

Development of CDMS-II Surface Event Rejection Techniques and Their Extensions to
Lower Energy Thresholds

A THESIS
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Thomas James Hofer

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Dr. Priscilla Cushman, Adviser

December 2014

© THOMAS JAMES HOFER 2014

UNIVERSITY OF MINNESOTA

THIS IS TO CERTIFY THAT I HAVE EXAMINED THIS COPY OF A DOCTORAL THESIS BY

Thomas James Hofer

AND HAVE FOUND THAT IT IS COMPLETE AND SATISFACTORY IN ALL RESPECTS,
AND THAT ANY AND ALL REVISIONS REQUIRED BY THE FINAL
EXAMINING COMMITTEE HAVE BEEN MADE.

Dr. Priscilla Cushman

NAME OF FACULTY ADVISER(S)

SIGNATURE OF FACULTY ADVISER(S)

December 11, 2014

DATE

GRADUATE SCHOOL

Acknowledgements

No Ph.D. is ever earned by a single person, and the work presented in this thesis is no different. Indeed, the list of people who have helped me reach this point goes back at least as far as my kindergarten teacher (and probably earlier).

In particular, to the many teachers who led me through my educational career - Mrs. Ricketts, Mrs. Widman, Mr. Conrad, Mrs. Coyne, Dr. Wetzell, Dr. Peterson, Dr. Beecken, Dr. Greenlee, Dr. Stein, Dr. Hoyt, and many, many more - I would never have even started this graduate school journey without your guidance.

The CDMS collaboration is as helpful and supportive a group of people as you're ever likely to meet. The local Minnesota group - Prisca, Vuk, Matt, Mark, Scott, Anthony, Jianjie, Allison, Hassan, D'Ann, Hannah, and all the others who came and went during my time here - truly, none of this would have been possible without you.

Classmates who start graduate school together are a little bit like war buddies - together we braved the minefields of Shankar, withstood the ravages of stat mech, slayed the giant known as Jackson, and sweated the GWE. Alex, Jan, Roxanne, Michael, Charles, Dan, Brian, and Mandy - I offer all the gratitude that a knowing look among friends can express.

For keeping me sane, and helping me remember that there is still a world outside of the dark matter direct detection community, I want to thank all my friends in the real world. Craig, Amanda, Justin, Chuck, Barret, Zack and Caleb - I wouldn't have made it without the many hours of board-gaming, video-gaming, show-watching, and kubb-playing you were willing to spend with me.

It almost goes without saying that none of this - literally *none* of this - could have happened without the love and support of my family. Mom, for being my first teacher; Dad, for always supporting my educational pursuits; Rachel, for teaching me how to drive against my will; Jordan, for teaching me more about myself than you probably know; and Jaime, for still liking me in spite of the childhood mistreatment. I love you all.

And finally, none of this means anything without you, Kia. We did it!

Dedication

To Kia Britt

Abstract

The CDMS-II phase of the Cryogenic Dark Matter Search, a dark matter direct-detection experiment, was operated at the Soudan Underground Laboratory from 2003 to 2008. The full payload consisted of 30 ZIP detectors, totaling approximately 1.1 kg of Si and 4.8 kg of Ge, operated at temperatures of ~ 50 mK. The ZIP detectors read out both ionization and phonon pulses from scatters within the crystals; channel segmentation and analysis of pulse timing parameters allowed effective fiducialization of the crystal volumes and background rejection sufficient to set world-leading limits at the times of their publications.

A full re-analysis of the CDMS-II data was motivated by an improvement in the event reconstruction algorithms which improved the resolution of ionization energy and timing information. The Ge data were re-analyzed using three distinct background-rejection techniques; the Si data from runs 125 - 128 were analyzed for the first time using the most successful of the techniques from the Ge re-analysis. The results of these analyses prompted a novel “mid-threshold” analysis, wherein energy thresholds were lowered but background rejection using phonon timing information was still maintained. This technique proved to have significant discrimination power, maintaining adequate signal acceptance and minimizing background leakage.

The primary background for CDMS-II analyses comes from surface events, whose poor ionization collection make them difficult to distinguish from true nuclear recoil events. The novel detector technology of SuperCDMS, the successor to CDMS-II, uses interleaved electrodes to achieve full ionization collection for events occurring at the top and bottom detector surfaces. This, along with dual-sided ionization and phonon instrumentation, allows for excellent fiducialization and relegates the surface-event rejection techniques of CDMS-II to a secondary level of background discrimination. Current and future SuperCDMS results hold great promise for mid- to low-mass WIMP-search results.

Contents

List of Tables	x
List of Figures	xi
1 The Problem of Dark Matter	1
1.1 Motivations	1
1.1.1 Astrophysical Evidence for Dark Matter	1
1.1.2 Cosmological Evidence for Dark Matter	4
1.2 The Case for WIMP Dark Matter	7
1.2.1 The WIMP Miracle	7
1.2.2 The LSP and Other WIMP Candidate Particles	8
1.3 Alternatives to WIMP Dark Matter	10
1.3.1 Particle Alternatives	10
1.3.2 MOND and Alternative Gravity	11
1.4 Experimental Searches for Dark Matter	12
1.4.1 Collider Searches	12
1.4.2 Indirect Detection	14
1.4.3 Direct Detection	17
2 The CDMS-II Experiment	20
2.1 ZIP Technology	20
2.1.1 Cryogenics and Experimental Apparatus	20
2.1.2 Detector Instrumentation and Signal Collection	24
2.2 Shielding	30
2.3 Sources and Calibration	32
2.3.1 ^{133}Ba Calibration	32

2.3.2	^{252}Cf Calibration	33
3	Data Processing and Event Reconstruction	35
3.1	The CDMS-II Data Pipeline	35
3.2	Data Acquisition System	35
3.2.1	Triggering	36
3.2.2	Pulse Digitization and Readout	37
3.3	Optimal Filter Pulse Fitting	38
3.3.1	Optimal Filter Theory	38
3.3.2	Implementation	40
3.4	RTFTWalk Routine	42
3.5	Energy Scale Calibration	45
3.5.1	Ionization Channel Energy Calibration	45
3.5.2	Phonon Channel Energy Calibration	46
3.5.3	Ionization Yield	48
3.6	Position Correction	50
3.6.1	Physical Origins of Position and Energy Dependence	50
3.6.2	Correction Mechanism	52
3.6.3	Results and Cross-checks	54
3.7	c58 Reprocessing	54
4	c58 Re-analysis: Event Selection	58
4.1	Analysis Philosophy	58
4.2	Signal Region Masking	59
4.3	Event Pre-selection	61
4.3.1	Time Period Removal	62
4.3.2	Event-by-event Removal	62

4.4	Fiducial Volume Selection	65
4.5	Yield Discrimination and Thresholds	68
4.6	Single Scatter Selection	71
5	c58 Re-analysis: High Threshold Phonon Timing Analysis	74
5.1	Philosophy of Phonon Timing Analysis	74
5.1.1	Exposure and Leakage Fits	75
5.1.2	Optimization of Phonon Timing Cut	76
5.1.3	Summary of Various Timing Analyses	79
5.2	Surface Event Leakage Estimation	81
5.2.1	Normalization of Leakage Estimates	81
5.2.2	Leakage Estimation Methods	83
5.3	Ge Timing Analysis	86
5.3.1	Classic Timing Analysis	86
5.3.2	5D χ^2 Timing Analysis	88
5.3.3	Neural Network Timing Analysis	91
5.4	Si Timing Analysis	102
5.5	Analysis Exposure and Efficiency	103
5.5.1	Ge Exposure and Efficiency	104
5.5.2	Si Exposure and Efficiency	104
6	Extension of Timing Analyses to Lower Thresholds	109
6.1	Mid-Threshold Analysis	109
6.1.1	Determination of Thresholds	110
6.1.2	Re-optimization of Ge 5D χ^2 Analysis	112
6.2	Background and Sensitivity Estimation	116
6.2.1	Efficiencies	116

6.2.2	Background Estimates and Modeling	116
6.2.3	Sensitivity Estimation	120
7	Candidates and Backgrounds	125
7.1	WIMP Candidate Events	125
7.1.1	High-Threshold Analysis Candidates	125
7.1.2	Mid-Threshold Analysis Candidates	126
7.2	A Catalog of Backgrounds	128
7.3	Background Estimation	130
7.3.1	Radiogenic Neutron Background	130
7.3.2	Cosmogenic Neutron Background	131
7.3.3	^{210}Pb - α Event Background	136
7.3.4	Surface Event Background	137
8	WIMP-search Results of the Re-analyzed c58 Data	140
8.1	Limit-setting in the Standard Halo Model	140
8.1.1	Models for WIMP Interaction	140
8.1.2	Astrophysical Models for Dark Matter	142
8.1.3	Optimum Interval Method	144
8.2	Limits	146
8.3	Further Investigation of Si NRSS Distributions	151
8.3.1	Individual Detector Cross-checks	153
8.3.2	Simultaneous Optimization Cross-checks	155
8.3.3	Conclusions	157
9	Conclusions	168
9.1	Results of the CDMS-II Experiment	168
9.2	Tension, Complementarity, and Resolution	169

9.3 The Future of CDMS	170
References	174
A Unmasked WIMP-search Data	184
A.1 Ge 5D χ^2 Analysis	185
A.2 Ge 5D χ^2 Re-optimized Analysis	198
A.3 Ge Classic Analysis	211
A.4 Ge Neural Network Analysis	224
A.5 Si 5D χ^2 Analysis	237

List of Tables

1	Density parameters of the components of the Λ CDM model.	5
2	Characteristics of the ZIP detectors.	23
3	Fitting routines used in CDMS-II data processing	36
4	List of detector livetimes.	63
5	Estimates of the total number of NRSS in the Ge detectors.	83
6	Estimates of the total number of NRSS in the Si detectors.	84
7	Masked SE leakage estimates for the Ge classic timing analysis.	87
8	Masked SE leakage estimates for the Ge 5D χ^2 timing analysis.	91
9	Summary of MLP nomenclature.	96
10	Masked SE leakage estimates for the Ge neural network timing analysis. . .	102
11	Preliminary SE leakage estimates for the Si 2D χ^2 phonon timing analysis. .	104
12	SAE of the Ge timing analyses.	105
13	Recoil energy thresholds for the MT analysis.	113
14	Leakage estimates and passage fractions for the MT timing analyses.	118
15	Candidate events for the HT analyses.	125
16	Candidate events for the MT analyses.	127
17	Candidate event counts by analysis.	127
18	Radiogenic neutron background estimates.	131
19	Counts from the cosmogenic neutron simulation for Ge.	133
20	Counts from the cosmogenic neutron simulation for Si.	133
21	Cosmogenic neutron background estimates.	136
22	Final cosmogenic neutron background estimates.	137
23	Final SE leakage estimates.	138
24	Total background estimates and observed candidate events.	139
25	Physical parameters of the standard halo model.	143

List of Figures

1	Rotation curves of spiral galaxies.	3
2	Distribution of mass and baryons in the Bullet cluster.	3
3	Correlation function from the SDSS.	6
4	CMB temperature anisotropy power spectrum.	6
5	Illustration of WIMP freeze out.	13
6	Simplified Feynman diagram of dark matter interactions with Standard Model particles.	13
7	Results from invisible channel dark matter searches.	16
8	Indirect detection results from <i>Fermi</i> -LAT observations.	16
9	Schematic of the cryogenic system for CDMS-II.	21
10	Physical layout of the ZIP detectors within the experimental volume.	25
11	Physical dimensions and orientation of the ZIP detectors.	25
12	Detail of phonon instrumentation.	29
13	Readout circuitry of the ionization channels.	29
14	Orientation of ionization channels.	31
15	Shielding layers within the icebox.	31
16	Schematic view of the scintillator panels.	32
17	Noise PSDs in ionization channels.	43
18	Example of the RTFTwalk routine.	43
19	Removal of RTFTwalk energy dependence using a variable Butterworth filter.	47
20	Illustration of the ionization energy calibration algorithm.	47
21	Illustration of the calibration of Q-outer events.	49
22	Illustration of the phonon energy calibration method.	49
23	Shrimp plot for T1Z5.	55
24	Example of a correlation plot for position correction diagnostics.	55

25	OFX event reconstruction for one of the 2010 candidate events.	56
26	Illustration of ER and NR yield bands.	60
27	Phonon pulse delay behavior.	60
28	Illustration of the χ^2 ionization reconstruction quality cut.	66
29	Visualization of the fiducial volume cut.	66
30	Ionization regions for T1Z5.	72
31	Example of yield and threshold cuts.	72
32	Example of leakage fitting for phonon timing cut optimization.	77
33	Example of the leakage scanning method.	80
34	Example of the equal slopes method.	80
35	Classic analysis signal region for detector T1Z5.	90
36	Example of mean and covariance matrix value fits.	90
37	Graphical representation of a neural network.	94
38	Example of an MLP neural network.	94
39	Example of neural network discrimination.	101
40	Neural network response visualized in the yield-recoil energy plane.	101
41	Efficiency of the Ge classic timing cut.	105
42	Efficiency of the Ge 5D χ^2 timing cut.	106
43	Efficiency of the Ge neural network timing cut.	106
44	Efficiencies of the Ge phonon timing cuts.	108
45	Efficiency of the Si 2D χ^2 timing cut.	108
46	Changes in expected leakage and WIMP signal as the ionization threshold is varied.	111
47	Comparison of WIMP spectra for different WIMP masses.	115
48	Problems with the covariance matrix element fits for detector T3Z6.	115
49	Kludge correction factor for the T3Z6 neutron covariance matrix.	117

50	Efficiencies for the Ge MT timing analyses.	117
51	Efficiencies for the Si MT timing analyses.	118
52	SE spectrum and fits for the Ge classic timing analysis.	122
53	Sensitivity estimates for the Ge classic cut using different background models.	122
54	Sensitivity estimates for the Ge no timing analysis using different background models.	123
55	Sensitivity estimates for the Ge classic timing cut using different ionization thresholds.	123
56	Sensitivity estimates for the Ge timing analyses.	124
57	Sensitivity estimates for the Si timing analyses.	124
58	WIMP-search data for detector T1Z5.	129
59	Results of fitting the radiogenic γ Monte Carlo to data.	132
60	Results of the radiogenic neutron simulation.	132
61	Differential event rates in the standard halo model.	147
62	Spin-independent results of the HT analyses.	147
63	Spin-independent results for the HT analysis of the full CDMS-II exposure.	148
64	Spin-independent results for the Ge MT classic timing analysis.	148
65	Spin-independent results for the Ge MT neural network timing analysis. . .	149
66	Spin-independent results for the Ge MT 5D χ^2 timing analysis.	149
67	Spin-independent results for the Ge MT analysis with no timing cut. . . .	150
68	Spin-independent results for the Si MT 2D χ^2 timing analysis.	150
69	Spin-independent results for the Si MT analysis with no timing cut. . . .	152
70	Spin-independent results for all MT analyses with recoil energy thresholds on.	152
71	Spin-independent results for all MT analyses with recoil energy thresholds off.	160
72	Primary limits from this work.	161
73	Timing distributions for the Si HT analysis.	162

74	Timing distributions for the Si MT analysis.	163
75	Timing, energy, and yield distributions for T4Z3.	164
76	Timing, energy, and yield distributions for T5Z3.	165
77	Results of the timing-cut scanning procedure for the HT Si timing analysis.	166
78	Results of the timing-cut scanning procedure for the MT Si timing analysis.	167
79	Current experimental limits and G2 sensitivity projections for DM experiments.	171
80	T1Z2 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	186
81	T1Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	187
82	T2Z3 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	188
83	T2Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	189
84	T3Z2 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	190
85	T3Z4 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	191
86	T3Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	192
87	T4Z2 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	193
88	T4Z4 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	194
89	T4Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	195
90	T5Z4 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	196
91	T5Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis.	197
92	T1Z2 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	199
93	T1Z5 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	200
94	T2Z3 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	201
95	T1Z2 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	202
96	T3Z4 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	203
97	T3Z6 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	204
98	T1Z2 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	205
99	T4Z4 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	206

100	T4Z5 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	207
101	T4Z6 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	208
102	T5Z4 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	209
103	T5Z5 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis.	210
104	T1Z2 WIMP-search data using the Ge classic MT timing analysis.	212
105	T1Z5 WIMP-search data using the Ge classic MT timing analysis.	213
106	T2Z3 WIMP-search data using the Ge classic MT timing analysis.	214
107	T2Z5 WIMP-search data using the Ge classic MT timing analysis.	215
108	T3Z2 WIMP-search data using the Ge classic MT timing analysis.	216
109	T3Z4 WIMP-search data using the Ge classic MT timing analysis.	217
110	T3Z5 WIMP-search data using the Ge classic MT timing analysis.	218
111	T4Z2 WIMP-search data using the Ge classic MT timing analysis.	219
112	T4Z4 WIMP-search data using the Ge classic MT timing analysis.	220
113	T4Z5 WIMP-search data using the Ge classic MT timing analysis.	221
114	T5Z4 WIMP-search data using the Ge classic MT timing analysis.	222
115	T5Z5 WIMP-search data using the Ge classic MT timing analysis.	223
116	T1Z2 WIMP-search data using the Ge neural network MT timing analysis.	225
117	T1Z5 WIMP-search data using the Ge neural network MT timing analysis.	226
118	T2Z3 WIMP-search data using the Ge neural network MT timing analysis.	227
119	T2Z5 WIMP-search data using the Ge neural network MT timing analysis.	228
120	T3Z2 WIMP-search data using the Ge neural network MT timing analysis.	229
121	T3Z4 WIMP-search data using the Ge neural network MT timing analysis.	230
122	T3Z5 WIMP-search data using the Ge neural network MT timing analysis.	231
123	T4Z2 WIMP-search data using the Ge neural network MT timing analysis.	232
124	T4Z4 WIMP-search data using the Ge neural network MT timing analysis.	233
125	T4Z5 WIMP-search data using the Ge neural network MT timing analysis.	234

126	T5Z4 WIMP-search data using the Ge neural network MT timing analysis.	235
127	T5Z5 WIMP-search data using the Ge neural network MT timing analysis.	236
128	T1Z4 WIMP-search data using the Si 2D χ^2 MT timing analysis.	238
129	T2Z1 WIMP-search data using the Si 2D χ^2 MT timing analysis.	239
130	T2Z2 WIMP-search data using the Si 2D χ^2 MT timing analysis.	240
131	T2Z4 WIMP-search data using the Si 2D χ^2 MT timing analysis.	241
132	T3Z3 WIMP-search data using the Si 2D χ^2 MT timing analysis.	242
133	T4Z1 WIMP-search data using the Si 2D χ^2 MT timing analysis.	243
134	T4Z3 WIMP-search data using the Si 2D χ^2 MT timing analysis.	244
135	T4Z3 WIMP-search data using the Si 2D χ^2 MT timing analysis, with R125 removed.	245
136	T5Z3 WIMP-search data using the Si 2D χ^2 MT timing analysis.	246

1 The Problem of Dark Matter

1.1 Motivations

A growing body of evidence from diverse sources indicates that a large fraction of the mass of the Universe is composed of a non-luminous, non-baryonic substance which has been dubbed “dark matter.” Dark matter represents a fascinating unsolved problem in modern physics, and it is no doubt linked to new physics beyond currently accepted models: if general relativity is correct, then the Standard Model of particle physics cannot explain dark matter; or if the Standard Model is correct, our understanding of gravity must be deficient. Whatever the true solution to the dark matter problem is, it will lead us to a deeper understanding of physics, and for that reason it is one of the most active areas of research in the physics community today.

1.1.1 Astrophysical Evidence for Dark Matter

The first evidence of a substantial non-luminous matter component of the Universe was observed in the Coma cluster by Zwicky in 1933[1]. The rotation velocity of a gravitationally-bound object at a radius r from a mass distribution is given by Eq. 1, where $M(r)$ is the enclosed mass at the radius of the object.

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

If the mass distribution of the cluster were determined exclusively by its distribution of luminous matter (i.e. stars), then at radii larger than the visible size of the cluster, the rotation velocity should behave as $v \sim r^{-1/2}$; instead it was observed by Zwicky, and continues to be observed in a variety of astrophysical structures[2], that rotational velocities tend to become constant at such radii (see Fig. 1), suggesting the presence of some non-luminous matter distribution with a density $\rho \sim r^{-2}$ that extends well beyond the luminous

matter distribution.

Spectroscopic measurements of the X-ray emissions of intergalactic gas in galaxy clusters[3, 4] also agree with the conclusion that a large fraction of their mass must be dark, with estimated gas mass fractions on the order of only 10%. These mass discrepancies are also supported by evidence from gravitational lensing, which can probe mass distributions from length scales of ~ 1 kpc (strong lensing[5]) to ~ 100 Mpc (weak lensing[6]). Some of the most compelling evidence for dark matter comes from a lensing analysis of the Bullet cluster, which not only supports the existence of dark matter, but also provides information on other properties of the dark matter itself.

The Bullet cluster is composed of two galaxy clusters in the process of colliding. In the absence of a dark matter component, the mass distribution of the cluster should trace its baryonic matter which is predominantly composed of an X-ray emitting intracluster plasma. However, weak lensing of the Bullet cluster shows that the the centers of the gravitational potential are clearly offset from the centers of the baryonic distribution (Fig. 2), at a significance of 8σ [7]; this phenomenon cannot be reproduced easily via models of modified gravity, although some have been proposed[8]. The Bullet cluster observation (and subsequent others like it[5]) also places limits on the self-interaction strength of the dark matter particles, since the dark matter distributions of the colliding clusters have passed through each other with minimal interaction.

N -body simulations of large-scale structure formation, such as Via Lactea[9] and Millenium[10], have indicated that astrophysical observations are well-reproduced by a model which contains a cosmological constant Λ (dark energy) and a cold (i.e. low velocity) dark matter (CDM) component; this model, known as the Λ CDM model, is able to explain many cosmological and astrophysical observations, and for this reason, the concept of collisionless CDM has become the predominantly accepted theory for the large-scale properties of the dark matter in our Universe.




Figure 1: Rotation curves of spiral galaxies[2]. The flattening of these curves at large radii is not expected if the matter density of the galaxies were completely determined by the luminous matter, and suggests a non-luminous mass component much larger, both in mass and dimension, than that from luminous matter.




Figure 2: Distribution of mass and baryons in the Bullet cluster[7]. The contours show the inferred mass distribution, based on weak lensing calculations, while the background heat map displays observed X-ray intensity which should trace the baryonic matter in the cluster.

1.1.2 Cosmological Evidence for Dark Matter

In addition to the evidence for CDM from galactic dynamics and structure formation, there is also strong cosmological evidence in support of the Λ CDM model. The observation of not only an expanding Universe[11], but an accelerating expansion of the Universe[12, 13], provides strong evidence for a cosmological constant Λ , interpreted as a dark energy with negative pressure, which drives this expansion. Despite several outstanding questions regarding the dark energy, including its composition and the apparent fine-tuning of the constant (naïve quantum field theory predictions are too large by a factor of 10^{120}), the evidence is strongly in favor of a dark energy component to the Universe.

The Λ CDM model can also be used to model the production of atoms in the early universe - a process known as Big Bang nucleosynthesis (BBN) - and thus to predict isotopic abundances at the present time. The predictions for lighter elements (^1H , ^2H , and ^4He) is in excellent agreement with predictions from BBN[14]; however, a significant discrepancy exists ($> 4\sigma$) between the observed and predicted values of the abundance of ^7Li . This so-called “lithium problem” is an outstanding problem in physics; its solution may lie in modifications to the Λ CDM model, or perhaps in new physics beyond the standard model.

An additional prediction of the Λ CDM model concerns the spatial correlation of the matter density of the Universe. At early times, overdense regions of the Universe were gravitationally attracted (a process dominated by the dark matter), but the heat of the photon-baryon gas exerted a pressure, resulting in expanding spherical sound waves in the gas. At recombination, when photons decoupled from the baryons, these baryon acoustic oscillations (BAO) were “frozen” at a radius of ~ 150 Mpc; meanwhile, the dark matter distribution remained focused on the initial overdense regions.

The density perturbations resulting from this process seeded structure formation; therefore observations of the spatial correlations of structures today should have BAO information encoded within them. The Λ CDM model predicts a correlation spectrum peaked at small

Quantity	Symbol	Observed Value
Baryon density	Ω_b	0.0456 ± 0.0015
Dark matter density	Ω_d	0.228 ± 0.013
Dark energy density	Ω_Λ	0.726 ± 0.015

Table 1: Density parameters of the components of the Λ CDM model[17], determined from WMAP, BAO, and supernovae observations. The total energy density of the universe $\Omega = 1.0050^{+0.0060}_{-0.0061}$ is consistent with unity, i.e., a flat universe; the total neutrino density, $\Omega_\nu \lesssim 0.015$, is too small to completely account for the dark matter.

distances (from the dark matter distribution), with a small acoustic peak at the characteristic length scale at recombination. Observations from the Sloan Digital Sky Survey (SDSS) display correlations in general agreement with these predictions[15], as shown in Fig. 3; this independent observation is another point in favor of a CDM component to the Universe.

This phenomenon is also imprinted on the cosmic microwave background (CMB), which is composed of those photons which decoupled from the baryon gas at recombination. The relative heights of the peaks in the CMB temperature anisotropy power spectrum (Fig. 4) are set by the relative amounts of dark matter and baryonic matter in the Universe; more dark matter tends to dampen these peaks, while more baryonic matter enhances the oscillations and magnifies the peaks. The high-resolution results from *Planck*[16] strongly favor the Λ CDM model, and put powerful constraints on the various parameters and matter components of the model.

All of these diverse observations comprise an overwhelming corpus of evidence which points toward a cold dark matter component to our Universe. Despite all this, the particle nature of dark matter remains elusive. Some compelling models exist, which are being actively probed by many dark matter experiments, but as yet none are definitively favored beyond theoretical considerations. The next sections outline possible particle candidates for dark matter, as well as the diverse experimental efforts being performed to test these theories.




Figure 3: Correlation function from the SDSS[15]. The falling shape at low distances is dominated by dark matter; the peak ~ 105 Mpc is the acoustic peak. The lowest solid curve is the Λ CDM prediction with no baryons (i.e., all dark matter); the other curves add baryons and adjust the dark matter content. The best-fit dark matter component is roughly consistent with other measurements.




Figure 4: CMB temperature anisotropy power spectrum, as measured by *Planck*[16]. The relative heights of the first several peaks directly probe the relative abundances of dark and baryonic matter in the early Universe; the Λ CDM model, which is used to generate the red best-fit curve, is in excellent agreement with the data.

1.2 The Case for WIMP Dark Matter

1.2.1 The WIMP Miracle

Since it is well-established that the matter content of the Universe is dominated by dark matter, the question naturally becomes: “what is the dark matter composed of?” Neutrinos seem like a possible component - they are electromagnetically dark, and neutrino oscillation experiments have established that they have a non-zero mass[18] - but observations indicate that their mass contribution to the Universe is much too small to account for the dark matter¹; furthermore they are relativistic (hot), so the corresponding arguments made in favor of CDM do not hold for them.

In the very early Universe, when temperatures were much greater than the mass of the dark matter particles, the dark matter particles would have been in chemical equilibrium with the other Standard Model particles via pair production and annihilation. The equilibrium rate of production and annihilation is given by Eq. 2, where σ is the dark matter annihilation cross-section, v is the relative velocity of the annihilating dark matter, n_{eq} is the equilibrium number density, and $\langle \rangle$ indicates an average over the thermal distribution of the dark matter.

$$R_{\text{eq}} = \langle \sigma v \rangle n_{\text{eq}} \quad (2)$$

However, as the Universe expanded and cooled below the mass of the dark matter particle, the dark matter production rate would decrease by a Boltzmann factor of $e^{-m/T}$ (where m is the mass of the dark matter particle) since only collisions in the high-energy tail of the Boltzmann distribution would have enough energy to produce dark matter particles. Simultaneously, with fewer “new” dark matter particles being produced, the expansion of the Universe lowered the number density n , which consequently lowered the annihilation and production rates of the dark matter. Eventually the mean free path of dark matter

¹Ironically, coherent neutrino scattering is expected to be a significant background for next-generation direct-detection experiments (Fig. 79), so at least *some* dark matter is likely to be detected soon.

annihilation and production became larger than the size of the Universe, at which point the comoving number density of dark matter particles became constant.

This process, known as chemical decoupling or “freeze-out”, produces a relic dark matter density which depends on the coupling $\langle\sigma v\rangle$, and it is illustrated in Fig. 5. Note that as $\langle\sigma v\rangle$ increases, the relic density decreases, and vice versa; this is because a more weakly-interacting dark matter particle would freeze out sooner, leaving a larger relic density. Interestingly, the freeze out happens at roughly $T \approx m/20$, meaning that such particles would naturally be non-relativistic (cold). In this model, the present relic density of dark matter is predicted[19] to be

$$\Omega_\chi h^2 \approx \frac{(10^{-37} \text{ cm}^2) \cdot c}{\langle\sigma v\rangle} \quad (3)$$

Surprisingly, a particle interacting on the weak scale and a mass above $\sim 2 \text{ GeV}/c^2$, provides just the right relic density to account for the dark matter[21, 22]; the moniker WIMP, for **weakly-interacting massive particle**, has been applied to this class of potential dark matter particles. This remarkable coincidence has sometimes been referred to as “the WIMP miracle” because a weak-scale interaction just so happens to fall out of seemingly unrelated cosmology and particle physics arguments.

1.2.2 The LSP and Other WIMP Candidate Particles

The Standard Model does not predict a stable, heavy, neutral, weakly-interacting particle that could be the WIMP; however, supersymmetry (SUSY) is a well-motivated extension to the Standard Model which predicts a fermion (boson) superpartner for each boson (fermion) in the Standard Model. In the minimal supersymmetric extension to the Standard Model (MSSM), R -parity is conserved which causes the lightest supersymmetric particle (LSP) to be stable; in addition, the LSP must almost certainly be electrically neutral and colorless,

otherwise it would have been observed due to its interaction with normal matter via the strong and electromagnetic forces[23].

The spin-0 sneutrino ($\tilde{\nu}$) is a possibility for the LSP, but direct detection experiments (see Sec. 1.4.3) strongly disfavor it[24]. The spin- $3/2$ gravitino ($\tilde{G}$) is likely to have gravity-scale (as opposed to weak-scale) interaction strengths, meaning that it probably would not have been produced thermally in the early Universe; there are some alternative scenarios for $\tilde{G}$ production, mostly via decays from the next-to-lightest supersymmetric particle (NLSP), that allow it to be a possible candidate for the WIMP[25].

By far the most studied WIMP candidate is the neutralino (χ), which is a linear combination (mass eigenstate) of the bino $\tilde{B}$, neutral wino $\tilde{W}^3$, and the two neutral higgsinos $\tilde{H}_1$ and $\tilde{H}_2$.

$$\chi = c_1 \tilde{B} + c_2 \tilde{W}^3 + c_3 \tilde{H}_1 + c_4 \tilde{H}_2 \quad (4)$$

In much of the interesting MSSM parameter space, the $\tilde{B}$ is the dominant component of χ . In the MSSM model, χ is also a Majorana fermion, which is an important characteristic both for determining its relic density from freeze out, and for making it potentially observable through indirect detection experiments (see Sec. 1.4.2).

Although the LSP provides a natural WIMP candidate, recent results from the LHC (including the discovery of the Higgs) put strong constraints on the allowed parameter space for the MSSM[26]. While MSSM is not completely ruled out, alternative SUSY theories, such as next-to-minimal SUSY (NMSSM), are receiving considerable attention[27].

In addition to the LSP, several non-SUSY WIMP candidates have been proposed. In theories with universal extra dimensions, Kaluza-Klein (KK) particles may be stable if some new KK parity is conserved. In this case, the lightest KK particle (LKP) would be stable, and the appropriate relic density can be reproduced at masses $\sim 100 \text{ GeV}/c^2 - 1 \text{ TeV}/c^2$ [28]. A heavy fourth-generation neutrino[29], the Little Higgs[30], and mirror dark matter[31] are all examples of possible non-SUSY WIMP candidates.

1.3 Alternatives to WIMP Dark Matter

1.3.1 Particle Alternatives

Despite the serendipity of the WIMP miracle, there are other non-WIMP particles which could explain the dark matter problem. If the WIMP miracle is correct but the resulting relic dark matter is only long-lived (but not completely stable), then the WIMPs could decay into other (stable) massive particles which decay even more weakly. These “superWIMPs” are usually postulated to be $\tilde{G}$, decaying from the NLSP, which can be identified as χ using the same WIMP miracle arguments as above.

Another potential dark matter candidate arises from the Peccei-Quinn solution to the strong CP problem[32]. The introduction of a pseudoparticle, dubbed the axion, adds a new spontaneously-broken global symmetry which dynamically relaxes CP -violation in our Universe to zero. Interestingly, axions can couple to photons in the presence of strong electromagnetic fields via the Primakoff effect[33]; because of this, axions may interact with valence electrons. Searches[34] have been performed for both relic axions and axions produced in the Sun (again via the Primakoff effect), setting strong limits on the photon and electron couplings of sub-keV/ c^2 axions.

Finally, sterile neutrinos (N) are hypothetical particles which only interact via gravity. They may arise as right-handed neutrinos, which are simple to add to the Standard Model - such a concept is theoretically well-motivated since neutrinos are the only Standard Model particles which are not observed in both chiralities. If the sterile neutrino mixes with the active neutrinos, then the sterile neutrino will be unstable and decay either via $N \rightarrow \nu\nu\bar{\nu}$ or, more rarely, $N \rightarrow \nu\gamma$; certain choices of active and sterile neutrino masses yield a lifetime for N of $\sim 10^{17}$ years[35], long enough to make it a metastable dark matter candidate. Astrophysical[36] and X-ray[37] observations have put constraints on sterile neutrino dark matter, but there remains significant available parameter space.

1.3.2 MOND and Alternative Gravity

The dark matter problem can be stated simply: the observed dynamics of large celestial objects does not conform to the Newtonian gravitational expectation for visible masses. The conclusion that non-luminous matter solves the problem necessarily presupposes that our theory of gravitational dynamics on large scales is correct. If we make the contrapositive argument - that the visible mass distributions are correct - then the problem must be solved by modifying our understanding of gravity. Such alternative gravity, or modified Newtonian dynamics (MOND), models have been extensively studied for some time.

Newtonian gravity with visible matter seems to be a good theory at distance scales from hundreds of km to tens of light years, but at galactic scales, rotation curves fail to obey Newtonian expectations (see Fig. 1). The canonical formulation of MOND[38] posits that below some acceleration scale $a_0 \approx 10^{-10} \text{ m/s}^2$, the gravitational dynamics change. The corresponding mathematical statement is given in Eq. 5, where $\tilde{\mu}(x) \rightarrow 1$ when $x \gg 1$ and $\tilde{\mu}(x) \rightarrow x$ when $x \ll 1$.

$$\tilde{\mu}\left(\frac{|\vec{a}|}{a_0}\right)\vec{a} = -\nabla\Phi_N \quad (5)$$

The MOND paradigm does resolve some problems quite elegantly, such as the Tully-Fisher relation $M \propto V_F^4$ which relates the mass of a galaxy to its asymptotic rotational velocity, and predicts galactic rotation curves with good accuracy[39]. However, at its core MOND is merely phenomenological; deeper physical theories must be constructed from which MOND can “fall out”.

The most popular alternative gravitational theory, which is covariant and reduces to general relativity (GR) in the limiting case, is **Tensor-Vector-Scalar** gravity (TeVeS)[40]. Experimental tests for gravitational lensing[41] and cosmology[42] in the TeVeS framework remain somewhat embryonic, but more stringent test should help distinguish between GR and alternative gravity theories.

1.4 Experimental Searches for Dark Matter

In order to detect particle dark matter and shed light on its composition, we must observe its interactions with Standard Model particles. In general, there are three complementary approaches one can take to observe these interactions (see Fig. 6): produce dark matter from high-energy collisions of Standard Model particles in a collider, mimicking their thermal production in the early Universe (collider searches); look for the decay and/or annihilation products of relic dark matter particles in the Universe today (indirect detection); or directly observe scatterings of dark matter particles off of standard model particles (direct detection). All three will be important in a future dark matter detection scenario to resolve degeneracies which may confound any one experimental paradigm.

1.4.1 Collider Searches

Experiments at the LHC, such as CMS and ATLAS, have constrained the dark matter phase space in two different ways: first, by searching “directly” for dark matter by looking for missing transverse momentum in pp collisions; or second, by using global constraints to restrict the available SUSY (and thus LSP WIMP) parameter space. The former is a more general approach (sometimes termed an “invisible channel search”) in that it is sensitive to any non-Standard Model dark particle, while the latter is of great interest at present because of the unexpectedly null SUSY results from the first run of the LHC.

ATLAS[43] and CMS[44] have both performed dark matter searches using the invisible channel technique. ATLAS searched for missing transverse momentum in events with hadronic showers, assuming a dark matter production mechanism of $pp \rightarrow \chi\bar{\chi} + X$, where X is a W or Z boson; the CMS analysis looked for final states with a photon and missing transverse momentum. The results of these searches (Fig. 7) show strong sensitivity to low-mass WIMPs, with a cut-off at high mass due to the limited amount of energy available for pair production. The floor of the exclusion curves is set by the integrated luminosity of




Figure 5: Illustration of chemical decoupling of dark matter particles[20]. As the temperature of the Universe decreases, the dark matter particles “freeze out” and settle at a constant relic density which depends on their interaction strength.




Figure 6: Simplified Feynman diagram of dark matter interactions with Standard Model particles. Via some new physics processes, WIMPs (χ) interact with Standard Model (SM) particles. Collider searches seek to produce WIMPs from SM particles, while indirect searches look for SM products of the inverse process; direct detection experiments instead look for scatterings of WIMPs off of SM particles.

the data.

The first run of the LHC discovered no evidence for SUSY, leading to strong constraints on the masses of several supersymmetric partner particles. However, because of the large parameter space of MSSM, NMSSM, and other SUSY variations, regions of the parameter space still exist which are compatible with present experimental limits[46, 47], although naturalness and fine-tuning problems make them less attractive than before. Complementary experimental searches, like direct detection searches down to the coherent neutrino scattering floor (see Sec. 9.2), will be important for a fuller understanding of the nature of SUSY.

1.4.2 Indirect Detection

Indirect detection is a technique which seeks to observe the annihilation and/or decay products of dark matter particles from astrophysical sources. Since the mechanisms for dark matter decay or annihilation are unknown, a variety of products must be searched for, including photons, neutrinos, and antimatter.

Perhaps the simplest dark matter products to look for are gamma rays (photons), since they suffer little absorption in the local Universe and propagate simply. For dark matter annihilation (of particular interest for the LSP, since it is a Majorana particle), the non-relativistic speeds of the dark matter means that the photons produced would be back-to-back with minimal spread in their energies. Thus, an annihilation scenario would be expected to produce sharp gamma-ray lines from regions with high dark matter densities (like the galactic center). For decay scenarios (such as for the meta-stable $\tilde{G}$), the galactic poles are a more promising region to search since backgrounds are lower and the expected spectra are less well-known.

Some recent measurements by the *Fermi*-LAT instrument have been analyzed as evidence for annihilating dark matter. The 110 and 130 GeV excesses, spatially centered on

the galactic center but also present in other regions, seems to be well-fit by a $\chi\bar{\chi} \rightarrow \gamma\gamma$ and $\chi\bar{\chi} \rightarrow \gamma Z$ interpretation[48], although additional data indicate that this may be a statistical fluke. Additionally, an excess near the galactic center at 1 - 3 GeV has been observed, and can be well-fit a dark matter particle ($m_\chi \sim 30 - 40 \text{ GeV}/c^2$) annihilating to $b\bar{b}$ [49]. Fig. 8 illustrates these results.

Another potentially observable signal may be from dark matter annihilation in the center of the Sun. As the Sun moves through the Milky Way's dark matter halo, some WIMPs may collide with solar nuclei and lose enough energy to become gravitationally bound. Over billions of years, a sufficient WIMP density may accumulate that the WIMP capture and annihilation rates reach an equilibrium. Although most annihilation products (including photons) would become trapped in the solar core, the Sun is mostly transparent to neutrinos, making them an excellent indirect detection search target. One advantage of this approach is that the WIMP capture rate is largely insensitive to the distribution or structure of dark matter in the galaxy, but is determined by the local dark matter density, average WIMP velocity, and escape velocity. This removes much of the astrophysical uncertainty associated with most other indirect and direct detection searches. Latest results from neutrino telescopes like IceCube[50] and ANTARES[51] report signals consistent with background, and set strong limits on spin-dependent WIMP-nucleon couplings.

Finally, WIMP annihilation in the galactic halo may lead to antimatter excesses ($\bar{p}$, e^+ , antinuclei, etc.) on top of known astrophysical sources for cosmic rays (CRs). Recent measurements of the CR positron fraction ($e^+/(e^+ + e^-)$) by PAMELA[52], *Fermi*-LAT[53], and AMS-02[54] indicate a spatially isotropic and featureless rise at energies above $\sim 10 \text{ GeV}$, which has been interpreted by some as an annihilation product from $\sim \text{TeV}/c^2$ WIMPs[55], but for a variety of reasons (e.g., the AMS-02 and *Fermi*-LAT dark matter allowed regions are incompatible at the 10^{-7} level) the dark matter interpretation is not considered conclusive.




Figure 7: Results from invisible channel dark matter searches by ATLAS[43] and CMS[44]. LEFT: Spin-independent results from ATLAS using the $pp \rightarrow \chi\bar{\chi} + X$ ($X = W$ or Z) production channel. Black squares (red triangles) are the results assuming that u and d quarks have opposite (same) sign couplings to WIMPs. RIGHT: Spin-independent results from CMS looking in the $pp \rightarrow \chi\bar{\chi} + \gamma$ production channel. The purple dashed limit is from CDF (Tevatron) data[45].




Figure 8: Indirect detection results from *Fermi*-LAT observations. LEFT: 110 and 130 GeV lines from the galactic center, which can be interpreted as $\chi\bar{\chi} \rightarrow \gamma\gamma$ and $\chi\bar{\chi} \rightarrow \gamma Z$ [48]. RIGHT: $\sim$ GeV excess, fit by a 35 GeV/ c^2 dark matter particle annihilating to $b\bar{b}$ [49].

1.4.3 Direct Detection

Direct detection experiments seek to observe the elastic scatters of WIMPs off of nuclei. By measuring the rate and energies of nuclear recoils, constraints can be placed on WIMP parameter space (see Sec. 8.1 for more details). Because of the extremely low rate of WIMP scatters, limiting backgrounds and maximizing exposure (i.e. target mass) are the primary obstacles to extending sensitivities in direct detection experiments. Furthermore, for kinematic reasons, different target materials will have better sensitivities in different WIMP mass regimes; in particular, lowering energy thresholds is the main challenge to extending sensitivities in the low WIMP-mass region of the parameter space.

Irreducible nuclear recoil backgrounds (e.g. neutrons) are mitigated by using radiopure materials and via shielding; for this reason many direct detection experiments can be found in deep underground laboratories. Many experiments remove other backgrounds by measuring recoil energies through two complementary channels; the partitioning of energy between these channels can then be used to discriminate between signals and backgrounds. The most common choices for these energy channels are ionization, scintillation, and heat measurement. A selection of experiments which utilize this approach, along with others which use alternative methods, will be surveyed below.

Some prominent direct-detection experiments which utilize two-channel energy measurements are LUX[56], XENON100[57], CRESST[58], EDELWEISS[59], and CDMS. LUX and XENON100 use a target material of liquid Xe (LXe), which is an efficient scintillator and ionizer for incident radiation at a relatively high temperature (~ 165 K). A liquid-gas dual phase set up is operated in a TPC, with a thin gas layer maintained above a large liquid volume. When an incident particle recoils in the liquid, it creates scintillation light, which is detected promptly, and ionization electrons, which are drifted toward the top of the TPC via an applied voltage. When these drifted electrons reach the gas phase, they produce a secondary electroluminescent light pulse. The relative heights and shapes of these two

light pulses can be used to discriminate between signal and background, and to estimate the depth of the interaction within the liquid. In general, liquid noble gas experiments are attractive for direct detection searches since they can be scaled up to large masses with relative ease, although the heavy nuclei and relatively high thresholds diminish their sensitivity to low-mass WIMPs. In particular, the next-generation LXe experiment LZ[60] will explore much of the high-mass WIMP parameter space down to the neutrino floor (see Sec. 9.2).

CRESST is a solid-state, multi-element experiment which utilizes CaWO_4 crystals as a target material, although the W dominates at high mass both due to its nuclear mass and the coherent enhancement of WIMP scattering for spin-independent interactions (see Sec. 8.1). The crystals are operated at cryogenic temperatures, allowing for heat measurements to be made simultaneously with the measurement of scintillation light. Multi-element solid-state experiments of this nature have sensitivity to a broad range of WIMP masses, due to the differing nuclear responses, but the inability to identify with certainty the type of recoiling nucleus can also introduce some systematic uncertainties.

CDMS (Ge and Si) and EDELWEISS (Ge only) are semiconductor experiments operated at cryogenic temperatures. Semiconductors are good ionizers, since radiation liberates valence electrons into the conduction band and creates paired hole quasiparticles which can be collected with an applied voltage across the crystals. At very low (~ 10 mK) temperatures, heat deposition can also be measured, allowing for discrimination of nuclear recoil signals from electron recoil backgrounds. Ge and Si are common target materials for semiconductor experiments, since they are single-element materials and can be produced economically with excellent purity; other materials, such as GaAs, have elicited some interest but technical difficulties have kept them from being seriously pursued. The necessarily modular nature of semiconductor dark matter detectors does increase detrimental surface effects, but the modularity is beneficial in that multiply-scattering events can be more easily tagged and

removed.

Some experiments eschew the dual-energy measurement technique, instead focusing on a single energy measurement to achieve low thresholds. For example, DAMA/LIBRA[61] and CoGeNT[62] read out only scintillation and ionization signals, respectively, and search for an annual modulation in their signals due to a changing WIMP flux during the year. In this regime, the absolute background rate is not important (as long as it is stable), since it adds only a DC offset to the modulating WIMP signal. Both of these experiments have observed annual modulation signals which could be interpreted as WIMP signals, but they are in strong conflict with other experimental results (see Sec. 9.2). Some two-channel experiments have also performed low-threshold analyses using only one energy channel; for example, the SuperCDMS experiment has also implemented a novel single-energy experimental mode called CDMSlite, wherein ionization measurements are sacrificed, but the phonon (heat) channel achieves thresholds of < 200 eV.

Finally, directional searches have also been proposed, which attempt to measure the “WIMP wind” caused by the relative motion of the solar system with respect to the dark matter halo. Rather than search for an annual modulation, such searches attempt to measure the modulation of signal due to the changing orientation of the detectors over a sidereal day. The most promising proposals are gas time-projection chambers (TPCs), which project ionization tracks in the target gas onto a 2-dimensional CCD image. Some prototype TPC experiments, such as DRIFT[63] and DMTPC[64], have been built, and larger implementations are planned.

2 The CDMS-II Experiment

2.1 ZIP Technology

The CDMS-II experiment was a direct detection experiment located in the Soudan Underground Laboratory in Soudan, Minnesota, USA. The laboratory is located at a depth of 780 m, with a water-equivalent depth of $\sim 2090\text{m}$ based on muon flux measurements[65]. Semiconductor detectors, operated at very low temperatures, were used for the simultaneous measurement of heat (through athermal phonon collection) and ionization (through electron-hole pair collection) signals from recoil events within the crystal lattice. In addition to two independent energy measurements, the depth within the detector at which particle recoils occurred can be inferred from pulse timing information. The CDMS-II detectors, henceforth referred to as “Z-sensitive Ionization and Phonon” (ZIP) detectors, derive their name from these unique characteristics.

2.1.1 Cryogenics and Experimental Apparatus

The ZIP detectors were operated at approximately 50 mK by a cryogenic cooling system. The primary cooling power at the lowest-temperature stage of the cryogenics system was provided by an Oxford Instruments 400S ^3He - ^4He dilution refrigerator, with a cooling power of $400\text{ }\mu\text{W}$ at 100 mK and a base temperature below 10 mK in the absence of a heat load. Since the refrigerator itself was not constructed out of radiopure materials, it was located outside the shielded experimental volume (known as the “icebox”) and cooling power was transferred to the experimental payload via a series of concentric copper pipes known as the fridge stem, or “F-stem”. Electronic readout from the experiment was routed through a separate egress (known as the “E-stem”) for connection to warm electronics and the DAQ system. Fig. 9 shows the physical layout of the apparatus, and the various temperature stages of the cryogenic system.




Figure 9: Schematic of the cryogenic system for CDMS-II. On the far left is the dilution refrigerator, which provides cooling power down to base temperature (black contour). The cold stem connects the refrigeration unit to the icebox in the center. Electronic readout occurs through the E-stem, at far right. Temperature stages between base temperature and room temperature (red) are shown: 77 K (yellow), 4 K (green), 600 mK (light blue), and 50 mK (dark blue).

In order to lower the heat load on the primary dilution refrigerator, several “step-down” temperature stages were used. The outermost and warmest of these stages were the liquid nitrogen (77 K) and liquid helium (4 K) shields which surrounded the refrigerator, F-stem, E-stem, and icebox. Inside the dilution refrigerator, a 1-K pot is used to condense the helium mixture as it enters the mixing chamber. Finally, there are two intermediate temperature stages (nominally 600 mK and 50 mK) which complete the stepping down to the base temperature stage.

Within the icebox, several layers of passive shielding surrounded the detectors themselves (see Sec. 2.2). The detector payload was arranged into 5 towers of detectors, each containing 6 detectors, for a total of 30 ZIP detectors. By target material, 19 were Ge and 11 were Si, for a total target mass of 4.8 kg and 1.1 kg, respectively. Fig. 10 and Tbl. 2 give details regarding the composition and placement of the 30 ZIP detectors.

The icebox was placed within an enclosure made out of scintillating panels which acted as an active muon veto (see Sec. 7). This veto, the refrigerator, and the electronics read-out hardware were all enclosed within a Faraday cage to protect against electromagnetic interference. The interior of the Faraday cage was maintained as a class-1000 (class-10,000 when workers are present) clean room to protect against airborne contamination of detector surfaces[65]. Experimental runs were indexed by periods of continuous base temperature operation; the full, five-tower operation of CDMS-II, spanning October 2006 to September 2008, comprised runs 123 to 128. Due to a change in the data processing between runs 124 and 125, the analysis in this document was performed only on runs 125 through 128, which will henceforth be referred to as “c58” (“combined 125-128”) for clarity.

Detector	Type	Mass (g)	Impurity (cm^{-3})
T1Z1	Ge n-type	230.5	$9.5 \times 10^{10} - 1.0 \times 10^{11}$
T1Z2	Ge n-type	227.6	$9.5 \times 10^{10} - 1.0 \times 10^{11}$
T1Z3	Ge n-type	219.3	$9.5 \times 10^{10} - 1.0 \times 10^{11}$
T1Z4	Si ?-type	104.6	unknown
T1Z5	Ge n-type	219.3	$9.5 \times 10^{10} - 1.0 \times 10^{11}$
T1Z6	Si ?-type	104.6	unknown
T2Z1	Si n-type	101.4	3×10^{16}
T2Z2	Si n-type	104.6	3×10^{16}
T2Z3	Ge n-type	219.3	$9.5 \times 10^{10} - 1.0 \times 10^{11}$
T2Z4	Si n-type	104.6	3×10^{16}
T2Z5	Ge n-type	238.9	$2.7 \times 10^{11} - 4.7 \times 10^{11}$
T2Z6	Si n-type	104.6	3×10^{16}
T3Z1	Si n-type	104.6	3×10^{16}
T3Z2	Ge n-type	231.2	$2.7 \times 10^{11} - 4.7 \times 10^{11}$
T3Z3	Si n-type	104.6	3×10^{16}
T3Z4	Ge p-type	238.9	$8.5 \times 10^{10} - 1.9 \times 10^{11}$
T3Z5	Ge p-type	238.9	$8.5 \times 10^{10} - 1.9 \times 10^{11}$
T3Z6	Ge n-type	231.7	$2.7 \times 10^{11} - 4.7 \times 10^{11}$
T4Z1	Si n-type	101.4	3×10^{16}
T4Z2	Ge p-type	238.9	$8.5 \times 10^{10} - 1.9 \times 10^{11}$
T4Z3	Si n-type	101.4	3×10^{16}
T4Z4	Ge p-type	234.6	$8.5 \times 10^{10} - 1.9 \times 10^{11}$
T4Z5	Ge p-type	231.9	$8.5 \times 10^{10} - 1.9 \times 10^{11}$
T4Z6	Ge p-type	238.9	$8.5 \times 10^{10} - 1.9 \times 10^{11}$
T5Z1	Ge n-type	224.5	$9.5 \times 10^{10} - 1.0 \times 10^{11}$
T5Z2	Ge p-type	229.5	$8.5 \times 10^{10} - 1.9 \times 10^{11}$
T5Z3	Si n-type	101.4	3×10^{16}
T5Z4	Ge n-type	224.5	$2.7 \times 10^{11} - 4.7 \times 10^{11}$
T5Z5	Ge p-type	234.8	$8.5 \times 10^{10} - 1.9 \times 10^{11}$
T5Z6	Ge n-type	223.6	$2.7 \times 10^{11} - 4.7 \times 10^{11}$

Table 2: Characteristics of the ZIP detectors. The naming system, T#Z#, indicates which tower in which the detector is located, and what position within the tower it occupies (see Fig. 10). Masses were computed from measured detector dimensions and densities, with uncertainties of $< \pm 0.5\%$. Impurity levels are taken from [66] (Ge) and [67] (Si).

2.1.2 Detector Instrumentation and Signal Collection

Much of the background-discrimination power of the ZIP detectors comes from the simultaneous measurement of heat and ionization energy on an event by event basis². Each ZIP was instrumented for phonon collection on one side and ionization collection on the other side, and a small bias (3 V for Ge, 4 V for Si) was maintained across the crystal to drift charge carriers for collection. Fig. 11 shows the geometrical placement of the ionization and phonon channels on the detector faces.

2.1.2.1 Phonon collection

Particles recoiling within the detectors will release high frequency (THz-scale) athermal phonons, either directly through a collision with a nucleus, or through energy released from the transport of charge carriers across the crystal bias voltage. Phonons propagate differently than many other particles, due to their unique dispersion properties which result from the discrete crystal medium through which they travel; this discrete medium allows phonons to undergo anharmonic decay, in which a single phonon decays into two lower-energy phonons, as well as the usual isotopic scattering (with no energy loss) off of atoms in the lattice. The scattering cross-sections are highly frequency dependent[69]:

$$\sigma_{\text{IS}} \propto \nu^4; \quad \sigma_{\text{AD}} \propto \nu^5 \quad (6)$$

where “IS” and “AD” denote isotopic and anharmonic processes, respectively.

The high-frequency phonons produced by the initial recoils³ thus have very short mean free paths, and rapidly down-convert via anharmonic decay to lower frequencies, until the scattering cross-section is low enough that the phonons become ballistic, with a mean free

²For an excellent and more thorough discussion of the ZIP detector technology, see [68].

³The energy spectrum of the phonons immediately after recoil is highly peaked near the Debye frequency, which is ~ 3 THz for Ge and ~ 6 THz for Si.




Figure 10: Physical layout of ZIP detectors within the experiment volume. TOP: tower layout; small labels along the edges of the hexagons indicate where electronic readout channels were placed for each of the 6 ZIPs within each tower. BOTTOM: Location of ZIP detectors within each tower. Labels on each ZIP indicate the name given to each crystal prior to installation. Crystals with labels beginning with “S” are Si, while those beginning with “G” are Ge.




Figure 11: LEFT: Physical dimensions of the ZIP detectors. The flats at top and bottom are required for electronic readout hardware. RIGHT: Orientation of the instrumented surfaces of the ZIP detectors (not to scale). Phonon channels (labelled “P”) are shown on the top face, and ionization channels (labelled “Q”) are shown on the bottom face. The X and Y axes are in the plane of the detector face.

path large compared to the crystal dimensions. These ballistic phonons travel in straight-line paths and reflect off of bare surfaces of the crystal. The speed of sound and the rate of anharmonic decay set a time-scale on the order of a few μs over which this quasi-diffusive process transports phonons from the initial site of recoil to the detector surfaces, where signals are collected and read out.

In addition to the primary, recoil-induced phonons, charge carriers drifting across the crystal can also produce phonons. These phonons, produced via the Neganov-Luke effect[70], are somewhat analogous to Cerenkov radiation[68]. For the fields applied in the CDMS-II experiment (several V/cm), charge carriers completely traverse the crystal on the order of a few 100ns. The total amount of energy converted from the charge carriers into the phonons is given by Eq. 7, where Q is the total ionization of the event and V is the applied voltage.

$$E_{\text{Luke}} = QV \quad (7)$$

The phonon collection instrumentation consists of patterned Al and W films deposited on one face of the detectors (Fig. 12). The superconducting Al, which is well below its critical temperature of $T_c = 1.2\text{K}$, absorbs incident phonons which break apart Cooper pairs, creating long-lived quasiparticles⁴. As these quasiparticles diffuse along the Al films, they become trapped in the Al-W overlap region due to its lower energy gap, and finally “funnelled” into the W films, which are operated as transition-edge sensors (TESs) whose signals are amplified and read out by SQUID arrays[71].

The entire Al-W phonon collection structures are known as “quasiparticle-assisted electro-thermal-feedback transition-edge-sensors” (QETs); in this case, the W film forms the transition-edge sensor (TES). Their general patterning and organization is shown in Fig. 12. Each of the 4 phonon channels consists of 37 square 5 mm cells connected in series electrically.

⁴This process is reminiscent of the production of electron-hole pairs in the semiconducting detector crystals.

Each cell contains many TESs to which quasiparticles are funneled by Al fins. The Al fins are intentionally designed *not* to have full coverage of the surface area of each cell, since phonons collected more than a few diffusion lengths away from the TES will lose energy through pair recombination⁵ within the Al. Since portions of the detector faces which are bare have nearly 100% reflection efficiency, energy collection can be improved dramatically by restricting Al coverage in smaller regions near the TES. The region near the outer edge of the crystal is patterned sparsely with Al to maintain a uniform electric field across the crystal while limiting phonon absorption.

2.1.2.2 Ionization collection

Since the detector materials used in CDMS-II are semiconductors, recoiling particles can elevate valence electrons into the conduction band, leaving behind holes in the valence band. The electrons and holes behave as negative and positive charge carriers, respectively, and an applied electric field will cause them to drift to opposite sides of the detector where they can be collected and read out.

The energy, ϵ , required to produce an electron-hole pair is 3.0 eV for Ge and 3.8 eV for Si. Thus, the ionization energy of an event can be related to the charge produced through Eq. 8, where e is the electron charge.

$$E_Q = \frac{Q\epsilon}{e} \quad (8)$$

The readout circuitry of the ionization channels is shown in Fig. 13. The bias voltage V_b drags the charge carriers across the detector (represented by C_d), and the ionization signal is capacitively coupled to a JFET amplifier. The output of the amplifier will be proportional to the deposited ionization energy, $V_f = Q/C_f \propto E_q$, which provides us with a

⁵Simulations have shown that fewer than 10% of phonons actually reach the TESs if Al is patterned at full coverage[68].

linear estimator of the ionization energy which can be accurately calibrated in an absolute sense (see Sec. 3.5). The time scale of these charge pulses is defined by the feedback circuit, with a characteristic time constant $R_f C_f \sim 40 \mu\text{s}$.

The physical layout of the ionization channels of the ZIP detectors is shown in Fig. 14. Al electrodes are patterned across one face of the detector, and act as the charge collection elements. Separate inner and outer electrodes are used to reject high-radius events which often have poor charge collection efficiency. The definition of the fiducial volume, defined through the rejection of events with measurable signals in the outer electrode, is discussed further in Sec. 4.

One important effect which can potentially depress ionization collect is charge carrier trapping. Impurities in the crystal structure create shallow energy “traps” for charge carriers in the valence-conduction band structure, with a characteristic depth of $\sim 10\text{meV}$. At room temperature ($kT \sim 20 \text{ meV}$) trapping is somewhat negligible, but at cryogenic temperatures ($kT_{base} \sim 5 \mu\text{eV}$) and low recoil energies (1 keV corresponds to only 250 to 350 electron-hole pairs), charge trapping due to impurity sites can be a significant effect. However, the low thermal energies of the charge carriers compared to the trap depth is also beneficial, since trapped charge carriers are unlikely to become dislodged, effectively “filling in” the charge traps. This metastable detector state can be consistently achieved by periodically exposing the detectors to photons from LEDs mounted on the detector housing, which liberate electron-hole pairs which diffuse throughout the crystal and fill the traps⁶. Short LED cycles (called LED “flashes”), a few minutes long, are performed periodically every few hours to ensure that detector neutralization is maintained. Longer LED cycles (called LED “bakes”) are performed infrequently on an as-needed basis.

⁶This procedure should properly be performed with no bias across the crystal, so that charge carriers are not accelerated across the traps and collected by the charge electrodes.




Figure 12: Detail of phonon instrumentation. Each of the four phonon channel quadrants consists of 37 cells (first inset) which are themselves composed of several QET structures. The QET structures (second inset) consist of Al “fins”, shown in grey, which collect athermal phonons and funnel them into the W sensing elements, shown in magenta. White space in the insets represents non-metallized crystal surfaces, which reflect athermal phonons with near-perfect efficiency.




Figure 13: Readout circuitry of the ionization channels. The detector, represented by C_d , is biased and coupled to a JFET amplifier via C_c , which shields the JFET from the bias voltage. The feedback elements define a characteristic timescale of $R_f C_f \sim 40 \mu\text{s}$ for the ionization pulses.

2.2 Shielding

CDMS-II is a counting experiment, and as such, background rejection is of paramount importance. While electron recoils can be recorded and rejected through analysis, singly-scattering neutrons are fundamentally indistinguishable from WIMP signals in the ZIP detectors, and therefore must not be allowed to recoil in the detectors at all, as far as possible. The neutron background has two primary sources: radioactive neutron decay from radioactive materials near the experimental apparatus (radiogenic neutrons); and neutrons produced in muon showers from cosmic rays (cosmogenic neutrons). These neutrons are removed through a combination of active and passive shielding.

Within the icebox, several layers of shielding protect the detectors from neutron backgrounds [65]. From innermost to outermost, these are: six concentric cylindrical Cu cans, heat-sunk to the various heat stages of the cryogenic system⁷; a 10 cm-thick layer of polyethylene; a 22.5 cm-thick layer of Pb, the inner 4.5 cm of which consists of very radiopure ancient Pb⁸; and another 40 cm-thick layer of polyethylene. Fig. 15 shows the physical layout of these layers of shielding within the icebox.

Cosmogenically-produced neutrons are primarily removed by tagging their parent muons with the active muon shield which surrounds the icebox⁹. 40 scintillating panels are arranged so that adjacent panels overlap, and gaps are positioned so that no direct line-of-sight to the detectors exists. Panels are placed surrounding, above, and below the icebox, as shown in Fig. 16. Events coincident with veto activity are read out normally, and events are even allowed to be triggered by veto activity; however, events whose veto activity occurred within a -185 to $+20$ μs window about the event trigger are removed from WIMP-search consideration.

⁷These Cu cans also effectively shield the detectors from external α and β radiation.

⁸This ancient Pb was recovered from a sunken ship off the coast of France; for more information see [65].

⁹Excellent descriptions of the muon veto system and explanations of background estimates are available in [73, 72].




Figure 14: Orientation of ionization channels (not to scale). Note that a $\sim 1\text{mm}$ annular gap exists between the inner and outer electrodes, in addition to a $\sim 0.8\text{mm}$ gap at the outer edge of the detector. Adapted from [72].




Figure 15: Shielding layers within the icebox. The polyethylene layers (white) moderate incoming fast neutrons, while Pb layers capture them. The innermost of the the 3 Pb layers is the ancient lead. Only the outermost Cu can is visible in this image; the remaining 5 were installed inside.




Figure 16: Schematic view of the 40 scintillator panels (magenta) which form the active muon veto. LEFT: overhead view; RIGHT: side view. The F-stem is the egress on the left, and the E-stem is the egress on the right. Figure from [73].

Estimation of the magnitudes of the neutron backgrounds, along with other backgrounds, is described in Sec. 7. The combination of active and passive shielding has effectively limited the neutron background to a subdominant level, a significant factor in the excellent results of the CDMS-II experiment.

2.3 Sources and Calibration

In order to understand the detector response to various event types, two radioactive calibration sources were used. The two sources used were ^{133}Ba , a γ source, and ^{252}Cf , and neutron source. Ba data is used primarily to study the properties of bulk and surface electron recoil events for background removal, while the Cf data serves to understand the sensitivity of the analysis techniques to nuclear recoils, which form the expected signal for the WIMP search.

2.3.1 ^{133}Ba Calibration

^{133}Ba is a γ emitter (half-life ~ 10.5 years) whose energy spectrum, almost entirely below 400 keV, is well-suited to the recoil energy range for a WIMP search (≤ 100 keV). The Ba

calibration serves three important functions: it provides a means for the absolute calibration of the electron recoil energy spectrum using the 3 prominent peaks between 300 and 400 keV (see Sec. 3.5); it provides a pure sample of surface electron recoils which can be used to tune background-removal cuts (see Sec. 5.1); and it provides information to study the stability of the ionization properties of the detectors over both long (week/month) and short (hour) time scales.

Two 4 μCi sources are used when taking Ba calibration data. Both sources are mounted on ~ 3 foot long rigid wires, which allow them to be easily inserted into plastic tubing which leads near the detector volume. The sources were placed on opposite sides of the experimental volume (the NE and SW source positions) in order to achieve as uniform an illumination on the detectors as possible.

^{133}Ba calibration series were taken with regularity, averaging ~ 4 series, approximately 3 hours each, per week. Due to the higher event rate, periodic “mini-flashes” were necessary to maintain neutralization. The high frequency of ^{133}Ba series was needed, both for granularity in the time-dependence of ionization properties, but also to achieve adequate statistics in the surface event population. Over the course of the c58 data period, approximately 123 million ^{133}Ba events were recorded, of which only a few thousand were low-ionization surface events.

2.3.2 ^{252}Cf Calibration

^{252}Cf (half-life ~ 2.6 years) decays via both α emission and spontaneous fission, releasing an average of 3.7 neutrons per spontaneous fission. Due to the high activity of the source, ^{252}Cf calibration runs cannot be performed at the same high frequency of the Ba runs. Additionally, neutrons from the ^{252}Cf source cause activation peaks in nearby materials, including the 10.4 and 1.3 keV lines in ^{71}Ge ; although these peaks have a relatively short half-life (11.4 days), their presence necessitates the removal of WIMP-search data to allow

them to decay away. Over the course of the c58 data period, 6 ^{252}Cf calibration runs were performed, acquiring a total of 1.6 million events (including both neutrons and γ events).

Calibration neutrons are important for two reasons: to define the WIMP-search region with a pure sample of nuclear recoil events; and to accurately estimate the neutron acceptance of the WIMP-search analysis. The measurement of the relative energy deposited in the ionization and phonon channels (the “quenching” or “yield” factors) as a function of energy is of extreme importance when reconstructing the true recoil energies of nuclear recoil events; this analysis is discussed in more detail in Sec. 4.

Several different ^{252}Cf series were taken using both the two source positions used for the Ba calibration (NE and SW) in addition to a third (NW) source position. Careful studies, accounting for the differing source positions and theoretical expectations for the quenching factors of Ge and Si, have been performed in order to validate the energy scale of nuclear recoils used in CDMS-II analyses[74].

3 Data Processing and Event Reconstruction

3.1 The CDMS-II Data Pipeline

The CDMS-II data acquisition system consists of a triggering and readout system. Once acquired, the data pulses are run through several processing algorithms which extract complementary estimates of the physical information (e.g. energies, delays) of the events themselves. After processing, these extracted quantities (known as RQs, or *reduced quantities*) are expressed in uncalibrated units, usually digitizer bins, which are later calibrated into RRQs (*relational reduced quantities*) which have meaningful physical units (e.g. keV, μ s). There are two primary processing techniques which are important for the WIMP-search analysis: the pulse template optimal filtering (OF) and risetime-falltime walk (RTFTwalk) routines. Other, simpler processing routines, such as simple time-domain fitting (F5), pulse integrals (INT) and pulse maximum (MAX), are useful at large (MeV) energy scales or for rough analyses, but these are ignored for the purposes of a careful WIMP-search analysis¹⁰. Tbl. 3 shows some of the relevant information of the fitting routines.

3.2 Data Acquisition System

The CDMS-II data acquisition system (DAQ) is the first link in the chain from the readout electronics (see Sec. 2.1.2) to the reconstructed variables used in data analysis. The DAQ consists primarily of two components: the triggering system, which determines when the signals in the readout chain should be recorded; and the digitization system, which records and stores the raw event pulses for later processing and analysis. A more technical description of the system is available in [75].

¹⁰The paranthetical abbreviations for these fitting routines are used as suffixes in the naming schemes of the RQs.

Routine	Pulses applied to	Summary	Abbreviation
Pulse template optimal filter	Phonon	Noise-weighted frequency-domain fitting	OF
Cross-talk corr. optimal filter	Ionization	OF with cross-talk correction	OFX
Risetime-falltime walk	Phonon	Filtered “walk-along” pulse timing	RTFT-walk
Saturated pulse fitting	Ionization	Simple time-domain fitting	F5
Pulse integral	Ionization, phonon	Baseline-subtracted pulse integral	INT
Pulse maximum	Ionization, phonon	Maximum bin of data pulse	MAX

Table 3: Fitting routines used in CDMS-II data processing. The first two are of predominant importance in the WIMP-search analysis. F5 quantities are useful for saturated pulses (common for ionization energies greater than ~ 2.5 MeV); the INT and MAX quantities require fewer computational resources, but are generally too crude for use in careful analysis.

3.2.1 Triggering

Raw pulses from each detector are fed into a custom-made Receiver/Trigger/Filter (RTF) board which performs the low-level operations of the DAQ. Triggering in CDMS-II is based on phonon signals. In particular, the phonon signals from each detector and channel are filtered through a 0.5 MHz low-pass filter and compared via hardware logic boards to tuned thresholds; the logical OR of these thresholds forms the global trigger in normal WIMP-search running. All channels of all detectors are read out in normal running, regardless of which detector(s) initiated the trigger, except during calibration running, when a selective readout scheme is used to limit the high data throughput. Although only the simple phonon threshold-based condition is used to trigger the experiment, other (unused) trigger information, based on different phonon and ionization channel thresholds, is recorded for consideration in later analysis.

In order to gauge the noise environment of the experiment, noise traces are periodically

read out during time periods when no experimental triggers are present; these “random triggers” have basic pulse-rejection algorithms applied to ensure that only pure noise traces are included. Each series begins and ends with 500 randomly-triggered events which are monitored in real time as a data quality measure. Random triggers are also interspersed during the series in such a way to maintain a constant ratio of random events to physics events. The only condition imposed on a randomly-triggered pulse (a “random” for short) is that no hardware triggers are issued during the readout time window. Using these randoms, each data series has a noise power spectral density (PSD) constructed which is used in pulse reconstruction algorithms (see Sec. 3.3). Furthermore, randoms can be useful in tuning and verifying certain analysis cuts - e.g., a pulse-reconstruction goodness-of-fit test, which should pass randoms at a very high rate.

Typical operation in WIMP-search mode results in a trigger rate of approximately 0.2 - 0.4 Hz, resulting in a very high ratio of detector livetime to deadtime. While hardware limitations put an upper limit of ~ 150 Hz on the trigger rate for calibration data, actual trigger rates during calibration data are typically much lower than this (~ 20 - 30 Hz).

3.2.2 Pulse Digitization and Readout

Triggered events are digitally recorded, channel by channel, by ADC digitizers at a sampling rate of 1.25 MHz. Each pulse is recorded for 2048 samples for a pulse length of 1.638 ms; 512 of the recorded bins are taken before the trigger for future analysis consideration. In addition to the detector channels, traces from the muon veto panels are digitized at a sampling rate of 5 MHz for 1024 bins ($204.8 \mu\text{s}$), 924 of which are taken before the trigger.

All of the digitized traces and trigger information from individual events are sent to a dedicated computer, known as *Event Builder*, which organizes the data into coherent “data events” and saves them into ROOT files[76]. Additionally, many other non-science variables in the datastream (e.g., fridge temperatures) are included in these files for robustness and

investigation of anomalous experimental behavior. Single WIMP-search events have a characteristic data size of ~ 1 MB, while calibration events (which employ selective readout) are only ~ 50 kB. Typical data throughput is on the order of 50 GB per day.

3.3 Optimal Filter Pulse Fitting

Pulse template optimal filtering is a computational method useful for answering the following question: given a noisy pulse with a well-understood expected shape (*template*) and a known noise environment, what is the best (i.e. least biased) estimate one can obtain for the amplitude and timing of the pulse?¹¹ Fourier analysis of the pulses in frequency space, along with the noise PSD obtained from random pulses taken during the series being processed, provide the basis for the OF processing routine.

3.3.1 Optimal Filter Theory

A data pulse, $v(t)$, can be written in the following form.

$$v(t) = A \cdot s(t) + n(t) \quad (9)$$

Here, $s(t)$ is the expected (ideal) pulse shape, A is the amplitude (proportional to the energy deposited; see Sec. 2.1.2), and $n(t)$ is the noise. For the case of a known ideal pulse (or template) and a known spectrum of Gaussian noise, it can be shown that optimal filtering gives an unbiased estimate of the pulse amplitude with the best possible resolution[78]. Insofar as these assumptions are true for the ZIP detectors, the optimal filter quantities are, in fact, optimally reconstructed.

Noise information is encoded in the noise power-spectral density function $J(f)$, defined

¹¹The canonical treatment of this method in CDMS is given in Appendix B of [77], from which much of the information in this section originates.

in terms of the noise autocorrelation function $R(t)$.

$$J(f) = \int_{-\infty}^{+\infty} R(t) e^{-i\omega t} dt \quad (10a)$$

$$R(t) = \lim_{T \rightarrow \infty} \frac{1}{2T} \int_{-T}^{+T} n(t') n(t + t') dt' \quad (10b)$$

Note that J is expressed as a function of $f = \omega/2\pi$. In order to properly account for the noise behavior, one transforms the pulse and template into their Fourier transformations, $\tilde{v}(f)$ and $\tilde{s}(f)$, respectively. In Fourier space, the fit χ^2 is defined in the following way.

$$\chi^2 \equiv \int_{-\infty}^{+\infty} df \frac{|\tilde{v}(f) - A \cdot \tilde{s}(f)|^2}{J(f)} \quad (11)$$

One can somewhat trivially minimize this χ^2 with respect to amplitude, obtaining the following estimator and variance.

$$\hat{A} = \frac{\int_{-\infty}^{+\infty} df \frac{\tilde{s}^*(f) \tilde{v}(f)}{J(f)}}{\int_{-\infty}^{+\infty} df \frac{|\tilde{s}(f)|^2}{J(f)}} \quad (12a)$$

$$\sigma_A^2 = \left[\int_{-\infty}^{+\infty} df \frac{|\tilde{s}(f)|^2}{J(f)} \right]^{-1} \quad (12b)$$

An important caveat for this formalism is that $s(t)$ is expected to encapsulate our (perfect) knowledge of the expected pulse, both its shape and any temporal offset it may have. In practice, the template shape is often very well known, while the time shift may not be (e.g., due to timing shifts relative to the event trigger due to position within the detector). The solution to this problem is to optimize over the time shift t_0 in addition to the pulse amplitude; mathematically, the condition for this optimization is shown in Eq.

13a in terms of the filter function ϕ , defined in Eq. 13b.

$$\left. \frac{\partial}{\partial t_0} [\phi * v](t_0) \right|_{t_0=\hat{t}_0} = 0 \quad (13a)$$

$$\tilde{\phi}(f) \equiv \frac{\tilde{s}^*(f)}{J(f)} \quad (13b)$$

Note that $\phi * v$ represents convolution. Note that the optimization condition in Eq. 13a cannot be written in a simple form like that of Eq. 12a. In practice, optimization of the time shift *à la* Eq. 13a requires significant computational time, and is rarely performed in that way.

3.3.2 Implementation

Since CDMS-II data pulses are taken for finite time periods with finite sampling frequencies, the continuous Fourier transforms described above must be appropriately changed into discrete fast Fourier transforms (FFTs)¹². Several FFTs must be computed during the course of the OF algorithm, but the dominant computational time comes from transforming the individual data pulses $\tilde{v}(f) = \mathcal{F}[v(t)]$; the pulse template FFTs $\tilde{s}(f)$ and the noise PSDs $J(f)$ need only be computed once per series.

The noise PSDs $J(f)$ are computed using the randoms taken during every series. A typical noise environment is shown in Fig. 17. The cut-off at low frequency is due to the finite sampling time of random triggers ($1/1.638\text{ms} \approx 600\text{Hz}$), while at high frequency, the cut-off is due to the Nyquist theorem, which states that spectral information above half the sampling frequency is lost ($\frac{1}{2} \times 1.25\text{ MHz} \approx 600\text{ MHz}$). The “roll-off” which occurs around 10^5 Hz can be thought of as a characteristic of the time scale of noise fluctuations in the channel.

The pulse templates are constructed empirically, but in a different fashion for the phonon

¹²It goes without saying that the correct factors of 2, π , and $\frac{1}{N}$ must be determined with great care.

and ionization pulses. In each case, separate templates are created for each detector and each of the four data runs. For phonons, a simple double exponential functional form is used, with the two time constants (one controlling the rising edge, and the other the falling edge of the pulse) measured from a subset of good pulses. For ionization pulses, templates are constructed by taking a sample of good pulses, aligning the start times, and averaging them. In addition to pulse templates for the inner and outer ionization channels, cross-talk templates are also produced (i.e. using pulses from the inner electrode when a pulse is present on the outer electrode, and vice versa) and used to correct for the effects of cross-talk between the ionization channels. In this cross-talk-corrected optimal filter routine (OFX), the pulse start times of the ionization channels are fixed to be identical.

As mentioned in the previous section, optimizing the delay of the template relative to the trigger is a computationally intensive task when done rigorously (Eq. 13a). Therefore, in practice, one scans over the delay parameter t_0 in units of digitizer bins ($0.8 \mu\text{s}$), optimizing the amplitude at each value and recording the χ^2 of the fit. Finally, the delay which produces the lowest χ^2 value is chosen as the delay, and the amplitude corresponding to that delay is used. To save computational cycles, events whose amplitudes are consistent with the noise of the channel¹³ are not optimized using the full minimum χ^2 but rather by selecting the delay which maximizes the amplitude. Care was taken to ensure that this threshold was set low enough so that any potentially useful physics events would be reconstructed using the more robust minimum χ^2 routine (see Sec. 3.7).

Performing the OF (OFX) routine on a pulse (or set of pulses) always results in 3 types of quantities: an amplitude for each pulse; a delay for each pulse (or a single delay for all pulses, in the case of the OFX routine); and a χ^2 value for the fit. Each of these is useful in different aspects of the analysis - amplitudes provide energy information, delays provide position information, and χ^2 values provide a handle on the fitting quality on an

¹³As measured by the simpler INT or MAX quantities.

event-by-event basis. Of course, the OF routine has drawbacks in certain areas (e.g. it fails for saturated pulses; it provides only a single quantity to describe the pulse shape), and so it is complemented by several other routines more suited for special cases.

3.4 RTFTWalk Routine

In addition to the delay quantity provided by the OF routine, often it is useful to know more about the shape of data pulses over the full pulse length. The “rise-time fall-time walk” (RTFTwalk) routine, as its name implies, provides timing information on both the rising and falling edges of the phonon pulses. It is theoretically quite simple - one smooths the pulses using a Butterworth filter, then “walks” along the pulse while recording when the pulse reaches certain thresholds on the rising and falling edges (e.g., 10%, 20%, 30% of the maximum value). In practice, there are some subtleties which complicate the procedure.

The basic procedure of the RTFTwalk routine is shown in Fig. 18. The four phonon traces are filtered, and a baseline is established for each channel by averaging over all of the pre-pulse bins¹⁴, and the pulse height is obtained by simply finding the maximum value of the filtered pulse. As the algorithm walks along the pulse, it finds the times at which the filtered pulse crosses predefined fractions of the maximum pulse height (10% to 90%, in steps of 10%) and records the crossing which is closest to the pulse maximum; the process is repeated after the maximum is reached, recording the times at which pulse falls below the same thresholds on the falling edge of the pulse.

The nominal frequency of the Butterworth filter is 50 kHz, but this is not an optimal value for all pulses. In particular, since the RTFTwalk routine always records the crossing closest to the pulse peak, it tends to be biased *slow* for risetimes and *fast* for falltimes; furthermore, this bias increases as the amplitude of pulses decreases with respect to the noise (e.g., phonon channel A in Fig. 18). To account for this, the frequency of the Butterworth

¹⁴The start of the pulse, shown by the red vertical line in the figure, is determined from the fast ionization pulse (top panel).




Figure 17: Noise PSDs in ionization channels. Figure adapted from [79].




Figure 18: Example of the RTFTwalk routine. Raw traces are in black (Q-inner channel at top, 4 phonon channels below) and filtered traces are in blue. The vertical red line is the start of pulse, determined from the ionization channel, and the red X on each phonon channel represents the 20% risetime point as determined by the RTFTwalk routine. Figure from [80].

filter is allowed to vary with amplitude. In particular, the Butterworth frequency was parameterized as in Eq. 14.

$$\nu_B = \alpha Q^\beta \quad (14)$$

Here, Q is the signal-to-noise ratio of the pulse, while α and β were fit parameters which were set in order to optimize the resolution and to remove energy dependence in the RTFTwalk quantities as much as possible. Both the variable-frequency and constant-frequency routines were used in the processing for c58; the improvement of the variable-frequency method is shown in Fig. 19.

Using the RTFTwalk outputs, several important quantities are constructed. There are many combinations of pulse timing that can be produced, but the most useful are the following:

- **pdel**: time difference between the ionization start time and the 20% point on the rising edge of the largest phonon pulse.
- **pminrt**: time difference between the 10% and 40% points on the rising edge of the largest phonon pulse.
- **ptopwidth**: time difference between the 80% points on the rising and falling edges of the largest phonon pulse.
- **pdelCF**, **pminrtCF**, **ptopwidthCF**: identical to the above quantities, but using a constant 50 kHz Butterworth filter.

Additionally, some position information can be obtained from the timing information of the various pulses. The phonon channel nearest to the recoil will have a faster risetime than the other channels (see Sec. 2.1). Using this information, the x and y coordinates of the recoil can be estimated in the following way. Define **pdel₀** as the **pdel** of the fastest of the four phonon channels; then define **pdel_x** and **pdel_y** as the **pdel** of the phonon channel

horizontally opposite and vertically opposite the fastest channel, respectively, using the geometry defined in Fig. 11. Then we can define the quantities $\mathbf{xdel}$ and $\mathbf{ydel}$ in the following way (up to an arbitrary scaling constant).

$$\mathbf{xdel} \propto \pm |\mathbf{pdel}_0 - \mathbf{pdel}_x| \quad (15a)$$

$$\mathbf{ydel} \propto \pm |\mathbf{pdel}_0 - \mathbf{pdel}_y| \quad (15b)$$

The signs in Eq. 15 are chosen so that the axis definitions in Fig. 11 are preserved.

3.5 Energy Scale Calibration

Pulse amplitudes are proportional to energy deposited in a particular channel, but absolute calibration of energy scales requires care. In particular, even if the absolute energy scales of the ionization and phonon channels are well known, the true recoil energy of a given event requires knowledge of the *type* of recoil which deposited the energy, since Luke effect corrections will be different for electron recoil (ER) and nuclear recoil (NR) events. Conversely, the partitioning of energy between the ionization and phonon channels provides a highly effective means of *inferring* the type of recoil responsible. For clarity, we differentiate the inferred recoil energy assuming an ER from that inferred assuming an NR, using the nomenclature keV_{ee} (“electron-recoil equivalent”) and keV_{nr} (“nuclear-recoil equivalent”) to make the distinction clear.

3.5.1 Ionization Channel Energy Calibration

^{133}Ba calibration data provides a very pure sample of ER events, and its spectrum has well-known peaks at 302.9, 356.0, and 383.8 keV. Due to the geometry of the ionization channels, peaks are easily visible on the inner electrode, but difficult to distinguish on the outer electrode. For this reason, the preliminary energy scale calibration is done using

only events which have a signal consistent with noise in the outer electrode (called Q-inner events). A spectrum is formed using the uncalibrated amplitudes from the OFX processing routine, and the three peaks are identified¹⁵, and gaussian fits are used to determine the location of the peaks. These fits are then used to define a linear scaling factor to convert the OFX amplitudes into calibrated energies¹⁶. This process is illustrated in Fig. 20.

With the Q-inner spectrum calibrated, energies in the outer electrode (the “Q-outer” spectrum) can be calibrated by examining events with shared energy between the two electrodes. Pure Q-outer events (events with signals consistent with noise in the inner electrode) rarely have well-resolved peaks in the ^{133}Ba spectrum, but often the full energy of an event will be recovered through partial collection in each electrode. The strongest of the ^{133}Ba peaks, at 356 keV, is prominently seen as a diagonal line in the plane of Q-inner / Q-outer energy (see Fig. 21). By diagonalizing this shared-charge line to 45° , the Q-outer energies can be accurately calibrated.

3.5.2 Phonon Channel Energy Calibration

Phonon channel calibrations are performed through a two-part process: an absolute calibration normalizing the phonon amplitudes to physical energies; and then relative calibrations to account for gain differences among the four channels. The former is performed through comparison with the already-calibrated ionization channels, while the latter makes use of the physical constraint that, on average, phonon power absorption should be isotropic across all detector channels.

^{133}Ba calibration data is used to perform the phonon channel calibrations, since it is known to consist of a pure sample of γ ER events. For these events, we impose the

¹⁵If they are visible. If a peak is not visible, it is not used in the calibration. The 356 keV line is visible in all detectors.

¹⁶This linear fit is constrained to pass through the origin, i.e. $E_Q = k \cdot Q$, where E_Q is the calibrated energy and Q is the raw amplitude of the ionization pulse.




Figure 19: Removal of the energy dependence of RTFTwalk quantities using a variable Butterworth filter. LEFT: 0% to 20% risetime using constant frequency (top) and variable frequency (bottom) method. RIGHT: 10% to 40% risetime using constant frequency (top) and variable frequency (bottom) method.




Figure 20: Illustration of the calibration algorithm using Q-inner ^{133}Ba events. Peaks are identified from the uncalibrated spectrum, and a straight-line calibration is obtained from these peaks. Figure from [81].

requirement that the ionization energy is equal to the total recoil energy¹⁷; this allows the phonon amplitude spectrum to be rescaled to the already-calibrated ionization spectrum.

Once the overall phonon energy calibrations are performed, the relative weights of each of the four phonon channels on each detector are allowed to float to account for gain differences. From phonon physics (see Sec. 2.1) we know that at late times in the phonon pulses, down-conversion causes the phonons to become ballistic, at which point power absorption becomes isotropic since all channels have equal Al coverage. To account for this, an iterative approach is used, so that the individual channel spectra are aligned. Furthermore, one can verify proper relative calibration by inspecting the xy position plots as estimated by the relative phonon energies - for a well-calibrated detector, such a plot should be perfectly square (see Fig. 22).

3.5.3 Ionization Yield

ERs and NRs deposit energy into the ionization and phonon channels in different proportions (see Sec. 4.1); this effect is quantified through the ionization yield¹⁸, defined as the ratio of the ionization energy to the recoil energy.

$$Y = \frac{E_Q}{E_R} \quad (16)$$

For ER events, we define the yield to be equal to unity; this provides the calibration scale for the phonon channels. The true recoil energy, E_R , of an event is related to the phonon energy E_P and the ionization energy E_Q via the Luke correction factor.

$$E_R = E_P - \frac{eV_b}{\epsilon} E_Q \quad (17)$$

Here, V_b is the bias voltage of the crystal and ϵ is the electron-hole pair production

¹⁷This can also be thought of as *defining* the ionization yield for γ s to be equal to unity.

¹⁸This is also commonly referred to simply as *yield*, and less frequently the *quenching factor*.




Figure 21: Illustration of the calibration of Q-outer events. The calibrated Q-inner spectrum (horizontal axis) and the population of shared-electrode 356 keV events (diagonal line) are used to fix the energy scale of the Q-outer spectrum (vertical axis). Figure adapted from [81].




Figure 22: Illustration of the phonon energy calibration method. Raw phonon amplitudes are summed by detector and matched to the ionization energy spectrum (dashed histograms); then individual channels are redistributed to match their peaks while maintaining the overall calibration (solid histograms). INSET: phonon energy-based position estimators, before relative calibrations (red points) and after (blue points). The squareness of the blue points indicate that relative calibrations are good. Figure adapted from [82].

energy (see Sec. 2.1.2). The bias voltages were tuned such that the factors of eV_b/ϵ are very nearly equal to 1; since $Y = 1$ for ER events, one can therefore use the rule of thumb that $E_R \approx \frac{1}{2}E_P$ for ER events.

If one has a good measurement of the yield as a function of energy (as in Fig. 26), one can infer the recoil energy using only the phonon energy, if one wishes¹⁹. The definition of this phonon-only recoil energy estimator, known as **prg**²⁰, follows trivially from Eqs. 16 and 17 under the “gamma assumption” (i.e., $Y = 1$).

$$\text{prg} = \frac{E_P}{1 + \frac{eV_b}{\epsilon}} \quad (18)$$

3.6 Position Correction

Many of the important phonon quantities used in the WIMP-search analysis are subtly dependent on the event’s position within the detectors, as well as the energy of the events. It is important to attempt to correct for these anomalies, since these instrumental variations can be larger than the physical difference between NRs (signal) and surface ERs (background) and thus can affect the definition and performance of the surface event rejection cut. The process of correcting for these variations is known as *position correction*, although energy corrections are also included.

3.6.1 Physical Origins of Position and Energy Dependence

There are many mechanisms through which position and energy dependence are introduced into the phonon quantities.

- **Timing estimator variation with energy**[83]

Higher energy events are more likely to produce primary phonons close to the phonon

¹⁹The yield for ER events is defined to be unity, so a “good measurement” of the yield is equivalent to having well-calibrated phonon signals.

²⁰For “**P**honon-only **R**ecoil energy using **G**amma assumption”.

sensors, which can produce local saturation effects in individual TESs. As more and more TESs are saturated at early times, the total phonon pulse is effectively temporally reweighted. Since many timing quantities are defined relative to the pulse maximum, this has the effect of biasing timing quantities *slower* as recoil energy increases. Furthermore, the highest-energy (primary) phonon channel is often used, and this channel will usually be most affected by saturation effects.

- **Energy estimator non-linearity**[84, 85]

Saturation effects also affect the estimation of total phonon energy, through mechanisms similar to the previous effect. Simply put, saturation effects flatten the phonon pulses at higher energies, leading to underestimation of the phonon energies with larger pulses. This effect can also be position dependent, since different channels (and indeed, different TESs) may saturate at different levels.

- **Timing estimator variation with xy position**

Phonon channels signals are effectively averages of the signals collected by the many TESs; therefore, events occurring directly under the centroid²¹ of a phonon channel will have systematically faster pulses than those which are nearer to the edges of the channel. Thus, one observes that events occurring near the center or the outer edge of the detector are reconstructed with systematically slower phonon timing quantities than those occurring at intermediate radius.

- **Energy estimator variation with radial position**

Because of the constraints of photolithographic patterning of the QETs, there are no TESs near the outer edge of the detector faces (see Fig. 12). Thus, events occurring at high radius will have proportionally more of their phonons collected by

²¹The centroid of a quarter unit circle is located at a radial distance of $\frac{4\sqrt{2}}{3\pi} \approx 0.6$.

non-instrumented Al than those occurring nearer to the detector center, systematically lowering the reconstructed phonon energy of these events.

- **Energy and timing estimator variability with instrumentation**

Despite great care during the detector fabrication, imperfections in the phonon collection system can lead to position dependence. Variation in the critical temperature across a channel can lead to variation in energy response - in fact, due to the inherently non-linear nature of TESs, such variations can have large effects. In addition, TES response is strongly dependent on the thicknesses of the Al and W films due to heat capacity effects. Finally, it is possible that certain regions of phonon channels could become shorted (either from fabrication or subsequent handling) and thus be rendered insensitive to absorbed phonons.

3.6.2 Correction Mechanism

The position, timing, and energy information of the phonon channels are encoded into 5 parameters which define the 5-dimensional space over which the position correction is performed. These 5 parameters are: `xppart` and `yppart` (energy partitioning in x and y directions); `xdel` and `ydel` (event position estimated from timing delays; see Sec. 3.4); and `prg` (phonon-only recoil energy estimate). The partition quantities `xppart` and `yppart` are simply defined using the geometry of Fig. 11.

$$\text{xppart} = \frac{(pC + pD) - (pA + pB)}{pA + pB + pC + pD} \quad (19a)$$

$$\text{yppart} = \frac{(pA + pD) - (pB + pC)}{pA + pB + pC + pD} \quad (19b)$$

Here, p^* represents the calibrated energy collected in each of the four phonon channels ($*$ = A, B, C, D). Additionally, the quantities `rppart` and `rdel` are defined as the Pythagorean length of the corresponding x and y quantities (e.g., $\text{rdel} = \sqrt{\text{xdel}^2 + \text{ydel}^2}$).

We construct the position vector $\vec{q}$ of an event within this position correction space using these 5 parameters as in Eq. 20, where L_d and L_e are normalization factors chosen to regularize the relative sizes of the different dimensions, and the “hatted” vectors are orthonormal basis vectors.

$$\vec{q} = (\text{xppart})\hat{i} + (\text{yppart})\hat{j} + \frac{\text{xdel}}{L_d}\hat{k} + \frac{\text{ydel}}{L_d}\hat{l} + \frac{\text{prg}}{L_e}\hat{m} \quad (20)$$

Once this space has been constructed, a lookup table is computed using an N nearest neighbors approach. The algorithm is performed in two steps, the first creating a clean sample of bulk γ events, and the second using those events to produce the lookup table used to correct the final data.

1. Data is separated into time periods of similar detector behavior; the 3 time periods used were R125, R126, and R127-8 combined.
2. Rough quality cuts and ER band definitions are used to select an initial γ sample from ^{133}Ba data.
3. An initial lookup table is created by recording, for each event in the γ sample, the average values, over the 64 nearest neighbors in the correction space, of the quantities to be corrected. This is the first-pass lookup table.
4. Using the first-pass lookup table, the γ sample is position-corrected by rescaling each quantity for each event so that the mean of its nearest-neighbor “ball” matches the mean of the entire sample: $\text{RQc} = \text{RQ} \cdot \frac{\langle \text{RQ} \rangle_{\text{all}}}{\langle \text{RQ} \rangle_{\text{NN}}}.$
5. More stringent cuts, including a tighter ER band cut, are then placed on this first-pass corrected γ sample in order to create a very clean γ sample.
6. A second-pass lookup table is computed using the pure γ sample.

7. The final position correction is then performed on all data (calibration and low-background) using the second-pass lookup table. Since the lookup table is composed of a subset of ^{133}Ba data, correction for events not in this subset is performed by locating the closest pure γ event in the 5-dimensional correction space, and using its nearest-neighbor ball for the correction.

Quantities which have undergone the position correction retain their original names, appended with a final “c” (e.g., `pdel` $\rightarrow$ `pdelc`).

3.6.3 Results and Cross-checks

Figs. 23 and 24 help illustrate the effect of the position correction algorithm. The “shrimp plot”, so named for its characteristic shape, is useful to ensure that the correction does not accidentally mix the correction for events in the “head” of the shrimp with those in the “body”.²² Correlation plots, like Fig. 24, were used to confirm the efficacy of the position correction algorithm.

The position correction marks the final step in the CDMS-II processing pipeline. With the completion of this step, the data are ready to be used for a WIMP-search analysis.

3.7 c58 Reprocessing

The original analysis of the c58 data was published in 2010[86], but following its publication it was discovered that the OFX algorithms had used the simpler maximum amplitude routine to compute the best fits for pulses with low ionization energies, rather than the more robust (and computationally intensive) minimum χ^2 routine (see Sec. 3.3). Although the effect of this change is marginal on the amplitude of the reconstructed pulses, the reconstructed start time of the ionization pulse is sensitively dependent on this change.

²²The folding in of the “head” (known as the *manifold degeneracy*) in the shrimp plots is somewhat remedied by removing events with signals in the outer electrode, since the degeneracy occurs at high radius.




Figure 23: Shrimp plot for T1Z5. The colormap shows values of `pdel` in the plane of `rdel` and `rppart`. LEFT: Before position correction, the position dependence of the `pdel` quantity is seen very strongly. RIGHT: Following position correction, the phonon quantity is no longer as strongly dependent on the position within the detector. Note the change of the color axis scale between the left and right panels.




Figure 24: Example of a correlation plot for position correction diagnostics. LEFT: Before correction, the phonon quantity `pdel` shows a strong unphysical correlation with the radial position estimator `rdel`. RIGHT: Following position correction, the correlation is visibly eliminated. The colors of the points refer to the primary phonon channel of the events (see legend).

Since the phonon timing quantities are defined with respect to the start of the ionization pulse, the surface event rejection analysis is also sensitive to the choice of reconstruction algorithm. In particular, one of the two candidate events from the 2010 analysis would not have been in the signal region had the more robust reconstruction been used (see Fig. 25).




Figure 25: OFX event reconstruction for one of the 2010 candidate events. LEFT: Reconstructed ionization amplitude vs. OFX delay (offset). RIGHT: Fit χ^2 per degree of freedom vs. OFX delay. The red vertical line represents the delay chosen using the maximum amplitude method, while the green line represents the delay chosen using the more robust minimum χ^2 method. The difference in reconstructed delays ($\sim 4\mu\text{s}$) is large enough to affect whether or not this event would have appeared in the signal region. Figure adapted from [87].

Of course, the effect of this reconstruction problem may not only remove existing signal region events, but could also introduce new events into the signal region. Furthermore, since the analysis itself is performed using distributions of calibration events and non-signal region low background data, changes to these distributions would have affected the setting of cuts at potentially every stage of the analysis. For these reasons, it was decided to perform a full reprocessing of the c58 data, implementing the minimum χ^2 OFX algorithm in a more conservative fashion. In particular, any pulse whose magnitude, as computed by the simpler maximum amplitude algorithm, exceeded a set threshold²³ was reconstructed using the full minimum χ^2 technique. In practice, this amplitude threshold corresponded

²³This threshold was set at a level such that only pulses consistent with noise would be reconstructed using the maximum amplitude shortcut.

to about 1-2 keV in calibrated ionization energy.

Reprocessing of the c58 data began in October 2010, and took approximately 4 months to complete, with position correction following shortly thereafter. To prevent confusion, the reprocessed data sets were given the moniker “c58R”. The vast majority of the analysis described in this thesis was performed on the c58R data, with only minor or simple studies carried over from the c58 analysis. Although the data quality and event preselection cuts follow the original 2010 c58 analysis rather closely, new techniques were developed for the surface event rejection cut (see Sec. 4), and a threshold-lowering analysis was undertaken to better understand the behavior of signal region events in the standard analysis (see Sec. 6).

4 c58 Re-analysis: Event Selection

4.1 Analysis Philosophy

Simultaneous measurement of phonon and ionization energy in ZIPs allows for discrimination between events recoiling off of nuclei from those recoiling off of electrons within the crystals. Since WIMPs are dark, their interactions will predominantly be with the nuclei²⁴; thus, this discrimination is of paramount importance for WIMP-search analysis with ZIPs.

NRs are known to deposit less ionization energy (as a fraction of total recoil energy) than ERs, a phenomenon known generally as “quenching”²⁵, and which is encapsulated in the yield parameter which was defined in Sec. 3.5.3. The yield provides the primary discrimination between signal region NR events and background ER events; bulk ER events can be rejected at the level of about $10^4 : 1$ using only the yield parameter. Note that yield is defined such that $E_R \approx \frac{1}{2}E_P$ for ER events (Eqs. 16 and 17); for NRs, using a nominal value of $Y \sim \frac{1}{3}$, the corresponding rule of thumb is $E_R \approx \frac{3}{4}E_P$.

CDMS-II data are most often displayed in the plane of recoil energy and yield. Fig. 26 illustrates how ER and NR yield bands can be defined in this plane, and how the yield quantity is insufficient to remove the surface event background from the NR band signal region.

Since the partitioning of energy between between the ionization and phonon channels is necessary for yield discrimination, any artificial depression of ionization collection can degrade discrimination and allow background contaminate the signal region. Surface effects can give rise to such a depression of ionization signals. First, events occurring at high radius often have charge carriers which diffuse toward the sidewalls, which are not instrumented;

²⁴Even if a WIMP did scatter off an electron, the mass mis-match between a $511 \text{ keV}/c^2$ electron and a GeV/c^2 -scale WIMP would produce recoil energies on the eV scale - not enough even to ionize even a single charge carrier.[88]

²⁵The quenching factor for CDMS-II detectors is measured to be $\sim 30\%$, broadly consistent with other measurements[89, 90] and with Lindhard theory[91, 92].

these events are easily rejected using the outer guard electrode (see Sec. 2.1.2.2). Second, recoils occurring within $\sim 10 \mu\text{m}$ of the top or bottom detector surfaces will have depressed charge collection due to charge carriers diffusing into the “wrong” electrodes, against the applied electric field of the crystal. These so-called surface events (SEs) comprise the most difficult background to reject; they are illustrated in Fig. 26. In particular, the timing properties of the phonon pulses possess some position dependence which allows the vertical location of a recoil to be reliably inferred. This capability is what lends the **Z** to the ZIP acronym.

As discussed in Sec. 2.1.2.1, phonons created in the crystal bulk tend to travel quasi-diffusively, i.e., they remain somewhat localized until anharmonic decay allows them to travel ballistically, a process which has a characteristic time of several μs . This naturally results in a time delay between the ionization pulses, which are prompt, and the phonon pulses. However, for recoils occurring very near the detector faces, the non-ballistic “ball” of phonons may be close enough to the phonon collection instrumentation that no delay, or a decreased delay, may be observed with respect to the ionization signal.

Fig. 27 illustrates dependence of phonon pulse shape on the depth of the recoil. While the phonon timing quantities are usually reasonably well-resolved due to the superior noise performance of the phonon sensors, ionization signals become difficult to reconstruct as recoil energies approach threshold²⁶. For this reason, several complementary algorithms were employed to estimate the relevant timing quantities, and are discussed in Secs. 3.3 and 3.4. The final analysis of the phonon pulse timing is discussed in full detail in Sec. 5.

4.2 Signal Region Masking

The use of strict signal region masking has always been a priority in CDMS WIMP-search analyses, since it eliminates the possibility of unconscious bias on the part of the experi-

²⁶Problematic ionization pulse reconstruction was the primary impetus for the re-analysis effort described in this thesis; see Sec. 3 for more details.




Figure 26: Illustration of ER and NR yield bands for detector and surface event contamination. ^{252}Cf neutrons are shown in green, with ^{133}Ba γ s are shown in blue. The black crosses are ^{133}Ba surface events which leak down to lower yield (surface events). The red dashed lines indicate the fitted $\pm 2\sigma$ NR and ER bands; the lowest red dashed line is a minimum threshold on ionization energy. The flaring of these bands at low energies is due to the definition of the yield as the ratio of two independent energy measurements.




Figure 27: Phonon pulse delay behavior. A fast phonon pulse (black) can be produced by a recoil near the detector surface, while bulk recoils will produce more slowly-rising phonon pulses (blue).

mentalist; however, the nature of a reanalysis renders the concept of traditional masking somewhat moot. Nevertheless, in the spirit of mirroring the 2010 analysis as closely as possible, the ostensible trappings of a masked analysis were maintained as much as was possible.

Calibration data requires no masking, but it is often useful to look at low background WIMP-search data which is outside of the signal region, henceforth referred to as *sideband data*. WIMP-search data was masked if it met all of the following criteria:

1. The event is veto anti-coincident - i.e., there is no veto signal within 50 μ s before event trigger, and the event is not triggered by the muon veto.
2. The event has negligible energy collected in the outer electrode, between -5 keV and +5 keV.
3. The event has a recoil energy above the phonon noise and less than 130 keV.
4. The event lies within the $\pm 3\sigma$ NR band.
5. The event is a single scatter - i.e., no other detectors have phonon energy greater than 6σ above the phonon noise.

These masking criteria were chosen to be intentionally loose, so that the masking would be as conservative as possible without removing an unreasonable amount of data. All low background data satisfying all of the above criteria were not used to tune the cuts for the WIMP-search analysis.

4.3 Event Pre-selection

Before the proper physics analysis can be undertaken, poor data must be removed so that a pure sample of events is available for the WIMP search. The cuts used to remove these poor data generally take on two forms - cuts which uniformly remove events over fixed

time periods, and cuts which employ event-by-event removal. The former are simple to account for, since they simply result in a reduced livetime, whereas the latter are assigned an efficiency (measured as a function of energy) to account for the non-uniformity of their event removal.

4.3.1 Time Period Removal

Bad time periods are selected through a variety of methods. A number of series had been marked as poor by experimental operators during the data collection, and were removed, along with time periods during which experimental settings (e.g. QET bias) were not stable[93]. A Kolmogorov-Smirnov (KS) test approach was used to check for series which had distributions of ionization and phonon quantities which varied drastically from the long-term normal behavior for those distributions.[94] Additionally, periods of extremely high noise on ionization channels[95], high trigger rates[96], and suspected He films[97] were identified and removed. Finally, data taken during time periods when the NuMI beam was active were removed.

Although many of these cuts are detector dependent, they affect most detectors in a similar fashion. After the removal of these time periods, the total livetime for the c58R analysis was approximately 200 live days. Tbl. 4 displays the raw livetime of the WIMP-search analysis after basic quality cuts are applied.

4.3.2 Event-by-event Removal

There are several event-by-event cuts that are implemented to ensure the quality of the data for the physics analysis.

- **Good baseline noise**[99]

If the baseline of a particular detector channel is misbehaving when a pulse is recorded²⁷,

²⁷Certain detectors appear to be sensitively coupled to the cryogenics systems, which can induce transient

Detector	T1	T2	T3	T4	T5
Z1	-	170.6	-	223.1	-
Z2	211.3	209.7	76.7	186.1	-
Z3	-	158.1	148.3	226.7	207.0
Z4	62.8	204.0	231.4	217.5	213.4
Z5	166.4	181.9	146.1	202.0	210.1
Z6	-	-	218.4	215.0	-

Table 4: Livetimes of detectors after application of the time-period removal and poor ionization baseline cuts. All numbers are given in units of days. Green cells indicate Ge detectors used in the final analysis, yellow cells indicate Si, and blank cells are used for detectors not used in the WIMP-search analysis. Numbers are taken from [98].

the pulse may be poorly reconstructed (affecting both amplitude and shape). For a channel with a good baseline, any variations in the digitizer bins taken before the trigger will come from the ambient (Gaussian) noise in the channel. Thus, events with poor baseline behavior can be rejected through analysis of these prepulse bins.

For every event, the prepulse bins are histogrammed and fit to a Gaussian. If the width of this fit is more than 4σ away from the mean of the width distribution for that channel, it is rejected. This analysis is performed separately on each of the ionization and phonon channels. Although this is an event-by-event cut, its behavior is not energy dependent. The ionization version of this cut does have some significant time dependence for some detectors, and is therefore treated as a livetime-removing cut; the phonon version has no such time dependence and is therefore treated with the efficiency interpretation. Its efficiency, which is constant with energy, was measured with WIMP-search sideband data to be $> 99\%$ on all detectors except T5Z5, for which it was 93%.

- **Physical phonon pulses**

There are two recurrent reconstruction problems with phonon pulses: pulses with noise of this type.

negative reconstructed amplitudes, and pulses with unphysical timing quantities. The former can be caused by a shifting prepulse baseline, too slow to be caught by the good baseline noise selection, or from “pileup” events, where two events occur close enough in time that the later pulse rides on the tail of the previous pulse. The latter phenomenon comes about due to peculiarities in the choice of time windows for the OF reconstruction algorithm. Both effects can be removed very effectively using simple cuts on phonon energy[100] and 20% phonon risetime[101], respectively. The efficiency of these cuts, measured on ^{252}Cf calibration data, is nearly always 100%.

- **Ionization pulse reconstruction quality[102]**

The OFX routine returns a χ^2 goodness-of-fit quantity for the ionization reconstruction of each event (see Sec. 3.3); this provides an important handle on the quality of the event reconstruction.²⁸ The distribution of the fit χ^2 quantities displays some energy dependence, so the cut is set in the plane of χ^2 vs. summed (inner and outer) ionization energy.

An illustration of this cut is shown in Fig. 28. The χ^2 quantity is binned into energy bins, and the cut position is chosen such that the ratio of the histogram at the cut position to that at the maximum value is equal to the same ratio for a Gaussian distribution with a $\mu + 3\sigma$ cut; therefore, we say that the cut is set at the 3σ Gaussian-equivalent level. The χ^2 distribution rises somewhat with energy, but the fact that it is centered near $\chi^2 = 4000$ at low energies is indicative of a good reconstruction, since the number of degrees of freedom ≈ 4096 , the number of bins in the digitized ionization pulses.

By virtue of the Gaussian-equivalent formulation of this cut, its efficiency is high and approximately constant above ~ 20 keV in total ionization energy, below which

²⁸Since the phonon timing quantities are dependent on the reconstructed start time of the ionization pulse, the quality of the ionization pulse reconstruction is important for phonon timing quantities as well.

it usually degrades to about 90%. The efficiencies for this cut were measured as a function of total ionization energy²⁹ on ^{252}Cf calibration data and fit to a functional form for each detector and run.

4.4 Fiducial Volume Selection

The detector fiducial volume is defined using the outer guard electrode as a veto against events occurring at high radius within the detectors.[103] In particular, the fiducial volume cut is set such that an event passes the cut (“is fiducialized”) if and only if the ionization signal in the outer electrode is consistent with noise (i.e., within $\pm 2\sigma$ of the noise mean). A visualization of this cut is shown in Fig. 29.

At high energies, reconstructed energies are more reliable and the cut has an efficiency of $\sim 85\%$, as one would expect from the detector geometry. At lower energies, even events which deposit non-zero energy in the outer channel will nonetheless be reconstructed with noise-consistent energy, leading to a degradation of the fiducial volume cut accompanied by a rise in efficiency.

The efficiency for the fiducial volume cut is measured using ^{252}Cf data, but care must be taken to correct for the pollution of the neutron sample by low-yield ERs. For well-fiducialized events, ER-NR yield separation is quite good; however, due to subtle effects in charge collection near the inner / outer electrode boundary (e.g., oblique propagation of charge carriers) low-yield γ leakage can be more significant for events with measurable outer electrode energy. For this reason, each detector is separated into 3 regions - “ear”, “funnel”, and “nose”³⁰, which roughly correspond to outer, shared, and inner ionization collection, respectively - wherein the γ correction is applied (see Fig. 30). These regions

²⁹To obtain the efficiency of this cut on NRs, one can easily transform the ionization energy to recoil energy using the measured NR yield and Eq. 16.

³⁰The *ear* and *funnel* nomenclature has been prevalent in CDMS for some time, but I attribute the naming of the *nose* region to extemporaneous comment by A.N. Villano.




Figure 28: Illustration of the χ^2 ionization reconstruction quality cut (adapted from [102]). LEFT: Distribution of the χ^2 variable in different energy bins. c58 data is shown in red (“Old”), while c58R data is shown in black (“New”). RIGHT: Cut boundary in the χ^2 vs. summed ionization energy plane. The black crosses are the cut positions set in the four energy bins, and the black line is a quadratic fit to these points. Events with χ^2 values below the boundary are accepted.




Figure 29: Visualization of the fiducial volume cut. Only events lying within $\pm 2\sigma$ of the noise mean for the outer electrode (between the pink lines) are fiducialized. The noise distribution was fit to a Gaussian in several energy bins using ^{133}Ba calibration data. Figure adapted from [103].

are prominent in the Ge detectors due to the crystal structure, but less prominent in the Si detectors.

The γ correction is performed by defining three yield bands: a $\pm 2\sigma$ NR band; a $\pm 2\sigma$ ER band; and a “between bands” (BB) region, at least 3σ above the NR band mean and 3σ below the ER band mean. Using these bands, one computes the correction as follows for each of the three ionization regions.

1. Count the ^{252}Cf events which lie in the between-bands region:

$$N_{\text{BB}}^{\text{Cf}}$$

2. Compute the ratio of ^{133}Ba events in the NR band to those between the bands:

$$\rho = \frac{N_{\text{NR}}^{\text{Ba}}}{N_{\text{BB}}^{\text{Ba}}}$$

3. Count the ^{252}Cf events which lie in the NR band:

$$N_{\text{NR}}^{\text{Cf}}$$

4. Estimate the number of ^{252}Cf neutrons leaking upward into the between-bands region, assuming a Gaussian yield distribution denoted by $f(Y)$:

$$\hat{N}_{\text{BB},n}^{\text{Cf}} = N_{\text{NR}}^{\text{Cf}} \times \frac{\int_{\text{BB}} f(Y) dY}{\int_{\text{NR}} f(Y) dY}$$

5. Estimate the number of true γ events between the bands in the ^{252}Cf data:

$$\hat{N}_{\text{BB},\gamma}^{\text{Cf}} = N_{\text{BB}}^{\text{Cf}} - \hat{N}_{\text{BB},n}^{\text{Cf}}$$

6. Estimate the leakage of true γ events in the ^{252}Cf NR band:

$$\hat{N}_{\text{NR},\gamma}^{\text{Cf}} = \hat{N}_{\text{BB},\gamma}^{\text{Cf}} \times \rho$$

7. Count the ^{252}Cf events in the NR band which pass the fiducial volume cut:

$$N_{\text{NR},\text{FV}}^{\text{Cf}}$$

8. Compute the fraction of ^{133}Ba events in the NR band which pass the fiducial volume cut:

$$\eta = \frac{N_{\text{NR},\text{FV}}^{\text{Ba}}}{N_{\text{NR}}^{\text{Ba}}}$$

The efficiency of the fiducial volume cut can then be computed, accounting for γ contamination, using Eq. 21. The index i represents summation over the 3 ionization regions defined in Fig. 30.

$$\text{Efficiency}_{\text{FV}} = \frac{\sum_i (N_{\text{NR},\text{FV}}^{\text{Cf}} - \eta \cdot \hat{N}_{\text{NR},\gamma}^{\text{Cf}})}{\sum_i (N_{\text{NR}}^{\text{Cf}} - \hat{N}_{\text{NR},\gamma}^{\text{Cf}})} \quad (21)$$

4.5 Yield Discrimination and Thresholds

The WIMP-search signal region is defined in the yield vs. recoil energy plane by 4 different cuts: the NR band cut; the below ER band cut; the ionization energy threshold cut; and the recoil energy threshold cut.³¹ When combined, these four cuts define a region of parameter space over which the ZIP detectors provide excellent background rejection. A schematic representation of these cuts is shown in Fig. 31.

• Nuclear Recoil Band Cut

³¹Because these four cuts define a unique 2-dimensional shape in the yield vs. recoil energy plane, they are sometimes referred to as *event topology* cuts.

The NR band is defined using ^{252}Cf calibration data.[104] The data are binned by energy, and a Gaussian is fit to the yield distribution in each bin. The mean and standard deviation are then fit as a function of energy, and these fits are used to define the band. The standard NR band cut is $\pm 2\sigma$, although this is sometimes altered to optimize the surface event rejection.

The efficiency of the NR band cut is measured on ^{252}Cf data, as a function of energy, by comparing the number of events within the $\pm 2\sigma$ band to the number of events within the $\pm 4\sigma$ band.

$$\text{Efficiency}_{\text{NR}} = \frac{N_{\pm 2\sigma}^{\text{Cf}}}{N_{\pm 4\sigma}^{\text{Cf}}}$$

The measured efficiency is roughly constant with energy, and generally agrees with the Gaussian expectation of $\sim 95\%$.

- **Below Electron Recoil Band Cut**

The ER band is defined using ^{133}Ba data in a method similar to that used for the NR band.[105] The below ER band cut rejects any events that are not at least 3σ below the mean of the ER band[106]; this cut is only more restrictive than the NR band cut at low energies where the yield bands begin to flare (see Fig. 31). Assuming that the ER band is Gaussian, this ensures bulk ER rejection of nearly 10^{-3} down to low energy.

The below ER cut has 100% efficiency on NRs at recoil energies above ~ 10 keV, but falls off quickly at low energies as it begins to cut into the NR band. The efficiency of this cut is measured on ^{252}Cf data, and was found to be well described by an error function.

- **Ionization Threshold Cut**

Since yield, the primary quantity for rejecting ER backgrounds, is ionization-based,

we set a threshold on ionization energy to ensure high signal-to-noise for the WIMP-search analysis. The noise behavior of the inner electrode is measured using randoms; this noise spectrum is then fit to a Gaussian³², and the standard cut is set at 4.5σ above the mean of the noise.[107]

Similar to the below ER band cut, the ionization energy threshold cut is only important at low recoil energies, where it begins to cut into the NR band (see Fig. 31). The efficiency of this cut is measured with ^{252}Cf data, and is well-described by an error function.

• Recoil Energy Threshold Cut

A recoil energy threshold is imposed on all detectors for two reasons: to remove the low-energy region where the yield bands flare and signal-to-noise is poor;³³ and to restrict our signal region to an area where sufficient calibration SE statistics exist from ^{133}Ba data to make accurate estimates of the background. A 5σ below ER band cut (instead of the standard 3σ) is used to define SEs from ^{133}Ba data, so statistics from this population begin to fall off at recoil energies higher than where one might expect from the standard definition of the signal region.

For the standard “high threshold” CDMS-II analysis, recoil energy thresholds were set conservatively in order to achieve a zero-background experiment. In practice, this led to 10 keV thresholds for all of the Ge detectors, and 7 keV thresholds for most of the Si detectors, although some were set higher. For the “mid threshold” analysis, the recoil and ionization energy thresholds were lowered and the zero-background criterion was dropped in order to increase sensitivity to low-mass WIMPs; for more details, see Sec.

³²Due to the use of the minimum χ^2 OF for random triggers, the noise blob for ionization channels is actually a double Gaussian, one peaked at positive energy and the other at negative energy. The Gaussian fit is performed on the positive energy peak.

³³The ionization threshold and below ER band cuts also perform this function, but the recoil energy threshold is more conservative.

6. The efficiency of the recoil energy threshold cut is trivial: it is simply a Heaviside step function centered on the threshold energy.

4.6 Single Scatter Selection

Because of the exceedingly small observed rate of WIMP scatters in direct detection experiments, it is logical to conclude that the rate of WIMP multiple scatters in the CDMS-II experiment is essentially zero. For this reason, we implement a single-scatter condition for the WIMP-search analysis, defined by requiring a large signal in one detector and negligible signals in all others (relative to the noise environment of each detector).³⁴

Since the phonon channels generally possess noise properties which are superior to the ionization channels, they are used to determine the multiplicity of a given event. In particular, the total phonon signal (summed over the four channels) is used as the energy estimator for the single-scatter selection. The only exception to this was detector T1Z1, which seemed to have problematic phonon channels; therefore its ionization channels were used for multiples tagging.[108] Note that all 30 detectors are used for multiples tagging, even though not all detectors are carried forward for WIMP-search analysis. The single-scatter cut is defined in the following way:

- The primary detector must have a signal at least 6σ above the noise mean.
- All other detectors must have a signal not more than 4σ above the noise mean.

The noise spectra are measured using randoms and fit to a Gaussian for every series. This allows the cut to vary with the noise behavior; since the condition on all other detectors is tighter than the primary detector condition, this is conservative.[109]

The efficiency of the single-scatter selection cut must be carefully defined. To be clear, what is meant by “efficiency” for this cut is the fraction of true single-scatter events which

³⁴Although having many detectors increases the number of surfaces (and thus the severity of surface effects), it also vastly improves multiples tagging as compared to monolithic experiments.




Figure 30: The ionization regions for Ge detector T1Z5, defined by eye. The “ear” (blue) consists of events which are primarily collected in the outer electrode; “nose” (green) events are collected in the inner electrode; “funnel” (red) events are shared. The horizontal (vertical) axis is the ionization yield computed using only the inner (outer) ionization energy.




Figure 31: Example of the yield and threshold cuts for detector T3Z4. The nuclear recoil band (magenta curves), the below electron recoil band cut (green curve), the ionization energy threshold (red curve), and the recoil energy threshold (cyan line) define the signal region (shaded yellow). Sparsified ^{133}Ba data is overlaid to illustrate bulk ER rejection, and also to highlight the problem of SE leakage.

are passed by the single-scatter selection. For true single-scatters, failure of this selection can only be due to statistical fluctuations of noise in other detectors. In order to measure the efficiency of this cut, the failure rate of random triggers in other detectors is measured; the efficiency is then just 1 minus this fraction. Obviously, this efficiency is independent of energy, and is $\sim 99\%$ for all detectors.

5 c58 Re-analysis: High Threshold Phonon Timing Analysis

5.1 Philosophy of Phonon Timing Analysis

The final and most important cut of the WIMP-search analysis is the phonon timing cut, which removes low-yield SEs, which mimic true NRs, from the signal region (see Sec. 4.1). The general strategy for dealing with these SEs is to examine the timing of the phonon pulses, since SEs have characteristically faster pulses than events in the bulk; however, how this analysis should be performed in detail is not straightforward. The reprocessing of the c58 data provided an opportunity to test several parallel timing analyses and compare their effectiveness.

To perform a timing analysis, one must first construct a discriminator variable, henceforth referred to as the “timing parameter”. Given such a timing parameter, the task becomes optimizing a cut on this variable such that the sensitivity to WIMPs is maximized. There is some ambiguity in this method which arises from the choice of WIMP mass at which to optimize the sensitivity; in general, maximum sensitivity for a given target material will be achieved at WIMP masses near the mass of the target nucleus, due to kinematic effects.

The timing cut threshold is determined separately for each detector; in addition, to encapsulate the energy dependence of the timing quantities, the phonon timing cut is set in 3 different recoil energy bins: less than 20 keV; 20 to 30 keV; and 30 to 100 keV. I will use a subscript z to index different detectors, while a subscript e will be used to index energy bins (where $e = 1, 2, 3$ refers to the 3 energy bins in the order they are defined above). Therefore, a timing cut is completely defined by the timing quantity and $3 \times N$ cut positions $\eta_{z,e}$, where N is the number of detectors used in the analysis.

In order to gauge the effectiveness of such a cut, it is important to understand both how well it removes background, as well as how much signal it removes. The background

removal is quantified with a leakage estimate, $L_{z,e}$, defined as the expected number of SEs which will pass the timing cut in a given detector and energy bin (see Sec. 5.2). The amount of signal acceptance is quantified with the *spectrum-averaged exposure* (SAE), which is a functional of the efficiency. Mathematically, the SAE is defined in Eq. 22.

$$\text{SAE} = \frac{\int \text{Eff}(E) \frac{dR}{dE} dE}{\int \frac{dR}{dE} dE} \quad (22)$$

Here, $\frac{dR}{dE}$ is the WIMP energy spectrum at a given WIMP mass[110] and the integral is taken over the energy range of the analysis (i.e., where the efficiency is not identically zero). These quantities, $\text{SAE}_{z,e}$ and $L_{z,e}$, are measured as functions of the cut positions $\eta_{z,e}$; with these quantities in hand, the optimization procedure essentially becomes a maximization problem for $\sum_{z,e} \text{SAE}_{z,e}$ with constraints on $\sum_{z,e} L_{z,e}$.

5.1.1 Exposure and Leakage Fits

To optimize the phonon timing cuts, it is important to understand how the leakage and exposure vary as the timing thresholds are adjusted. The SAE, which is computed using Eq. 22, requires measurements of the efficiency using ^{252}Cf data, while the leakage is estimated using SEs from ^{133}Ba data (see Sec. 5.2). Since the calibration data, especially the ^{133}Ba , is statistics-limited and the tails of the distributions are sparsely populated, the CDFs of the calculated SAE and leakage are fit to functional forms for the purposes of optimizing the timing cut.

An example of these fits is shown in Fig. 32. Generalized Pareto functions, whose functional form is shown in Eq. 23, are used to fit the tails, defined as the final 10 data points from calibration data, while Gaussian kernel smoothing is used for the bulk of the distributions. Finally, the smoothed bulk and the fitted tails are stitched together using a

cubic spline interpolation where they meet.

$$f(x) = \begin{cases} (1 + \xi \frac{x-\mu}{\sigma})^{-1/\xi}, & \frac{x-\mu}{\sigma} \geq 0 \\ 1, & \frac{x-\mu}{\sigma} < 0 \end{cases} \quad (23)$$

Despite the drawbacks of using smoothed and extrapolated fits to gauge the performance and optimization of the phonon timing cuts, it is important that we do so for two reasons: (1) if extrapolations could not be used, we could never reliably tune the phonon timing cut past the farthest outlier point available from the calibration data; and (2) the slope-matching optimization method cannot be implemented without differentiable functions for the SAE and leakage as a function of timing threshold (see Sec. 5.1.2).

It is important to note that, although the use of fits does introduce systematics into the *optimization* procedures for the phonon timing cut, these systematics are confined to the optimization - in other words, once the cut is set (by whatever method), any systematics involved in the optimization do not carry over into the final analysis using that cut.

5.1.2 Optimization of Phonon Timing Cut

There are two optimization methods that were used for the various c58R phonon timing analyses: the *leakage-scanning* method, and the *slope-matching* method.

5.1.2.1 Leakage-scanning method

The leakage-scanning method makes use of an objective function, defined in Eq. 24.

$$f(\vec{\eta}) = \left(1 - \frac{\sum_{z,e} \text{SAE}_{z,e}(\eta_{z,e})}{(\sum_{z,e} \text{SAE}_{z,e})_{\max}}\right)^2 + \alpha \left(1 - \frac{\sum_{z,e} L_{z,e}(\eta_{z,e})}{L_{\text{target}}}\right)^2 \quad (24)$$

The vector $\vec{\eta}$ represents the complete set of values $\eta_{z,e}$ which define the timing cut, and




Figure 32: Example of leakage fitting for phonon timing cut optimization. A similar method is performed for the SAE. The survival function of calibration SEs from ^{133}Ba data (blue) is smoothed for the bulk of the distribution, and fit to a Pareto function for the final 10 points of the tail. The fit is used to extrapolate the leakage behavior beyond the calibration data points. INSET: The red vertical line denotes the cross-over point from bulk smoothing (to the left) to tail fitting (to the right). The discontinuity between the smoothed and fit data is patched with a cubic spline, resulting in the black curve.

α is a weighting factor which determines the relative importance assigned to maximizing the exposure versus adhering to the target leakage; a value of $\alpha = 10$ was used for the analyses in this thesis. Given an input value of L_{target} , the set of timing thresholds $\vec{\eta}$ which simultaneously minimize the objective function f defines the optimized phonon timing cut.

The minimization of the objective function is done iteratively. To begin, each energy-bin-detector combination has its timing threshold set such that the leakage is equally apportioned. Using this as an initial guess, the objective function is properly minimized. If any detector takes up an inordinate amount of leakage (more than 15% of the excess leakage above the target), it is discarded (its timing threshold is set to infinity) and the process is repeated until no more detectors are discarded.

Since this method requires an input value for L_{target} , one scans over values of target leakage (as implied by the name of the method) to find the best cut. For the analyses in this thesis, the target leakage was scanned from 0.1 to 1.2 in increments of 0.1; in past analyses, $L_{\text{target}} = 0.5$ was chosen beforehand as the optimal value³⁵, although this value is allowed to float in these analyses.

By following the above optimization procedure, one effectively obtains the SAE as a function of target leakage. Using this information, we construct the counting rate sensitivity estimator S_{90} in Eq. 25.

$$S_{90}(L_{\text{target}}) = \frac{N_{90}(L(L_{\text{target}}))}{\text{SAE}(L_{\text{target}})} \quad (25)$$

Here, $N_{90}(n)$ is the 90% upper limit on a Poisson random variable given an observed number of counts n , and $L(L_{\text{target}})$ and $\text{SAE}(L_{\text{target}})$ are the expected leakage and SAE summed over all detectors and energy bins when optimized at a target leakage L_{target} . The optimized cut is chosen to be the one which minimizes the sensitivity estimator S_{90} (see Fig. 33).

³⁵This was done because for 0.5 expected leakage events, one expects a background-free experiment most of the time ($\sim 60\%$ assuming Poisson statistics).

5.1.2.2 Slope-matching Optimization Method

The slope-matching method uses the same characteristic quantities, SAE and L , but an analytical criterion is used to find the optimal cut thresholds, rather than the numerical approach taken in the leakage-scanning method. However, this analytical method necessarily relies more heavily on the fits of SAE and L with respect to timing threshold, so systematics from this fitting are more important (see Sec. 5.1.1).

Given fits for the SAE vs. cut threshold and L versus cut threshold, one can stitch them together to obtain SAE as a function of L for each of the detectors and energy bins. The optimization criterion for the slope-matching method is this: *“for a given target leakage, the rate of change of the exposure with respect to leakage in each detector and energy bin must be identical.”* Since the total target leakage is fixed, this is optimal since small redistributions in leakage will not increase the SAE. Mathematically, the criterion is shown in Eq. 26, where C represents the equal-slope value, which can be allowed to float.

$$\frac{d\text{SAE}_{z,e}}{dL_{z,e}}(\eta_{z,e}) = C \quad \text{for all } z, e \quad (26)$$

Fig. 34 illustrates the use of this method. The horizontal line represents the choice of a value for C . One can scan over values of C to obtain a curve of sensitivity vs. L_{target} . As with the leakage-scanning method, the cut which corresponds to the target leakage with the minimum value of S_{90} is chosen as the optimal phonon timing cut.

5.1.3 Summary of Various Timing Analyses

Three timing analyses were performed on the Ge data: a simple timing analysis called the “classic” analysis, which is identical in design to that performed during the original 2010 analysis[86]; a machine-learning-based approach, called the “neural network” analysis; and a distribution comparison approach, called the “5D χ^2 ” analysis. The Si data were analyzed




Figure 33: Example of the leakage scanning method. The blue curve (left axis) shows the spectrum-averaged exposure as a function of the target leakage; using this information, the 90% counting rate sensitivity is computed (red curve, right axis). The phonon timing cut corresponding to the target leakage which minimizes the rate sensitivity is the optimal one. The example shown is for the Ge neural network analysis.




Figure 34: Example of the equal slopes method. LEFT: Each curve represents a different energy-bin-detector combination, and the horizontal line corresponds to the choice of C . The timing thresholds are fixed by the intersection of the horizontal line with the different curves. RIGHT: The green curve is constructed by scanning over values of C , and the point which minimizes the counting rate sensitivity corresponds to the optimized phonon timing cut. Figure adapted from [111].

later[112], and utilized a modified version of the 5D χ^2 method using only two variables, and thus it was called the “2D χ^2 ” analysis. The remainder of this section describes the high threshold, zero leakage analyses performed with these four timing cuts. Sec. 6 details how these analyses were extended to lower thresholds.

5.2 Surface Event Leakage Estimation

SE leakage estimation is important for both the optimization of the phonon timing cut (before unmasking the signal region) and for the interpretation of results (after unmasking; see Sec. 7.3.4). Initial leakage estimates must be computed using calibration and sideband data, but after unmasking, WIMP-search data can be used to compute a final leakage estimate.

There are two major pieces of information which go into the leakage estimation procedure: an estimate of the total SE event background in the absence of a phonon timing cut (the normalization), and the passage fraction of the phonon timing cut. The normalization is outlined in Sec. 5.2.1, while the different methods for computing the passage fraction are discussed in Sec. 5.2.2.

5.2.1 Normalization of Leakage Estimates

The fundamental idea used to estimate the number of events in the signal region is to look at the unmasked data from the first 5-tower runs of CDMS-II (runs 123 and 124, also known as c34) and scale them appropriately to the c58 data. 3 techniques were developed, each with their own combinations of statistical and systematic uncertainties.

For simplicity, we adopt the following abbreviations for event classes pertinent to these calculations. Unless stated otherwise, they refer to WIMP-search data.

- **NRSS: nuclear recoil single scatters.** Signal-region events which scatter in only one detector.

- **NRMS**: nuclear recoil **m**ultiple **s**catters. Signal-region events which scatter in multiple detectors.
- **WBSS**: wide **b**and **s**ingle **s**catters. Events outside the signal region (but also not in the ER band) which scatter in only one detector.
- **WBMS**: wide **b**and **m**ultiple **s**catters. Events outside the signal region (but also not in the ER band) which scatter in multiple detectors.

Using these definitions, we define the three estimates for the normalization as follows.

1. Scale the observed c34 NRSS counts by the live-time ratio of c58 to c34.

$$N_1 = \text{NRSS}_{c34} \times \frac{T_{c58}}{T_{c34}}$$

2. Scale the observed c58 NRMS counts by the ratio of c34 NRSS to NRMS.

$$N_2 = \text{NRMS}_{c58} \times \frac{\text{NRSS}_{c34}}{\text{NRMS}_{c34}}$$

3. Scale the observed c58 NRMS counts by the ratio of c58 WBSS and NRSS to WBMS and NRMS.

$$N_3 = \text{NRMS}_{c58} \times \frac{\text{NRSS}_{c58} + \text{WBSS}_{c58}}{\text{NRMS}_{c58} + \text{WBMS}_{c58}}$$

Tbls. 5 and 6 display the results of these calculations applied to the detectors carried forward for WIMP-search analysis (statistical uncertainties only), using the high-threshold energy ranges. N_1 has poor statistics, and its primary systematic uncertainty is due changing detector performance between c34 and c58. N_2 has better statistics, but introduces a systematic uncertainty due to differing behavior between single and multiple scatters. N_3 has the best statistics, but also suffers from significant systematic effects due to possible

Detector	N_1	N_2	N_3
T1Z2	15.30 ± 4.40	33.00 ± 14.83	11.48 ± 2.65
T1Z5	7.70 ± 2.10	12.89 ± 5.31	22.08 ± 3.90
T2Z3	7.50 ± 2.10	13.16 ± 6.87	17.24 ± 3.99
T2Z5	8.60 ± 2.40	4.40 ± 3.24	19.64 ± 3.83
T3Z2	2.20 ± 0.50	1.81 ± 1.87	11.89 ± 3.15
T3Z4	6.50 ± 1.30	7.31 ± 4.03	12.04 ± 2.42
T3Z5	4.00 ± 0.80	1.72 ± 1.80	4.83 ± 1.53
T3Z6	20.90 ± 4.80	35.44 ± 17.24	95.11 ± 20.95
T4Z2	5.40 ± 1.10	6.53 ± 3.60	12.46 ± 2.74
T4Z4	6.30 ± 1.30	2.49 ± 1.55	7.11 ± 1.84
T4Z5	5.80 ± 1.20	6.49 ± 3.23	12.90 ± 3.05
T4Z6	20.70 ± 4.70	24.14 ± 8.57	80.09 ± 14.99
T5Z4	6.10 ± 1.20	5.45 ± 3.31	18.73 ± 3.66
T5Z5	6.10 ± 1.20	7.31 ± 4.03	9.35 ± 2.34
Sum	123.10 ± 9.43	162.15 ± 27.43	334.97 ± 27.81

Table 5: Estimates of the total number of NRSS in the c58 WIMP-search data for the Ge detectors.

discrepancies between single and multiple scatters, and between wide band and signal region events. It is also not tied to c34 data counts in any way.

N_2 was chosen as the most useful estimator for the leakage normalization, due to its good mix of systematics and statistics. In particular, the systematics of N_1 (assuming identical detector performance between c34 and c58) and N_3 (wideband / signal region and singles / multiples) were deemed too problematic.

5.2.2 Leakage Estimation Methods

5.2.2.1 Method 1

The simplest way to estimate the SE leakage is to multiply the expected total number of NRSS (see Sec. 5.2.1) by the passage fraction of WIMP-search NRMS. This is called Method 1 (or Method 1 WS, to indicate that WIMP-search data is used to compute the pass-fail

Detector	N_1	N_2	N_3
T1Z4	8.95 ± 3.38	8.68 ± 4.03	7.48 ± 4.68
T2Z1	44.22 ± 5.39	35.88 ± 8.55	38.52 ± 11.16
T2Z2	8.55 ± 2.18	6.00 ± 1.97	9.46 ± 3.16
T2Z4	21.96 ± 3.23	27.57 ± 5.90	11.20 ± 2.97
T3Z3	27.47 ± 5.92	30.13 ± 8.16	18.53 ± 4.98
T4Z1	18.18 ± 2.81	16.47 ± 5.04	14.61 ± 4.67
T4Z3	6.67 ± 2.75	5.13 ± 2.21	7.09 ± 3.00
T5Z3	23.71 ± 4.31	16.11 ± 3.81	10.94 ± 2.67
Sum	159.69 ± 11.17	145.97 ± 15.47	117.83 ± 15.10

Table 6: Estimates of the total number of NRSS in the c58 WIMP-search data for the Si detectors. Information from [113].

ratios). The method is given mathematically in Eq. 27.

$$n_z = N_z \times \frac{b_z}{B_z} \quad (27)$$

Here, z indicates the detector, N represents the total NRSS estimated as in Sec. 5.2.1, and B represents actual NRMS counts from data. Upper case letters indicate the number of events before the phonon timing cut is applied, and lower case letters indicate the number of events which pass the phonon timing cut.

Although Method 1 is straightforward, it suffers from several systematics, such as variation in passage fraction with recoil energy, and differences in passage fractions for SEs on the two faces of the detectors. Another method was therefore developed to correct for these systematics.

5.2.2.2 Methods 2 and 2++

To correct for differences in passage fraction for the two faces of the detector, we require a sample of surface events which can be positively identified as having occurred on a particular

face of the detector. To do this, we use multiple scatter events which also scatter in a nearest-neighbor detector above or below the detector of interest. These are called *face-tagged nearest-neighbor multiples*; we denote the two faces of the detector as either *P-side* or *Q-side*, referring to whether that face of the detector is instrumented for phonon or ionization collection (see Sec. 2.1). In the interest of adequate statistics, ^{133}Ba calibration data below the ER band is used to define the SE population.

Using P and Q to denote the number of P-side and Q-side multiples before the phonon timing cut, and p and q to denote the number which pass the phonon timing cut, we define Method 2 using Eq. 28.

$$n_z = N_z \left[f_p \times \frac{p_z}{P_z} + f_q \times \frac{q_z}{Q_z} \right] \quad (28)$$

The scale factors f_p and f_q are the fraction of events, measured using WIMP-search sidebands, which occur on each face of the detector³⁶, ensuring that the final estimate is properly scaled.

To account for energy systematics in the leakage estimation, we define an extension of Method 2, called Method 2++, which bins the data into 3 energy bins: less than 20 keV, between 20 and 30 keV, and between 30 and 100 keV. We index the energy binning by e , and refer to the three bins as $e = l$ (low energy), $e = m$ (mid energy), and $e = h$ (high energy). Using this nomenclature, Method 2++ is defined using Eq. 29.

$$n_z = N_z \sum_{e=l,m,h} \left[f_{e,p} \times \frac{p_{z,e}}{P_{z,e}} + f_{e,q} \times \frac{q_{z,e}}{Q_{z,e}} \right] \quad (29)$$

5.2.2.3 Special Considerations for Endcap Detectors

Endcap detectors - those at either the top or bottom of a stack - necessarily have one side for which face-tagged multiples cannot be determined, because no neighboring detector

³⁶Note that $f_p + f_q = 1$.

exists on that side. To account for this in Methods 2 and 2++, all NRSS are labelled as face-tagged multiples for the “bare” face. Since it has been empirically observed that Q-side events tend to leak at a higher rate than P-side events, this method is conservative for bottom endcap detectors (which have a bare Q face), since it slightly overestimates the number of Q-side events, but biases the leakage estimates low for top endcap detectors (which have a bare P face) for the opposite reason.

5.3 Ge Timing Analysis

5.3.1 Classic Timing Analysis

The classic timing analysis³⁷ utilizes a simple linear combination of the phonon pulse rise-time and the phonon pulse delay relative to the ionization pulse. Mathematically, the timing quantity is defined as in Eq. 30.

$$t_{\text{classic}} = \text{pminrtCFc} + \text{pdelCFc} \quad (30)$$

For more information on the definition of these quantities and the reconstruction methods used, see Sec. 3.

The classic timing cut was tuned using the leakage-scanning method outlined in Sec. 5.1.2.1 using a target leakage of 0.5 expected events. No energy binning was used for the timing thresholds (i.e., each detector had only a single timing threshold). The two Ge endcap detectors, T3Z6 and T4Z6, were assigned very high thresholds due to their high expected leakage, and therefore contributed negligible exposure to the analysis.

In addition to the minimum threshold on t_{classic} set by the optimization method, neutron consistency cuts were also defined on a detector-by-detector basis to ensure that no extreme outlier events would inadvertently pass the cut. The first neutron consistency cut was a

³⁷This analysis was performed by Dr. Jianjie Zhang.[114]

Method	Leakage Estimate
M1 WS	1.140 ± 1.586
M2++ WS	0.180 ± 1.586
M2++ Ba	0.563 ± 0.151
Best Fit	0.565 ± 0.150

Table 7: Masked SE leakage estimates for the Ge classic phonon timing analysis. The poor statistics of the WIMP-search sidebands leads to extremely large statistical errors; this also causes the best-fit leakage estimate to be heavily influenced by the Method 2++ Ba estimate. Information taken from [115].

maximum threshold on t_{classic} , ensuring that no unphysically fast pulses would pass the timing cut. It was set at the 99.5% quantile of the neutron t_{classic} distribution (i.e., it removed 0.5% of the neutron sample). Finally, a consistency cut was set on the difference of the two timing quantities ($t_{\text{diff}} = \text{pminrtCFc} - \text{pdelCFc}$) at the 0.5% and 99.5% quantile levels, ensuring that no event could have drastically different timing behavior between the two quantities (e.g., very fast delay but very slow rise time).

Fig. 35 illustrates how the cut is defined in the plane of the two timing quantities. The efficiency of the cut was measured on ^{252}Cf neutrons as a function of energy. The efficiency varied by detector, generally lying in the range of 40% - 70% and slowly changing with energy (see Sec. 5.5).

After tuning the cut, initial leakage estimates using masked data were computed using three different methods (see Sec. 5.2): Method 1 on WIMP-search data (M1 WS); Method 2++ on WIMP-search data (M2++ WS); and Method 2++ on ^{133}Ba calibration data (M2++ Ba). These estimates, along with a best-fit value from the three estimates, is given in Tbl. 7.

5.3.2 5D χ^2 Timing Analysis

The 5D χ^2 timing analysis³⁸ utilizes five phonon timing quantities to compare each event to the measured distributions for signal (neutrons from ^{252}Cf calibration) and for background SEs (from ^{133}Ba calibration). The five timing quantities chosen for this analysis were:

- `pminrtc`
- `pdelc`
- `pminrtCFc`
- `pdelCFc`
- `pminrt4070c`

The distributions of these quantities can be represented by their means μ_j and covariance matrix elements C_{jk} ; using these, one can construct a goodness-of-fit χ^2 parameter using Eq. 31.

$$\chi^2 = \sum_{jk} (x_j - \mu_j) C_{jk}^{-1} (x_k - \mu_k) \quad (31)$$

The distributions of three different event classes were used for this analysis: neutrons, P-side SEs, and Q-side SEs. For each distribution, the means and covariance matrix elements were measured as functions of recoil energy and fit (Fig. 36); the fits are then evaluated at the recoil energy of the events whose χ^2 is being evaluated. The functional forms used to fit the means and covariances are given in Eq. 32.

$$\mu(E) = A_1 + A_2 E^2 + A_3 \sqrt{E} \quad (32a)$$

$$\sigma^2(E) = B_1 + \frac{B_2}{E^2} \quad (32b)$$

³⁸This analysis was performed by Dr. Joseph Kiveni.[116]

The χ_n^2 quantity gives information about how much the event in question “looks like” a neutron (i.e., signal) event; a similar quantity, describing how much the event “looks like” a background event, is constructed by taking the minimum of the P-side and Q-side χ^2 values: $\chi_b^2 = \min(\chi_p^2, \chi_q^2)$. The timing quantity of the 5D χ^2 timing analysis is defined as the difference between the neutron and background χ^2 values - higher values are more background-like, while lower values are more neutron-like.

$$t_{5D\chi^2} = \chi_b^2 - \chi_n^2 \quad (33)$$

For each detector, a neutron consistency cut is defined, requiring that events have a neutron χ^2 value no greater than the 90% quantile value for the ^{252}Cf neutron distribution. Without this requirement, it could be possible for events to pass the phonon timing cut because they were more neutron-like than background-like, without actually being neutron-like in an absolute sense.

The actual timing thresholds were set in three energy bins per detector using the slopes-matching optimization method (see Sec. 5.1.2.2). Additionally, the NR band width was simultaneously optimized with the timing thresholds. An optimized NR band of $\mu_{-1.9\sigma}^{+1.8\sigma}$ was found to maximize the sensitivity of the analysis, slightly tighter than the standard band of $\mu \pm 2\sigma$.

The efficiency of the 5D χ^2 timing cut was measured on ^{252}Cf calibration neutrons as a function of energy. The optimization procedure tends to set the cut more tightly in the low energy bins, since the yield discrimination is poorer and expected background is higher, resulting in a depressed efficiency at low recoil energy. At high energy, the efficiency is slowly varying with energy and high, with a maximum value ranging from 50% to 80% from detector to detector. Tbl. 8 gives the initial leakage estimates using masked data for the 5D χ^2 timing cut.




Figure 35: Classic analysis signal region for detector T1Z5. The neutron consistency cuts are shown as dashed blue lines, while the timing threshold itself is shown in red. Cyan points are ^{252}Cf calibration neutrons, and magenta points are SEs from ^{133}Ba calibration data. The region enclosed by these four lines is the “signal box”.




Figure 36: Example of mean and covariance matrix value fits. TOP: Measured means of the five timing quantities as a function of energy. P-side and Q-side SEs are shown in green and red, respectively; neutrons are shown in orange. BOTTOM: Measured covariances among the five timing quantities. Coloration is identical to that of the top panel. Figure from [87].

Method	Leakage Estimate
M1 WS	$2.26^{+1.31}_{-2.69}$
M2++ WS	$1.03^{+0.56}_{-2.45}$
M2++ Ba	0.55 ± 0.11
Best Fit	$0.568^{+0.106}_{-0.109}$

Table 8: Masked SE leakage estimates for the Ge 5D χ^2 timing analysis. The poor statistics of the WIMP-search sidebands leads to extremely large statistical errors; this also causes the best-fit leakage estimate to be heavily influenced by the Method 2++ Ba estimate. Information taken from [117].

5.3.3 Neural Network Timing Analysis

The neural network timing analysis is a machine-learning-based timing analysis, with a principal component analysis pre-processing step.³⁹ Four phonon timing quantities were chosen as the input space for this analysis.

- `pdelc`
- `pminrtc`
- `ptopwidth`
- `PPintPSD5070`

The first three are discussed in Sec. 3; the fourth quantity is the power spectral density of the largest phonon pulse, integrated from 50 to 70 kHz. These quantities were chosen by examining their 1-dimensional histograms for SEs and neutrons, and choosing those which had the most significant separation. Additionally, quantities were chosen with orthogonality in mind - i.e., quantities containing redundant and/or correlated information were not preferred (e.g. `pdelc` or `pdelCFc`, but not both).

³⁹To my knowledge, the first person to apply the idea of a neural network technique to CDMS data was Dr. Xinjie Qiu[118]; many of his original techniques were adapted in this work.

5.3.3.1 Mathematics of Neural Networks

The basic building block of a neural network is the *neuron*, defined by an activation function ϕ . ϕ is a function of a single variable, representing the response of a biological neuron to a stimulus. The activation function chosen for this analysis is a logistic sigmoid, which maps the domain $(-\infty, \infty)$ onto the range $(0, 1)$.

$$\phi(x) = \frac{1}{1 + e^{-x}} \quad (34)$$

A neural network is composed of many neurons, with outputs fed as inputs into other neurons via linear weights. The organization of outputs and inputs is called the *topology* of the network. Fig. 37 graphically illustrates the concept of topology for neural networks: inputs are on the left, outputs on the right; neurons are represented by blue circles; the black lines represent the topological connections among the neurons, with a linear weight attached to each of the connections. Although they are not shown explicitly in the illustration, each neuron also has a constant bias term feeding into it which is added to the weighted linear combinations coming from other neurons.

Clearly, the set of topologically-distinct neural networks is infinite in size. For pattern recognition applications, a specific form of neural network known as the multi-layer perceptron (MLP) is used. Fig. 38 shows the graphical representation of an MLP - the internal neurons are divided into layers (called *hidden layers*), each of which feeds forward into the next layer. The topology of an MLP is completely defined by the number of inputs, outputs, hidden layers, and neurons in each hidden layer.

Feed-forward networks, such as the MLP, are of special importance for pattern recognition applications because they can be trained using a method known as “backward propagation of errors”, or backpropagation for short. This is performed with a training sample of events with assigned target values for the network output. For the purposes of this timing

analysis, a single output MLP is used, with a target value of 1 assigned to ^{252}Cf neutrons and a target value of 0 assigned to SEs from ^{133}Ba calibration data.⁴⁰ As is common for machine learning applications, the available calibration data was randomly divided in half; one half was used as the training sample, while the other half was used as a test sample to ascertain the quality of the network training in an unbiased manner.

Backpropagation is a standard technique in neural network analysis;⁴¹ it is performed iteratively until a certain tolerance is met. In order to properly explain it, we must introduce some nomenclature, which is summarized in Tbl. 9. Let us denote the output of neuron i in layer k by y_i^k , and denote the weight connecting the output of this same neuron i , in layer k , to another neuron j in the next level $k+1$ by w_{ij}^k .⁴² Finally, let us denote the input into a neuron i in layer k by v_i^k . By the definitions given, one can see that v_i^k is given in terms of the outputs and weights according to Eq. 35.

$$v_i^k = \sum_j w_{ij}^{k-1} y_j^{k-1} \quad (35)$$

The final output of the MLP, denoted by y (no superscript, but possibly a subscript), can be written in terms of the MLP inputs, denoted by x_i , according to Eq. 36, where N is the number of hidden layers of the MLP. Note that the weights for layer N have only a single subscript; this is because they all feed into a single neuron to form the output layer (see Fig. 38).

$$y = \phi \left(\sum_i w_i^N \phi \left(w_{ij}^{N-1} \phi \left(\sum_k w_{jk}^{N-2} \dots \phi \left(\sum_m w_{lm}^0 \phi(x_l) \right) \right) \right) \right) \quad (36)$$

Now, given a training set of events, with a target values d_i assigned to each of the M

⁴⁰Note that 0 and 1 are extreme values for the neural network output, due to the definition of the activation function ϕ .

⁴¹For a pedagogical approach, see e.g. [119].

⁴²The bias term, which is not connected to any other neurons, is usually denoted by a subscript $i = 0$.




Figure 37: Graphical representation of a neural network. The four lines on the left represent the inputs to the neural network; the two lines on the right represent the outputs. Neurons are represented by blue circles, and the linear weights by the black lines.




Figure 38: Example of an MLP neural network. This example has 4 inputs, 1 output, and 3 hidden layers with 5 neurons each. An MLP is a feed-forward network, meaning that no layer of neurons receives inputs from layers further along the chain - i.e., each layer of neuron outputs is causal with respect to the previous layers.

training events, we define the error ε of the MLP using the sum-of-squares technique given in Eq. 37.

$$\varepsilon = \frac{1}{2} \sum_{i=1}^M (y_i - d_i)^2 \quad (37)$$

We now have the mathematical machinery needed to define the backpropagation MLP training algorithm.

1. Try an initial guess for the MLP weights w_{ij}^k .
2. Compute the output y_i for each of the training events. This is the *forward propagation* step.
3. Compute the error ε of the output.
4. Make corrections to the linear weights in order to decrease the error. Using the gradient descent method, one finds that the change in weights is given by Eq. 38.

$$\Delta w_{ij}^k = -\eta \frac{\partial \varepsilon}{\partial v_i^k} y_j^k \quad (38)$$

Here, η is an adjustable called the “learning rate”, which is tuned to ensure convergence. This is the *backward propagation* step - the final error is used to correct the intermediate weights.

5. Use the Δw_{ij}^k to update the weights; return to step 2 and iterate until a tolerance on ε is satisfied.

Using Eq. 37, it is straightforward to show that the partial derivative in Eq. 38 is given by Eq. 39, where $'$ denotes differentiation.

$$\frac{\partial \varepsilon}{\partial v_i^k} = \begin{cases} \phi'(v_i^k) \sum_{j=1}^M (d_j - y_j) & \text{output layer } (k = N) \\ \phi'(v_i^k) \sum_l \frac{\partial \varepsilon}{\partial v_l^k} w_{il}^k & \text{hidden layers } (k \neq N) \end{cases} \quad (39)$$

Quantity	Explanation	Relevant Eq.
ϕ	Neuron activation function	Eq. 34
x_i	i^{th} MLP input	
y	Final MLP output	Eq. 36
y_i^k	Output of i^{th} neuron in layer k	
w_{ij}^k	Weight connecting neuron i in layer k to neuron j in layer $k + 1$	
v_i^k	Input to neuron i in layer k	Eq. 35
d_i	Target value of the i^{th} training event	
ε	MLP error with respect to the training sample	Eq. 37
η	Learning rate	Eq. 38

Table 9: Summary of MLP nomenclature.

Once the training is complete, the weights are fixed and new data can be “run through” the network. A network response near 1 indicates that the event is neutron-like, and a response near 0 indicates that it is SE-like. The output of the trained networks is therefore a useful quantity with which to perform background discrimination; this forms the basis for the neural network timing cut.

The method outlined here utilizes gradient descent as the error-minimization algorithm, since it is one of the simpler methods to understand conceptually. In practice, a more robust and faster-converging method known as the scaled-conjugate gradient method[120] was used; it varies only in how the weight updates Δw_{ij}^k are computed in each iteration of the process.

5.3.3.2 Principal Component Analysis

For a single hidden layer MLP (such as is used in this analysis; see Sec. 5.3.3.3), the

computational complexity scales quickly with the number of inputs to the neural network.⁴³ For this reason, we desire to limit the number of input parameters while still maintaining high discrimination power.

A principal component analysis (PCA) is a statistical method which takes N possibly correlated input vectors and returns N new vectors which are linear combinations of the original vectors.[121] The transformation, which is orthogonal, is defined such that, in the transformed coordinate system, the projection of the data onto the first coordinate (or first principal component) gives the maximum variance; the projection onto the second principal component gives the second-greatest variance; and so on.

Observations in the N -dimensional vector space are represented by N -by-1 vectors, $\vec{x}_i$ (where i runs from 1 to M , the total number of observations). The N -by- N transformation matrix $\mathbf{P}$ takes vectors $\vec{x}_i$ from the original vector space and returns vectors $\vec{y}_i$ in the principal component space via Eq. 40.

$$\vec{y}_i = \mathbf{P}\vec{x}_i \quad (40)$$

Furthermore, we can represent the entire set of M observations with a matrix $\mathbf{X}$ whose columns are composed of the vectors $\vec{x}_i$. Analogously, we can construct the analogous vector for the transformed coordinate system, which we denote $\mathbf{Y} = \mathbf{P}\mathbf{X}$. In this representation, the k^{th} row of $\mathbf{P}$ (denoted $\mathbf{P}_{(k)}$) represents the linear weights which must be applied to the N columns of $\mathbf{X}$ to obtain the k^{th} principal component (denoted $\mathbf{Y}_{(k)}$).

By definition, the variance of the first principal component $\mathbf{Y}_{(1)}$ is maximized. If we impose that the mean of each of the N components of $\mathbf{X}$ is zero (which trivially ensures that the mean of the N components of $\mathbf{Y}$ are also zero), the variance of $\mathbf{Y}_{(1)}$ is given by

⁴³Although the number of weights feeding from the input layer to the hidden layer increases only linearly with the number of inputs, the computational time needed to compute the Δ_{ij}^k terms scales quadratically, since each term requires derivatives with respect to the other hidden neurons (see Eq. 39).

Eq. 41.

$$\text{var}(\mathbf{Y}_{(1)}) = \sum_i (\mathbf{Y}_{(1),i})^2 = \mathbf{Y}_{(1)} \mathbf{Y}_{(1)}^T = \frac{\mathbf{P}_{(1)} \mathbf{X} \mathbf{X}^T \mathbf{P}_{(1)}^T}{\mathbf{P}_{(1)} \mathbf{P}_{(1)}^T} \quad (41)$$

The denominator in the final expression is simply equal to unity (by the orthogonality requirement), but it is included since the corresponding quantity is a Rayleigh quotient. A standard result of linear algebra is that the quantity on the right-hand side of Eq. 41 is maximized when $\mathbf{P}_{(1)}$ is equal to the eigenvector of $\mathbf{X} \mathbf{X}^T$ which has the largest eigenvalue, and thus we identify this eigenvector as the transformation vector corresponding to the first principal component.

With the first principal component in hand, the second principal component can be found by constructing a new matrix $\mathbf{X}'$ which has the first principal component subtracted away.

$$\mathbf{X}' = \mathbf{X} - \mathbf{P}_{(1)}^T \mathbf{P}_{(1)} \mathbf{X} \quad (42)$$

The second principal component of $\mathbf{X}$ is simply the first principal component of $\mathbf{X}'$, computed in the same way as above. This process is applied repeatedly to obtain all of the principal components. It can be shown that each $\mathbf{P}_{(i)}$ is simply the eigenvector of $\mathbf{X} \mathbf{X}^T$ corresponding to the i^{th} -highest eigenvalue. Thus, the problem of finding the first principal component can be reduced to an eigenvector computation, for which there is no shortage of numerical techniques.

Our purpose for using the PCA transformation is to take the 4-dimensional space of timing parameters, and reduce it into a 2-dimensional space (composed of the first two principal components) which retains much of the discrimination power. Because the PCA technique chooses the components based on their variance, it is important that all input quantities (which are measured in units like μs and $\text{V}\sqrt{\text{Hz}}$) be normalized such that they are comparable. Therefore, in addition to the PCA requirement that the input vectors be mean-subtracted, the input distributions are also normalized to have unit variance.

There is an important caveat regarding the use of the PCA as a tool to choose variables with maximum discrimination. The two calibration distributions - ^{252}Cf and ^{133}Ba as proxies for signal and background, respectively - are not treated differently in any way by the PCA technique. It is an **assumption** that the linear combinations of the timing quantities which have the largest variance are also those which have maximum discrimination power. However, this assumption is well founded if the two distributions are separated enough that most of the variance of their union is due to their separation, and not to the intrinsic variances of the separate distributions. Therefore, the use of PCA to select discrimination parameters is valid insofar as the distributions of the initial timing quantities obey this condition.

5.3.3.3 Implementation

The PCA technique, including mean-subtraction, variance normalization, and transformation matrix computation, was performed separately for each detector and run. The first and second principal components were then used as inputs for the neural networks. In order to take into account recoil energy systematics within the timing parameters, separate neural networks were trained for high energy (> 30 keV) and low energy (≤ 30 keV) event populations within each detector-run combination. In total, this amounted to 102 separate neural networks which were trained for this analysis.⁴⁴

The neural networks used for this analysis were 2 input, 1 output MLPs, with 1 hidden layer containing 30 neurons. The implementation and training of these networks was performed using the NETLAB package[122] for MATLAB. The available sample of ^{252}Cf neutrons and ^{133}Ba SEs was randomly divided in half; one half was used as a training sample, while the other was withheld as a testing sample to gauge performance.

⁴⁴This is less than $14 \times 4 \times 2 = 112$ networks one would expect for 14 detectors and 4 runs, but certain detectors had some runs removed from WIMP-search consideration, which accounts for the discrepancy.

Due to the choice of activation function ϕ (Eq. 34), the final output of each neural network (denoted y) lies on the interval $(0, 1)$; for the purposes of setting a timing cut, we define the timing parameter for this analysis to be given by Eq. 43.

$$t_{\text{nnet}} = \phi^{-1}(y) = -\ln\left(\frac{1}{y} - 1\right) \quad (43)$$

t_{nnet} lies on the interval $(-\infty, \infty)$; this definition of timing parameter is more useful for examining the distributions of neural network output, since the distributions of y tend to be heavily bunched near the extremes of 0 (SE-like events) and 1 (neutron-like events) and the use of t_{nnet} has the effect of un-bunching them.

Although separate networks are trained for low- and high-energy events, a single timing threshold is set for each detector and run combination (i.e., the same timing threshold is applied to all events, regardless of energy). However, the network through which an event is propagated depends on its energy (e.g., a 25 keV event in a given detector would be sent through the corresponding low-energy network, while a 35 keV event would be sent through the high-energy network).

Neural network contours, as shown in Fig. 39, give a visual representation of how the network classifies events from the input parameter space. Setting a threshold in the histogram is equivalent to selecting a contour in the principal component plane; the contour corresponding to the chosen cut threshold is known as the “decision boundary”. Fig. 40 displays an example of the response of a neural network to neutrons and SEs from the test sample.

The leakage-scanning method (see Sec. 5.1.2.1) was used to optimize the neural network timing cut. The efficiency was measured as a function of energy using ^{252}Cf calibration neutrons. The efficiency generally decreased slowly with energy, providing the highest efficiency of the three timing analyses at energies near the recoil energy threshold; it varied by detector, generally falling between 50% - 80% near 10 keV.




Figure 39: Example of neural network discrimination for low-energy events from detector T3Z4, run 125. LEFT: test sample points in the principal component plane, with contours of constant t_{nnet} overlaid. Blue points are ^{252}Cf neutrons; red points are ^{133}Ba SEs. RIGHT: histogram of t_{nnet} . Good separation between neutrons (blue) and SEs (red) is observed. Setting a cut in the bottom plot is equivalent to choosing a contour on the top plot.




Figure 40: Neural network response visualized in the yield-recoil energy plane for detector T3Z4, run 125. Test sample ^{133}Ba SEs (LEFT) and ^{252}Cf neutrons (RIGHT) are shown with a colormap representing y , the neural network response. Colder (bluer) points are classified by the neural network as more SE-like, while warmer (redder) events are classified as more signal-like. The vertical dashed line demarcates low- and high-energy events, which are classified by separate networks.

Method	Leakage Estimate
M1 WS	2.113 ± 1.641
M2++ WS	0.363 ± 1.491
M2++ Ba	0.492 ± 0.304
Best Fit	0.539 ± 0.293

Table 10: Masked SE leakage estimates for the Ge neural network phonon timing analysis. The poor statistics of the WIMP-search sidebands leads to extremely large statistical errors; this also causes the best-fit leakage estimate to be heavily influenced by the Method 2++ Ba estimate.

Initial leakage estimates using masked data were computed using three different methods (see Sec. 5.2): Method 1 on WIMP-search data (M1 WS); Method 2++ on WIMP-search data (M2++ WS); and Method 2++ on ^{133}Ba calibration data (M2++ Ba). These estimates, along with a best-fit value from the three estimates, is given in Tbl. 10.

5.4 Si Timing Analysis

The Si phonon timing analysis was performed using only a single cut - a 2-dimensional version of the χ^2 cut⁴⁵ described in Sec. 5.3. Since the Si timing analysis was performed after the Ge timing analysis was completed,⁴⁶ a 2D χ^2 approach was chosen because of the excellent results of the 5D χ^2 Ge timing cut over a broad range of WIMP masses (see Sec. 8).

The parameters used for the Si 2D χ^2 timing analysis were `pminrtc` and `pdelc` - the same as were used for the classic timing analysis. The χ^2 quantities for this analysis are defined in the same way as those for the Ge 5D χ^2 analysis. The slope-matching method (see Sec. 5.1.2.2) was used to optimize this cut, with one important difference: the cut threshold in the lowest energy bin (< 20 keV) was optimized for a $15 \text{ GeV}/c^2$ WIMP mass; the middle energy bin ($20 - 30$ keV) for $30 \text{ GeV}/c^2$; and the highest energy bin (> 30 keV)

⁴⁵The Si timing analysis was performed by Dr. Kevin McCarthy.[112]

⁴⁶The Si data were not used for WIMP-search analysis until the c58R reprocessing.

for 60 GeV/ c^2 .

As with the Ge version of the cut, the discrimination parameter is the difference between the background and neutron χ^2 quantities.

$$t_{2D\chi^2} = \chi_b^2 - \chi_n^2 \quad (44)$$

The cut is defined in two parts: a consistency cut on χ_n^2 , and a rejection cut on $t_{2D\chi^2}$. Since leakage fractions tend to increase at lower recoil energies, the cut was optimized to be tighter in the low energy bins, similar to what was seen in the Ge 5D χ^2 analysis. For two detectors, the timing threshold was set at infinity (i.e. the bin was removed from WIMP search): T3Z3, below 30 keV; and T4Z1, below 20 keV.

The efficiency of the Si 2D χ^2 timing cut was measured on ^{252}Cf neutrons as a function of energy. As expected, the efficiency decreases at low recoil energies, and levels off at a roughly constant value above ~ 30 keV. The efficiency varies by detector, generally falling between 60% and 85% at high energy.

After tuning the cut, preliminary leakage estimates using masked data were computed using three different methods (see Sec. 5.2): Method 1 on WIMP-search data (M1 WS); Method 2++ on WIMP-search data (M2++ WS); and Method 2++ on ^{133}Ba calibration data (M2++ Ba). These estimates, along with a best-fit value from the three estimates, is given in Tbl. 11.

5.5 Analysis Exposure and Efficiency

The raw exposure of these analyses is defined to be the product of the detector masses and live-time, summed over all detectors, where the live-time is computed after all time-period based event pre-selection cuts have been applied (see Sec. 4.3). The raw exposure of the c58R Ge analysis was 612.2 kg-day; for the c58R Si analysis, it was 140.2 kg-day.[112]

Method	Leakage Estimate
M1 WS	0.824 ± 0.592
M2++ WS	1.230 ± 0.849
M2++ Ba	0.301 ± 0.178
Best Fit	0.412 ± 0.167

Table 11: Preliminary SE leakage estimates for the Si 2D χ^2 phonon timing analysis. The poor statistics of the WIMP-search sidebands leads to extremely large statistical errors; this also causes the best-fit leakage estimate to be heavily influenced by the Method 2++ Ba estimate. Information taken from [123].

5.5.1 Ge Exposure and Efficiency

Figs. 41, 42, and 43 display the total efficiencies of the Ge classic, 5D χ^2 , neural network timing analyses, respectively. Fig. 44 displays them on a single plot for easy comparison. The curves indicate the effects on the raw exposure (black horizontal line) by the successive application of the trigger efficiency, event pre-selection cuts, fiducial volume selection, NR band cut, and timing cut. The vertical dashed line indicates the phonon threshold for the high-threshold analysis - 10 keV for all Ge detectors. The left axis displays the efficiency, while the right axis is normalized to the full exposure of the analysis.

Tbl. 12 displays the final SAE values (see Eq. 22) for the 3 Ge phonon timing analyses at the nominal WIMP mass of 60 GeV/ c^2 . The high SAE of the 5D χ^2 analysis is indicative of its discrimination power at high energy, although the higher efficiencies of the classic and neural network timing cuts at low energy indicate their superior power at low WIMP masses. Because of the high SAE of the 5D χ^2 analysis, it was chosen prior to unmasking to be the primary analysis for the c58R Ge exposure.

5.5.2 Si Exposure and Efficiency

Fig. 45 displays the total efficiency of the Si 2D χ^2 timing analysis. The stairsteps in the total efficiency curve represent the different recoil energy thresholds for different detectors

Analysis	SAE (kg-day)	WIMP mass (GeV/c^2)	Relevant section
Classic	220	60	Sec. 5.3.1
5D χ^2	250	60	Sec. 5.3.2
Neural network	216	60	Sec. 5.3.3

Table 12: SAE of the 3 Ge phonon timing analyses. The superior performance of the 5D χ^2 analysis resulted in its selection as the primary analysis for the c58R Ge exposure.




Figure 41: Efficiency of the Ge classic phonon timing cut. The raw exposure of 612.2 kg-day is represented by the horizontal black line. The other curves indicate the effect of the successive applications of the trigger efficiency (cyan), event pre-selection (green), fiducial volume cut (red), yield cuts (blue), and classic timing cut (magenta).




Figure 42: Efficiency of the Ge 5D χ^2 phonon timing cut. The raw exposure of 612.2 kg-day is represented by the horizontal black line. The other curves indicate the effect of the successive applications of the trigger efficiency (cyan), event pre-selection (green), fiducial volume cut (red), yield cuts (blue), and 5D χ^2 timing cut (magenta).




Figure 43: Efficiency of the Ge neural network phonon timing cut. The raw exposure of 612.2 kg-day is represented by the horizontal black line. The other curves indicate the effect of the successive applications of the trigger efficiency (cyan), event pre-selection (green), fiducial volume cut (red), yield cuts (blue), and neural network timing cut (magenta).

and runs. At a WIMP mass of $15 \text{ GeV}/c^2$, the SAE is 22.0 kg-day .⁴⁷

⁴⁷Note that SAE values are not directly comparable for different WIMP masses or target nuclei.




Figure 44: Efficiencies of the Ge phonon timing cuts. The 5D χ^2 cut (blue dashed) has very good efficiency at high energy, while the classic (magenta dash-dot) and neural network (red solid) cuts perform better at low energies.




Figure 45: Efficiency of the Si 2D χ^2 phonon timing cut. Figure from [112].

6 Extension of Timing Analyses to Lower Thresholds

6.1 Mid-Threshold Analysis

Historically, the CDMS collaboration has performed two types of analyses: a high-threshold-style (HT) analysis, which employs (as the name implies) high energy thresholds and a phonon timing cut in order to remove SE backgrounds; and a low-threshold-style (LT) analysis, which pushes energy thresholds down to noise-limited levels and makes no attempt to remove background through phonon timing analysis (and thus suffers no corresponding loss in efficiency). The former type of analysis nominally aims for zero background events, while the latter accepts backgrounds while attempting to quantify and understand them.

HT analyses are generally quite successful at setting strong limits at high WIMP masses, since at high energies the yield quantity and the phonon timing quantities are well-resolved and backgrounds can be effectively removed. On the other hand, LT analyses often perform better at low WIMP masses since their low recoil energy thresholds make them sensitive to the spectra of low-mass WIMPs. However, since LT analyses have significant backgrounds, their limits are weakened either by the statistical (and systematic) penalties involved with background subtraction, or by conservatively assuming all candidates are WIMPs.

It was suggested by Dr. Richard Schnee that a blended approach, combining elements of both HT and LT analyses, may provide improved sensitivity in the ~ 5 -10 keV energy range.⁴⁸ In this so-called “mid-threshold” (MT) analysis, energy thresholds are moderately relaxed (but not as far as in a LT analysis), while phonon timing cuts still provide some background rejection. At intermediate energies, the rejection power of the timing cuts outweighs their corresponding efficiency losses.

The HT timing cuts were not re-tuned for the MT analyses. Instead, the cut positions from the lowest energy bin were used for the data below the original recoil energy thresholds.

⁴⁸The original motivation of this extension was as a cross-check of the potential WIMP signal observed in the c58 Si analysis.[124]

However, after unmasking the MT analyses, a MT reoptimization of the 5D χ^2 Ge analysis was performed to see what sensitivity improvements were possible. This analysis is described in Sec. 6.1.2.

6.1.1 Determination of Thresholds

At low recoil energies, the NR band is pinched off by the flaring of the ionization threshold and below-ER cuts (see Fig. 31); we define the “pinch-off energy” as the recoil energy at which the below-ER and ionization threshold cuts intersect in the yield plane; even in the absence of an explicit recoil energy threshold, the signal region is completely excluded (“pinched off”) below the pinch-off energy by the action of those two cuts. Therefore, in order to extend the signal region to lower recoil energies, either or both of these cuts must be relaxed. The below-ER and ionization thresholds are set at 3σ and 4.5σ , respectively, for the HT analysis. Due to the high rate of bulk γ s incident on the detectors compared to those in the zero-yield band, loosening the ionization threshold cut is preferred.⁴⁹

The ionization threshold was varied from 4.5σ down to 3σ to determine its effect on thresholds and expected leakage. Leakage estimates using ^{133}Ba data indicated a small increase in expected leakage, although these changes were completely within statistical errors. Furthermore, decreasing the ionization threshold from 4.5σ to 3σ lowered the pinch-off energy by between 0.5 and 1 keV for each detector (down to as low as 3 keV). Ionization thresholds of 3σ , 3.5σ , 4σ , and 4.5σ were considered, and background simulations were used to determine the optimal value (see Sec. 6.2); in the end, the lowest threshold (3σ) maximized the projected sensitivity, and was chosen for all analyses.

In order to obtain reasonable estimates of the SE background, we require a population of calibration SEs which populate the same parameter space as our signal region. However, the

⁴⁹After unmasking, it was discovered that the Ge neural network timing cut had a high passage fraction of zero-yield events (see Sec. 7.1.2); in retrospect, loosening the ionization threshold cut for that analysis contributed more leakage than loosening the below-ER cut would have.




Figure 46: Changes in expected leakage (top panels) and WIMP signal (bottom panels) for the Ge (left panels) and Si (right panels) MT analyses as the ionization threshold is varied. Lowering the threshold to 3σ has an insignificant effect on expected leakage, but expected WIMP signal benefits greatly from the improved efficiency at low recoil energies. For expected WIMP signal calculations, $m_\chi = 8.6 \text{ GeV}/c^2$ and $\sigma_{\text{SI}} = 1.9 \times 10^{-41} \text{ cm}^2$.

selection criterion for calibration SEs from ^{133}Ba data requires that they be at least 5σ below the ER band (to prevent bulk γ contamination in the ^{133}Ba calibration SE population), rather than the 3σ cut used for WIMP search event selection. The pinch-off energy of the 5σ below-ER cut is higher than that of the 3σ cut, usually by ~ 2 keV; therefore, if no explicit recoil energy threshold is imposed, there is a piece of the signal region for which we cannot estimate the SE background with ^{133}Ba data.

For this reason, we define a recoil energy threshold for the MT analysis at the point where the 5σ below-ER cut intersects with the ionization energy threshold cut. Results will be presented with the recoil energy threshold on and off, but to be conservative, the analysis was performed as though the thresholds would always be enforced. Tbl. 13 displays the MT analysis thresholds computed in this way for each detector and run.

6.1.2 Re-optimization of Ge 5D χ^2 Analysis

In light of the excellent high WIMP-mass sensitivity of the 5D χ^2 HT analysis (see Sec. 5.3.2), an effort was made to re-optimize⁵⁰ that timing cut for the MT analysis using a lower target WIMP mass: $8 \text{ GeV}/c^2$, rather than the $60 \text{ GeV}/c^2$ which was used in the HT optimization. For simplicity, a single energy bin from 5 to 15 keV was used, since recoil energies much higher than that do not contribute significantly to low mass sensitivity (see Fig. 47).

The mean and covariance matrix fits from the HT analysis were used for the MT re-optimization, as well as the same neutron consistency cut (see Sec. 5.3.2). However, some problems were observed at low recoil energies, due to the nature of the covariance matrix element fits.

An empirically-measured covariance matrix obeys some basic relations, including having a non-negative determinant and the Cauchy-Schwartz inequality (Eq. 45). However, the

⁵⁰This analysis was performed by Dr. Joseph M. Kiveni.

Detector	Material	R125	R126	R127	R128
T1Z2	Ge	5.51	5.78	6.25	6.15
T1Z4	Si	8.43	-	-	-
T1Z5	Ge	4.50	4.14	-	-
T2Z1	Si	4.51	4.61	5.40	5.18
T2Z2	Si	5.09	5.61	5.40	5.94
T2Z3	Ge	8.50	8.02	8.61	8.56
T2Z4	Si	4.71	4.55	4.40	5.38
T2Z5	Ge	6.64	6.80	6.96	6.71
T3Z2	Ge	7.30	7.22	7.84	8.16
T3Z3	Si	6.24	6.12	-	6.43
T3Z4	Ge	4.98	5.08	5.31	5.64
T3Z5	Ge	10.02	10.03	11.13	10.72
T3Z6	Ge	5.14	5.43	5.61	5.45
T4Z1	Si	7.75	7.87	8.56	8.41
T4Z2	Ge	6.72	7.16	8.19	8.38
T4Z3	Si	13.42	6.20	6.08	5.88
T4Z4	Ge	6.38	6.08	6.73	7.12
T4Z5	Ge	8.05	8.09	8.42	7.75
T4Z6	Ge	7.54	7.50	7.91	7.71
T5Z3	Si	5.07	5.43	4.75	5.06
T5Z4	Ge	6.97	7.00	7.83	7.50
T5Z5	Ge	8.03	8.48	8.99	8.85

Table 13: Recoil energy thresholds for the MT analysis (in keV).

functional forms used to compute the χ^2 quantities are simply fit to empirically measured matrix element values, and the functional forms themselves are not guaranteed to satisfy the definition of a true covariance matrix at all energies. Since the matrix element fits begin to diverge at low energies, and must be extrapolated from the more reliable data points at high energies (see Fig. 36), this effect was not seen above the HT recoil energy threshold of 10 keV.

$$|\sigma(x, y)| \leq \sigma(x)\sigma(y) \quad (45)$$

We observe that the matrix element fits are good at high energies - i.e., the elements obey the properties of a covariance matrix. However, if the fits become poor at low energies - in particular, if $\det(C) < 0$ - then problems arise, because $\det(C) = 0$ at some energy. Because the computation of the χ^2 quantities (Eq. 31) involves inverting C , and $C^{-1} \propto 1/\det(C)$, poorly fit covariance matrix elements at low energies can lead to singularities in the χ^2 quantities themselves. This problem is illustrated in Fig. 48.

Rather than take on the problem of re-fitting the covariance matrix elements (perhaps using the method of Lagrange multipliers to enforce conditions on the covariance matrix), a simpler kludge was employed to “fix” the matrix elements. Recall that the functional form used to fit the covariance matrix elements is given by Eq. 32; the term responsible for its divergence at low energies is the second term, B_2/E^2 . We attempt to correct the fits by introducing a kludge factor, $\kappa \geq 1$, by which each off-diagonal B_2 constant is rescaled.

$$B_2 \rightarrow \frac{B_2}{\kappa} \quad \text{for all } \sigma_{ij} \ (i \neq j) \quad (46)$$

Note that the same factor κ is used to rescale all of the off-diagonal fit parameters for a given covariance matrix.

The effect of the kludge factor on the matrix determinants is illustrated in Fig. 49. The smallest value of the kludge factor which pushes the singularity below threshold is chosen to




Figure 47: Comparison of WIMP spectra recoiling on Ge nuclei for different WIMP masses, using the Lewin & Smith formulation[110] and standard halo assumptions.[125, 126] The HT analyses were optimized to a $60 \text{ GeV}/c^2$ WIMP mass (black); a re-optimized version of the 5D χ^2 Ge analysis was tuned for $8 \text{ GeV}/c^2$ (blue). Intermediate WIMP masses of 10 (red), 15 (green), and 30 (magenta) GeV/c^2 are also shown.




Figure 48: Problems with the covariance matrix element fits for detector T3Z6. LEFT: Distribution of χ_n^2 versus energy, for ^{133}Ba SEs (blue) and ^{252}Cf neutrons (green). RIGHT: Determinant of the neutron covariance matrix, using fits to the matrix elements. The point at which the determinant passes through zero corresponds to the sharp discontinuity in the distribution.

“correct” the covariance matrix. This correction was performed by eye; no correction factor exceeded 1% (i.e. $\kappa < 1.01$), and 6 of the detectors required no correction at all ($\kappa = 1$).⁵¹

With the kludge factor correction in place, the timing cut was optimized using the slope-matching method for the lower 8 GeV/ c^2 WIMP mass. As with the HT 5D χ^2 analysis, the width of the NR band was simultaneously optimized; the NR band was set to $\mu_{-1.8\sigma}^{+1.5\sigma}$, tighter than the standard $\mu \pm 2\sigma$ NR band used for the other analyses.

6.2 Background and Sensitivity Estimation

6.2.1 Efficiencies

Since the efficiencies for the Ge HT analyses were only important down to the 10 keV recoil energy threshold, the efficiencies were re-measured and re-fit using ^{252}Cf neutrons down to lower energies. Figs. 50 and 51 show the efficiencies of the timing cuts for the Ge and Si MT analyses, respectively, with the ‘no timing cut’ efficiency shown for comparison.

6.2.2 Background Estimates and Modeling

Backgrounds were estimated using Method 2, as outlined in Sec. 5.2; however, here we estimate the leakages only below 20 keV⁵², so no energy binning (i.e. M2++) is used. Tbl. 14 displays the leakage passage fractions of the ^{133}Ba multiple scatters, and the leakage estimates which follow therefrom.

Using these background estimates, simple background models were constructed by fitting the SE spectrum and fixing the normalizations using the leakage estimates obtained in Tbl. 14. For robustness, 6 different background models were utilized and compared.

- **Model 0:** flat SE spectrum (no fitting required).

⁵¹The Si 2D χ^2 analysis, which was also extended to lower recoil energy thresholds, was found not to suffer from such singularities at the MT energy thresholds.

⁵²The MT analyses were performed after unmasking the HT analyses; thus it was already known that no high energy (>20 keV) leakage events would be observed.




Figure 49: Kludge correction factor for the T3Z6 neutron covariance matrix. The determinant of the neutron covariance matrix is plotted versus energy for a variety of kludge factors (labelled “c” in the legend). A kludge factor of 1.0017 was chosen for this detector (red dashed line), since it pushes the singularity below the recoil energy threshold of 5 keV. Figure adapted from [127].




Figure 50: Efficiencies for the Ge MT timing analyses: classic (red), neural network (green), re-optimized 5D χ^2 (blue), and no timing cut (black). Dashed lines indicate 1σ error bar fits. LEFT: Example of efficiencies for detector T3Z4. The four stair-steps correspond to the different recoil energy thresholds for the four runs of c58. RIGHT: Final efficiencies of the timing analyses, using the exposure-weighted average over all detectors.




Figure 51: Efficiencies for the Si MT timing analyses: 2D χ^2 (blue) and no timing cut (red). Dashed lines indicate 1σ error bar fits. LEFT: Example of efficiencies for detector T2Z4. The four stair-steps correspond to the different recoil energy thresholds for the four runs of c58. RIGHT: Final efficiencies of the timing analyses, using the exposure-weighted average over all detectors.

Analysis	N_2	P-side passage fraction	Q-side passage fraction	M2 leakage estimate
Ge Classic	128.98 ± 21.38	$51/3248$	$31/2113$	$1.52^{+0.38}_{-0.37}$
Ge 5D χ^2	82.05 ± 14.47	$110/1775$	$53/1094$	$2.86^{+0.62}_{-0.61}$
Ge Neural Network	128.98 ± 21.38	$52/3248$	$24/2113$	$1.06^{+0.28}_{-0.28}$
Ge no timing	128.98 ± 21.38	$3248/3248$	$2113/2113$	$128.98^{+21.38}_{-21.38}$
Si 2D χ^2	112.12 ± 16.66	$7/529$	$7/983$	$1.40^{+0.63}_{-0.58}$
Si no timing	112.12 ± 16.66	$529/529$	$983/983$	$112.12^{+16.66}_{-16.66}$

Table 14: Leakage estimates and passage fractions for the MT timing analyses. The computation of N_2 is described in Sec. 5.2.1; the differences between the values for the Ge 5D χ^2 analysis and the other Ge analyses is due to the fact that its upper energy threshold was set at 15 keV, rather than 20 keV, and that its NR band was optimized more tightly than for the other Ge analyses. The M2 leakage estimate calculation is described in Eq. 28.

- **Model 1:** linear SE spectrum: $f(x) = c_1 \cdot x + c_2$, $c_1 < 0$.
- **Model 1+:** same as Model 1 with 1σ overweighting (see below).
- **Model 2:** exponential spectrum: $f(x) = c_1 \cdot e^{c_2 x}$, $c_2 < 0$.
- **Model 2+:** same as Model 2 with 1σ overweighting (see below).
- **Model 3:** same spectral shape as an 8 GeV/ c^2 WIMP (no fitting required).

Fits were performed with ^{133}Ba SEs which pass the timing cut in question, co-added over all Ge or Si detectors for statistics, from 4 to 20 keV with 1 keV bins. In Models 1+ and 2+, **overweighting** is a method for “steepening” the fitted spectra by adding 1σ upward Poisson fluctuations to three lowest non-zero bins, and subtracting 1σ downward Poisson fluctuations to all other higher energy bins; this causes the background fits to be more heavily weighted towards low energies, which is conservative when estimating sensitivity to low-mass WIMPs.

Models 1 and 2, along with their overweighted counterparts 1+ and 2+, are realistic models to the observed SE spectra and fit the data well; Models 0 and 3 are included as extreme “bookends” to estimate systematics, since they are conservative for high- and low-mass WIMPs, respectively. Since the recoil energy thresholds are known exactly, their effects can be introduced by means of a sum of Heaviside step functions at the appropriate energies, appropriately weighted by the respective exposures of the detector-run combinations; this function is given in Eq. 47. Here, z and r index the detectors and runs, MT is the exposure, and $E_{z,r}^T$ is the recoil energy threshold for a given detector and run.

$$T(E) = \frac{\sum_{z,r} (MT)_{z,r} H(E - E_{z,r}^T)}{\sum_{z,r} (MT)_{z,r}} \quad (47)$$

The product of this threshold-step function and the fitted SE spectra form the SE energy distribution of the background model. The results of this procedure are shown in Fig. 52 for

the Ge classic timing analysis.

6.2.3 Sensitivity Estimation

Given an estimate of the total magnitude of the SE background (Tbl. 14) and a model for the spectral shape of this background (Fig. 52), one can simulate a large number of toy experiments (assuming no signal) to estimate the sensitivity of the analyses. The simulation procedure is given as follows.

1. Using the posterior distribution of the expected number of leakage events f_λ , randomly sample a value λ_i (where i denotes the i^{th} simulation).
2. Using λ_i as the parameter for a Poisson distribution, randomly sample an “observed” (integer) number of leakage events N_i .
3. Randomly sample N_i events from the SE spectrum. This is the background for the i^{th} experiment.

Armed with a large number of simulated experiments and measured efficiencies, we can utilize the optimal interval method to compute 90% confidence limits on the WIMP-nucleon coupling as a function of WIMP mass (see Sec. 8.1). These can then be used to form confidence intervals on the expected exclusion limits of the analysis.

10-20 simulated experiments were thrown for each combination of timing analysis, background model, and ionization energy threshold. Fig. 53 shows the different sensitivities obtained when the background models are varied for the Ge classic analysis; they are strongly overlapping over a large WIMP mass range, indicating the robustness of the background modeling approaches. This robustness was seen in all of the timing cuts; however, with no timing cut applied, very large differences were observed between Model 3 (the most conservative) and the others, as seen in Fig. 54.

Fig. 55 illustrates how the sensitivity changes when the ionization threshold is varied from 4.5σ to 3σ - although the improvement in sensitivity is within error bars at high WIMP masses, the sensitivity improvement at low WIMP masses was significant enough for a 3σ ionization threshold to be chosen.⁵³

Finally, the various timing analyses were compared so that a primary analysis could be chosen before the WIMP-search data were unmasked. Figs. 56 and 57 show these comparisons for the Ge and Si timing analyses, respectively. For the Ge analyses, the neural network and classic timing cuts provide the strongest sensitivities; over most of the WIMP mass range, their sensitivities are within errors of each other. Although the no-timing analysis is strongest at low WIMP masses, its large modeling systematics caused it to be discounted. In the end, **the classic timing cut was chosen as the primary Ge MT analysis** because of its strong sensitivity. It was chosen over the neural network analysis because the construction of its timing parameter is considerably more straightforward, and this “legibility” is valuable when interpreting candidate events and final results.

For the Si analysis, the 2D χ^2 analysis outperforms the no-timing analysis except at low WIMP masses, but the modeling systematics seen in the Ge no-timing analysis are also present in Si; therefore, **the 2D χ^2 timing cut was chosen as the primary Si MT analysis.**

⁵³A similar improvement was seen in all Ge and Si timing analyses, so a 3σ ionization threshold was employed for all MT analyses.




Figure 52: SE spectrum and fits for the Ge classic timing analysis. The measured ^{133}Ba spectrum of SEs which pass the Ge classic timing cut (thick black) is well approximated by Models 1 (red), 1+ (magenta), 2 (blue), and 2+ (cyan); Models 0 (black dashed) and 3 (green) are more extreme models, and are used to gauge the magnitude of systematics in the background modeling. All curves are normalized so that they integrate to unity.




Figure 53: Sensitivity estimates for the Ge classic timing cut using different background models. Shaded regions indicate the 68% confidence intervals, and solid lines show median sensitivity for Models 0 (red), 1 (green), 1+ (magenta), 2 (cyan), 2+ (yellow), and 3 (black). An ionization threshold of 3σ is used.




Figure 54: Sensitivity estimates for the Ge no timing analysis using different background models. The sensitivity estimate for the most conservative model, Model 3 (black), strongly conflicts with the estimates of the other models. An ionization threshold of 3σ is used.




Figure 55: Sensitivity estimates for the Ge classic timing cut using different ionization thresholds. Lowering the ionization threshold from 4.5σ (cyan) to 4σ (magenta), 3.5σ (green), and 3σ (red) progressively improves the sensitivity at low WIMP masses. Model 2+ is used for the background.




Figure 56: Sensitivity estimates for the Ge timing analyses, using Model 2+. The classic (blue) and neural network (green) have good sensitivities over the full WIMP mass range; the no-timing analysis (black) does best at low WIMP masses, but its sensitivity estimates suffer from large systematics, and is therefore discounted. The re-optimized 5D χ^2 cut (red) does poorly, largely on account of its poor efficiency at low energies. An ionization threshold of 3σ is used.




Figure 57: Sensitivity estimates for the Si timing analyses, using Model 2+. The 2D χ^2 analysis (blue) outperforms the no-timing analysis (red) at all except low WIMP masses, but the sensitivity estimates for the no-timing analysis suffer from large systematics, and is therefore discounted. An ionization threshold of 3σ is used.

Candidate	Detector	Date (Run)	E_R [keV]	Passing Analyses
G1	T1Z5 (Ge)	27 Oct 2007 (R125)	12.3	Classic, neur. net.
G2	T2Z3 (Ge)	30 May 2008 (R127)	10.8	Classic, re-opt. 5D χ^2
S1	T4Z3 (Si)	1 Jul 2008 (R127)	9.5	2D χ^2
S2	T4Z3 (Si)	6 Sep 2008 (R128)	12.3	2D χ^2
S3	T5Z3 (Si)	14 Mar 2008 (R126)	8.2	2D χ^2

Table 15: Candidate events for the high-threshold analyses. Event G1 was one of the two candidate events from the original analysis of the Ge data[86]; the other candidate from that analysis was easily rejected by all of the Ge analyses in this thesis. G2 is only a candidate for the re-optimized low-mass 5D χ^2 mid-threshold analysis; the high-threshold version of that analysis produced no candidate events.

7 Candidates and Backgrounds

7.1 WIMP Candidate Events

7.1.1 High-Threshold Analysis Candidates

The c58 Ge high-threshold timing analyses were unmasked simultaneously in July 2012, and the Si analysis was unmasked on Christmas Day of the same year. The Ge analyses resulted in 0, 2, and 1 candidate events for the 5D χ^2 , classic, and neural network analyses, respectively; the Si analysis had 3 candidate events. Information about these events is shown in Tbl. 15.

The re-analysis of these data was originally motivated by the poor ionization pulse reconstruction of one of the two candidate events in the original analysis (see Sec. 3.7). That poorly-reconstructed event was rejected by a wide margin by all of the Ge analyses presented in this thesis; the second of the original candidate events (which was well-reconstructed) is present in the re-analysis as candidate G1, which passed the classic and neural network Ge analyses.

The candidate events were inspected to ensure that they were well-behaved and passed all event selection cuts in a reasonable fashion. Detector behavior, raw pulse quality, pulse

reconstruction, and event selection were all determined to be normal[128, 129]; therefore all candidate events are considered to be good.

7.1.2 Mid-Threshold Analysis Candidates

There is some subtlety in the identification of the candidate events for the mid-threshold analyses. For the three analyses which were simply extensions of their high-threshold counterparts (Ge neural network, Ge classic, and Si 2D χ^2), the set of mid-threshold candidate events necessarily contains the high-threshold candidates as a subset. The re-optimized Ge 5D χ^2 mid-threshold analysis is not a simple extension of a high-threshold cut; therefore its set of candidates is independent of the set of candidates for the high-threshold analysis.

Furthermore, recoil energy thresholds were enforced for the mid-threshold analyses to exclude parts of the signal region for which the SE leakage could not be reliably estimated (see Sec. 6.1.1). For completeness, candidate events below these thresholds but which still pass all other cuts are displayed as well, although it is understood that no good estimate of the background for these events is available, and the primary results (chosen before unmasking) will be those for which the recoil energy thresholds are enforced. The candidate events are summarized in Tbls. 16; candidate event counts for all analyses are shown in Tbl. 17.

The neural network analysis stands out among the MT analyses for its significantly higher counts - 6 additional events for the recoil threshold “on”, and 9 more for the recoil threshold “off”. Investigating these events leads to the conclusion that the neural network timing cut passes low-ionization events at a much higher rate than any of the other analyses; as such events begin to pollute the signal region at lower recoil energies, the corresponding leakage increases substantially. An example of this phenomenon is illustrated in Fig. 58. Full plots of the unmasked WIMP-search data for all detectors and analyses can be found in Appendix A. This is an example of the unexpected behavior one may receive from “black-

Candidate	Detector	Date (Run)	E_R [keV]	Passing Analyses	Threshold
G3	T1Z5 (Ge)	14 Feb 2008 (R126)	5.7	Classic, neur. net.	+
G4	T4Z4 (Ge)	10 Apr 2008 (R126)	7.6	Classic, neur. net.	+
G5	T2Z5 (Ge)	6 Nov 2007 (R125)	3.5	Classic	-
G6	T4Z5 (Ge)	24 Aug 2008 (R128)	7.3	Classic, neur. net.	-
G7	T3Z2 (Ge)	18 Aug 2008 (R128)	9.7	Neur. net.	+
G8	T3Z2 (Ge)	23 Aug 2008 (R128)	8.8	Neur. net.	+
G9	T3Z3 (Ge)	31 Aug 2008 (R128)	9.5	Neur. net.	+
G10	T3Z4 (Ge)	21 Feb 2008 (R126)	7.5	Neur. net.	+
G11	T1Z2 (Ge)	16 Aug 2008 (R128)	5.8	Neur. net.	-
G12	T1Z5 (Ge)	3 Aug 2007 (R125)	3.9	Neur. net.	-
G13	T1Z5 (Ge)	6 Nov 2007 (R125)	3.5	Neur. net.	-
G14	T1Z5 (Ge)	16 Dec 2007 (R125)	3.6	Neur. net.	-
G15	T2Z5 (Ge)	29 Jul 2008 (R127)	5.1	Neur. net.	-
G16	T3Z4 (Ge)	26 Jun 2008 (R127)	3.9	Neur. net.	-
G17	T4Z2 (Ge)	5 Jun 2008 (R127)	6.9	Neur. net.	-
G18	T4Z4 (Ge)	14 Feb 2008 (R126)	5.4	Neur. net.	-
G19	T3Z6 (Ge)	1 Jun 2008 (R127)	12.7	5D χ^2	+
S4	T2Z2 (Si)	18 Apr 2008 (R126)	4.9	2D χ^2	-

Table 16: Candidate events for the mid-threshold analyses. Here, “5D χ^2 ” refers exclusively to the low-mass re-optimized version used for the MT analysis. The final column indicates whether the candidate was above (+) or below (-) the recoil energy threshold. Note that for all analyses except the re-optimized Ge 5D χ^2 , the HT candidates are a subset of the MT candidates; although not shown here, candidate G2 from Tbl. 15 is a candidate for the re-optimized Ge 5D χ^2 MT analysis.

Analysis	HT Candidates	MT Candidates (threshold on)	MT Candidates (threshold off)
Ge classic	2	4	6
Ge neural network	1	7	16
Ge 5D χ^2 (HT)	0	-	-
Ge 5D χ^2 (MT)	-	2	2
Si 2D χ^2	3	3	4

Table 17: Candidate event counts for the various analyses. Recoil threshold *on* or *off* refers to whether the number of candidates excludes or includes, respectively, events below the recoil energy thresholds. These counts include HT candidates for the relevant analyses.

box” analyses such as neural networks, which helps to justify our choice of the classic timing analysis as the primary Ge MT analysis due to its better legibility compared to the neural network analysis.

7.2 A Catalog of Backgrounds

There are several classes of background events, which are itemized below and whose magnitudes are estimated in the following sections.

- **Radiogenic Neutrons**

These neutrons are produced in the natural radioactive decays of the rock and material which surround the experiment. Passive shielding (see Sec. 2.2) protects the experiment from these neutrons, although contamination of the shielding itself contributes to this background.

- **Cosmogenic Neutrons**

These neutrons are the results of muon showers from cosmic rays. The rock overburden of the cavern provides a $\sim 10^5$ reduction in muon flux. The active muon veto, described in Sec. 2.2, tags the parent muons and is used to reject events coincident with these muons, but untagged muon events may contain neutrons in their showers which contaminate the signal region.

- **^{210}Pb - α Events**

^{222}Rn contamination on detector surfaces produces a non-negligible amount of surface-implanted ^{210}Pb , which subsequently decays to ^{206}Pb through ^{210}Bi and ^{210}Po . The β produced by the latter decay (half-life of ~ 5 days) is a likely to appear as a surface electron recoil background event. Additionally, the recoiling ^{206}Pb which accompanies the α also constitute a background component. If the coincident α is not detected by any other detector, the ^{206}Pb nucleus is not tagged and appears as a low-yield event,




Figure 58: WIMP-search data for Ge detector T1Z5. Shaded curves indicate the position of the nuclear recoil (magenta) and below-ER (green) bands, and the ionization (red) and recoil (cyan) energy thresholds. Widths of the curves indicate run-to-run variations. Signal region events for the HT analysis are shown as red points; MT signal region events with the recoil energy threshold “on” and “off” are shown as green and brown points, respectively. Events passing the timing cuts are circled in black. The classic (TOP LEFT) and re-optimized 5D χ^2 (TOP RIGHT) cuts pass zero-yield events at much lower rates than the neural network cut (BOTTOM).

similar to the WIMP signal. Note that the β background is included with the rest of the surface event background for the purposes of estimating the magnitude of the background.

- **Surface Events**

This background, which the phonon timing cuts are specifically designed to remove, are ERs whose ionization signals are depressed due to surface effects such that their low yield mimics that of true NRs. See Sec. 5.1 for a more complete discussion.

7.3 Background Estimation

To properly interpret the results of a WIMP-search analysis, an accurate accounting of the magnitude of the backgrounds must be known. A combination of simulation- and data-based approaches is used to estimate the magnitudes of the various background components for the analyses described in this thesis.

7.3.1 Radiogenic Neutron Background

To estimate the radiogenic neutron background, a two-stage Monte Carlo simulation was used[130]. The primary contaminants were assumed to be daughters of the ^{232}Th and ^{238}U decay chains in the Cu cans, detector housings, and the inner polyethylene and Pb shielding. Additionally, contamination from ^{40}K , ^{60}Co , and ^{222}Rn daughters were included for several of the simulation components. Two distinct simulations were performed using the Geant4 simulation package[131]: one throwing γ events only, and another throwing neutron events. The γ simulation is used to create an electron recoil spectrum extending to high energies (~ 3 MeV), which can be compared to actual CDMS-II data. A best fit to the data is performed by allowing the levels of individual contaminants in the various shielding elements to float. The results of this fitting are shown in Fig. 59.

High-threshold Analyses	
Ge 5D χ^2	0.0284 ± 0.0001 (stat.) ± 0.0035 (syst.)
Ge Classic	0.0249 ± 0.0001 (stat.) ± 0.0031 (syst.)
Ge Neural Network	0.0245 ± 0.0001 (stat.) ± 0.0030 (syst.)
Si 2D χ^2	0.0159 ± 0.0001 (stat.) ± 0.0019 (syst.)
Mid-threshold Analyses	
Ge Re-optimized 5D χ^2	0.0299 ± 0.0001 (stat.) ± 0.0041 (syst.)
Ge Classic	0.0285 ± 0.0001 (stat.) ± 0.0039 (syst.)
Ge Neural Network	0.0288 ± 0.0001 (stat.) ± 0.0039 (syst.)
Si 2D χ^2	0.0160 ± 0.0001 (stat.) ± 0.0020 (syst.)

Table 18: Radiogenic neutron background estimates for the high- and mid-threshold timing analyses. Note that the mid-threshold extensions introduce negligible increases for this background.

Once the contaminant levels have been obtained from fitting the simulated γ spectrum to data, these levels can be used as inputs to the neutron Monte Carlo simulations. The simulated radiogenic neutron spectra for the Ge and Si detectors are shown in Fig. 60. The total expected radiogenic background is obtained by multiplying these spectra with the total efficiencies for each analysis, integrating, and normalizing to the appropriate exposure. The results of this calculation are shown in Tbl. 18. The quoted systematics are obtained by taking 8 reasonable variations to the γ simulation fit which result in small increases to the fit χ^2 per degree of freedom.[132]

7.3.2 Cosmogenic Neutron Background

Background estimates for cosmogenic neutrons are obtained with the aid of a Geant4 Monte Carlo simulation, the geometry of which includes a model of the CDMS-II experimental setup complete with active and passive shielding elements. Simulated muon primaries are created using the MUSUN package[133] based on slant path data from measurements for the Soudan2[134] and MINOS[18] experiments. Analysis of the simulated data is used to tag nuclear recoil events in the detectors, which are further tagged as either single-




Figure 59: Results of fitting the radiogenic γ Monte Carlo to data, with prominent activation peaks highlighted. The simulation (red) and data (black) show good agreement. The contamination values obtained from this best-fit are used as inputs to the neutron Monte Carlo, which produces a radiogenic neutron spectrum used to estimate the magnitude of the background.




Figure 60: Results of radiogenic neutron simulation[132], using the inputs obtained from the γ simulation fits. LEFT: spectrum for the Ge detectors used in the WIMP-search analysis; the spectrum after the timing cut efficiency is shown for the 5D χ^2 (blue), classic (red), and neural network (green). RIGHT: spectrum for the Si WIMP-search detectors. The spectrum after the 2D χ^2 timing cut efficiency is applied is shown in blue.

	Vetoed	Unvetoed	Total
Singles	1409	9	1418
Multiples	4151	12	4163

Table 19: Counts from the cosmogenic neutron simulation for Ge[135].

	Vetoed	Unvetoed	Total
Singles	579	5	584
Multiples	2513	8	2521

Table 20: Counts from the cosmogenic neutron simulation for Si[135].

or multiple-scatters, and as either vetoed (i.e., depositing measurable energy in the veto panels) or unvetoed. In the actual WIMP-search data, observed vetoed nuclear recoils can be counted and used to normalize the simulation results and obtain a leakage estimate. As an alternative, the simulation data can be normalized to the appropriate exposure to obtain a simulation-only leakage estimate. Tbls. 19 and 20 display the counts obtained from the available simulation data, which corresponds to 66.01 live-years of data for Ge and Si[135].

7.3.2.1 Data-normalized background estimates

Using the counts from simulations, there are several possible ratios by which the observed counts can be scaled. Two possible methods, which were determined to be the most reliable, were carried forward. The methods differ in that the first method uses counts from both single- and multiple-scatters, while the second method uses only single-scatters. The trade-off between better statistics (method 1) and fewer systematics (method 2) requires consideration, and in the event of significant disparity between the two methods, method 2 ought to be used.

Eqs. 48 show the mathematical form of the two methods described above, where “S”

or “M” in the subscript indicates single- or multiple-scatters, respectively. Typically, the data counts N^{data} are on the order of 10 or fewer (although this number is different for the different phonon timing analyses; see Sec. 8), and low statistics also plague the unvetoes counts from simulation; therefore, care was taken to properly account for the asymmetric error bars of the final background estimates.

$$N_1 = N_{S+M}^{\text{data}} \times \frac{N_{S,\text{unvetoes}}^{\text{sim}}}{N_{S+M,\text{vetoes}}^{\text{sim}}} \quad (48a)$$

$$N_2 = N_S^{\text{data}} \times \frac{N_{S,\text{unvetoes}}^{\text{sim}}}{N_{S,\text{vetoes}}^{\text{sim}}} \quad (48b)$$

7.3.2.2 Simulation-only background estimates

The cosmogenic neutron background can also be estimated in a third way, directly from the simulated number of leakage events and scaled by the relevant factors to match the actual experimental exposure. This is shown in Eq. 49.

$$N_3 = N_{S,\text{unvetoes}}^{\text{sim}} \times \frac{(MT)_{\text{actual}}}{(MT)_{\text{sim}}} \times \epsilon \times S \quad (49)$$

Here, MT is the exposure, ϵ is neutron efficiency (dependent on analysis), and S is a scaling factor which encapsulates the discrepancy between the simulation and data with respect to detector livetime. This correction is necessary because some of the detectors were not active during the entire c58 period, while in simulation, all detectors are assumed to be active for the entire duration of the simulation. This scaling factor is estimated using the following method.

Using the detectors which had non-zero exposure in each of the 4 data runs, the ratio of the exposure in each detector in each run, relative to run 125, can be computed, and

is found to be quite consistent amongst (non-zero exposure) detectors. Averaging over all detectors, the following scaling factors are obtained.

$$s_{125} = 1 \quad (50a)$$

$$s_{126} = 0.584 \pm 0.046 \quad (50b)$$

$$s_{127} = 0.420 \pm 0.037 \quad (50c)$$

$$s_{128} = 0.215 \pm 0.018 \quad (50d)$$

Using this, the expected exposure of a detector d during a run r in which it was not active⁵⁴ is given by Eq. 51.

$$(MT)_{d,r}^{\text{exp.}} = (MT)_{d,125}^{\text{meas.}} \times s_r \quad (51)$$

Finally, the scaling factor S can be computed using Eq. 52. Sums over the primed variables are done only over those detectors and runs for which they were not active; sums over unprimed variables are over all detectors and runs.

$$S = \frac{\sum_{d,r} (MT)_{d,r}^{\text{meas.}}}{\sum_{d,r} (MT)_{d,r}^{\text{meas.}} + \sum_{d',r'} (MT)_{d',r'}^{\text{exp.}}} \quad (52)$$

The result of this calculation for Ge is $S_{\text{Ge}} = 0.956 \pm 0.002$, and for Si is $S_{\text{Si}} = 0.924 \pm 0.006$.

7.3.2.3 Results

Tbl. 21 displays the results of the three cosmogenic neutrons background estimation

⁵⁴All detectors under consideration in this analysis were active in run 125.

High-threshold Analyses			
	N_1	N_2	N_3
Ge 5D χ^2	$0.0188^{+0.0152}_{-0.0075}$	$0.0360^{+0.0387}_{-0.0169}$	0.0214 ± 0.0081
Ge Classic	$0.0188^{+0.0152}_{-0.0075}$	$0.0283^{+0.0332}_{-0.0141}$	0.0187 ± 0.0071
Ge Neural Network	$0.0205^{+0.0161}_{-0.0080}$	$0.0283^{+0.0332}_{-0.0141}$	0.0183 ± 0.0069
Si 2D χ^2	$0.0075^{+0.0090}_{-0.0037}$	$0.0063^{+0.0181}_{-0.0048}$	0.0129 ± 0.0058
Mid-threshold Analyses			
	N_1	N_2	N_3
Ge 5D χ^2	$0.0195^{+0.0139}_{-0.0071}$	$0.0280^{+0.0280}_{-0.0132}$	0.0218 ± 0.0073
Ge Classic	$0.0195^{+0.0139}_{-0.0071}$	$0.0281^{+0.0279}_{-0.0127}$	0.0235 ± 0.0079
Ge Neural Network	$0.0257^{+0.0172}_{-0.0090}$	$0.0220^{+0.0243}_{-0.0106}$	0.0237 ± 0.0079
Si 2D χ^2	$0.0068^{+0.0082}_{-0.0034}$	$0.0053^{+0.0153}_{-0.0040}$	0.0120 ± 0.0054

Table 21: Cosmogenic neutron background estimates for the high- and mid-threshold timing analyses, using various methods.

methods outlined in Secs. 7.3.2.1 and 7.3.2.2. Since the simulation-only method (N_3) is usually the most conservative estimate, has the best statistics, and is in agreement with the other methods, we quote it as the final estimate for the cosmogenic background; the spread of the three estimates is then quoted as an estimate of the systematic uncertainty in this estimate, which is displayed in Tbl. 22.

7.3.3 ^{210}Pb - α Event Background

Any recoiling ^{206}Pb nucleus should be accompanied by a correlated high-energy (5.3 MeV cutoff) α ; therefore, by counting the rate of these high-energy α s we also obtain an estimate of the rate of ^{206}Pb nuclei. Unfortunately, since many of these Pb-recoil events deposit very little ionization energy, the resolution on the ionization pulse start time is poor; this has the effect of rendering the relative phonon start time quantity effectively random and diminishing the discrimination power of the phonon timing cuts on these events.

Using the measured rate of α events in detectors, and conservatively estimating the phonon timing cut passage fraction of Pb-recoil events, a conservative upper limit on this

High-threshold Analyses	
	Combined Estimate
Ge 5D χ^2	0.0214 ± 0.0081 (stat.) ± 0.0086 (syst.)
Ge Classic	0.0187 ± 0.0071 (stat.) ± 0.0096 (syst.)
Ge Neural Network	0.0183 ± 0.0069 (stat.) ± 0.0100 (syst.)
Si 2D χ^2	0.0129 ± 0.0058 (stat.) ± 0.0033 (syst.)
Mid-threshold Analyses	
	Combined Estimate
Ge 5D χ^2	0.0218 ± 0.0073 (stat.) ± 0.0043 (syst.)
Ge Classic	0.0235 ± 0.0079 (stat.) ± 0.0044 (syst.)
Ge Neural Network	0.0237 ± 0.0079 (stat.) ± 0.0019 (syst.)
Si 2D χ^2	0.0120 ± 0.0054 (stat.) ± 0.0034 (syst.)

Table 22: Final cosmogenic neutron background estimates for the high- and mid-threshold timing analyses.

background has been set at 0.174 events[136] for all analyses.

7.3.4 Surface Event Background

The methods for estimating the leakage due to SEs was outlined in Sec. 5.2, since such estimates are necessary for the timing cut tuning algorithms. Whereas those estimates utilized calibration and WIMP-search sideband data, a final estimate can be made using WIMP-search signal region events which fail the timing cuts.

There are two differences between how the final estimates are computed, and how they were computed in Tbls. 8 - 11: the timing cut passage fractions (number passing out of total) in Eqs.27 and 29 are replaced with pass-fail ratios, and the normalizations (which were previously estimated; see Sec. 5.2.1) are replaced by the observed number of failing WIMP-search signal region singles for each analysis. The results of this estimation are shown in Tbl. 23.

Finally, Tbl. 24 shows the combined background estimate, assuming the backgrounds are uncorrelated, compared to the observed candidates for each analysis.

High-threshold Analyses				
	Prelim. Est.	Est. NRSS	Fail Timing	Final Est.
Ge 5D χ^2	0.57 ± 0.11	162.15 ± 27.43	176	$1.19^{+0.23}_{-0.21}$
Ge Classic	0.57 ± 0.15	162.15 ± 27.43	175	$0.64^{+0.17}_{-0.15}$
Ge Neural Network	0.54 ± 0.29	162.15 ± 27.43	177	$0.87^{+0.24}_{-0.21}$
Si 2D χ^2	0.41 ± 0.17	145.97 ± 15.47	88	$0.54^{+0.20}_{-0.14}$
Mid-threshold Analyses				
	Prelim. Est.	Est. NRSS	Obs. NRSS	Final Est.
Ge 5D χ^2 (8 GeV re-opt.)	1.22 ± 0.24	79.67 ± 14.02	128	1.82 ± 0.31
Ge Classic	1.50 ± 0.25	92.48 ± 15.81	84	1.48 ± 0.20
Ge Neural Network	1.14 ± 0.20	92.48 ± 15.81	82	1.39 ± 0.21
Si 2D χ^2	1.14 ± 0.36	112.12 ± 16.67	133	1.13 ± 0.31

Table 23: Final SE leakage estimates. The preliminary estimates were obtained without using masked signal region data. The observed number of events failing the timing cut (from the unmasked data) replaced the NR single-scatter estimates for the final estimate. Only a lower energy range (< 15 keV for re-optimized Ge 5D χ^2 , < 20 keV for others) is shown for the mid-threshold analyses, since the analyses are identical at higher energies.

High-threshold Analyses

	Radio. neut.	Cosmo. neut.	Pb- α	SE	Total	# Cands.
Ge 5D χ^2	0.0284 ± 0.0001	0.0214 ± 0.0081	< 0.174	$1.19^{+0.23}_{-0.21}$	1.33 ± 0.23	0
Ge Classic	0.0249 ± 0.0001	0.0187 ± 0.0071	< 0.174	$0.64^{+0.17}_{-0.15}$	0.77 ± 0.17	2
Ge Neural Network	0.0245 ± 0.0001	0.0183 ± 0.0069	< 0.174	$0.87^{+0.24}_{-0.21}$	1.00 ± 0.24	1
Si 2D χ^2	0.0159 ± 0.0001	0.0129 ± 0.0058	< 0.174	$0.54^{+0.20}_{-0.14}$	0.66 ± 0.18	3

Mid-threshold Analyses

	Radio. neut.	Cosmo. neut.	Pb- α	SE	Total	# Cands.
Ge 5D χ^2	0.0299 ± 0.0001	0.0218 ± 0.0073	< 0.174	1.82 ± 0.31	1.96 ± 0.32	2
Ge Classic	0.0285 ± 0.0001	0.0235 ± 0.0079	< 0.174	1.48 ± 0.20	1.62 ± 0.21	4
Ge Neural Network	0.0288 ± 0.0001	0.0237 ± 0.0079	< 0.174	1.39 ± 0.21	1.53 ± 0.22	7
Si 2D χ^2	0.0160 ± 0.0001	0.0120 ± 0.0054	< 0.174	1.13 ± 0.31	1.25 ± 0.32	3

Table 24: Total background estimates, combined assuming no correlation, and observed candidate events. The Ge 5D χ^2 MT analysis is the one re-optimized for 8 GeV/ c^2 . Uncertainties shown are statistical only. Since reliable background estimates for the MT analyses only exist when recoil thresholds are on, candidate events listed refer to that scenario.

8 WIMP-search Results of the Re-analyzed c58 Data

8.1 Limit-setting in the Standard Halo Model

Once an analysis has been unmasked, and a set of candidate events is obtained, it is standard to interpret the results in the context of excluding (in the case of a null result) or favoring (in the case of a large signal) regions of the WIMP mass - coupling parameter space. Such a process is necessarily model dependent, both in the sense that one requires a model for the coupling of the WIMPs to the detector material, and that the astrophysical model used to describe the large-scale distribution of the dark matter will affect the results. In general, there are several ingredients which go into this process:

1. Information about how the analysis performed accepts WIMPs - i.e., the efficiency.
2. A set of WIMP candidate events and their recoil energies.
3. A model for the interaction of the WIMPs with the standard model particles in the detector.
4. A model for the astrophysical behavior of the WIMP distribution.
5. (Optional) A reliable model for the backgrounds expected for the analysis.

The final optional component is only necessary if one wishes to either set a confidence region (i.e., interpret the results as a WIMP detection), or set a limit using background subtraction. Obviously, either of these methods will be subject to the systematics of the background model; the most conservative approach is to assume that all candidates are real WIMP events and use the events to set an upper limit on the WIMP interaction.

8.1.1 Models for WIMP Interaction

In order to put constraints on the WIMP interaction with the target nuclei, some model for the interaction - specifically, a model for the interaction cross-section - must be pre-

scribed. The simplest model assumes the interaction spin-independent (SI), and therefore neutrons and protons contribute coherently to the cross-section, which goes as A^2 , where A is the number of nucleons. Alternatively, one can construct a spin-dependent (SD) interaction, wherein paired nucleons contribute to the cross-section with opposite signs, effectively negating their contributions. In the SD model, only unpaired neutrons and protons interact with the WIMPs, restricting the number of eligible isotopes and also eliminating the coherent A^2 dependence enjoyed by the spin-independent model. In general, SI interactions are best probed by large nuclei due to the A^2 cross-section scaling, while experiments probing SD interactions generally favor lighter targets for kinematic reasons.

Both of these models are widely used, and the relevant cross-sections, nuclear form factors, and numerical corrections have been computed and are well-understood (see, e.g., [110] and [137]). Owing to their simplicity, these models are widely used in the reporting of WIMP-search results; recently, however, the SI model has come to be used nearly exclusively by several direct-detection collaborations, including CDMS, in large part due to the recent prominence of heavy-target experiments, and the fact that SI cross-sections can be probed many orders of magnitude deeper than SD cross-sections due to the A^2 amplification and corresponding kinematic enhancements. In any case, the results in this thesis will be presented in the standard SI framework⁵⁵ for consistency with the broader direct-detection community.

Despite the hegemony of the SI/SD framework, recent work has shown that these two models fail to span the possible space of WIMP-nucleus interactions. In particular, an effective field theory (EFT) approach shows that there are six independent couplings which can be constructed[138], three of which are velocity-dependent; additionally, the finite size of the target nuclei, and inelastic nuclear excitations (which are accounted for through a form factor in the SI/SD formulation) are more naturally accounted for the EFT model.

⁵⁵This is the framework used in Secs. 5 and 6 for the optimization of the phonon timing cuts.

Despite the promising nature of these models, few comprehensive analyses using them have been undertaken (e.g., [139]). This is partially due to the novelty of the approach, but also the fact that some of the relevant nuclear physics for certain target isotopes is poorly understood; unfortunately, Si is one such element for which the necessary information is lacking.

Finally, the WIMP-nucleon cross-section may contain an inelastic scattering contribution via the creation of an excited WIMP state[140]; this scenario is important if the mass splitting, δ , is on the order of the expected WIMP recoil energies (10s or 100s of keV) because this sets a cut-off below which this inelastic scattering cannot occur. In particular, this model has been proposed as a way to resolve the tension between the claimed DAMA/LIBRA signal and other experiments which exclude it.[141] Although it remains an interesting possibility, the inelastic dark matter scenario will not be considered further here.

8.1.2 Astrophysical Models for Dark Matter

The astrophysical properties of the dark matter halo, and of the Milky Way galaxy, affect both the normalization and the shape of the WIMP spectrum. Given a model physics of the WIMP-nucleon interaction (Sec. 8.1.1), one requires only the mass and velocity distribution of the Earth relative to the halo to completely determine the WIMP recoil energy spectrum. The relevant parameters described below, and summarized in Tbl. 25.

The observed WIMP event rate will be directly proportional to the local mass density of dark matter. The standard value, obtained assuming a spherical halo with a cusp, is $\rho_0 = (0.30 \pm 0.05) \text{ GeV/cm}^3$ [142]; although we only require this value at one location (i.e., the Earth), its measurement requires assumptions about the mass distribution, thus introducing some systematic uncertainty into these measurements.

The shape of the recoil spectrum is determined by the velocity distribution of the

Quantity	Symbol	Accepted Value	Source
Local D.M. Density	ρ_0	$(0.30 \pm 0.05) \text{ GeV/cm}^3$	[142]
Local circular speed	v_c	$(220 \pm 20) \text{ km/s}$	[143]
Galactic escape velocity	v_{esc}	$(544^{+64}_{-46}) \text{ km/s}$	[145]
Peculiar solar speed	$v_{\odot}$	$(13.4 \pm 0.8) \text{ km/s}$	[144]
Earth orbital speed	v_E	29.8 km/s	[146]

Table 25: Physical parameters of the standard halo model, and their commonly accepted values.

WIMPs. In the standard halo model (SHM), an isotropic Gaussian velocity distribution is assumed, as in Eq. 53.

$$f(\vec{v}) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{|\vec{v}|^2}{2\sigma^2}\right) \quad (53)$$

The velocity dispersion is given by $\sigma\sqrt{3/2}v_c$, where v_c is the local circular velocity relative to the galactic center. Using the standard value for this quantity, $v_c = (220 \pm 20) \text{ km/s}$ [143], one obtains $\sigma = 269 \text{ km/s}$. The peculiar speed of the Sun, $v_{\odot} = (13.4 \pm 0.8) \text{ km/s}$ [144], and the orbital speed of the Earth, $v_E = 29.8 \text{ km/s}$, complete the description of the Earth's motion relative to the halo.

The SHM velocity distribution described by Eq. 53 formally extends to infinite velocity, but (assuming the WIMP distribution has reached a steady state) particles exceeding the local escape velocity $v_{\text{esc}} = \sqrt{\frac{2GM}{R}}$ will have been ejected and will not contribute to the velocity distribution; thus $f(\vec{v})$ must be manually truncated above the measured escape velocity, whose standard value is $v_{\text{esc}} = (544^{+64}_{-46}) \text{ km/s}$ [145].

For the SHM, the differential WIMP event rate can be written approximately as in Eq. 54[110], where F^2 is a nuclear form factor, R_0 is the rate in the $E_R \rightarrow 0$ limit, and E_c is the characteristic energy scale. Differential rate curves for different WIMP masses and target materials are shown in Fig. 61.

$$\frac{dR}{dE_R} \approx R_0 F^2(E_R) \exp\left(-\frac{E_R}{E_c}\right) \quad (54)$$

8.1.3 Optimum Interval Method

Using a Bayesian or Feldman-Cousins[147] likelihood analysis, one can use the known shape of the WIMP differential spectrum to set a confidence interval for the spin-independent cross-section, as a function of WIMP mass. In this way, one can construct confidence intervals which describe regions of parameter space which are favored by the experimental results.⁵⁶ Such analyses, however, are difficult because they require an understanding of the experimental backgrounds which is both accurate and complete. If one is not confident in their background description, or if the observed number of events is consistent with the background expectation, then the most conservative interpretation of results is to assume that all events are signal, and compute an upper limit on the cross-section under this assumption.

The method by which one computes the upper limit must be carefully chosen, so that the limit is as constraining as the data allows, without being biased. The simplest possible method is to count the number of candidate events in some energy range, compute the 90% Poisson upper limit on that number, and use that to compute a limit on the rate (and thus the cross-section) of WIMPs. Besides throwing away spectral information (which makes the limit overly conservative), this method is also biased because it depends sensitively on the choice of energy range used. Therefore we desire a method which both takes into account the spectral information of the candidates, but which also removes the arbitrary choice of the energy range.

Yellin's optimum interval method[148] satisfies these criteria. It is a general method which requires two inputs: the energies E_i of each of the N candidate events, and the energy spectrum of the expected signal, $\frac{dN}{dE}$; given the shape of the spectrum (which is set by the WIMP mass and astrophysical parameters described above), the method sets an

⁵⁶The analysis which produced the favored region for the CDMS-II Si result[124] used a Bayesian likelihood analysis; the difficulty in background modeling led the authors to specifically state that the result is not robust enough to claim detection.

upper limit on the normalization of the spectrum (i.e., the cross-section).

In the optimum interval method, we define an “ n -event gap” as an energy interval, bounded by either candidate events or the upper or lower energy thresholds, within which there are n candidate events. Thus, for an experiment with N candidate events, one would find $(N + 1)$ 0-event gaps, (N) 1-event gaps, $(N - 1)$ 2-event gaps..., and one N -event gap. For each of these gaps, one computes the expected signal, $\langle S \rangle = \int_{E_1}^{E_2} \frac{dN}{dE} dE$, and compares it with the observed n events to compute a probability $C_n(\langle S \rangle)$ that the expected signal spectrum would produce more events than were observed. For N candidate events, there will be $\frac{1}{2}(N + 1)(N + 2)$ total gaps for which this quantity must be computed; if the maximum of all these probabilities is $C_{\max}$, then the assumed cross-section is rejected at the $C_{\max}$ confidence level. If one desires a particular confidence level beforehand (90% is common), then one adjusts the cross-section until the maximum $C_n(\langle S \rangle)$ is equal to that value.

In simpler terms, the optimum interval method looks for the least probable gap in the candidate event spectrum and uses it to set a limit on the expected signal. Although at first glance this may seem biased, since the energy interval used is set by the data, the opposite is actually true - because the method seeks out intervals which have low event counts relative to expected signal, it necessarily avoids regions for which there are large backgrounds. The method is especially well-suited to experiments with both small signals and small backgrounds, making it a natural fit for the CDMS-II experiment.

In practice, limit curves are set by scanning over a range of WIMP masses and setting a confidence limit on the cross-section at each mass; all limit curves shown in this thesis are produced in this way, at a 90% confidence level.

8.2 Limits

The results of the high-threshold phonon timing analyses are shown in Fig. 62. The Ge results are shown as upper limits, with the 5D χ^2 limit shown as a thick line to represent that it was chosen to be the primary analysis for the Ge high-threshold exposure. The Si result is shown both as an upper limit, and as a closed contour based on a likelihood analysis of the 3 Si candidate events.[124]

CDMS-II data was also analyzed in a similar fashion for runs 118-119[149]⁵⁷ and runs 123-124.[150, 151]. It is straightforward to combine these exposures and results to obtain limit curves for these combined exposures; these results are shown in Fig. 63.

The limits for the MT analyses are presented here in several different ways. In Figs. 64 - 69, the limits for each individual MT analysis are displayed with the recoil energy thresholds turned on and off, with the sensitivity projections from Sec. 6.2 shown for comparison. Figs. 70 and 71 present the limits for all of the MT analyses with the recoil thresholds on and off, respectively. Recall that the Ge classic and Si 2D χ^2 with recoil thresholds on were chosen prior to unmasking to be the primary analyses.

In general, the limits obtained with the recoil thresholds on agree reasonably well with the simple sensitivity estimates obtained from background model 2+; the obvious exceptions are the Ge results with no timing cuts applied (Fig. 67), whose limits are much worse than those predicted from simulations. In fact, for this analysis, even the most conservative model (model 3) does not accurately predict this behavior. Although the models show some utility for tuning the analysis, they are too simple to be relied upon in any substantial manner.

Another, more subtle mis-projection can be seen in the limits for the Si 2D χ^2 analysis (Fig. 68). The limit obtained is higher than the projection, but not significantly so; however, the fact that the higher limits tend to skirt around the best-fit region can be interpreted

⁵⁷Runs 118 and 119 were run with only ZIP towers 1 and 2 (12 detectors).




Figure 61: Differential event rates in the SHM for $\sigma_{\text{SI}} = 10^{-42} \text{ cm}^2$, with a $60 \text{ GeV}/c^2$ WIMP (solid lines) and an $8 \text{ GeV}/c^2$ WIMP (dashed lines). Xe (green) has the highest maximum rate in both cases, due to its large nucleus, but kinematics disfavor it for lower WIMP masses. Conversely, Si (red) has the lowest overall rate but its spectrum extends to higher energies for lower WIMP masses. Ge (blue) enjoys intermediate behavior.




Figure 62: Spin-independent results of the high-threshold analyses for Ge and Si. The primary Ge result - the $5\text{D } \chi^2$ analysis (solid black) - gives strong results at high mass, but the neural network (red solid) analysis outperforms it at lower masses (see inset); the classic analysis result (blue dashed) is intermediate at high mass. The Si results are presented as both an upper limit (magenta solid), and as a best-fit region (68% green, 90% yellow; black cross is the best-fit point).




Figure 63: Spin-independent results of the high-threshold analysis of the full CDMS-II exposure. The limit curves are colored the same as in Fig. 62; the Si best-fit region is unchanged, and is shown for reference. The Ge limits extend to lower mass (compared to c58R-only results) and converge because of the 7 keV threshold of runs 118-119; all other Ge runs had a 10 keV threshold.




Figure 64: Spin-independent results for the Ge MT classic timing analysis, which was chosen as the primary Ge MT analysis. The blue curve is the limit with the recoil threshold turned on; the red curve is with the threshold off. The shaded band displays the 90% confidence interval on the expected sensitivity (estimated with the recoil threshold on). The HT Si region is shown for comparison.




Figure 65: Spin-independent results for the Ge MT neural network timing analysis. The blue curve is the limit with the recoil threshold turned on; the red curve is with the threshold off. The shaded band displays the 90% confidence interval on the expected sensitivity (estimated with the recoil threshold on). The HT Si region is shown for comparison.




Figure 66: Spin-independent results for the Ge MT 5D χ^2 timing analysis, re-optimized for a WIMP mass of $8 \text{ GeV}/c^2$. The blue curve is the limit with the recoil threshold turned on; the red curve is with the threshold off. The shaded band displays the 90% confidence interval on the expected sensitivity (estimated with the recoil threshold on). The HT Si region is shown for comparison.




Figure 67: Spin-independent results for the Ge MT analysis with no timing cut applied. The blue curve is the limit with the recoil threshold turned on; the red curve is with the threshold off. The shaded band displays the 90% confidence interval on the expected sensitivity (estimated with the recoil threshold on). The HT Si region is shown for comparison.




Figure 68: Spin-independent results for the Si MT 2D χ^2 timing analysis, which was chosen as the primary Si MT analysis. The blue curve is the limit with the recoil threshold turned on; the red curve is with the threshold off. The shaded band displays the 90% confidence interval on the expected sensitivity (estimated with the recoil threshold on). The HT Si region is shown for comparison.

in one of two ways: some persistent background, misunderstood background is mimicking a $\sim 8 - 10 \text{ GeV}/c^2$ WIMP spectrum; or a real WIMP signal is present. The strong tension with the Ge results (and other experiments; see Fig. 72) supports the former interpretation, at least within the standard spin-independent interaction framework.

With the recoil thresholds on, the Ge classic and neural network analyses produce similar limits, despite the neural network producing more candidates. This illustrates the interplay between efficiency and rejection power; if the recoil thresholds are removed, the neural network gains 9 more candidates and its superior efficiency is not enough to keep the classic analysis from overtaking it. The poor limits set by the Ge and Si analyses with no timing cut applied provide confirmation that timing information still provides useful discrimination power down to low recoil energies.

The results of both the HT and MT analyses are summarized in Fig. 72, which shows the primary HT and MT results for the Ge and Si analyses, along with results from other relevant experiments for context. The MT analysis was more successful for Ge than for Si, possibly due to the unknown Si background discussed above. The tension among (and, for CDMS, within) experiments and target materials remains unresolved; a more nuanced model of dark matter interactions (e.g., an EFT framework) could help resolve this tension, but at present this avenue remains relatively unexplored. More discussion about the performance of the analyses can be found in Sec. 9.

8.3 Further Investigation of Si NRSS Distributions

The most important result of the c58 reanalysis is the three candidate events observed in the Si detectors, and their possible interpretation as a signal from an $8.6 \text{ GeV}/c^2$ WIMP (see Sec. 8). In light of the importance of these candidate events, further investigation into the distributions of their relevant parameters may help us to better understand how plausible their interpretation as a WIMP signal is. These checks are performed in two ways: first,




Figure 69: Spin-independent results for the Si MT analysis with no timing cut applied. The blue curve is the limit with the recoil threshold turned on; the red curve is with the threshold off. The shaded band displays the 90% confidence interval on the expected sensitivity (estimated with the recoil threshold on). The HT Si region is shown for comparison.




Figure 70: Spin-independent results for all MT analyses, with recoil energy thresholds turned on: Ge classic (thick blue), Ge neural network (green), Ge 5D χ^2 (red), Ge no timing cut (cyan), Si 2D χ^2 (thick magenta), and Si no timing cut (yellow). The HT Si region is shown for comparison.

on a detector-by-detector basis, where distributions from calibration data are compared to those observed in WIMP-search data; and second, in a more global manner, where the timing cut optimization routine is used to systematically vary the placement of the timing cut and event distributions are compared across all detectors simultaneously. The former checks provide useful insight into previously unaccounted-for systematics between the ^{133}Ba calibration data and WIMP-search data, while the latter are useful for understanding how leakage-fitting systematics can affect the optimization of the final timing cut.

8.3.1 Individual Detector Cross-checks

The level at which the timing cut is set, in each detector and energy bin, is informed by the distributions of ^{252}Cf neutrons and low-yield ^{133}Ba ERs, which stand in as proxies for WIMP signals and SE backgrounds, respectively. Furthermore, leakage fits based on the ^{133}Ba distributions are used in the timing cut optimization routine, adding another possible layer of systematics into the placement of the final timing cut.

Since Si 2D χ^2 timing cut was set for the HT analysis⁵⁸, it is of primary interest how the calibration distributions compare to the WIMP-search data in this regime. Fig. 73 shows how these distributions compare in the lowest energy bin (< 20 keV) for each detector. The raw counts of WIMP-search signal-region events are shown in blue, and raw neutron counts from ^{252}Cf are shown in green. The SE curve, in blue, is constructed from low-yield ^{133}Ba multiple scatters in nearest-neighbor detectors; they are added together with the appropriate P -side and Q -side scaling factors, and normalized according to the technique described in Sec. 5.2.1. The fit to the SE distribution is shown with a dashed black line, and the optimized cut position is designated with a vertical dashed magenta line.

It is readily apparent that most of the detectors show good agreement between the WIMP-search distributions and those from ^{133}Ba ; however, the two detectors in which

⁵⁸Recall that the timing cut was not re-optimized for the MT analysis.

candidates were observed (T4Z3 and T5Z3) are notably discrepant.⁵⁹ Despite the attempt to fit the tails of the ^{133}Ba distributions as conservatively as was reasonable, the fits drastically underestimate the observed WIMP-search distributions in the tails of these two detectors, although this is apparently due to fundamental discrepancies between the distributions themselves, and not a result of poor or biased fitting.

Since no new candidate events were observed when the Si 2D χ^2 timing analysis was extended into the MT range, we should also examine how the timing distributions change when this extension is performed. Fig. 74 shows these distributions in the MT analysis range, along with a new fit to the MT ^{133}Ba data. These fits were intentionally made to be less conservative in the tails than those for the HT analysis, in order to study how systematics in the fits propagate into the final timing cut (see Sec. 8.3.2).

The compatibility between the ^{133}Ba distributions and those from WIMP-search data seems to improve for almost all detectors (see, e.g., T4Z1) when the analysis is extended into the MT range. In particular, the distributions for T5Z3 are now in excellent agreement; however, those for T4Z3 are not substantially improved. This behavior indicates that there may be detector-dependent systematics which may involve correlations between energy, timing, and possibly yield.

The signal region for the WIMP-search data is defined in the yield-energy plane using a $\pm 2\sigma$ NR band and a 3σ below-ER band cut. However, the SE population is drawn from ^{133}Ba data in a much larger region of the yield-energy plane, including all events with yield between 0.1 and 0.7 which are also at least 5σ below the ER band.⁶⁰ Although the expanded SE signal region definition ensures better statistics, it also means that the ^{133}Ba calibration data will be preferentially weighted to higher yields than the events in the final WIMP-search signal region.

⁵⁹ Another detector, T2Z2, is also rather bad in this regard, but the low WIMP-search statistics (two events) and somewhat serendipitous placement of the timing cut just beyond the furthest outlier make it less interesting for the present discussion.

⁶⁰ A threshold on ionization energy is also applied in each case.

Figs. 75 and 76 display the ^{133}Ba and WIMP-search distributions for T4Z3 and T5Z3, respectively. One immediately notices that the events added in the extension to the MT range (red events) are clearly located in a lower-yield region than those in the HT range (blue events). For T5Z3 (and several of the other detectors), there is clearly a correlation between lower yield and higher (more NR-like) timing, but this correlation is absent in T4Z3 - although whether this is due to a poor statistics or real detector physics is difficult to distinguish. Note that, although almost all of the WIMP-search data fall below a yield of 0.4, much of the ^{133}Ba data lies above this value. The fact that low-yield (less than 0.4) events in T5Z3 tend to have more NR-like timing for MT events than HT events with similar yields may indicate that there is an energy dependence as well, but these effects are hard to disentangle.

8.3.2 Simultaneous Optimization Cross-checks

In the previous section, we have examined how the low-energy distributions from calibration and WIMP-search data compare on a detector-by-detector basis. However, the actual timing cut was set using a method (the slope-matching method; see Sec. 5.1.2.2) which optimized the sensitivity of the cut over all detectors simultaneously. Thus, simply looking at each detector and noting how close or far the cut positions were from WIMP-search outliers may be misleading, because the cut positions in each detector may not respond linearly (or even continuously) to small changes in the optimization routine.

For this reason, it may be instructive to systematically vary the cut positions in all detectors simultaneously by modifying the equal-slopes optimization routine to scan across a range of possible timing cuts. This can be achieved by setting the equal-slope to a wide range of values, rather than just the optimized one, and computing expected leakages, observed candidate events, and expected WIMP signal at each possible timing cut. For simplicity, this technique (henceforth referred to as "scanning the timing cut") is performed

only in the lowest energy bin (< 20 keV) for each of the eight detectors.

For the HT analysis, this scan is implemented quite simply using the existing fits of exposure and leakage vs. timing parameter, which are combined to construct the the $\frac{d\text{Exposure}}{d\text{Leakage}}$ curves, which are then used to fix the cut positions in each detector (see Fig. 34). As one scans across slopes, one can systematically tighten (higher slope) or loosen (lower slope) the timing cut. For the MT analysis, the leakage curves were re-fit (shown in Fig. 74) but the exposure fits remained the same. As mentioned in the previous section, the leakage fits used for the scanning of the MT timing cut were intentionally chosen to be less conservative in the tails of the distributions and more faithful in the bulk, allowing us to observe the effect of such a systematic in the placement of the timing cut.

At each point over which the timing cut is scanned, several quantities are computed: the number of WIMP-search signal region events which pass this timing cut; the expected SE leakage from the leakage fits; the expected SE leakage directly from the scaled ^{133}Ba data (obtained by reading off the value of the red curve in Figs. 73 and 74); and the number of ^{252}Cf calibration neutrons which pass the cut (obtained by reading off the value of the green curve in the aforementioned figures).

The acceptance of neutrons into the signal region can be interpreted as a proxy for WIMP signals, but it must be normalized properly if it is to stand in as an estimate of the expected WIMP signal. To do this, the total number of WIMPs expected at a given WIMP mass (say, 10 or 10 GeV/ c^2) and some nominal cross-section (10^{-42} cm 2 is customary) is computed using the measured efficiencies of the Si 2D χ^2 timing cut. Since, for a fixed WIMP mass, the rate is simply proportional to the cross-section, we can scale this number for any arbitrary cross-section by a simple factor of $\frac{\sigma_{\text{SI}}}{10^{-42} \text{ cm}^2}$. We use the 90% confidence region of the Si HT likelihood analysis[124] to define upper- and lower-limit cross-sections for any WIMP mass within the confidence region.

Once we have a range of the expected number of WIMPs for a given WIMP mass, these

can be used to normalize the neutron distributions from the timing cut scanning technique. Specifically, since the expected number of WIMPs was computed for the optimized timing cut, we should rescale the neutron distribution such that the point which most closely corresponds to the optimized timing cut (i.e., within rounding errors) such that it is equal to the expected number of WIMPs computed previously.

The results of these timing-cut scans are shown in Figs. 77 (for the HT analysis) and 78 (for the MT analysis). There are several interesting features in these plots which elucidate how the aforementioned distribution and fit systematics propagate into the cut optimization algorithm, which will be discussed in the next section.

8.3.3 Conclusions

As an initial observation, it is obvious that the leakage estimates (whether directly from ^{133}Ba data or fits) are not particularly consistent with the observed WIMP-search event distributions in the vicinity of the optimized timing cut. As the scans progress to higher leakages, we see that the agreement between the leakage estimates and observations improve, but this region (with expected leakages much greater than 1) is not a regime in which we would ever expect a reasonable timing cut to be placed. The fact that the WIMP-normalized ^{252}Cf neutron distributions match the observations is not too surprising, since the WIMP parameter space of interest from which this curve is constructed was fit to the three observed candidate events, but nonetheless we can attempt to answer a few relevant questions regarding these event distributions: *Can the observed WIMP-search distributions be explained well with background distributions?* and *How different would the timing cut had to have been to exclude some or all of the 3 candidate events?*

The first interesting observation which can be drawn from the results of the timing cut scanning procedure relates to how the fitting of ^{133}Ba data affects the final placement of the timing cut. This behavior can be explored by comparing the red curves (the direct ^{133}Ba

distributions) to the green curves (from fits made to the red curves). In the HT analysis (Fig. 77), the fits were intentionally chosen to be quite conservative in the tails of the distributions; this effect can be seen in how the fits overestimate the true ^{133}Ba distribution at low leakage values. When the tails were fit more faithfully, as in the MT analysis (Fig. 78), there is a consistent low bias in the fits with respect to the true distribution. This behavior can be understood by noting that the optimization routine has no information about the quality of the fits being used, so it will tend to preferentially select detectors with underestimated fits (since they maximize sensitivity) and ignore overestimated detectors.

The conservativeness of the HT fits resulted in the final timing cut being set more tightly than it would have been based on ^{133}Ba distributions alone, yet still not tightly enough to exclude all candidate events. A closer look at individual detector distributions (Figs. 73 and 74) for T5Z3 reveals that, perhaps with a slightly more conservative leakage fit or a more inclusive SE definition in the ^{133}Ba data, the cut position could have been set far enough out to exclude the single candidate in that detector. T4Z3, which had two candidate events, cannot be explained in the same way - the fits were made as conservatively as was reasonably possible,⁶¹ and the yield and energy systematics which are apparent on other detectors are not as pronounced in this one. In the end, we can make a reasonable case for the T5Z3 candidate as an “unlucky” outlier, but the two candidate events in T4Z3 would likely have passed any reasonable timing cut which was tuned using ^{133}Ba calibration data.

As a final observation for Figs. 77 and 78, note that there is no timing cut optimization in the region interest for which the ^{133}Ba distributions overestimate the true number of observations. While the yield-timing systematics discussed in the previous section may account for this to some extent by underpopulating the tails of the ^{133}Ba distributions, another mechanism exists via the normalization of the calibration distributions which is

⁶¹The fit for T4Z3 in the MT analysis is much less conservative, which causes the two candidates to be nearly un-removable under any reasonable timing cut optimization, and causing the black curve to bottom out at 2 events on the left side of Fig. 78.

estimated from c34 data (see Sec. 5.2.1). Even though the total number of expected NRSS in the WIMP-search signal region appears to be unbiased (Tbl. 23) when summed over all detectors, fluctuations on a detector-by-detector level will tend to systematically bias the optimization routine in a manner similar to that described earlier for the ^{133}Ba fits. In particular, if the normalization for a certain detector is overestimated, the cut position in that detector will tend to be set more tightly due to the high number of expected events; similarly, an under-normalized detector will have a looser cut position due to its low number of expected events. Note that this mechanism will cause a systematic underestimation of the true number of candidate events even if there is good agreement between the shapes of the calibration and WIMP-search timing distributions.

In the final accounting, it appears that only the candidate event in T5Z3 can be truly thought of as “leakage”, i.e., a marginal outlier that just barely seems to sneak past a reasonably-defined cut. The two candidates in T4Z3 are far enough removed from the calibration distribution that they cannot be thought of as arising from the same distribution as the ^{133}Ba data. The leakage-scanning method shines a light on a few systematic biases present in the timing cut optimization routine which will always tend to allow more leakage than estimates would suggest - namely, that the routine will always be a “foolish optimist”, prioritizing detectors with underestimated leakages. This problem (which could also be called the “too good to be true” problem) could be ameliorated through the use of a feedback mechanism which conveys information about the quality of the fits and normalizations during the optimization procedure; alternatively, a blind estimate of the magnitude of this effect could be added as a systematic uncertainty to the total leakage estimate.




Figure 71: Spin-independent results for all MT analyses, with recoil energy thresholds turned off: Ge classic (thick blue), Ge neural network (green), Ge 5D χ^2 (red), Ge no timing cut (cyan), Si 2D χ^2 (thick magenta), and Si no timing cut (yellow). The HT Si region is shown for comparison.




Figure 72: Primary limits from this work, along with other relevant limits. Primary results for Ge (thick blue) and Si (thick red) are shown for the HT (solid) and MT (dashed) analyses. Confidence intervals are shown for c58R Si (yellow 90%, green 68%, black cross best-fit point)[124], CRESST-II[58] (light blue 2σ , dark blue 1σ), DAMA/LIBRA[61] (light red 99%, dark red 90%), and CoGeNT[62] (brown). 90% exclusion limits are plotted for EDELWEISS[59] (cyan), XENON10 S2-only[152] (dark green), LUX[56] (magenta), CDMSlite[153] (purple), and SuperCDMS low-threshold[154] (brown). TOP: WIMP-search results up to $1000 \text{ GeV}/c^2$. BOTTOM: zoom in for low-mass results.




Figure 73: Timing distributions of calibration and WIMP-search data for the HT analysis. See text for the color scheme. Note that the horizontal axis (timing quantity) is shifted by 10 units so that it can be plotted on a log-scale. All distributions shown are survival functions.




Figure 74: Timing distributions of calibration and WIMP-search data for the MT analysis. See text for the color scheme. Note that the horizontal axis (timing quantity) is shifted by 10 units so that it can be plotted on a log-scale. All distributions shown are survival functions.




Figure 75: Timing, energy, and yield distributions for T4Z3. Yield vs. energy plots are shown on the left, while yield vs. timing are shown on the right. The top panels display ^{133}Ba data, while the bottom panels display WIMP-search data.




Figure 76: Timing, energy, and yield distributions for T5Z3. Yield vs. energy plots are shown on the left, while yield vs. timing are shown on the right. The top panels display ^{133}Ba data, while the bottom panels display WIMP-search data.




Figure 77: Results of the timing-cut scanning procedure for the HT Si timing analysis. The horizontal axis is the summed expected leakage from the ^{133}Ba fits, which serves to index the optimization of the timing cut. The different curves represent the number of observed WIMP-search candidate events (black); the ^{252}Cf neutron distribution normalized to a $10 \text{ GeV}/c^2$ WIMP (blue curve and band); the ^{133}Ba distribution normalized to WIMP-search data (red); and the expected leakage from fits (green). The vertical magenta line corresponds to the optimized timing cut.




Figure 78: Results of the timing-cut scanning procedure for the MT Si timing analysis. The horizontal axis is the summed expected leakage from the ^{133}Ba fits, which serves to index the optimization of the timing cut. The different curves represent the number of observed WIMP-search candidate events (black); the ^{252}Cf neutron distribution normalized to a $10\text{ GeV}/c^2$ WIMP (blue curve and band); the ^{133}Ba distribution normalized to WIMP-search data (red); and the expected leakage from fits (green). The vertical magenta line corresponds to the optimized timing cut.

9 Conclusions

9.1 Results of the CDMS-II Experiment

The c58 dataset was the final WIMP-search data obtained by the CDMS-II experiment. The primary HT analysis of the full 1159 kg-day Ge exposure resulted in 2 candidate events[149, 150], and set a minimum upper limit of $\sigma_{\text{SI}} < 1.96 \times 10^{-44} \text{ cm}^2$ for a WIMP mass of $64 \text{ GeV}/c^2$ (Fig. 63), very close to the anticipated sensitivity of $\sim 10^{-44} \text{ cm}^2$ [155]. The Si HT analyses, which resulted in 3 candidate events over a total exposure of 232 kg-day[149, 151, 124], set competitive limits at low WIMP masses ($< 6 \text{ GeV}/c^2$) and provided hints of a low mass WIMP signal.

The re-analysis of the c58 data proved to be an invaluable opportunity to improve the data analysis techniques of the CDMS collaboration, and to achieve improved WIMP-search reach with two target materials. The MT analyses demonstrate that surface event rejection using phonon timing information retains useful discrimination power to lower recoil energies, even though such analyses can no longer be considered background-free. At high energies, where phonon timing information is well-resolved, a χ^2 comparison technique significantly improves on the performance of the classic timing discriminator, although it does not appear to be well-suited to low-energy analysis where phonon resolutions degrade.

This re-analysis has also demonstrated the potential strengths and weaknesses of a neural network background rejection technique. The Ge neural network analysis achieved the best limits at low WIMP masses among the three HT techniques, largely due to its superior signal passage fraction. One of its empirical limitations was exposed, however, in the MT analysis, where an extremely high leakage of zero-yield events into the signal region was observed. This is instructive of the general danger of any “black box” analysis technique: namely, it may be particularly susceptible to unexpected pathologies because of its complicated internal workings and multi-layered complexity. In any case, machine learning techniques

may be useful in future CDMS analyses if such liabilities can be understood and mitigated appropriately.

Finally, the exciting hint of a WIMP signal observed in the c58R Si analysis has helped to increase interest in the possibility of a low-mass WIMP; in fact, it was the original impetus for the undertaking of the MT analyses described in this thesis. The jumbled state of affairs presently seen in the parameter space between 5 and 10 GeV/ c^2 , although frustrating to those tasked with making sense of it all, will no doubt be seen as an exciting chapter along the way to our ultimate understanding of the problem of dark matter.

9.2 Tension, Complementarity, and Resolution

The polite word for the strongly-contradictory results seen in several different direct-detection experiments (Fig. 72) is “tension”. As discussed in Sec. 8.1.1, many frameworks exist which can potentially reconcile this tension, namely EFT models, inelastic DM, and even multi-component dark matter. Regardless, more data is needed to settle the issue, and never before has the importance of target complementarity been more apparent.

If the couplings of dark matter particles to nucleons is a combination of several components (e.g. spin-dependent, velocity-dependent, etc.), or if several dark matter particles exist, then different target materials could possess different cross-sections to WIMPs. Furthermore, significant parameter space remains unexplored both at high and low WIMP masses, and due to the kinematics of nuclear recoils, few target materials have good sensitivity to both regions. Thus, the second generation (G2) of dark matter direct detection experiments will consist of a variety of target materials which have sensitivity to WIMP masses from hundreds of MeV/ c^2 to tens of TeV/ c^2 .

Furthermore, in the event of a discovery scenario, scientific rigor demands that consistent, reproducible signals be observed in several experiments. Multi-target experiments, like SuperCDMS, are particularly well-suited for this, especially if the payload deployment

scheme is flexible enough to respond to potential signals in other experiments. Finally, G2 experiments are expected to begin to suffer from irreducible backgrounds from coherent neutrino scattering[156] (CNS; see Fig. 79), which put a hard limit (often called the “neutrino floor”) on conventional direct-detection techniques.⁶² As G2 experiments map out the available WIMP parameter space, the new data should resolve the current low-mass tension, either with detection(s) or null results, and hopefully provide guidance for future dark matter searches.

9.3 The Future of CDMS

Following the end of data taking with CDMS-II in 2009, the electronic and cryogenic infrastructure of the experiment was adapted for the use of iZIP detectors for the SuperCDMS Soudan phase of the CDMS experiment. As with any low-background counting experiment, increases in sensitivity must be achieved in a two-fold manner: increasing exposure, along with mitigating backgrounds. SuperCDMS implements significant improvements in background rejection and a larger experimental payload, which translate into excellent sensitivity for future WIMP searches.

The superiority of SuperCDMS can be thought of as an improvement in the fiducial volume of the detectors. The SuperCDMS detectors, known as “interleaved ZIPs” or “iZIPs”, are about 2.5 times larger⁶³ than the CDMS-II ZIPs, which mitigates some surface effects simply through a higher ratio of volume to surface area. Additionally, full phonon and ionization readout instrumentation is patterned on both sides of each detector, increasing the position information and even allowing phonon-based fiducial volume cuts to be implemented.

The real power of the iZIP technology, and the characteristic which lends the “i” to

⁶²Although the neutrino floor presents an irreducible background to all experiments, the neutrino energy spectrum is well-known, and spectral analysis will still provide WIMP-search reach below these levels.

⁶³The current iZIP design, deployed at SuperCDMS Soudan, has the same diameter but is 2.5 times thicker than the CDMS-II ZIPs; future designs call for increasing the diameter of the detectors as well.




Figure 79: Current experimental limits (solid lines) and G2 sensitivity projections (dashed lines) for a variety of experiments. The coherent neutrino scattering floor (thick orange dashed) presents a practical limit to conventional WIMP direct-detection experiments. The SuperCDMS SNOLAB experiment, the successor to CDMS-II, has excellent and unique sensitivity to very low-mass WIMPs.

its name, is the interleaved electrode pattern on each detector face which enables “native” rejection of surface events, which are the dominant background in CDMS-II, at a level of at least 1.7×10^{-5} . [157] SuperCDMS has already produced important WIMP-search results at the Soudan laboratory [154, 153], where it is re-using the cryogenic apparatus of CDMS-II; a future upgrade at the deeper SNOLAB facility is expected to probe a large region of low-mass WIMP parameter space, and has excellent prospects for possible discovery.

The full 9 kg payload of SuperCDMS Soudan consists exclusively of Ge iZIPs, but Si iZIPs are being developed for future deployment. In addition to its considerable physics reach, SuperCDMS Soudan also serves as a proof-of-concept for the G2 SuperCDMS SNO-LAB experiment, which will feature larger payloads, lower backgrounds, and a fully re-designed cryogenics and DAQ system.

Early results from the SuperCDMS Soudan experiment have produced strong limits at low WIMP masses, excluding new parameter space which was heretofore unexplored by direct detection experiments. A low-threshold analysis [154] of the first WIMP-search dataset used a boosted decision tree (BDT) machine learning analysis tool to perform background rejection with good success, although some of the problems common to machine learning algorithms (e.g. unexpected background pathologies, “black box” behavior) were observed, as in the c58R neural network analysis. Additionally, one iZIP was operated at high bias voltage (~ 70 V), amplifying the Luke phonon signal and improving noise performance, although dual energy measurement, and thus yield discrimination, must be sacrificed. An analysis on one month of data running in this mode (known as CDMSlite, for “low ionization threshold experiment”) achieved a recoil energy threshold of 170 eV_{ee} and set strong limits below WIMP masses of 4 GeV/ c^2 [153]. Future analyses from SuperCDMS Soudan are projected to produce world-leading limits at low WIMP masses, and provide complementarity to the strong limits from LUX [56] at high WIMP masses.

SuperCDMS SNOLAB has been selected by the US Department of Energy, along with

the Xe-based LZ experiment, as a primary G2 direct detection experiment which will provide unique sensitivity to the low WIMP-mass parameter space (see Fig. 79). An initial deployment of 42 iZIPs, with a target mass of ~ 50 kg, will feature a mix of Ge and Si iZIPs, including several detectors dedicated to CDMSlite running; the cryogenics and shielding will be designed for the payload to be incrementally augmented to 400 kg. This provides flexibility to strategically adjust payloads to respond to potential signals, and the dual target materials provide excellent complementarity in a WIMP detection scenario.

The future of the dark matter community is very exciting, and SuperCDMS SNOLAB is poised to be a strong player, especially in the low WIMP-mass regime. Regardless of whether a detection occurs, fully exploring the WIMP parameter space above the CNS background floor, and the potential first direct detection of CNS itself, will be important scientific results in their own rights. With either positive or null results, we will have a more clear and complete understanding of the nature of dark matter in our Universe.

References

- [1] F. Zwicky. Republication of: The redshift of extragalactic nebulae. *General Relativity and Gravitation*, 41(1):207–224, November 2008.
- [2] Yoshiaki Sofue and Vera Rubin. Rotation curves of spiral galaxies. *Annual Review of Astronomy and Astrophysics*, 39(1):137–174, 2001.
- [3] Aaron D. Lewis, David A. Buote, and John T. Stocke. Chandra observations of a2029: The dark matter profile down to below 0.01r_{vir} in an unusually relaxed cluster. *The Astrophysical Journal*, 586(1):135, 2003.
- [4] B. J. Maughan, et al. An xmm-newton observation of the massive relaxed galaxy cluster clj1226.9+3332 at $z = 0.98$. *Mon. Not. R. Astron. Soc.*, 351:1193–1203, 2004.
- [5] J.-P. Kneib, et al. Hubble space telescope observations of the lensing cluster abell 2218. *The Astrophysical Journal*, 471(2):643, 1996.
- [6] Erin S. Sheldon, et al. Cross-correlation weak lensing of sdss galaxy clusters. iii. mass-to-light ratios. *The Astrophysical Journal*, 703(2):2232, 2009.
- [7] Douglas Clowe, et al. A direct empirical proof of the existence of dark matter. *Astrophys.J.*, 648:L109–L113, 2006.
- [8] Garry W. Angus, et al. On the proof of dark matter, the law of gravity, and the mass of neutrinos. *The Astrophysical Journal Letters*, 654(1):L13, 2007.
- [9] Marcel Zemp. The structure of cold dark matter halos: Recent insights from high resolution simulations. *Modern Physics Letters A*, 24(29):2291–2305, 2009.
- [10] Michael Boylan-Kolchin, et al. Resolving cosmic structure formation with the millenium-ii simulation. *Mon. Not. Roy. Astron. Soc.*, 398:1150–1164, 2009.
- [11] E. Hubble. A relation between distance and radial velocity among extra-galactic nebulae. *Proceedings of the National Academy of Science*, 15:168–173, March 1929.
- [12] Adam G. Riess et al. Observational evidence from supernovae for an accelerating universe and a cosmological constant. *Astron.J.*, 116:1009–1038, 1998.
- [13] S. Perlmutter, et al. Measurements of Ω and Ω_m from 42 high-redshift supernovae. *The Astrophysical Journal*, 517(2):565, 1999.
- [14] Richard H Cyburt, Brian D Fields, and Keith A Olive. An update on the big bang nucleosynthesis prediction for ^7Li : the problem worsens. *Journal of Cosmology and Astroparticle Physics*, 2008(11):012, 2008.

- [15] Daniel J. Eisenstein, et al. Detection of the baryon acoustic peak in the large-scale correlation function of sdss luminous red galaxies. *The Astrophysical Journal*, 633(2):560, 2005.
- [16] Planck Collaboration, et al. Planck 2013 results. xvi. cosmological parameters. *ArXiv e-prints*, March 2013.
- [17] G. Hinshaw, et al. Five-year wilkinson microwave anisotropy probe observations: Data processing, sky maps, and basic results. *The Astrophysical Journal Supplement Series*, 180:225–245, February 2009.
- [18] Justin Evans. The MINOS experiment: results and prospects. *Adv.High Energy Phys.*, 2013:182537, 2013.
- [19] J. Beringer, et al. Review of particle physics. *Phys. Rev. D*, 86:010001, Jul 2012.
- [20] Edward Kolb and Michael Turner. *The Early Universe*. Frontiers in Physics. Westview Press, 1994.
- [21] Benjamin W. Lee and Steven Weinberg. Cosmological lower bound on heavy-neutrino masses. *Phys. Rev. Lett.*, 39:165–168, Jul 1977.
- [22] P. Hut. Limits on masses and number of neutral weakly interacting particles. *Physics Letters B*, 69(1):85 – 88, 1977.
- [23] T. K. Hemmick, et al. Search for low- z nuclei containing massive stable particles. *Phys. Rev. D*, 41:2074–2080, Apr 1990.
- [24] Toby Falk, Keith A. Olive, and Mark Srednicki. Heavy sneutrinos as dark matter. *Physics Letters B*, 339(3):248 – 251, 1994.
- [25] Frank Daniel Steffen. Gravitino dark matter and cosmological constraints. *JCAP*, 0609:001, 2006.
- [26] Howard Baer, Vernon Barger, and Azar Mustafayev. Neutralino dark matter in msugra/cmssm with a 125 gev light higgs scalar. *JHEP*, 1205:091, 2012.
- [27] Howard Baer and Jenny List. Post-lhc8 supersymmetry benchmark points for ilc physics. *Phys. Rev. D*, 88:055004, Sep 2013.
- [28] Hsin-Chia Cheng, Jonathan L. Feng, and Konstantin T. Matchev. Kaluza-klein dark matter. *Phys.Rev.Lett.*, 89:211301, 2002.
- [29] Kimmo Kainulainen, Kimmo Tuominen, and Jussi Virkajarvi. The wimp of a minimal technicolor theory. *Phys.Rev.*, D75:085003, 2007.

- [30] Lei Wang, Jin Min Yang, and Jingya Zhu. Dark matter in the little higgs model under current experimental constraints from the lhc, planck, and xenon data. *Phys.Rev.*, D88(7):075018, 2013.
- [31] R. Foot. Mirror dark matter: Cosmology, galaxy structure and direct detection. *Int.J.Mod.Phys.*, A29:1430013, 2014.
- [32] R. D. Peccei and Helen R. Quinn. CP conservation in the presence of pseudoparticles. *Phys. Rev. Lett.*, 38:1440–1443, Jun 1977.
- [33] R.J. Creswick, et al. Theory for the direct detection of solar axions by coherent primakoff conversion in germanium detectors. *Physics Letters B*, 427(34):235 – 240, 1998.
- [34] Z. Ahmed, et al. Search for axions with the cdms experiment. *Phys. Rev. Lett.*, 103:141802, Oct 2009.
- [35] A.D. Dolgov and S.H. Hansen. Massive sterile neutrinos as warm dark matter. *Astroparticle Physics*, 16(3):339 – 344, 2002.
- [36] Alexey Boyarsky, Oleg Ruchayskiy, and Dmytro Iakubovskyi. A lower bound on the mass of dark matter particles. *JCAP*, 0903:005, 2009.
- [37] Alexey Boyarsky, et al. Where to find a dark matter sterile neutrino? *Phys.Rev.Lett.*, 97:261302, 2006.
- [38] M. Milgrom. A modification of the newtonian dynamics as a possible alternative to the hidden mass hypothesis. *Astrophys.J.*, 270:365–370, July 1983.
- [39] R.H. Sanders and E. Noordermeer. Confrontation of mond with the rotation curves of early-type disc galaxies. *Mon.Not.Roy.Astron.Soc.*, 379:702–710, 2007.
- [40] Jacob D. Bekenstein. Relativistic gravitation theory for the mond paradigm. *Phys.Rev.*, D70:083509, 2004.
- [41] I. Ferreras, et al. Confronting mond and teves with strong gravitational lensing over galactic scales: An extended survey. *Phys.Rev.D*, 86(8):083507, October 2012.
- [42] C. Skordis. Topical review: The tensor-vector-scalar theory and its cosmology. *Classical and Quantum Gravity*, 26(14):143001, July 2009.
- [43] Georges Aad et al. Search for dark matter in events with a hadronically decaying w or z boson and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ tev with the atlas detector. *Phys.Rev.Lett.*, 112(4):041802, 2014.
- [44] Serguei Chatrchyan et al. Search for dark matter and large extra dimensions in pp collisions yielding a photon and missing transverse energy. *Phys.Rev.Lett.*, 108:261803, 2012.

- [45] T. Aaltonen et al. A search for dark matter in events with one jet and missing transverse energy in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ tev. *Phys.Rev.Lett.*, 108:211804, 2012.
- [46] Nathaniel Craig. The state of supersymmetry after run i of the lhc. 2013.
- [47] Tao Han, Zhen Liu, and Shufang Su. Light neutralino dark matter: Direct/indirect detection and collider searches. 2014.
- [48] Elmo Tempel, Andi Hektor, and Martti Raidal. Fermi 130 gev gamma-ray excess and dark matter annihilation in sub-haloes and in the galactic centre. *JCAP*, 1209:032, 2012.
- [49] Tansu Daylan, et al. The characterization of the gamma-ray signal from the central milky way: A compelling case for annihilating dark matter. 2014.
- [50] IceCube collaboration, et al. Search for dark matter annihilations in the sun with the 79-string icecube detector. *ArXiv e-prints*, December 2012.
- [51] S. Adrian-Martinez et al. First results on dark matter annihilation in the sun using the antares neutrino telescope. *JCAP*, 1311:032, 2013.
- [52] O. Adriani, et al. A statistical procedure for the identification of positrons in the PAMELA experiment. *Astroparticle Physics*, 34(1):1 – 11, 2010.
- [53] M. Ackermann, et al. Measurement of separate cosmic-ray electron and positron spectra with the fermi large area telescope. *Phys. Rev. Lett.*, 108:011103, Jan 2012.
- [54] M. Aguilar, et al. First result from the alpha magnetic spectrometer on the international space station: Precision measurement of the positron fraction in primary cosmic rays of 0.5–350 gev. *Phys. Rev. Lett.*, 110:141102, Apr 2013.
- [55] Hong-Bo Jin, Yue-Liang Wu, and Yu-Feng Zhou. Implications of the first ams-02 measurement for dark matter annihilation and decay. *JCAP*, 1311:026, 2013.
- [56] S. Akerib, D. et al. First results from the lux dark matter experiment at the sanford underground research facility. *Phys. Rev. Lett.*, 112:091303, Mar 2014.
- [57] Xenon100 Collaboration, et al. The xenon100 dark matter experiment. *Astroparticle Physics*, 35:573–590, April 2012.
- [58] G. Angloher, et al. Results from 730 kgdays of the cressst-ii dark matter search. *The European Physical Journal C*, 72(4), 2012.
- [59] E. Armengaud, et al. Final results of the edelweiss-ii wimp search using a 4-kg array of cryogenic germanium detectors with interleaved electrodes. *Physics Letters B*, 702(5):329 – 335, 2011.
- [60] D. C. Malling, et al. After lux: The lz program. *ArXiv e-prints*, October 2011.

- [61] R. Bernabei, et al. New results from dama/libra. *The European Physical Journal C*, 67(1-2):39–49, 2010.
- [62] C. E. Aalseth, et al. Search for an annual modulation in a p-type point contact germanium dark matter detector. *Phys. Rev. Lett.*, 107:141301, Sep 2011.
- [63] E. Daw, et al. The drift dark matter experiments. *ArXiv e-prints*, October 2011.
- [64] J.P. Lopez, et al. Background rejection in the dmtpc dark matter search using charge signals. *Nucl.Instrum.Meth.*, A696:121–128, 2012.
- [65] D. S. Akerib, et al. Exclusion limits on the wimp-nucleon cross section from the first run of the cryogenic dark matter search in the soudan underground laboratory. *Phys. Rev. D*, 72:052009, Sep 2005.
- [66] L. Novak. Ge boule quotes. Ebook note.
- [67] P. Brink. Collaboration email.
- [68] Tarek Saab. *Search for Weakly Interacting Massive Particles with the Cryogenic Dark Matter Search Experiment*. PhD thesis, Stanford University, August 2002.
- [69] S.W. Leman, et al. Phonon quasidiffusion in cryogenic dark matter search large germanium detectors. *Chinese Journal of Physics*, 49:349, February 2011.
- [70] P.N. Luke. Voltage-assisted calorimetric ionization detector. *Journal of Applied Physics*, 64:6858, December 1988.
- [71] K.D. Irwin, et al. A self-biasing cryogenic particle detector utilizing electrothermal feedback and a squid readout. *IEEE Transactions on Applied Superconductivity*, 5:2690, June 1995.
- [72] A. Reisetter. *Results from the Two-Tower Run of the Cryogenic Dark Matter Search*. PhD thesis, University of Minnesota, August 2005.
- [73] M. Fritts. *Background Characterization and Discrimination in the Final Analysis of the CDMS II Phase of the Cryogenic Dark Matter Search*. PhD thesis, University of Minnesota, February 2011.
- [74] CDMS Collaboration. Nuclear-recoil energy scale in cdms ii silicon dark-matter detectors. *Nucl. Inst. and Methods*, to appear.
- [75] D.A. Bauer, et al. The CDMS II data acquisition system. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 638(1):127 – 133, 2011.

- [76] I. Antcheva, et al. Root a c++ framework for petabyte data storage, statistical analysis and visualization. *Computer Physics Communications*, 180(12):2499 – 2512, 2009. 40 YEARS OF CPC: A celebratory issue focused on quality software for high performance, grid and novel computing architectures.
- [77] Sunil Golwala. *Exclusion Limits on the WIMP-Nucleon Elastic-Scattering Cross Section from the Cryogenic Dark Matter Search*. PhD thesis, University of California at Berkeley, Fall 2000.
- [78] J. Filippini. Optimal filtering. Ebook note.
- [79] R.W. Ogburn. *A Search For Particle Dark Matter Using Cryogenic Germanium and Silicon Detectors in the One- and Two-Tower Runs of CDMS-II at Soudan*. PhD thesis, Stanford University, April 2008.
- [80] Scott Hertel. Adjustments to the rtftwalk routine for c58. Ebook note.
- [81] Adam J. Anderson. R133 charge calibration v2. Ebook note.
- [82] K. McCarthy. Comparison of individual vs overall calibration of runs 125-128. Ebook note.
- [83] M. Pyle. Possible explanations of energy dependence seen in timing parameters. Ebook note.
- [84] M. Pyle. Limiting energy dependence in corrected quantities. Ebook note.
- [85] B. Serfass and M. Pyle. Preliminary comparison of yield resolution between pipefitter and darkpipe for t1z3. Ebook note.
- [86] The CDMS II Collaboration. Dark matter search results from the cdms ii experiment. *Science*, 327(5973):1619–1621, 2010.
- [87] The CDMS II Collaboration. Improved wimp-search reach of the cdms ii germanium data. In preparation.
- [88] Jeffrey P. Filippini. *A Search for WIMP Dark Matter Using the First Five-Tower Run of the Cryogenic Dark Matter Search*. PhD thesis, University of California, Berkeley, Fall 2008.
- [89] D. Barker and D.-M. Mei. Germanium detector response to nuclear recoils in searching for dark matter. *Astroparticle Physics*, 38(0):1 – 6, 2012.
- [90] Z. Ahmed et al. Understanding the energy scale for nuclear recoils in the cdms dark matter detector. *Nuclear Instruments and Methods*, In preparation.
- [91] J. Lindhard, M. Scharff, and H. E. Schiøtt. Range concepts and heavy ion ranges. *Mat. Fys. Medd. Dan. Vid. Selsk.*, 33, 1963.

- [92] J. Lindhard, et al. Integral equations governing radiation effects. *Mat. Fys. Medd. Dan. Vid. Selsk.*, 33, 1963.
- [93] Kevin McCarthy. Stable tuning regions in run 125-128. Ebook note.
- [94] Scott Hertel, et al. c58 post-run ks test (wimp-search data). Ebook note.
- [95] D. Moore. Time intervals with high charge noise in r125-128. Ebook note.
- [96] Joseph Manungu. Trigger rate analysis for the background data run 125 to 128. Ebook note.
- [97] M. Kos. He film studies. Ebook note.
- [98] Thomas Hofer. Analysis cut efficiencies for c58r. Ebook note.
- [99] D. Moore. cpstd and cqstd cuts for c58. Ebook note.
- [100] Joel Sander. Negative phonon events. Ebook note.
- [101] Kevin McCarthy. Check of cgoodstarttime for run 125-128. Ebook note.
- [102] J. Zhang. Charge χ^2 cut and efficiencies for the reprocessed run 125-128 data. Ebook note.
- [103] Danielle Speller. Definition and efficiencies of cqin.c58r.m. Ebook note.
- [104] Jeter Hall. c58 nuclear recoil bands: v1. Ebook note.
- [105] Tobias Bruch. Definition of the ER bands and yield distributions for mean and gradient corrected data. Ebook note.
- [106] Tobias Bruch. Definition of cbelower. Ebook note.
- [107] Danielle Speller. Charge thresholds: cqthresh.c58r.m. Ebook note.
- [108] Joseph M. Kiveni. Charge noise distribution in qiofvols and qoofvols for t1z1. Ebook note.
- [109] Danielle Speller. csingle.c58r definition and efficiencies. Ebook note.
- [110] J. D. Lewin and P. F. Smith. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. *Astroparticle Physics*, 6:87–112, 1996.
- [111] Joseph M. Kiveni. Optimizing 5d chisq for a 8 gev wimp mass. Ebook note.
- [112] Kevin A. McCarthy. *Detector Simulation and WIMP Search Analysis for the Cryogenic Dark Matter Search Experiment*. PhD thesis, Massachusetts Institute of Technology, June 2013.

- [113] Kevin McCarthy. Leakage considerations in c58r si low-background. Ebook note.
- [114] Jianjie Zhang. *A Dark Matter Search Using the Final CDMS-II Data and 100 mm SuperCDMS Germanium Detector Ionization Test*. PhD thesis, University of Minnesota, July 2014.
- [115] Jianjie Zhang. C58 ge reanalysis timing cuts, wimp-search sideband leakages, and unblinding results. Ebook note.
- [116] Joseph M. Kiveni. *A Search for WIMP Dark Matter using an Optimized Chi-square Technique on the Final Data from the Cryogenic Dark Matter Search Experiment (CDMS II)*. PhD thesis, University of Minnesota, December 2012.
- [117] Joseph M. Kiveni. c58r timing cuts analysis using a 5d optimized energy dependent chi square. Ebook note.
- [118] Xinjie Qiu. *Advanced Analysis and Background Techniques for the Cryogenic Dark Matter Search*. PhD thesis, University of Minnesota, December 2009.
- [119] Christopher M. Bishop. *Neural Networks for Pattern Recognition*. Oxford University Press, Inc., New York, NY, USA, 1995.
- [120] Martin Fodsløtte Møller. A scaled conjugate gradient algorithm for fast supervised learning. *Neural Networks*, 6(4):525 – 533, 1993.
- [121] Karl Pearson. Liii. on lines and planes of closest fit to systems of points in space. *Philosophical Magazine Series 6*, 2(11):559–572, 1901.
- [122] I. Nabney. *NETLAB: Algorithms for Pattern Recognition*. Advances in Computer Vision and Pattern Recognition. Springer, 2002.
- [123] Kevin McCarthy. Timing χ^2 cut for c58r si analysis - cross-checks and leakage estimates from open ws data. Ebook note.
- [124] R. Agnese, et al. Silicon detector dark matter results from the final exposure of cdms ii. *Phys. Rev. Lett.*, 111:251301, Dec 2013.
- [125] Jo Bovy and Scott Tremaine. On the local dark matter density. *The Astrophysical Journal*, 756(1):89, 2012.
- [126] Martin C. Smith, et al. The rave survey: constraining the local galactic escape speed. *Monthly Notices of the Royal Astronomical Society*, 379(2):755–772, 2007.
- [127] Joseph M. Kiveni. Singularity removal in low threshold 5d χ^2 timing for $m_\chi = 8 \text{ geV}/c^2$. Ebook note.
- [128] K. McCarthy. c58r si analysis - post-unblinding events and checks. Ebook note.

- [129] Jianjie Zhang. c58 ge reanalysis wimp candidates and near-miss events. Ebook note.
- [130] Mark Pepin. c58r neutron background estimation. Ebook note.
- [131] S. Agostinelli, et al. Geant4a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250 – 303, 2003.
- [132] Mark Pepin. c58r background estimation. Ebook note.
- [133] V.A. Kudryavtsev. Muon simulation codes music and musun for underground physics. *Computer Physics Communications*, 180(3):339 – 346, 2009.
- [134] Jonathan L. Thron. The Soudan-II Proton Decay Experiment. *Nucl.Instrum.Meth.*, A283:642–645, 1989.
- [135] Thomas Hofer. Updated cosmogenic background estimates for c58r ge and si mt and ht analyses. Ebook note.
- [136] K. McCarthy. c58r alpha/lead recoil studies. Ebook note.
- [137] J. Engel. Nuclear form factors for the scattering of weakly interacting massive particles. *Physics Letters B*, 264:114–119, July 1991.
- [138] Nikhil Anand, A. Liam Fitzpatrick, and W.C. Haxton. Model-independent analyses of dark-matter particle interactions. 2014.
- [139] Riccardo Catena and Paolo Gondolo. Global fits of the dark matter-nucleon effective interactions. 2014.
- [140] David Smith and Neal Weiner. Inelastic dark matter. *Phys. Rev. D*, 64:043502, Jul 2001.
- [141] Spencer Chang, Aaron Pierce, and Neal Weiner. Using the energy spectrum at dama/libra to probe light dark matter. *Phys.Rev.*, D79:115011, 2009.
- [142] Lawrence M. Widrow, Brent Pym, and John Dubinski. Dynamical blueprints for galaxies. *Astrophys.J.*, 679:1239–1259, 2008.
- [143] F. J. Kerr and D. Lynden-Bell. Review of galactic constants. *Mon. Not. Roy. Astron. Soc.*, 221:1023–1038, 1986.
- [144] Walter Dehnen and James Binney. Local stellar kinematics from hipparcos data. *Mon. Not. Roy. Astron. Soc.*, 298:387–394, 1998.
- [145] Martin C. Smith, et al. The rave survey: Constraining the local galactic escape speed. *Mon. Not. Roy. Astron. Soc.*, 379:755–772, 2007.

- [146] Graciela Gelmini and Paolo Gondolo. Wimp annual modulation with opposite phase in late-infall halo models. *Phys.Rev.*, D64:023504, 2001.
- [147] Gary J. Feldman and Robert D. Cousins. Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, Apr 1998.
- [148] S. Yellin. Finding an upper limit in the presence of an unknown background. *Phys. Rev. D*, 66:032005, Aug 2002.
- [149] D. S. Akerib, et al. Limits on spin-independent interactions of weakly interacting massive particles with nucleons from the two-tower run of the cryogenic dark matter search. *Phys. Rev. Lett.*, 96:011302, Jan 2006.
- [150] Z. Ahmed, et al. Search for weakly interacting massive particles with the first five-tower data from the cryogenic dark matter search at the soudan underground laboratory. *Phys. Rev. Lett.*, 102:011301, Jan 2009.
- [151] R. Agnese, et al. Silicon detector results from the first five-tower run of cdms ii. *Phys. Rev. D*, 88:031104, Aug 2013.
- [152] J. Angle, et al. Search for light dark matter in xenon10 data. *Phys. Rev. Lett.*, 107:051301, Jul 2011.
- [153] R. Agnese, et al. Search for low-mass weakly interacting massive particles using voltage-assisted calorimetric ionization detection in the supercdms experiment. *Phys. Rev. Lett.*, 112:041302, Jan 2014.
- [154] R. Agnese, et al. Search for low-mass weakly interacting massive particles with supercdms. *Phys. Rev. Lett.*, 112:241302, Jun 2014.
- [155] D. Abrams, et al. Present status of the cryogenic dark matter search (cdms ii) experiment. *Physica B Condensed Matter*, 329:1590–1591, May 2003.
- [156] J. Billard, L. Strigari, and E. Figueroa-Feliciano. Implication of neutrino backgrounds on the reach of next generation dark matter direct detection experiments. *Phys.Rev.*, D89:023524, 2014.
- [157] R. Agnese, et al. Demonstration of surface electron rejection with interleaved germanium detectors for dark matter searches. *Applied Physics Letters*, 103(16):–, 2013.

A Unmasked WIMP-search Data

This Appendix displays all of the unmasked WIMP-search data for the Ge and Si MT analyses and, by extension, for the HT analyses as well. The data are displayed graphically, with each figure showing the data for a particular detector and timing analysis. The color coding for each figure is as follows.

- **Green bands** indicate the $\pm 3\sigma$ edges of the ER band (Sec. 4.5); the widths of the bands indicate the run-to-run variation over the c58 data-taking period. The lower edge of this band defines the below ER band cut.
- **Magenta bands** indicate the $\pm 2\sigma$ edges of the NR band (Sec. 4.5); the widths of the bands indicate the run-to-run variation over the c58 data-taking period.
- **Light red bands** indicate the $\mu + 3\sigma$ ionization threshold (Sec. 4.5); the widths of the bands indicate the run-to-run variation over the c58 data-taking period. The ionization threshold was $\mu + 3\sigma$ for the MT analyses, but $\mu + 4.5\sigma$ for the HT analyses.
- **Cyan bands** indicate the recoil energy thresholds for the MT analyses (Sec. 6.1.1); the widths of the bands indicate the run-to-run variation over the c58 data-taking period. For MT analyses, the thresholds were generally 10 keV for Ge and 7 keV for Si.
- **Blue points** indicate events which pass all basic event pre-selection cuts (Sec. 4.3) and the single-scatter condition (Sec. 4.6).
- **Red points** indicate events which pass all basic even pre-selection cuts, the single-scatter condition, and are within the HT signal region (Sec. 4.5).
- **Dark green points** indicate events which pass basic event pre-selection cuts, the single-scatter condition, and are within the MT signal region (Sec. 6.1) but are not within the HT signal region.

- **Brown points** indicate events which pass all the same conditions as the dark green points, but fail the MT recoil energy thresholds (see Sec. 7.1.2).
- **Black circles** around a point indicate that the event passed the timing cut.

A.1 Ge 5D χ^2 Analysis

Figs. 80 - 91 display the unmasked WIMP-search data for the Ge 5D χ^2 timing analysis (Sec. 5.3.2).




Figure 80: T1Z2 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 81: T1Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 82: T2Z3 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 83: T2Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 84: T3Z2 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 85: T3Z4 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 86: T3Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 87: T4Z2 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 88: T4Z4 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 89: T4Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 90: T5Z4 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.




Figure 91: T5Z5 WIMP-search data using the Ge 5D χ^2 MT timing analysis. No candidates are observed.

A.2 Ge 5D χ^2 Re-optimized Analysis

Figs. 92 - 103 display the unmasked WIMP-search data for the Ge 5D χ^2 re-optimized timing analysis (Sec. 6.1.2).




Figure 92: T1Z2 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 93: T1Z5 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 94: T2Z3 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. 1 candidate (G2) is observed.




Figure 95: T2Z5 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 96: T3Z4 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 97: T3Z6 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. 1 candidate (G19) is observed.




Figure 98: T4Z2 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 99: T4Z4 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 100: T4Z5 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 101: T4Z6 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 102: T5Z4 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.




Figure 103: T5Z5 WIMP-search data using the Ge 5D χ^2 re-optimized MT timing analysis. No candidates are observed.

A.3 Ge Classic Analysis

Figs. 104 - 115 display the unmasked WIMP-search data for the Ge classic timing analysis (Sec. 5.3.1).




Figure 104: T1Z2 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.




Figure 105: T1Z5 WIMP-search data using the Ge classic MT timing analysis. 2 candidates (G1 and G3) are observed.




Figure 106: T2Z3 WIMP-search data using the Ge classic MT timing analysis. 1 candidate (G2) is observed.




Figure 107: T2Z5 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.




Figure 108: T3Z2 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.




Figure 109: T3Z4 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.




Figure 110: T3Z5 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.




Figure 111: T4Z2 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.




Figure 112: T4Z4 WIMP-search data using the Ge classic MT timing analysis. 1 candidate (G4) is observed.




Figure 113: T4Z5 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.




Figure 114: T5Z4 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.




Figure 115: T5Z5 WIMP-search data using the Ge classic MT timing analysis. No candidates are observed.

A.4 Ge Neural Network Analysis

Figs. 116 - 127 display the unmasked WIMP-search data for the Ge neural network timing analysis (Sec. 5.3.3).




Figure 116: T1Z2 WIMP-search data using the Ge neural network MT timing analysis. 1 candidate (G11) is observed.




Figure 117: T1Z5 WIMP-search data using the Ge neural network MT timing analysis. 5 candidates (G1, G3, G12, G13, and G14) are observed.




Figure 118: T2Z3 WIMP-search data using the Ge neural network MT timing analysis. No candidates are observed.




Figure 119: T2Z5 WIMP-search data using the Ge neural network MT timing analysis. 1 candidate (G15) is observed.




Figure 120: T3Z2 WIMP-search data using the Ge neural network MT timing analysis. 3 candidates (G7, G8, and G9) are observed.




Figure 121: T3Z4 WIMP-search data using the Ge neural network MT timing analysis. 2 candidates (G10 and G16) are observed.




Figure 122: T3Z5 WIMP-search data using the Ge neural network MT timing analysis. No candidates are observed.




Figure 123: T4Z2 WIMP-search data using the Ge neural network MT timing analysis. 1 candidate (G17) is observed.




Figure 124: T4Z4 WIMP-search data using the Ge neural network MT timing analysis. 2 candidates (G4 and G18) are observed.




Figure 125: T4Z5 WIMP-search data using the Ge neural network MT timing analysis. 1 candidate (G6) is observed.




Figure 126: T5Z4 WIMP-search data using the Ge neural network MT timing analysis. No candidates are observed.




Figure 127: T5Z5 WIMP-search data using the Ge neural network MT timing analysis. No candidates are observed.

A.5 Si 5D χ^2 Analysis

Figs. 128 - 136 display the unmasked WIMP-search data for the Si 2D χ^2 timing analysis (Sec. 5.4).




Figure 128: T1Z4 WIMP-search data using the Si 2D χ^2 MT timing analysis. No candidates are observed.




Figure 129: T2Z1 WIMP-search data using the Si 2D χ^2 MT timing analysis. No candidates are observed.




Figure 130: T2Z2 WIMP-search data using the Si 2D χ^2 MT timing analysis. 1 candidate (S4) is observed.




Figure 131: T2Z4 WIMP-search data using the Si 2D χ^2 MT timing analysis. No candidates are observed.




Figure 132: T3Z3 WIMP-search data using the Si 2D χ^2 MT timing analysis. No candidates are observed.




Figure 133: T4Z1 WIMP-search data using the Si 2D χ^2 MT timing analysis. No candidates are observed.




Figure 134: T4Z3 WIMP-search data using the Si 2D χ^2 MT timing analysis. 2 candidates (S1 and S2) are observed.




Figure 135: T4Z3 WIMP-search data using the Si 2D χ^2 MT timing analysis, with R125 removed. 2 candidates (S1 and S2) are observed.




Figure 136: T5Z3 WIMP-search data using the Si 2D χ^2 MT timing analysis. 1 candidate (S3) is observed.