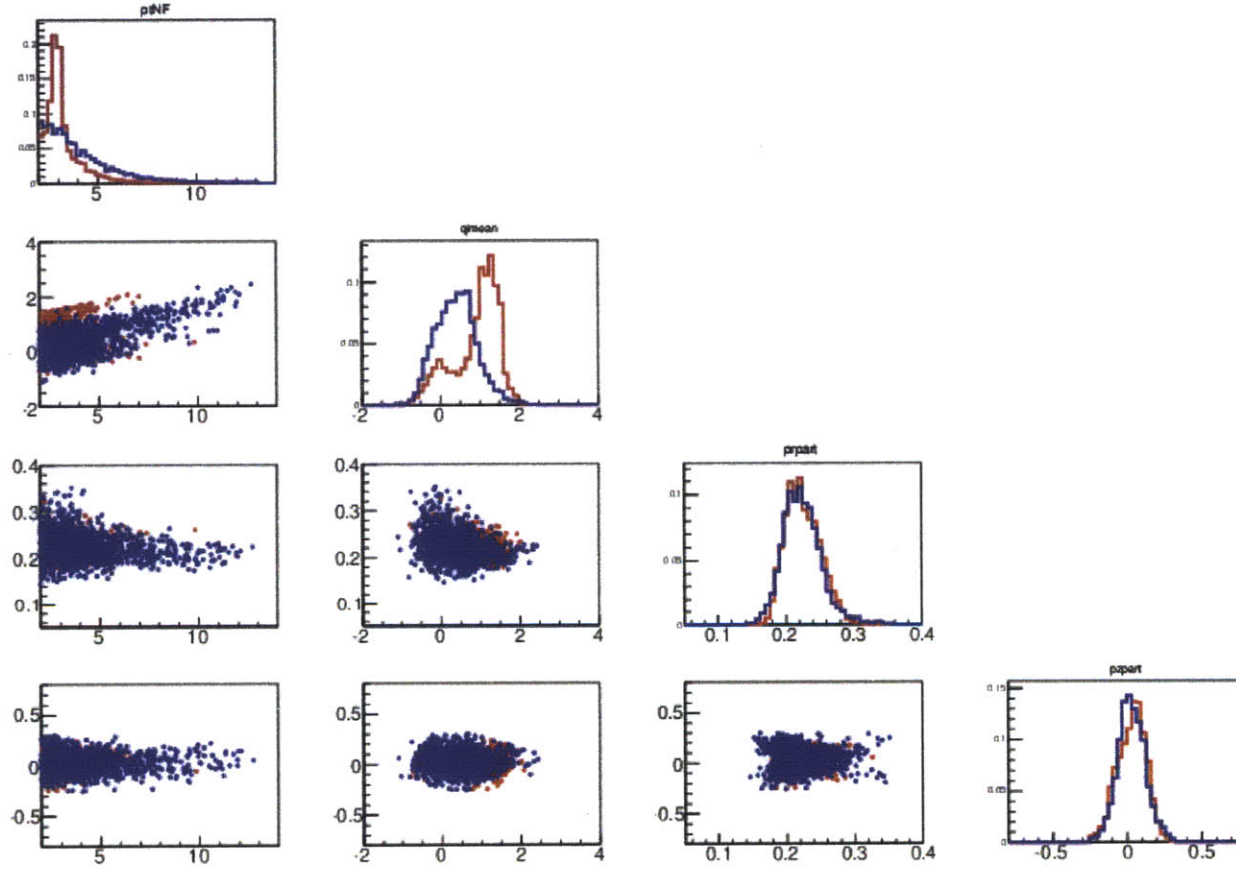


Figure 5-32: Input variables to the boosted decision tree for signal model (*blue*) and background model (*red*), on detector T2Z2. The signal model has been weighted to represent a 10 GeV WIMP. (Courtesy of Julien Billard)

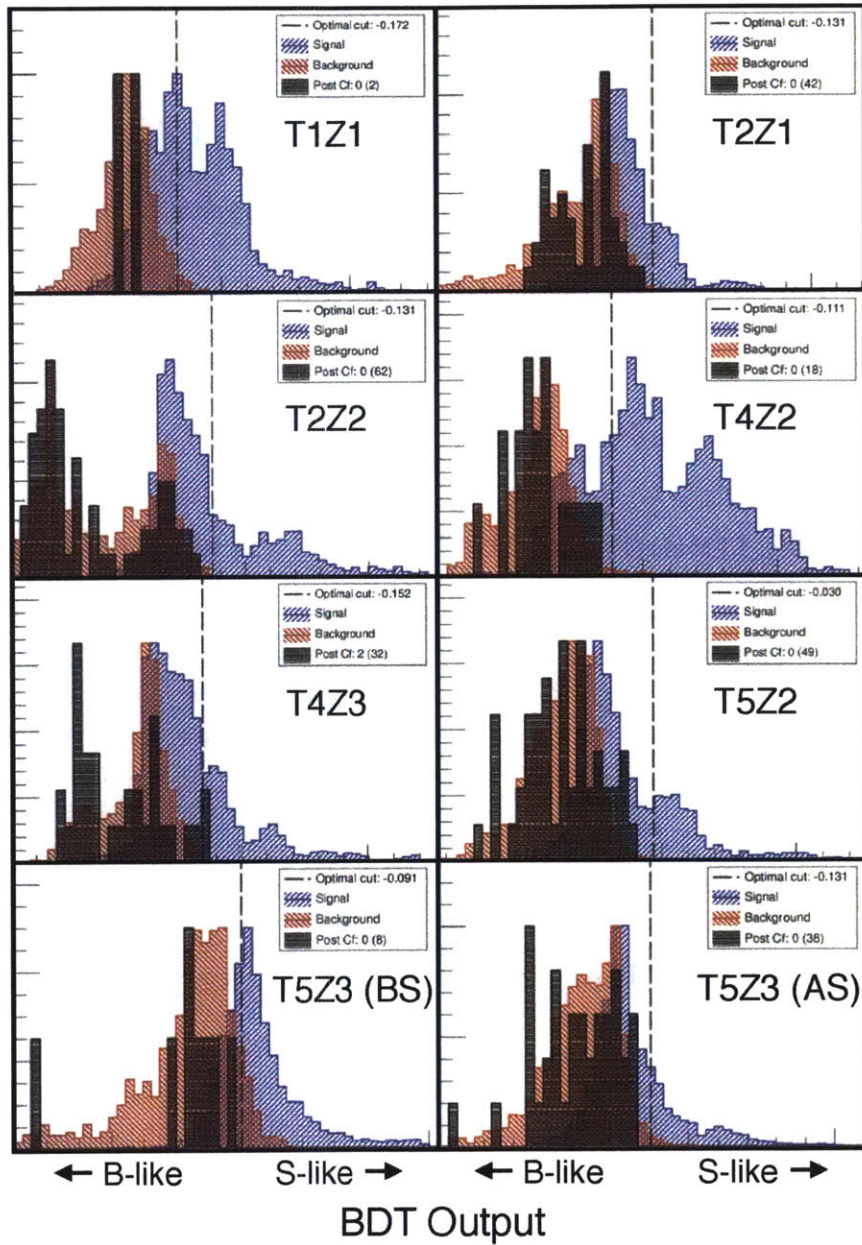


Figure 5-33: Pdfs of BDT output trained on a 10 GeV WIMP. Blue distributions show the signal model, red distributions show the background model, and black distributions show data from the open post-Cf period. Dashed line shows cut values (accepting data to the right) that optimize the expected 90% C.L. upper limit for each detector individually. (Courtesy of Julien Billard)

the background model. The cdfs of these distributions are shown in Figure 5-34 for the BDT optimized to a 10 GeV WIMP. While the agreement is not perfect, it is reasonable, given the statistical limitations of the post-Cf dataset. This agreement is a powerful test of the background model because it is addressing whether the data and background model agree in the regions of parameter space that have the most discrimination between signal and background.

We perform a simple cut-and-count analysis on the BDT output, so the BDT cut position must be optimized for each detector. This is done by minimizing the expected 90% C.L. upper limit, as in the cut-based approach. Unlike the cut-based approach, there is only 1 cut position for each detector, so it is computationally possible to adjust all cuts simultaneously to optimize the sensitivity. The sensitivity parameter as a function of cut position is shown in Figure 5-35. The sensitivity improves significantly as the cut position moves in the positive direction past the majority of the background. As the cut position continues to move in the positive direction, the sensitivity begins to slowly degrade because of a loss of acceptance.

5.10 Background Estimates and Systematics

After fixing the discrimination cuts that define the signal region, we perform estimate the background and its systematic uncertainty in this region. We discuss four main sources of systematic uncertainty in this section: the statistical uncertainty in the normalization of each of the background components, as discussed in Section 5.7.1, the systematic difference between the normalization methods for the ^{210}Pb background, the uncertainty in the low-energy yield of sidewall betas, the effect of finite Monte Carlo statistics. These systematic uncertainties are not always gaussian, so we combine them by Monte Carlo simulation to derive the final background estimate.

5.10.1 Statistical Uncertainty from Background Normalizations

In all cases, there are a sufficient number of events in the normalization region of each background that the gaussian approximation to Poisson statistics is valid. Table 5.11 summarizes the relative uncertainty from all backgrounds on all detectors, which are combined by Monte Carlo simulation. Uncertainties from L-shell lines, gammas, Qouter surface events, and Qinner surface events are mutually uncorrelated because their normalizations are computed on independent samples. The uncertainties on the subcomponents of the Qouter surface events and Qinner surface events have perfect correlation because each of these two groups of backgrounds are normalized using a single sideband component.

5.10.2 ^{210}Pb Normalization

It would be ideal to estimate the performance of these two methods on data from a calibration source, such as a detector housing implanted with a large amount of

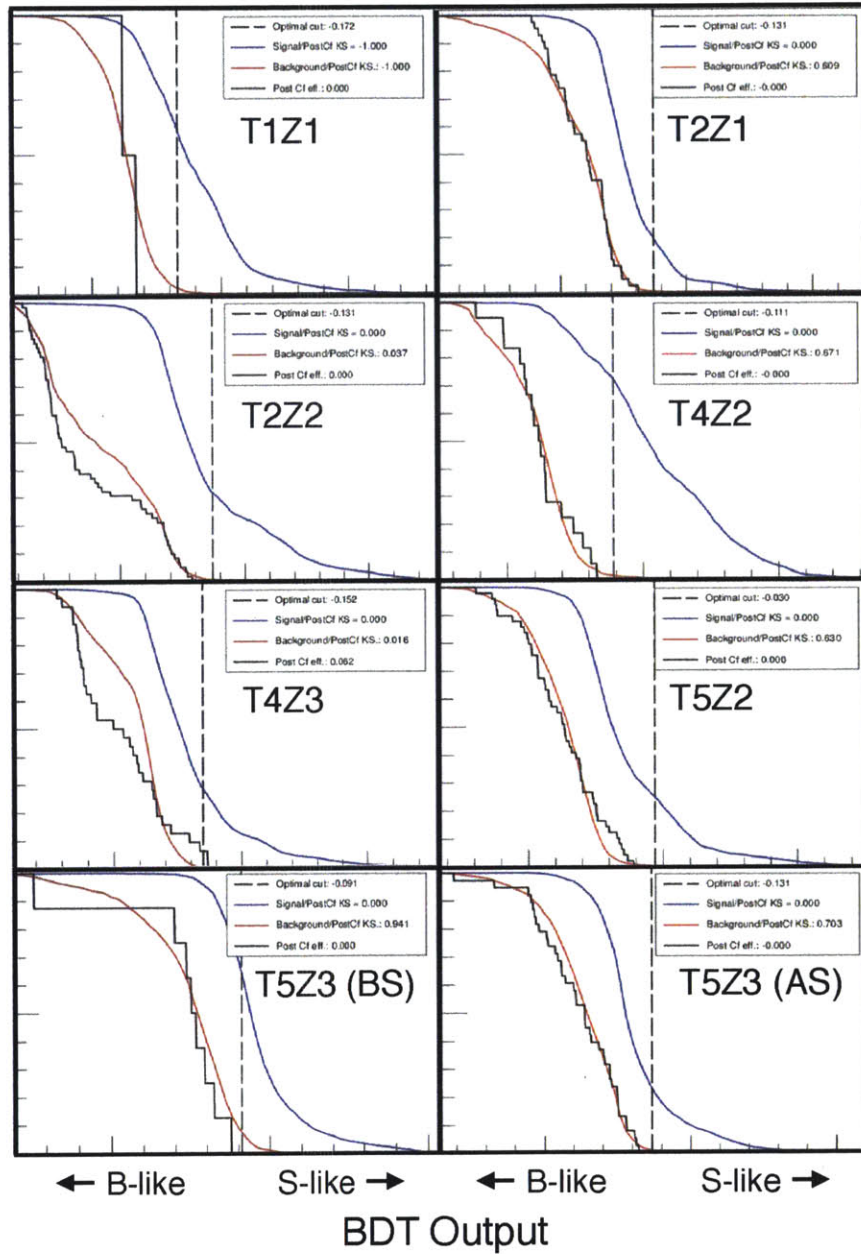


Figure 5-34: Cdfs of BDT output trained on a 10 GeV WIMP. Blue lines show the signal model, red lines show the background model, and black lines show data from the open post-Cf period. Dashed line shows cut values (accepting data to the right) that optimize the expected 90% C.L. upper limit for each detector individually. (Courtesy of Julien Billard)

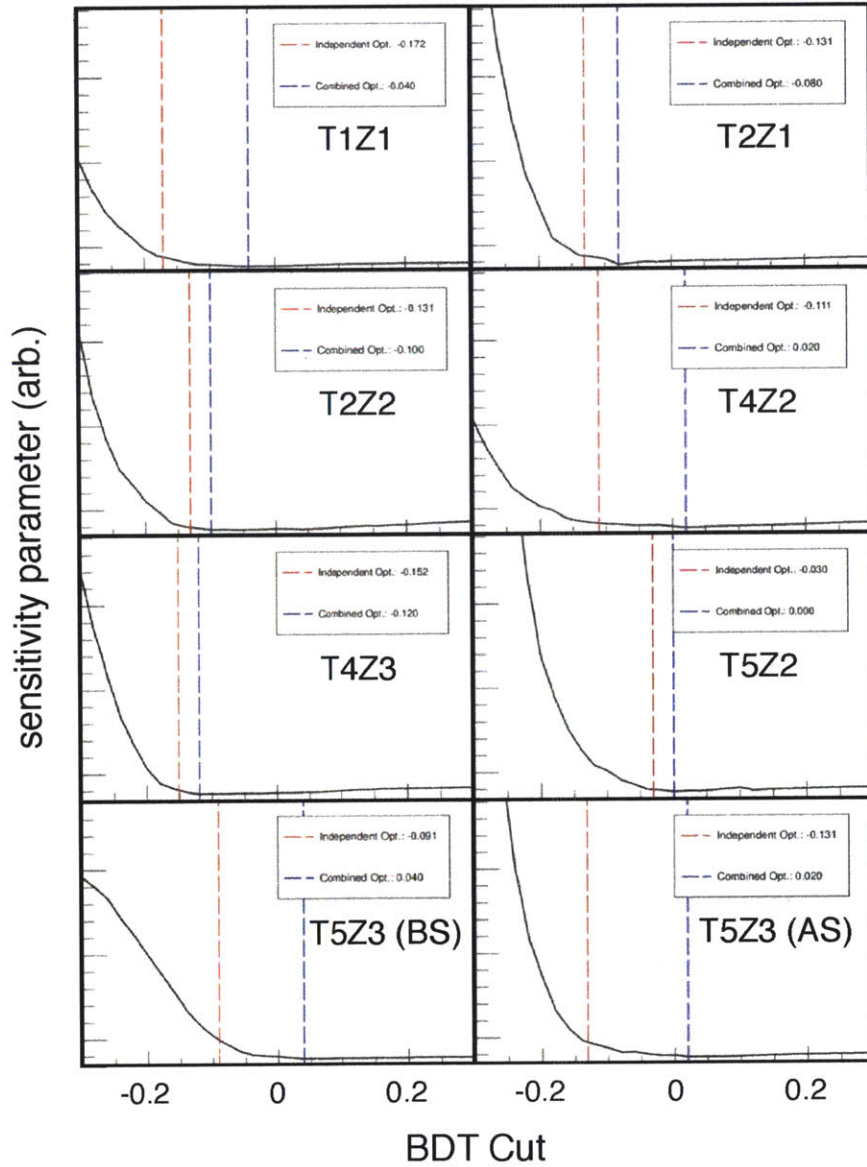


Figure 5-35: Sensitivity parameter function along axis of BDT cut position in each detector. Black lines show the sensitivity function, blue dashed lines show the optimal cut positions on each detector, while red dashed lines show the cut positions when the optimization is performed on each detector individually. The simultaneous cut tends to favor a harsher cut on all detectors, but the exact magnitude depends on the overall discrimination power of each detector. (Courtesy of Julien Billard)

Detector	L-shell	gamma	^{210}Pb Qout	^{210}Bi Qout	^{206}Pb Qout	^{210}Pb Qin	^{210}Bi Qin
T1Z1	33.6%	8.3%	4.5%	4.5%	4.5%	12.5%	12.5%
T2Z1	28.3%	8.7%	4.6%	4.6%	4.6%	13.2%	13.2%
T2Z2	19.3%	12.0%	8.8%	8.8%	8.8%	26.7%	26.7%
T4Z2	16.5%	13.6%	4.4%	4.4%	4.4%	22.9%	22.9%
T4Z3	27.9%	7.9%	4.6%	4.6%	4.6%	18.0%	18.0%
T5Z2	21.3%	8.9%	5.0%	5.0%	5.0%	33.3%	33.3%
T5Z3 (BS)	31.7%	17.7%	4.5%	4.5%	4.5%	7.9%	7.9%
T5Z3 (AS)	31.4%	8.1%	4.5%	4.5%	4.5%	7.9%	7.9%

Table 5.11: Statistical uncertainty of the normalization of each background component. Errors are taken to be gaussian.

^{210}Pb . Since such a source was not tested before Run 133, we have no alternative but to absorb the difference between the two background estimates as a systematic uncertainty. This systematic is treated in the following manner. The background model is first normalized using the low-yield counting method, for simplicity. The systematic uncertainty takes the form of a correction factor f to the low-yield counting method. The normalizations from the low-yield counting method and the alpha counting method each have an associated uncertainty from Poisson statistics. The pdf of the systematic uncertainty of f is taken to be a uniform distribution between the value for the alpha counting method and the value for the low-yield counting method. The statistical and systematic uncertainties are convolved by Monte Carlo simulation to generate the final pdf for the ^{210}Pb chain normalization uncertainty, as shown in Figure 5-36 and Table 5.12. This uncertainty applies equally to each of the three components of this background model.

Given the very large systematic uncertainty in the ^{210}Pb chain background normalization, it is natural to wonder whether any of the gamma or conversion electron lines or the ^{206}Pb endpoint could be used to provide a better calibration for the overall normalization. These possibilities were not explored for this first analysis of SuperCDMS data, but were explored in considerable detail in internal SuperCDMS notes not included in this thesis. The basic conclusion of these studies is that the high-energy data spectrum shows important differences from simulations, precluding its use as a means of normalizing the simulation.

5.10.3 Yield Shape Normalization

One of the key limitations of the detector pulse simulation is the assumption that the response of the detector is invariant in energy. Although this is a good approximation, variation with energy of the ionization yield of the backgrounds could produce very large changes in the background rate passing the charge fiducial volume. Increasing the ionization yield of surface events, for example, causes them to be much more likely to be rejected by the charge symmetry requirement. The ionization yield of surface




Figure 5-36: Pdf of scale factor for ^{210}Pb background, incorporating systematic and statistical uncertainties. Value on the horizontal axis is the factor by which the ^{210}Pb background normalization is scaled, relative to the nominal value estimated from the low-yield counting method. The orange dashed line shows the average of the alpha counting and low-yield counting methods, the red distribution shows the statistical uncertainty on the average of the two methods, the blue distribution shows the assumed distribution of the systematic error on the scale factor (uniform between the two measurements), and the green distribution is the combined statistical and systematic uncertainty.

Detector	scale factor rel. to low-yield method	range of scale factor systematic
T1Z1	0.950 ± 0.046	[0.900, 1]
T2Z1	0.818 ± 0.123	[0.635, 1]
T2Z2	1.57 ± 0.388	[1, 2.15]
T4Z2	0.954 ± 0.043	[0.907, 1]
T4Z3	0.861 ± 0.094	[0.723, 1]
T5Z2	0.818 ± 0.122	[0.637, 1]
T5Z3 (BS)	0.873 ± 0.090	[0.747, 1]
T5Z3 (AS)	0.843 ± 0.106	[0.685, 1]

Table 5.12: Scale factor incorporating the statistical and systematic uncertainty on ^{210}Pb background normalization. The numerical value of the scale factor is the value by which the normalization from the low-yield method (see text) must be multiplied to incorporate the systematically different measurement from the alpha method. Uncertainties in the scale factor are the $\pm 1\sigma$ quantiles from the mean of the pdfs shown in Figure 5-36. The range of the scale factor systematic shown in the third column shows the endpoints of the blue distribution in Figure 5-36, which is the difference between the alpha method and the low-yield method.

events is known to decrease as the penetration depth into the detector decreases. Since low-energy surface events tend to penetrate less deeply, we expect a correlation between event energy and ionization yield.

Of all the backgrounds we consider, uncertainty in the ionization yield only affects betas that interact in the detector sidewall. Bulk electron recoils are known from Ba calibration data and the L-shell activation lines to have a yield of 1, even at very low energies. Templates for the ^{206}Pb sidewall events typically have a yield around 0.1. In the region of interest of 2-13.1 keV total phonon energy, this is low enough that ^{206}Pb sidewall events have negligible ionization energy. Betas from the decay of ^{210}Pb and ^{210}Bi , however, stop in the outer few microns of the detector and are subject to some ionization losses. Data from surface event calibration sources on the top and bottom surfaces suggest that ionization yield losses are small on the detector faces. But there is no calibration data for sidewall betas, so we cannot actually measure this drop in yield as a function of energy. On the other hand, we do know that the low-energy beta ionization yield can neither be lower than the ionization yield assumed in the pulse simulation for ^{206}Pb sidewall nuclei, nor can it be greater than the ionization yield for high-energy betas, as we have nominally assumed in the pulse simulation. We use these two extremes to bound the systematic uncertainty from the sidewall beta yield.

To parameterize the yield uncertainty, we first compute the background under both the low-yield and high-yield hypotheses for the sidewall betas. The low-yield hypothesis produces a lower background estimate than the high-yield hypothesis because the additional charge energy under the high-yield hypothesis increases the probability of being rejected by radial charge cuts `cQin1_v53_LT` and `cQin2_v53_LT`. The

Detector	difference in background between low-yield and high-yield
T1Z1	0.022
T2Z1	0.390
T2Z2	0.832
T4Z2	0.006
T4Z3	1.052
T5Z2	0.600
T5Z3 (BS)	0.028
T5Z3 (AS)	0.103

Table 5.13: Difference in background in BDT signal region between low-yield and high-yield hypotheses for the sidewall betas.

difference between these two background hypotheses is shown in Table 5.12. There are large differences across detectors, which is primarily driven by differences in threshold: detectors with lower thresholds tend to have higher sidewall event background because they are sensitive to energies with lower discrimination. We chose to parameterize the systematic uncertainty between these two background estimates with a half-gaussian distribution. The peak of the half gaussian is at the high-yield background estimate, and the 2σ point of the half gaussian is at the low-yield background estimate. While this parameterization is somewhat arbitrary, it captures following features of the background:

- it cannot be smaller than it is in the high-yield hypothesis
- it is unlikely to be much larger than in the low-yield hypothesis
- it is probably closer to the high-yield hypothesis than the low-yield one

In addition, we assume that this systematic uncertainty is *perfectly correlated on all detectors*. This is justified because the yield decrease is a function of electric field and charge physics at the detector surface, which should be fairly consistent across detectors.

5.10.4 Finite Monte Carlo Statistics

If the total number of Monte Carlo events passing the final event selection is small, then there is another systematic uncertainty due to the finite number of samples used to estimate the background. This uncertainty encodes the fact that repeated Monte Carlo simulations of the same size could produce different results because of statistical fluctuations. Although this type of systematic does not seem to be typically incorporated into physics analyses, we estimate it for completeness and because the weights for different Monte Carlo events used to estimate the background can vary significantly: even if the background is estimated with a large number of Monte Carlo




Figure 5-37: Half gaussian distribution used to model systematic uncertainty due to yield of sidewall betas. The “high-yield” hypothesis corresponds to the background expected if the yield of low-energy sidewall betas is the same as at high energy, while the “low-yield” hypothesis corresponds to the background expected if the yield is the same as ^{206}Pb sidewall events. The functional form of the systematic is arbitrary, but captures the essential features of our belief about the background. The systematic is assumed to be perfectly correlated on all detectors in the sense that if the background is 1σ above the high-yield hypothesis on one detector, it is 1σ above on all detectors.

events, the bulk of the background may be set by a smaller number of events with higher weights.

There is no analytic way of calculating the standard error of the background estimate based on the Monte Carlo events that lie in the selection region. One option is to use the *bootstrap*, which provides completely nonparametric estimates of standard errors that are comparable to maximum likelihood errors in situations where the latter can be computed [105]. In the bootstrap method, if we simulate N Monte Carlo events, then we resample N events from our Monte Carlo with replacement, and use the resampled data to compute a new background estimate (because we are resampling with replacement, some events will be duplicated). We repeat this procedure many times to build up a distribution of the background estimate. The bootstrap method does not appear to be in widespread use in high-energy physics, although Ref. [110] performs detailed studies of applications similar to ours.

5.10.5 Combined Estimates

Table 5.14 shows the background estimates after combining the systematic uncertainties described above by Monte Carlo simulation. Figure 5-38 also shows the PDF of the background when coadded across all detectors. This is the estimate most relevant for the analysis because we compute limits on the coadded events and exposure.




Figure 5-38: Distribution of total background, coadded across all detector, in signal region due to statistical and systematic uncertainties. The blue distribution shows the statistical uncertainty from the normalization samples, green includes the systematic uncertainty from the two methods for normalizing the ^{210}Pb chain, purple includes the systematic uncertainty from finite Monte Carlo statistics, and orange includes the systematic uncertainty of the sidewall beta yield. The median total background coadded across all detectors is $6.1 - 0.9 + 1.2$ events (median, 16% and 84% quantiles).

Detector	Total background
T1Z1	$0.03 - 0.01 + 0.02$
T2Z1	$1.39 - 0.22 + 0.24$
T2Z2	$1.77 - 0.31 + 0.36$
T4Z2	$0.04 - 0.02 + 0.02$
T4Z3	$1.65 - 0.30 + 0.40$
T5Z2	$1.08 - 0.26 + 0.31$
T5Z3 (BS)	$0.04 - 0.01 + 0.01$
T5Z3 (AS)	$0.09 - 0.04 + 0.05$
Coadded	$6.07 - 0.87 + 1.17$

Table 5.14: Background by detector and coadded across detectors. Values are the median and distance to 16% and 84% quantiles.

5.11 Selection Criteria Efficiencies

The efficiency of the BDT-based discrimination cut and nuclear recoil consistency requirements is estimated based on the fraction of nuclear recoils in ^{252}Cf calibration data that pass these requirements. The final analysis efficiency is the product of these efficiencies with the quality cut efficiency and trigger requirement efficiency, which are assumed to be uncorrelated. The raw passage rates of the ^{252}Cf data must be corrected for several systematic differences between WIMPs and the calibration data, described below.

- **Gamma contamination:** ^{252}Cf neutron sources produce gammas over a wide range of energies from 0-10 MeV [111]. At low energies, the number of gammas is much smaller than the neutrons, but their presence must still be accounted for. Since the electron recoil band contains a very pure sample of electron recoils, a template distribution constructed from ^{133}Ba data is normalized in the electron recoil band and used to estimate and subtract the contamination from gammas.
- **Multiple scattering of neutrons:** Unlike WIMPs, neutrons have a reasonable probability of scattering multiple times inside a single detector. If a WIMP scatters once in the inner part of the detector and once in the outer part, it will be rejected by the strict charge fiducial volume requirements. This *geometrical* effect causes the neutron acceptance to be lower than the WIMP acceptance. This phenomenon is also convoluted with a *resolution* effect: with poor resolution, we sometimes will not resolve the second scatter in the outer charge channel, so resolution acts to “damp” the bias of the geometrical effect. The method for correcting the neutron acceptance is quite complicated and was worked out by Elías Lopez Asamar.

The geometrical effect is estimated from Geant4 simulations of the ^{252}Cf calibration source and the detector array. Events removed by the fiducial volume

cuts are assumed to occur in a sub-cylinder of the detector,⁷ whose size is varied until the acceptance of neutrons in the simulation matches the data. The resolution effect is a correction to the geometrical effect, which is estimated using the same pulse simulation of sidewall events employed in the background modeling.

Figure 5-39 shows both the corrected and uncorrected efficiencies of the discrimination and WIMP consistency requirements by detector. The magnitude of the correction is typically about 10% or so, though it varies between different detectors and energies. The multiple-scattering correction procedure introduces both statistical and systematic uncertainties into the nuclear recoil acceptance. These are propagated by Monte Carlo simulation, together with the statistical uncertainties in the trigger efficiency described in Section . Because a large number of uncorrelated uncertainties are averaged together when detectors are coadded, the final uncertainty of the total acceptance is very small. Figure 5-40 shows the final acceptance as a function of total phonon energy for each stage of event selection.

⁷Strictly speaking, this is incorrect because oblique propagation of electrons causes the volume removed by the charge fiducial volume to be non-cylindrical near the electron collection side. Quantifying this effect truly requires the physics-based DMC simulation, a topic that was explored in some interesting early studies of the iZIP detectors by Scott Hertel (e.g. §12.2 of [50]).

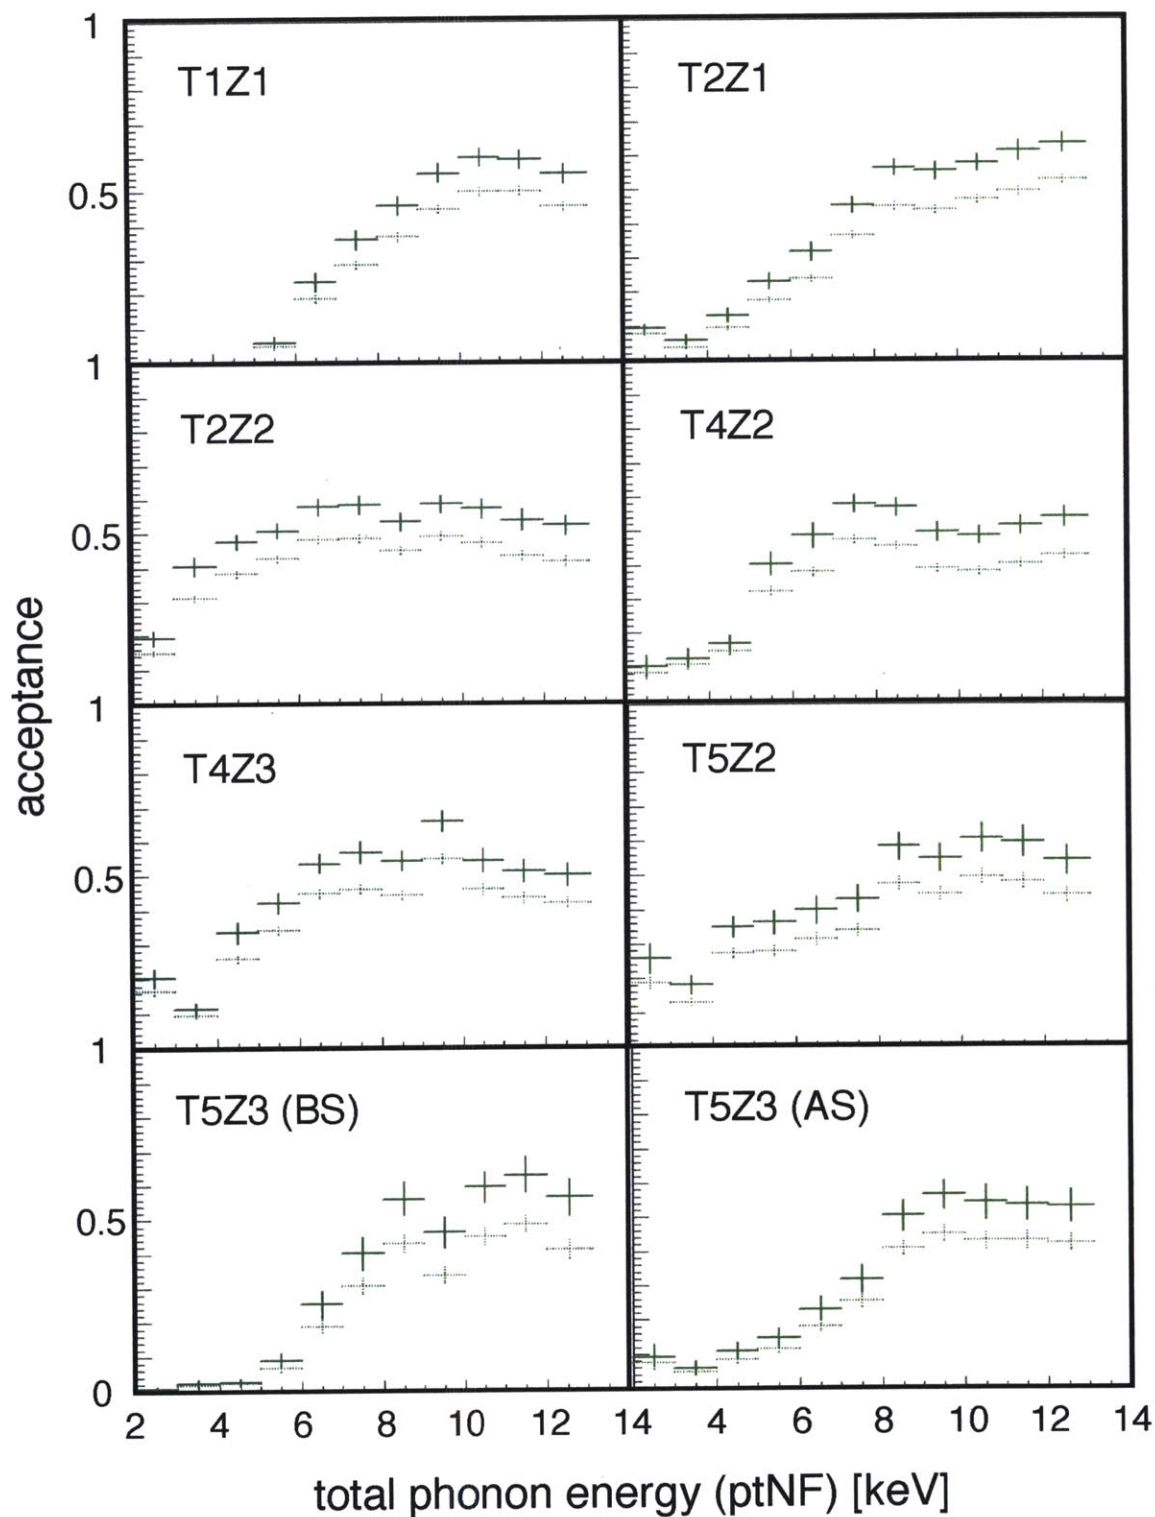


Figure 5-39: Efficiency as a function of total phonon energy of BDT discrimination cuts and WIMP consistency cuts. Dotted error bars are the raw acceptance of ^{252}Cf calibration data with statistical errors. Solid error bars are the final acceptance, corrected for both multiple scattering of neutrons within a single detector and gammas that are present in calibration data. (Courtesy of Elías Lopez Asamar)




Figure 5-40: Cumulative efficiencies after sequential application of each stage of event selection. From top to bottom, these are data-quality criteria, trigger and analysis thresholds, preselection criteria, and BDT discrimination with 68% C.L. (stat. + syst.) uncertainty band. The preselection and BDT selection efficiencies are interpolated from measurements in 1 keV bins, though this choice of interpolation was confirmed to introduce a negligible change in the WIMP exclusion. Steps are due to time-dependent analysis thresholds for individual detectors. For illustrative purposes, the Lindhard nuclear-recoil energy is provided as an approximate nuclear-recoil energy scale.

Chapter 6

Constraints on Light WIMPs

Chapter 5 describes the data, event selection criteria, and analysis structure used to search for WIMP-induced nuclear recoils. In this chapter, we apply the analysis to the data and use the results to set constraints on the WIMP interactions, focusing on the low-mass 1-20 GeV/ c^2 region where experimental anomalies have arisen. In the following sections we unblind the data (Section 6.1), assess the compatibility with the background model (Section 6.2), and derive constraints on the spin-independent WIMP-nucleon cross section (Section 6.3), annual modulation (Section 6.4), and study the effect of astrophysical uncertainties (Section 6.5).

6.1 Candidate Events

The unblinded WIMP-search data are shown in the plane of total phonon energy and charge energy in Figure 6-1, summed across detectors. Eleven events pass the final WIMP event selection criteria, compared to the total background expectation of $6.1 - 0.9 + 1.2$, with energies listed individually in Table 6.1. The number of candidate events is somewhat above the background expectation, although the effect is not particularly significant when Poisson fluctuations are accounted for. Just before applying the BDT discrimination cut, there are 1218 events passing the WIMP consistency requirements, compared to 1183 expected from the background model, so the overall event rate is at an expected level. The BDT alone also achieves an overall rejection of better than 1:100, which is impressive given the low energies of this analysis. The grey points in Figure 6-1 shows the basic structure expected from WIMP-search background: a strong L-shell activation line near 1.3 keVee, a zero-charge band from sidewall events rising exponentially with lower energy, and an electron recoil band at high ionization yield.

6.2 Comparison with Background Expectations

Candidate events appear to generally be high-quality physics events. The distribution of events is fairly uniform in time. One of the events on T2Z2 has a slightly unusual pulse shape that suggests that it is a glitch event, but the energy is so low (2.09 keV)




Figure 6-1: Small gray dots are all veto-anticoincident single-scatter events within the ionization-partition fiducial volume that pass the data-quality selection criteria. Large encircled shapes are the 11 candidate events. Overlapping shaded regions (from light to dark) are the 95% confidence contours expected for 5, 7, 10 and 15 GeV/ c^2 WIMPs, after application of all selection criteria. The three highest-energy events occur on detector T5Z3, which has a shorted ionization guard. The band of events above the expected signal contours corresponds to bulk electron recoils, including the 1.3 keV activation line at a total phonon energy of ~ 3 keV. High-radius events near the detector sidewalls form the wide band of events with near-zero ionization energy. For illustrative purposes, the Lindhard nuclear-recoil energy is provided as an approximate nuclear-recoil energy scale.

Detector	Candidate energies [keV _{nr}]	Expected background	Average 50% threshold [keV _{nr}]
T1Z1	—	$0.03^{+0.01}_{-0.01}$	4.6
T2Z1	1.7, 1.8	$1.4^{+0.2}_{-0.2}$	1.5
T2Z2	1.9, 2.7	$1.8^{+0.4}_{-0.3}$	1.8
T4Z2	—	$0.04^{+0.02}_{-0.02}$	4.7
T4Z3	—	$1.7^{+0.4}_{-0.3}$	1.7
T5Z2	1.9, 2.3, 3.0, 5.8	$1.1^{+0.3}_{-0.3}$	2.0
T5Z3	7.0, 7.8, 9.4	$0.13^{+0.06}_{-0.04}$	1.7

Table 6.1: Energies of candidate events in each detector, labeled by tower (first number) and position within tower from top to bottom (second number). Expected background is based on the model used to train the BDT and includes the estimated systematic uncertainty. Differences in expected background across detectors reflect different trigger thresholds and background event rates. Event energies are calculated using the measured mean ionization energy for nuclear recoils.

that we cannot make conclusive judgements. With the exception of detector T5Z3, the data is in fair agreement with the expectations of the background model. When the events in the signal region on each detector are coadded together, the p -value for the background-only rate is 0.07 when systematic uncertainties on the background model are accounted for. When the data are split by detector, the agreement remains reasonable, except on T5Z3. Figure 6-2 shows the observed background on each detector, compared with the predictions of the model. The observation is obviously in excellent agreement on the T1, T2, and T4 detectors. While the T5Z2 rate is somewhat higher than expected ($1.1^{+0.3}_{-0.3}$ versus 4 observed), the p -value of 0.033 for this observation is not unreasonable low. This value is much higher when we account for the fact that there is more than one detector in the analysis, also known as the *trials factor*. For an analysis with M detectors the probability of observing a p -value at least as extreme as p_0 on one or more detectors is $p' = 1 - (1 - p_0)^M$. The p -value for T5Z2 is 0.18 when corrected for the trials factor of the 6 detectors besides T5Z3, on which we expect the background model to be reliable.

The situation on T5Z3 is considerably more complicated than the other detectors. The probability of observing 3 events when 0.13 are expected is negligible, even when the trials factor is included, and all of the background is expected below 5 keV total phonon energy. The background model is clearly not correct on this detector. The reason for this discrepancy is not necessarily understood in a fully quantitative manner, but it is strongly believed to be due to the shorted outer ionization electrode. Figure 6-3 shows a cross-sectional view of the electric field inside a detector with a shorted electric field. The important point to notice is that electric field lines go from the outer radial wall of the detector to the inner charge electrodes on side 1 which has the shorted channel. Events which occur on the sidewall can then produce a signal in the inner charge channel, thereby faking an inner event. One might object that these




Figure 6-2: Observed event rates by detector with $\pm 1\sigma$ Poisson error bars (black), compared with mean background model rate (red band). The thickness of the model prediction is the $\pm 1\sigma$ interval of the systematic uncertainty on the background model.

events should be removed by the charge symmetry requirement `cQisym_v53_LT` since the inner charge is only produced on one side, but at low energies noise fluctuations could potentially cause these events to pass the event selection requirements.

Figure 6-4 shows direct evidence for this interpretation. Unlike fully functional detectors, T5Z3 has a large population of events with both large *S1* inner charge energy and large *S2* outer charge energy. Since events with nonzero *S2* outer charge energy are primarily sidewall events, the existence of this population shows that sidewall events produce an abnormally large signal in the inner charge channels of a detector like T5Z3 with a shorted charge electrode as we would anticipate from the electric field configuration.

The pulse simulation used in the background model accounts for the fact that there is no information from shorted detector channels, but it produces incorrect predictions if the events used as templates are not representative of the true background. This likely occurred because one of the selection criteria for the sidewall event templates is that sidewall events have no detectable charge energy on an inner electrode (cuts `cQout1_v53_LT` and `cQout2_v53_LT`, described in Section 5.7.2). This requirement clearly makes sense for a fully functional detector, but it produces a bias in T5Z3 if sidewall events can produce a significant signal on the inner charge electrodes.

Finally, it is important to emphasize that the failure of the background model to account for the electric field of T5Z3 does *not* affect the validity of our limit. It was determined before unblinding that limits would be computed without background subtraction, conservatively treating all events as WIMPs. While excess background

Electric Field & Potential for $Q_{in} = \pm 2 \text{ V}$ and $Q_{out} = 2 / 0$




Figure 6-3: Cross-sectional view of simulation of axisymmetric E-field model with charge channel QOS1 shorted to ground, as in the SuperCDMS experiment. Red lines show the direction of the electric field, while blue lines show the equipotential contours. Note that in the bottom right-hand corner, near channel QOS1, electric field lines from the outer radial surface of the detector connect with the inner charge channel. This provides the mechanism by which sidewall events can produce a signal on the inner charge electrode and appear as good bulk events. (Courtesy of Blas Cabrera)




Figure 6-4: Comparison of S2 outer charge energy versus S2 inner charge energy between a fully functional detector (T2Z1, left panel) and T5Z3 with its shorted S1 outer electrode (right panel). Red points fail the charge fiducial volume, while blue points pass. There are two noticeable differences between the two distributions. First, the S2 inner charge energy has a bimodal distribution near zero, which is an unimportant artifact of the optimal filter algorithm on detectors with shorted channels. Second, unlike fully functional detectors, T5Z3 has a large population of events with both large S2 outer charge energy *and* large S1 inner charge energy.

on one detector complicates the interpretation of the consistency with the background-only hypothesis, its only effect is to slightly degrade the limit relative to the expected sensitivity. The uncertainty in the detector response to these sidewall backgrounds was, after all, the reason that this search was constructed as an exclusion analysis.

The overall agreement of the background model with the data, summed across detectors is still reasonable, as shown in Figure 6-5 for the $10 \text{ GeV}/c^2$ BDT. Summing over detectors somewhat smooths out systematic errors that exist on individual detectors. Once the systematic uncertainty on the background model is included, the p -values for agreement between the background model and data in the preselection region range between 0.08 and 0.26 for the BDTs trained to each WIMP mass.

6.3 Constraints on Spin-Independent Dark Matter

We set limits on spin-independent WIMP interactions using the optimum interval method, conservatively treating all candidate events as potential WIMPs [97]. The optimum interval method exploits some information about the shape of the observed events to set a limit that is better than a simple Poisson limit, but does so without making any assumption about the shape of the background. The method sets the limit using events only in the energy interval which produces the strongest limit, while including the appropriate statistical penalty for the freedom to make this choice. All detectors are coadded together, treated as a single monolithic detector, as assumed during the discrimination cut optimization. Variants of the optimum interval method

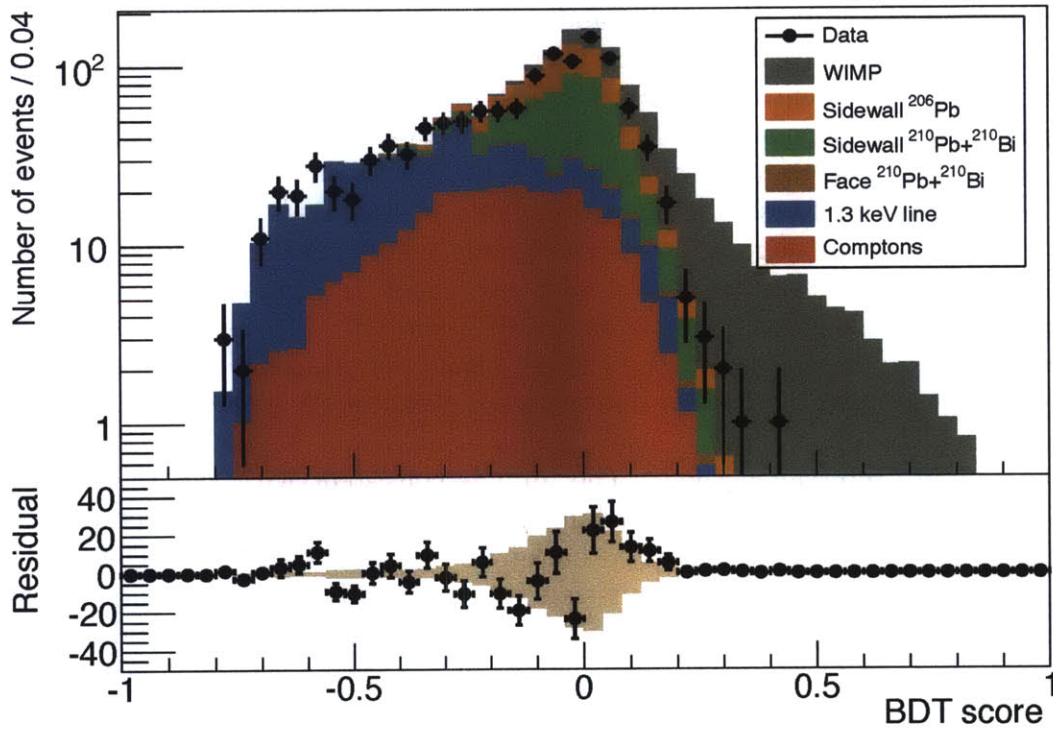


Figure 6-5: Top: Stacked histogram showing the components of the background model passing the preselection criteria, summed over all detectors (neutron backgrounds are negligible and not included). For comparison, a $10 \text{ GeV}/c^2$ WIMP with cross section $6 \times 10^{-42} \text{ cm}^2$ is shown on top of the total background. Events passing preselection criteria are overlaid (markers with statistical errors). A p-value statistic comparing the data to background model is 14% for this selection. Bottom: Difference between the data and the background expectation. Tan bars indicate the systematic uncertainty (68% C.L.) on the background estimate. Each component of the background model was computed prior to unblinding and was not fit or rescaled to match the data.

exist which only use data from a single detector with the most constraining interval [112], as was used in [80]. In addition to being simpler to understand and implement, the coadding method produces a stronger limit when the background is fairly low and similar between detectors, since restricting to a single detector reduces the total exposure. To avoid statistical bias (e.g. flip-flopping [113]), the decision to set an upper limit and the decision to use the optimum interval method were made prior to unblinding.

For each WIMP mass, an upper limit is set on the number signal events using the optimum interval method. We make standard assumptions about the WIMP halo and particle physics, as described in detail in Section 2.2.1. Following [114], we use a galactic dark matter density of $\rho = 0.3 \text{ GeV}/c^2 \text{ cm}^{-3}$, a Maxwellian DM speed distribution with a dispersion of 220 km s^{-1} , and an escape velocity of 544 km s^{-1} [44]. We take the WIMP to have equal coupling to protons and neutrons $f^p/f^n = 1$, so that the cross section for scattering on the entire nucleus scales as A^2 .

The solid black line in Figure 6-6 shows the upper limit on the spin-independent WIMP-nucleon cross section, compared with exclusions and putative signals from other experiments. The green bands show the 68% and 95% C.L. range of expected sensitivity, computed before unblinding. These bands were constructed by sampling pseudo-experiments from the background model of events passing the BDT analysis, then calculating the upper limit of each pseudo-experiment at each WIMP mass using the optimum interval method. The limit is consistent with the 1σ band of the sensitivity below about $9 \text{ GeV}/c^2$, but worsens relative to the expectation at higher masses because of the unexpected background at high energies in T5Z3. The narrow grey band surrounding the limit shows the 95% C.L. range of limits due to systematic uncertainties in the event selection efficiency, trigger efficiency, and nuclear recoil energy scale. It is relatively small, except at the lowest energies, largely due to the fact that the data is coadded over detectors. Since most of the systematic uncertainties are uncorrelated by detector, the addition tends to average out the fluctuations on individual detectors.

Under the standard assumptions, all low-mass anomalies are strongly excluded by this result. Although varying the halo parameters and the isospin coupling of the WIMP may relax some tension, the tension with the CoGeNT experiment cannot be alleviated. Indeed if a WIMP existed with the CoGeNT best-fit mass and cross-section, ~ 200 events would have been observed in the signal region of this analysis. For standard WIMP physics, a similar number of events would have been produced by the CDMS II Si best-fit parameters. Barring a catastrophic error in our analysis, a WIMP interpretation of CoGeNT is severely excluded. As we will show in Section 6.5, WIMP interpretations of CDMS II Si are constrained independent of astrophysical assumptions.

6.4 Constraints on Annual Modulation

In recent years, experiments such as CoGeNT [120] and DAMA/LIBRA [84] have reported an annual modulation in their event rates, which could be interpreted as




Figure 6-6: The 90% confidence upper limit (solid black) based on all observed events is shown with 95% C.L. systematic uncertainty band (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. The disagreement between the limit and sensitivity at high WIMP mass is due to the events in T5Z3, whose occurrence weakens the limit in this range. Closed contours shown are CDMS II Si [82] (*dotted blue*, 90% C.L.), CoGeNT [83] (*yellow*, 90% C.L.), CRESST-II [85] (*dashed pink*, 95% C.L.), and DAMA/LIBRA [115] (*dash-dotted tan*, 90% C.L.). 90% C.L. exclusion limits shown are CDMS II Ge [116] (*dotted dark red*), CDMS II Ge low-threshold [80] (*dashed-dotted red*), CDMSlite [117] (*solid dark red*), LUX [48] (*solid green*), XENON10 S2-only [86, 118] (*dashed dark green*), and EDELWEISS low-threshold [119] (*dashed orange*).

dark matter. These interpretations are obviously in strong tension with the null results from measurements of overall background rates in experiments such as LUX and SuperCDMS, as well as a dedicated annual modulation search using CDMS II data [50, 121]. It is nevertheless still interesting to investigate annual modulation of this SuperCDMS dataset. Although we only analyze 9 months of data, a preliminary annual modulation analysis will allow us to study potential issues that would be relevant to a future search with more data.

In this section, we perform two alternative analyses. The first (hereafter called “maximum modulation” method) uses the 11 events that pass the final BDT event selection of the primary WIMP search, and then simply sets Poisson upper limits on the modulation, conservatively assuming that 1.) all observed events are WIMPs, and 2.) that there is 100% modulation of the signal rate. This trivial approach will provide the most conservative possible upper limit on the modulation signal. The second analysis is a more conventional unbinned maximum likelihood analysis using a more inclusive dataset (hereafter called “maximum likelihood” method). Although the maximum modulation method relies on extremely crude assumptions, unlike the maximum likelihood method, it exploits the excellent background discrimination provided by the BDT analysis. On the other hand, the maximum modulation method analyzes an event sample with higher WIMP acceptance, so could potentially compensate for the lower purity of its event sample.

6.4.1 Maximum Modulation Method

In this section, we make the overly conservative assumption that all of the events in the BDT signal region are due to a modulating component of the overall event rate. This assumption will produce the most conservative (worst) limits on the modulation rate, but allow us to take advantage of the excellent discrimination offered by the BDT analysis. For simplicity, we coadd all detectors before setting limits and we set the limits on the modulation in individual recoil energy bins. The reason for binning in energy is to aid the comparison with results from CDMS II, DAMA, and CoGeNT, which estimate their modulation in bins of recoil energy. The generic functional form for the rate as a function of time in the j th energy bin is given by

$$R^j(t) = [R_{mod}^j(\cos(\omega t + \phi^j) + 1) + R_{DC}^j] \times \left(\sum_{k=1}^{N_{det}} M \epsilon^{j,k}(t) f^k(t) \right), \quad (6.1)$$

R_{mod}^j = modulating amplitude for energy bin j (evt kg⁻¹ d⁻¹ keVnr⁻¹)

ω = frequency of modulation, fixed to 1 year

ϕ_j = phase of modulation

R_{DC}^j = constant “DC” rate

M = detector mass

$\epsilon^{j,k}(t)$ = preselection cut efficiency as a function of time for energy bin j and detector k

$f^k(t)$ = live time fraction as a function of time for detector k

Since the live time and efficiency information cannot be measured on timescales smaller than an individual data series, it is convenient to actually approximate the continuous-time expression by a functional form that is binned by data series (typically 3 hours). Doing this, and setting $R_{DC}^j = 0$ for maximum modulation, we find that the number of events in expected in series t_i in energy bin j is

$$N^j(t_i) = R_{mod}^j [\cos(\omega t_i + \phi^j) + 1] \times \left(\sum_{k=1}^{N_{det}} M \epsilon^{j,k}(t_i) L_i^k \right), \quad (6.2)$$

where L_i^k is the livetime for detector k and series t_i . If we observe a total of N_j events in the data for energy bin j coadded over all series, then we can set a Poisson upper limit of N_{90}^j on the rate of events in the energy bin. Summing equation (6.2) over all series and substituting N_{90}^j allows us to solve easily for the limit on the maximum modulation amplitude

$$R_{mod,90}^j = \frac{N_{90}^j}{\sum_{i=1}^{N_{series}} \sum_{k=1}^{N_{det}} [\cos(\omega t_i + \phi^j) + 1] \times M \epsilon^{j,k}(t_i) L_i^k}. \quad (6.3)$$

Upper limits on the number of events are shown in Figure 6-10, while the upper limits on the modulation rate are shown in Figure 6-10.

6.4.2 Unbinned Maximum Likelihood

The standard method searching for annual modulation involves performing fits in the time domain or constructing a periodogram (see Ref. [122]). We set limits in this section by performing maximum likelihood fits, unbinned in time, but binned in energy as above. Doing a separate fit at each energy instead of a full 2D fit in energy and time also means that we do not need to specialize to a particular DM model or WIMP mass. The number of events in data series t_i is given by

$$N^j(t_i) = [R_{mod}^j (\cos(\omega t_i + \phi^j) + 1) + R_{DC}^j] \times \left(\sum_{k=1}^{N_{det}} M \epsilon^{j,k}(t_i) L_i^k \right), \quad (6.4)$$

where L_i^k is the lifetime in time bin i for detector k .

This analysis is analogous to fitting an energy spectrum with a signal component




Figure 6-7: Upper limits on observed counts used in maximum modulation method for annual modulation constraints. Black points are 1σ error bars, while red points are 90% C.L. upper limits. Note that some of the observed events from the BDT analysis lie below the 2 keVnr threshold imposed for this study.

and background component, except that now our signal and background components are functions of time instead of functions of energy. The signal and background PDFs are given by a simple modulating component and a flat background, respectively

$$p_{mod}(t_i) = \frac{1}{\mathcal{N}_{mod}(\phi)} [\cos(\omega t_i + \phi^j) + 1] \times \left(\sum_{k=1}^{N_{det}} M \epsilon^{j,k}(t_i) L_i^k \right) \quad (6.5)$$

$$p_{DC}(t_i) = \frac{1}{\mathcal{N}_{DC}} \times \left(\sum_{k=1}^{N_{det}} M \epsilon^{j,k}(t_i) L_i^k \right), \quad (6.6)$$

where $\mathcal{N}_{mod}(\phi)$ and \mathcal{N}_{DC} are normalization constants. The normalization constant for the modulating component depends on the phase because the total analysis time is not a multiple of the 1-year modulation period. The combined PDF for the data is then given by

$$f(t_i) = \frac{\mu_{mod}(R_{mod})p_{mod}(t_i) + \mu_{DC}(R_{DC})p_{DC}(t_i)}{(\mu_{mod}(R_{mod}) + \mu_{DC}(R_{DC}))}, \quad (6.7)$$

where μ_{mod} is the number of events in the modulating component of the data, and μ_{DC} is the number of events in the constant component. Both of these terms are functions of the underlying modulating and constant event rates R_{mod} and R_{DC} .

We use the method of extended maximum likelihood to constrain the rate. Using

the total PDF of equation (6.7), the extended likelihood function for the data is given by

$$\begin{aligned} \mathcal{L}(\phi, R_{mod}, R_{DC} | \{t_i\}, N) &= \frac{(\mu_{mod} + \mu_{DC})^N}{N!} e^{-(\mu_{mod} + \mu_{DC})} \\ &\times \prod_{i=1}^N \frac{\mu_{mod} p_{mod}(t_i) + \mu_{DC} p_{DC}(t_i)}{(\mu_{mod} + \mu_{DC})}. \end{aligned} \quad (6.8)$$

As is customary, we minimize twice the negative logarithm of the likelihood, discarding constant factors that do not affect the shape of the likelihood

$$\begin{aligned} L(\phi, R_{mod}, R_{DC}) &= -2 \log \mathcal{L} \\ &= -2N \log(\mu_{mod} + \mu_{DC}) + 2(\mu_{mod} + \mu_{DC}) \\ &\quad - 2 \sum_{i=1}^N \log \left(\frac{\mu_{mod} p_{mod}(t_i) + \mu_{DC} p_{DC}(t_i)}{(\mu_{mod} + \mu_{DC})} \right) \\ &= 2(\mu_{mod} + \mu_{DC}) - 2 \sum_{i=1}^N \log (\mu_{mod} p_{mod}(t_i) + \mu_{DC} p_{DC}(t_i)). \end{aligned} \quad (6.9)$$

The parameter R_{DC} is a nuisance parameter, while ϕ and R_{mod} are the parameters of interest. We construct confidence intervals using the profile likelihood test statistic

$$q(\phi, R_{mod}) = -2 \log \lambda(\phi, R_{mod}) \quad (6.10)$$

where

$$\lambda(\phi, R_{mod}) = \frac{\max_{R_{DC}} \mathcal{L}(\phi, R_{mod}, R_{DC})}{\mathcal{L}(\hat{\phi}, \hat{R}_{mod}, \hat{R}_{DC})}, \quad (6.11)$$

and quantities with hats are the ML estimators. In the profile construction, the test statistic q is a function of the parameters of interest ϕ and R_{mod} . We maximize the likelihood over the nuisance parameter at each point (ϕ, R_{mod}) to calculate the profile likelihood function and then the profile likelihood test statistic. The test statistic q is asymptotically distributed according to a χ^2 distribution with n degrees of freedom, where n is the number of parameters of interest.

The fits in each energy bin are shown in Figure 6-8, while limits in the plane of modulation rate and phase are shown in Figure 6-9. Some mild modulation is preferred in a few energy bins, but not systematically across all energies. This mild preference could be due to unquantified systematics or perhaps simply statistical fluctuations. Limits from the maximum likelihood method are compared with limits from the maximum modulation method in Figure 6-10. The maximum modulation method performs significantly better than the maximum likelihood fits, and both methods provide stringent constraints that exclude the low-energy modulation fraction in the CoGeNT experiment. This may seem an academic exercise because the BDT analysis of the overall rate clearly excludes a WIMP interpretation of CoGeNT. If one were skeptical of the efficiency estimation of the BDT analysis, however, the

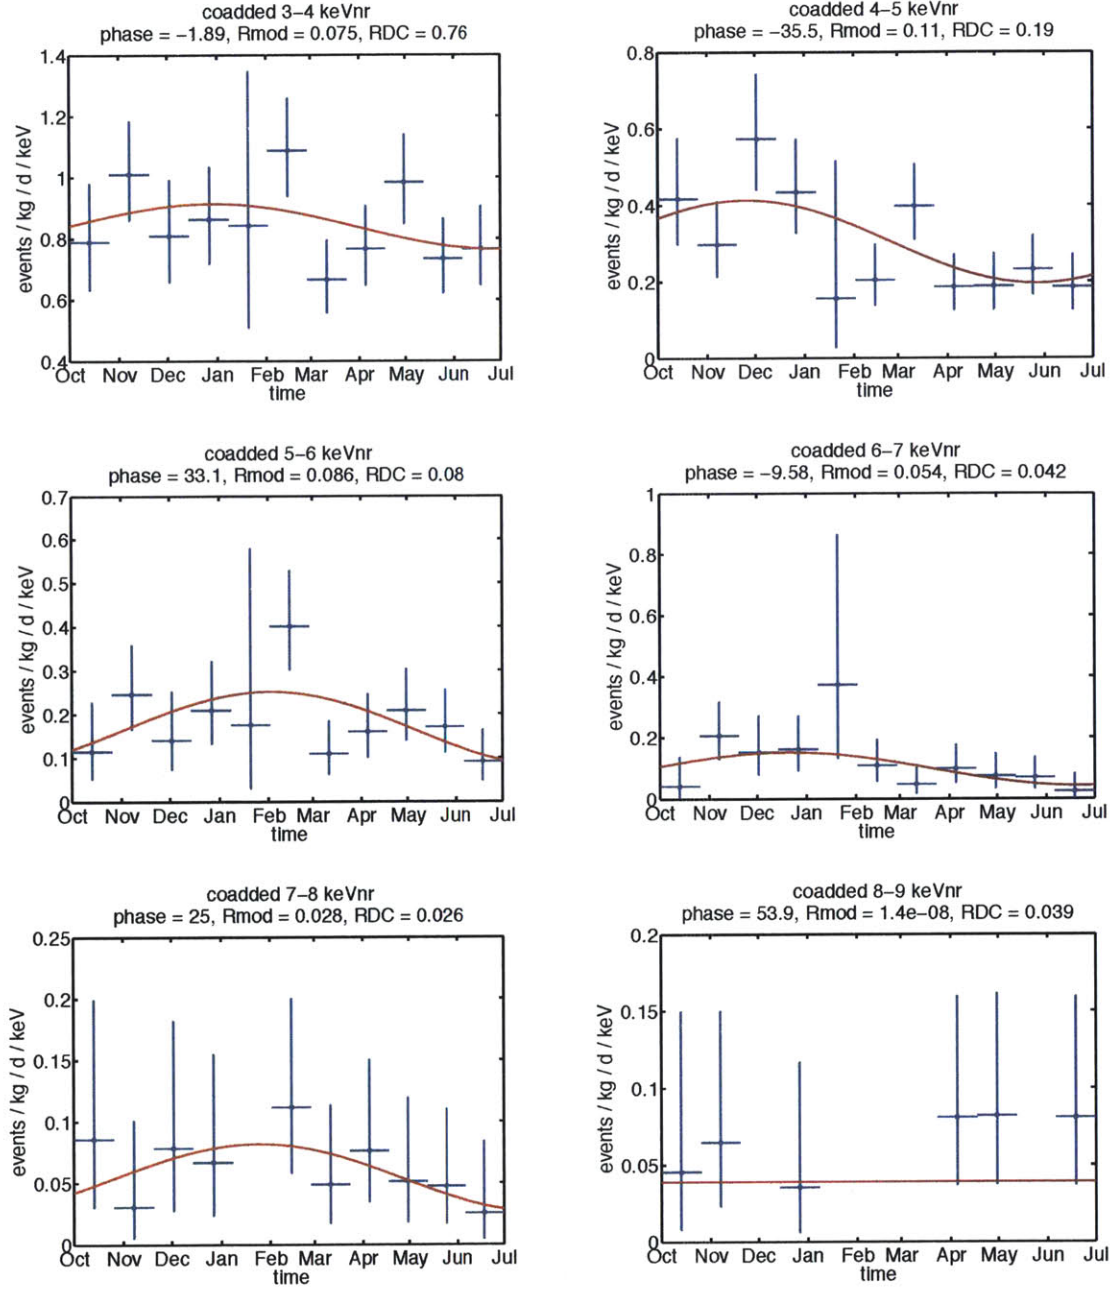


Figure 6-8: Maximum likelihood fits for an annually modulating event rate to coadded WIMP-search data in 1 keVnr energy bins. Note that fits are not performed in the 2-3 keVnr energy bin because of the presence of the L-shell activation lines, which are not included in the model, and which have an exponential time-dependence not described by the model.

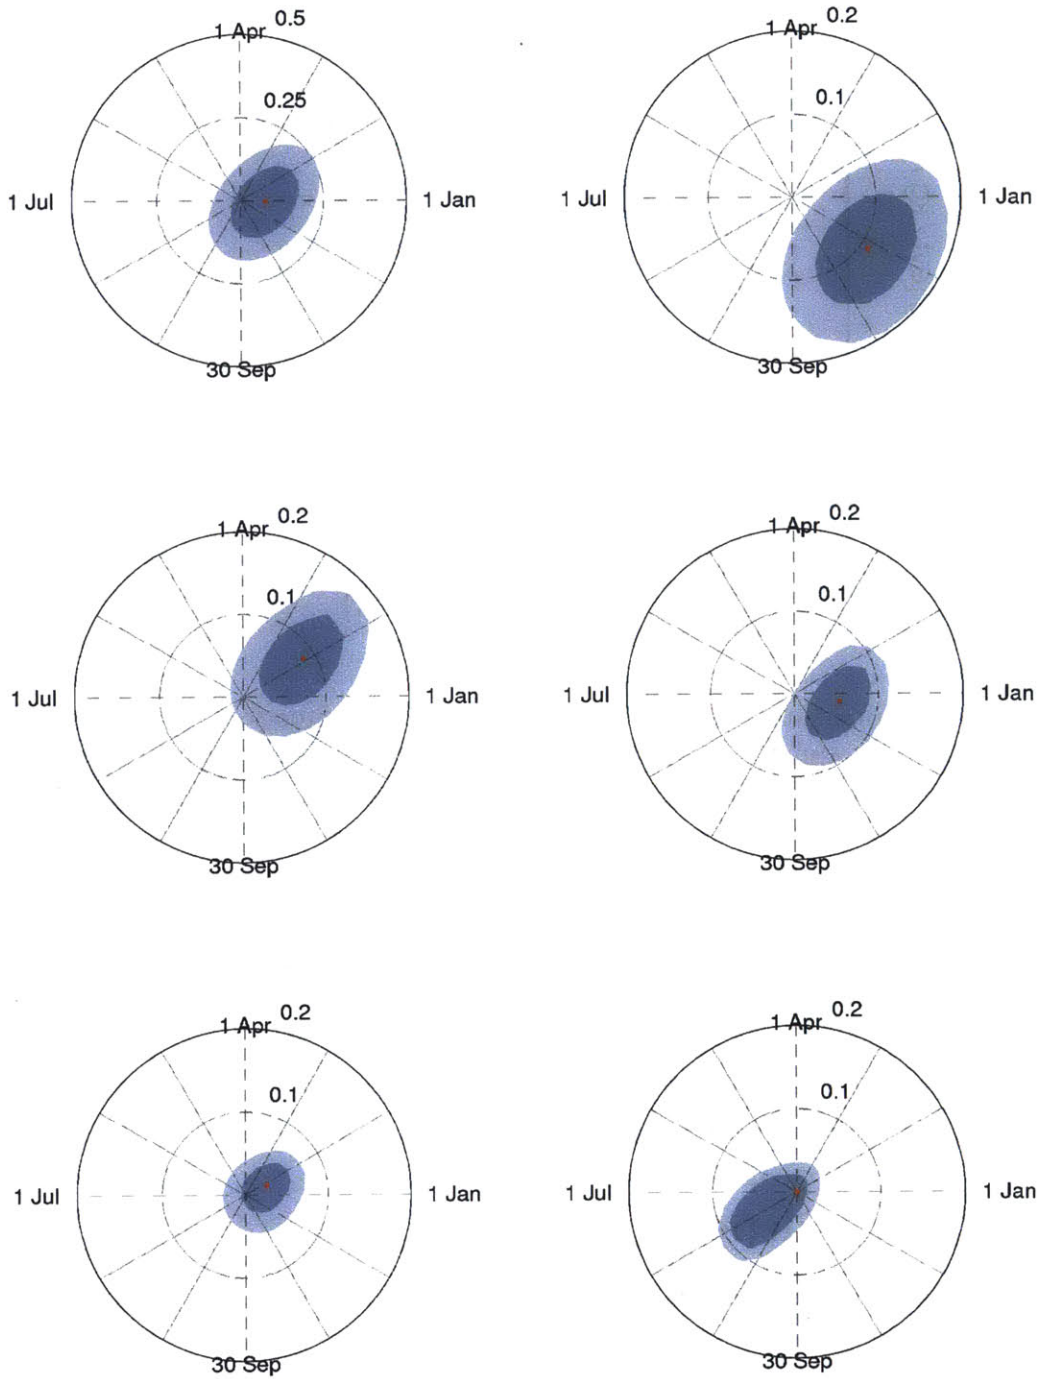


Figure 6-9: Maximum likelihood confidence regions in the plane of modulation amplitude and phase from fits to WIMP-search data in 1 keVnr energy bins. Polar plots show modulation amplitude as the radial coordinate and cosine phase as the azimuthal coordinate. Note that fits are not performed in the 2-3 keVnr energy bin because of the presence of the L-shell activation lines, which are not included in the model, and which have an exponential time-dependence not described by the model.




Figure 6-10: Upper limits on the modulating rate as a function of recoil energy from the maximum modulation method (*blue*) and ML method (*black*), compared with modulation estimated in CDMS II and CoGeNT data from [121]. CDMS II constraints are strong at high energies, but the better discrimination at lower energies in available in SuperCDMS enables stronger constraints at low energies, which strongly exclude CoGeNT.

maximum modulation annual modulation analysis shows that CDMS is still in tension with CoGeNT. The maximum likelihood method analyzes a sample of data that includes all nuclear recoils that pass quality and fiducial volume requirements that have fairly high efficiency ($\sim 50\%$).

6.5 Halo-Independent Constraints

As discussed in Section 2.2.1, there is significant uncertainty in the structure of the DM halo, which may significantly affect both overall limits on the WIMP-nucleon cross section as well as the relative position of limits between experiments using different targets. A method was originally proposed in [122] to perform comparisons between separate experiments *independent of assumptions about the DM halo*. The basic scheme of this method has since been generalized to handle detector resolution [123] and arbitrary forms of DM interactions [124], as well as serve as a way of probing the halo structure itself in both binned [125] and unbinned approaches [126], and to statistically quantify disagreement between experiments with different targets [127]. These methods typically involve adding degrees of freedom to the analysis, which

invariably eases tension between experimental results, though incompatibilities can still be found in this more general framework. In this section, we briefly outline one of the more recent methods [126] for extracting halo-independent bounds, as a particularly clear and elegant example of these methods. We then discuss how SuperCDMS data can exclude the signal hint from the Si detectors the CDMS II experiment in a fully halo-independent context over a wide range of WIMP masses, as discussed in the forthcoming paper [128].

6.5.1 Unbinned Halo-Independent Limits

We briefly review here the unbinned technique for halo-independent analyses of direct detection experiments described in [126]. Recall that for spin-independent DM-nuclear scattering, the differential event rate at a direct detection experiment is given by

$$\frac{dR}{dE_R} = \frac{N_A \rho_\chi \sigma_n m_n}{2m_\chi \mu_{n\chi}^2} C_T^2(A, Z) \int dE'_R G(E_R, E'_R) \epsilon(E'_R) F^2(E'_R) g(v_{min}(E'_R)), \quad (6.12)$$

where m_χ is the DM mass, m_n the nucleon mass, $\mu_{n\chi}$ the nucleon-DM reduced mass, σ_n the DM-nucleon scattering cross-section, ρ_χ the local DM density, N_A is Avogadro's number, $F(E_R)$ is the nuclear form factor which accounts for loss of coherence as the DM resolves sub-nuclear distance scales, $C_T(A, Z) = (f_p/f_n Z + (A - Z))$ is the usual coherent DM-nucleus coupling factor, $\epsilon(E_R)$ is the detector efficiency, and $G(E_R, E'_R)$ is the detector resolution function.

The integral over the velocity distribution (hereafter called the “halo integral”) is

$$g(v_{min}) = \int_{v_{min}}^{\infty} \frac{f(\mathbf{v} + \mathbf{v}_E)}{v} d^3v, \quad (6.13)$$

where $f(\mathbf{v})$ is the unknown DM velocity distribution and \mathbf{v}_E is the Earth's velocity, both in the Galactic frame. For the sake of simplicity, we ignore the small time dependence of the Earth's velocity in the Galactic frame.¹ The lower limit of the halo integral, v_{min} , is the recoil energy-dependent minimum DM velocity required to produce a nuclear recoil E_R , and depends on the kinematics of the interaction. For elastic scattering $v_{min}(E_R)$ is given by

$$v_{min}(E_R) = \sqrt{\frac{m_N E_R}{2\mu_{N\chi}^2}}, \quad (6.14)$$

from basic elastic scattering kinematics.

As first pointed out in [129, 130], since the DM velocity distribution $f(\mathbf{v})$ is positive semi-definite, $g(v_{min})$ must be a monotonically decreasing function of v_{min} . The idea

¹Note that we also integrate the halo to infinity, rather than truncating at the galactic escape velocity. This is done explicitly to avoid any assumption about the halo structure. If there is a finite escape velocity, however, this will be reflected by a reconstructed value of $g(v_{min}) \rightarrow 1$ at finite v_{min} .

of halo-independent techniques is to *float* the values of $g(v_{min})$ in either a binned or unbinned fit to derive constraints on the halo itself. In practice, this is typically done by multiplying the halo integral by factors common to all experiments that set the overall scale of the rate, such as

$$\tilde{g}(v_{min}) = \frac{\rho_\chi \sigma_n}{m_\chi} g(v_{min}). \quad (6.15)$$

In a typical DM analysis in the variables of interest (m, σ) , a signal produces closed contours around single point, and an exclusion produces a band above which points in parameter space are excluded at some confidence level (e.g. Figure 6-6). In a halo-independent analysis, the “parameter space” is no longer a fixed set of variables, but rather an infinite-dimensional set of *functions* $\{\tilde{g}(v_{min})\}$. A signal produces a subset of these functions that are *allowed* at some confidence interval, while an exclusion is a subset of functions that are *not allowed* at some confidence interval.

The impracticality of working in an infinite-dimensional parameter space can be overcome by either binning in v_{min} , or using the trivial but elegant observation of [126] that the unbinned likelihood function only depends on the energies of the events actually observed. For example, an extended likelihood function for an experiment with background and N_O observed events is

$$\mathcal{L} = \frac{e^{-N_E}}{N_O!} \prod_{i=1}^{N_O} \left(\frac{dR_{bg}}{dE_R} + \frac{dR}{dE_R} \right) \Big|_{E_R=E_i}, \quad (6.16)$$

where N_E is the total number of expected events

$$N_E = \int_{E_{min}}^{E_{max}} \left(\frac{dR_{bg}}{dE_R} + \frac{dR}{dE_R} \right) dE_R. \quad (6.17)$$

Clearly \mathcal{L} only depends on the values of the halo integral $\tilde{g}_i = \tilde{g}(v_{min}(E_i))$ at the observed events. The only effect of the particular values of $\tilde{g}(v_{min}(E))$ at $E \neq E_i$ is to change the total number of expected events N_E . But the likelihood is maximized when N_E is minimized. When we combine this statement with the requirement that $\tilde{g}(v_{min})$ be monotone decreasing, we conclude immediately that the maximum likelihood estimator of $\tilde{g}(v_{min})$ must be a decreasing series of N_O step functions.² If the value of $\tilde{g}(v_{min})$ were increased in between the steps, the product in equation (6.16) would remain unchanged, but N_E would increase, reducing the likelihood. We can therefore write the likelihood function as

$$\mathcal{L} \equiv \mathcal{L}(\{\tilde{g}_i\}, \{v_{min,i}\} | \{E_{R,i}\}). \quad (6.18)$$

Compatibility between different experiments can be assessed quantitatively via the joint likelihood techniques advocated in [127], or a more qualitative graphical method for displaying exclusions and signals.

²It is nonobvious that this statement holds in the case of finite resolution, a fact that is proved rigorously in [126].

- **Signals:** When experimental data is fit with a signal component, one obtains a best-fit halo function corresponding to the step heights and positions of \tilde{g} that maximize the likelihood. One can then vary the step heights and positions to generate new halo functions whose likelihood can be evaluated using equation (6.16). Monte Carlo simulation suggests that the distribution of the likelihood ratio test statistic

$$q(\lambda(\{\tilde{g}_i\}, \{v_{min,i}\}) = -2 \log(\lambda(\{\tilde{g}_i\}, \{v_{min,i}\})) \quad (6.19)$$

with

$$\lambda(\{\tilde{g}_i\}, \{v_{min,i}\}) = \frac{\mathcal{L}(\{\tilde{g}_i\}, \{v_{min,i}\})}{\mathcal{L}(\{\hat{g}_i\}, \{\hat{v}_{min,i}\})} \quad (6.20)$$

is approximately distributed as χ^2 with $2N - 1$ degrees of freedom: $2N$ arises from the position and height of each of the steps, while 1 dof is subtracted for the monotonicity constraint. Depending on the distribution of the observed events, steps may “merge” together, with degenerate heights or positions. The number of “effective” degrees of freedom may be further reduced if bins merge together.

The heights and positions of the halo steps can be varied, for example by Markov Chain Monte Carlo, the likelihood can be evaluated for each halo, and this halo is then accepted or rejected if its test statistic q is less than the critical value for confidence level $1 - \alpha$. One way of visualizing the halo functions is then to draw the envelope of accepted halos. Any halo with a point outside the envelope must therefore be excluded at $1 - \alpha$.

- **Exclusions:** To derive halo-independent exclusion curves, note that for a given point v_0 on the v_{min} -axis, the halo integral which gives the least number of events consistent with monotonicity is

$$\tilde{g}(v_{min}) = \tilde{g}_0 \Theta(v_0 - v_{min}) \quad (6.21)$$

Setting an upper bound on the number of events using this choice of halo integral, one derives the weakest (i.e. most conservative) possible bounds on $\tilde{g}(v_0)$. Thus, by sweeping over v_0 in the v_{min} range of an experiment, one can build up an exclusion contour which at every x -coordinate v_0 sets the most conservative bounds on the halo integral: the true halo integral must lie entirely below this exclusion contour for every v_{min} . Such an exclusion curve is extremely robust, as any positive signal so excluded can be said to be excluded for *all* DM halos.

When carrying out a halo-independent analysis, as mentioned above, a particular DM mass must be assumed. The constraints on, or predictions for, $g(v)$ for another DM mass can be determined by simple rescalings [126]. For a DM mass, m_χ , a point (v_{min}, \tilde{g}) is mapped to a new point (v'_{min}, \tilde{g}') for DM mass $m_{\chi'}$, by

$$(v_{min}, \tilde{g}; m_\chi) \rightarrow \left(\frac{\mu_{N\chi}}{\mu_{N\chi'}} v_{min}, \frac{\mu_{n\chi'}^2}{\mu_{n\chi}^2} \tilde{g}; m_{\chi'} \right) \quad (6.22)$$

Since the mapping depends on the mass of the nuclear target, this will shift different detectors by differing amounts.

6.5.2 Comparing XENON10, LUX, and CDMS II Si with SuperCDMS

To demonstrate the methods above for comparing positive signals with exclusions, we consider the case of the SuperCDMS, LUX, and XENON10 limits and the preferred halo for the Si detectors of the CDMS II experiment.

The SuperCDMS collaboration has publicly released data from the analysis of 577 kg d of exposure on their seven Ge detectors with the lowest energy thresholds [117, 131], described in this chapter and the previous one. Although the observed events are consistent with the background estimate, we did not perform a background subtraction before calculating the limit, conservatively assuming that all observed events are potential DM scattering. Using the efficiency, ionization yield, and energies of the 11 events from the data release, we set 90% C.L. limits using the “Pmax” method [97].

The CDMS II Si analysis observed three events in 140.2 kg d of exposure. We follow the analysis approach used in [126], taking a gaussian detector resolution of 0.2 keV, the acceptance from [82], and backgrounds from [132]. We furthermore assume that the distribution of ΔL is χ^2 corresponding to five degrees of freedom: one for each step position and height in the halo, minus one for the monotonicity constraint. The χ^2 assumption is a reasonable approximation to the sampling distribution of ΔL [126], a more precise determination could be made by a Monte Carlo simulation, varying the number of events in each pseudoexperiment. A value of $\Delta L = 9.2$ then corresponds to 90% C.L. We perform Markov Chain Monte Carlo (MCMC) sampling of the 6-dimensional parameter space, and use halos with $\Delta L < 9.2$ to calculate the envelope preferred by the CDMS II Si result.

For comparison, we also compute limits for LUX and XENON10 at a fixed dark matter mass. Similar to [126], we use events and ionization yield \mathcal{Q}_y from the XENON10 S2-only analysis [86]. We conservatively assume that the ionization yield drops to zero below 1.4 keVnr because the signal acceptance cannot be reliably estimated below this energy. The energy resolution is taken to be $\Delta E_R = E_R / \sqrt{E_R \mathcal{Q}_y(E_R)}$, while the acceptance is 95% on an exposure of 15 kg d. Yellin’s “Pmax” method [97] is used to set limits at 90% C.L., using the conservative assumption that all observed events are potential DM scatterings. Following [126], we deduce Poisson upper limits from the results of the LUX experiment [48]. The background distributions for the LUX experiment have not been made public, so we cannot compute limits based on the profile likelihood ratio test statistic as done by the LUX collaboration. Because no events were observed in LUX in the low-energy nuclear recoil band, we instead use the 90% upper limit for an observation of 0 events for a Poisson process with no background, taking the efficiency from [48]. This simplification gives reasonable agreement with the limits reported by LUX in the low-mass region.³

³See [133, 134] for more details on reconstructing the published LUX results.




Figure 6-11: Exclusion limits from XENON10 (dot-dashed green), LUX (solid purple), SuperCDMS (dotted orange) overlaid on the preferred region (90% C.L.) for CDMS II Si (shaded blue) and best fit halo (dashed blue) in v_{min} -space. **Left:** v_{min} -space plot for 10 GeV DM. **Right:** range of SuperCDMS limits scaled with respect to 10 GeV DM using the factor appropriate for Si (shaded orange) (analogous to Fig. 6-12, right), overlaid with the preferred region for CDMS II Si (shaded blue).

In the left panel of Fig. 6-11 we compare limits from XENON10, LUX, and SuperCDMS, for 10 GeV DM in the usual v_{min} -space plot. As pointed out in [126], the limits in v_{min} space exhibit a simple scaling with the DM mass,

$$v'_{min}(E_R) = \frac{\mu_{N\chi}}{\mu_{N\chi'}} v_{min}(E_R). \quad (6.23)$$

This can be visualized in a single plot, such as the right panel of Fig. 6-11, which shows the range of limits for SuperCDMS, when referred to 10 GeV DM using the scaling factor appropriate for Ge. Such a plot will contain, for every possible mass, the correct positioning of the CDMS-Si preferred region relative to SuperCDMS. For sufficiently large DM masses, the exclusion limit from SuperCDMS drops below the lower boundary of the CDMS II Si contour, indicating strong tension between the inferred dark matter halos of the two experiments. The mass at which the SuperCDMS exclusion drops below the lower limit of the CDMS II Si preferred region, is the mass above which SuperCDMS excludes CDMS II Si. In our analysis, SuperCDMS excludes CDMS II Si for $m_\chi \gtrsim 11$ GeV.

The dark matter mass scaling relation introduced in equation (6.23) for comparing null results can be used to concisely illustrate the mass ranges for which Xe-based and Ge-based experiments give stronger results. In the left panel of Fig. 6-12 we compare limits from XENON10, LUX, and SuperCDMS, for a 10 GeV DM. This can be visualized in a single plot, such as the right panel of Fig. 6-12, which shows the range of limits for each experiment, when referred to 10 GeV DM using the




Figure 6-12: Exclusion limits from XENON10 (dot-dashed green), LUX (solid purple), and SuperCDMS (dotted orange) in v_{min} -space. **Left:** v_{min} -space exclusions for 10 GeV DM. Kinks and slight deviations from monotonicity are present in the SuperCDMS limit, occurring when the “Pmax” method switches the interval used to calculate the exclusion. Similar discrete artifacts are not present in LUX because an observation of zero events is used to calculate the limit, nor in XENON10 because the experiment observed a large number of events. **Right:** limits with v_{min} calculated for a range of DM masses and scaled to $m_\chi = 10 \text{ GeV}$ using the factor appropriate for Ge $\mu_{Ge,\chi}/\mu_{Ge,\chi'}$, where $\mu_{Ge,\chi}$ is the reduced mass for 10 GeV DM and Ge target. This has the effect of mapping SuperCDMS limits to a single line (dotted orange), independent of mass, while limits from Xe experiments are mapped to different locations depending on the DM mass. Dark bands correspond to the range of limits with masses 5-100 GeV, while the light bands show the range of limits for all possible masses, which facilitates comparison of the relative strengths of the limits for different DM masses.

scaling factor appropriate for Ge. Such a plot will contain, for each possible mass, the correct positioning of LUX or XENON10 limits relative to SuperCDMS. It is manifestly clear from the plot that SuperCDMS is always stronger than XENON10 at small DM masses, and that SuperCDMS and LUX are complementary across all DM masses due to the higher v_{min} threshold of LUX. More quantitatively, we find that SuperCDMS is stronger than XENON10 for all v_{min} when $m_\chi \lesssim 59 \text{ GeV}$.

6.6 Future Directions

The analysis described here leaves a variety of questions unanswered and much room for improvement. Although great effort was made to improve the sophistication of the background model relative to CDMS II analyses, we suffered from the lack of calibration data for the main background, the ^{210}Pb sidewall events. To better understand

this issue, the group of Wolfgang Rau at Queens University has built a ^{210}Pb sidewall “source” by contaminating a CDMS detector housing with radon. As of this writing, data with the source is currently being taken by the Berkeley group. It is important that these calibrations are performed with various ionization channels shorted, to simulate the compromised detectors that we have in the experiment. Major progress in more advanced simulations has also been made by improving and tuning the DMC. Although the fine structure of the DMC output is still being validated against calibration data, the results currently look very promising, nearly matching the level of accuracy of the pulse simulation in some cases. And finally, a likelihood analysis of the data analyzed here is now nearing completion, based on a much more detailed pulse simulation that improves on many of the limitations of the background model described in the previous chapters. This will provide somewhat better sensitivity than we achieved with our optimal interval limit.

While these developments unfortunately were not completed soon enough to make it into this thesis, they are still very important for the future of SuperCDMS. SuperCDMS is now designing a much larger experiment at SNOLAB. It is very valuable to develop and perfect robust analysis tools using the Soudan data, which will allow data from SNOLAB to be analyzed more easily. While the improvements above will not lead to enormous increases in sensitivity, they will enable SuperCDMS at SNOLAB to proceed more smoothly and with fewer uncertainties.

Chapter 7

Coherent Neutrino Scattering

Detectors designed for dark matter searches currently provide a unique combination of excellent, few-100 eV energy resolution, extremely low intrinsic backgrounds, and large masses. The rate of improvement in each of these characteristics is also uniquely rapid. Detectors for SuperCDMS at SNOLAB are anticipated to have a total phonon threshold of 350 eV in low-voltage mode and 70 eV in HV mode. Manipulation of the critical temperature and material properties could lead to even better resolution and lower thresholds for cryogenic experiments like SuperCDMS, EDELWEISS, and CRESST [49, 70]. Meanwhile, noble liquid detectors are fulfilling their promise of unrivaled scalability, radiopurity, and self-shielding. The XENON1T collaboration will soon take data with a detector ten times as massive than LUX [135, 136], while LZ will commission a detector with 7T of active mass just a couple years after that [137, 138]. While Argon-based detectors suffer from the intrinsic radioactivity of the long-lived naturally occurring isotope ^{39}Ar , the DEAP collaboration is nevertheless currently commissioning a 3.6T Ar detector [139].

The exponential improvement in these multiple facets DM detector performance begs the question of whether there are alternate applications of these technologies. Detectors for neutrino physics, in particular, often have similar technical requirements to DM detectors: large size, low-background, and good energy resolution. While DM detectors still cannot match the scale of the largest neutrino detectors, such as SuperKamiokande [140] or IceCube with its 1 km³ detector volume [141], the superior energy resolution and energy thresholds of DM detectors enable them to probe low-energy neutrino interactions that have not been measured by existing experiments. The unmeasured process of coherent-neutrino nucleus scattering (CNS) presents one opportunity, which could be sensitive to several beyond-the-standard-model (BSM) effects such as nonstandard neutrino interactions (NSI) and sterile neutrino oscillations. Measurements of low-energy neutrino electron scattering, as another example, could also improve constraints on neutrino magnetic moments.

This section sketches a few novel possibilities for applications of DM-style detectors to low-energy neutrino physics. We review the relevant physics processes and discuss several previous attempts to apply DM detectors to neutrino physics. We then focus on the detection of CNS and its applications with several configurations of neutrino sources. This section contains an abridged compilation of work presented in [142–144],

which were collaborations between several authors.

7.1 Coherent Neutrino Scattering

Coherent neutrino-nucleus scattering is a well-predicted standard model process sensitive to a variety of BSM effects. The interaction involves the scattering of a neutrino with a nucleus by the weak neutral current $\nu + A \rightarrow \nu + A$. At low momentum transfer, the scattering is coherent in the sense that all nucleons participate in the interaction, much like the forward coherence observed in electron scattering $e^- + A \rightarrow e^- + A$ [145]. The CNS cross section depends on the number of neutrons, N , and protons, Z , of the target material with mass M . If T is the recoil energy of the interaction and the incoming neutrino has energy E_ν , the differential cross section is given by

$$\frac{d\sigma}{dT} = \frac{G_F^2}{4\pi} Q_W^2 M \left(1 - \frac{MT}{2E_\nu^2}\right) F(Q^2)^2. \quad (7.1)$$

where G_F and $Q_W = N - (1 - 4 \sin^2 \theta_W)Z$ are the precisely-known Fermi constant and weak charge, respectively [146]. The form factor, $F(Q^2)$, dominates the $\sim 5\%$ cross section uncertainty. As the cross section is well predicted, coherent elastic neutrino-nucleus scattering is an ideal source to search for new physics in the neutrino sector. A cross section measurement with $\sim 10\%$ uncertainty will result in an uncertainty on $\sin^2 \theta_W$ of $\sim 5\%$ [147]. While this uncertainty is large compared to existing and planned precision atomic parity violation and Møller scattering measurements, a discrepancy from the Standard Model prediction already observed in the neutrino sector by the NuTeV experiment [148] motivates more neutrino-based measurements. Notably, a $\sin^2 \theta_W$ measurement with coherent neutrino-nucleus scattering would be at $Q \sim 0.04$ GeV/c, well away from all previous neutrino scattering measurements (including NuTeV's at $Q \sim 4$ GeV/c).

A coherent neutrino-nucleus scattering cross section measurement agreement within 10% uncertainty of the Standard Model prediction will result in limits on non-standard neutrino interactions (NSI) which improve on the present ones by more than an order of magnitude [147, 149, 150]. The low- Q existing and planned precision measurements mentioned above are not sensitive to new physics unique to neutrino interactions. NSI terms can enter the Standard Model Lagrangian through an extra term,

$$\mathcal{L}_{\text{eff}}^{\text{NSI}} = -\varepsilon_{\alpha\beta}^{fP} 2\sqrt{2}G_F (\bar{\nu}_\alpha \gamma_\rho L \nu_\beta) (\bar{f} \gamma^\rho P f) \quad (7.2)$$

where f is a first generation Standard Model fermion, e , u , or d and $P = L$ or R [151]. These $\varepsilon_{\alpha\beta}^{fP}$ terms can appear due to a range of sources, including incorporating neutrino mass into the Standard Model [152] and supersymmetry [153]. The ε_{ee} and $\varepsilon_{e\tau}$ terms, in particular, are poorly constrained by existing measurements. Measuring a coherent scattering cross section in disagreement with the Standard Model expectation could be an indication of NSI. In the case that a cross section discrepancy is observed, multiple nuclear targets could be employed in order to disentangle effects from NSI, a $\sin^2 \theta_W$ anomaly (e.g. consistent with NuTeV), and/or nuclear physics.

Characterizing neutrino coherent scattering is also essential to the understanding of supernova evolution as the energy carried away by neutrinos comprises $\sim 99\%$ of the supernova’s total energy and the coherent channel’s cross section exceeds all others by at least an order of magnitude in the relevant energy region. In a stellar core collapse, the density of the electron/nucleus plasma at the core can reach $>10^{12}$ g/cm³. At these densities, a 20 MeV neutrino’s mean free path is on the order of 0.5 km [154] with the opacity in nucleus-rich regions dominated by coherent scattering. Along with being relevant for supernova evolution in general, the coherent cross section may also affect the supernova neutrino signals expected on Earth.

The coherent process is important for supernova burst neutrino detection as well, providing information about all flavors of neutrinos—not just $\nu_e/\bar{\nu}_e$. For coherent neutrino-nucleus scattering specifically, unlike other channels of flavor-blind neutral current interactions, the nuclear recoil energy is proportional to neutrino energy due to the elastic nature of the interaction. Such information could be combined with charged current $\nu_e/\bar{\nu}_e$ interaction information from other sources to develop a complete picture of oscillations with supernova neutrinos. Note that approximately seven neutrino-nucleus coherent events in one ton of Ar during a ten second window for a galactic core-collapse supernova at 10 kpc are expected with a recoil energy threshold of 5 keV [155]. Although the detectors discussed below are probably too small to provide a sizable sample of supernova burst neutrino-nucleus coherent scatters (unless the supernova is very close), an accelerator-based measurement to confirm the predicted interaction cross section would prove valuable to a next-generation coherent neutrino scattering experiment’s supernova burst neutrino measurement.

7.2 Coherent Neutrino Scattering in Dark Matter Detectors

In this section, we study the prospects for detecting CNS using a future ton-scale Ge-based DM detector in combination with potential future neutrino sources. Such an detector is envisioned to be located at a deep underground site, with the neutrino source on the surface. Section 7.2.1 provides our assumptions for a few example experiments and the raw rates in these detectors with various exposures and as a function of envisioned energy thresholds and baseline lengths. The effect of coherent events as a background for WIMP interactions is also considered in Section 7.2.2. Then, we consider the scenario in which a detector module is brought within tens of meters of the neutrino source and used to obtain a high statistics sample of events for coherent neutrino physics in Section 7.2.3.

7.2.1 Detection

The coherent neutrino-nucleus cross section favors very low recoil energies, in the few-to-tens of keV range. This is well below the threshold of the most sensitive recent and existing large-scale low energy neutrino detectors, like SNO [156], Borexino [157], and KamLAND [158], which explains why this relatively high cross section process

has not yet been observed. Dark matter detectors, on the other-hand, have energy thresholds in the ~ 10 keV range and lower. As such, these detectors are potentially ideal coherent neutrino scattering detectors if given a sufficiently large target mass and neutrino flux. A ton-scale dark matter detector at its nominal depth underground of 1-2 km, in combination with an intense decay-at-rest (DAR) neutrino source, could discover the coherent neutrino scattering process and provide a much needed consistency check of any dark matter limit/detection claim.

It is worth noting that dark matter detectors could observe ^8B solar neutrino coherent events, as has been pointed out in Ref. [159]. The ^8B rate depends on the detector threshold and material, with more events for lighter target nuclei and lower thresholds. This solar neutrino signal becomes negligible for high- A targets when the low-energy recoil threshold of the experiment is between 5 and 10 keV or more. There are zero solar coherent events expected with a 5/10 keV threshold for the targets (Xe, Ge/Ar) and exposures considered here. In contrast, the DAR accelerator source described in Section 7.2.2 will produce a significant number of events within a well-defined time window when the accelerator is on. These events can therefore only be considered a WIMP background during those times. In terms of a coherent neutrino physics measurement, the accelerator-based case has the luxury of an *in-situ* background measurement (when the source is off) in addition to the higher-energy, above-threshold nuclear recoils.

7.2.2 Rates and Time to Discovery

In order to be reasonably concrete, we study a set of experimental designs inspired by proposals for the Deep Underground Science and Engineering Laboratory (DUSEL). We note that the detector designs are not very different from those under consideration at other underground laboratories and that the results can be easily scaled. For the neutrino source, we assume a DAR configuration produced by high intensity cyclotrons which are now under development [160].

DAR neutrinos are known to be an excellent source for neutrino-nucleus coherent scattering experiments [161, 162]. The neutrinos are produced with relatively low energies (< 52.8 MeV), a range where coherent neutrino scattering dominates all other cross sections by about an order of magnitude. Ref. [160] calls for a design with 800 MeV protons, accelerated via high intensity cyclotron(s), which impinge on a carbon target and facilitate the production and eventual decay of pions: $\pi^+ \rightarrow \mu^+ \nu_\mu$, followed by $\mu^+ \rightarrow e^+ \bar{\nu}_\mu \nu_e$. A DAR source flux profile is shown in Fig. 7-1.

High intensity DAR sources are being proposed for CP violation searches involving Gd-doped ultra-large water Cerenkov detectors at underground science laboratories [163–165]. The design for this search [160] calls for multiple accelerator sites at varying distances from ultra-large water Cerenkov detector(s). The DAE δ ALUS cyclotron-based near accelerator is proposed to run with a duty factor between 13% and 20% [160], with an average of 1 MW of power in either case. Each 1 MW accelerator will provide 4×10^{22} neutrinos of each flavor per year produced as an isotropic flux within the time window, expected to be about 67 ms out of 500 ms for the near accelerator with the current design [166]. The absolute normalization of

the neutrino flux, determined with electron-neutrino elastic scattering ($\nu_e e^- \rightarrow \nu_e e^-$) as measured by ultra-large water Cerenkov detector(s), will have a systematic uncertainty of 1% with dominant contributions from the cross section and energy scale uncertainties [163]. The statistical uncertainty on the flux depends on the run period, but is expected to be on the order of 1% as well. The near accelerator site is envisioned at or near the surface of the laboratory with the other cyclotrons located many kilometers away. Note that the far sites will produce insignificant coherent rates due to the $1/r^2$ dependence of the flux. However, the near accelerator can provide a significant event rate, during the 13% beam-on time, for detectors which are sufficiently close and large. Examples of other physics opportunities with this near accelerator are discussed in Refs. [160, 167, 168].

In order to provide realistic calculations, we examine three dark matter experiments which are drawn from the designs of GEODM [169], LZ [137], and MAX [170]. These experiments use germanium, xenon, and argon as their targets, respectively. Note that neon is also commonly considered as an alternative target medium in the noble liquid detectors mentioned. We assume that the accelerator and beam dump are located at or near the surface. As GEODM is proposed for the DUSEL 7400 ft level and LZ/MAX are proposed for the 4800 ft level, we simply consider baseline lengths of 2.3 km and 1.5 km, respectively. The rates for each target are calculated for a ton-year fiducial exposure since the design of each detector is still under consideration.

As discussed above, the coherent neutrino-nucleus interaction takes place at very low recoil energies. Fig. 7-2 shows the distribution of recoil energies for a DAR source with ^{20}Ne , ^{40}Ar , ^{76}Ge , and ^{132}Xe . The experimental rates will strongly depend upon the recoil energy threshold for reconstructed events, T_{min} . As the exact values of this cut for the various detectors are unknown, we consider five possible values of T_{min} . For the aforementioned targets, we find the rates given in Table I, where we assume 100% efficiency for detecting events in the time-window above the threshold T_{min} .

We note that the coherent event rates for dark matter detectors at their nominal depths underground are in the 0-35 events/ton/year range depending on target, baseline, energy threshold, and unrealistically assuming 100% detection efficiency. This is too low to be competitive with presently used neutron sources for detector calibration. Neutrons are adequate for energy calibration, despite their propensity to multiple scatter and activate the detector. However, demonstrating a measured excess between beam-on and beam-off times corresponding to the expected neutrino signal




Figure 7-1: Energy distribution of neutrinos in a DAR source, from Ref [163].

Events/ton/year at distance	For target	T_{min}				
		0 keV	5 keV	10 keV	20 keV	30 keV
1.5 km	^{40}Ar	11.1	9.1	7.5	4.9	3.1
	^{132}Xe	36.4	16.3	6.6	1.1	0.1
	$^{76}\text{Ge}^\dagger$	21.9	14.6	9.4	3.5	1.4
2.3 km	^{76}Ge	9.3	6.2	4.0	1.5	0.6

Table 7.1: Coherent neutrino scattering events/ton/year (with the accelerator running at 1 MW with a 13% duty factor) for various detector layouts and thresholds. The rates reported assume 100% detection efficiency. [†]The present plan is for the GEODM (^{76}Ge -based) baseline to be 2.3 km—although 1.5 km is included for completeness.




Figure 7-2: Recoil energy distributions for coherent scattering 1.5 km from a DAR neutrino source for Ne, Ar, Ge, and Xe. The rates reported assume 100% detection efficiency.

would be a valuable consistency check for any dark matter detector. Such a verification would certainly aid the corroboration of any new dark matter limit/detection claim between DUSEL dark matter experiments.

Detection at a Ton-Scale Dark Matter Experiment

Next, we consider one of the nuclear targets mentioned above (^{76}Ge) in more detail. GEODM is a proposed ton-scale dark matter detector [169] based on the cryogenic Ge crystal technology used in the CDMS experiment. The target design for GEODM is an array of 300 ~ 5 kg Ge crystals operated at 40 mK with a total target mass of ~ 1500 kg. Interaction events in an individual crystal produce populations of athermal phonons and electron-hole pairs which are measured by various phonon and ionization sensors lithographically patterned on the crystal surfaces. The ratio of ionization to phonon signals for an event is a powerful discriminator between electron and nuclear recoils. The signals also enable precise determination of the position and energy of each event, which allow volume and energy cuts. This information is used to set the number of electron recoils that can pass the cuts and pose as nuclear recoils.

	Baseline	Optimistic
Threshold	10 keV	5 keV
Efficiency	See Fig. 7-3	See Fig. 7-3
Energy resolution near threshold	300 eV	25 eV
Neutron background events/(4.5 ton·year)	0.13	0.29
Surface background events/(4.5 ton·year)	1.0	1.0
Fiducial mass uncertainty	5%	1%

Table 7.2: The GEODM detector scenarios considered. The neutron background expectation is after efficiency corrections.

This electron recoil “leakage” into the nuclear recoil band constitutes one source of background events. There is also a background from muon- and radiogenic-induced neutrons, which is controlled through the use of radio-pure materials and passive and active shields to be < 0.15 events/ton/year.

The efficiency of GEODM depends critically on the cuts used to achieve a target leakage background. This is largely a function of the future detector’s performance, which we need to estimate. In light of this, we consider a “baseline” scenario and an “optimistic” scenario for the detector parameters, summarized in Table II. Fig. 7-3 shows the recoil energy distribution in a detector with two hypothetical efficiency curves as a function of recoil energy. The baseline scenario has a 10 keV nuclear recoil energy threshold with the efficiency rising linearly to 0.3 at 20 keV. We additionally assume an energy resolution near threshold of 300 eV and a fiducial mass uncertainty of 5%. These parameters are consistent with the performance of the Ge detectors used in the CDMS II experiment [171]. The optimistic scenario assumes that refinement of the detectors improves all of these parameters. In the optimistic scenario, we assume a 5 keV nuclear recoil energy threshold and an efficiency that rises linearly to 0.6 at 20 keV. We also assume an energy resolution near threshold of 25 eV and a fiducial mass uncertainty of 1%. In all cases, unless otherwise specified, we assume a raw exposure of 4.5 ton·year before the efficiency curve is applied. We also assume a leakage background of 1 event per 4.5 ton·year raw exposure with a spectrum of $e^{-E_{\text{recoil}}/10 \text{ keV}}$, and a neutron background of 0.15 events per ton·year of exposure with a spectrum of $e^{-E_{\text{recoil}}/50 \text{ keV}}$ before the efficiency curve is applied. Events are required to have recoil energies less than 100 keV to be considered nuclear recoils, so we do not use any background events with energies exceeding 100 keV. This neutron spectrum is convolved with each efficiency curve and scaled by the exposure to obtain the expected distribution of neutron events for each scenario. For the baseline scenario we expect 0.13 total neutron events, and for the optimistic scenario we expect 0.29 total neutron events with a 4.5 ton·year exposure.

Using the baseline efficiency and energy threshold, we find that a coherent neutrino rate of 0.8 events/ton/year is expected in a ^{76}Ge -based detector at a baseline of 2.3 km. GEODM expects an overall leakage rate of about 1 background surface leakage event with a 4.5 ton·year exposure, obtained by adjusting efficiency/fiducial volume cuts to get to that number. In addition, there is a neutron-induced nuclear recoil background (absolute rate convolved with detection efficiency) of 0.13 events/(4.5 ton·year) and

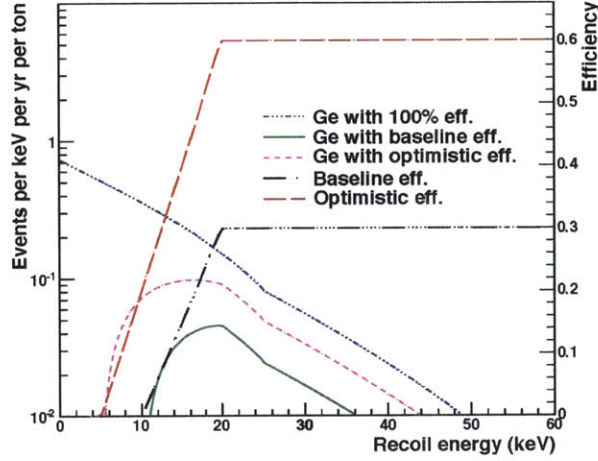


Figure 7-3: Recoil energy distribution for coherent neutrino scattering on Ge at 2.3 km from a DAR neutrino source, before and after the effect of two hypothetical efficiency curves.

0.29 events/(4.5 ton-year) for the baseline and optimistic scenarios, respectively. The total expected background rate during the 13% of beam-on time would therefore be 0.15 events/(4.5 ton-year) for the baseline and 0.17 events/(4.5 ton-year) for the optimistic scenario. The assumptions for GEODM as a coherent neutrino detector are shown in Table III. Based on these considerations and assuming no WIMP “background”, one can see that an experiment like GEODM could find evidence for coherent scattering in a ~ 4.5 ton-year exposure with 3-4 signal events above a background expectation of 0.15 events. The probability for 4 observed events to be completely due to background, with a background expectation of 0.15 events, is $\sim 0.002\%$. The signal rate and evidence/discovery timeline would be quickly improved in the case that the baseline efficiency estimate, especially in the low energy region (and possibly below 10 keV), is too conservative. With the optimistic energy threshold and efficiency scenario, we expect a coherent rate of 2.0 events/ton/year. Under the same background assumption as above, we find that GEODM will discover coherent scattering in a ~ 2 ton-year exposure with 4 signal events above a background expectation of 0.07 events. The probability for 4 observed events to be completely due to background, with a background expectation of 0.07 events, is $\sim 0.00009\%$.

It is worth noting that the absolute (100% efficiency and 100% on-time) solar coherent neutrino interaction rate on a ^{76}Ge target above 5 keV is expected to be 0.079 events/ton/year. In either efficiency scenario and with 13% beam-on time, the solar coherent “background” rate is negligible.

The effect of the coherent background on WIMP sensitivity

Since coherent neutrino scattering is an irreducible background for WIMP searches, the presence of a neutrino source near a dark matter experiment will reduce the sensitivity of the experiment. In calculating a WIMP-nucleon cross section limit, one

GEODM Assumptions	
Scenarios considered	"Baseline" and "Optimistic"
ν source	$4 \times 10^{22} \nu/\text{flavor}/\text{year}$ w/ 13% duty factor
ν flux uncertainty	2%
Distance from ν source	2.3 km
Exposure	4.5 ton·year

Table 7.3: The assumptions used in the text for coherent neutrino detection with GEODM deep underground.

can either use data from only the 87% of the time when the neutrino source is off or data from the 87% beam-off time and the 13% beam-on time. Since background events reduce the sensitivity of a limit in the optimum interval method [97], using the combined exposure may result in a worse limit than using only the beam-off exposure. Fig. 7-4(a) and Fig. 7-4(b) present a comparison of limits for the baseline and optimistic GEODM detector scenarios. The limits assume a GEODM-style detector with a raw exposure of 4.5 ton·year before efficiency cuts and the same assumptions as in the previous section. In particular, we use the hypothetical efficiency curves shown in Fig. 7-3, a neutron rate of 0.15 events/ton/year before efficiency convolution, and a total surface-event leakage of 1 event per 4.5 ton·year raw exposure. We use the recoil energy distribution from Fig. 7-3 for neutrino events in the 13% beam-on time and a detector at 2.3 km from a DAR neutrino source.

Using these parameters, we randomly generate 100 realizations of the background events and compute the average of the upper limits on the WIMP-nucleon cross section for each. In this study, the beam-off data actually has greater sensitivity than the combined beam-on and beam-off data. Using the beam-off data with 87% exposure only, the mean upper limit at maximum sensitivity is $4.0 \cdot 10^{-47} \text{cm}^2$ (baseline) or $1.9 \cdot 10^{-47} \text{cm}^2$ (optimistic). The combined data give a limit of $5.5 \cdot 10^{-47} \text{cm}^2$ (baseline) or $3.6 \cdot 10^{-47} \text{cm}^2$ (optimistic). The results for the 87% exposure obviously also have worse sensitivity than 100% exposure without any beam-on time. If there were no beam-on time, the upper limit at maximal sensitivity would be $3.3 \cdot 10^{-47} \text{cm}^2$ (baseline) or $1.7 \cdot 10^{-47} \text{cm}^2$ (optimistic).

Note that when the number of background events is zero, the sensitivity to WIMPs scales as the reciprocal of the exposure. This occurs because the expected number of events in an experiment is proportional to the product of the cross section σ and the exposure E ,

$$\mu \propto \sigma E. \quad (7.3)$$

A limit at confidence level C is obtained by determining the expected number of events μ such that there is a probability C of observing zero events. Assuming Poisson statistics the expected number of events is $\mu = -\log(1 - C)$. It is thus apparent from Eq. (7.3) that the cross section limit σ corresponding to confidence level C , scales as the reciprocal of the exposure E . In the presence of a nonzero background that is proportional to exposure, the limit will scale more slowly than the reciprocal of the exposure. Fixed backgrounds set by fiducial cuts may produce more complicated




Figure 7-4: Expected average limits for the WIMP-nucleon cross section for a 4.5 ton-year exposure, assuming no WIMP signal and calculated using the optimum interval method. All limits include neutron and surface event leakage background events. The '13% exposure w/ ν ' limit also includes background events from coherent neutrino scattering, while '13% exposure w/ no ν ', '87% exposure w/ no ν ', and '100% exposure w/ no ν ' do not. The 'Combined 87% w/ no ν + 13% w/ ν ' limit is obtained by combining the events and exposures from the '13% exposure w/ ν ' limit and the '87% exposure w/ no ν ' limit, and treating it as a single experiment. Each limit is calculated 100 times with background events randomly drawn from their distributions. The resulting limits are averaged and averages are shown in this figure.

scaling behavior.

7.2.3 Physics with a Detector Close to the Neutrino Source

A suitable detector within tens of meters of the stopped pion source could gather a rather large sample of elastic neutrino coherent scatters for physics studies. The 300 ft adit at DUSEL could provide direct tunnel access to an envisioned experimental site for such a detector, with the cyclotron located just outside the tunnel. As discussed earlier, a coherent cross section measurement is sensitive to a number of physics possibilities. For simplicity, we consider a flux-integrated total cross section measurement as our figure of merit. However, the shape of the cross section as a function of energy is also interesting, especially in the case that a measured total cross section is inconsistent with expectation.

Although many low-threshold nuclear-recoil-sensitive detector technologies could work, for concreteness we consider the specific example of a set of GEODM-derived detectors with 16.7 kg raw mass, 20 m away from the stopped-pion neutrino source. With the optimistic efficiency estimate (Figure 7-3) and a fiducial mass of 10 kg, we expect a detected coherent rate of 0.74 events/(10 kg·day) within the 13% beam-on time window. The background rate design goal for such an experiment would be <0.1 events/(10 kg·day) within the same window. This rate seems reasonable with ~300 ft of rock shielding along with modest passive/active shielding immediately surrounding the detector for prompt cosmic-ray-induced background attenuation/tagging. The radiogenic background is assumed negligible, at a rate consistent with GEODM deep underground. The uncertainty on the non-beam related background estimate will easily be statistics-dominated with an *in-situ* background measurement during beam-off. We also assume that there are no background events from neutrons produced by the DAR neutrino source at a 20 m baseline with 17 m rock and 3 m Fe shielding. This assumption was justified by performing a Geant4 [8] simulation of an isotropic neutron source along a 20 m baseline, fitting the flux at various distances from the neutron source to the functional form

$$F(z) = \frac{Ae^{-z/\lambda}}{z^2},$$

where A and λ are fit parameters and z is the distance along the baseline. The rate from this fit was extrapolated to 20 m distance and a 50 kg·year exposure, and the number of neutron events was found to be negligible. The neutron flux and spectrum used were taken from the SNS source.

A 50 kg·year exposure with the previously described experimental design would yield about 1350 coherent events, assuming a coherent cross section consistent with the Standard Model. With an optimistic 1% uncertainty on the target mass, 0.1 background events/(10 kg·day) with statistical-only error, 2% absolute flux normalization uncertainty, and a 0.5% uncertainty on the energy resolution near threshold, a flux-averaged total cross section measurement with <5% (statistical and systematic) uncertainty would be achieved. The assumptions that went into the event rate and

GEODM Module Close to the ν Source Assumptions	
Scenario considered	"Optimistic"
Source	$4 \times 10^{22} \nu/\text{flavor}/\text{year}$ w/ 13% duty factor
ν flux uncertainty	2%
Distance from ν source	20 m
Exposure	50 kg·year
Background rate	0.1 events/(10 kg·day) in beam window

Table 7.4: The assumptions used in the text for coherent neutrino detection with a GEODM module close to the ν source.

cross section measurement uncertainty estimates are summarized in Table IV.

7.2.4 Conclusions

Coherent elastic neutrino-nucleus scattering has never been observed. Relevant for supernova evolution, supernova-burst neutrino detection, probing non-standard neutrino interactions, and measuring $\sin^2 \theta_W$ with neutrinos at low- Q , the process is very well predicted by the Standard Model and confirmation of the $\sim 5\%$ precision theoretical cross section prediction is needed.

Dark matter detectors can double as coherent neutrino scattering experiments as the products of WIMP and coherent scattering interactions are predicted to be nearly identical. In the case of a decay-at-rest neutrino source and a suite of dark matter experiments at the same site, the deep underground detectors there would merely need to receive a beam timing signal in order to participate in the coherent search. Furthermore, these detectors would receive a free dark matter detection consistency check in the form of non-WIMP rare events in a well-known time window—all with a modest cost to the WIMP analysis/exposure. In both optimistic and baseline detection scenarios, the best limit on the WIMP-nucleon cross section uses only data from the period when the DAR source is off. Depending on the efficiency scenario, the cross section limit may be only about 12% weaker than if no neutrino source were present.

A coherent neutrino interaction discovery in GEODM could be achieved with a 2 ton·year exposure. About 2.0 detected coherent neutrino events/ton/year over a background of 0.03 events/ton/year are expected in a GEODM-style detector at a 2.3 km baseline, given optimistic assumptions for energy threshold and detection efficiency. Even in a conservative (baseline) scenario, with energy threshold and efficiency reasonably consistent with CDMS II, evidence for coherent neutrino scattering could be obtained with a 4.5 ton·year exposure. In addition, a 10 kg fiducial mass GEODM-derived detector brought within 20 m of the neutrino source could collect about 1350 events with a 50 kg·year exposure. Such a sample would be good for a $< 5\%$ flux-averaged total cross section measurement uncertainty and significant tests of the Standard Model.

7.3 Oscillation of eV-Scale Sterile Neutrinos

Sterile neutrino models have been invoked to explain a series of intriguing oscillation signals at $\Delta m^2 \sim 1 \text{ eV}^2$ [172–174]. These signals have relied on neutrino detection through charged current interactions. In the case of charged current appearance, the signal is interpreted as an active flavor oscillating to another active flavor, which can occur at these high Δm^2 values if one or more neutrino mass states with $m_4, \dots \sim 1 \text{ eV}$ is added to the neutrino mass spectrum. The extra mass states are assumed to participate in neutrino oscillations, and must therefore be small admixtures of weakly-interacting neutrino flavor states, with the remaining flavor composition being *sterile* (i.e. non-weakly-interacting). In the case of charged current disappearance, the signal is interpreted as arising from active-flavor neutrino (e, μ, τ) oscillation to any other neutrino flavor (e, μ, τ , or s , with s being the sterile flavor).

The oscillation probabilities for appearance and disappearance through charged current searches are expressed as functions of the active flavor content of the extra mass eigenstate(s) [172, 173]. In this paper, we assume that only one such extra mass state, m_4 , exists. In that case, the oscillation probabilities are given by

$$P(\nu_\alpha \rightarrow \nu_{\beta \neq \alpha}) = 4|U_{\alpha 4}|^2|U_{\beta 4}|^2 \sin^2(1.27\Delta m_{41}^2 L/E) \quad (7.4)$$

in the case of active appearance searches, and

$$P(\nu_\alpha \rightarrow \nu_{\not{\alpha}}) = 4|U_{\alpha 4}|^2(1 - |U_{\alpha 4}|^2) \sin^2(1.27\Delta m_{41}^2 L/E) \quad (7.5)$$

in the case of active disappearance searches, where $\alpha, \beta = e, \mu, \tau$; $\not{\alpha}$ corresponds to all flavors other than α , including active and sterile; $|U_{\alpha 4}|^2$ corresponds to the α -flavor content of the fourth mass eigenstate; and L and E represent the neutrino travel distance and energy, respectively. Note that neither search case is purely sensitive to the sterile neutrino content of the extra neutrino mass state, $|U_{s4}|^2$. In this paper, we discuss a strictly neutral current search using coherent neutrino scattering that allows for pure active-to-sterile oscillation sensitivity.

7.4 Sterile Searches at DAR Neutrino Source

An active-to-sterile neutrino oscillation search is motivated in Section 7.4.1. We describe an experimental design which makes use of a high intensity pion- and muon-decay-at-rest (DAR) neutrino source in Section 7.4.2. The coherent scattering process is introduced and the experimental design is discussed in Section 7.4.3. Sensitivities to neutrino oscillations at $\Delta m^2 \sim 1 \text{ eV}^2$ are shown in Section 7.4.4.

7.4.1 Motivation for an Active-to-Sterile Oscillation Search

A decade ago, sterile neutrino oscillation models were largely motivated by the LSND anomaly [172, 175–178]. This result presented a 3.8σ excess of $\bar{\nu}_e$ events consistent with $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations described by Eq. 7.4 at $\Delta m^2 \sim 1 \text{ eV}^2$ and $\sin^2 2\theta_{\mu e} =$

$4|U_{e4}|^2|U_{\mu4}|^2 \sim 0.003$. The apparent appearance signal is thus interpreted as indirect evidence for at least one additional neutrino carrying the ability to mix with active flavors. Being mostly sterile, an additional neutrino avoids conflict with measurements of the Z invisible width [179] (characteristic of three weakly-interacting light neutrino states) and the three-neutrino oscillation model established by solar [156, 180], reactor [181], and atmospheric/accelerator [182–184] experiments.

The LSND signal was not present in a similar but less sensitive $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillation search by the KARMEN experiment [185]. More recently, however, the MiniBooNE experiment [186] has explored the $\Delta m^2 \sim 1 \text{ eV}^2$ parameter space and yielded a number of interesting results. MiniBooNE features a higher beam energy and larger distance than LSND but preserves the L/E oscillation probability dependence, allowing for an independent cross check of the signal. In searching for ν_e appearance in a pure ν_μ beam, MiniBooNE has excluded $\nu_\mu \rightarrow \nu_e$ oscillations in the LSND Δm^2 range at the 90% CL [187]. However, MiniBooNE’s search for $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations in “anti-neutrino-mode” is only consistent with the no oscillation hypothesis at the 0.5% level [188]. The anti-neutrino result is consistent with LSND and $\bar{\nu}_\mu \rightarrow \bar{\nu}_e$ oscillations in the $\Delta m^2 = 0.1 - 1.0 \text{ eV}^2$ range. The statistics-limited measurement is expected to improve with additional data being taken through at least 2012.

Recently, further results for $\bar{\nu}_e$ disappearance at high Δm^2 have been reported from short-baseline reactor anti-neutrino experiments. More specifically, a re-analysis of the anti-neutrino spectra produced by fission products in a reactor core [189] has led to an effect termed “the reactor anti-neutrino anomaly”, where the ratio of the observed anti-neutrino rate to the predicted rate deviates below unity at 98.6% CL [190]. This can be interpreted as disappearance according to Eq. 7.5, where charged current interactions of active flavors other than e are kinematically forbidden, and/or where the oscillation was into a non-interacting sterile neutrino. Assuming CPT conservation, which requires that Eq. 7.5 holds for both neutrinos and anti-neutrinos, the strongest limits on $\bar{\nu}_e$ disappearance come from a joint analysis of KARMEN and LSND $\nu_e + {}^{12}\text{C} \rightarrow {}^{12}\text{N}_{gs} + e^-$ scattering events, analyzed for evidence of ν_e disappearance [191]. The reactor-anomaly signal is found to be marginally consistent with the KARMEN and LSND ν_e disappearance results.

The above experiments feature a single source, single detector design. An alternative approach is a near-far detector configuration, where the measured flux in the near detector replaces the first-principles flux prediction. A near-far design removes a significant source of uncertainty due to the flux prediction, especially if the detectors are built to be nearly identical. Using the near-far technique, the CDHS [192], CCFR [193], and SciBooNE/MiniBooNE [194] experiments have probed neutrino disappearance at $\Delta m^2 \sim 1 \text{ eV}^2$ using ν_μ charged current interactions. Among the recent near-far comparison data sets, the MINOS experiment has set the only limits on active-to-sterile oscillations using neutral current interactions [195].

7.4.2 The Neutrino Source

A DAR neutrino source can be employed to search for active-to-sterile neutrino oscillations through the neutral current coherent scattering interaction. A search for

active-to-sterile oscillations is envisioned with a series of measurements at different values of L from the DAR source. In our design, a cyclotron directs a proton beam to two graphite targets embedded in a single iron shield. As the DAR neutrino flavor content and energy distribution is driven by the weak interaction, the well understood flux emitted isotropically from each target will be effectively identical, barring oscillations, at each baseline L . We consider two dark matter detector technologies; a germanium-based CDMS-style detector and a liquid argon-based one similar to the CLEAN [196] and CLEAR designs.

Targeting to allow multiple baselines

The experiment described here could use one of these DAE δ ALUS cyclotrons, described in Section 7.2, combined with a dual-target configuration as a neutrino source. In a baseline scenario, the cyclotron-based beam will be diverted between the two targets so that no target receives more than 1 MW average power. The beam will be directed at 90° with respect to the detector, so as to minimize DIF backgrounds. Notably, a multi-target design can also be exploited for a charged current neutrino interaction oscillation measurement with a common detector and multiple baselines. The main technical issue in the two-target cyclotron design is maintaining a good vacuum in the two-prong extraction line. The beam will be “painted” across the face of each target in order to prevent hot spots in the graphite, an effect which will dominate the ± 25 cm uncertainty on the experimental L from each neutrino source. The targets will be arranged in a row enveloped within a single iron shield, with the detector located 20 m downstream of the near target and 40 m downstream of the far target. This configuration has been found to provide the best overall sensitivity to the LSND allowed region.

The analysis below exploits the L dependence of neutrino oscillations. Therefore, the flux of protons on each target must be well understood in time; standard proton beam monitors allow a 0.5% measurement precision. The absolute neutrino flux is less important, as sensitivity to the oscillation signal depends on relative detected rates at the various distances. The systematic uncertainty associated with the flux normalization is 10% if there is no large water or oil detector available and 1.1% if such a detector does exist. A high statistics ν -electron scattering measurement at a large water detector provides a precise determination of the flux normalization.

7.4.3 Experimental Setup

The envisioned experimental setup is consistent with the current DAE δ ALUS accelerator proposal and follows a realistic detector design. A single DAE δ ALUS cyclotron will produce 4×10^{22} ν /flavor/year running with a duty cycle between 13% and 20%. A duty cycle of 13% and a physics run exposure of five total years are assumed here. With baselines of 20 m and 40 m, the beam time exposure distribution at the two baselines is optimal in a 1 : 4 ratio: one cycle to near (20 m), four cycles to far (40 m). Instantaneous cycling between targets is important for target cooling and removes systematics between near and far baselines associated with detector changes

ν source	$4 \times 10^{22} \nu/\text{flavor}/\text{year}$
Duty factor	13%
Baseline correlation	0.99
ν flux norm. uncertainty	1.5%
Uncorr. sys. uncertainty	0.5%
Distances from ν source	20 m, 40 m
Exposure	5 years: 1 near, 4 far
Depth	300 ft

Table 7.5: The experimental configuration assumptions.

over time. The accelerator and detector location is envisioned inside an adit leading into a sharp 300 ft rise at the Sanford Research Facility at Homestake, in South Dakota. The neutrino flux normalization uncertainty at each baseline is conservatively expected at 1.5%. We assume the flux has been constrained to this level by an independent measurement of ν -electron scattering with a large water-based Cerenkov detector also assumed to be in operation at Sanford Labs. The 1.5% uncertainty estimate takes into consideration the theoretical uncertainty in the ν -electron scattering cross section and the statistics achievable with a large water detector. The flux normalization correlation coefficient between the near and far baselines is conservatively set to 0.99, its deviation from unity being dominated by differences between the two beam dumps. An uncorrelated systematic uncertainty of 0.5% at each baseline, is also included. The general experimental assumptions can be seen in Table 7.5.

We also consider a “dedicated” physics run scenario in which the duty factor is raised from 13% to 50% for all five years. With the instantaneous power achievable remaining constant, this change leads to an average power increase of a factor of 4. Steady-state and beam-related backgrounds also increase by this factor in a dedicated scenario. The nominal duty factor of 13% is driven by the requirement that the various DAE δ ALUS accelerator baseline beam windows do not overlap in time. A dedicated scenario is possible in consideration of maintaining sufficient target cooling and the phased DAE δ ALUS deployment timeline. With a 13% duty factor only required when all baselines have operational accelerators, a longer duty factor and higher average power seems possible in DAE δ ALUS single-baseline-only operation. Note that although only two targets are required for the experimental design described here, supplementing the beamline with more targets can ensure optimal use of beam time in consideration of cooling requirements and ultimately increase neutrino oscillation sensitivity.

Germanium detector – signal and backgrounds

A 100 kg active mass of germanium is considered for the experiment described here, using detectors similar to SuperCDMS iZIPs. The detection efficiency above a 10 keV threshold is set to 0.67 with a 3% energy resolution near the threshold. These assumptions are reasonably conservative, based on our experience with the analysis in Chapter 5.

Two classes of background events are considered for a germanium experiment:

1. *Misidentified electronic recoils* - Electronic recoils can be produced by photons and beta particles interacting with the active detection medium. Misidentification of such events is particularly problematic near the detector surfaces, where the collection of electron-hole pairs is suppressed and discrimination is less effective. We assume a rate of radiogenic background detection (~ 2 events/year) above 10 keV, which is consistent with the iZIP performance studies described in Chapter 3.
2. *Cosmogenic neutrons* - Single scatter neutrons can produce a signal identical to a coherent neutrino scattering event, and the rate of these events would be significant at a shallow site. As a point of reference for surface experiments, the CDMS experiment located at the Stanford Underground Facility with 16 m.w.e. of overburden measured a neutron background of 0.67 events/(kg·day) [197]. This figure could be significantly reduced with additional active and passive shielding and the larger overburden envisioned for the DAE δ ALUS site. A cosmogenic-induced background of 0.1 detected events/(10 kg·day), after correcting for efficiency and during beam-on, is assumed. This value is considered a design goal and can be met with a 300 ft overburden and modest active and passive shielding.

In this study, the estimated radiogenic and cosmogenic background rates are distributed evenly across the germanium nuclear recoil energy range considered, 10 keV to 100 keV. We neglect other sources of background, such as radiogenic neutrons from the rock cavern surrounding the experiment. Dark matter experiments with large exposure, such as CDMS II, have typically not been limited by this source of background.

Liquid argon detector – signal and backgrounds

A single phase liquid argon detector can be used to detect the scintillation light created by WIMP- or coherent neutrino-induced nuclear recoils. Such detectors employ a large, homogeneous liquid argon volume surrounded by photomultiplier tubes (PMTs). Inner detector surfaces as well as the PMTs themselves are usually covered in a wavelength shifting substance which converts the 128 nm scintillation light into the visible spectrum for detection.

A 456 kg active mass of liquid argon with a flat efficiency of 0.50 above a 30 keV energy threshold is considered for the experiment described here. The detection volume and efficiency are consistent with the proposed CLEAR design [162]. An 18% energy resolution near threshold is used, assuming resolution slightly worse than what would be expected from photoelectron Poisson statistics [198], 6 photoelectrons/keVee light collection, and a quenching factor of 0.25 [199].

There are three primary sources of background that are considered for a single phase liquid argon detector:

1. *Cosmogenic neutrons* - The muon-induced neutron background is contingent on the geometry of the site, overburden, and active/passive detector shielding. Muon events and muon-induced neutrons can be vetoed with high efficiency and low detector dead time in a liquid argon detector near the surface [162]. The target design cosmogenic background is 0.1 detected events/(10 kg·day).
2. *^{39}Ar contamination* - ^{39}Ar is a naturally occurring radioisotope with an isotopic abundance of $^{39}\text{Ar}/\text{Ar} = 8 \times 10^{-16}$, corresponding to a specific activity of 1.01 Bq/kg [200]. The isotope is a beta emitter with an energy endpoint of 565 keV.

Pulse-shape discrimination (PSD) can be used to separate the ^{39}Ar -induced electronic recoils from the nuclear recoils produced in WIMP and coherent neutrino scattering events [201]. The electronic recoil contamination (ERC) of the nuclear recoils decreases exponentially as the number of photoelectrons detected increases. Refs. [202, 203] measure the ERC for PSD in the single phase liquid argon DEAP-1 detector (4.85 photoelectrons per keVee) and also provide a "theoretical" Monte Carlo estimate of the ERC attainable for an ideal detector with 4π PMT coverage and 6 photoelectrons/keVee. Both scenarios correspond to a 50% efficiency for nuclear recoil detection in the fiducial volume. Note that, according to Ref. [204], the MicroCLEAN experiment has achieved 6 photoelectrons/keVee sensitivity. The abundant ^{39}Ar background could further be alleviated with the use of depleted argon from underground sources, which has an isotopic abundance of ^{39}Ar that is $<5\%$ of natural argon at the surface [205]. Figure 7-5 shows the rate of ^{39}Ar after PSD with 13% on-time, for two assumptions of ERC reported in Ref. [202, 203]. The theoretical ERC with non-depleted liquid argon is employed for this study.

3. *Surface contamination* - Radioactive impurities on the detector surfaces can decay and contribute to the background. These surface backgrounds have been measured in the DEAP-1 detector and were found to have an activity of 1.3×10^{-4} Bq [206]. Depending on the origin of these events, the scaling and resulting background prediction can differ. If the events are due exclusively to ^{210}Pb surface contamination, the DEAP-1 figure can be scaled by detector surface area to yield 1.3×10^4 events/year. However, this rate may be substantially reduced by the use of cleaner materials, scrubbing of the surface, and fiducialization. A surface background contamination of 100 detected events/year is assumed here and can be considered a design goal.

The 30 keV energy threshold employed here is larger than the oft-chosen 20 keV threshold in single phase liquid argon detectors in order to mitigate the steeply falling ^{39}Ar contamination. If ^{39}Ar discrimination improves in a future design, adjusting the threshold to (e.g.) 20 keV would allow a 60% larger signal sample. In this study, the estimated surface and cosmogenic background rates are distributed evenly across the argon nucleus recoil energy range considered, 30 keV to 200 keV.

One additional possibility that would significantly reduce the non-beam-related background would be to use a pulsed source of neutrinos, such as at the Spallation




Figure 7-5: The expected ^{39}Ar background energy spectrum under two sets of assumptions. The line labeled “Measured” corresponds to an ERC that was obtained in a detector with 4.85 photoelectrons/keVee. The line labeled “Theoretical” is the ERC simulated in an ideal detector with 6 photoelectrons/keVee and represents the background ERC used for this study. Both lines correspond to a 50% efficiency for detecting nuclear recoils in the fiducial volume.

Neutron Source (SNS). The SNS produces protons in very short bunches of <750 nsec at a rate of about 60 Hz, so that the time window for expected signal events is a small fraction of the total running time. Combining a pulsed DAR beam structure with a liquid argon detector was previously proposed by the CLEAR experiment [162], allowing them to claim an additional rejection of 6×10^{-4} for steady-state, non-beam-related backgrounds using a timing cut. Although the DAE δ ALUS proposal does not include this timing structure, the experimental concept described here could be employed at other facilities.

The detector-specific assumptions are summarized in Table 7.6 and the expected signal and background rates are shown in Fig. 7-6.

Neutron flux from the source

DAR sources produce a large flux of neutrons, arising from spallation reactions of protons with the beam dump material. For the DAR source considered here, the neutrons have energies up to 800 MeV. In a 1 MW beam, the neutron production rate is $\sim 10^{16}/\text{s}$ and attenuation lengths may be as high as tens of centimeters. Single scatter neutrons can produce elastic recoils in the detector volume that are indistinguishable from coherent neutrino scattering on an event-by-event basis. Moreover, because the neutron flux is attenuated by matter, underestimating the neutron background in the detector could mimic a deviation from the $1/r^2$ -dependence of the neutrino flux, similar to what is expected for neutrino disappearance. It is therefore essential to

	^{76}Ge	^{40}Ar
Active mass	100 kg	456 kg
Efficiency	0.67 (flat)	0.50 (flat)
Threshold	10 keV	30 keV
$\frac{\Delta E}{E}$ at threshold	3%	18%
Radiogenic background	2/year	See text
Cosmogenic background	0.1/(10 kg·day)	0.1/(10 kg·day)
Beam-related background	0/year	0/year

Table 7.6: The assumptions relevant for the specific detector technologies considered.




Figure 7-6: The expected non-oscillated signal and total background rates at a 20 m baseline for the two detector technologies considered in the baseline physics run scenario. The rates correspond to what is expected for one full year of near-target-only operation at 13% duty factor.

locate the detector far enough from the source that the beam-related neutron flux is negligible.

A precise estimate of the neutron flux would require detailed knowledge of the experimental site, beam configuration, and shielding. The neutron flux is estimated with a Monte Carlo simulation of the experimental geometry consistent with the DAE δ ALUS proposal and several simplifying assumptions. Instead of simulating the passage of neutrons through the beam dump shielding, we simply assume that a cubic shield with sides of length 6 m is sufficient to reduce the escaping neutron flux to a level consistent with safety regulations. Also, we assume that this cube of shielding is adjacent to a rock (SiO_2) cliff. The maximum permissible annual dose for workers in a restricted area with a neutron beam is 100 mRem. The neutron flux escaping the shielding is set to a rate equivalent to an exposure of 100 mRem in 40 hours.

Using a Geant4-based simulation [8], neutrons are injected at the edge of the shielding cube. The neutrons are simulated in energy bins from 0-30 MeV. The flux

is tallied at 20 cm intervals into the rock cliff, and the fluxes beyond 1 m into the cliff are fit to the functional form

$$\Phi(z) = \frac{Ae^{-z/\lambda}}{z^2}, \quad (7.6)$$

where A and λ are fit parameters, and z is the distance from each flux tally point to the DAR source. The neutron fluxes are in reasonable agreement with this functional form. The fit function is then used to extrapolate the flux to a full year of running and larger distances from the source. A simulation is also employed to estimate the fraction of incident neutrons that produce single-scatter nuclear recoils in the detection volume. Less than 0.2 beam-related events are expected per year for a 456 kg liquid argon detector at a 12 m baseline. The beam-related background at 20 m from the source, the shortest relevant detector baseline considered here, is therefore assumed to be negligible.

7.4.4 Measurement Strategy and Sensitivity

Overall strategy

Neutrino oscillations depend upon neutrino energy and distance traveled. Since the neutrino energy cannot be reconstructed precisely with the coherent interaction, our sensitivity to the oscillatory behavior arises mainly from L , a value which is well determined by the location of the target being used at any given time and its distance to the common detector. In the case that a disappearance signal is detected, the target exposure priorities for the two baselines can be optimized to maximize sensitivity.

The purely neutral current experiment described is sensitive to the effective disappearance of all three types of neutrinos present in the beam, ν_μ , $\bar{\nu}_\mu$, and ν_e , into ν_s . We assume this disappearance can be approximated by a two-neutrino oscillation driven by a Δm^2 in the LSND allowed region, and that the oscillation probability under the approximation is the same for neutrinos and anti-neutrinos. The baselines for the experiment, 20 m and 40 m, have been chosen in order to provide the best sensitivity to the LSND allowed parameter space, given the neutrino energy spectrum of each flavor in the beam. The experiment described here provides indirect sensitivity to the LSND allowed parameter space by simultaneously measuring terms describing the amplitude of active neutrino mixing to a sterile neutrino: $4|U_{e4}|^2|U_{s4}|^2$ in the case of ν_e in the beam, and $4|U_{\mu4}|^2|U_{s4}|^2$ in the case of ν_μ and $\bar{\nu}_\mu$ in the beam. These terms are then translated to the appearance amplitude measured by LSND, $\sin^2 2\theta_{\mu e} = 4|U_{e4}|^2|U_{\mu4}|^2$. Sensitivity to $\sin^2 2\theta_{\mu e}$, along with simultaneous sensitivity to $\sin^2 2\theta_{ee} = 4|U_{e4}|^2(1 - |U_{e4}|^2)$ and $\sin^2 2\theta_{\mu\mu} = 4|U_{\mu4}|^2(1 - |U_{\mu4}|^2)$, are considered the figures of merit here, as they can be easily compared to existing charged current appearance and disappearance measurements. Of course, distinguishing between $\sin^2 2\theta_{ee}$ and $\sin^2 2\theta_{\mu\mu}$ in the case of an observed disappearance is not possible in a flavor-blind experiment. Therefore, we rely on marginalizing over the full parameter space of $|U_{\mu4}|$ and $|U_{e4}|$ explored, in the most conservative case possible, when drawing sensitivity contours for each case.

The sensitivity to any particular set of oscillation parameters is obtained by simul-

taneously fitting the expected flavor-summed coherent signal events as a function of recoil energy at the near and far baselines. The events at each baseline are distributed among bins of nuclear recoil energy (1 bin/10 keV); however, the sensitivity results are largely insensitive to the number of recoil energy bins used in the comparison.

Sensitivities

The signal predictions are evaluated for each set of oscillation parameters, $\Delta m_{41}^2 \equiv \Delta m^2$, $|U_{\mu 4}|$, and $|U_{e 4}|$. A χ^2 is calculated by comparing the oscillations-predicted spectra, including backgrounds, to the no-oscillations prediction.

The χ^2 is constructed as

$$\chi^2 = \sum_{i,j=1}^{N_{\text{bins}}} (P_i - N_i)(P_j - N_j) M_{ij}^{-1} , \quad (7.7)$$

where i and j denote the energy bins at the near and far baselines, respectively; P_i is the oscillations-predicted event spectrum as a function of $N_{\text{bins}} = 1, \dots, 10, 11, \dots, 20$ bins, corresponding to (e.g.) 10 energy bins for the two baselines appended side by side; N_i is the corresponding no-oscillations spectrum; and M_{ij}^{-1} is the inverse covariance matrix including statistical and systematic uncertainties and normalization systematic correlations between the two baselines and different recoil energy bins. Note that the background contributions to P_i and N_i cancel. The background-contributed statistical uncertainty, however, is accounted for in M_{ij} . The background contribution can be measured with high statistics during beam-off cycles, and so systematic uncertainties associated with background are small relative to statistical uncertainties.

The oscillations-predicted spectra, P_i , are obtained by summing over all neutrino flavors predicted in each recoil energy bin of the unoscillated spectrum, and reweighting each neutrino according to its flavor $\alpha = e, \mu$ by the following “active” survival probability

$$\begin{aligned} P(\nu_\alpha \rightarrow \nu_{\text{active}}) &= 1 - P(\nu_\alpha \rightarrow \nu_s) \\ &= 1 - \sin^2 2\theta_{\alpha s} \sin^2(1.27 \Delta m^2 L/E) , \end{aligned} \quad (7.8)$$

where ν_{active} can be any active state including ν_α , and $\sin^2 2\theta_{\alpha s} = 4|U_{\alpha 4}|^2|U_{s 4}|^2$. By unitarity assumptions, $|U_{s 4}|^2$ is a function of $\sum_{\alpha=e,\mu,\tau} |U_{\alpha 4}|^2$,

$$|U_{s 4}|^2 = 1 - \sum_{\alpha=e,\mu,\tau} |U_{\alpha 4}|^2 . \quad (7.9)$$

During the fit, we vary $|U_{e 4}|$, $|U_{\mu 4}|$, and Δm^2 . For simplicity, however, we assume $|U_{\tau 4}| = 0$. Note that a non-zero $|U_{\tau 4}|$ would increase the active survival probability for any given $|U_{e 4}|$ and $|U_{\mu 4}|$, and would therefore make this search slightly less sensitive to oscillations in terms of $\sin^2 2\theta_{\mu e}$. On the other hand, if non-zero $|U_{e 4}|$ and $|U_{\mu 4}|$ were to be established independently by other short baseline experiments, the type of

neutral current search outlined in this paper may offer sensitivity to $U_{\tau 4}$, depending on the sizes of $|U_{e4}|$, $|U_{\mu 4}|$ and $|U_{\tau 4}|$.

Figures 7-7, 7-8 and 7-9 show the expected sensitivity to the LSND allowed region with a germanium detector in the baseline and dedicated physics run scenarios and an argon detector in the baseline scenario, respectively. In obtaining the sensitivity curves, the 3D search grid is reduced from $(\Delta m^2, |U_{e4}|^2, |U_{\mu 4}|^2)$ to a 2D space of Δm^2 and $\sin^2 2\theta_{\mu e} = 4|U_{e4}|^2|U_{\mu 4}|^2$. Note that a non-zero $\sin^2 2\theta_{\mu e}$ requires both ν_e and ν_μ disappearance.

The $\sin^2 2\theta_{\mu e}$ sensitivity curves are obtained using a raster scan in Δm^2 space. That is, each curve maps out the maximum $\sin^2 2\theta_{\mu e} = 4|U_{e4}|^2|U_{\mu 4}|^2$ which satisfies $\chi^2 \leq \Delta\chi_{cut}^2$ at a given confidence level, for each point in Δm^2 . The 90%, 99%, and 3σ confidence level curves shown in this paper correspond to $\Delta\chi_{cut}^2 = 1.64, 6.63$, and 9.00 for a one degree of freedom, one-sided raster scan (90%), and a one degree of freedom, two-sided raster scan (99% and 3σ), respectively.




Figure 7-7: Sensitivity to the LSND 90% CL allowed parameter space with a germanium-based detector under the baseline physics run scenario.

Figure 7-10 shows the oscillation sensitivity for a germanium detector in terms of the disappearance amplitudes which would be accessible in charged current searches, $\sin^2 2\theta_{ee} = 4|U_{e4}|^2(1 - |U_{e4}|^2)$ and $\sin^2 2\theta_{\mu\mu} = 4|U_{\mu 4}|^2(1 - |U_{\mu 4}|^2)$ overlaid with the region allowed by LSND at 90% CL, assuming the LSND best-fit $\Delta m^2 = 1.2 \text{ eV}^2$. The curves are obtained using a one-sided raster scan in $\sin^2 2\theta_{ee}$ with the $\Delta\chi_{cut}^2$ values defined above. The figure also shows the approximate region of $\sin^2 2\theta_{ee}$ values allowed at 90% CL by fits to the reactor anomaly and gallium experiment calibration data sets in Ref. [190]. The “reactor” allowed contour is for $\Delta m^2 \gtrsim 1.5 \text{ eV}^2$ and is relatively independent of Δm^2 in this region. As a reference, limits on $\sin^2 2\theta_{\mu\mu}$ from the MINOS neutral-current oscillation search correspond to $\sin^2 2\theta_{\mu\mu} < 0.1$ at 90%




Figure 7-8: Sensitivity to the LSND 90% CL allowed parameter space with a germanium-based detector under the dedicated physics run scenario.

CL, for $\Delta m^2 = 1.2 \text{ eV}^2$ [195].

Figures 7-7 and 7-9 show that, despite the difference in fiducial mass, the 100 kg germanium detector performs slightly better than the 456 kg liquid argon one. The difference is in part due to the difference in nuclear recoil energy threshold; 10 keV for germanium, 30 keV for argon. This emphasizes the fact that a low detector energy threshold is important for obtaining a high-statistics sample of coherent neutrino scattering events as the rate is dominated by events with very low energy recoils ($\lesssim 10 \text{ keV}$).

In a baseline physics run scenario, an experiment featuring a germanium- or argon-based detector can exclude the LSND best-fit mass splitting ($\Delta m^2 = 1.2 \text{ eV}^2$) at 3.8σ or 3.4σ , respectively. The LSND best-fit mass splitting is excluded at 4.8σ in the dedicated, germanium-based physics run scenario considered. For sensitivity in terms of $\sin^2 2\theta_{ee}$ and $\sin^2 2\theta_{\mu\mu}$, a germanium-based experiment in the baseline scenario could exclude nearly all of the available 90% CL LSND parameter space at the 3σ level and large portions of the available reactor anomaly allowed region, assuming $\Delta m^2 \sim 1.2 \text{ eV}^2$.

7.4.5 Conclusions

This paper has described a method to search for active-to-sterile neutrino oscillations at relatively short baselines using neutral current coherent neutrino-nucleus scattering. Detection of such a process could definitively establish the existence of sterile neutrinos and measure their mixing parameters.

An experiment that relies on the high statistics detection of an as-yet-undetected




Figure 7-9: Sensitivity to the LSND 90% CL allowed parameter space with an argon-based detector under the baseline physics run scenario.

process is obviously difficult. However, all of the technology required for such an experiment either exists or has been proposed with realistic assumptions. A cyclotron-based proton beam can be directed to a set of targets, producing a low energy neutrino source with multiple baselines. This allows a measurement of the distance dependence of an oscillation signal without moving detectors or instrumenting multiple devices. Both a germanium-based detector inspired by the CDMS design and a liquid argon detector inspired by the proposed CLEAR experiment would be effective for performing these measurements.

Along with relevance in understanding Type II supernova evolution and supernova neutrino detection, coherent neutrino-nucleus scattering can provide sensitivity to non-standard neutrino interactions, the weak mixing angle, and, as shown in this paper, neutrino oscillations at $\Delta m^2 \sim 1 \text{ eV}^2$. Depending on the detector technology and run scenario, the experiment described is sensitive to the LSND best-fit mass splitting at the level of $3\text{--}5\sigma$ and can probe large regions of the LSND and reactor anomaly allowed regions. The experiment offers a pure and unique analysis of neutrino oscillations that is complementary to charged current-based appearance and disappearance searches.

7.5 Sterile Searches with Electron-Capture Neutrino Sources

The use of intense neutrino sources to probe sterile neutrinos has been proposed previously in the literature [207–210]. The difficulty with all such detection schemes is the




Figure 7-10: Sensitivity to disappearance amplitudes accessible with charged current searches, assuming the LSND best-fit $\Delta m^2 = 1.2 \text{ eV}^2$. The sensitivity corresponds to a germanium-based detector under the baseline physics run scenario. The LSND band represents the 90% CL allowed values of $\sin^2 2\theta_{\mu e}$ at $\Delta m^2 = 1.2 \text{ eV}^2$. “Reactor” refers to the result reported in Ref. [190] and indicates the range of $\sin^2 2\theta_{ee}$ values preferred by a joint fit to reactor and gallium experiment calibration measurements. The reactor result is nearly independent of Δm^2 , for Δm^2 values near and above 1.5 eV^2 .

low energy threshold necessary to detect the signature nuclear recoil. Such difficulties are circumvented by either resorting to targets with low mass numbers—considerably lowering the cross-section amplitude and requiring large mass detectors—or by looking instead at the charged current reaction using higher energy neutrinos. In this section, we discuss a low energy threshold detector based on cryogenic bolometers that has the capability of reaching recoil energy thresholds as low as 10 eV. Such detectors re-open the door to neutral current coherent scattering as a method for sterile neutrino detection.

In our study, we will mainly consider mono-energetic electron capture sources, all of which have neutrino energies below 1 MeV. The maximum momentum transfer for such sources is $|q_{\text{max}}| \leq 2E_\nu \ll 2 \text{ MeV}$. Since the form factor $F(q^2) \rightarrow 1$ for cases where the scale of the momentum probe is much larger than the size of the nucleus, we can safely ignore this correction factor for our analysis.

The maximum kinetic energy imparted on the nuclear recoil depends on the neutrino energy and the mass of the recoil target:

$$T_{\text{max}} \leq \frac{E_\nu}{1 + \frac{M_A}{2E_\nu}} \quad (7.10)$$

For a silicon target at 1 MeV, that implies a maximum kinetic energy of about 50 eV. For a germanium target the maximum kinetic energy would be around 20 eV. Such low kinetic energies are why detection of the process has been so elusive to date. The fraction of events that is detectable by a given experiment depends crucially on the inherent threshold of the detector. For a monochromatic source of energy E_ν , the effective cross-section can be written as:

$$\bar{\sigma} = \int_{T_0}^{T_{\max}} \frac{d\sigma}{dT}(E_\nu) \cdot dT \quad (7.11)$$

$$\bar{\sigma} = \sigma_0(E_\nu) \cdot f(E_\nu, T_0) \quad (7.12)$$

where $\sigma_0(E_\nu) \equiv \frac{G_F^2}{4\pi} E_\nu^2 Q_W^2$ is the total integrated cross-section assuming no energy threshold and $f(E_\nu, T_0)$ represents the fraction of events above a given threshold energy, T_0 . In the limit that $E_\nu \ll M_A$, the fraction of events above threshold can be written as:

$$f(E_\nu, T_0) = (1 - \frac{T_0}{T_{\max}})^2 \quad (7.13)$$

Any detector hoping to detect such a signal with sufficient statistics must achieve as low a recoil threshold as possible.

7.5.1 The ^{37}Ar Source

Oscillometry-based measurements benefit greatly from the use of mono-energetic neutrino sources, since it reduces the measurement to a pure flux-versus-distance analysis. Low energy electron capture sources provide the most effective and clean source of such neutrinos available to date [211]. A number of such neutrino sources have been considered in the literature; a few of them are listed in Table 7.7. Historically, two such high intensity source have been produced for neutrino studies: a ^{51}Cr source, used by the SAGE and GALLEX experiments [212, 213], and an ^{37}Ar gaseous source used in conjunction with the SAGE experiment [214].

Source	Half-Life	Progeny	Production	E_ν	Gamma (?)
^{37}Ar	35.04 days	^{37}Cl	$^{40}\text{Ca}(n,\alpha)^{37}\text{Ar}$	811 keV (90.2%) 813 keV (9.8%)	inner brems
^{51}Cr	27.70 days	^{51}V	n capture on ^{50}Cr	747 keV (81.6%) 427 keV (9%) 752 keV (8.5%)	320 keV γ
^{65}Zn	244 days	^{65}Cu	n capture on ^{64}Zn	1343 keV (49.3%) 227 keV (50.7%)	1.1 MeV γ

Table 7.7: List of properties of selected electron capture neutrino sources.

The ^{37}Ar source is perhaps the most ideal with respect to a future coherent-scattering measurement, for a number of reasons:

- ^{37}Ar produces a very high-energy, near mono-energetic neutrino (90.2% at 811 keV, 9.8% at 813 keV).
- With the exception of inner bremsstrahlung photons, almost all the energy is carried away by neutrinos, facilitating shielding and enabling the source to be extremely compact.
- Extremely high production yield per reactor target.

The SAGE collaboration successfully produced such a source with a total activity of about 400 kCi to be used in conjunction with their gallium solar neutrino detector. The source was also very compact, extending 14 cm in length and 8 cm in diameter, including shielding [215]. Further reduction in size might be possible, even with increased activity, making ^{37}Ar an ideal portable neutrino source.

Despite its clear advantages as a source and its historical precedent, production of such sources is less than ideal. The reaction process by which it is generated ($^{40}\text{Ca}(n,\alpha)^{37}\text{Ar}$) requires a high fast neutron flux above 2 MeV, an energy regime where few reactors operate [216, 217]. Production also requires large amounts of CaO and processing in nitric oxide, which makes post-production handling difficult. Far less complex to produce is ^{51}Cr , which requires only thermal neutrons capturing on ^{50}Cr . However, as a source, the high energy gamma produced from the decay of the excited state of ^{51}V imposes more shielding requirements. As such, intense ^{51}Cr may be less ideal for this investigation, but still worth considering given the advantages in producing the required activity.

With its high energy neutrino emission, ^{65}Zn is also an attractive source for consideration [218]. However, its 1.1 MeV gamma emission complicates the shielding, so this source is not considered further.

7.5.2 The Detector

The detector requirements for this experiment are extremely challenging. Due to the low energy of the neutrinos (≤ 1 MeV), the recoil energy deposited in the target is in the order of tens of eV, while the minimum mass needed is hundreds of kilograms. Methods of determining the energy deposition from particle interactions in a target include measuring the ionization, the scintillation, and/or the phonon excitations in the material. For nuclear recoils of tens of eV, the fraction of the energy deposited by the scattering event that produces free or conduction band electrons (the quenching factor) is unknown at these energies, and is expected to be very low (could be zero for some materials). Thus any readout scheme involving ionization channels will be at a severe disadvantage. Similar uncertainties hold for the scintillation yield from nuclear recoils at these energies. An additional problem for both ionization and scintillation readout is that the energy required to create a single electron, electron-hole pair, or scintillation photon from a nuclear recoil in most liquid or solid targets is a few eV for

ionization and tens of eV for scintillation. Thus, even if any quanta were produced, Poisson statistics would make the measurement of the energy of any given recoil event fairly poor. We have therefore focused our attention to the measurement of phonons created in the interaction. With mean energies of the order of μeV , thermal phonons provide high statistics at 10 eV and sample the full energy of the recoil with no quenching effects.

Detector Design and Expected Performance

This section focuses on the details of our design choices for this experiment. For more background material on low temperature detectors, we suggest the following: a comprehensive overview of the applications can be found in [219]; specifics on operation and capabilities are reviewed in [220], and details of TES physics are reviewed in [221].

The threshold for a bolometer is a function of its baseline energy resolution. A dimensionless measure of the sensitivity of a resistive thermometer at a temperature T and resistance R is the quantity α , defined as $\alpha \equiv \frac{T}{R} \frac{dR}{dT}$. The energy resolution of a TES bolometer is approximately [222]

$$\Delta E_{\text{rms}} = \sigma_E \approx \sqrt{\frac{4k_B T^2 C_{\text{tot}}}{\alpha}} \sqrt{\frac{\beta + 1}{2}}, \quad (7.14)$$

where k_B is the Boltzmann constant, C_{tot} is the total heat capacity of the bolometer, and β is the exponent of the temperature dependance of the thermal conductivity between the bolometer and the refrigerator. To unambiguously detect events above the noise from the detector, we set the experimental threshold to $7.5 \sigma_E$. For a 10 eV threshold, we then need a detector with $\sigma_E < 1.33$ eV, or expressed in terms of the full width at half maximum, $\Delta E_{\text{FWHM}} = 2\sqrt{2 \ln 2} \sigma_E < 3.14$ eV.

Assuming a conduction path to the cold bath of the refrigerator dominated by Kapitza resistance, $\beta = 4$, and with a temperature $T = 15$ mK, a 10 eV threshold could be attained with a heat capacity $C_{\text{tot}} \leq 200$ pJ/K. However, this model is not complete, as it assumes a perfectly isothermal bolometer. In practice, the various internal heat capacity systems of the bolometer are decoupled from each other through internal conductances, and thermalization times of each separate heat capacity must also be taken into account. These internal decouplings introduce various sources of noise, degrading the energy resolution of the bolometer and consequently requiring a smaller heat capacity to attain the desired threshold.

Fig. 7-11 shows a schematic of the model. The bolometer is connected to the cold bath at temperature T_b through a weak thermal conductance G_{pb} . The total heat capacity can be described by $C_{\text{tot}} = C_{\text{Si}} + C_{\text{TES}} + C_{\text{excess}}$, where $C_{\text{Si}} \propto T^3$ is the theoretical heat capacity of Si given by Debye theory, $C_{\text{TES}} \propto T$ is the TES heat capacity dominated by the metal electron system, and C_{excess} is the heat capacity of impurity bands and two-level systems in the crystal. The TES phonon system is assumed to be at the same temperature as the silicon phonon system, since the sub-micron thickness of the TES makes it incapable of sustaining its own thermal phonon




Figure 7-11: Schematic of bolometer model. The refrigerator acts as a cold bath at temperature T_b . The Si heat capacity is connected by a thermal conductance G_{pb} to the bath. The TES is connected by the electron-phonon conductance G_{ep} to the Si. A potential excess heat capacity with its coupling are shown in dashed outlines. For this study we have assumed C_{excess} and/or G_{excess} can be made small enough to become negligible.

population. The TES electron system is coupled to the phonon system through its electron-phonon coupling conductance G_{ep} .

There are two reasons for choosing a TES as the thermometer for this design. Firstly, in order to get the high energy resolution, the thermometer must be very well coupled to the absorber. The use of a Mo/Au TES uses the high electron-phonon coupling in Au to achieve this (G_{ep} in Fig. 7-11). Second, we want very good control of the the heat capacity of the bolometer. In a TES, the heat capacity is dominated by the electron system and is orders of magnitude above the TES phonon system's heat capacity. This makes the purity requirements on the TES from a heat capacity perspective fairly relaxed. Indeed, heat capacity measurements in TES devices routinely fall within expectations calculated from bulk elemental properties [223].

The measurements of the heat capacity of high-purity Si by Knaak and Meißner [224] suggest that our low required heat capacities are achievable. In our design, we benefit from the fact that our TES thermometer makes up about half of the total heat capacity, which allows us to tolerate some excess heat capacity from the silicon target. We also have the option to optimize using less mass per bolometer, and trading off

overall experiment mass for lower threshold per detector.

For this study, we will assume the excess heat capacity is negligible, and optimize the heat capacity of the thermometer C_{TES} , the electron-phonon coupling G_{ep} , and the Si heat capacity C_{Si} to obtain the desired 10 eV threshold with the highest possible target mass. We make the following assumptions:

- Each detector is a Si cube ranging in mass from 20–100 g. The heat capacity is determined from Debye theory.
- The conductance between the Si and the cold bath, G_{pb} , can be engineered to give a desired value. The value is chosen to give a thermal impulse response time of 50 ms as measured by the thermometer readout.
- The thermometer is a Mo/Au TES bilayer with a superconducting transition engineered to a specific temperature between 10–100 mK. Mo/Au TES X-ray detectors have achieved resolutions of $\Delta E_{\text{FWHM}} = 2$ eV [225].
- The TES heat capacity and electron-phonon coupling are taken from the literature and are a function of the chosen volume of the TES and the temperature.

Given these general assumptions, several combinations of detector mass and transition temperature were tested for both Si and Ge targets, scaling the TES volume to obtain the best energy resolution, following the theoretical framework of [69]. The TES volume is a compromise between two competing interests: having a small TES heat capacity, and having a fast thermal link between the TES and the Si or Ge target. The optimum volume corresponds to a TES heat capacity that is roughly equal to the target. An important quantity is the ratio $G_{\text{ep}}/G_{\text{pb}}$. As long as this ratio is $\gtrsim 100$, the TES remains in quasi-thermal equilibrium with the target throughout a pulse (except for the initial athermal phase on the order of 1 ms). If needed, one can make G_{pb} smaller and gain energy resolution (and lower threshold) at the expense of slower signals.

The results of our models are shown in Fig. 7-12. The plotted threshold is calculated as $7.5 \sigma_E$. Due to practical limitations in refrigeration and considering the readout necessary for the size of the experiment, we focus on a transition temperature of 15 mK, with the refrigerator base temperature at 7.5 mK. At this temperature, a low-energy threshold of 10 eV can be obtained with bolometers with 50 g of Si or 20 g of Ge. A 50 g Si target sees about twice the rate of neutrino coherent recoil events as a 20 g Ge target when both have a 10 eV threshold. We will thus focus on Si. Model parameters for the Si detector are given in Table 7.5.2.

The natural decay time of the bolometer $C_{\text{tot}}/G_{\text{pb}} = 436$ ms. Electro-thermal feedback [222] from the TES speeds up the response time to roughly 50 ms. The TES readout is bandwidth-limited by an inductor which critically damps the system, causing a further speedup in the response. The decay time of recoil events is reduced from the 50 ms decay with no inductor to a 30 ms decay with the inductor. Fig. 7-13 shows a simulation of 10–50 eV neutrino coherent scatters in a 50 g Si bolometer. The




Figure 7-12: Calculation of threshold for Si and Ge targets from 20–100 g at different operating temperatures. The lowest line for each target material is the 20 g line. The model (see Fig 7-11) takes into account the heat capacity of the TES and the target, the internal thermal fluctuation noise between the target and the TES thermometer, the electronics noise, the Johnson noise from the TES and its bias resistor, and the phonon noise between the target to the bath. The volume of the TES was scaled to give the best energy resolution at 15 mK. The horizontal dashed line marks the desired 10 eV threshold, corresponding to an energy resolution $\Delta E_{\text{FWHM}} = 3.14$ eV. The vertical dashed line marks the desired operating temperature of 15 mK. For Si, a 50 g target meets the requirements. For Ge, a 20 g target meets the requirements.

Parameter	Value	Units	Description
C_{Si}	43.3	pJ/K	Debye heat capacity
C_{TES}	31.1	pJ/K	TES electron heat capacity
G_{ep}	29.3	nW/K	TES-Si thermal conductance
G_{pb}	0.17	nW/K	Si-bath thermal conductance
T_b	7.5	mK	Cold bath temperature
T_c	15	mK	TES temperature
R_o	3	m Ω	Quiescent TES resistance
I_o	14.1	μA	Quiescent TES current
P_o	0.6	pW	Quiescent TES power
$\alpha = \frac{T_c}{R_o} \frac{dR}{dT}$	50	-	TES sensitivity
τ_o	436.2	ms	Natural decay time $C_{\text{tot}}/G_{\text{pb}}$
τ_{eff}	51.1	ms	Response time with TES speedup
τ_{decay}	29.2	ms	Decay time with readout circuit
L	30	μH	Readout inductance

Table 7.8: Model parameters for a 50 g Si target coupled to a Mo/Au TES operated at 15 mK. The Si target is a 28 mm cube, and the TES is an 25 mm \times 2 mm film 600 nm thick deposited on the Si surface. The energy resolution for this model is 3 eV FWHM, with a 10 eV threshold. Pulses from this model are shown in Fig 7-13.

pulses are clearly separated from the noise, and the energies of the different events are clearly separated by eye.

The total heat capacity is on the order of 460 keV/mK, and given a transition width of around 1 mK for a TES, we estimate that the bolometer will have fairly linear response up to hundreds of keV. Higher energies will have a non-linear response but will retain significant energy resolution. This large energy bandwidth will help understand the background in our experiments, and enable other rare event searches such as limits on the neutrino magnetic moment and dark matter interactions.

Multiplexing readout schemes for transition-edge sensors are now a mature technology being developed for many astronomical applications, for example [226], and 10,000 channel systems with time constants similar to this application are already in operation [227]. Schemes for even larger multiplexing gains are in development [228]. Given the slow time constants of this application, a 10,000 channel multiplexer design carries a fairly low risk and would allow 500 kg of Si to be instrumented.

A concept for a 500 kg payload is shown in Fig. 7-14. The 10,000 Si bolometers are arranged in a column of dimensions 0.42 (dia.) \times 2.0 (length) meters inside a dilution refrigerator suspended from a vibration isolation mount. Passive or active shielding surrounds the refrigerator. The exact shape and type of shielding will be determined at a later time. A cylindrical bore, perhaps 10 cm or less in diameter, is removed from the shield and allows the ^{37}Ar source, mounted on a radio-pure translation mechanism, to be moved to different positions along the side of the array. Periodic movement of the source throughout the measurement sequence allows each detector to sample multiple baselines, enables cross-calibration among detectors, and aids in




Figure 7-13: Simulated current readout for 10–50 eV recoils using the model parameters in Table 7.5.2. The current has been multiplied by -1 to make the pulses positive. Noise sources modeled are: the phonon noise between the target to the bath, the internal thermal fluctuation noise between the target and the TES thermometer, the Johnson noise from the TES and its bias resistor, and the electronics noise. The modeled 10 eV pulses are clearly separated from the noise.

background subtraction. The minimum distance from the source to a bolometer is assumed to be ~ 10 cm.

Detector Backgrounds

The detectors described in the previous section will be sensitive to several sources of background in the recoil energy range of 10-50 eV. Unfortunately, it is difficult to estimate accurately the rate of events from each of these sources, and the levels expected are currently unknown. Although we do not have a quantitative understanding of the backgrounds in this regime, it is important to note that backgrounds can be measured and subtracted using data taken when the neutrino source is not in place. We qualitatively consider several sources of background which we expect to be present in the energy range 10-50 eV:

- *Radiogenic impurities:* Most radioactive decay products have energies in the range of hundreds of keV to tens of MeV and will be clearly distinguishable from the neutrino signal. Radiogenic impurities may still contribute to the background in two primary ways. First, many common impurities produce gamma rays that can interact with material by the photoelectric effect or Compton scattering and produce background events by the mechanisms described below. These gammas commonly arise from the U and Th chains and also from ^{40}K and ^{60}Co . Second, electrons from beta decay isotopes, such as tritium, may have arbitrarily small energies and therefore can produce electron recoils in the




Figure 7-14: Conceptual schematic of the experimental setup for a bolometric measurement of coherent scattering from a high-intensity ^{37}Ar neutrino source. An array of 10,000 Si bolometers is arranged in a column of dimensions 0.42 (dia.) \times 2.0 (length) meters (shown in green) inside a dilution refrigerator suspended from a vibration isolation mount. Each Si bolometer has a mass of 50 g for a total active mass of 500 kg. Appropriate passive or active shielding surrounds the refrigerator. A cylindrical bore in the shield allows the ^{37}Ar source, mounted on a translation mechanism, to be moved to different positions along the side of the array. Periodic movement of the source throughout the measurement sequence allows each detector to sample multiple baselines, enables cross-calibration among detectors, and aids in background subtraction. The minimum distance from the source to a bolometer is assumed to be ~ 10 cm.

signal region. Nuclear recoils from decays at the detector surface, in which the electron is undetected, may also deposit small amounts of energy.

- *Compton scattering:* Photons from radioactivity and atomic transitions in the detector material or housing may Compton scatter once at a shallow angle in the detector. Since there is no discrimination between electronic and nuclear recoils, such shallow scattering would be indistinguishable from the neutrino signal. While the rate of these events is obviously dependent on the level of radioactive contamination, we expect kinematics to strongly suppress the rate of these events.
- *Photoelectrons:* Photons produced in the detector or housing may produce photoelectrons in the detector material, which could be ejected. Recoils from such events could produce small energy depositions in the energy region of interest. Low-energy secondaries from high-energy gammas produced in the detector or housing may impinge on other inactive material in the experiment and eject low-energy photoelectrons that could strike a detector.
- *Photons from atomic relaxation transitions from the surrounding surfaces:* Photons from atomic transitions are of roughly the correct energy to produce some background events near a 10 eV threshold in a Si detector. Copper, a good material for the detector housing for example, has 250 atomic lines with energies in the range 10-50 eV. The rate expected from such events is very difficult to quantify, and would depend on the amount of low-energy radiation present in the cryostat to excite these transitions.
- *Neutrons:* Conventional methods used to reduce and model the neutron flux in dark matter experiments can be used. Running the experiment at large overburden, a muon veto can be used to veto cosmogenic neutrons with high efficiency. The background rate of neutrons from muons that miss the veto can also be estimated with simulation. CDMS, for example, is able to achieve an unvetoes neutron rate of < 0.1 events/kg/year in the energy range 10keV to 100keV. Since the cross section for elastic scattering of neutrons on Si is fairly constant down to low energies, we do not expect the neutron background to be significant in our energy region of interest.
- *Neutrino-electron scattering:* In addition to scattering coherently off nuclei, the neutrinos will also scatter off electrons in the detector material. The cross section for this process is lower than the cross section for coherent neutrino scattering, and the recoil spectrum extends up to hundreds of keV for ^{37}Ar neutrinos. Since this background is also well-predicted by the standard model, we expect this to be a small contribution that may be reliably subtracted.

Excluding unknown backgrounds due to atomic transitions, the Compton scattering and photoelectron backgrounds are expected to dominate. Using the raw rate of events in CDMS, we can conservatively estimate the background rate due to these two sources. A good detector (250 g) in CDMS sees a raw rate of 0.001 Hz between

1–200 keV. If we conservatively assume that these events all lie in the range 1–10 keV and furthermore are attributable only to Compton scattering and photoelectric effect, then the rate of events is 38.4 events/kg/day/keV, assuming a flat spectrum. Assuming that the spectrum is flat down to 10 eV, we would see 1.54 events/kg/day in the energy range 10–50 eV. In reality, we do not expect the spectrum to be flat down to low energies. Low-energy gammas from k- and l-shell electron captures are much more likely to be absorbed by the photoelectric effect, causing their full energy to be measured in the detector, and suppressing their background at lower energies. Only a small fraction of these produce ejected photoelectrons which could leave small energy depositions in the detector as described above. Compton scattering is very suppressed in the 1–10 keV range, and only dominates above 60 keV. The fraction of the background seen in CDMS detectors that lies above 60 keV is very small. For these reasons, the figure of 1.54 events/kg/day between 10–50 eV is a very conservative estimate of the Compton and photoelectron background. For the sake of simplicity, we assume a 1 events/kg/day background in our energy window of interest.

Detector Calibration

Detector-to-detector variations constitute the most serious uncertainty in this measurement. Differences in the yield at different distance could quickly mask –or worse yield a false-positive on– the oscillation signal. Part of the problem can be solved by calibrating the efficiency of the detectors using a neutron calibration source. An attractive candidate would be to use a ${}^3\text{H}({}^3\text{H}, 2n){}^4\text{He}$ compact pulsable di-neutron source, such as used in the petroleum industry [229]. The neutrons produced from this source should yield a relatively flat neutron energy spectrum down to very low kinetic energies, which is ideal for studying detector acceptance and to verify the $1/r^2$ response of the array.

As shown in Figure 7-14, our approach is to remove the dependency of the oscillation measurement on the response of a particular detector. The source is placed on a movable platform, and moved along the Si array throughout the measurement. Over the course of the experiment, each detector samples multiple baselines, and can be cross-calibrated with other detectors at each baseline to remove the individual detector response differences. Detector variations are essentially constrained by the in-situ measurement.

It will also be important to calibrate the detector response to low energy photons and electrons. Recent advances in solid state UV diodes make well-tuned eV photon sources readily available. UV diodes ranging from 255 nm to 350 nm (3.5 eV - 5.2 eV) with sub-eV resolution are now commercially available.

7.5.3 Sensitivity and Outlook

Having discussed in detail both the source and the detector, we can now examine the signal in such an apparatus. For a monochromatic and isotropic source with activity $R(t)$ encapsulated in some volume V_S , the signal rate as a function of time t is given by the expression:

$$S(t) = \sum_i R(t) \cdot \sigma_0(E_\nu) \cdot f(E_\nu, T_0) \cdot \frac{N_A}{A} \cdot \rho_t \int \frac{dV_s}{V_s} \int \frac{P(E_\nu, r_{st})}{4\pi r_{st}^2} dV_{T,i} \quad (7.15)$$

where N_A is Avogadro's number, ρ_t is the target density, $dV_{T,i}$ is the differential target volume of a single detector, and r_{st} is the source-target distance. The sum is taken over all discrete detectors available for the measurement. In the approximation of a point source, Equation 7.15 reverts to the more familiar form:

$$S(t) = \sum_i R(t) \cdot \sigma_0(E_\nu) \cdot f(E_\nu, T_0) \cdot \frac{N_A}{A} \cdot M_t \cdot \frac{P(E_\nu, \bar{r}_i)}{4\pi \bar{r}_i^2} \quad (7.16)$$

where \bar{r}_i now is the average-weighted distance from the source to the individual detectors and M_t is the mass of each detector. In the limit where the measurement time is much greater than the source half-life, the total number of accumulated signal events is given by $N \simeq S(t_0) \cdot \tau_{1/2} / \ln 2$, where τ is the half-life of the neutrino source. Extending measurements well beyond the peak source activity has the added benefit of reducing the statistical uncertainty on the background.

For a monoenergetic source, the oscillation signal is all encoded within the spatial distribution of events. A deviation from the expected r^{-2} dependence could constitute a possible oscillation signal. For the case where there is only one additional neutrino, the oscillation probability is given by the neutrino oscillation formula:

$$P(E_\nu, r) = 1 - \sin^2(2\theta_S) \sin^2(1.27 \Delta m_S^2 \frac{r}{E_\nu}) \quad (7.17)$$

where $\sin^2(2\theta_S)$ is the amplitude to oscillate to the sterile state, and Δm_S^2 represents the sterile mass splitting. In this case, E_ν is measured in units of MeV, r in meters, and Δm_S^2 in eV^2 . For simplicity, we look at the simple 3+1 neutrino model, where the oscillation is to just one additional sterile neutrino.

We use simulated data from a mock experiment to determine the potential sensitivity to sterile neutrinos. We consider a compact 5 MCi ^{37}Ar source to be used in conjunction with a 500 kg silicon array. We consider a total exposure of 300 days in order to extract both signal and background rates. Parameters relevant for the fit are listed in Table 7.5.3. For comparison, we also list the parameters for a germanium array with similar number of detectors and energy threshold. Due to the lower mass per detector needed to achieve the lower threshold and the lower recoil energies of neutrinos off the heavier germanium nucleus, a germanium array will achieve a signal rate that is about half of the silicon array.

For such an experiment, we also consider a number of systematic errors:

- *Source Strength:* The SAGE collaboration used a variety of techniques in order to determine the final ^{37}Ar activity, including gas volume, gas mass, calorimetry, direct counting and isotopic dilution. Any one of these methods in isolation achieved a $\pm 1\%$ accuracy, while in conjunction the total uncertainty was less than $\pm 0.5\%$. In this study, we assume a conservative $\pm 1\%$ uncertainty on the

Parameter	Detector Type	
Detector Material	Si	Ge
Atomic Number	28	72.6
$\sigma_0(E_\nu)$ (10^{-42} cm ²)	0.44	3.82
T_{\max}	50.3 eV	19.4 eV
Threshold	10 eV	
$f(E_\nu, T_0)$ (see Eq. (7.13))	64.2%	23.6%
Detector cube size	28 mm	15.5 mm
Detector Mass	50 g	20 g
Number of Detectors	10,000	
Total Mass	500 kg	200 kg
Yield at 10 cm ($\text{kg}^{-1}\text{day}^{-1}\text{MCi}^{-1}$)	15.28	19.0
Signal Rate at 10 cm	3.82 day ⁻¹	1.90 day ⁻¹

Table 7.9: List of relevant source and detector parameters used for sensitivity analysis. The signal rate is quoted for a *single* detector located 10 cm away from the center of a 5 MCi (185 PBq) ³⁷Ar source.

source strength. Since the source uncertainty applies to all detectors globally, it has minimal impact on the oscillometry measurement.

- *Cross-section:* The cross-section uncertainty, much like the source strength uncertainty, is a global uncertainty and has little impact on our oscillometry extraction. Its uncertainty would nominally be dominated by the uncertainty in the form factor, but at such exchange momenta the effect is expected to be small. We therefore assume a $\pm 1\%$ global uncertainty due to the cross-section.
- *Vertex Resolution:* The bolometric detector in this experimental design is composed of 10,000 silicon or germanium absorbers instrumented with a single thermometer. The dimensions of these absorber cubes are 28 and 15.5 mm per side for silicon and germanium, respectively. These dimensions are smaller than the source itself (assumed to have a radius of 4 cm), thus the vertex resolution is dominated solely by the extension of the source. This effect is incorporated into our analysis.
- *Detector Variations:* Detector variations are kept under control via the series of in-situ and ex-situ calibration measurements discussed in the previous section. Using the movable source depicted in Fig. 7-14, one should be able to calibrate the detector variations to about $\pm 2\%$. The global uncertainty, which also depends on fiducial volume dependence, overall efficiency, etc., is estimated to be $\pm 5\%$.
- *Detector Backgrounds:* With detector-to-detector variations calibrated away, the main challenge for such a measurement remains the number of detector backgrounds that accumulate during the measurement. As discussed above, the

question of what the background will be between 10–50 eV is hard to estimate at this point, and more work needs to be done to enable a credible estimate. For this study, we assume a total background activity of 1 event/kg/day in the signal region of interest. The signal-to-noise ratio should scale roughly as the square root of the number of background events. The dependence of the accuracy of the measurement as a function of then signal-to-noise ratio is shown in Fig. 7-15. Measurements taken with background levels below 1 event/kg/day are essentially systematics dominated.

- *Source-Induced Backgrounds:* Any backgrounds that stem directly from the ^{37}Ar source may potentially dilute the sensitivity of the measurement, since they, too, would exhibit a $1/r^2$ behavior.

Though the majority of the energy from the decay of ^{37}Ar is removed by neutrinos, a fraction of the energy is carried away from recoils and internal-bremsstrahlung photons. The SAGE source effectively reduced this contribution to less than 0.2%. The electron capture process primarily produces gammas and Auger electrons at 2-3 keV. At this energy, the range for electrons in the continuous-slowing-down-approximation (CSDA range) in lead is $\sim 8 \times 10^{-5}$ cm [230]. For photons, the attenuation length is $\sim 2 \times 10^{-5}$ cm [96]. Assuming exponential attenuation in both cases, as little as 1 cm of ancient Pb would provide more than 10^4 attenuation lengths of shielding—more than sufficient to eliminate leakage of the gammas and Auger electrons expected from the source.

The SAGE group has produced an extremely pure argon source, with less than 0.4% of the volume having ^{39}Ar contamination. Mass spectrometry of the source found no significant amount of radioactive material besides ^{39}Ar and ^{37}Ar [214, 215]. Given the extremely long half-life of ^{39}Ar and the vastly different signature (β -decay), we believe this is a negligible background source. Consider, for example, an ^{39}Ar contamination of 1% in a 5 MCi ^{37}Ar source. If the source had a total specific activity of 92.7 kCi g^{-1} , as measured in SAGE, the activity of ^{39}Ar would be $\sim 26 \text{ Ci}$. The beta decay of ^{39}Ar has an endpoint of 565 keV, and the CSDA range for 500 keV electrons in Pb is $\sim 0.03 \text{ cm}$ [230]. Assuming exponential attenuation of electrons, about 1.5 cm of Pb would likely provide sufficient shielding to suppress all backgrounds from the ^{39}Ar . For our analysis, we assume that the contribution from source-induced backgrounds is negligible.

The situation is less favorable in a ^{51}Cr source. The decay of ^{51}Cr produces a 320 keV gamma from the decay of the excited ^{51}V daughter in 10% of decays. The attenuation length of gammas of this energy in Pb is 0.2 cm. To shield the entire flux of gammas from a 5 MCi source would require at least 10 cm of Pb.

- *Other Neutrino Interactions:* One of the isotopes of germanium (^{71}Ge) has a low enough threshold to allow ν_e charged current scattering. However, with a threshold energy of 321 keV, the outgoing electron will have recoil energy far above the energy region of interest. Hence, its contribution to the overall

background can be considered negligible. Charged current interactions on silicon all have thresholds above 1 MeV, hence they do not contribute to the background activity. As discussed previously, other charged-current interactions, such as $\nu_e e^-$ elastic scattering, are highly suppressed. As such, their contribution is also expected to be negligible.




Figure 7-15: Plot of the relative signal error versus signal-to-noise ratio $\frac{S}{\sqrt{B}}$ (S represents signal strength, B represents background counts) for a 500-kg Si array exposed to a 5 MCi ^{37}Ar for 300 days. This array configuration and source intensity yields approximately $S \simeq 54,000$ total signal events. Arrow indicates signal-to-noise ratio corresponding to 1 background event/kg/day.

A summary of the relevant systematic uncertainties are listed in Table 7.5.3. A simple χ^2 -fit is used to estimate the sensitivity of the proposed Si and Ge arrays. The data extracted from the entire array is first fit as a function of time in order to extract the overall source strength and background (see, for example, Figure 7-16). The background-subtracted signal is then fit to the oscillation formula of Eq. 7.17. The analysis uses both shape and rate to determine the sensitivity to sterile neutrinos. In the case of the shape+rate analysis, an additional penalty term is added to the likelihood from the overall flux measurement.

Results for a 500 kg Si detector array are shown in Fig. 7-18(a). The distortion caused by a non-zero sterile mixing is statistically distinguishable in the measured distance profile (see Figure 7-17(a)). As can also be seen from the figure, the array is not necessarily fully optimized for a given oscillation length scale. Such optimization can proceed once the parameter space for sterile neutrinos is further constrained by ongoing and future neutrino experiments. Nevertheless, for the bulk of the region of $\Delta m_{\tilde{S}}^2 = 1 - 10 \text{ eV}^2$ and $\sin(2\theta_S)^2 \geq 0.08$, typically preferred from the reactor data, is ruled out at the 90% C.L. If the best fit solution from the reactor anomaly is viable,

Source	Systematic	
	Global	Shape Only
Source Strength	$\pm 1\%$	-
Cross-section	$\pm 1\%$	-
Detector Variation	$\pm 2\%$	$\pm 2\%$
Absolute Efficiency	$\pm 5\%$	-
Source-Induced Background	$< 1\%$	$< 1\%$
Vertex Resolution	± 2.8 cm	± 2.8 cm
Source Extent	± 4 cm	± 4 cm
Total Systematic	$\pm 5.5\%$	$\pm 2\%$
Statistical (Whole Array)	$\pm 1\%$	

Table 7.10: List of systematic uncertainties expected for a low-threshold germanium detector array. Uncertainties are listed for both shape+rate and shape only analysis.

then the measurement should be detectable at the 99% C.L. (see Fig. 7-17(b)). It is possible to also conduct a shape-only analysis. Most of the sensitivity to sterile oscillations is retained for Δm^2_5 masses below 10 eV².

For comparison, we also consider an equivalent Ge array with a total mass of 200 kg. These results are shown in Fig. 7-18(b). Finally, for completeness we also show the detector sensitivity for the Ge and Si arrays using an equivalent ⁵¹Cr radioactive neutrino source (Figures 7-18(c) and 7-18(d)). In general, the reduced source energy decreases the available statistics, so a relatively stronger source needs to be considered in such a case.

7.5.4 Summary

We have outlined the possibility of probing the existence of sterile neutrinos using coherent scattering on a bolometric array. Such a method could provide the most direct test of oscillations to sterile neutrinos. With the advent of low threshold detectors and the use of intense neutrino sources, such an experiment appears feasible with our current technology. Such a program is also very complimentary to any existing dark matter search.

Even in the absence of sterile neutrinos, the experiment as described in this letter can make other important measurements. Most prominently, such an experiment may constitute the first observation of coherent scattering. For a 500 kg detector, it should be able to make a $\simeq 5\%$ measurement on the overall cross-section, pending on the absolute calibration of the efficiency. For an isoscalar target, such as silicon, this provides a direct measurement of the weak mixing angle at momentum transfer as low as 1 MeV.




Figure 7-16: Distribution of events on a 500 kg Si array as a function of time of source deployment. Source considered here is a 5 MCi ^{37}Ar electron-capture source.







Figure 7-17: Left: Ratio of data and Monte Carlo for a simulated neutrino oscillation signal ($\Delta m_S^2 = 1.5 \text{ eV}^2$, $\sin(2\theta_S)^2 = 0.15$) as a function of source distance from a 5 MCi ^{37}Ar neutrino source and a 500 kg Si-array. Right: Likelihood contour curves for same signal after 300 days of data taking. Contour levels are shown at 90% (blue), 95% (green), and 99% (red).




(a) Silicon - ^{37}Ar



(b) Germanium - ^{37}Ar



(c) Silicon - ^{51}Cr



(d) Germanium - ^{51}Cr

Figure 7-18: Likelihood contours for a 300-day run on a 500 kg Si array (left) and 200 kg Ge (right) array exposed to a 5 MCi ^{37}Ar (top) and ^{51}Cr (bottom) source, using both shape and rate information. Confidence levels in all plots are shown at 90% (blue), 95% (green), and 99% (red). Statistical and systematic errors are included in the signal analysis.

Chapter 8

Constraints on Sterile Neutrino Dark Matter

8.1 Introduction

A variety of dark matter models predict photon production via dark matter decay, annihilation, or de-excitation. Some of these models predict mono-energetic photons with energies in the 1-100 keV range, prompting recent searches for lines in existing X-ray data from the *XMM-Newton* and *Chandra* observatories. Due to the well-understood atomic physics in this energy range and the ability to check the morphology of a potential signal against expectations from galactic dark matter halos, X-ray lines could provide unambiguous evidence for some models of astrophysical dark matter.

There has been heightened interest in dark matter searches in the X-ray band following claims of an unidentified X-ray line seen in both galaxy and galaxy cluster observations. [7] analyzed *XMM-Newton* observations of 73 stacked galaxy clusters and found an excess line with energy around 3.56 keV. The line is present at the $> 3\sigma$ level in three separate subsamples of data from both the MOS and PN instruments, and they also detect it in *Chandra* observations of the Perseus cluster. [231] reported a $> 3\sigma$ excess around 3.53 keV in their spectral fits of *XMM-Newton* observations of the Perseus cluster and the Andromeda galaxy. In both cases, the width of the measured excess is determined by the *XMM-Newton* and *Chandra* instrument response. Analysis of *XMM-Newton* observations of the Milky Way Galactic Center (MW GC) by [3] finds a formal 5.7σ excess at the expected energy. However, the complexity of the GC makes modeling the background difficult, and because of the instrumental resolution of the observation they cannot rule out the possibility of the excess coming from K XVIII emission.

A vigorous search has ensued, with various reports of non-detections: [232] in the GC, [233] in the GC and M31, [234] in galaxies and galaxy groups, and [235] in dwarf spheroidal galaxies. There has been some debate as to how to best fit the continuum and of what the allowed flux of astrophysical lines (primarily from K, Cl, and Ar) in the pertinent energy range should be [7, 233].

[236] reproduce the line in Perseus using *Suzaku* but find its spatial distribution in tension with expectations for decaying dark matter, and further do not detect the expected scaled emission in either the Coma, Virgo or Ophiuchus clusters. [237] performed a morphological study of the continuum-subtracted excess emission at 3.5 keV in the GC and Perseus. They find the GC excess spatial distribution incompatible with the expected DM distribution, and strongly correlated with the morphology of atomic lines from Ar and Ca with energies between 3–4 keV. The Perseus emission is correlated most strongly with the cool core emission, confirming the tension presented in the *Suzaku* Perseus observation.

This set of observations demonstrates the challenges in searching for X-ray lines with the current observatories. New instruments are needed to improve the sensitivity of searches for X-ray line emission from dark matter. [6] studied the optimal characteristics of a mission dedicated to searches of diffuse line emission in the X-ray regime. They pointed out that the main determinants of instrument sensitivity are *grasp* $\equiv A_{\text{eff}}\Omega_{\text{FOV}}$ (effective area \times field of view, also referred to as *étendue*) and *energy resolution* $\Delta E/E$. As a “prototype” demonstration, Boyarsky et al. calculated the limits on the sterile neutrino mixing angle $\sin^2 2\theta$ for data from the third flight of the X-ray Quantum Calorimeter (XQC) sounding rocket payload [238] for sterile neutrino masses between 0.4–2 keV.

Existing X-ray telescopes tend to have comparatively small fields of view (e.g. a few tens of arcminutes for instruments on *XMM-Newton* and *Chandra*) with insufficient energy resolution to resolve closely spaced weak spectral lines (e.g. ~ 100 eV FWHM at 2 keV for the EPIC camera on *XMM-Newton*). Because of our location within the dark matter halo of the MW, the sterile neutrino decay is an *all-sky signal*, so sensitivity can be improved by increasing the FOV. Discrimination of a signal against atomic lines can also be significantly improved using the superior energy resolution available with X-ray microcalorimeters. The combination of large-FOV with high spectral resolution is achieved in existing microcalorimeter sounding rocket payloads. Although the exposure from a typical sounding rocket flight is less than 300 s, the sensitivity of these short observations can be competitive with deep *XMM-Newton* observations of the GC. The upcoming SXS microcalorimeter instrument onboard ASTRO-H [239] will have excellent < 7 eV resolution, but its narrow $3' \times 3'$ FOV limits its sensitivity to the all-sky signal expected from sterile neutrino decay in the MW. Wide-FOV sounding rocket observations are therefore complementary to the deep (~ 1 Msec) observations of the cores of galaxy clusters, galaxies, and dwarf spheroidals that ASTRO-H will perform [240].

In this paper we set limits on decaying sterile neutrino dark matter using a new dataset from the XQC sounding rocket in order to demonstrate the reach and analysis of large-FOV observations, and we discuss the sensitivity and optimization of future observations using new instruments, such as the Micro-X detector. Section 8.2 discusses the sterile neutrino signal and estimates the signal and background of large-FOV observations for a putative signal. In Section 8.3 we describe the XQC instrument, present an analysis of data taken during the 5th flight of XQC, and place limits on the sterile neutrino mixing angle $\sin^2 2\theta$ for sterile neutrino masses of 4–10 keV. Finally, Sections 8.4 and 8.5 estimate the sensitivity of observations with the future

Micro-X payload by constructing a detailed background model based on ROSAT observations and analyzing mock data sets.

8.2 keV Dark Matter with Rockets

8.2.1 Dark Matter Interpretations of X-ray Lines

Well-motivated dark matter models can produce X-ray lines through decay, de-excitation, or annihilation. Perhaps the best-known scenario is that of keV-mass sterile neutrinos [241–243], although a large number of models have been proposed following the observations of the 3.56 keV line. Such models include axions, axinos, exciting dark matter, gravitinos, moduli, and WIMPs, among others [233].

Sterile neutrinos and other models that produce photons by particle decay predict a flux per volume element that scales as the dark matter density (ρ). Other models, such as eXciting dark matter [244–246], that require two dark matter particles to interact predict a line flux per volume element that scales as the dark matter density squared (ρ^2). More complex scenarios, involving eXciting dark matter with a primordial population in an excited state [245], can furthermore produce fluxes that scale as ρ^α with $1 < \alpha < 2$. A non-linear scaling implies a much smaller flux from lower density systems like dwarf spheroidal galaxies and could ease the tension with non-observations of the 3.5 keV line in these systems [235]. In this paper we will use sterile neutrinos as our benchmark model, although we present our results as a line flux limit from a particular target, which can be translated into a constraint or signal in any of these models.

Sterile neutrinos with masses in the ~ 1 –100 keV range may contribute to the dark matter relic density if they are produced in the early universe. Two well-studied production mechanisms are non-resonant oscillation of active neutrinos [22] and resonant oscillation via the MSW effect [247]. The existence of sterile neutrinos is additionally motivated by neutrino oscillation data, which could be explained by adding sterile right-handed neutrinos to the standard model, such as in the ν MSM scenario [248]. Although they must possess cosmological lifetimes in order to contribute to the dark matter relic density, sterile neutrinos may decay to a photon and active neutrino via a loop-suppressed process mediated by oscillation between the active and sterile states. The rate for this process is given by [249]

$$\Gamma = \frac{9\alpha G_F^2 m_s^5 \sin^2 2\theta}{1024\pi^4} \quad (8.1)$$

$$= (1.38 \times 10^{-29} \text{ s}^{-1}) \left(\frac{\sin^2 2\theta}{10^{-7}} \right) \left(\frac{m_s}{1 \text{ keV}} \right)^5, \quad (8.2)$$

where m_s is the sterile neutrino mass and θ is the mixing angle between the active and sterile states.

Limits on $\sin^2 2\theta$ depend on both the observed flux and the fraction of dark matter comprised by sterile neutrinos. For simplicity, limits are typically quoted under the assumption that sterile neutrinos comprise all of the dark matter. The X-ray fluxes




Figure 8-1: Example of common DM halo profiles using NFW [1] (*solid*), Einasto [2] (*dashed*), and cored Burkart [1] (*dotted*) parameterizations.

reported by positive observations in stacked galaxy clusters, M31, Perseus, and the MW GC correspond roughly to $\sin^2 2\theta$ of 10^{-11} to 10^{-10} . [246] point out that with this level of mixing, the Dodelson-Widrow mechanism produces an insufficient abundance of sterile neutrinos to comprise all of the dark matter; however, a slightly larger mixing angle and smaller sterile neutrino fraction of the dark matter can resolve this tension.

8.2.2 Dark Matter Signal for Large FOV Observations

The flux expected from decay of sterile neutrinos in the MW halo is proportional to the integral of the DM density along the line-of-sight and over the field of view

$$\mathcal{F} = \frac{\Gamma}{m_s} \frac{1}{4\pi} \int_{FOV} \int_0^\infty \rho(r(\ell, \psi)) d\ell d\Omega, \quad (8.3)$$

where in the integral of the dark matter profile density $\rho(r)$, the parameter ℓ is the distance along the line of sight, r is the distance from the GC, and the angular integral is taken over the field of view of the instrument. The distance from the GC is related to the line-of-sight distance by

$$r(\ell, \psi) = \sqrt{\ell^2 + d^2 - 2\ell d \cos \psi}, \quad (8.4)$$

where d is the distance of the earth from the GC and ψ is the opening angle from the GC.

Sounding rocket flights observe for a few hundred seconds, and the effective areas for the instruments considered in this paper are on the order of 1 cm^2 . *XMM-Newton* has made observations on the order of a megasecond, with an effective area at 3.5 keV




Figure 8-2: Ratio of sterile neutrino decay signal in a large FOV to a 14'-radius FOV (*XMM-Newton*) around the galactic center as a function of the FOV half-opening angle. Orange curves are for an observation centered at the GC, while blue curves are for an observation centered on the XQC field of $(b, l) = (165^\circ, -5^\circ)$. Line style corresponds to different DM profiles: NFW [1] (*solid*), Einasto [2] (*dashed*), and cored Burkart [1] (*dotted*) parameterizations. The Micro-X FOV is given by the dashed vertical line, while the dotted vertical line shows the FOV of XQC.

of around 200 cm^2 for each MOS detector. Sounding rockets, however, can observe a much larger FOV than *XMM-Newton*. In order to compare between different FOV observations, a DM halo must be assumed, and we show several representative profiles in Fig. 8-1. In Fig. 8-2 we show in orange the ratio of the expected sterile neutrino decay rate from the central 14' radius of the GC (*XMM-Newton*'s FOV) to a larger FOV also centered on the GC. In blue we show the ratio between the *XMM-Newton* GC observation and a different field near the MW anti-center at Galactic coordinates $l = 165^\circ, b = -5^\circ$, observed by the 5th flight of the XQC (discussed in the next section). With a sufficiently large FOV, signal rates 3 to 4 orders of magnitude larger than the *XMM-Newton* GC observation are attainable. Furthermore, the higher energy resolution of the microcalorimeter instruments in sounding rockets can cut the continuum background per energy bin by over an order of magnitude when compared to the CCD energy resolution of *XMM-Newton*.

To get a feel for the potential sensitivity of these observations, consider a hypothetical instrument with a 1 cm^2 effective area, 3 eV FWHM resolution, and 20° radius FOV at 3.55 keV, which observes the GC for 300 s (see Table 8.1). If the flux reported by [3] over a 14' radius FOV is due to decaying sterile neutrino dark matter with the NFW profile of Fig. 8-1, then our hypothetical 20° field would expect a scaled signal flux of $6.1 \times 10^{-2} \text{ cm}^{-2} \text{ s}^{-1}$, 2000 times higher than the *XMM-Newton* observation. A 300 s observation would measure 17.8 total events in the X-ray line. At 3.5 keV, the background model described in Section 8.4.1 predicts a flux of $4.5 \text{ cm}^{-2} \text{ s}^{-1} \text{ keV}^{-1}$.

reference flux (in 14' of GC)	$2.9 \times 10^{-5} \text{ cm}^{-2} \text{ s}^{-1}$
scaled flux (in 20° of GC)	$6.1 \times 10^{-2} \text{ cm}^{-2} \text{ s}^{-1}$
effective area at 3.55 keV	1 cm ²
exposure time	300 s
resolution (FWHM)	3 eV
signal events (in 20° of GC)	18.2
bg. rate at 3.55 keV (see §8.4.1)	$4.5 \text{ cm}^{-2} \text{ s}^{-1} \text{ keV}^{-1}$
bg. events in signal window	6.7
median signal significance	5.6 σ

Table 8.1: Basic signal and rates expected for a hypothetical observation of the GC using a microcalorimeter on a sounding rocket, assuming a fiducial signal flux from [3].

The background in a 5.1 eV ($\pm 2\sigma_E$) window would be 6.7 events. In spite of the small statistics, the median significance of the putative signal above the continuum background would be 5.6σ in this short observation.

For comparison, in the [3] analysis of 1.4 Ms of *XMM-Newton* data we estimate around 7,500 signal counts in the claimed 3.54 keV line in each MOS detector. In that same resolution element, there are upwards of 500,000 background counts. With a signal-to-noise ratio of 0.015, the authors use the *XMM-Newton* high statistics measurement to detect such a small signal at high formal significance (5.7σ), but doing so depends on an accurate model of their background and minimal systematic errors.

8.3 Analysis of XQC Data

Having laid out the basic strategy, we now focus on existing data from X-Ray Quantum Calorimeter. The XQC payload is a mature flight system with 6 flights between 1995 and 2014 [250]. The XQC spectrometer is an array of 36 microcalorimeters using ion-implanted semiconductor thermistors each coupled to a $2 \text{ mm} \times 2 \text{ mm} \times 0.96 \text{ }\mu\text{m}$ HgTe absorber on a $14 \text{ }\mu\text{m}$ -thick Si substrate, with total area of 1.44 cm^2 . The energy resolution below 1 keV is 11 eV FWHM, although due to position dependence the resolution degrades to 23 eV FWHM at 3.3 keV. The microcalorimeter array is mounted inside a cryogenic system which uses pumped He as a 1.5 K bath for an Adiabatic Demagnetization Refrigerator (ADR), which is coupled to the detector assembly and cools it to $\sim 50 \text{ mK}$ temperatures. To survive the launch vibrations while cold, the cryogenic system is suspended with vibration insulators from the skin of the rocket, and the resonant frequencies of the system are designed to minimize coupling of skin vibrations to the detectors during launch. The FOV of XQC is 1 sr, subtending a 32.3° radius in the sky.

Data from an observation centered at the Galactic coordinates of $l = 90^\circ, b = 60^\circ$ during the 3rd flight of XQC were first presented by [238]. [6] used this data to




Figure 8-3: Spectrum of XQC data overlaid with fitted total background model (*solid blue*). Dashed lines show background model components, consisting of a power law continuum from the diffuse X-ray background (*dashed cyan*), a power law continuum from the Crab (*dashed red*), cosmic rays (*dashed green*), and lines from the ^{41}Ca calibration source onboard the instrument (*dashed orange*). The calibration source produces lines at 3.31 keV and 3.59 keV from $K\alpha$ and $K\beta$ transitions of potassium, while the two broad peaks at lower energies are due to $K\alpha$ and $K\beta$ X-rays that interact in the Si substrate of the HgTe absorbers and experience energy losses due to charge trapping. The flat continuum visible below 2.5 keV in the calibration spectrum is due to source events in which the photoelectron escapes the absorber. The bottom panel shows residuals between the data and the total background model. Only 25% of events above a threshold of ~ 1.8 keV are saved in the data stream because of bandwidth constraints. The gradual turn-on of this threshold makes the efficiency difficult to determine in the 1-2 keV range, so we restrict the analysis to energies above 2 keV.

constrain the decay of sterile neutrino dark matter, and their results are reproduced in Figure 8-11. Their analysis did not perform background subtraction and was limited to data below ~ 1 keV.

We perform a new analysis that develops a background model for the data between 2.0 keV and 5.0 keV, and then uses the data and background model to constrain the flux of an unidentified line in this interval. The use of background subtraction and higher-energy data from a more recent flight of XQC are the main improvements over [6]. We analyze a partial data set from the fifth flight of the XQC rocket, which flew 2011 November 06 at 08:00 UT as flight 36.364UH from the White Sands Missile Range. It obtained about five minutes of on-target data at altitudes above 160 km. The field of view was centered at the Galactic coordinates of $l = 165^\circ, b = -5^\circ$, close to the galactic anti-center. A total of 200 s of on-target data was analyzed on 29 functional pixels. After a very conservative quality cut to remove pixels and time periods with unstable event rates, 2551 pixel \cdot s remain on 24 pixels, for an effective exposure of 106 s per pixel. Data from other XQC flights is also being reprocessed, and the combination of these data sets will increase the total exposure by a factor of a few in a future analysis.

Fig. 8-3 shows the XQC data above 2.0 keV. The spectrum contains a power law continuum with strong lines at 3.31 keV and 3.59 keV from $K\alpha$ and $K\beta$ transitions of potassium, respectively. These lines arise from a ^{41}Ca source which provides continuous calibration during the flight, and is used to correct gain fluctuations. X-rays incident on XQC may be absorbed in either the HgTe absorber or its Si substrate. The photons absorbed in the HgTe are efficiently thermalized, while those absorbed in the substrate experience an energy loss of about 15-20%. Potassium $K\alpha$ and $K\beta$ events absorbed in the substrate form the two broad peaks centered at 2.80 keV and 3.0 keV, below the corresponding lines due to absorption in the HgTe. The relative intensity of the full-energy peak in the absorber and the second peak from events in the substrate is determined by the relative absorption efficiencies for X-rays in the two detector elements, shown in Fig. 8-4. The efficiency of the Si rises rapidly above 1.0 keV, and becomes comparable to the HgTe efficiency above 3.5 keV.

The overall strategy of the analysis is to perform an exclusion on the rate above background of an unidentified line centered at each energy between 2.0 keV to 5.0 keV. To do this, we first fit a background model to the data incorporating the important spectral features described above. We then use the model to estimate the background in a sliding energy window, allowing us to set an upper limit on the expected flux from an unidentified line, as a function of energy. Modeling of the energy spectrum becomes more complex at energies below 2.0 keV due to weak atomic lines from thermal emission. At energies significantly above 5 keV the energy scale could become nonlinear and the detection efficiency drops due to saturation. For simplicity, we restrict to a conservative energy window of 2.0-5.0 keV, although this range could likely be expanded in future analyses.

The background model and its components are shown in Fig. 8-3. The model incorporates the lines from the ^{41}Ca source, a continuum from events in which photo-electrons escape the absorber, a power law continuum from the Crab [251], a power law continuum from the diffuse X-ray background [252], and component from cos-




Figure 8-4: Efficiency of XQC for detecting X-rays as a function of energy. Curves show the efficiency of absorption in the HgTe absorber (*dashed orange*), and the efficiency of absorption in the Si substrate (*solid blue*). Events absorbed in the Si substrate lose 15-20% of their energy due to charge trapping, but otherwise appear as good pulses (see text for discussion).

mic rays. The power law models are properly corrected for the suppressed energy measurement when X-rays interact in the Si substrate, as well as for the efficiency of X-ray detection shown in Figure 8-4. Fits are performed using the RooFit software package based on the Minuit numerical minimizer [253]. We use the method of extended unbinned maximum likelihood [254], which is more stable than binned fits in a low-statistics setting. More details on the model and statistical methodology are provided in the Appendix.

For each energy E_0 , we construct a window $[E_0 - 2\sigma, E_0 + 2\sigma]$ of two standard deviations in energy resolution which we use to set upper limits on any flux above the modeled background. The background model is fit to all data outside this signal window, and then extrapolated into the window. The background model in the window is then integrated to obtain the background rate b with uncertainty σ_b propagated from the uncertainty on the fit parameters. An upper limit is then set on the rate of signal above background in the window, for a Poisson process. Limits are set using the profile likelihood test statistic described in [255], which incorporates the uncertainty on the background rate in the window. The critical value of the test statistic for the desired confidence level is calculated exactly from Monte Carlo simulations rather than using the asymptotic distribution of the test statistic. This is important at higher energies, where the low statistics cause the test statistic to differ from its asymptotic distribution. The upper limits are set using the RooStats software package [256]. Although the “sliding window” approach does not exploit the signal and background shape within the energy window, there is little loss of information because of the very low statistics of a potential signal.

Fig. 8-5 shows the limits on the flux of an unidentified line as a function of the line energy, calculated using the limit-setting procedure described above with the




Figure 8-5: Upper limit at 95% CL on the flux of an unidentified line in the XQC spectrum (*black line*), as a function of the line energy. Bands show the $\pm 1\sigma$ (*green band*) and $\pm 2\sigma$ (*yellow band*) range of the expected limit from the background-only hypothesis. Inset shows the energy range around the putative 3.5 keV line of [3], referred to the XQC field of view using the NFW profile from Fig. 8-2.

Data	DM profile	Flux [$\text{cm}^{-2} \text{s}^{-1}$]
XQC (this work)	N/A	$< 17 \times 10^{-2}$ (95% CL)
Expected line flux in	NFW	$(2.03 \pm 0.35) \times 10^{-2}$
XQC field scaled from [3]	Einasto	$(1.95 \pm 0.34) \times 10^{-2}$
	Burkart	$(11.7 \pm 0.2) \times 10^{-2}$

Table 8.2: Flux limits on a line at 3.53 keV from XQC data in this work, compared with the expected flux in the XQC field from scaling the result of the galactic center observation of [3], using the same fiducial DM profiles as Fig. 8-2.

background model. At 3.53 keV, we set an upper limit on the flux of an unidentified line of $0.17 \text{ cm}^{-2} \text{s}^{-1}$ at 95% CL. The flux reported in the GC by [3], when referred to the XQC field using fiducial DM profiles is listed in Table 8.2. The inset of Fig. 8-5 also shows a strong downward fluctuation of the limit at 3.55 keV, consistent with the lower edge of the $\pm 2\sigma$ band of expected limits. While this fluctuation may be purely random, it could also be caused by a very slight difference in the lower tail of the $K\beta$ line between the flight data and the calibration template constructed using ground calibration data. Future flights of XQC would clearly benefit by choosing a calibration source with lines further from the signal region near 3.5 keV. Although the XQC data do not exclude the Boyarsky best-fit flux, the upper limit is close in the case of a cored DM profile, such as Burkart. It is furthermore important to emphasize that XQC achieves this result with merely ~ 106 s of data on only a subset of its pixels. This underscores the value of combining additional datasets from other flights: additional statistics will both improve the limit and enable more robust background modeling.

8.4 Estimates for Future Observations

The XQC limits shown in the previous section are photon-limited; more observation time will result in better sensitivity. We also plan future observations in several fields including the GC to both increase sensitivity to potential undiscovered sterile neutrino lines and provide a definitive test of the nature of the line found by [3]. These future observations will benefit from higher resolution microcalorimeter arrays, such as those developed for the Micro-X rocket payload.

The Micro-X payload is a new system based on the XQC design, with new detectors and readout to allow for a larger array of higher-resolution microcalorimeters [257]. Although Micro-X was designed to be used with a 2.1 m X-ray optic, for dark matter searches it will be reconfigured to fly without it, using an optical stop like XQC to set its field of view. The Micro-X detector consists of an array of 128 microcalorimeters using Transition-Edge Sensor (TES) thermometers each coupled to a $0.6 \text{ mm} \times 0.6 \text{ mm}$ BiAu absorber. For the dark matter flight, new absorbers with 0.9 mm per side and thickness of $3 \mu\text{m}$ (Bi) + $0.7 \mu\text{m}$ (Au) will be used, with total area of 1 cm^2 . The absorption efficiency of X-rays in the the detector is shown in




Figure 8-6: Total efficiency of Micro-X for X-ray detection as a function of energy. The total area of the detector is 1 cm².

Fig. 8-6. The BiAu absorbers allow for better thermalization and should minimize position dependence, allowing Micro-X to retain its design resolution of 3 eV FWHM at 3.5 keV.

In this section we estimate the sensitivity of a potential GC observation with the Micro-X payload. A GC observation is the most direct comparison to [3], and the higher energy resolution of the Micro-X instrument results in a lower background from continuum emission and better discrimination between unexpected lines and those coming from atomic transitions in the observed plasma. Micro-X is designed to be coupled to a mirror, so for this calculation we use a 0.38 sr, 20° radius FOV as an estimate of the achievable FOV without the Micro-X optics. A future redesign of the optical aperture of the cryostat could increase the FOV to 1 sr.

8.4.1 Backgrounds

The first step in estimating the sensitivity of a potential MW GC observation with Micro-X is constructing a background model of the complex emission from the large FOV. For this we have used data from the *ROSAT* Survey Diffuse Background Maps to obtain the contribution from diffuse emission, and *ROSAT* Bright Source Catalogue [258] to estimate the contribution from point sources in the field. The total background model is shown (in black) in Figure 8-7 and discussed below.

GC Diffuse Background

The *ROSAT* data for finding the diffuse component of the backgrounds was accessed via the HEASARC X-ray Background Tool v2.5. This tool provides the integrated count rate over regions of the sky (up to a maximum radius of 10°) in various energy bins seen by the *ROSAT* PSPC instrument. The assumed FOV for a future Micro-X GC observation has a radius of 20° while the X-ray Background Tool can only select areas with a maximum radius of 10°, so we combined data using a tiling of fields.




Figure 8-7: Expected X-ray background from a 20° radius region centered on the Galactic Center. The total background spectrum is shown in black, the expected emission from the brightest low mass X-ray binaries (as shown in Table 8.3) in green, and the power-law and thermal components from the diffuse background emission are shown in red and blue respectively.

The tiling pattern consists of a 10° cone in the center, surrounded by 9 regions with radii of 5° to fill in the the surrounding annulus, and 9 more regions each with radii of 1° and 1.659° to fill in the gaps. The outer part of this set of observation regions covers 86.3% of the outer annulus, and all regions together cover 89.7% of the FOV. To make up for the missing regions, we take the mean flux in the sampled regions of the outer annulus and assume that it matches the mean flux for the entire annulus. For each line of sight, we convert the observed count rate in the *ROSAT* band into a photon flux using the WebPIMMS portable mission simulator, also from HEASARC, and combining them we obtain the total photon flux in the ROSAT band from all diffuse emission and unresolved point sources in the FOV of interest.

The spectral model is constructed using the components from observations of the Galactic Center (GC) and Galactic Ridge (GR) taken with Suzaku [259]. The integrated spectrum is the sum GC, GR and emission from latitudes $b \geq |4.1^\circ|$. Each of the Galactic components is a combination of emission from two thermal plasmas with different temperatures and metallicities, and a powerlaw component to account for the cosmic X-ray background (CXB), all absorbed through the appropriate column densities. The thermal plasmas are simulated in XSPEC using the *APEC* model from AtomDB version 3.0.1 [260, 261]. In addition to the high temperature emission from the GC and GR regions, we have added 2 thermal plasmas with low T and low metallicity ($Z = 0.05$) which accounts for common components on the entire Galactic plane. The CXB contribution from directions off the Galactic plane is also added as an absorbed powerlaw. All parameters are taken from the work by [259], and the overall normalization is adjusted to match the photon flux observed with ROSAT. The total thermal emission is shown in red in Fig. 8-7, and the integrated CXB emission in presented in blue.

Background from Bright Sources

The RASS-BSC includes 558 sources within 20° of the GC, which have been left out of the RASS diffuse emission maps. The 12 brightest sources within the region of interest, all low mass X-ray binaries (LMXBs), account for 80% of the total count rate from the BSC, and we have listed them in Table 8.3. Their spectra is modeled using the emission and absorption parameters obtained through X-ray observations. Note that the observed flux from these sources will depend on their actual state at the time of observation, so there is some inherent uncertainty in this background estimate. Since the remaining portion of the bright source contribution is dominated by LMXBs, we scale the total spectrum from the top 12 sources to match the total count rate as observed by *ROSAT*. The integrated BSC spectrum is shown in green in Fig. 8-7.

GC Combined Background

Fig. 8-7 shows the total background model for a 20° radius field centered on the GC, which includes the diffuse and point source emission discussed above. We show the total flux in black, the expected emission from bright sources (as shown in Table 8.3)




Figure 8-8: Background from a 20° radius region centered on the GC (as in Fig. 8-7), in the 3.3-3.7 keV range. The total background spectrum is shown in black, emission from resolved bright sources (from the RASS-BSC) in green, and the CXB and thermal components from the diffuse background emission are shown in red and blue respectively. Emission lines from ions Ar XVIII, S XVI, Cl XVI, Cl XVII, K XVIII, and Ar XVII are also shown. The abundance of Cl and K have been set to solar values, using the tables from [4].

Table 8.3: LMXBs included in the background model for the Micro-X observation.

ROSAT Name	Associated Name	Count Rate (counts/s)	N_H (10^{22} cm^2)	$F_{0.5-2 \text{ keV}}$ ($10^{-9} \text{ ergs/s/cm}^2$)	$F_{2-10 \text{ keV}}$	Reference
1RXS J173143.6-165736	GX 9+9	144.5	0.15	1.3448	4.5127	[262]
1RXS J182340.5-302137	4U 1820-30	127	0.078	1.3219	5.188	[263]
1RXS J170544.6-362527	GX 349+02	91.82	0.7	1.8785	12.075	[262]
1RXS J173858.1-442659	4U 1735?44	61.88	0.185	1.1257	4.5544	[262]
1RXS J175840.1-334828	4U 1755-33	31.08	—	—	—	[264]
1RXS J180132.3-203132	GX 9+1	30.92	1.447	0.71644	19.997	[265]
1RXS J173602.0-272541	GS 1732-273	29.79	0.67	0.82566	1.2021	[266]
1RXS J181601.2-140213	GX 17+2	24.82	3.18	0.6595	15.843	[267]
1RXS J170855.6-440653	4U 1705?44	20.37	1.5	0.5785	6.4055	[262]
1RXS J174755.8-263352	GX?3+1	19.2	1.7	0.42601	4.4799	[268]
1RXS J180108.7-250444	GX 5-1	17.78	2.8	0.6671	41.512	[269]
1RXS J173413.0-260527	KS 1731-260	14.09	1.08	0.40815	2.9953	[270]

in green, and the power-law and thermal components from the diffuse background emission are shown in blue and red respectively. Fig. 8-8 shows these different components but zooms in on the 3.3 keV to 3.7 keV region (line colors are consistent with those of Fig. 8-7). The continuum in this energy band is dominated by the emission from the BSC LMXBs, yet some thermal line emission is also significant. The relevant ions in this energy range are Ar XVIII, S XVI, Cl XVI, Cl XVII, K XVIII, and Ar XVII. We used the abundances table of [4] and set the abundance of Cl and K to those values.

Other Backgrounds

In a realistic flight, Micro-X requires an on-board calibration source similar to the one that produces the potassium lines in Fig. 8-3. The ^{41}Ca source used by XQC is not ideal for searching for a line at 3.5 keV because the $\text{K}\beta$ line is at an energy similar to the signal. We consider an alternate calibration source consisting of an ^{55}Fe source that produces fluorescence X-rays by illuminating a NaCl wafer. A kapton filter could block Auger electrons from the NaCl and the ~ 1 keV X-rays from Na fluorescence, leaving only the $\text{K}\alpha$ (2.62 keV) and $\text{K}\beta$ (2.82 keV) lines from Cl fluorescence, as well as a small number of back-scattered ^{55}Fe X-rays. We simulate the energy spectrum observed by Micro-X using Geant4 [8], and add it to the astrophysical X-ray spectrum, assuming a calibration source rate of 1 Hz / pixel.

Cosmic rays will also produce a background in Micro-X. Most cosmic ray primaries are protons with energies around a few GeV. These act as minimum-ionizing particles producing a broad continuum of energies in the detector, but peaking in the signal region around 3-4 keV. We simulate the cosmic ray energy spectrum using Geant4, and add this as a background component in the background spectrum. With a total rate of about 1 Hz in MicroX, cosmic rays comprise a very small fraction of the total background.

8.4.2 Signal and Sensitivity Estimates

A mock observation of the GC with and without the [3] line is shown in Fig. 8-9. A candidate line of this strength in the Micro-X observation would also be stronger relative to atomic lines than in *XMM-Newton* and *Chandra* observations. In fact, since the background model for the GC is dominated by a blackbody continuum from low-mass X-ray binaries, no significant atomic lines are expected to be visible above the continuum in the energy range of interest. The strongest lines between 3.5 and 3.6 keV come from K XVIII and Cl XVII with expected counts of less than 1 event in the Micro-X observation, so if a line of this strength was detected at this energy it would be in strong tension with a “standard” astrophysical origin.

Using the background model and the methodology of our analysis of XQC data, we can estimate the sensitivity to an unidentified line over the entire energy range available in observations by Micro-X. We consider measurements centered on the GC. While the Micro-X constraints are qualitatively similar to those of XQC, they are more stringent because of the significantly higher spectral resolution of Micro-X.




Figure 8-9: Mock data in the energy range of interest for the 3.5 keV line, with (*red*) and without (*black*) a signal, with background model (*blue line*) and signal model (*dashed green line*) overlaid. Note that the excellent spectral resolution of Micro-X provides significant separation of a signal line from nearby atomic lines.




Figure 8-10: Expected limit from an observation of a 20° field around the GC by Micro-X. Black dashed line shows the median expected 95% CL upper limit, while the green and yellow bands are the $\pm 1\sigma$ (*green band*) and $\pm 2\sigma$ (*yellow band*) ranges of the expected upper limits. The expected upper limit rises at high energies because of the falling efficiency to detect x-rays. The red point is the flux of [3], extrapolated to the Micro-X field of view using an NFW profile. Gray bands overlay strong calibration source lines.

Using the same analysis approach that we applied to XQC, we compute the expected upper limit on the flux of an unidentified line as a function of energy, under the background-only hypothesis. The resulting limit is shown in Fig. 8-10. Fluctuations of the limit at low energies are due to the presence of atomic lines. At higher energies, large jumps in the limit are caused by the small number of background events in the signal region, while small fluctuations are due to finite statistics in the Monte Carlo simulation used to set the upper limit.

8.5 Sterile Neutrino Interpretation

The flux limits obtained from XQC data and projected for Micro-X can be translated into constraints on models of dark matter. Although the literature contains a range of models that could produce an X-ray line, we consider a decaying sterile neutrino as a benchmark model because it has been extensively discussed as the source of the 3.5 keV excess. Using equations (8.2) and (8.3) for the flux of a decaying sterile neutrino, we translate the limits of Fig. 8-5 and Fig. 8-10 into limits on the sterile neutrino mass m_s and mixing angle $\sin^2 2\theta$. In addition to showing the limit expected from an observation of the galactic center, Fig. 8-11 also shows the limit expected from a Micro-X observation of the field of view observed by XQC. For this projection, we take the background to be the diffuse cosmic X-ray background measured by the XQC observation in Fig. 8-3, scaled to the smaller Micro-X field of view, and without the Crab component. The Crab can easily be avoided with the smaller field of view of Micro-X.

Although the XQC limit is not strong enough to provide a robust exclusion of the parameters inferred by [7], a Micro-X observation of the GC could provide a significantly stronger constraint because of its better energy resolution and larger exposure, and because of the larger signal strength in the GC. These limits obviously depend on the structure of the dark matter halo, with cored profiles producing significantly stronger constraints in the XQC field than NFW-like profiles. In addition, the wide-FOV observations are complementary to narrow-FOV observations of distant galaxies and clusters because they directly address whether an unidentified line is present as an all-sky signal in the MW. This confirmation is crucial for distinguishing an atomic interpretation from an exotic DM one.

8.6 Conclusion

Microcalorimeters onboard sounding rockets have the ability to place competitive bounds on keV sterile neutrinos or other dark matter models whose flux scales linearly with dark matter density. We have analyzed a subset of the data acquired during the 5th flight of the XQC payload corresponding to an effective exposure of 106 s on 24 pixels and placed a upper limit on keV sterile neutrinos between 4-10 keV which demonstrates the prospect for future observations with this type of instrument. A study of a future observation of the Galactic center with the Micro-X payload would




Figure 8-11: Constraints on decaying sterile neutrino dark matter, assuming that sterile neutrinos comprise all of the DM in the NFW profile of [1]. Limits include the XQC observation analyzed in this work (*black*), Micro-X median expectation from an observation of the XQC field (1), Micro-X median expectation from an observation of the GC (2), constraints from M31 [5] (*shaded orange*), and constraints from the previous analysis of XQC data by [6] (*gray*). The putative signal of [7] is also shown (*red point*).

have sensitivity to new parameter space in the $(m_s, \sin^2 2\theta)$ sterile neutrino plane. Optimizations of the pointing direction, FOV, and resolution, along with repeated observations would all increase the sensitivity of this technique in the future.

8.A Statistical Model for XQC

We use the method of unbinned extended maximum likelihood to fit the XQC data. The unbinned method produces equivalent results to binned χ^2 fits in the limit of high statistics, but produces more reliable fits in low statistics settings where many bins would have few or zero events [109]. The likelihood function for our background model has the form

$$\mathcal{L}(\mathbf{S}; \{E_i\}) = \frac{e^{-\mu} \mu^N}{N!} \prod_{i=1}^N \left[\sum_{k=1}^7 \frac{\mu_k}{\mu} P_k(\mathbf{S}_k; E_i) \right], \quad (8.5)$$

where the product is taken over the N total events in the observation, and the sum is taken over each of 7 components of the background model. The values μ_k are the estimated number of events in each of the background components and $\mu \equiv \sum_k \mu_k$. The PDFs for each component are the $P_k(\mathbf{S}_k; E_i)$, which are functions of energy and depend on the vector of parameters \mathbf{S}_k (with $\mathbf{S} = \cup_k \mathbf{S}_k$). The functional forms for each PDF are listed in Table 8.4, and the corresponding parameters are in Table 8.5. Background PDFs which have a fixed template shape, such as the cosmic rays, do not have any parameters, so \mathbf{S}_k is an empty set. The Poisson term (first) is the extended likelihood term, which constrains the total expected event rate by the number of observed counts.

The components of the background PDF are summarized in Table 8.4, and the key model parameters are contained in Table 8.5. The calibration lines from interactions in the HgTe absorber are modeled using a gaussian kernel density estimate (KDE) based on calibration data taken in a lab after launch. An alternate model using Voigt profiles for each of the calibration lines also produces a reasonable fit, but the KDE-based model was chosen because it agrees better in the low-energy tails of the calibration lines. A similar KDE does not accurately describe the corresponding lines in the substrate, around 2.80 keV and 3.0 keV. These are more reliably modeled by gaussian PDFs with fitted means and a common fitted width. Events that interact in the substrate have suppressed energy because of charge trapping effects that depend on the neutralization state of the Si, so differences between flight and calibration data are not surprising. Lastly, the photoelectron produced by an X-ray interaction in the HgTe absorber escapes the active detector volume in about $\sim 5\%$ of events. We model this by a flat distribution extending from zero energy to the $K\alpha$ line, whose normalization is fixed to 5% of the number of events in the $K\alpha$ and $K\beta$ absorber lines.

We model the X-ray continuum with two power laws. One describes the flux from the Crab, using canonical parameters [251]. The other describes the diffuse X-ray background, using parameters from the Chandra deep field measurement of [252]. Before fitting, the power laws must be weighted by the efficiency of both the HgTe

k	PDF component	Functional form
1	potassium $K\alpha$, $K\beta$ cal. events in HgTe	KDE based on pre-flight calibration
2	potassium $K\alpha$ cal. events in Si substrate	gaussian, with fitted mean and width
3	potassium $K\beta$ cal. events in Si substrate	gaussian, with fitted mean and width
4	photoelectrons from cal. source escaping from absorber	uniform from 0 keV to potassium $K\alpha$ energy
5	Crab nebula	power law
6	cosmic XRB	power law
7	cosmic rays	spectrum derived from Geant4 simulation of protons with power law above 1 GeV ($\alpha = 2.7$)

Table 8.4: Components of XQC background PDF in equation (8.5).

k	Parameter	Value
1	number of $K\alpha$, $K\beta$ cal. events in HgTe	1281 ± 36
2	number of $K\alpha$ cal. events in Si substrate	722 ± 29
2	mean measured energy of $K\alpha$ cal. line in Si substrate	2.784 ± 0.006 keV
2 & 3	width of measured $K\alpha$, $K\beta$ cal. lines in Si substrate	0.061 ± 0.002 keV
3	number of $K\beta$ cal. events in Si substrate	221 ± 18
3	mean measured energy of $K\beta$ cal. events in Si substrate	3.020 ± 0.006 keV
4	number of cal. source events with escaping electrons in fit range	25 (fixed 5% escape fraction)
5	Crab spectral index [251]	2.1 (fixed)
	Crab flux at 1 keV [251]	9.7 ± 0.5 photons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1}$ (fixed)
	number of Crab events	155 (fixed)
6	cosmic XRB spectral index [252]	1.4 (fixed)
	constraint on cosmic XRB flux at 1 keV [252]	10.9 ± 1.3 photons $\text{cm}^{-2} \text{s}^{-1} \text{keV}^{-1} \text{sr}^{-1}$ (fixed)
	number of cosmic XRB events	369 (fixed)
7	number of cosmic rays in fit range	5.2 (fixed)

Table 8.5: Parameters of the XQC background PDF in equation (8.5).




Figure 8-12: Simulated energy spectrum in the XQC (*blue*) and Micro-X (*green*) X-ray absorbers. Geant4 [8] is used to simulate a power law distribution of primary cosmic ray protons [9] impinging isotropically on the two absorbers. Both distributions are typical for minimum-ionizing particles. The mean energy deposited in XQC is larger than in Micro-X because of the additional 15 μm Si substrate of the HgTe absorber present in XQC but not in Micro-X.

absorber and the Si substrate. Since the events in the Si substrate appear below their true energy, this component must be shifted to lower energies by a similar fractional energy loss as the calibration lines. The resulting PDF for the continuum is given by

$$P_{\text{power}}(E) = \epsilon^{\text{HgTe}}(E)E^{-\alpha} + \epsilon^{\text{subst.}}(E/k) \left(\frac{E}{k}\right)^{-\alpha}, \quad (8.6)$$

where $k = E_{K\alpha}^{\text{subst.}}/3.31 \text{ keV}$ is an estimate of the fractional energy loss of the $K\alpha$ calibration line. Note that this parameterization implicitly assumes that the charge trapping process in the substrate is energy-independent. Since the Crab lies 19.5° off the observation axis, the geometrical acceptance of the detector to X-rays from the Crab is 94.2% of the effective area for on-axis events. Because of a similar geometrical effect, the acceptance of the diffuse X-rays is 92.7% of the effective area for on-axis events.

The background component due to cosmic rays is only about 2 events in the 2-5 keV window. We obtain the spectral shape of cosmic rays by simulating protons with a typical power law spectrum [9] impinging on the XQC absorber and substrate. Since protons are minimum ionizing particles, this is approximately a Landau distribution with a peak around 7 keV, as shown in Figure 8-12.

Bibliography

- [1] F. Nesti and P. Salucci, J. Cosmol. Astropart. Phys. **07**, 016 (2013).
- [2] N. Bernal and S. Palomares-Ruiz, J. Cosmol. Astropart. Phys. **01**, 006 (2011).
- [3] A. Boyarsky *et al.*, (2014), arXiv:1408.2503.
- [4] E. Anders and N. Grevesse, Geochim. Cosmochim. Acta **53**, 197 (1989).
- [5] S. Horiuchi *et al.*, Phys. Rev. D **89**, 025017 (2014).
- [6] A. Boyarsky *et al.*, Astropart. Phys. **28**, 303 (2006).
- [7] E. Bulbul *et al.*, (2014), arXiv:1402.2301v1.
- [8] S. Agostinelli *et al.*, Nucl. Instrum. Meth. A **506**, 250 (2003).
- [9] P. Papini, C. Grimani, and S. A. Stephens, Nuovo Cim. C **19**, 367 (1996).
- [10] R. Wald, *General Relativity* (University of Chicago Press, Chicago, 1984).
- [11] S. Weinberg, *Cosmology* (Oxford University Press, Oxford, 2008).
- [12] V. C. Rubin and W. K. Ford, Astrophys. J. **159**, 379 (1970).
- [13] Y. Sofue and V. Rubin, Annu. Rev. Astron. Astrophys. **39**, 137 (2001).
- [14] Y. Sofue, ASP Conf. Ser. **197**, 85 (2000).
- [15] F. Zwicky, Astrophys. J. **86**, 217 (1937).
- [16] D. Clowe *et al.*, Astrophys. J. **648**, 109 (2006).
- [17] A. G. Riess *et al.*, Astron. J. **116**, 1009 (1998).
- [18] S. Perlmutter *et al.*, Astrophys. J. **517**, 565 (1999).
- [19] P. Ade *et al.*, Astron. Astrophys. **571**, A15 (2014).
- [20] K. A. Olive *et al.*, Chinese Phys. C **38**, 090001 (2014).
- [21] L. D. Duffy and K. Van Bibber, New J. Phys. **11**, 105008 (2009).

- [22] S. Dodelson and L. M. Widrow, Phys. Rev. Lett. **72**, 17 (1994).
- [23] K. Griest, A. M. Cieplak, and M. J. Lehner, Phys. Rev. Lett. **111**, 181302 (2013).
- [24] G. Jungman, M. Kamionkowski, and K. Griest, Phys. Rep. **267**, 195 (1996).
- [25] J. L. Feng, Ann. Rev. Astron. Astrophys. **48**, 495 (2010).
- [26] J. Ellis *et al.*, Nucl. Phys. B **238**, 453 (1984).
- [27] O. Buchmueller *et al.*, Eur. Phys. J. C **74**, 2922 (2014).
- [28] C. Boehm *et al.*, J. High Energy Phys. **2013**, 113 (2013).
- [29] D. B. Kaplan, Phys. Rev. Lett. **68**, 741 (1992).
- [30] D. Kaplan, M. Luty, and K. Zurek, Phys. Rev. D **79**, 115016 (2009).
- [31] K. M. Zurek, Phys. Rep. **537**, 91 (2014).
- [32] S. D. McDermott *et al.*, Phys. Rev. D **85**, 023519 (2012).
- [33] M. Pospelov, A. Ritz, and M. Voloshin, Phys. Lett. B **662**, 53 (2008).
- [34] D. G. Cerdeño and A. M. Green, Direct detection of WIMPs, in *Particle Dark Matter: Observations, Models and Searches*, edited by G. Bertone, p. 347, Cambridge, 2010.
- [35] C. A. Bertulani, *Nuclear Physics in a Nutshell* (Princeton University Press, Princeton, 2007).
- [36] J. D. Lewin and P. F. Smith, Astropart. Phys. **6**, 87 (1996).
- [37] J. L. Feng *et al.*, Phys. Lett. B **703**, 124 (2011).
- [38] J. L. Feng, J. Kumar, and D. Sanford, Phys. Rev. D **88**, 015021 (2013).
- [39] A. M. Green, Mod. Phys. Lett. A **27**, 1230004 (2012).
- [40] L. M. Widrow and J. Dubinski, Astrophys. J. **697**, 1239 (2008).
- [41] R. Catena and P. Ullio, J. Cosmol. Astropart. Phys. **2010**, 004 (2010).
- [42] M. Pato *et al.*, Phys. Rev. D **82**, 023531 (2010), arXiv:1006.1322.
- [43] F. Kerr and D. Lynden-Bell, Mon. Not. R. Astron. Soc. **221**, 1023 (1986).
- [44] M. C. Smith *et al.*, Mon. Not. R. Astron. Soc. **379**, 755 (2007).
- [45] Y.-Y. Mao *et al.*, Astrophys. J. **764**, 35 (2013).
- [46] G. Gelmini and P. Gondolo, Phys. Rev. D **64**, 023504 (2000).

- [47] D. Smith and N. Weiner, Phys. Rev. D **64**, 043502 (2001).
- [48] D. Akerib *et al.*, Phys. Rev. Lett. **112**, 091303 (2014).
- [49] M. Pyle, *Optimizing the Design and Analysis of Cryogenic Semiconductor Dark Matter Detectors for Maximum Sensitivity*, PhD, Stanford University, 2012.
- [50] S. A. Hertel, *Advancing the Search for Dark Matter from CDMS II to Super-CDMS*, PhD, Massachusetts Institute of Technology, 2012.
- [51] B. Neganov and V. Trofimov, Otkrytia i Izobreteniya **146**, 215 (1985).
- [52] P. N. Luke, J. Appl. Phys. **64**, 6858 (1988).
- [53] C. Kittel, *Introduction to Solid State Physics*, 7th ed. (Wiley, New York, 1996).
- [54] R. Pehl *et al.*, Nucl. Instrum. Methods **59**, 45 (1968).
- [55] J. Lindhard *et al.*, Mat. Fys. Medd. Dan. Vid. Selsk. **33** (1963).
- [56] P. Barbeau, *Neutrino and astroparticle physics with P-type point contact high purity germanium detectors*, PhD, University of Chicago, 2009.
- [57] A. K. Soma *et al.*, arXiv:1411.4802v1.
- [58] L. Pincherle, *Electronic Energy Bands in Solids* (Macdonald & Co., London, 1971).
- [59] B. Cabrera *et al.*, (2010), arXiv:1004.1233.
- [60] V. Mandic, *First Results from the Cryogenic Dark Matter Search Experiment at the Deep Site*, PhD, University of California, Berkeley, 2004.
- [61] R. Agnese *et al.*, Appl. Phys. Lett. **103**, 164105 (2013).
- [62] P. Y. Yu and M. Cardona, *Fundamentals of Semiconductors*, 4 ed. (Springer, Berlin, 2010).
- [63] N. W. Ashcroft and D. N. Mermin, *Solid State Physics* (Harcourt, Fort Worth, 1976).
- [64] S. W. Leman, Rev. Sci. Instrum. **83** (2012), arXiv:1109.1193.
- [65] K. McCarthy, *Detector simulation and WIMP search analysis for the Cryogenic Dark Matter Search experiment*, PhD, Massachusetts Institute of Technology, 2013.
- [66] G. Wang, J. Appl. Phys. **107**, 094504 (2010).
- [67] M. Tinkham, *Introduction to Superconductivity*, 2 ed. (Dover, Mineola, 2006).

- [68] K. D. Irwin and G. C. Hilton, Transition-Edge Sensors, in *Cryog. Part. Detect.*, edited by C. Enss, chap. 3, pp. 63–149, Springer, Berlin, 2005.
- [69] E. Figueroa-Feliciano, J. Appl. Phys. **99**, 114513 (2006).
- [70] M. Pyle, E. Figueroa-Feliciano, and B. Sadoulet, (2015), arXiv:1503.01200v2.
- [71] S. Golwala, *Exclusion Limits on the WIMP-Nucleon Elastic-Scattering Cross Section from the Cryogenic Dark Matter Search*, PhD, University of California, Berkeley, 2000.
- [72] J. Filippini, *A Search for WIMP Dark Matter Using the First Five-Tower Run of the Cryogenic Dark Matter Search*, PhD, University of California, Berkeley, 2008.
- [73] J. M. Carmona *et al.*, Astropart. Phys. **21**, 523 (2004).
- [74] Z. Ahmed *et al.*, Phys. Rev. D **72** (2005).
- [75] C. Enss and S. Hunklinger, *Low-Temperature Physics* (Springer, Berlin, 2005).
- [76] D. Seitz, CDMS Cold Hardware Website.
- [77] K. M. Sundqvist *et al.*, J. Low Temp. Phys. **176**, 188 (2014).
- [78] D. S. Akerib *et al.*, Phys. Rev. D **82**, 122004 (2010).
- [79] R. A. Bunker, *A Low-threshold Analysis of Data from the Cryogenic Dark Matter Search Experiment*, PhD, University of California, Santa Barbara, 2011.
- [80] Z. Ahmed *et al.*, Phys. Rev. Lett. **106**, 131302 (2011).
- [81] D. C. Moore, *A Search for Low-Mass Dark Matter with the Cryogenic Dark Matter Search and the Development of Highly Multiplexed Phonon-Mediated Particle Detectors Thesis*, PhD, California Institute of Technology, 2012.
- [82] R. Agnese *et al.*, Phys. Rev. Lett. **111**, 251301 (2013).
- [83] C. E. Aalseth *et al.*, Phys. Rev. D **88**, 012002 (2013).
- [84] R. Bernabei *et al.*, Eur. Phys. J. C **73**, 2648 (2013).
- [85] G. Angloher *et al.*, Eur. Phys. J. C **72**, 1971 (2012).
- [86] J. Angle *et al.*, Phys. Rev. Lett. **107**, 051301 (2011).
- [87] E. Aprile *et al.*, Phys. Rev. Lett. **109**, 181301 (2012).
- [88] A. L. Fitzpatrick *et al.*, J. Cosmol. Astropart. Phys. **02**, 004 (2013).
- [89] A. Thompson *et al.*, *X-ray Data Booklet*, 2 ed. (Lawrence Berkeley National Laboratory, University of California, 2009).

- [90] S. Cebrián *et al.*, *Astropart. Phys.* **33**, 316 (2010).
- [91] J. Bahcall, *Phys. Rev.* **132**, 362 (1963).
- [92] R. Agnese *et al.*, arXiv:1410.1003v1.
- [93] M. H. Mendenhall and R. A. Weller, *Nucl. Instrum. Meth. B* **227**, 420 (2005).
- [94] P. Redl, *J. Low Temp. Phys.* **176**, 937 (2014).
- [95] J. F. Ziegler, M. D. Ziegler, and J. P. Biersack, *Nucl. Instrum. Meth. B* **268**, 1818 (2010).
- [96] J. Hubbell and S. Seltzer, *Tables of X-Ray Mass Attenuation Coefficients and Mass-Energy Absorption Coefficients* (version 1.4), 2004.
- [97] S. Yellin, *Phys. Rev. D* **66**, 032005 (2002).
- [98] R. H. Byrd, J. C. Gilbert, and J. Nocedal, *Math. Program.* **89**, 149 (2000).
- [99] J. Nelder and R. Mead, *Comput. J.* **7**, 308 (1964).
- [100] P. van Laarhoven and E. Aarts, *Simulated Annealing: Theory and Applications* (Springer, Berlin, 1987).
- [101] W. Härdle, *Smoothing Techniques with Implementation in S* (Springer, New York, 1990).
- [102] C. Bishop, *Pattern Recognition and Machine Learning* (Springer, New York, 2006).
- [103] B. Silverman, *Density Estimation for Statistics and Data Analysis* (Chapman and Hall, New York, 1986).
- [104] Y. Coadou, *EPJ Web Conf.* **55**, 1 (2013).
- [105] T. Hastie, R. Tibshirani, and J. Friedmann, *The Elements of Statistical Learning*, 2nd ed. (Springer, New York, 2009).
- [106] R. E. Schapire, *Mach. Learn.* **5**, 197 (1990).
- [107] Y. Freund and R. E. Schapire, *J. Comput. Syst. Sci.* **55**, 119 (1997).
- [108] H. Voss *et al.*, TMVA, the Toolkit for Multivariate Data Analysis with ROOT, in *XI Int. Work. Adv. Comput. Anal. Tech. Phys. Res.*, p. 40, 2007.
- [109] F. James, *Statistical Methods in Experimental Physics*, 2nd ed. (World Scientific, Singapore, 2006).
- [110] R. Barlow, *Application of the Bootstrap Resampling Technique to Particle Physics Experiments*, 2000.

- [111] R. Khabaz, *Ann. Nucl. Energy* **46**, 76 (2012).
- [112] S. Yellin, (2011), arXiv:1105.2928.
- [113] G. J. Feldman and R. D. Cousins, *Phys. Rev. D* **57**, 3873 (1998).
- [114] Z. Ahmed *et al.*, *Phys. Rev. D* **83**, 112002 (2011).
- [115] C. Savage *et al.*, *J. Cosmol. Astropart. Phys.* **04**, 010 (2009).
- [116] CDMS II Collaboration, *Science* **327**, 1619 (2010).
- [117] R. Agnese *et al.*, *Phys. Rev. Lett.* **112**, 041302 (2014).
- [118] J. Angle *et al.*, *Phys. Rev. Lett.* **110**, 249901 (2013).
- [119] E. Armengaud *et al.*, *Phys. Rev. D* **86**, 051701 (2012).
- [120] C. E. Aalseth *et al.*, (2014), arXiv:1401.3295.
- [121] Z. Ahmed *et al.*, (2012), arXiv:1203.1309.
- [122] P. J. Fox, J. Kopp, M. Lisanti, and N. Weiner, *Phys. Rev. D* **85**, 1 (2012).
- [123] P. Gondolo and G. B. Gelmini, *J. Cosmol. Astropart. Phys.* **12**, 015 (2012).
- [124] E. D. Nobile *et al.*, *J. Cosmol. Astropart. Phys.* **10**, 048 (2013).
- [125] B. Feldstein and F. Kahlhoefer, *J. Cosmol. Astropart. Phys.* **08**, 065 (2014).
- [126] P. J. Fox, Y. Kahn, and M. McCullough, *J. Cosmol. Astropart. Phys.* **010**, 076 (2014).
- [127] B. Feldstein and F. Kahlhoefer, *J. Cosmol. Astropart. Phys.* **12**, 052 (2014).
- [128] A. J. Anderson, P. J. Fox, Y. Kahn, and M. McCullough, (2015), arXiv:1504.03333.
- [129] P. J. Fox, G. D. Kribs, and T. M. P. Tait, *Phys. Rev. D* **83**, 34007 (2011).
- [130] P. J. Fox; J. Liu, and N. Weiner, *Phys. Rev. D* **83**, 103514 (2011).
- [131] SuperCDMS Collaboration, Public Release #1 of Low-Energy SuperCDMS Data.
- [132] K. McCarthy, Dark Matter Search Results from the Silicon Detectors of the Cryogenic Dark Matter Search Experiment, APS April Meeting, 2013.
- [133] E. Del Nobile *et al.*, *J. Cosmol. Astropart. Phys.* **03**, 014 (2014), arXiv:1311.4247.
- [134] P. J. Fox *et al.*, *Phys. Rev. D* **89**, 103526 (2014).

- [135] XENON1T Collaboration, E. Aprile, The XENON1T Dark Matter Search Experiment, in *Proceedings of Dark Matter 2014*, 2012, arXiv:1206.6288.
- [136] G. Plante, The Multi-Ton XENON Program at the Gran Sasso Laboratory, Dark Matter 2014, 2014.
- [137] D. Mailing, After LUX : The LZ Program, in *Proceedings of DPF-2011 Conf.*, 2011, arXiv:1110.0103v2.
- [138] H. Araujo, Mining for WIMPs: The LUX-ZEPLIN (LZ) Experiment, TIPP 2014, 2014.
- [139] M. Batygov *et al.*, DEAP-3600 Dark Matter Search, in *Proceedings of ICHEP*, 2014, arXiv:1410.7673v1.
- [140] S. Fukuda *et al.*, Nucl. Instrum. Meth. A **501**, 418 (2003).
- [141] T. Gaisser and F. Halzen, Annu. Rev. Nucl. Part. Sci. **64**, 101 (2014).
- [142] A. J. Anderson, J. M. Conrad, E. Figueroa-Feliciano, K. Scholberg, and J. Spitz, Phys. Rev. D **84**, 013008 (2011).
- [143] A. J. Anderson *et al.*, Phys. Rev. D **86**, 013004 (2012).
- [144] J. A. Formaggio, E. Figueroa-Feliciano, and A. J. Anderson, Phys. Rev. D **85**, 013009 (2012).
- [145] D. Z. Freedman, Phys. Rev. D **9**, 1389 (1974).
- [146] J. A. Formaggio and G. P. Zeller, Rev. Mod. Phys. **84**, 1307 (2012).
- [147] K. Scholberg, Phys. Rev. D **73**, 033005 (2006), 0511042.
- [148] G. P. Zeller *et al.*, Phys. Rev. Lett. **88**, 091802 (2002).
- [149] J. Barranco, O. G. Miranda, and T. I. Rashba, J. High Energy Phys. **12**, 021 (2005).
- [150] J. Barranco, O. G. Miranda, and T. I. Rashba, Phys. Rev. D **76**, 073008 (2007).
- [151] S. Davidson *et al.*, J. High Energy Phys. **03**, 011 (2003).
- [152] J. W. F. Valle and J. Schechter, Phys. Rev. D **22**, 2227 (1980).
- [153] V. Barger, G. F. Giudice, and T. Han, Phys. Rev. D **40**, 2987 (1989).
- [154] C. J. Horowitz, Phys. Rev. D **55**, 4577 (1997).
- [155] C. J. Horowitz, K. J. Coakley, and D. N. McKinsey, Phys. Rev. D **68**, 023005 (2003).
- [156] SNO Collaboration, B. Aharmim *et al.*, Phys. Rev. C **81**, 055504 (2010).

- [157] C. Arpesella *et al.*, Phys. Rev. Lett. **101**, 091302 (2008).
- [158] S. Abe *et al.*, Phys. Rev. Lett. **100**, 221803 (2008).
- [159] J. Monroe and P. Fisher, Phys. Rev. D **76**, 033007 (2007).
- [160] J. Alonso *et al.*, (2010), arXiv:1006.0260.
- [161] K. Scholberg, J. Phys. Conf. Ser. **136**, 042044 (2008).
- [162] K. Scholberg *et al.*, (2009), arXiv:0910.1989.
- [163] J. M. Conrad and M. H. Shaevitz, Phys. Rev. Lett. **104**, 141802 (2010).
- [164] DAE δ ALUS Collaboration, J. Alonso *et al.*, (2010), arXiv:1008.4967.
- [165] S. K. Agarwalla and P. Huber, Phys. Lett. B **696**, 359 (2011).
- [166] J. M. Conrad, The DAE δ ALUS Experiment, in *Proceedings of Neutrino 2010*, 2010, arXiv:1012.4853v1.
- [167] S. K. Agarwalla *et al.*, J. High Energy Phys. **2011**, 099 (2011).
- [168] S. K. Agarwalla and P. Huber, J. High Energy Phys. **2011**, 059 (2011).
- [169] E. Figueroa-Feliciano, AIP Conf. Proc. **1200**, 959 (2010).
- [170] C. Galbiati and R. Purtschert, J. Phys. Conf. Ser. **120**, 042015 (2008).
- [171] Z. Ahmed *et al.*, Phys. Rev. Lett. **102**, 011301 (2009).
- [172] M. Sorel, J. M. Conrad, and M. H. Shaevitz, Phys. Rev. D **70**, 073004 (2004).
- [173] G. Karagiorgi *et al.*, Phys. Rev. D **80**, 073001 (2009).
- [174] J. Kopp, M. Maltoni, and T. Schwetz, Phys. Rev. Lett. **107**, 091801 (2011).
- [175] A. Aguilar *et al.*, Phys. Rev. D **64**, 112007 (2001).
- [176] O. L. G. Peres and A. Y. Smirnov, Nucl. Phys. B **599**, 3 (2001).
- [177] A. Strumia, Phys. Lett. B **539**, 91 (2002).
- [178] W. Grimus and T. Schwetz, Eur. Phys. J. C **20**, 1 (2001).
- [179] ALEPH, DELPHI, L3, OPAL, and SLD Collaborations, Phys. Rep. **427**, 257 (2006).
- [180] K. Abe *et al.*, Phys. Rev. D **83**, 052010 (2011).
- [181] A. Gando *et al.*, Phys. Rev. D **83**, 052002 (2011).
- [182] W. W. M. Allison *et al.*, Phys. Rev. D **72**, 1 (2005).

- [183] P. Adamson *et al.*, Phys. Rev. Lett. **107**, 021801 (2011).
- [184] P. Adamson *et al.*, Phys. Rev. Lett. **106**, 181801 (2011).
- [185] B. Armbruster *et al.*, Phys. Rev. D **65**, 112001 (2002).
- [186] A. A. Aguilar-Arevalo *et al.*, Nucl. Instrum. Meth. A **599**, 28 (2009).
- [187] A. A. Aguilar-Arevalo *et al.*, Phys. Rev. Lett. **102**, 101802 (2009), arXiv:0812.2243.
- [188] A. A. Aguilar-Arevalo *et al.*, Phys. Rev. Lett. **105**, 181801 (2010).
- [189] T. A. Mueller *et al.*, Phys. Rev. C **83**, 054615 (2011).
- [190] G. Mention *et al.*, Phys. Rev. D **83**, 073006 (2011).
- [191] J. M. Conrad and M. H. Shaevitz, Phys. Rev. D **85**, 013017 (2012).
- [192] F. Dydak *et al.*, Phys. Lett. B **134**, 281 (1984).
- [193] I. E. Stockdale *et al.*, Phys. Rev. Lett. **52**, 1384 (1984).
- [194] MiniBooNE and SciBooNE Collaborations, Phys. Rev. D **85**, 032007 (2011).
- [195] P. Adamson *et al.*, Phys. Rev. Lett. **107**, 011802 (2011), arXiv:1104.3922.
- [196] D. N. McKinsey and K. J. Coakley, Astropart. Phys. **22**, 355 (2005).
- [197] D. S. Akerib *et al.*, Phys. Rev. D **68**, 082002 (2003).
- [198] W. H. Lippincott *et al.*, Phys. Rev. C **81**, 045803 (2010).
- [199] D. Gastler *et al.*, Phys. Rev. C **85**, 065811 (2012).
- [200] P. Benetti *et al.*, Nucl. Instrum. Meth. A **574**, 83 (2007).
- [201] M. G. Boulay and A. Hime, Astropart. Phys. **25**, 179 (2006).
- [202] W. H. Lippincott *et al.*, Phys. Rev. C **78**, 035801 (2008).
- [203] W. H. Lippincott *et al.*, Phys. Rev. C **81**, 039901(E) (2010).
- [204] A. Hime, (2011), arXiv:1110.1005v1.
- [205] H. Back *et al.*, Phys. Procedia **37**, 1105 (2012).
- [206] B. Cai *et al.*, AIP Conf. Proc. **1338** (2011).
- [207] J. Vergados, Y. Giomataris, and Y. Novikov, Nucl. Phys. B **854**, 54 (2012).
- [208] J. D. Vergados, Y. Giomataris, and Y. N. Novikov, J. Phys. Conf. Ser. **259**, 12100 (2010).

- [209] S. K. Agarwalla and R. S. Raghavan, (2010), arXiv:1011.4509.
- [210] C. Grieb, J. M. Link, and R. S. Raghavan, Phys. Rev. D **75**, 93006 (2007).
- [211] W. C. Haxton, Phys. Rev. **C38**, 2474 (1988).
- [212] J. N. Abdurashitov *et al.*, Phys. Rev. C **59**, 2246 (1999).
- [213] P. Anselmann *et al.*, Phys. Lett. **B342**, 440 (1995).
- [214] J. N. Abdurashitov *et al.*, Phys. Rev. **C73**, 45805 (2006).
- [215] V. Barsanov *et al.*, Phys. Atom. Nucl. **70**, 300 (2007).
- [216] H. Michael, R. Wolffe, and S. M. Qaim, Int. J. Appl. Radiat. Isot. **35**, 813 (1984).
- [217] B. Cleveland, private communication, 2011.
- [218] R. T. Kouzes, M. M. Lowry, and C. L. Bennett, Phys. Rev. C **29**, 2343 (1984).
- [219] N. E. Booth, B. Cabrera, and E. Fiorini, Annu. Rev. Nucl. Part. Sci. **46**, 471 (1996).
- [220] C. Enss and D. McCammon, J. Low Temp. Phys. **151**, 5 (2008).
- [221] B. Cabrera, J. Low Temp. Phys. **151**, 82 (2008).
- [222] K. D. Irwin, Appl. Phys. Lett. **66**, 1998 (1995).
- [223] M. A. Lindeman *et al.*, J. Low Temp. Phys. **151**, 190 (2008).
- [224] W. Knaak and M. Meissner, Time-Resolved Specific Heat Measurements on High-Purity Single Crystals Between 50 mK and 1 K, in *Seventeenth Int. Conf. Low-Temperature Phys.*, edited by U. Eckren, A. Schmid, W. Weber, and H. Wühl, pp. 667–668, 1984.
- [225] S. R. Bandler *et al.*, J. Low Temp. Phys. **151**, 400 (2008).
- [226] C. D. Reintsema *et al.*, Electronics for a Next-Generation SQUID-Based Time-Domain Multiplexing System, in *LTD13*, pp. 237–240, AIP Conference Proceedings, 2009.
- [227] D. Bintley *et al.*, Proceedings of SPIE, in *Millimeter, Submillimeter, Far-Infrared Detect. Instrum. Astron. V*, p. 774106, SPIE, 2010.
- [228] M. D. Niemack *et al.*, Appl. Phys. Lett. **96**, 163509 (2010).
- [229] Z. Zhou, private communication, 2007.

- [230] M. J. Berger, J. S. Coursey, M. A. Zucker, and J. Chang, ESTAR, PSTAR, and ASTAR: Computer Programs for Calculating Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions (version 1.2.3).
- [231] A. Boyarsky *et al.*, (2014), arXiv:1402.4119.
- [232] S. Riemer-Sorensen, (2014), arXiv:1405.7943.
- [233] T. E. Jeltema and S. Profumo, (2014), arXiv:1408.4388.
- [234] M. E. Anderson, E. Churazov, and J. N. Bregman, (2014), arXiv:1408.4115.
- [235] D. Malyshev, A. Neronov, and D. Eckert, Phys. Rev. D **103506**, 1 (2014).
- [236] O. Urban *et al.*, (2014), arXiv:1411.0050v1.
- [237] E. Carlson, D. Hooper, and T. Linden, (2014), arXiv:1409.1572.
- [238] D. McCammon *et al.*, Astrophys. J. **576**, 188 (2002), 0205012.
- [239] T. Takahashi *et al.*, The ASTRO-H X-ray astronomy satellite, in *Soc. Photo-Opt. Instru.*, , Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series Vol. 9144, p. 25, 2014, arXiv:1412.1356.
- [240] T. Kitayama *et al.*, (2014), arXiv:1412.1176.
- [241] K. Abazajian, G. M. Fuller, and W. H. Tucker, Astrophys. J. **562**, 593 (2001).
- [242] T. Asaka, S. Blanchet, and M. Shaposhnikov, Phys. Lett. B **631**, 151 (2005).
- [243] A. Boyarsky *et al.*, Phys. Rev. Lett. **97**, 261302 (2006).
- [244] D. P. Finkbeiner and N. Weiner, Phys. Rev. D **76**, 083519 (2007).
- [245] D. P. Finkbeiner and N. Weiner, (2014), arXiv:1402.6671.
- [246] A. Berlin, A. Difranzo, and D. Hooper, (2015), arXiv:1501.03496v1.
- [247] X. Shi and G. Fuller, Phys. Rev. Lett. **82**, 2832 (1999).
- [248] T. Asaka and M. Shaposhnikov, Phys. Lett. B **620**, 17 (2005).
- [249] P. B. Pal and L. Wolfenstein, Phys. Rev. D **25**, 766 (1982).
- [250] S. G. Crowder *et al.*, Astrophys. J. **758**, 143 (2012).
- [251] K. Mori *et al.*, Astrophys. J. **609**, 186 (2004).
- [252] R. C. Hickox and M. Markevitch, Astrophys. J. **645**, 95 (2006).
- [253] W. Verkerke and D. Kirkby, (2003), arXiv:physics/0306116.
- [254] R. Barlow, Nucl. Instrum. Meth. A **297**, 496 (1990).

- [255] G. Cowan *et al.*, Eur. Phys. J. C **71**, 1 (2011).
- [256] L. Moneta *et al.*, (2011), arXiv:1009.1003v2.
- [257] S. N. T. Heine *et al.*, J. Low Temp. Phys. **176**, 1082 (2014).
- [258] W. Voges *et al.*, Astron. Astrophys. **349**, 389 (1999).
- [259] H. Uchiyama, M. Nobukawa, T. G. Tsuru, and K. Koyama, Publ. Astron. Soc. Jpn. **65**, 19 (2013), 1209.0067.
- [260] R. K. Smith *et al.*, Astrophys. J. **556**, L91 (2001).
- [261] A. R. Foster *et al.*, Astrophys. J. **756**, 128 (2012).
- [262] C. Ng *et al.*, Astron. Astrophys. **522**, A96 (2010).
- [263] E. Costantini *et al.*, Astron. Astrophys. **539**, A32 (2012).
- [264] L. Angelini and N. E. White, Astrophys. J. **586**, L71 (2003).
- [265] R. Iaria *et al.*, Astrophys. J. **439**, 575 (2005).
- [266] S. Yamauchi and E. Nakamura, Publ. Astron. Soc. Jpn. **56**, 803 (2004).
- [267] E. M. Cackett *et al.*, Astrophys. J. **690**, 1847 (2009).
- [268] S. Piraino *et al.*, Astron. Astrophys. **542**, L27 (2012).
- [269] Y. Ueda, K. Mitsuda, H. Murakami, and K. Matsushita, Astrophys. J. **620**, 274 (2005).
- [270] T. Narita, J. E. Grindlay, and D. Barret, Astrophys. J. **547**, 420 (2001).