

Results from the Two-Tower Run of the Cryogenic Dark Matter Search

The Cryogenic Dark Matter Search has completed two runs at the Soudan Underground Laboratory. In the second, two towers of detectors were operated from March to August 2004. CDMS used Ge and Si ZIP (Z-sensitive, Ionization, and Phonon) detectors, operated at 50mK, to look for Weakly Interacting Massive Particles (WIMPs) which may make up most of the dark matter in our universe. These detectors are surrounded by lead and polyethylene shielding as well as an active muon veto. These shields, as well as the overburden of Soudan rock, provide a low background environment for the detectors.

The ZIP detectors record the ratio of ionization signal to phonon signal to discriminate between nuclear recoils, characteristic of WIMPs and neutrons, and electron recoils, characteristic of gamma and beta backgrounds. They also provide timing information from the four phonon channels that is used to reject surface events, for which ionization collection is poor. A blind analysis, defined using calibration data taken in situ throughout the run, provides a definition of the WIMP signal region by rejecting backgrounds. This analysis applied to the WIMP search data gives a limit on the spin independent WIMP-nucleon cross-section that is an order of magnitude lower than any other experiment has published.

UNIVERSITY OF MINNESOTA

This is to certify that I have examined this bound copy of a doctoral thesis by

Angela Jean Reisetter

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.

Professor Priscilla Cushman
(Faculty Adviser)

GRADUATE SCHOOL

**Results from the Two-Tower Run of the Cryogenic Dark
Matter Search**

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

Angela Jean Reisetter

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

August, 2005

Results from the Two-Tower Run of the Cryogenic Dark Matter Search

by Angela Jean Reisetter

Under the supervision of Professor Priscilla Cushman

ABSTRACT

The Cryogenic Dark Matter Search has completed two runs at the Soudan Underground Laboratory. In the second, two towers of detectors were operated from March to August 2004. CDMS used Ge and Si ZIP (Z-sensitive, Ionization, and Phonon) detectors, operated at 50mK, to look for Weakly Interacting Massive Particles (WIMPs) which may make up most of the dark matter in our universe. These detectors are surrounded by lead and polyethylene shielding as well as an active muon veto. These shields, as well as the overburden of Soudan rock, provide a low background environment for the detectors.

The ZIP detectors record the ratio of ionization signal to phonon signal to discriminate between nuclear recoils, characteristic of WIMPs and neutrons, and electron recoils, characteristic of gamma and beta backgrounds. They also provide timing information from the four phonon channels that is used to reject surface events, for which ionization collection is poor. A blind analysis, defined using calibration data taken in situ throughout the run, provides a definition of the WIMP signal region by rejecting backgrounds. This analysis applied to the WIMP search data gives a limit on the spin independent WIMP-nucleon cross-section that is an order of magnitude lower than any other experiment has published.

Acknowledgements

The idea of having some space here to say whatever I like is tantalizing, but I'll go ahead and do the conventional thing and thank those who helped me through this whole process known as grad school. My advisor, Prisca Cushman, and the CDMS postdoc here, Long Duong, generously supported me in every way throughout this process, from my first summer on, and I owe a great deal to both of them. I've learned most of what I know about being an experimentalist and a member of a collaboration from these two. My husband, Benjamin Skroch, was also indispensable. Without his help, I surely would've managed by now to precipitate some disaster in my living quarters through inattention as well as tear all my hair out and grind my teeth down to their roots. For all that support, I will gladly put up with that all-to-often asked question "So, ya find any dark matter today?" for a good long while to come. Maybe someday the answer will be "yes".

In addition to these three who I saw nearly every day I was in town, I learned a great deal from other members of the CDMS collaboration. Richard Schnee taught me how to carry out analysis of the CDMS data, and I picked up all kinds of tidbits about physics, life, and dilution fridges from those I shared shifts with at Soudan. I won't pretend to be all-inclusive here, but these folks included Vuk Mandic, Clarence Chang, Blas Cabrera, Bernard Sadoulet, Nader Mirabolfathi, Walter Ogburn, Jim Beaty, Dan Bauer, Betty Young, Paul Brink, Bruno Serfass, and Dan Akerib, in no particular order whatsoever. And then there's Joel Sander, who feels my PRL pain.

I also want to say that the CDMS collaboration as a whole is a great collaboration to work with, in a field that is not necessarily blessed with an excess of such functional groups. The members of the collaboration work hard to maintain the open atmosphere that pervades our environments, even when that means too many teleconferences.

And then, because I promised Erin I would, I would like to say that my cousins made it all possible for me. Actually, I think I promised I would say that when I got famous, but I'm guessing this will be the only real opportunity I have to say anything of the kind.

Seriously, though, I've been lucky enough to have a great deal of support in and outside of the astrophysics field, and my extended family is a great example of that. My grandfather, Earl Clare Ralya, who passed away while I was on shift at Soudan in Oct 2003, made me promise not to quit in the middle of this thing, with his quiet insistence on doing it right and getting it done. My maternal grandmother, Naomi Reisetter, was the one who dismissed my initial reservations

about going to grad school with such disgust that I was no longer willing to entertain them. My parents, who moved here to the Twin Cities when I was halfway through, have helped me in every way they could, despite the fact that the very word “cosmogenic” makes my father snicker and reach for his Spaceman Spiff shades. My sister has helped out with all kinds of activities while insisting that I goof off more often, which is very good for me. Visiting younger cousins have made it justifiable that I go to too many productions of the “Nutcracker”, visit the zoo, and get into the Children’s Museum without getting funny looks. Chatting with family always restores my sense of perspective which occasionally gets lost in staring at computer monitors for too long.

In a ridiculously grand and impossibly general gesture, I want to thank my friends and acquaintances here in the Twin Cities. This is probably the most livable place I’ve ever lived in, and I’ve met all kinds of great people. Many of them are members of the Baha’i community here, who are always interested in my work, even when I’m not. We’ve studied together and worshiped together and celebrated together, sometimes for no particular reason. There are too many people in this category to start naming them, so I’ll keep my general gesture general. Some of these acquaintances are physicists, professors and fellow graduate students, and many are musicians, who met Benjamin before they met me, but befriended me anyway, generously enough. In short, there are many wonderful, intellectually curious people that have comforted, supported, and just befriended me throughout the process, and I’m grateful.

Table of Contents

Abstract	i
Acknowledgements	ii
List of Tables	vi
List of Figures	vii
1 Introduction to the Search for Dark Matter	1
1.1 Evolution of the Standard Cosmological Model	1
1.2 Theoretical Dark Matter Candidates	7
1.3 Experimental searches for WIMPs	10
1.3.1 Indirect WIMP Searches	10
1.3.2 Direct WIMP Searches	10
1.3.3 Limit Comparisons	13
2 CDMS II Detectors and Experimental Setup	15
2.1 Detector Description	15
2.1.1 Ionization Measurement	16
2.1.2 Phonon Measurement	19
2.2 Tower Construction	23
2.3 Warm Electronics and DAQ	26
2.4 Experimental Setup and Shielding	28
2.5 Cryogenics System	29
3 Data Processing in Run 119	33
3.1 Run 119 Live Time	34
3.2 Data Monitoring	35
3.3 Event Reconstruction	38
3.4 Corrections for ZIP Non-uniformities	39
3.5 ZIP Energy and Position Resolutions	41

4	Analysis: Backgrounds and Rejection	46
4.1	Gammas and Band Definitions	47
4.2	Neutrons and the Veto	53
4.2.1	Veto data analysis	54
4.2.2	Neutron background simulations	55
4.3	Betas and Surface Event Rejection	57
4.3.1	Defining Analysis Cuts	58
4.3.2	Beta Sources	66
5	Analysis: Methods and Results	70
5.1	Event Selection and Signal Region Definition	70
5.1.1	Event Selection	70
5.1.2	Fiducial Volume Definition	72
5.1.3	Signal Region Definition	74
5.2	Efficiencies and Expected Backgrounds	77
5.2.1	Overall Analysis Efficiencies	77
5.2.2	Expected Background Calculation	80
5.3	Results	82
5.3.1	Unmasking	82
5.3.2	WIMP Search Limits	86
5.4	Future Prospects	87
	Appendix A. Characterization of Surface Contamination	92
A.1	Description of Tests	92
A.1.1	Results	95
	Appendix B. Details of Cosmogenics Simulations	99
	References	105

List of Tables

4.1	Fraction of betas passing the timing cut in three data sets: first half of Ba data (half used to define the cuts), second half of Ba data (half used to test cuts), and WIMP search multiples. A beta by definition has an ionization yield between 0.1 and the bottom of the 5 sigma electron recoil band.	65
4.2	Alpha rates, in events/cm/day for runs 118 and 119.	68
5.1	Events passing all analysis cuts but surface event rejection (timing) cuts and expected backgrounds, by detector. The two different beta populations used for the two methods of background calculation are described in the text.	81
5.2	Expected beta background and statistical error in Ge and Si using all WIMP-search beta multiples, only nuclear recoil band WIMP-search beta multiples, and the weighted average of the two.	82

List of Figures

1.1	Rotational curve of a typical galaxy. The dashed line is the expected Keplerian rotational curve calculated from the visible disk matter. The dotted line is the contribution of gas in the galaxy. The dot-dashed line is the calculated contribution from a spherical halo of dark matter. Figure taken from [3].	2
1.2	Superimposed rotational curves from several different galaxies of different types, and a zoomed in look at lower radii. The flat rotational curve at large radii is a shared feature. Figure taken from [5].	3
1.3	WIMP-nucleon cross section vs WIMP mass for several different recently published theories.	9
1.4	Recoil energy distributions for 3 typical WIMP masses. Figure taken from [11]. . .	11
1.5	Spin independent upper limits on cross sections, including the allowed region from DAMA, superimposed on the same theoretical parameter space as in Fig. 1.3, zoomed in on the experimental limits.	14
2.1	Ionization sensors on the bottom face of a ZIP. The inner electrode is considered the fiducial volume.	16
2.2	Schematic of ionization circuit. From [3].	17
2.3	A typical ionization channel pulse. This event took place in T1Z3, and is a 50 keV electron recoil event. Like all events selected as good events, this event deposits nearly all ionization energy in the inner electrode.	17
2.4	Phonon sensors on top face of a ZIP.	20
2.5	Side view of layers deposited during fabrication of the phonon sensors.	21
2.6	Schematic of the phonon detection circuit.	23
2.7	A typical phonon pulse, all four channels. This is the same event as in Fig. 2.3. It's located in quadrant B, which is the trace with the largest amplitude.	24
2.8	A side view diagram of a CDMS tower mounted in the icebox at Soudan. From [25].	25
2.9	CDMS II data acquisition system, warm electronics.	27
2.10	CDMS setup in the Soudan Underground Laboratory. From [25].	28
2.11	CDMS shielding surrounding the icebox, side view and top view. From [3]	30
3.1	Configuration of detectors in the two towers.	33

3.2	Run 119 accumulated live time vs. calendar days.	34
3.3	A typical noise spectrum for run 119, T1Z3.	36
3.4	Online monitoring example for Tower 1, printed at the end of a Ba calibration run (140413_0856). Plots are, left to right, histogram of summed ionization energy, ionization vs phonon energy, the same plot zoomed in, 2 position plots, and ionization/phonon vs y-position.	37
3.5	Phonon energy vs ionization energy (left) before and (right) after phonon position correction, T2Z1.	40
3.6	Charge energy vs y position for (upper)T2Z3 and (lower) T2Z5. Plots on the left are of the uncorrected quantity before correction, in Volts, and on the right are of the corrected quantity after correction and scaling to keV. The smooth variation in T2Z3 is typical of detectors. The unusual region in T2Z5 at $y < -20$, marked with a vertical line, is excluded from further analysis.	43
3.7	Spectra of ionization energy in all 6 Ge detectors of run 119. These lines are not visible in Si detectors. Best resolution is in T1Z5 and T2Z3. T2Z5 has poor resolution due to the area of bad ionization collection.	44
3.8	Position plots for detector T2Z3. Position using timing information is left, using energy distribution is right.	45
4.1	Definition of bands in detector T2Z3. The lower band is the nuclear recoil (NR) band and the upper one is the electron recoil (ER) band. Cf data is used to define the NR band with a $\pm 2\text{-}\sigma$ width around the mean, and Ba data is similarly used to define the $2\text{-}\sigma$ ER band.	48
4.2	Surface event rate (events/day) vs dataset index for 5 summed Ge detectors (T1Z2, T1Z3, T1Z5, T2Z3, and T2Z5). The four ^{252}Cf calibrations are shown as black dotted lines. The rise in rates near index 70 corresponds to a decrease in old air purge. Dataset numbers correspond to late March 26 at 0 and August 8 at 128. The dates of the calibrations are May 7, June 15, July 12, and August 2.	50
4.3	(a)A comparison of nuclear recoil band means in the 5 good Ge detectors with Lindhard theory. (b) A comparison of the nuclear recoil band means in the 4 good Si detectors.	51
4.4	(a)Nuclear recoil band efficiencies from neutrons in Cf and (b) Electron recoil band efficiencies from gammas in Ba, shown in the 5 good Ge detectors. Error bars are added for the representative detector T2Z3.	52
4.5	A typical veto scintillator spectrum. The peak just above threshold is ambient gammas; the second, wider peak is due to muons. The excess at extreme right is a saturation effect in the digitizer.	55
4.6	Log-log neutron spectrum entering the cavern, as obtained from Geant4 simulation. The energies of 19,577 neutrons are plotted; average energy is 45.33 MeV.	56

4.7	Pictorial definition of phonon delay and phonon risetime. The y-axis is the fraction of the maximum of the phonon pulse (.1=10%). The faster-rising pulse is from the ionization channel; larger pulse is from phonon quadrant B.	58
4.8	Plot of phonon delay vs. ionization yield in Tower 1. Dark dots are gammas from Ba, x's are betas from Ba, and grey dots are neutrons from Cf.	59
4.9	Plot of phonon delay vs. ionization yield in Tower 2. Dark dots are gammas from Ba, x's are betas from Ba, and grey dots are neutrons from Cf.	60
4.10	Phonon delay (μs) vs. recoil energy (keV) for towers 1 and 2. All data is from Ba calibrations. Dots are gammas, x's are betas. Note that the mean of the distribution does not vary with energy, but the width does.	62
4.11	Plot of energy adjusted phonon risetime (μs) vs. phonon delay (μs) in Tower 1. Dots are neutrons from Cf, and x's are betas from Ba. Location of timing cuts are shown in black solid lines. Events that pass the cuts are in the upper right region where most neutrons lie.	63
4.12	Plot of energy-adjusted phonon risetime (μs) vs. phonon delay (μs) in Tower 2. Dots are neutrons from Cf, and x's are betas from Ba. Location of timing cuts are shown in black solid lines. Events that pass the cuts are in the upper right region where most neutrons lie.	64
4.13	Ionization energy (keV) vs integrated phonon energy (keV) for T1Z3, using WIMP-search data. Alphas, marked with xs, are high-energy, low-yield events, well separated from the high-energy gammas in the electron recoil band (marked by dots). Alphas are produced by decays in the Rn chain. At these high energies, phonon energies are obtained by integrating the traces, rather than with the standard optimal filter, and ionization energies are obtained using a five-parameter fit if the digitizers saturate. Solid lines indicate the cuts used to select alphas, and the dashed line labels the 5.3 MeV endpoint for the alphas.	67
4.14	Ratio of alpha rate in detectors to beta rate in detectors, as a function of detector. Detector numbers run 1 to 12, with detector 7 corresponding to T2Z1.	69
5.1	KS test significances for simulated data sets with (a) 0% variability and (b) 15% variability in the tested data set.	71
5.2	Outer electrode energy vs inner electrode energy for (a) T1 and (b)T2. All axes are in keV. The inner electrode cuts are shown as black dashed lines.	73
5.3	Cartoon of ionization electrodes. There is a 1 mm gap between the inner and outer electrodes, and .8mm of uncovered surface at the detector's edge.	74

5.4	Illustration of different ionization collection populations. Events below the dashed line (an approximation) are inner electrode events, used in the analysis. On the diagonal are events shared by both electrodes. Circled gap events, also called funnel events for the shape of the distribution, have suppressed ionization collection and occur in the gap between the two electrodes. The curl events have negative pulses in the inner electrode, most commonly caused by cross talk in the electronics between the outer and inner channels. The tails of these two circles distributions have a small probability of leaking into the inner electrode population and must be considered in the inner electrode cut efficiency.	75
5.5	Inner electrode selection efficiency (fraction of total population) vs recoil energy (keV) for all detectors. Results from two methods of calculation are illustrated; they are consistent with one another in every case.. Data are from [54] and [53]. .	76
5.6	WIMP recoil efficiencies of analysis cuts, applied progressively. Efficiencies are averaged over all Ge or Si detectors. See text for explanation of the cuts as they are applied.	78
5.7	WIMP recoil efficiencies of primary analysis, including all cuts. Efficiencies are averaged over detectors. Behavior at low energies is dominated by timing cut efficiencies, and behavior at high energies is dominated by inner electrode cut efficiencies, which are worse in Ge than in Si.	79
5.8	Events passing all cuts except timing cuts (dots) and all cuts including primary timing cut (x's) for Tower 1. Vertical axis is ionization yield and horizontal axis is recoil energy (keV). Missing detectors were not analyzed in this final stage due to known performance problems.	83
5.9	Events passing all cuts except timing cuts (dots) and all cuts including primary timing cut (x's) for Tower 2. Vertical axis is ionization yield and horizontal axis is recoil energy (keV). Missing detectors were not analyzed in this final stage due to known performance problems.	84
5.10	Yield vs time for the data set 140616_1450, detector T2Z5. All events are shown; single events are Xs. Large, dark Xs are events that passed the timing cut with energies between 7–10 keV. Large, lighter Xs pass are above 10 keV and pass all cuts but the timing cut; the additional circle indicated that the event passes the timing cut. The circled event is the 10.488 keV event that passes all cuts. It occurred in the middle of a period when the gamma band in this detector had clearly drooped away from 1. After the LED flash, the detector recovered.	85
5.11	Spin-independent limits from R119 (dashed) and combined advance R119 and R118 (solid), compared to other limits and theories. Standard halo assumptions are used.	89
5.12	Spin-independent limits from R119 Si data. Two curves from Fig. 5.11 are included: the combined R118 and R119 Ge limit and the ZEPLIN-I limit, both of which are vertical at a WIMP mass of 10 GeV. Standard halo assumptions are used.	90

5.13	Spin dependent allowed regions and limits for (a) pure WIMP-neutron coupling and (b) pure WIMP-proton coupling. Results in black are new CDMS results, Ge in solid and Si in dot-dash. The dashed curve is combined Ge R118 and R119 analysis, and the grey (neutron plot only) is from the last CDMS run at the shallow SUF site. In red is the allowed DAMA region, and (b) includes limits from CRESST (blue crosses), SIMPLE (green circles), and SuperKamiokande (dark green stars). Figure from [57], produced by [41].	91
A.1	RBS spectrum of 3 CDMS wafers: 03B90 has only am-Si, 04a16 has am-Si and Al, and 04a12 has am-Si, Al, and W. See text for interpretation.	93
A.2	PIXE spectrum of 04A12, which has am-Si, Al, and W. See text for interpretation.	94
A.3	(above) Auger scan of surface of CDMS wafer 03A36, which has W, Al, and am-Si. (below) Depth profile obtained using Auger and sputtering. From [61].	96
A.4	(above) Zoomed-in RBS scan of sample 03B90 showing a clear contamination peak around the K energy, but also near the Ca and Ar energies. (b) PIXE spectrum of a long run (180 μ C) on 04A36, a sample with only W deposited on its surface (no Al). The contamination is identifiable as Ar.	97
A.5	High-resolution ICPMS scan around Sb masses of sample 04A16, a sample with am-Si and Al deposited. Natural Sb masses of 121 and 123 are clearly visible, as is the peak of the radioactive 125 isotope.	98
B.1	Muon kinetic energy spectrum in the Soudan experimental cavern. Average energy is 225 GeV.	100
B.2	Muon cos (azimuthal angle) distribution in the Soudan experimental cavern. A value of -1 is straight down. Average value is -0.81.	100
B.3	Neutron kinetic energy spectrum in the Soudan experimental cavern. Average energy is 41 MeV.	101
B.4	Pion kinetic energy spectrum in the Soudan experimental cavern. Average energy is 2.4 GeV.	101
B.5	Gamma kinetic energy spectrum in the Soudan experimental cavern. Average energy is 35 MeV.	103
B.6	Distance (m) between neutrons and muons that enter the cavern in the same shower event.	103
B.7	Simulated spectrum of energy (MeV) deposited in events in veto panel 1, a panel on the top of the CDMS shield.	104

Chapter 1

Introduction to the Search for Dark Matter

1.1 Evolution of the Standard Cosmological Model

Fritz Zwicky first proposed the existence of unseen, extra matter in the universe in a 1933 paper [1], in which he observed that the velocities of galaxies in clusters were too high; the clusters should have flown apart because there was not enough matter to keep them gravitationally clustered. His solution, the proposed mystery dark matter, lay fairly unexplored until it was revived in the 1970s, most prominently by Vera Rubin and her colleagues [2]. Fig. 1.1 shows a typical galaxy's rotational curve and what was expected. Using more advanced and reliable observational techniques, astronomers have cataloged the motions of hundreds of galaxies and plotted their rotation curves, some of which are shown in Fig. 1.2. The contribution of visible matter should have caused the velocities of objects at large radii to fall off more rapidly than they did, as shown by the dashed line representing Keplerian orbits in Fig. 1.1. The simplest explanation was a spherical cloud of invisible mass, extending far beyond the radius of the visible matter, thus making the velocity of visible objects roughly constant versus radius.

Rotational curves of galaxies, as well as velocities on larger scales, such as galaxy clusters, continued to be collected and published through the 1980s and into the 90s. As the evidence mounted, it was clear that even more extra mass was needed at larger scales and that there was unlikely to be any concentration of visible matter without some accompanying concentration of invisible matter. These data combined to tell us very little about the matter itself, however: only that it (a) had mass, (b) was dark (did not interact electromagnetically), and (c) was distributed rather evenly in association with visible matter. The phrase "dark matter halo" is used commonly to refer to these distributions of invisible matter associated with visible matter clumps. The most recent studies of the gravitational effects of a dark matter halo and its distribution come from the




Figure 1.1: Rotational curve of a typical galaxy. The dashed line is the expected Keplerian rotational curve calculated from the visible disk matter. The dotted line is the contribution of gas in the galaxy. The dot-dashed line is the calculated contribution from a spherical halo of dark matter. Figure taken from [3].

SLOAN Digital Sky Survey (SDSS) in New Mexico [6].

Historically, rotational curves were the first hint of the existence of dark matter, and they remain the simplest justification for its existence. In the last two decades, however, observational data from many sources have provided additional evidence for the existence of dark matter. They also offer more detailed information about its composition.

MACHOs and Nucleosynthesis

The simplest solution to the dark matter mystery seemed to be ordinary matter that we couldn't see. Candidates included large rocky planets, brown dwarfs, clouds of H or He gas, and black holes (not exactly ordinary, but at least we know they exist). Looking for clouds of gas was fairly straightforward, because we know its properties well. Several experiments, using instruments like the Hubble Telescope and the Hopkins Ultraviolet Telescope, looked for these clouds of gas and found enough to make up a only a small fraction of the dark matter, less than 10% [4].

The other two kinds of ordinary matter were lumped together in one category of objects, called MAssive Compact Halo Objects (MACHOs). These are harder to look for, but because they are massive and compact, they bend passing light rays toward them. The result is a lensing effect, because of which we on earth may see a temporary brightening of one distant, bright object. The light rays from the object in the background are bent by the dark object in the foreground, and

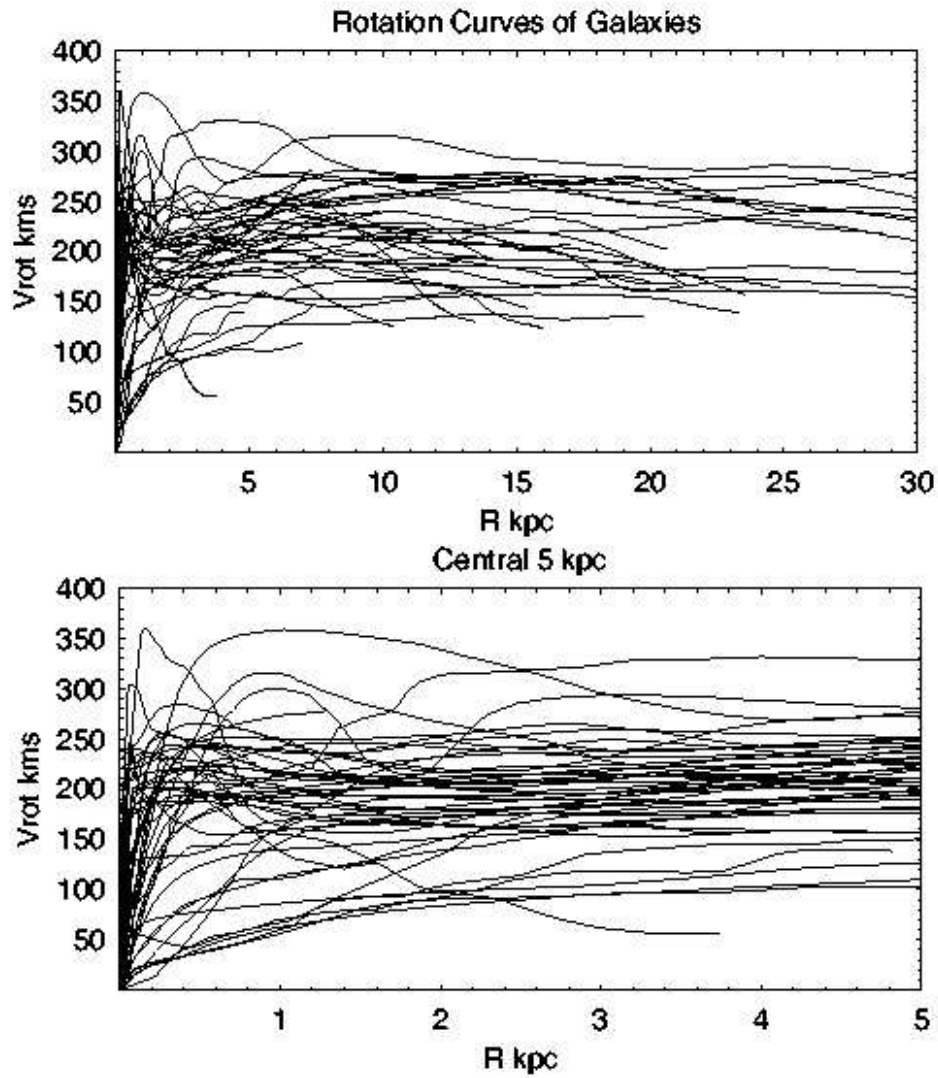


Figure 1.2: Superimposed rotational curves from several different galaxies of different types, and a zoomed in look at lower radii. The flat rotational curve at large radii is a shared feature. Figure taken from [5].

that is how we detect the presence of the dark object.

Many experiments, including most notably EROS [7] and MACHO [8], have searched for these MACHOs, and they have found dozens of these dark objects. However, the density of the object is so low that they can make up at most 25% of the dark matter observed in the universe.

Models of how elements were formed in the early universe, generally referred to as Big Bang Nucleosynthesis (BBN), give further limits on what the dark matter, even these compact halo objects, can be made of. Observational experiments based on these models indicated that baryons, the particles that make up gas and rock and the matter that we know, can make up at most 15% of the matter in the universe.

These searches have constrained the possible distribution and composition of dark matter; it is not concentrated in relatively few very massive objects, but it is far less compact. Also, most of it is not made of particles we've seen before, but something entirely new.

The Destiny and Composition of the Universe

The quest to discover the ultimate destiny of the universe was not launched in answer to the dark matter mystery, but it is inexorably tied up with it. Sophisticated observational studies necessarily address both questions together.

The Big Bang and inflation theory, well-accepted parts of the standard cosmological model give us an expanding universe and address the question of how the universe began. One of the fundamental questions in this study is how much matter and energy there is in the universe. Observers set out to measure the density of the universe, and express the contribution of any component as a fraction of the critical density for a flat, matter-only universe, or $\Omega_x = \rho_x / \rho_{critical}$, where

$$\rho_{critical} = \frac{3H_0^2}{8\pi G} = 1.88 \times 10^{-26} h^2 \text{kgm}^{-3} = 10.5 h^2 \text{keVc}^{-2} \text{cm}^{-3} [12] \quad (1.1)$$

where H_0 is the present value of the Hubble constant, h is the dimensionless form of H_0 in units of 100km/s/Mpc, which is roughly $0.7 \pm 5\%$. An Ω_{total} value of 1 would indicate a flat universe.

Many experiments put limits on the density of the universe, most notably by the following two methods. Type 1A supernova events at high red shift can be used as standard candles. Since these supernovae are very similar in spectra and time development, they are used to determine cosmological characteristics of their local environments. The observational experiments Supernova Cosmology Project (SCP) and High-Z Supernova Search Team both announced in 1998 that the expansion of the universe is accelerating, and have since published additional studies refining their measurements of the expansion rate. They constrain a combined Ω_m and Ω_{energy} to give the accelerating expansion they measure.

The second method studies the Cosmic Microwave Background (CMB), the radiation picture of the 300,000-year-old universe. In 1982, Peebles wrote a paper in which he postulated that some

form of weakly interacting dark matter particles at the time the CMB was created was necessary to account for structure formation in the universe, given how smooth the CMB was at the time [14]. If this were true, the CMB could not be truly smooth in temperature, and he predicted the form of the power spectrum of anisotropies in the CMB. He concluded that "The expected background temperature fluctuations are well below present observational limits." [14] If they could be measured, though, they would address both questions of history (how did galaxies form?) and the future (what is the density of mass in the universe?).

Several experiments have now measured the power spectrum of the anisotropies in the CMB. The COBE experiment in 1996 made a major breakthrough by measuring the magnitude of the fluctuations, published in [15]. Many experiments with greater sensitivity followed, including balloon-based experiments like BOOMERanG and space-based experiments like WMAP [17, 16]. These experiments constrain the total density of the universe, indicating a flat geometry.

A combination of all applicable methods of measuring the density of the universe gives a picture of the universe that is becoming an accepted part of the Standard Cosmological Model. In [22], a combined analysis of WMAP, supernova, Big Bang Nucleosynthesis, and Sloan Digital Sky Survey data gives

$$\Omega_{total} = 1.012^{+0.018}_{-0.022} \quad (1.2)$$

$$\Omega_{matter} = \Omega_{darkmatter} + \Omega_{baryons} = 0.317^{+0.053}_{-0.045} \quad (1.3)$$

$$\Omega_{darkenergy} = 0.695^{+0.034}_{-0.037} \quad (1.4)$$

$$h^2\Omega_{baryons} = 0.023^{+0.013}_{-0.011}; \Omega_{baryons} \approx .0528 \quad (1.5)$$

$$h^2\Omega_{darkmatter} = 0.115 \pm 0.012; \Omega_{darkmatter} \approx 0.2640 \quad (1.6)$$

These data answer many questions. One is that the total density of the universe seems to be very close to the critical density. Also, that matter contributes only about 30% of that density; the rest is an even more mysterious quantity called dark energy. Of the dark matter, roughly 80% is non-baryonic, something altogether different from the material that makes up what we know of the universe, i.e. stars, rocks, us, etc. The data combine to give us specific characteristics of that dark matter, as well as its density. It clusters with visible matter, but does not form large dense bodies as visible matter does; that is, it is most likely a cloud of some kind of non-baryonic particle. It has existed since the CMB was created, so it is very stable. To act as the gravitational seeds during the formation of galaxies, it cannot have been or be relativistic in kinetic energy (it must be "cold"); high energy dark matter could not have helped to form galaxies. It still has to be weakly interacting and massive. The data pose more questions than answers, though. The uncomfortable conclusion is that we do not know what 95% of the universe is.

There are alternatives to the Standard Cosmological Model, as there are to any theory. A clear alternative to the existence of dark matter is the possibility that we simply do not understand gravity on large scales yet, and it contrives to produce what we see with only the matter of which

we are already aware. The most commonly mentioned alternative to dark matter is Modified Newtonian Dynamics (MOND), and while it explains rotational curves at most radii, Sloan Digital Sky Survey data at very high radii disagree with MOND theory, and WMAP data are more consistent with dark matter theory than with MOND [6, 12, 20].

Standard Dark Matter Parameters

Dark matter, an idea older and a little less daunting than dark energy, has several known characteristics now that experiments searching directly for dark matter must take into account. To estimate a dark matter interaction rate in a detector in the Earth's local environment, we use

$$\frac{dN}{dt} = \phi \sigma_n \frac{\sigma_A}{\sigma_n} N_{target} \quad (1.7)$$

where ϕ is the local dark matter flux, σ_n and σ_A are the WIMP-nucleus and WIMP-nucleon cross sections, respectively, and N_{target} is the number of target nuclei per volume. Eq. 1.7 takes this particular form because for an experiment, it is σ_n that is most easily calculated, and the ratio σ_n/σ_A is a separate theoretical calculation.

The local WIMP flux, ϕ , is calculable from the observational data described above, and depends mainly on the local dark matter density and velocity distribution. A common calculation of the local dark matter density uses experimental rotation curves from the Milky Way, and is reported in [23] as $0.3 \text{ GeV}/c^2/\text{cm}^3$. The velocity distribution under simple assumptions, based on a spherical halo in thermal equilibrium, in [23] becomes

$$\frac{dP(v)}{v^2 dv} = \frac{1}{(\pi v_0^2)^3/2} e^{-\frac{v^2}{v_0^2}} \quad (1.8)$$

where $v_0 \approx 220 \text{ km/s}$, typical of speeds of objects in our galactic neighborhood. Different theoretical assumptions about the geometry and characteristics of the halo can change the average value of this velocity distribution within a range of 170-270 km/s, but 220km/s is currently used as the standard value [12]. The motion of the Earth around the sun as it travels through the galaxy and therefore the local dark matter halo will either boost or suppress this velocity depending on the time of year, but it will not affect the average.

The cross section σ_n is not known, but the ratio of nuclear and nucleon cross sections is calculated in [23]. The ratio depends heavily on whether the interaction is spin-dependent (SD) or spin-independent (SI), so two different ratios are used simultaneously:

$$\frac{\sigma_A}{\sigma_n} = \frac{\mu_A^2}{\mu_n^2} A^2 (SI) // \frac{\sigma_A}{\sigma_n} = \frac{\mu_A^2}{\mu_n^2} C J(J+1) (SD) \quad (1.9)$$

where μ is the invariant mass of the WIMP-nucleus or WIMP-nucleon system and C is a factor, generally less than 1, that depends on the structure of the nucleus.

1.2 Theoretical Dark Matter Candidates

By making some simple assumptions about the geometry of the dark matter halo, we can calculate all of the necessary quantities to estimate an interaction rate in detectors except two: the WIMP mass and the WIMP-nucleus cross section. The astrophysical observations do not tell us much about these two parameters. In order to limit the otherwise infinite parameter space, we use theory to suggest probable candidates for the dark matter constituents.

One clear candidate for a component of dark matter is the neutrino, a non-baryonic, very light particle known to exist, have mass, and interact weakly, as well as be distributed smoothly in the galaxy. There are three flavors of neutrinos in the Standard Model of Physics, and neutrino physics is a very active area of study because the masses of these neutrinos are unknown, and there are possibilities for additional types outside the Standard Model. The standard neutrinos are limited by experiment to have masses on the order of 1 eV or less, and would be a very small contributor to the total dark matter density, at a maximum of $\Omega_\nu h^2 < 0.07$ [40]. More importantly, standard neutrinos are relativistic, and therefore do not meet the requirement that the dark matter be cold. Some theories predict non-standard, sterile, and heavy neutrinos, but these are constrained to be heavier than 45 GeV. Early dark matter searches were sensitive to the corresponding parameter space and did not see a signal. With new WMAP results, sterile, heavy neutrino theories are heavily constrained [12, 40]. Thus, it seems unlikely that neutrinos are a significant component of dark matter.

With the elimination of neutrinos as a major constituent of dark matter, there are no other candidates that both fit the requirements for cold dark matter and are known to exist. Two significant theoretical particles emerge as leading possibilities for dark matter constituents. One is the axion, a boson proposed to solve the strong CP problem of Quantum Chromodynamics, and further appear in superstring theories. Axions would be very light, at most 10^{-4} eV, cold, and stable. Their cross section with ordinary matter is postulated to be very small, on the order of 10^{-20} pb. There are numerous experiments searching for axions, but none have reported a positive signal. For example, ADMX is a cavity based experiment at LLNL which recently posted sensitive limits on axion flow in the mass range around $2 \mu\text{eV}$ [27]. The experiment is upgrading for further sensitivity.

The other theoretical candidate is loosely called a Weakly Interacting Massive Particle (WIMP), and it is defined as the lightest, stable particle predicted by Supersymmetric extensions to the Standard Model. Supersymmetric theories have been proposed to address several questions in the Standard Model of Physics, most prominently the gage hierarchy problem. This problem arises when radiative corrections to particle masses are calculated. For fermions, these corrections are small, and the contributions of these corrections are logarithmically divergent. However, for bosons (the gage bosons, which are scalars), these contributions become very large, diverging quadratically, which disagrees with data. The solution provided by supersymmetry is to propose the existence of partner particles for all known particles: a boson partner for each fermion and a fermion partner for each boson. This arrangement cancels the quadratic divergences. References [25] and [24]

contain good overviews of Supersymmetry. The lightest predicted neutral particles include the fermion partners to neutral gage bosons in the Standard Model, generally called neutralinos.

In the minimal extension of the Standard Model (minimal Supersymmetry, or MSSM), the lightest supersymmetric particle (LSP) is expected to be a mixing of four neutralinos. The neutral gage bosons in the Standard Model are only three in number: the photon (electromagnetic), the Z^0 boson (weak), and the Higgs boson. (The Z^0 and the photon are themselves mixings of two neutral bosons associated with symmetries, the B and the W^3 ; supersymmetric theories generally refer to the B and the W^3 rather than to the photon and the Z^0 .) MSSM proposes an additional Higgs boson and four neutralinos: the Bino (partner to the B), the Wino (partner to the W^3), and two Higgsino states. Theories differ on how and whether these four neutralinos might mix. The LSP is most likely a mixing of these four, and the nature of the mixing will determine its cross sections and mass. Therefore, there is a wide, but not infinite, range of WIMP masses and WIMP-nucleon cross-sections defined by MSSM and the theories that stem from it.

MSSM introduces several unknown parameters; efforts to constrain these parameters include constrained MSSM (CMSSM), which assumes gaugino and scalar mass unification at GUT scales, and minimal super-gravity (mSUGRA), which also incorporates gravity at the GUT scale (gravity is not included in the Standard Model, another significant inadequacy) [25, 26, 40]. The parameter spaces allowed by several theories in these two frameworks will be included on plots throughout this thesis. Some recent theoretical calculations are plotted in Fig. 1.3; theories are either regions or a set of likely points, denoted by Xs. It should be noted that, while the WIMP mass varies over a few orders of magnitude, the more daunting range is in the cross section calculation, ranging over more than 10 orders of magnitude in a relatively small mass range.

There are several other theoretical candidates for dark matter constituents, including most commonly those resulting from Kaluza-Klein extra dimensional theory, axinos, and WIMPZILLAS (non-thermal WIMPS $10^{10}\times$ more massive than standard WIMPs), but these are generally less natural solutions and far more difficult to detect than WIMPs or axions([12]). Experiments are focused on these two primary candidates, and we will focus mainly here on WIMPs.

The SUSY theories for WIMP mass and WIMP-nucleus cross section give ranges for the final two variables necessary in Eq. 1.7, allowing experimentalists to design detectors and look for WIMPs, and the resulting experiments are the subject of the next section. However, it should be remembered that the theories merely point to a likely parameter space for dark matter WIMPs; discovery of dark matter in this parameter space does not necessarily mean that it is supersymmetric, nor would discovery of a supersymmetric particle in an accelerator mean that it is the major component of dark matter. The overlapping regions of dark matter searches and supersymmetry signatures will thus provide complementary information.




Figure 1.3: WIMP-nucleon cross section vs WIMP mass for several different recently published theories.

1.3 Experimental searches for WIMPs

The experiments surveyed here fall into two main categories, indirect and direct searches, with a heavy emphasis on the latter. Indirect searches detect the products of WIMP-WIMP annihilation, which is most likely to occur in mass concentrations such as the earth, the sun, and the center of the galaxy. These experiments make few assumptions about the dark matter halo except that the WIMP density is greater where visible matter density is greater, and that the WIMPs are supersymmetric and have predictable annihilation products. Direct searches use underground detectors to look for WIMP-nucleon interactions as the earth moves through the local halo. These experiments must make assumptions about the dark matter halo distribution and velocity to predict their sensitivities. However, direct searches have more control over the backgrounds in their environments, as opposed to indirect searches, which use celestial objects as their laboratories.

1.3.1 Indirect WIMP Searches

Decay chains from WIMP-WIMP annihilations are various, including lepton-antilepton, quark-antiquark, and gauge boson pairs and their subsequent decays. The final decay products are primarily very energetic leptons and hadrons.

The first channels to be thoroughly explored were neutrino daughters from neutralino self-annihilation. Experiments such as MACRO, AMANDA, and Super-Kamiokande, have examined neutrino data and set upper limits on neutrino fluxes coming from the center of the earth and sun. However, calculations of the velocity distribution of halo WIMPs have indicated that WIMP capture in the Earth is suppressed; thus discovery through that channel is unlikely [28]. Several experiments have also looked for antiprotons and positrons, such as the balloon-borne BESS [19] and the High Energy Antimatter Telescope (HEAT) [18], but distinguishing a WIMP annihilation signal from the cosmic ray interaction background is very difficult in these experiments. Data from BESS could be interpreted as excluding WIMP masses above 100 GeV under certain assumptions.

High energy gamma rays from the galactic center are another possible signature for WIMPs. Data from EGRET indicated excess gamma rays at high galactic latitudes and above 1 GeV that may be interpreted as resulting from dark matter annihilation under some assumptions [21], but that data is inconclusive at best. Further exploration of this phenomenon by the Gamma-ray Large Area Space Telescope (GLAST) may help to clarify this matter [28].

1.3.2 Direct WIMP Searches

There are many internationally experiments searching for direct WIMP interactions using sensitive detectors. The three main categories consist of: 1) NaI experiments, mostly in the past, which have poor background rejection and look for annual modulation of the WIMP signal, 2) current cryogenic experiments sensitive to nuclear recoils with background rejection, and 3) current and future liquid-based techniques looking for nuclear recoils with background rejection.

WIMPs scattering elastically on nuclei produce a low energy recoil spectrum and thus these experiments' thresholds must be quite low. This requirement poses a challenge, because many of the experimental techniques were borrowed from high-energy particle physics. Fig. 1.4 is a plot of recoil energy distributions for three WIMP masses in Ge, a typical target. Lower mass WIMPs give lower energy spectra; a 10 GeV WIMP would be very difficult to detect with a 10 keV threshold. Most WIMP searches focus on the recoil energies below 100 keV, the range shown in the figure, because this range is appropriate for the more likely WIMP masses of 50 and 100 GeV. The main two challenges for any direct detection experiment are low energy sensitivity and background rejection.




Figure 1.4: Recoil energy distributions for 3 typical WIMP masses. Figure taken from [11].

NaI experiments

Two main experiments, DAMA in the Gran Sasso tunnel and the DM collaboration in the Boulby Mine, UK, used NaI crystals to detect WIMPs, originally with the intent to use pulse-shape discrimination to identify a WIMP signal, beginning in the 1990s. However, the pulse shape method, one borrowed from high energy particle physics, proved inadequate for low energies essential for the WIMP search. Instead, these experiments used their data to look for the effect of the earth's annual revolution around the sun, theoretically limited to a 7% effect.

In 1998, DAMA reported an annual modulation signal with one year of data, and they have since published similar results from seven years of data [29, 30]. The signal they reported corresponds to a rate of approximately 1 nuclear recoil/kg/day in Ge based experiments, and the signal claim was

quickly followed by null results in experiments, such as EDELWEISS and CDMS, that excluded the WIMP mass and cross section allowed by the DAMA signal [31, 37]. More recent measurements have only strengthened those exclusions. Using unusual assumptions about the dark matter halo, it is still possible, though difficult, to reconcile the DAMA dark matter discovery claim with other experiments' exclusion limits. As other experiments become more sensitive, however, this becomes more and more difficult, and it seems likely that the annual modulation may be due to some non-WIMP source. DAMA is expanding and updating their experiment under the name LIBRA.

The DRIFT experiment, a gaseous ion drift chamber which is not NaI based and therefore subject to different systematic effects, may help in untangling the annual modulation puzzle. It is direction sensitive and has significant background rejection. It has been running since 2002 in the Boulby mine in the UK [33], but has yet to report results.

Solid State Experiments

Currently, solid state experiments are setting the most sensitive upper limits on WIMP mass cross sections. Generally, they are based on very cold solid state detectors and measure two physical quantities to discriminate nuclear recoils, in which incident particles scatter off nuclei in the crystal, from electron recoils, in which they scatter off of valence electrons in the crystal. These recoils deposit energy in three forms: scintillation, ionization, and phonon (the vibrations of the recoil in the lattice). Experiments measure either scintillation and phonon signals or ionization and phonon signals. Nuclear recoils produce lower ratios of scintillation or ionization to phonon energy than do electron recoils, and this fact can be used to reject electromagnetic backgrounds, the most prominent background in most experiments. These experiments are at deep underground sites, like the NaI experiments, and thus shielded from the cosmogenic muon and neutron flux. Experimentalists have been successful in adapting high energy techniques to recoil energy thresholds of a few keV.

CRESST, in the Gran Sasso tunnel, measures scintillation and phonons using a sapphire, or Ca_2WO_4 , target. The collaboration is currently working to decrease its energy threshold, and suffers from significant uncertainty in the quenching factor for light in Ca and W. However, it has set competitive upper limits on WIMP-nucleus cross sections using both the rejection method described above and the physical differences between the Ca and W in the target [34].

The Cryogenic Dark Matter Search (CDMS) at the Soudan Underground Laboratory in Minnesota, measures ionization and phonons using Ge and Si targets. The phonons measured are athermal, meaning a faster signal rise and relaxation time. They also record position information of the recoils, allowing rejection of events near the surfaces of the detectors. The collaboration has reported results from two data runs in that location and is currently running increased mass in the lab [35, 36].

EDELWEISS, in the Frejus Tunnel on the France-Italy border, also measures ionization and phonons using Ge targets. The phonons measured in this experiment are thermalized, and are detected by temperature fluctuations in the crystals themselves. The collaborations has recently completed a significant increase in detector mass and improvement in instrumentation and will

begin running soon [37].

Liquid Experiments

With the exception of the ZEPLIN experiment in the Boulby mine, most liquid experiments are in the proposal and development stages and have not yet published competitive results. In general, these experiments use targets of large masses of liquid Xe or Ne, detecting scintillation and ionization signals for background discrimination. Many use two phases, liquid and gas, to amplify the signal from the recoils in the liquid, and position information to avoid the edges of the detectors.

ZEPLIN, in the Boulby mine in the UK, is a multi-stage liquid Xe experiment which measures scintillation. The first of these stages, in which the detector was basically a U-shaped pipe of Xe with a PMT at either end, resulted in competitive cross section limits in both the spin dependent and spin independent cases [38]. This experiment will eventually scale up to the 1 ton scale.

The XENON experiment, in the US, uses a two stage Xe detector with CsI cathodes immersed in the liquid to measure scintillation. The gas phase amplifies the ionization signal. XENON has presented first results from running a 10kg prototype detector above ground [39], and is in the process of moving to a deep site.

1.3.3 Limit Comparisons

Fig. 1.5 is a snapshot of spin-independent experimental limits, using standard halo assumptions, published before the results reported in this thesis. Included are the 3σ allowed region from DAMA and limits from several experiments mentioned above. The limits themselves are the 90% C.L. exclusion limits at each mass; parameter space above these lines are inconsistent with data. The fairly sharp edges of the limits at low WIMP mass are defined by the thresholds of the experiments; most experiments are not sensitive to WIMPs with masses below 10 GeV. The slow rise at high WIMP masses is due to the fact that halos of heavier WIMPs would be less dense; thus an experiment with a fixed exposure is less sensitive to heavier WIMPs. The first CDMS limit from Soudan is the most sensitive in the spin independent case, but it only begins to exclude a small portion of the CMSSM parameter space. The experimental limits combine to exclude a significant portion of unconstrained MSSM parameter space, as well as the DAMA signal interpreted with standard halo assumptions.

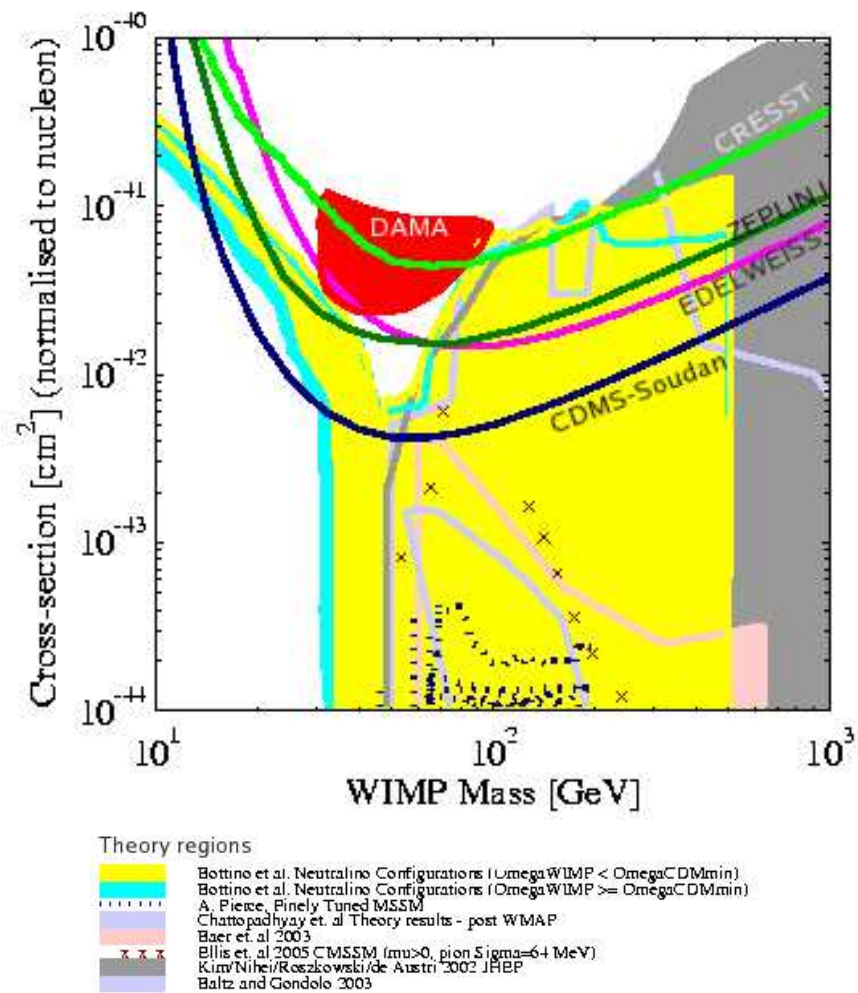


Figure 1.5: Spin independent upper limits on cross sections, including the allowed region from DAMA, superimposed on the same theoretical parameter space as in Fig. 1.3, zoomed in on the experimental limits.

Chapter 2

CDMS II Detectors and Experimental Setup

The Cryogenic Dark Matter Search II, housed at the Soudan Underground Laboratory in northern Minnesota, is a direct search for WIMPs. CDMS II uses Z-sensitive Ionization and Phonon detectors (ZIPs), operated at temperatures around 50 mK, to simultaneously record athermal phonon and ionization signals from any particle scattering in the crystals. Because CDMS II is a rare event search, it requires a low background environment and uses both passive and active shielding. Two runs have been completed in the Soudan Laboratory: run 118 with six ZIPs and run 119 with twelve.

2.1 Detector Description

ZIP detectors are 7.81 cm in diameter and 1 cm thick, made of either Ge or Si. When an incident particle scatters in the crystal, it causes two physically different phenomena. First, it ionizes electrons in the semi-conductor. A constant applied electric field is set up across the ZIP to collect these electrons. Second, any scattering that takes place in the crystal shakes the lattice, producing localized vibrations, or phonons, that propagate through the lattice. Phonon sensors on the face of ZIP detectors then convert these vibrations to electronic signals. These two measurements, ionization and phonon, form the core of CDMS's data. The phonon sensors are segmented into four quadrants, and the ionization sensor into inner and outer electrodes, so ZIPs also provide position information in the event reconstruction. These features make the ZIPs well-suited to rejecting electromagnetic background particles that would otherwise mask any WIMP signal.

2.1.1 Ionization Measurement

A constant electric field is applied across each detector from top to bottom. The ionization grounding grid is interlaced with the phonon detector circuitry on the top face, and the charge electrodes are on the bottom of each ZIP. There are two concentric electrodes, an inner electrode and an outer guard ring, as shown in Fig. 2.1. The outer guard ring, with a width of just 1.9 mm, serves as a veto for any event that deposits energy near the outer edge of the detector, where the electric field is not well defined. When the insulating crystals are at 50 mK, it takes about 3 eV to break an electron-hole pair in Ge and 3.8 eV in Si. The electric fields across the detectors are therefore set at -3 V/cm for the Ge detectors and -4 V/cm for the Si detectors. Fig. 2.2 shows the biasing and amplification circuit for the ionization channels.

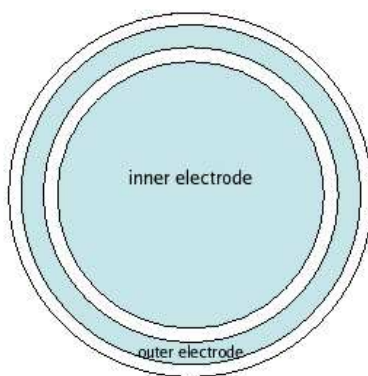


Figure 2.1: Ionization sensors on the bottom face of a ZIP. The inner electrode is considered the fiducial volume.

Fig. 2.3 is a typical ionization signal. Two traces are shown, one for each electrode. Note that most of the energy for this pulse, and for all good pulses, is in the inner electrode. The full digitization window is 1.64 ms ($2048\text{ bins} \times 0.8\mu\text{s}$), and the the ionization pulse itself rises from baseline to peak in just a couple microseconds. The decaying exponential after the pulse maximum has a time constant of approximately $40\mu\text{s}$, being fixed uniquely by the time constant of the feedback circuit.

Three phenomena in the detectors complicate what would otherwise be a simple measurement of the ionization energy deposited in the crystal. The first can be seen in the square grounding grid geometry shown in Fig. 2.4. Because the grounding electrode is deposited in the square geometry on the face of the detectors and the detectors are round, there are areas near the outer guard ring, but still clearly inside of it, that have poor grounding. There is a layer of amorphous Si on the faces of each crystal, which should have provided a moderately good grounding plate for these areas of

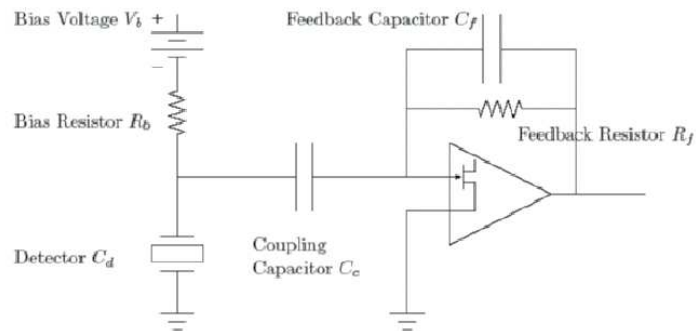


Figure 2.2: Schematic of ionization circuit. From [3].

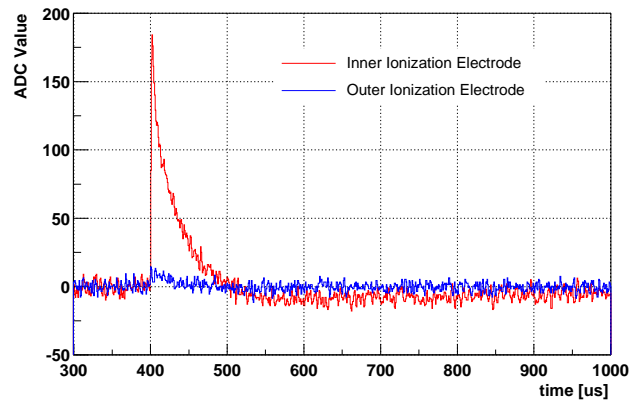


Figure 2.3: A typical ionization channel pulse. This event took place in T1Z3, and is a 50 keV electron recoil event. Like all events selected as good events, this event deposits nearly all ionization energy in the inner electrode.

the detector. However, the amorphous Si did not prove to be a good conductor in practice, and events occurring in these areas of the detectors had low charge collection efficiency. In run 119, this issue was addressed in analysis, but design changes were made to subsequent ZIPs, and future detectors (including towers 3, 4, and 5) will not suffer from this problem.

The second is surface events, another geometry-related issue. At the detector surface is a layer, approximately $30\mu\text{m}$ thick, in which electrons or holes can travel backwards to the wrong electrode, due to its proximity to the electrode and the shape of the electric field in these regions. The net effect is a suppressed ionization signal. The amorphous Si on the surface of the detectors serves to mitigate this effect, but it still poses one of the most challenging issues to CDMS. The problem of rejecting these surface events caused by electrons, which are most often produced by photons or β -decays from surface contamination, is addressed in detail in Chapter 4.

The third complication is related to impurities in the crystals. Ideally, freed electron-hole pairs drift unencumbered to the electrodes on the faces of the detectors. However, if the crystals are biased for a long period of time, they develop localized excesses of charge. These are the natural result of the fact that any crystal will have a small number of impurities that have either donor or acceptor states close to the conduction band. Such impurities can form traps for free electrons traveling through the crystal. When the detectors are initially cooled down, these impurity sites are fully ionized and interfere with charge collection. Before data can be taken, each detector is “baked” using LEDs mounted near each face. It is flashed repeatedly with photons, and these photons produce electron-hole pairs in the crystal, which then fill the acceptor and donor sites in the impurities. Once these sites are filled, signal electrons can drift through the crystal and arrive at the electrodes. The detector is then said to be fully neutralized; the neutralization process takes several hours for each Ge detector, and a few days for Si detectors.

When the detectors have been continuously biased for a long period of time, the electrons in the impurity sites drain away, leaving electron traps behind and again degrading charge collection. To ensure good charge collection, data taking is stopped on a regular basis, then the ZIPs are grounded, and the LEDs are flashed to restore neutralization. Neutralization was more closely monitored, and ZIPs more often flashed, in run 119 than in the previous run.

Si detectors more easily lose neutralization than do Ge detectors, and studies of run 118 data showed that Si detectors, especially T1Z6, began to show loss of neutralization after about 4.5 hours of continuous biasing. Therefore, the data acquisition system was programmed to stop data taking every four hours, flash the LEDs, wait for the detectors to return to base temperature, and resume data taking automatically. Additionally, the detectors were all flashed at least twice a day during periods of cryogen transfers, when data taking was suspended anyway. Usually, the icebox took between 15–30 minutes to return to base temperature after LED flashing, during which the innermost can could reach temperatures as high as 350 mK.

2.1.2 Phonon Measurement

A ZIP detector records the signals from four phonon channels, each of which covers 25% of the detector face, as shown in Fig. 2.4. The channels and the quadrants they correspond to are called A, B, C, and D throughout this discussion. When a particle of any kind scatters in the crystal, the dominant energy deposited takes the form of phonons, and so this signal is larger than the ionization signal. Since the ionization energy is also recorded in the phonon sensors once the electrons and holes reach the electrodes, the total phonon signal includes all of the originally deposited energy. The measurement and readout of the phonon energy is much more complicated compared to those of the ionization energy. A detailed explanation follows.

Hardware Fabrication

The phonon sensors are composed of Transition Edge Sensors (TESs), small strips of tungsten on the surface of the detector, as shown in Fig. 2.4. However, to describe the entire process of phonon collection, it is useful to first summarize detector fabrication, focusing only on the top of the detector, where the phonon detection is performed. Most of the fabrication is performed at the Stanford Nanofabrication Facility, and the basic steps are as follows:

- A 40 nm thick amorphous Si film is applied evenly over the face of the detector. All depositions are performed with a CD magnetron sputtering machine, Balzers 450.
- A 300nm thick layer of Al is deposited. This Al is the material used for phonon collection (fins) and the electrical connections between TESs (rails).
- A 35nm thick layer of W is deposited. This is the first layer of the TESs.
- Al fins are carved out. A photoresist is applied and exposed to an Ultratech stepper-aligner to mask desired areas. The machine can only expose a $5 \times 5 \text{ mm}^2$ area, so a grid of cells this size is made on the detector surface. Exposed areas are then etched away. This step forms the Al fins and rails shown in Fig. 2.4. The Al fins cover most of the area of the detector. Each cell has 224 Al fins, each $380\mu\text{m} \times 50\mu\text{m}$. An Al rail connects all cells in a quadrant in parallel. The remaining Al forms a grid for the grounding plane of the ionization signal.
- An additional layer of 35nm of W is applied. This layer completes the overlap regions with the Al fins (shown in Fig. 2.5 and will form the portion of the TES mounted on the am-Si, where the previous layer of W has been etched away.
- The surface is again masked, exposed, and etched to form the tungsten TES sensors on the surface of the detector. These are also shown in Fig. 2.4. There is one TES for every 8 Al fins, a total of 28 in each cell, and each is $250\mu\text{m} \times 1\mu\text{m}$. A side view of the final sensor is shown in Fig. 2.5.

The four quadrants of the phonon side each have 37 cells of this kind, for a total of 1036 TESs and 8288 Al fins per quadrant.

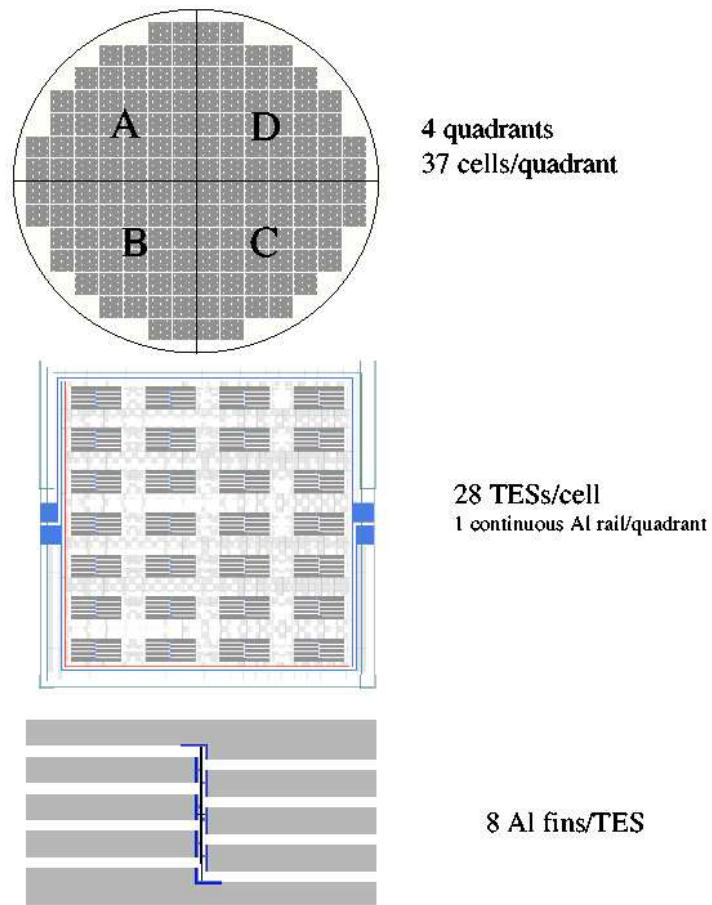


Figure 2.4: Phonon sensors on top face of a ZIP.




Figure 2.5: Side view of layers deposited during fabrication of the phonon sensors.

Phonon Sources and Detection

The phonons collected in the Al fins on the surface of the detector come from several sources:

- **The initial scattering interaction.** The phonons that are produced at the initial scattering are very high energy (approximately 3 THz in Ge and 6 THz in Si). These high frequency phonons scatter frequently and travel slowly from the site toward the edges of the detector. The edges of this phonon cloud travel at approximately 1/3 the speed of sound in the material. Along the way, they propagate through quasi-diffusion, dominated by two scattering processes: isotope scattering (in which the phonons scatter elastically with Ge or Si isotopes and change direction) and anharmonic decay (a splitting of one high energy phonon into two or more lower energy phonons, with a time scale of a fraction of a microsecond).
- **Down-converted ballistic phonons from the initial scattering interaction.** On a time scale of a few microseconds, the initial high frequency phonons can attain low enough energies (less than 1 TeV) such that their mean free path is comparable to the size of the detector. They are then ballistic, bouncing from detector edge to edge until they are absorbed in the Al fins. This process is greatly enhanced if the original scatter takes place fairly close to a detector face.
- **Neganov-Luke phonons.** These phonons are produced as the electron-hole pairs drift across the crystal. They are ballistic, and their total energy is equal to the work done to drift the charges across the crystal to the electrode.

- **Ionization recombination.** Once the electrons and holes reach the electrodes, they recombine as pairs in the Al, and this energy is recorded in the Al fins. These phonons are originally high energy but rapidly down convert because they are already at the surface of the detector.

When these phonons are absorbed into the Al covering most of the face of the detector, about 30% of their energy is transferred to the Al. At 50 mK, the Al is superconducting and therefore a sea of electrons grouped in Cooper pairs. The Cooper pairs will break up if they absorb more than $340 \mu\text{eV}$, the energy of the superconducting energy gap of Al, which is equivalent to a phonon frequency of 82 GHz. When incident phonons deposit more than this threshold energy, they break up Cooper pairs, and the resulting quasiparticles diffuse into the nearby tungsten TESs. Due to the difference in the superconducting energy gap in Al and W, quasiparticles that lose energy become trapped in the W and cannot back-diffuse.

Detection Circuitry

The W strips are operated as electro-thermal feedback TESs (Transition Edge Sensors). They are maintained at a temperature in the middle of their superconducting transition region (critical temperature, or T_c) by a positive feedback circuit using current biasing, shown in Fig. 2.6. The absorption of quasiparticles from the Al causes a large change in the resistance, and therefore current, of the TES, which is read out by a SQUID (superconducting quantum interference device) array. At the same time, the increased resistance in the TES decreases the current in the biasing circuit, reducing Joule heating, and the TES cools back down to its operating temperature.

The SQUID array readout mechanism was chosen for its low noise characteristics. The amplification is a factor of 10, the ratio between the number of turns in the feedback coil to the input coil.

It would have in principle been possible to collect the phonons directly with W covering the whole of the detector. However, that much W on the detectors would have increased the inductance of the readout circuitry enough to make operation of the electro thermal feedback system nearly impossible. The design of TESs assisted by Al quasi-particle traps is therefore necessary.

The performance of this rather complicated signal production mechanism varied greatly with detector. Due to the fact that each quadrant had several hundred TESs and therefore several hundred SQUIDS to readout, performance could greatly vary across the face of one detector and from detector to detector. Gradients in T_c across individual detectors resulted in variations in detected energy as a function of position. As discussed later in this chapter, all detectors suffer mildly from this issue. Of the twelve detectors operated for run 119, T1Z1 and T2Z1, the top detector in each tower, suffered such poor performance due to T_c gradients that they were excluded from the final analysis.

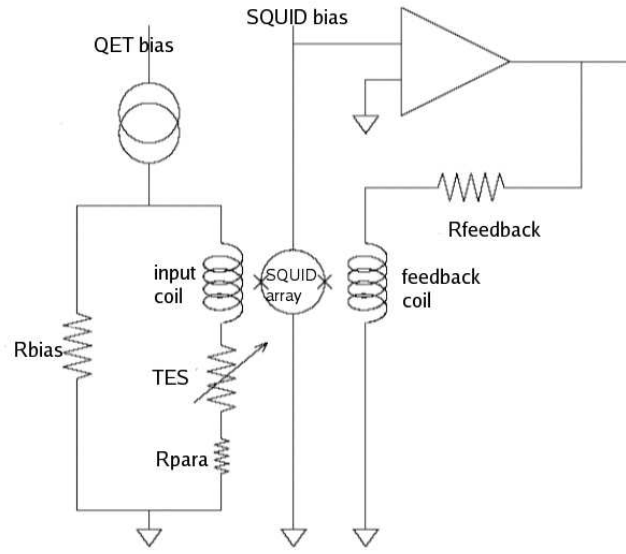


Figure 2.6: Schematic of the phonon detection circuit.

Phonon Signal

Fig. 2.7 shows four digitized phonon traces from a typical electron recoil in a ZIP. The time scale of the phonon signal is much slower than the ionization signal because of the time required for quasi-diffusion. The risetimes of the phonon signals are governed mainly by event position and the physics of phonon propagation in the crystal. The long fall times are mainly due to electronics, most prominently the thermal relaxation of the TESs.

The relative timing and energy data in each quadrant are rich in position and physics information. However, using this information requires that each quadrant respond similarly for a recoil of any given energy and type. This parity is achieved in two ways: first, by hardware tuning of the biasing currents in the SQUET (SQUID and QET) circuit, and second, by software corrections as a preliminary analysis step.

2.2 Tower Construction

The ZIPs are mounted in two hexagonal towers of six detectors each. Each tower consists of four sections, each at a different temperature and mounted to the appropriate icebox can. Fig. 2.8 is a side view diagram of a sample tower mounted in the icebox at Soudan. The ZIPs are at the bottom of the tower, in the 10 mK stage (actually operated at 50 mK). They are individually housed in hexagonal copper cases, but the spatial separation between them is only 3 mm, with a minimum

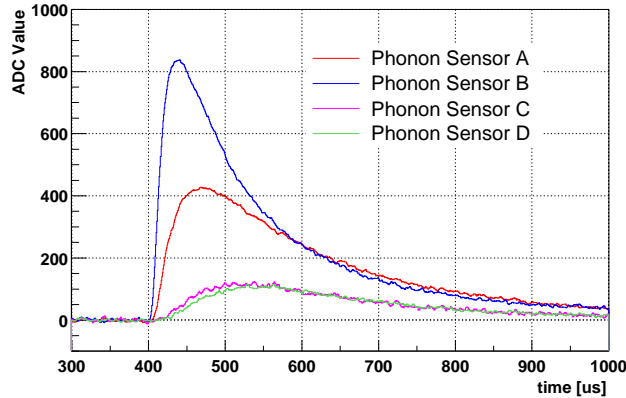


Figure 2.7: A typical phonon pulse, all four channels. This is the same event as in Fig. 2.3. It's located in quadrant B, which is the trace with the largest amplitude.

of intervening material in order to enhance the probability of multiple scatters between ZIPs.

On each side of the hexagonal tower is mounted a copper plate carrying the coaxial cables between an electronics card at the top of the tower to a detector down below; this is what demands a hexagonal geometry for a six detector tower. Each detector is therefore rotated 60° from the one above it, with the y position maximum defined as the location of the coaxial cable for that detector. This rotation should be remembered when examining position diagrams of events in the crystals; the relative positions of quadrants A–D differ from detector to detector.

At the top of the tower are the electronics cards, the cold electronics for the experiment. The SQUIDs that read out the phonon signals are operable at 4 K, but they are heat sunk to the 600 mK icebox can to reduce SQUID noise. The FET circuitry for the ionization measurement is mounted to the 4 K icebox can because their 4.5 mW of power dissipation can be tolerated at this temperature. The combination of the SQUID circuitry and the FET circuitry is collectively referred to as the SQUET card and is fabricated as one unit. The greatest heat loads are the graphite rods that house the connecting wires that must be heat sunk to the stages of the fridge. There are four such stages: the 10 mK (mixing chamber), where the ZIPs operate, the 50 mK (cold plate), the 600 mK (still), where the SQUID operate, and the 4 K (liquid helium), where the FETs operate.

At the top of the tower, flex signal cables 1 inch wide and 10 feet long, called striplines, carry signals through the shielding and out to the data acquisition system as discussed below. One stripline per detector is used.




Figure 2.8: A side view diagram of a CDMS tower mounted in the icebox at Soudan. From [25].

2.3 Warm Electronics and DAQ

Fig. 2.9 is a schematic of the data acquisition (DAQ) system, traced from the detectors themselves through to the computers in the electronics room. The striplines run from the towers through the shielding in a copper stem. The 50 signal traces of each stripline are 0.005" copper between two layers of .005" kapton, copper ground plates, and 0.001" kapton. The striplines are threaded through a few copper blocks en route. One of these copper blocks, called the gamma shield, includes a physical, angular shift so that there is no line of sight path for background particles to reach the detectors through the stem. On either side of the gamma shield are the He and N intercepts, named for the temperature of the cryogenic stage to which they are thermally coupled, warming the striplines from 4K inside to room temperature outside the shielding. The warm ends of the striplines are connected through 50-pin D-connector ports on a large metal box (called the electronics box, or e-box) to new cables. The purpose of this box is to facilitate the transfer of signal from the inner environment, which is under vacuum, to the external environment in the RF shielded room that encloses the experiment itself.

The signals are then routed through front end boards (FEBs), one board for each ZIP. GPIB commands are sent to FEB's, which control voltage and current biasing for the ionization and charge channels, LED biasing, and any heating necessary to maintain operational temperatures. The FEBs also house second stage amplifiers for the ionization and phonon signals. Amplified signals from the FEBs are routed through twisted pair shielded cables that penetrate the RF shielding and connect to Receiver-Trigger-Filter (RTF) boards in the electronics room. The RTF filter the signals using both high-pass and low-pass filters before passing the signals on to the digitizers. The RTF boards also sum the four phonon signals, as well as the two ionization channels, and determine whether either the phonon or ionization signals are above threshold. If so, a digital pulse is sent to the Trigger Logic Board, which triggers recording of the event. Throughout run 119, the experiment was set to trigger on the phonon channels, regardless of signals in the ionization channels.

The signals from the 40 veto panels are passed directly upstairs through coaxial cables to their own discrimination and digitizer circuitry. The signals are stretched and integrated using custom circuitry and then discriminated. A coincidence of two or more veto panels with signals above threshold is used to trigger the DAQ readout, regardless of detector signal. Upon a global trigger from the TLB, ZIP and veto traces, as well as the discriminated logic signals, are recorded.

The DAQ is driven by six linux computers. The *Tower* PC is devoted to reading out signals from the digitizers for both towers of detectors. The *Vetocrate* PC serves the same purpose for the veto panels. The *Event Builder* PC gives commands for run states, assembles event information from the first two, packages the data into cdms format, and writes out data to a fiberdisk. The fourth PC, *Datasrv*, reads data from the fiberdisk, archives the raw data to tape and copies it to an analysis farm housed above ground in a secondary control room. The *Monitor* PC has two primary purposes: to simply monitor all available variables regarding the detectors, veto panels, and related fridge information, and to run the GPIB interface to control the FEB and RTF boards.

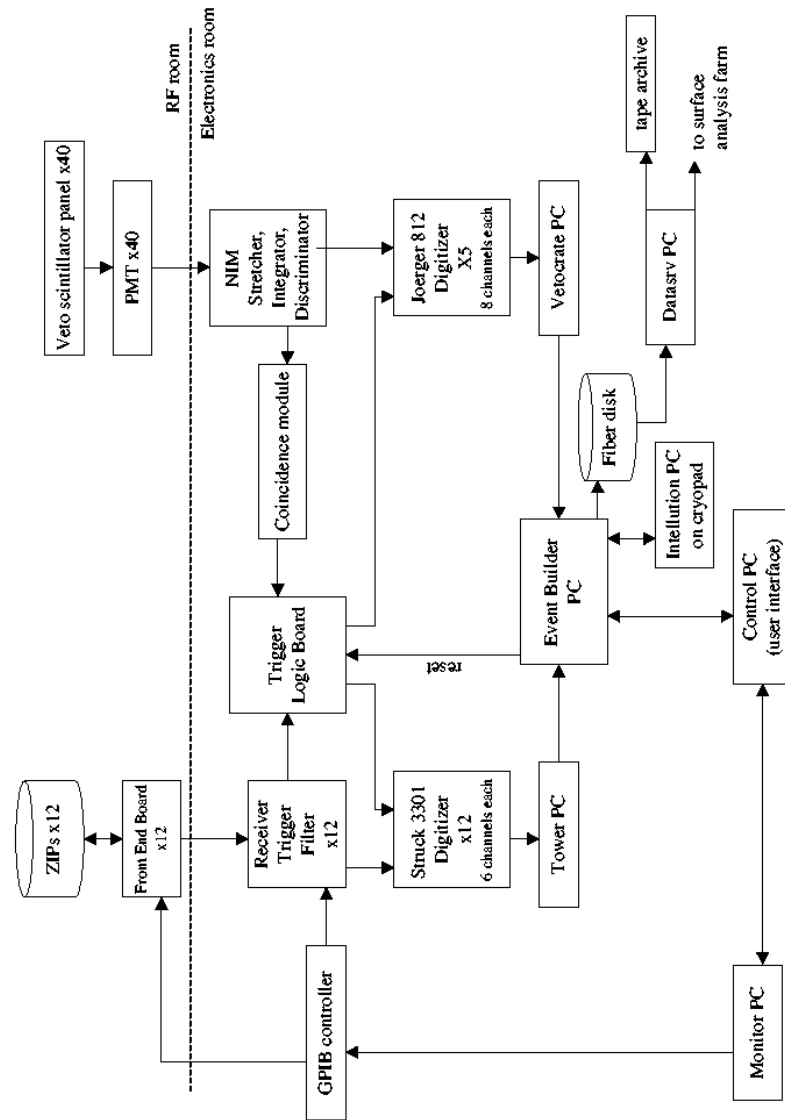


Figure 2.9: CDMS II data acquisition system, warm electronics.

The *Control* PC is the primary user machine which coordinates and monitors all java processes, and sends interlock commands to the event builder. Both *Datasrv* and *Control* are on external networks, as well as the local electronics room network, and can be accessed from the surface, making the entire data acquisition system operable from the surface building, or, indeed, to any user anywhere with a secure connection.

CDMS's analysis farm is housed in the surface building, and data is transferred there via a fast Ethernet cable strung for that purpose through the length of the mine shaft. The surface building houses 6 linux machines used for the analysis farm as well as three additional computers to remotely control both the cryogenic and DAQ systems.

2.4 Experimental Setup and Shielding




Figure 2.10: CDMS setup in the Soudan Underground Laboratory. From [25].

The CDMS II experiment is located in an RF shielded, class-1000 clean room on level 27 of the Soudan mine. This level also houses the MINOS experiment, in an adjacent tunnel. The CDMS tunnel was previously the location of the Soudan II proton decay experiment. Fig. 2.10 is a side view of the experimental setup in the cavern. To the north of the experimental setup itself, on the ground floor, is the cryopad, which houses supply dewars, plumbing, and remote control and monitoring computers for the dilution refrigerator. To the south of the RF shielded room is an ante chamber used as a work and storage room for clean hardware. It is a preparation area for items or persons that need to enter the clean room environment. Above the ante chamber is the electronics

room, where the data acquisition and hardware control systems for the ZIP detectors and the muon veto are housed. These four rooms make up the working space of the CDMS experiment.

Fig. 2.11 is a schematic of the shielding that houses the Soudan detectors inside the RF shielded room. All materials in the shielding have been screened for low radioactivity. From the outside of the experimental setup to the detectors on the inside, the layers are:

- Active muon veto. The muon veto consists of 40 panels of plastic scintillator, approximately 31" x 54" x 2" with additional overlapping regions, wrapped around the cylindrical experiment. The light signals produced in each panel are read out by one photomultiplier tube. The electronics thresholds were set individually on each PMT. The overall veto rate was roughly 600Hz, mostly due to ambient gammas. The muon rate in the veto was approximately 1 muon/hour.
- Polyethylene shielding. Approximately 50 cm of polyethylene bricks lie just inside of the muon veto. This shield was designed to moderate neutrons produced in the cavern's walls. Neutrons originating from radioactivity within the rock are sufficiently moderated so that any energy they deposit in the detectors is below threshold.
- Lead bricks. Just inside of the polyethylene is a 17 cm thick layer of lead brick. This is to reduce the gamma background produced both outside the experimental set up and in the polyethylene itself.
- Ancient lead bricks. An additional layer of ancient lead brick 5 cm thick reduces any gamma background from radioactivity in the larger lead layer. This lead was extracted from a French ship that sank in the 19th century and has low intrinsic radioactivity due to ²¹⁰Pb.
- Inner polyethylene shielding. A thin inner layer of polyethylene, only 8 cm thick, lies closest to the cryogenic cans as additional moderation for any neutrons produced in the lead layers.
- Copper. The copper cans that house the detector towers themselves make up the icebox, each associated with a cryogenic stage, but also supply approximately 0.5 cm copper shielding. There are six cans, the outermost with a radius of 28.3 cm and the innermost with a radius of 15.5 cm.

2.5 Cryogenics System

The ZIP detectors used in the CDMS experiment operate at less than 50 mK. To attain this operating temperature, we use an Oxford Kelvinox 400-S dilution refrigerator, designed to provide 400 μ W of cooling power at 100 mK. This fridge sits in the RF shielded room, but external to the shielding of the experiment, as it is not made from low radioactivity materials. Concentric stems of low activity copper thermally couple the icebox layers to corresponding thermal stages of the fridge. The outermost can is at room temperature, and the layers, proceeding inward, are at




Figure 2.11: CDMS shielding surrounding the icebox, side view and top view. From [3]

77 K (liquid N), 4 K (liquid He), 600 mK (still temperature), 50 mK (cold plate temperature), and 10 mK (mixing chamber temperature). With the heat load of the two towers, each of these cans operated slightly warmer than designed without load. For example, the inner (mixing chamber) can operated at 50 mK and the cold plate can at approximately 80 mK. The copper stems are designed such that there is no line-of-sight trajectory from the fridge to the detectors.

The fridge is connected via plumbing and cables through the north wall of the RF room to the cryopad. On the cryopad are two control and monitoring computers. One computer, a Windows NT machine, runs Intellution software and controls the plumbing external to the $^3\text{He}/^4\text{He}$ mix circulation loop, including plumbing for cryogen transfers. The software controls two external hardware units: an industrial strength control and monitoring unit (APACs), and an AC bridge LR700 (from Linear Research) with an internal multiplexer. The APACs primarily controls the interfacing between Intellution and the experiment, such as the opening and closing of valves, and the LR700 reads all thermometry data from the fridge and icebox into the computer for monitoring purposes. Additionally, the Intellution computer has two digital connections to the data acquisition system in the electronics room, allowing for computer-controlled interlock between data taking and cryogenic activity. The Intellution software features remote access as well, allowing the transfers to be controlled from a few other locations, most notably a building on the surface at Soudan that is accessible at all times. This remote control capability allows for 24-hour operations outside of the 7:30am–5:30pm access hours of the Soudan Laboratory.

The second control computer on the cryopad, an Apple G3, runs LabView software provided by Oxford with the fridge. This LabView program is essentially a copy of the control panel on the Intelligent Gas Handling (IGH) system, also provided by Oxford. The IGH has internal valves that control the path of the He mix through the circulation loop (including the dilution unit, the core of the fridge), as well as switches for the pumps associated with that loop. This software also provides monitoring of limited fridge temperatures, pressures, and cryogen levels online, accessible to all members of the experiment online. This computer additionally runs "Remote Desktop" software, which is used to control this computer, and therefore the IGH, from the surface building. can also be controlled from the surface building.

The CDMS II dilution refrigerator required two fills per day of the liquid N and He baths during run 119. The high cryogen consumption rate was due to an outer vacuum chamber (OVC) leak to the He bath. Much of the work required to understand the leak and reach a stable running state occurred before run 118. In August of 2003, the He bath was inadvertently emptied. Before filling it, we poured a small amount of liquid nitrogen into the bath. When the He bath was then filled, this N formed a solid plug at the bottom of the bath, where the leak was believed to be. As hoped, this plug partially blocked the leak and slowed the leak rate considerably, reducing the frequency of fills and unexpected warm ups. This plug was still in place in 2004, when run 119 took place. The remaining leak, however, did keep He and N boil off rates high enough to require two transfers per day. At all times, there were two dewars each of He and N connected to the fridge plumbing. The He dewars were specially ordered 350L dewars and each one lasted through a little over 3 fills

on average, or a day and a half. Each N dewar lasted for 4 transfers on average. Therefore, to get through a normal 2 day weekend, the Intellution software had to recognize when one dewar was empty and switch, most often mid-transfer, to the other dewar, which required each dewar being placed on a scale read out by Intellution. In addition to forcing transfers more often than usual, the OVC leak also required that the transfers be done at fairly low pressure, just below 10psi.

This running state was fairly stable, but every 3 to 4 weeks the fridge did experience a base temperature excursion, during which the temperature of the inner icebox would reach a couple hundred mK. These excursions usually were triggered by the softening of the OVC vacuum during transfers, which warmed the bath walls. Liquid He would boil rapidly out of the bath, and the He mix would rush out of the fridge and into the storage keg. Usually these excursions were small enough that the performance of the detectors was not affected. However, it usually required about a day to recover from these events, gradually re-introducing the mix into the fridge.

One change in the cryogenic setup between run 118 and run 119 was the addition of an external He trap into the circulation loop. In normal running, the He mix is run through a long finger shaped chamber, kept at 4K, to freeze out all materials that may have contaminated the mix. During run 118, this trap was often clogged and had to be cleaned. The cleaning process required that circulation be stopped for a short period, and data could not be taken during these times, adversely affecting the live time of the experiment. In run 119, an extra He dewar was added to the cryopad, and the mixture was run through a trap submerged in that dewar before entering the internal He trap in the fridge. The new trap was advantageous because it was larger and therefore clogged less often. Cleaning the external He cold trap was also easier since it could be bypassed through a manifold. Data taking was thus uninterrupted while cleaning this trap.

Chapter 3

Data Processing in Run 119

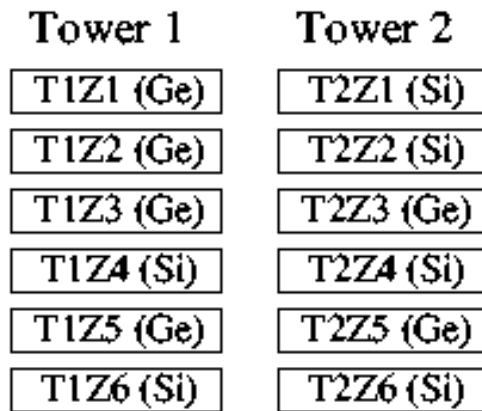


Figure 3.1: Configuration of detectors in the two towers.

The two-tower run at Soudan, run 119, lasted for approximately four months in 2004. During this time, two towers of six ZIP detectors each were operated. The distribution of Ge and Si ZIP detectors within the towers is shown in Fig. 3.1. Once signals from both towers were digitized, the DAQ system and analysis farm combined to monitor the detectors and package the data for the final analysis, making the necessary corrections in software.

3.1 Run 119 Live Time

The two-tower run, run 119, was the second completed data run at Soudan. During the first, run 118, only one tower was operational, and that tower had been extensively studied in prior data runs at the Stanford Underground Facility, a shallow site. Run 118 produced the world's most sensitive WIMP-search limit at that time. In run 119, a second tower of detectors, never before operated in a low background environment, was added. While run 118 was in effect a test of the Soudan deep site, run 119 was the test of a new tower of detectors and advanced analysis techniques. The fridge was not warmed up between runs 118 and 119, but was maintained at base temperature. The new run number, 119, reflects the turning on of Tower 2, which was mounted inside the icebox during run 118.

Run 119 lasted from March 24, 2004 to August 8, 2004. Fig. 3.2 is a plot of the accumulated WIMP-search live time versus calendar days. A total accumulation of 74.5 live days in 138 calendar days gives an efficiency of 54%.

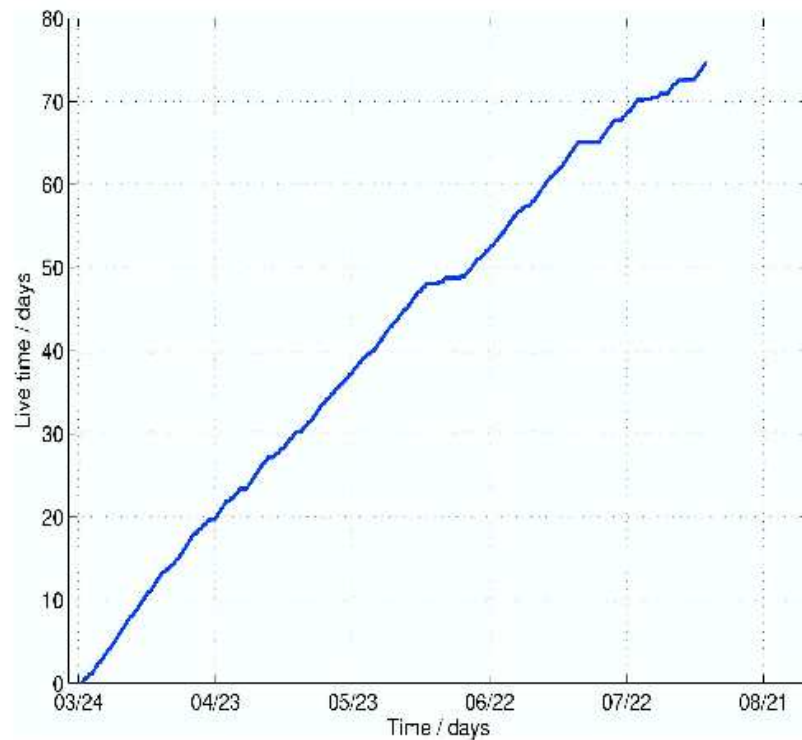


Figure 3.2: Run 119 accumulated live time vs. calendar days.

Operational reasons for decreased livetime include:

- Calibration runs. Starting in mid-April, at least 4 hours daily of Ba calibration data was taken. This consistent effort produced a total of roughly 8 million Ba calibration events over the course of the run, evenly distributed in time. In addition, there were 4 separate Cf calibration runs. In the last 2 weeks of data taking, several runs experimenting with different settings, such as detector biases, were performed, decreasing the efficiency during that time.
- Cryogen transfers. Twice a day, data taking was halted while cryogens were transferred to the fridge. During this time, the LEDs were flashed to neutralize the detectors. Each transfer was about an hour in duration.
- Baseline temperature excursions. These events, described in more detail above in the cryogenics section, occurred approximately once a month (twice a month at the end of the run) and cost about a day of live time each.
- Old air purge times. The area around the icebox, within the shield, was continually flushed with air from canisters that had aged for at least two weeks, to lower the radon levels in that cavity. When this purge was at times inadvertently stopped, the beta rate from background radon decays could rise, so data during these times, most prevalent in the last week of the run, was not included in the live time sum.

The regular use of the Ba source resulted in more than twice the number of calibration events per detector compared to run 118, with a much smoother distribution in time. This increase in statistics accommodated thorough testing of analysis techniques. Also, the WIMP-search data live time was approximately 1.5 times the live time for run 118. With the 1.5 times increase in the Ge mass, these data were expected to produce a factor of roughly 2 increase in WIMP-search sensitivity compared to run 118.

3.2 Data Monitoring

The first level of data monitoring is implemented as part of the DAQ system. The MATLAB-based analysis farm in the surface building provides several methods of data integrity checks, but these are somewhat delayed from the actual data taking. Operators receive three kinds of immediate feedback from the DAQ fast online diagnostics and monitoring system.

Run Condition Warnings

A java module running on the *Datasrv* PC collects the following information about the run conditions:

- muon veto panel rates, currents, and biasing voltages
- phonon offset voltages, rates, threshold settings

- ionization rates, threshold settings

Information on trigger rates and threshold voltages provide monitoring of the general health of the electronics system as well as some aspects of detector performance. These data are sampled once a minute, and the java module can display the latest 24 hours of data to both local and remote operators. Additionally, this java routine issues operational warning to the java process runcontrol, which displays these messages for the user. Runcontrol responds to certain warnings. For instance, if the SQUID circuitry for a detector quadrant develops a non-zero DC offset voltage, runcontrol will automatically pause data taking, reset that channel, and resume data taking.

Noise Monitoring




Figure 3.3: A typical noise spectrum for run 119, T1Z3.

At the beginning of each run, the DAQ automatically triggers 500 random events and displays the fast fourier transform noise spectrum from those events (any pulse in the events is first removed). Users could easily compare the new noise spectrum with spectra from earlier runs and

immediately address any abnormalities. Toward the end of the run, automatic noise filter checking was implemented. Fig. 3.3 is a typical, good noise spectrum from a run 119 data set in June, including the four phonon channels and two ionization channels for one detector. The roll-off at high frequencies is due to the low-pass filter implemented in the RTF boards before the signal is digitized. If one of these randomly generated events is accidentally coincident with a pulse, the low-frequency noise will increase substantially; these noise spectra do not include any above-threshold pulses. The high-frequency noise spikes were associated with electronics and pumps in the RF shielded room and did not adversely affect data quality.

Online Data Monitoring

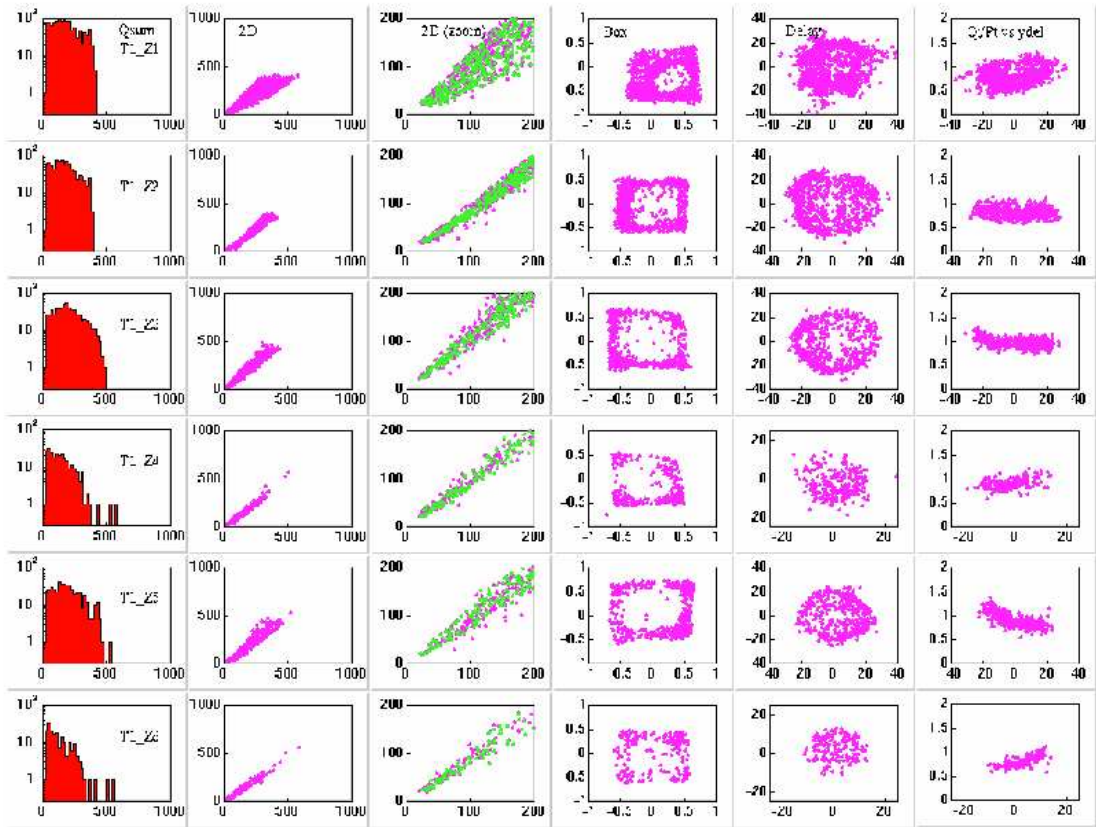


Figure 3.4: Online monitoring example for Tower 1, printed at the end of a Ba calibration run (140413_0856). Plots are, left to right, histogram of summed ionization energy, ionization vs phonon energy, the same plot zoomed in, 2 position plots, and ionization/phonon vs y-position.

A ROOT-based pulse-height analysis ran continuously on the data as it was being taken,

providing a real time evaluation of detector performance. This monitoring was rather extensive and included the following features:

- Data run information, including time started, time elapsed, total number of events recorded, how many files generated, etc.
- Real-time traces of all channels, so that users could watch pulses in each detector and veto panel, looking for under-performing channels.
- Energy and event number histograms for all detectors and veto panels.
- Analysis quantities based on crude integration of unprocessed and unfiltered data, including ionization and phonon energies and the ratio between the two, the sum of the two ionization channels, and position information using energy distributions as well as timing information among the quadrants. The position information assisted operators to set biases for the QETs and SQUIDs for each detector, to balance the quadrants so that a relatively isotropic distribution of events results in an isotropic position plot, as well as check for the success of neutralization efforts, without waiting for output from the analysis farm.

A sample of the ROOT-based analysis quantity display for Tower 1 is shown in Fig. 3.4. Events were added to the plots in real time. In this sample, there are six kinds of plots for each detector. The leftmost is a histogram of the energy summed over both ionization channels. This plot shows the Ba gamma energy spectrum, including the endpoint, and, in the Ge detectors when zoomed in, the 356 keV line that characterizes the source. The next two plots are summed ionization energy versus summed phonon energy on two different scales. If charge collection is poor in the ZIP, the slope may not be unity, or significant numbers of events may begin to “droop” out of the main population. T1Z1, the detector represented by the topmost row of plots, has a T_c gradient, so the phonon energy resolution is not as good in this detector as in the others, and it has a wider band in this plot. The next two plots are position plots, used to check that the detector response is not strongly position dependent. T1Z1, shows strangely shaped position plots because of its T_c gradient. The other detectors can be tuned to keep the quadrant responses symmetrical. The last plot is ionization yield (ionization energy / phonon energy) versus y position. If the detector is not well neutralized, events will begin to droop out of the main band at a yield of unity. The first events to droop will most likely be those at negative y position, far from the LED used to neutralize the detectors.

3.3 Event Reconstruction

The data processing done at the DAQ level underground is optimized for speed. A more thorough automated analysis is implemented on the analysis farm at the surface building. This automated analysis is MATLAB-based and involves two major steps. The first converts the raw traces to matrices of physical quantities, including timing, energy, and correlations between quantities. This

first step is called DarkPipe. This CDMS analysis package runs within MATLAB, using both MATLAB functions and C functions imported into the MATLAB environment. The second step, called PipeCleaner, performs several corrections to the physical quantities output from DarkPipe. The corrections account for some of the position variation across the detectors and also convert the electronically meaningful quantities from DarkPipe (digitizer bins, for example) into physically meaningful units (μs , for example).

DarkPipe reconstructs the event energy from the raw traces using a few different methods. The primary method uses optimal filtering in the frequency domain in both phonon and ionization channels [42]. Traces are transformed into frequency space and fit using a pulse template generated from data with well-defined, good experimental conditions and averaged over energy and position. The optimal filter algorithm weights each trace using the noise spectrum generated at the beginning of each data run, usually no more than a few hours before the event itself. Secondly, DarkPipe uses a five parameter fit to ionization pulses in the time domain to estimate energies for high energy events that exceed the range of the digitizer. DarkPipe also records the simple time-domain integrated energy for phonon pulses. Unless otherwise noted, all energy reconstruction used in the analysis are from the optimal filtering algorithm.

Most of the timing information for both ionization and phonon pulses is obtained using a simple algorithm. The raw traces are first filtered using a 2-pole Butterworth digital low pass filter. The algorithm then finds the maximum of the pulse in the time domain, defined as the 100% value. It then walks forward and back from that point, and evaluates, for instance, the time at which the pulse hits 10% of the maximum on the rising edge and the 80% on the falling edge. For ionization pulses, only the time of the 20% on the rising edge is recorded, but for phonon pulses several points on both the rising and falling edges are recorded. This algorithm begins to fail for low energy pulses, below 10 keV, because of noise. However, it is the primary method for obtaining times for the analysis. The optimal filter algorithm also returns an event time, but as it is from a template, does not provide information about the physics that affects the shape of the leading edge of the phonon pulse. This optimal filter time is used in the analysis of low energy pulses to confirm that the simple algorithm has identified the correct maximum.

DarkPipe also returns information on timing and energies in the veto, which are from the time domain. Because the pulses are from an integrator circuit, a simple pulseheight reconstruction is used. The veto panel energies are only useful in discriminating between muons and background gammas interacting in the scintillator panels.

3.4 Corrections for ZIP Non-uniformities

PipeCleaner converts the electronically meaningful quantities from DarkPipe into physically meaningful quantities (e.g. energies in keV, timing in microseconds). It also calculates a number of quantities derived from DarkPipe, including various ratios of energies and timing variables. PipeCleaner also performs what are commonly called position corrections on the data from the

ZIPs. The phonon position correction accounts for any T_c gradients in the TESs on the ZIP face. Using the high statistics Ba data, PipeCleaner smoothes out variations in pulse height and timing using a population of nearest-neighbor events to correct each individual event. The correction algorithm is explained extensively in [3]. This process affects every detector in energy resolution, but the difference is particularly marked in detectors that have large T_c gradients, including T1Z1 and T2Z1. Fig. 3.5 is an illustration of phonon versus ionization energy before and after the phonon position correction in T2Z1.




Figure 3.5: Phonon energy vs ionization energy (left) before and (right) after phonon position correction, T2Z1.

The ionization correction, while mathematically far simpler than the phonon correction, is less well-understood. Most detectors have a sinusoidally shaped variation in energy from high y to low y position in the detector, as shown in the top "before" picture in Fig. 3.6. The variation in the x -direction is less prominent. The cause of this variation is most likely electrode design, but exactly how it comes about is unclear. This correction helped us define a region of T2Z5 to be questionable. Instead of a smoothly varying energy across the detector, its behavior changes

sharply at $y \approx -20$, as shown in the bottom of Fig. 3.6. The line is not well-defined below $y = -20$, indicating poor energy resolution, so that region of T2Z5 was excluded from subsequent analysis.

3.5 ZIP Energy and Position Resolutions

There are two main measures of detector performance: energy resolution and position resolution. Energy resolution in CDMS detectors is better in the ionization channels than in the phonon channels. Fig. 3.7 is an illustration of the ^{133}Ba lines in all 6 Ge detectors from run 119. These lines are not visible in Si individual detectors, as the Si is not massive enough to stop gammas of those energies. Si spectra are compared to simulation using the endpoint as the main feature. It is these features that are used to calculate the calibration constants used in PipeCleaner. In the Ge detectors, the width of the 356 keV line ranges from 2.5 keV to 10 keV depending on the detector. In the WIMP-search data, neutron capture in Ge crystals gives a line at 10.4 keV to check energy resolution at low energies, and the width of that line ranges from .3 to .6 keV, indicating that resolution and calibration are linear in energy.

Phonon energies are actually normalized to ionization energies in PipeCleaner. That is, for the gamma events in the Ba calibration data sets, the total energy measured in the phonon channels is scaled to be equal to the energy in the ionization channels. The phonon channels are not individually calibrated to lines in the calibration data. The lines are still used to estimate energy resolution, however. In the Ge detectors, the 10.4 keV line has widths ranging from .32 keV to 1.0 keV, with T1Z1 an abnormally wide detector. The 356 line width ranges from 12.4 keV to 34.4 keV in the 5 good Ge detectors (excluding T1Z1). Energy resolution is worse in phonon channels than in ionization channels. However, it is notable that the Ge detectors in Tower 2 have on average better resolution than the Tower 1 Ge detectors.

Position resolution in the ZIP detectors is a less straightforward measurement than energy resolution, because it varies with position. Generally speaking, the position resolution is best near the centers of each quadrant, and worse at the boundaries between quadrants, especially at the center of the detector. CDMS uses several different measurements of event position, the most common of which are time delay and energy partition.

Fig. 3.8 shows a typical detector's position information in both of these constructions. The time delay based position reconstruction has units of microseconds, and is determined by the difference in time between the event in the local quadrant, where the pulse rises the fastest, and when the signal arrives in the nearest quadrant in x or y . For instance, for an event in quadrant B, the x_{delay} coordinate is the difference of the start times in B and C, and the y_{delay} is the difference in the start times in B and A. The time in D is not important for an event local to quadrant B. This position parameter is the most commonly used in CDMS analysis. Masking tests in testing facilities have shown the position resolution to be better than 1 mm at the centers of the quadrants.

An alternative position construction uses the fact the most of the energy arrives in the local quadrant, and the farther away from a quadrant an event takes place, the less energy arrives in

that quadrant. No matter which quadrant is local to the event, the energy x position is, in terms of energies, $(D + C - A - B)/totalE$ and y position is $(A + D - B - C)/totalE$. As is shown in Fig. 3.8, this is not a very realistic reconstruction of the position, as the resulting shape of the detector is rather square, and the physical detectors are cylindrical. Also, almost all events are clustered in an outer rim, with very few in the center. This method of reconstruction is mainly used as a check on how the delay quantities are behaving for individual events.

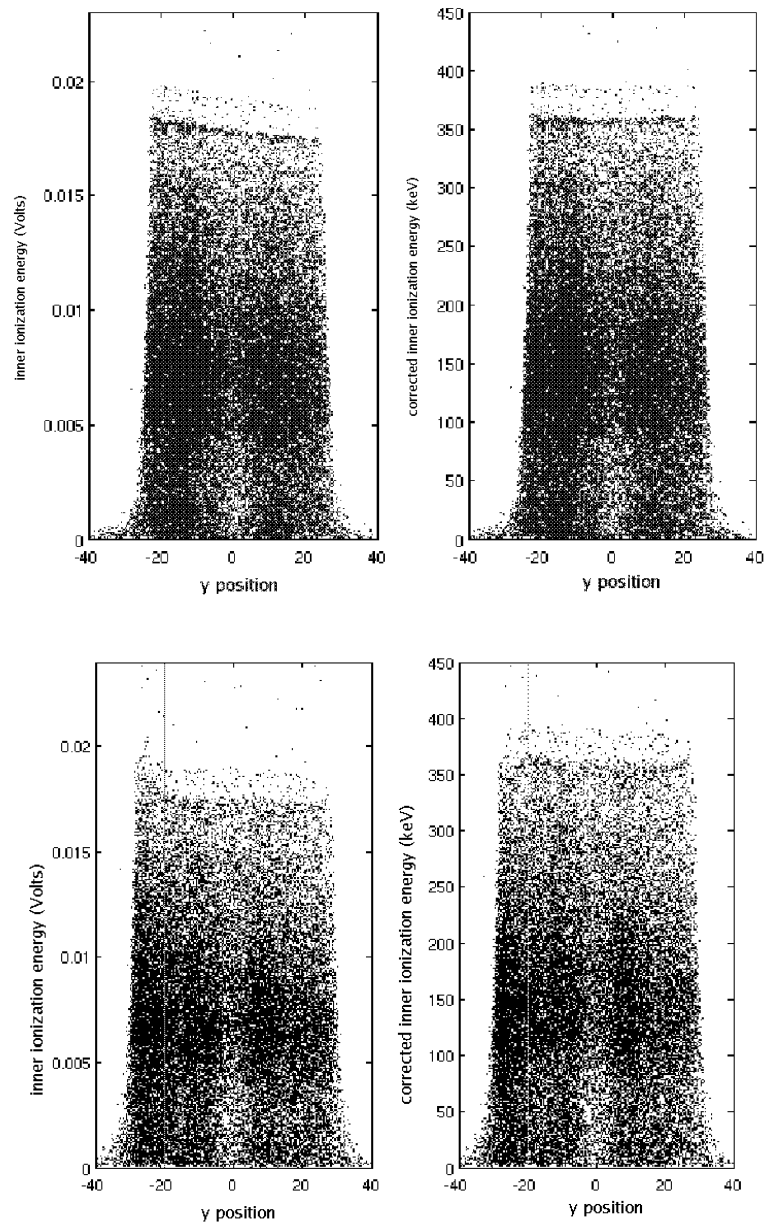


Figure 3.6: Charge energy vs y position for (upper) T2Z3 and (lower) T2Z5. Plots on the left are of the uncorrected quantity before correction, in Volts, and on the right are of the corrected quantity after correction and scaling to keV. The smooth variation in T2Z3 is typical of detectors. The unusual region in T2Z5 at $y < -20$, marked with a vertical line, is excluded from further analysis.




Figure 3.7: Spectra of ionization energy in all 6 Ge detectors of run 119. These lines are not visible in Si detectors. Best resolution is in T1Z5 and T2Z3. T2Z5 has poor resolution due to the area of bad ionization collection.




Figure 3.8: Position plots for detector T2Z3. Position using timing information is left, using energy distribution is right.

Chapter 4

Analysis: Backgrounds and Rejection

In order to be sensitive to WIMPs, CDMS must have a low background. This chapter will detail the known backgrounds in the two-tower run and the hardware and analysis techniques employed to reject them. The primary backgrounds are gammas, neutrons, and surface events, or betas, which dominate the expected background after cuts.

Cuts designed to reject these backgrounds define the CDMS signal region, which is bounded by several different quantities. The WIMP-search data (not calibration data) was masked at the beginning of the analysis using several cuts. Any event that was a single scatter, had a low ionization yield (ionization energy to phonon energy ratio), was veto anti-coincident, and occurred inside the fiducial volume of its detector was masked from use by both an automatic routine (which did not function well) and users themselves (who functioned considerably better). Events that failed any one of these loose selection criteria were known from the outset to be background and outside the signal region. This precaution allowed us to do our analysis without biasing ourselves, and, at the same time, look at data known to be background. A tighter, more accurate definition of the signal region was developed using the techniques described in this chapter, calibration data, and WIMP-search data far from the signal region.

Gammas are the most common background in CDMS, occurring at a rate of a few hundred events per kg-day in most detectors. The ZIP detectors were designed to reject gammas, which produce electron recoils in the crystals rather than nuclear recoils, as WIMPs would do. Analysis techniques reject gammas with very high efficiency. Gammas come from many sources in the Soudan Underground Laboratory, including radioactive decays in the rock surrounding the cavern, materials in the cavern, and the CDMS shield itself.

Neutrons are rare events in the ZIP detectors, but nevertheless a subject of intense study because they cause nuclear recoils in the crystals. They therefore closely mimic the WIMP signals.

When Tower 1 operated at the shallow site at the Stanford Underground Facility, cosmogenic neutrons penetrated the scintillating veto undetected, created dozens of nuclear recoils in the ZIPs, and were therefore the dominant background after analysis cuts. At the deep site in Soudan, where neutrons are no longer a dominant background, the primary neutron source is still cosmogenic neutrons interacting either in the rock surrounding the cavern or in the shield materials. Neutrons from natural radioactivity in the rocks are sufficiently moderated by the polyethelene in the shield that they are not a measurable background.

Surface events, most often caused by betas, are currently the dominant background in CDMS after analysis cuts. Beta sources include gammas scattering in material near the ZIPs and knocking off a beta, and surface contamination, either Rn or other radioactive elements, on the ZIPs themselves. Although beta events occur at a far lower rate than gammas, approximately one event per day in each detector, they are more challenging to reject. They interact near the surface of the detector where the charge collection is poor. The resulting electron recoils appear like nuclear recoils because of the reduced ionization collection. Analysis techniques incorporating phonon timing information have been developed to reject these surface events, but this rejection is not as efficient as either neutron or gamma rejection, so betas remain the dominant background after cuts.

4.1 Gammas and Band Definitions

Although the lead surrounding the detectors reduces the gamma rate by factor on the order of 2×10^4 , additional gamma rejection is achieved through analysis techniques. Most current and future direct WIMP searches use the fact that gammas scatter off electrons in the crystal to tag and reject them. We define the ionization yield, or the ratio of ionization to phonon signals, to be 1 for an electron recoil. The yield is most accurately described as the ratio of the total energy of the event as inferred from the ionization signal to the total energy as inferred from the phonon signal, assuming the event is a gamma. It is *not*, as might be implied by the name, the fraction of the total energy that is deposited through ionization. It is the ratio of two different measurements of the total energy of the event, normalized to one for gammas. When neutrons and WIMPs recoil off the nuclei in the crystals, less ionization energy is measured, giving a mismatch in the total energy measured through the ionization and in phonon channels and therefore a lower ionization yield. The assumption that the event is a gamma has failed, and the event is classified as a nuclear recoil.

Fig. 4.1 shows the distribution of gammas and neutrons in ionization yield vs. recoil energy. The gammas shown are from a ^{133}Ba source, and the neutrons from a ^{252}Cf source. The band at a yield of unity is the electron-recoil band, where most of the Ba data lie. The neutrons from Cf form the nuclear-recoil band near a yield of 0.3. The Cf source produces gammas as well as neutrons, but those data points are plotted under the Ba data and therefore cannot be seen in the figure. The data points well outside these two bands, scattered throughout the plot, are surface




Figure 4.1: Definition of bands in detector T2Z3. The lower band is the nuclear recoil (NR) band and the upper one is the electron recoil (ER) band. Cf data is used to define the NR band with a $\pm 2\text{-}\sigma$ width around the mean, and Ba data is similarly used to define the $2\text{-}\sigma$ ER band.

events, which are caused by betas incident on the surface of the detector. In calibration data, most of these betas are products of Compton scatters of the gammas from the source near the detectors. They are said to be “drooping” from the electron-recoil band because their low yields are due to poor ionization collection rather than to the physics of the recoil.

For the purpose of the CDMS analysis, background particles are referred to in terms of the populations shown in Fig. 4.1: “gammas” populate the ER band, “neutrons” populate the NR band, and “betas” are outside of these bands. In most cases, these labels agree with the standard physical particles to which they refer. Gammas in the analysis correspond to photon-triggered events most of the time. However, there are exceptions to this rule, and it must be remembered that the labeling system of the analysis may include other incident particles which fall inside the same yield band. For instance, $\sim 20\%$ of surface electron events are classified as gammas in the

analysis because they lie in the electron recoil band. Any gamma that falls outside of electron-recoil band will be called a beta in the analysis. Events from Cf that lie in the nuclear recoil band are called neutrons; however, it is possible in Cf data as well as in Ba that a few betas may coincidentally lie in the nuclear recoil band.

To define the bands, data were binned in bins of 10–40 keV and the yield distributions were fitted with a Gaussian function. The fitted Gaussian means vs energy were then fitted to the functional form

$$yield = aE_r^{b-1} \quad (4.1)$$

and the standard deviations with the form

$$\sigma E_r = c(aE_r^b) + d, \quad (4.2)$$

where E_r is the recoil energy, and a, b, c, and d are parameters of the fit[25]. Typical fit values were in the range of a=1, b=1, c=0.04 and d=0.5 for the electron recoil band, and a=0.1, b=1, c=0.1, and 0.3 for the nuclear recoil band. We use 2 sigma bands centered around the fitted means to define a gamma (using the electron recoil band) or the signal region (using the nuclear recoil band).

The nuclear recoil band fits suffer from low statistics in the highest energy bins (60-100 keV), but statistics for most other bins are good. This problem of low statistics for neutrons at higher energies is not one that can be corrected by simply increasing statistics in the Cf data. Simulations and experience have verified that inserting the Cf source beyond the shielding for long periods of time risks long term activation of the material surrounding the detectors, especially the copper in the cryogenic cans. For this reason, Cf calibrations were limited to 2 hours each. Fig. 4.2 shows the pattern of elevated low yield event rate after each Cf calibration run; data runs with these elevated rates were excluded from the final WIMP-search data set. Activation was carefully monitored, and Cf runs were limited to two hours each. To compensate for low statistics in the highest energy bin, the data from all Ge detectors in this bin were combined in the Gaussian fit for those energy bins. Only T1Z3 was excluded from the sum because the mean of its nuclear recoil band had qualitatively different behavior for both neutrons and gammas. Once the Gaussian fit parameters were obtained for the highest energy bin, the bands were fit on a detector-by-detector basis using the functional forms above.

The functional forms of yield versus energy come from Lindhard theory, which is a calculation of the ideal ionization energy to total energy response for WIMPs recoiling within the crystals [43, 44]. Fig. 4.3 is a comparison of the means of the nuclear recoil bands in our detectors and Lindhard theory. Note that the measured ionization yield agrees with theory at low energies and is higher than theory predicts for larger recoil energies. Part of that difference is that our calibration is based on neutron data rather than WIMP data. Neutrons with kinetic energies above a few hundred keV can inelastically scatter within the detector, producing gammas; WIMPs would not scatter inelastically. A typical kinetic energy for Cf neutrons is 1 MeV, at which energy inelastic

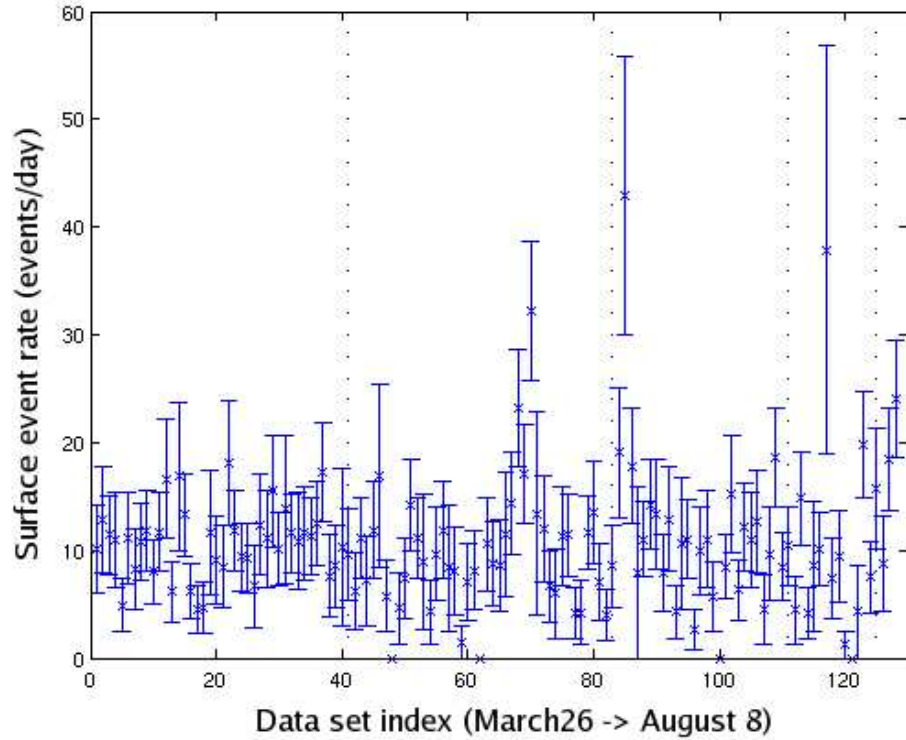


Figure 4.2: Surface event rate (events/day) vs dataset index for 5 summed Ge detectors (T1Z2, T1Z3, T1Z5, T2Z3, and T2Z5). The four ^{252}Cf calibrations are shown as black dotted lines. The rise in rates near index 70 corresponds to a decrease in old air purge. Dataset numbers correspond to late March 26 at 0 and August 8 at 128. The dates of the calibrations are May 7, June 15, July 12, and August 2.

scatters account for nearly 50% of the total scatters [45]. The gammas raise the measured ionization yield of these events. This scattering process therefore contributes to the error in the nuclear recoil selection efficiency, as discussed in Chapter 5. An additional reason for departures from theory is the detector performance with respect to tuning, thresholds, and noise performance. The similarity among most detectors, both Ge and Si, indicate that this is a small effect, the exception being T1Z3, which has a nuclear-recoil band mean significantly higher than the other Ge detectors at low energies. This detector has a similarly shaped electron-recoil band, indicating a difference in overall detector performance. Nevertheless, it should be noted that in all detectors the differences between the data means and Lindhard theory at recoil energies greater than 15 keV are smaller than the 1σ statistical uncertainty in the bands themselves. Below 15 keV, only T1Z3 differs statistically from theory.




Figure 4.3: (a) A comparison of nuclear recoil band means in the 5 good Ge detectors with Lindhard theory. (b) A comparison of the nuclear recoil band means in the 4 good Si detectors.





Figure 4.4: (a) Nuclear recoil band efficiencies from neutrons in Cf and (b) Electron recoil band efficiencies from gammas in Ba, shown in the 5 good Ge detectors. Error bars are added for the representative detector T2Z3.

The performance of the bands can be measured using selection efficiencies in the calibration data, i.e. the fraction of gammas (neutrons) from the Ba (Cf) source that lie in the 2σ electron (nuclear) recoil band. Fig. 4.4 shows the efficiencies for gammas and neutrons in the Ge detectors.

The most important efficiencies for our results are the nuclear recoil acceptance efficiencies. However, gamma acceptance efficiencies are useful in detector characterization and identification of any systematic differences between calibration and WIMP search data, since only they can be calculated in both datasets. These data suggest that there is a small systematic difference in the means of the gamma bands between Ba data and WIMP data; the WIMP search data tend to have slightly lower gamma bands. This is most likely due to the rate, and in the future analysis simulations will be used to try to quantify by how much, if any, the nuclear recoil bands should be extended downward due to this systematic difference as well as any systematic differences between neutrons and WIMPs. Also, neither the electron recoil nor the nuclear recoil bands are perfectly Gaussian; however, this should matter little due to the functional energy smoothing done.

Despite all these possible difficulties, the electron and nuclear bands and their efficiencies in all data sets are still quite good and relatively robust. They remain our primary analysis discrimination technique to reject the gamma background. Selecting the 2σ nuclear recoil band as part of the signal region rejects over 99.9% of all gammas. We do not calculate expected leakage of gammas into our signal region, because any gamma that droops out of the electron recoil band and into the signal region is labeled as a beta surface event (see Section 4.3) in our analysis of both the calibration data and the WIMP search data, and it would therefore be included in the beta background estimate.

4.2 Neutrons and the Veto

The cosmogenic neutron background was the primary motivator for CDMS II for going to a deep site. At Soudan, the detectors are shielded by 780m of rock, or 2090 meters water equivalent. As a result, the muon flux is reduced by a factor of 50,000 from the surface rate. Or, to put it in more useful terms, from one muon per second passing through an outstretched human hand to roughly one and a half muons per day. Muons can produce neutrons both inside our shield (these are easily vetoed) and outside our shield, in the Soudan rock (these are more difficult to veto).

Another source of neutrons is radioactivity in the rock surrounding the experimental cavern. Soudan rock contains trace elements of U/Th which produce neutrons in their decay chains. Simulations show, however, that in order to penetrate the 50 cm of polyethylene surrounding our detectors and produce an above-threshold recoil in a ZIP, a neutron must have at least 50 MeV kinetic energy. The ambient neutrons in Soudan from the rock have a maximum kinetic energy of ~ 6 MeV and thus do not contribute to our expected neutron background.

Besides passive shielding to reduce neutron backgrounds, CDMS also uses some analysis techniques to reject neutron-generated nuclear recoil events. For example, WIMPs are far less likely than neutrons to multiply scatter among ZIPs, so we define a “singles cut” by selecting only events

that have above-threshold phonon energy in exactly one ZIP. For this cut, the threshold is taken to be six sigma above the total phonon noise signal (summed over the four phonon channels). The relevant efficiency for this cut is 1-P, where P is the fraction of single events we mistakenly identify as multiples because of above threshold noise fluctuations in the detectors. The efficiency is therefore energy and detector independent, and is 99.98%. Another neutron rejection cut depends on the fact that the spin-independent WIMP-nucleon cross section is 6 times larger in Ge than in Si; the cross-sections in Ge and Si are comparable for neutrons. We simply choose the Ge ZIPs as our primary WIMP detectors. For spin dependent assumptions and at very low WIMP masses, Si is as sensitive as Ge. However, at first pass, the Si primarily serves as an multiple-detector event veto for the Ge in our main result.

4.2.1 Veto data analysis

The most effective means we have to reject neutrons is tagging using the muon veto. Statistically, neutrons produced in the shield by muons are far more numerous, by more than an order of magnitude, than so-called external, or punch-through, neutrons produced outside the shield that penetrate to the ZIPs. The veto cut rejects any event in the ZIPs that occurs within $50\ \mu\text{s}$ after above-threshold activity in the scintillating veto. At 600 Hz, most of this activity in the veto is caused by ambient gammas, not muons.

The muon rate in the detectors, $\sim 4\text{--}6$ muons per day for run 119, is sufficiently low to make veto cut definitions using WIMP search data at Soudan difficult. We instead rely on data from the shallow site at Stanford, where the rate was orders of magnitude higher, to define the veto cut. These data show that a dead time of $40\ \mu\text{s}$ was sufficient to veto 99.95% of the muons that hit the detectors. In fact, roughly $15\ \mu\text{s}$ was sufficient for muon coincident events in all detectors except T1Z1, which has a significant TES gradient and slow reaction time to recoils. The $50\ \mu\text{s}$ cut is therefore very conservative. This cut tags all muon-like events in the ZIPs as veto-coincident.

In the current run, an additional detail was added to the definition of the veto cut. Muons, as minimum ionizing particles, deposit more energy in the veto panels as they pass through than do the ambient gammas. Fig. 4.5 shows a typical veto panel spectrum. The background at low energies, exponentially rising toward zero, is the gamma spectrum. The wide peak at higher energies, above approximately 100 pC, is the muon spectrum. The new veto cut takes advantage of the good separation to veto a longer period of time when the event in the scintillator looks like a muon, specifically $185\ \mu\text{s}$ [46]. However, the cut on gamma activity is retained because simulations show that such activity can be caused by cosmogenic showers without the parent muon and can serve to veto external neutrons that might otherwise be missed. Therefore, for a gamma-like event in the scintillator, the vetoed time was still $50\ \mu\text{s}$. The addition of the muon-specific cut had little effect on the overall live time of the experiment because the muon rate was so low (a few a day). The relevant efficiency for this cut in the signal region is 1-P, where P is the fraction of non-muon related events that are mistakenly labeled as muon coincident due to ambient activity in the veto. Taking into account the 600Hz veto rate (almost all gamma activity) and the 50us vetoed time




Figure 4.5: A typical veto scintillator spectrum. The peak just above threshold is ambient gammas; the second, wider peak is due to muons. The excess at extreme right is a saturation effect in the digitizer.

associated with each event, the efficiency of this cut is $96.97\% \pm .01\%$. This efficiency is detector and energy independent.

In the first run at Soudan, no signal candidates were rejected by this cut; that is, we observed no neutrons in the detector coincident with muons in the veto. In run 119, there is one such event at the bottom of Tower 2: a double shared by T2Z5 and T2Z6. This is roughly consistent with expectations from simulations. There are no observed muon anti-coincident multiples in the data set. We use these two results plus results from simulations to estimate the number of neutrons expected to leak into the signal region, which includes only muon anti-coincident single events.

4.2.2 Neutron background simulations

As noted, we need an estimate of the efficiency of the scintillating veto in tagging external neutrons. The cross section of neutrons in plastic scintillator is not very large, so the efficiency of directly tagging the neutrons with the veto is $<30\%$. However, since most neutrons are accompanied by particles from the same cosmogenic shower, and these daughter particles are more likely to trigger the veto, the tagging efficiency is much better.

Several simulations were performed in order to estimate the fraction of external neutrons that are vetoed. FLUKA, Geant4 based on that FLUKA, and an independent Geant4 simulation were all done. In both FLUKA and Geant4, muons from the Soudan spectrum were propagated toward

the experimental apparatus through a minimum of 10 m of Soudan rock. The resulting showers were traced through to the Soudan veto. Different simulations, all in Geant4, tracked those showers through the shielding around the detectors to the ZIPs themselves. At Minnesota, the study was done using Geant4 for both stages of the simulation.

To begin an event in the UMN simulation, a muon is randomly generated in a plane 10 m above the experimental cavern and propagated through rock. Muons and their associated showers that reach the cavern are recorded. Details regarding the codes used, the geometry details, and particle energy spectra and angular distributions can all be found in Appendix B. Neutrons are the shower particles of interest, because they are the background particles that are difficult to distinguish from WIMPs. Fig. 4.6 shows the spectrum of neutrons as they enter the cavern. These neutrons have been produced by muons in the rock surrounding the cavern in the simulation. The average energy is 45.33 MeV. Neutrons with energies below approximately 50 MeV are not energetic enough to penetrate CDMS shielding and have above-threshold recoil energy in the ZIPs. It is therefore the higher energy neutrons that are the background for the experiment.




Figure 4.6: Log-log neutron spectrum entering the cavern, as obtained from Geant4 simulation. The energies of 19,577 neutrons are plotted; average energy is 45.33 MeV.

The Geant4 geometry for the CDMS experimental setup, from the outer veto panels to the innermost polyethelene shielding, is centered at (0, 2 m, -0.4 m) in that cavern. The details of the icebox and detector towers are not necessary to do basic neutron veto studies, and so they were not included in the simulation. Any shower particles that deposit energy in either the outer veto or the inner polyethelene (there are several layers of lead and poly between the two) are recorded. Any

neutron in either volume is recorded whether or not it directly deposits energy. This arrangement makes calculations of veto fractions for neutrons, both internal and external, straightforward to calculate. It also allows for the full propagation of events from the original cosmogenic muon to the area surrounding the detectors in one simulation. Vectors containing energy, position, momentum and timing information are written out to a MATLAB file for each step.

The wholistic, large-scale nature of this simulation limits statistics. For 8M muons thrown, only 270 neutrons hit the scintillating veto with energies above 50 MeV. Analysis of the MATLAB simulation output indicates that 46% of the neutrons that hit the veto are vetoed by a muon hitting the veto within the hardware timing window. If all shower particles' energy depositions in the veto panels are added, the veto fraction of neutrons in the outer veto increases to $78\% \pm 6\%$.

The FLUKA simulation produces a smaller overall veto fraction of $73\% \pm 2\%$ [51]. These veto fractions are consistent with one another. The small differences in these numbers can be partially attributed to calculation method and partially to differences in the physics files in FLUKA and Geant4. The Geant4 simulation uses physics files and the scintillator in the geometry to calculate energies deposited in the veto. The FLUKA simulation, lacking the CDMS setup geometry in the cavern itself, used particle energies and theory to calculate how much energy would likely be deposited in the veto. Also, the neutron production differs slightly between the two physics codes such that Geant4 produces fewer neutrons at high energies than does FLUKA. Work is continuing to advance both simulation methods, but they are consistent and give a reasonably good measure of the veto fraction of external neutrons.

Taking the differences in the simulations into account, we use 70% for the fraction of external neutrons that are vetoed either by their own interaction or by their associated shower particles. This was the veto fraction agreed upon by the collaboration to be used in expected background calculations and was chosen to give a conservative (slightly high) estimate. Nearly all (>99%) internal neutrons are vetoed. Using the facts that there are more than $10\times$ more internals than externals in simulations, and only one vetoed neutron event was observed in the data, 0.03 anti-coincident neutrons are expected in the WIMP-search data. Using another method based on the simulations and accumulated live time alone, we estimate a slightly higher neutron background: 0.05 events in our signal region in Ge and 0.05 events in Si. We use the higher estimate, but the two are consistent.

4.3 Betas and Surface Event Rejection

Although neutrons were the dominant background for CDMS at the shallow Stanford site, betas are the dominant background at Soudan. While betas cause electron recoils and might therefore be expected to be rejected along with the gammas, those from external sources, such as surface contamination, typically do not penetrate into the detector and instead recoil in the first $40\ \mu\text{m}$ of the surface of the detector. This surface layer is an area of poor charge collection, a so-called dead layer. Because the ionization signal is suppressed, so is the ionization yield, which is used

to discriminate between electron recoils and nuclear recoils. The betas droop out of the electron recoil band and sparsely populate the region of lower yield, including the two sigma nuclear recoil band which defines the signal region.

4.3.1 Defining Analysis Cuts

Historically, the first discrimination parameters were based on phonon signal timing, and these parameters remain central to the surface event analysis. Fig. 4.7 shows the definition of these two parameters. The first is phonon delay, the time between the fast ionization signal and the slower phonon signal. In each case, the start time of the pulse is defined as the time at which the signal is at 20% of the maximum. The relevant phonon channel is the local event quadrant (for the pulse in Fig. 4.7, it is quadrant B). The second parameter is the phonon risetime, with is a measure of how quickly the phonon signal in the event quadrant rises. The time difference between the time at 10% of the max to the time at 40% of the max is used.

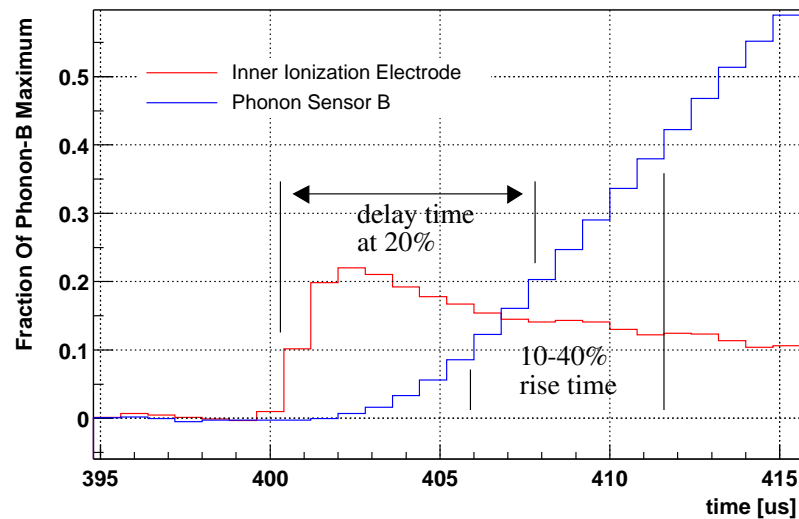


Figure 4.7: Pictorial definition of phonon delay and phonon risetime. The y-axis is the fraction of the maximum of the phonon pulse (.1=10%). The faster-rising pulse is from the ionization channel; larger pulse is from phonon quadrant B.

Fig. 4.8 and Fig. 4.9 show the discrimination power of the phonon delay. The neutrons and gammas are well localized in yield, in the nuclear recoil and electron recoil bands, respectively. The betas form a population spread between these two regions, but mainly at low phonon delay. Physically, the proximity of the event to the surface of the detector shortens the time required for

the phonons to become ballistic, and therefore speeds the response of the local phonon channel. Note that it is the proximity of the event to the surface, not the proximity of the event to the phonon sensors, that matters. Calibrations performed at the test facility at Berkeley demonstrated that betas incident on the charge electrode side of the detector, opposite the phonon sensors, also show a faster phonon response time. Similarly, betas generally have faster phonon risetimes than do bulk events, especially bulk nuclear recoils.




Figure 4.8: Plot of phonon delay vs. ionization yield in Tower 1. Dark dots are gammas from Ba, x's are betas from Ba, and grey dots are neutrons from Cf.

The timing cut relies on only these two timing parameters. In the previous run's analysis, the timing cut was dependent on the same two parameters, and was an energy dependent cut. For




Figure 4.9: Plot of phonon delay vs. ionization yield in Tower 2. Dark dots are gammas from Ba, x's are betas from Ba, and grey dots are neutrons from Cf.

this analysis, it was decided that the cut should be energy independent for two reasons: first, that the parameters themselves show no energy dependence, i.e. the means of the parameters are flat in energy and the widths are constant except for a noise-induced widening at energies less than 20 keV. Secondly, binning in energy limits how strict the cut can be, since each placement of each cut sets an expected background in that bin, so the expected background rises with the number of energy bins. Fig. 4.10 shows the distribution of phonon delay versus energy for a typical detector using ~ 3.5 M events from Ba calibration data. To begin, a fit was done to the upper edge of the beta distribution for each parameter in each detector. The fit was of the form $C + \sqrt{A + \frac{B}{E^2}}$, a dispersion relation where E is the energy of the recoil. Each event was then assigned a value corresponding to how far away from that edge the event lay, i.e. the value D in the equation

$$TIME = C + D \left(\sqrt{A + \frac{B}{E^2}} \right) \quad (4.3)$$

where A, B, and C are fixed from the fit to the upper edge. Events on the upper edge of the beta distribution therefore have a value of 1, outliers above it greater than one, and events below the mean a negative value of D. The result is two parameters, the energy-adjusted phonon delay and risetime, which have constant beta outliers in energy.

The timing cut was then defined on these two energy-adjusted parameters. Fig. 4.11 and Fig. 4.12 are two dimensional plots showing the distribution of neutrons and betas in these parameters. The timing cut is actually a combination of two cuts designed to accept neutrons and reject betas. First, a loose cut on the difference of the two energy-adjusted timing parameters was performed: it is a data pre-selection cut accepting the four sigma width of the neutron distribution, shown as black lines. Events outside of this cut have uncorrelated timing parameters and are therefore outside of the signal region. Secondly, the main timing cut is set as a value on the sum of energy adjusted phonon delay and risetime. This second timing cut is also shown in Fig. 4.11 and Fig. 4.12, as a black line perpendicular to the other two.

The actual value of the second cut was set on a detector-by-detector basis and tuned to pass the desired number of total beta events. We needed to determine the optimal number of betas for the cut to accept that still allowed the cut to accept as many neutrons as possible. Determining this number was a challenge, since the signal region in the WIMP-search data was masked, and we therefore did not know the number of events in the region before the timing cut. Using data from the previous run's detectors T1Z2, T1Z3, and T1Z5, as well as the kg-d exposure in both runs, we estimated that there would be roughly a factor of 100 difference between the number of passing betas in our Ba data set and the number of passing background betas in the WIMP-search signal region. In setting the timing cut, we therefore allowed 4 or 5 passing beta events in each detector, expecting a beta background in the WIMP-search data in the neighborhood of 0.3 events. The total number of passing betas, in detectors T1Z2, T1Z3, T1Z5, T2Z3, and T2Z5 combined, was 26 events.

The timing cut was tested by applying it to the other half of the Ba data, which had been

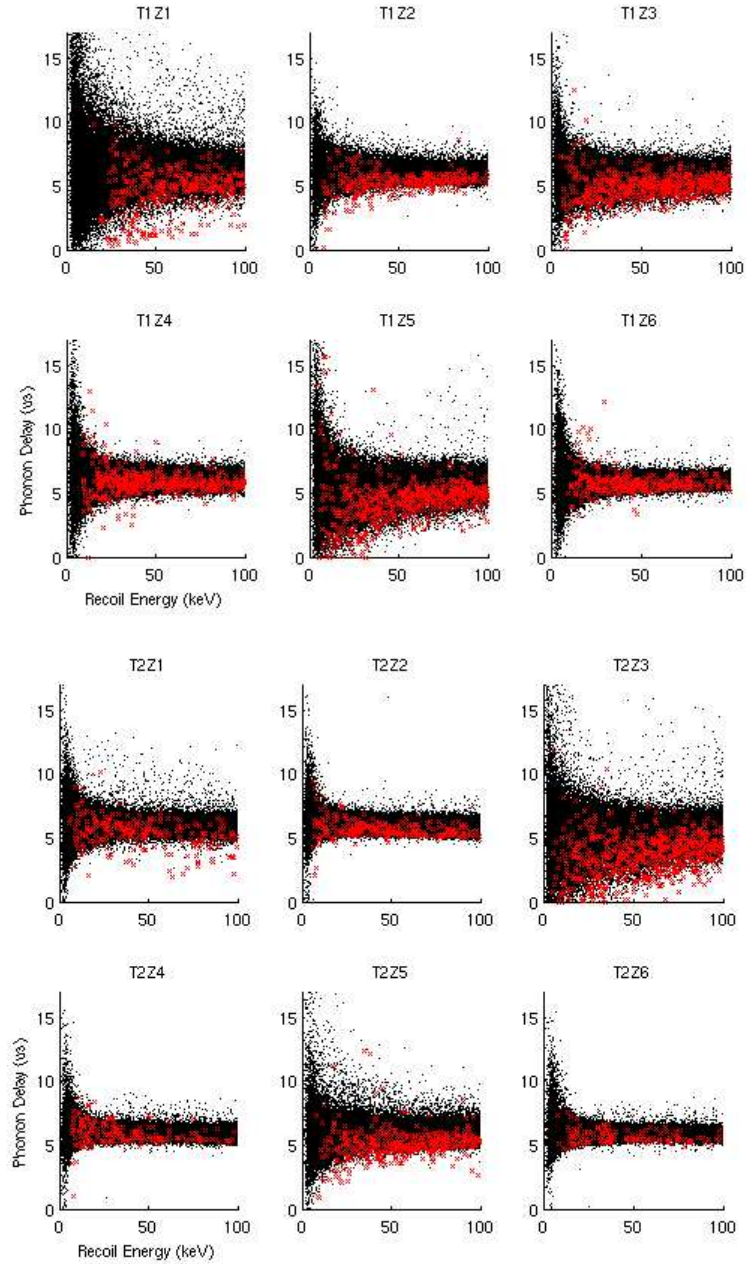


Figure 4.10: Phonon delay (μs) vs. recoil energy (keV) for towers 1 and 2. All data is from Ba calibrations. Dots are gammas, x's are betas. Note that the mean of the distribution does not vary with energy, but the width does.




Figure 4.11: Plot of energy adjusted phonon risetime (μs) vs. phonon delay (μs) in Tower 1. Dots are neutrons from Cf, and x's are betas from Ba. Location of timing cuts are shown in black solid lines. Events that pass the cuts are in the upper right region where most neutrons lie.




Figure 4.12: Plot of energy-adjusted phonon risetime (μs) vs. phonon delay (μs) in Tower 2. Dots are neutrons from Cf, and x's are betas from Ba. Location of timing cuts are shown in black solid lines. Events that pass the cuts are in the upper right region where most neutrons lie.

unexamined for the definition of the cut. The application of the cut passed 26 beta events, as predicted. However, the distribution of these events in the second half of the Ba was not as smooth as in the first half. Most notably, fewer betas passed in Tower 2 than expected, and T1Z5 in particular had excess passing betas. Table 4.1 shows the fraction of passing beta events in both Ba data sets in all detectors. The systematic differences between these two data sets seem to be due to systematic shifts in behavior of the detectors on a timescale of a few hours, which was the time scale with which the two halves of the Ba data were divided. After this was discovered, the Ba data was again split up on an event-by-event basis, every other event, and no such systematic difference between the two data sets was observed [52]. Because the WIMP search data itself was taken a few hours at a time, the systematic difference between the original two halves of the data is a reasonable measure of our systematic uncertainty. This is discussed in more detail in Chapter 5.

Table 4.1: Fraction of betas passing the timing cut in three data sets: first half of Ba data (half used to define the cuts), second half of Ba data (half used to test cuts), and WIMP search multiples. A beta by definition has an ionization yield between 0.1 and the bottom of the 5 sigma electron recoil band.

Detector	First half Ba	Second half Ba	WIMP search multiples
T1Z2 (Ge)	$\frac{4}{294} = 0.014$	$\frac{4}{336} = 0.013$	$\frac{3}{68} = 0.044$
T1Z3 (Ge)	$\frac{4}{551} = 0.007$	$\frac{8}{615} = 0.013$	$\frac{0}{98} = 0$
T1Z4 (Si)	$\frac{9}{444} = 0.020$	$\frac{5}{138} = 0.013$	$\frac{1}{71} = 0.014$
T1Z5 (Ge)	$\frac{5}{603} = 0.008$	$\frac{12}{639} = 0.019$	$\frac{1}{148} = 0.007$
T1Z6 (Si)	$\frac{4}{302} = 0.019$	$\frac{4}{318} = 0.013$	$\frac{5}{154} = 0.033$
T2Z1 (Si)	$\frac{6}{242} = 0.025$	$\frac{4}{221} = 0.023$	$\frac{11}{56} = 0.196$
T2Z2 (Si)	$\frac{4}{269} = 0.015$	$\frac{8}{231} = 0.035$	$\frac{1}{60} = 0.017$
T2Z3 (Ge)	$\frac{5}{656} = 0.008$	$\frac{2}{697} = 0.003$	$\frac{3}{161} = 0.019$
T2Z4 (Si)	$\frac{4}{176} = 0.023$	$\frac{4}{138} = 0.029$	$\frac{0}{56} = 0$
T2Z5 (Ge)	$\frac{4}{409} = 0.010$	$\frac{0}{352} = 0$	$\frac{7}{146} = 0.048$
T2Z6 (Si)	$\frac{4}{133} = 0.030$	$\frac{1}{140} = 0.007$	$\frac{3}{59} = 0.051$

A second, and perhaps more important, test of the timing cut was its application to multiples in the WIMP search data. Since multiple scatters are by definition outside the signal region, they were accessible while the signal region was still masked. We expect that multiples in the WIMP search data are from the same source as the single scatter background, so they make a good test population for the timing cut. Table 4.1 also shows the fraction of passing betas in all detectors in the WIMP search multiples. The result is markedly different for this population than the Ba.

The means of the beta population's timing parameters in each detector shift between Ba data and WIMP search data. In the detectors in which the mean shifts negatively, fewer betas pass the timing cuts. In detectors in which the mean shifts positively, the largest such shift being in T2Z5, more betas pass the cut than are expected.

Such a shift in mean is not entirely unexpected. The timing parameters vary with position within the ZIP and somewhat with energy as well. The betas in the Ba data are produced by gamma Compton scatters in the vicinity of the detectors, and those gammas are from a point source. The betas in the Ba calibration data are therefore clustered in the phonon quadrant closest to the source. In contrast, the betas in the WIMP search data most likely come from surface contamination spread across the surface of the detectors, giving a more uniform position distribution. Also, these surface beta emitters have a lower energy distribution than betas in the Ba source. The differences in position and energy produce a shift in the timing parameters. This result led us to use the WIMP search multiples data to calculate the number of single-scatter betas expected to leak into the signal region in the WIMP search singles. This is different from what was done in run 118, in which the Ba beta population was used, and what we had originally intended to do in run 119.

4.3.2 Beta Sources

The analysis techniques used to reject surface events do not rely heavily on knowing the source of those events. However, beta events are our dominant background, so knowing their physical sources will lead to better control of these backgrounds in the future. Several studies were done to characterize the beta background in order to control it better.

One beta source known to be present comes from the gamma background. Gammas can kick electrons free in the crystal to yield a surface event signal in the sensors. Simulations show that the rate of these betas is roughly 1 per 25,000 incident gammas [47]. This approximation predicts roughly 1 beta per detector in the course of the run, or about 1% of the observed beta background.

The rest of the betas come from surface contamination of the detectors themselves. The first contributor to the beta contamination is radon contamination. Tower 1, the tower which was first run at the Stanford Underground Facility, was exposed on several occasions to air, especially the radon rich air in the tunnels at Stanford. Its history was not as well documented as later towers in terms of the time it was kept in environments with differing radon levels. Tower 2 was more carefully tracked and guarded against high radon environments. Subsequent towers are also expected to be cleaner in this respect.

Midway through the one-tower run at Soudan, it was noticed that the beta rates were higher than expected, and a purge using two week old air was started through the area immediately surrounding the copper cryogenic cans housing the detectors. The gamma rate dropped by approximately a factor of 5 over all detectors. Correspondingly, the beta rate decreased by less than a factor of two in most detectors. The excess betas most likely had been products of gamma scattering near the ZIPs. The drop in rates indicated that the area around the detectors, at least,

was significantly contaminated with radon. The old air purge continued throughout the two tower run.

To measure the radon contamination on the detectors, we studied the alphas in the WIMP search data. Alphas are well outside the signal region, scattering at high energies with an endpoint of 5.3 MeV, and are therefore not considered a dangerous background. Fig. 4.13 shows the alpha population for T1Z3, well separated from the electron-recoil band. Alphas are expected to be produced as part of the radon decay chain, from the decay of ^{210}Po to ^{206}Pb . The Po itself is produced in the ^{210}Pb decay to ^{210}Bi , which decays with a half life of 5 days to ^{210}Po , so the beta rate from this interaction should be roughly the same as the alpha rate overall. The alpha rate in each detector is shown in Table 4.2. The detection rate of alphas is taken to be approximately 100%, whereas the detection rate of betas is much lower. Simulations show the expected beta detection rate to be approximately 80% [47].




Figure 4.13: Ionization energy (keV) vs integrated phonon energy (keV) for T1Z3, using WIMP-search data. Alphas, marked with xs, are high-energy, low-yield events, well separated from the high-energy gammas in the electron recoil band (marked by dots). Alphas are produced by decays in the Rn chain. At these high energies, phonon energies are obtained by integrating the traces, rather than with the standard optimal filter, and ionization energies are obtained using a five-parameter fit if the digitizers saturate. Solid lines indicate the cuts used to select alphas, and the dashed line labels the 5.3 MeV endpoint for the alphas.

The ratio of alpha to total beta rates in each detector in the current run are shown in Fig. 4.14. They are approximately on the same scale, and so we conclude that radon contamination is a significant contributor to the beta rate; most likely 30% to 50% of the surface event background is due to radon contamination. Because each subsequent tower installed in Soudan has lower radon exposure, we expect that beta rates will be lower in future runs.

Table 4.2: Alpha rates, in events/cm/day for runs 118 and 119.

Detector	Run 118	Run 119
T1Z1 (Ge)	1.209×10^{-2}	4.70×10^{-3}
T1Z2 (Ge)	0.930×10^{-2}	3.60×10^{-3}
T1Z3 (Ge)	0.651×10^{-2}	3.98×10^{-3}
T1Z4 (Si)	0.791×10^{-2}	9.48×10^{-3}
T1Z5 (Ge)	1.442×10^{-2}	6.10×10^{-3}
T1Z6 (Si)	1.814×10^{-2}	6.10×10^{-3}
T2Z1 (Si)	N/A	1.11×10^{-3}
T2Z2 (Si)	N/A	3.40×10^{-3}
T2Z1 (Si)	N/A	7.20×10^{-3}
T2Z1 (Si)	N/A	6.45×10^{-3}
T2Z1 (Si)	N/A	9.29×10^{-3}
T2Z1 (Si)	N/A	7.58×10^{-3}

There are several other possible sources of beta contamination. To search for these, several tests were performed at the University of Minnesota and at Case Western Reserve University. Test wafers were fabricated alongside actual ZIPs at the Stanford Nanofabrication Facility and sent to the two testing sites. Several analyses of the surface elements of the wafers were performed, and these tests and their results are detailed in Appendix A.

The contaminants that were studied in detail are C, K, and Sb. The C contamination is due to exposure of the wafers to air, comprises less than a monolayer, and is not deposited during fabrication. It is not expected to contribute significantly to the beta background rate. Whether or not there is ^{40}K contamination on the detectors is still an open question. Some tests were only able to set an upper limit on K contamination of the wafers, whereas more sensitive tests found trace amounts of ^{39}K , but are not sensitive to the beta emitter of interest. However, with these limits, it is unlikely that K contamination can account for all of the remaining beta background seen in data. Mass spectroscopy tests have identified ^{125}Sb , an anthropogenic beta-emitter, on the surface of all test wafers. Preliminary calculations have shown that this could be a significant contributor to the beta background, enough to account for the background not due to Rn.

Operational steps have already been taken to achieve lower Rn contamination on tower 3, 4, and 5. The beta backgrounds measured with these towers in the future will inform us as to the level of the remaining background and aid in identifying it. Tests at Case and UMN also continue to




Figure 4.14: Ratio of alpha rate in detectors to beta rate in detectors, as a function of detector. Detector numbers run 1 to 12, with detector 7 corresponding to T2Z1.

study the nature of the surface contamination of the ZIPs. Identifying the source of the additional beta contamination and reducing it will further limit the background in CDMS.

Chapter 5

Analysis: Methods and Results

5.1 Event Selection and Signal Region Definition

5.1.1 Event Selection

We began the analysis using standard data quality statistical checks, specifically Kolmogorov-Smirnov tests. The tests compared each data series to a reference, composed of three good series taken from the beginning, middle, and late times in the run. The KS tests included a battery of 14 tests, both one- and two-dimensional, for each of the 12 detectors, a total of 168 KS tests scores for each data set. Figure 5.1 shows the KS test distributions for a data set that is consistent with the template and one that is not. The inconsistent data set has a KS distribution skewed toward 0. The selection criteria was that any data sets with 15 or fewer test scores less than 0.01 of the 168 was a good data set. The most commonly failed tests were the noise level consistency checks in both ionization and phonon channels. After unmasking the WIMP-search data, we did decide that in future these tests will be somehow weighted, in that failing some is more problematic than failing others. For this analysis, the tests were all weighted equally.

There were a series of additional cuts used to ensure that events in the data sets that passed the KS test cut were usable. These consisted of:

- Operational cuts, including cuts on data that were too close to cryogenic transfers, power outages, or gpib errors, that had bad optimal filter files, or for which errors occurred during automated analysis.
- χ^2 cut on the sum of the ionization channels. This cut was very important in the previous Soudan run, during which a thermostatically controlled heater fan on the top of the dilution fridge periodically introduced large noise pickup into the electronics and affected the quality of the ionization signal. A better scheme using a DC fan with ceramic casing was implemented in run 119. This cut also serves to reject pile-up events, in which two pulses occur in the




Figure 5.1: KS test significances for simulated data sets with (a) 0% variability and (b) 15% variability in the tested data set.

150ms trace window. These events, however, are extremely rare in the WIMP-search data.

- A cut on the pre-pulse baseline standard deviation on the phonon channels. This is a largely redundant cut, looking for increased noise in the phonon channels, whereas the chi-squared cut looks for increased noise in the ionization channels.
- A cut on drastically increased beta rates in a data set, which may point to operational difficulties, such as lack of neutralization, or real increased rates, such as those immediately following the Cf calibrations.
- Ionization Threshold. This cut required that the signal in the inner electrode be 3.85 sigma above noise. Simple simulations showed that, given the statistical level of our data, no noise

event is likely to be mistaken for a real pulse above this level. This cutoff must be lower than the phonon threshold cut to allow for events at low energy with low yield.

These data selection cuts reduced the raw exposure by $\sim 5\%$. The data preselected for analysis using these cuts totaled of 74.5 live days.

5.1.2 Fiducial Volume Definition

In addition to these data quality selection cuts, a fiducial volume cut was defined. The goal is to ensure that selected events occur in the detector volume encompassed by the inner electrode, where the electric field is well defined and the field lines end on the electrodes and not on the side of the detector, where accumulated charge may be lost. The cut requires that the signal in the outer ionization electrode be within noise. Specifically, the Ba data was split up in bins 10 or 20 keV wide in the inner electrode energy, and the outer electrode noise band is fit with a Gaussian in each bin. Then independent linear fits of outer electrode energy versus inner electrode energy were done to the upper and lower boundaries of the 2 sigma bands. The results are shown in Fig. 5.2.

The efficiency of the inner electrode cut is surprisingly complicated to calculate. To do so, we begin by selecting a population of events smoothly distributed throughout the detector volume with characteristics similar to the signal. Ba data, though it is used to evaluate the performance of the detector and define the cut, would give incorrect efficiencies because the point source is located to the side of the towers of detectors. ^{252}Cf neutrons were chosen as the best population to estimate efficiencies because the neutrons scatter sufficiently before hitting the detectors to give a relatively isotropic distribution in the ZIPs.

The difficulty arises mainly from two phenomena in charge collection, one of which is related to the region between the inner and outer electrodes, a 1 mm gap in coverage, and another is related to electronics cross talk. Fig. 5.4 is a plot of yield calculated from the outer electrode versus the yield calculated from the inner electrode. Two main population, labeled as *curl* and *gap*, represent abnormal behavior. Events in the *gap* region occur in the gap in coverage between the inner and outer electrodes. The charge collection in this region is not complete, and the tail of this distribution extends down into the inner electrode event region. Events in the *curl* region, named for its shape, have negative pulses in the inner electrode, most commonly caused by electronics cross-talk. Again, the tail of this distribution can populate the inner electrode event region.

These events with abnormal charge collection should appear neither in the numerator or the denominator of the passage fraction of this cut. Two independent calculations were done of the inner electrode efficiencies in which estimates of these tail contributions to the overall statistics were estimated and subtracted out. The results are shown in Fig. 5.5. Although these two estimators were very similar, in each energy bin and each detector the lowest efficiency was taken as the final number. Such a choice was conservative in that it slightly reduced our exposure.

This data-based approach is a significant change from the previous run. In run 118, a simulation




Figure 5.2: Outer electrode energy vs inner electrode energy for (a) T1 and (b)T2. All axes are in keV. The inner electrode cuts are shown as black dashed lines.




Figure 5.3: Cartoon of ionization electrodes. There is a 1 mm gap between the inner and outer electrodes, and .8mm of uncovered surface at the detector's edge.

was used to estimate the efficiency of the inner electrode cut because a simple calculation of the efficiency using data gave a higher estimate than did simulations. However, great care was taken in run 119 to use the data and make the appropriate subtractions rather than a simulation. It should be noted that these efficiencies are significantly lower, by 5 to 10% depending on energy, for this run than they were in run 118: this cut rejects more neutrons than its predecessor.

5.1.3 Signal Region Definition

In addition to data selection cuts, we used background rejection techniques to define the signal region for a WIMP in our detectors. Some of these cuts have been discussed in more detail in Chapter 4. They are:

- Single event. This cut relies on theory in that WIMPs are very unlikely to multiply scatter in our experimental apparatus. It requires that exactly one detector have a signal in the sum of the four phonon channels that is more than 6 sigma above the noise level.
- Phonon Threshold. Calculations using the ionization threshold simulation data show that the phonon energy threshold should be 7 keV to ensure that <10% of misconstrued ionization pulses lead to gammas appearing in the nuclear recoil band [55]. This threshold and the ionization threshold combine to give a conservative low energy threshold. Between 7 keV and 10 keV, however, the electron and nuclear recoil bands are close enough together and statistics




Figure 5.4: Illustration of different ionization collection populations. Events below the dashed line (an approximation) are inner electrode events, used in the analysis. On the diagonal are events shared by both electrodes. Circled gap events, also called funnel events for the shape of the distribution, have suppressed ionization collection and occur in the gap between the two electrodes. The curl events have negative pulses in the inner electrode, most commonly caused by cross talk in the electronics between the outer and inner channels. The tails of these two circles distributions have a small probability of leaking into the inner electrode population and must be considered in the inner electrode cut efficiency.

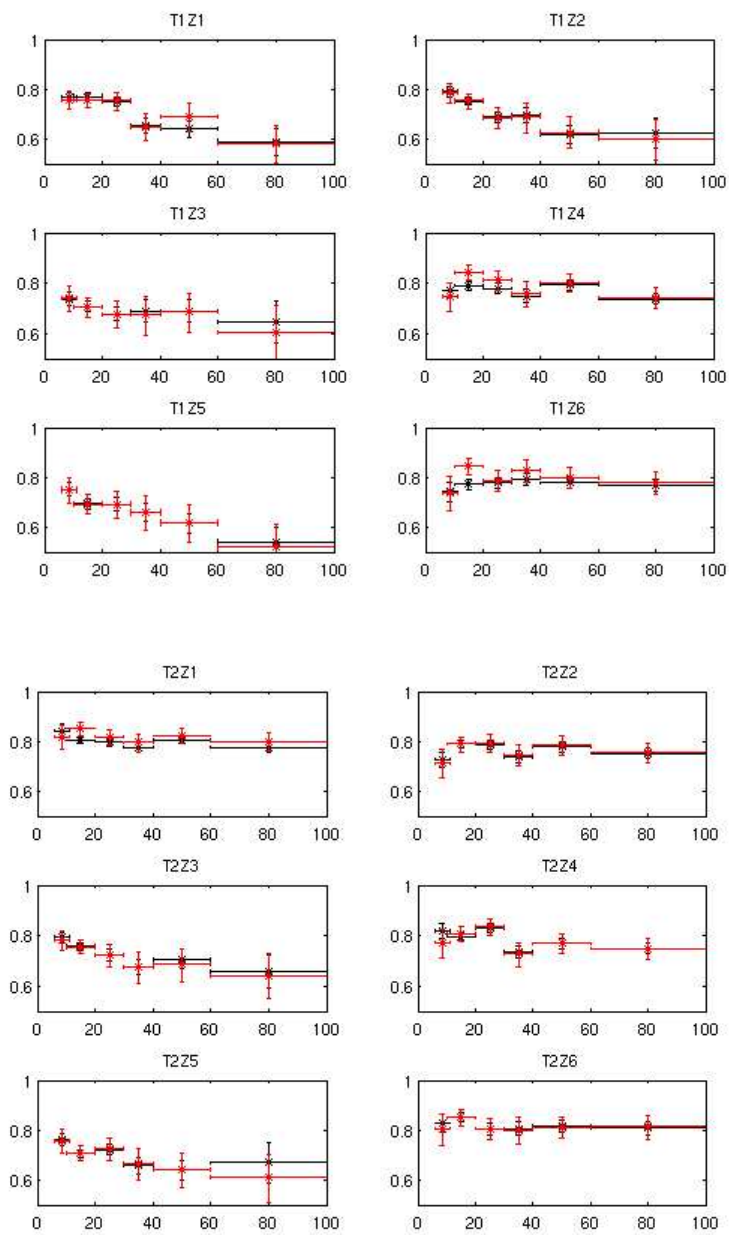


Figure 5.5: Inner electrode selection efficiency (fraction of total population) vs recoil energy (keV) for all detectors. Results from two methods of calculation are illustrated; they are consistent with one another in every case.. Data are from [54] and [53].

levels low enough that determining beta contamination is very difficult. The decision was made that our phonon threshold would be 10 keV, but efficiencies and cuts would be defined down to 7 keV as well while the data was still masked. All cuts performed well in this lowest energy range, and the 7 keV threshold was used for some calculations, where it will be specified. Unless otherwise specified, the phonon threshold used is 10 keV.

- Veto anti-coincidence. This cut requires that there has been no above-threshold signal at all in the muon veto for 50 μ s before an event, and no muon-like activity in the veto for 184 μ s. This is to strictly eliminate cosmogenically generated events in our detectors.
- Nuclear recoil. This cut requires that the events have an ionization yield in the 2 sigma band defined by neutrons, centered at approximately 0.3. This cut rejects gammas and most beta backgrounds and is described in detail in Chapter 4.
- Timing cuts. Several independent versions of these cuts were defined while the signal region was masked. All require that the phonon pulse have timing parameters sufficiently slow to identify it as a bulk recoil rather than a surface event. Some also have energy sharing requirements. The timing cut used in this analysis is a one-dimensional cut on the sum of two phonon timing parameters.

5.2 Efficiencies and Expected Backgrounds

5.2.1 Overall Analysis Efficiencies

The overall WIMP-search efficiencies of the analyses using all of the above cuts were calculated in energy bins between 7–100 keV. The efficiencies of the ionization threshold, phonon threshold, and singles cuts were determined using simulations and DAQ triggered, pulse-less data from the WIMP-search data set. The veto anti-coincidence efficiency was calculated mathematically. The nuclear recoil, inner electrode, and timing cut efficiencies were calculated using neutron data from Cf calibrations. The acceptance fraction of WIMP-like nuclear recoils calculated after each level of cuts were applied to Ge and Si detectors is illustrated in Fig. 5.6. First the data selection cuts are applied, with greater than 90% acceptance. Then the fiducial cut, cQin, is applied. Then the signal region cuts are applied in two stages: first, the cut selecting single-detector events and the cut selecting the nuclear recoil band in yield. Then the timing cuts and the cut rejecting the poor region of T2Z5 are applied. Clearly, the timing cut causes the low efficiency at low energies.

To obtain the final sensitivity, the efficiencies are interpolated to get a smoothly varying function in energy. The results in Ge and Si are shown in Fig. 5.7. In determining the approximate systematic error on these efficiencies (statistical error is quite small), several factors are involved. The dominant one is the placement of the nuclear recoil band, defined using neutrons. It is possible that this band is too high, given the theoretical differences between WIMPs and neutrons discussed in Chapter 4. Neutrons can inelastically scatter, producing gammas and raising the yield of events




Figure 5.6: WIMP recoil efficiencies of analysis cuts, applied progressively. Efficiencies are averaged over all Ge or Si detectors. See text for explanation of the cuts as they are applied.




Figure 5.7: WIMP recoil efficiencies of primary analysis, including all cuts. Efficiencies are averaged over detectors. Behavior at low energies is dominated by timing cut efficiencies, and behavior at high energies is dominated by inner electrode cut efficiencies, which are worse in Ge than in Si.

at high energies. It is therefore possible that the efficiency uncertainties extend down as much as 15% below the averaged value in the figure.

One effect that lowers the expected efficiencies that is not an analysis cut is a DAQ efficiency. Even at high energies, the DAQ does not always trigger and digitize when an event occurs in the detectors. This effect was discovered by comparing digital trigger information to the information from analog pulses for multiple events. When there are energy deposits in two or more detectors, but a detector is missing the digital tag to say it is triggered, the DAQ has missed a trigger. Using this multiples information, the percentage of missed single triggers is estimated. Above 15 keV, this rate of missed triggers ranges from 1 to 3% depending on the detector. This problem is linked to the Receiver-Trigger-Filter boards associated with each detector and has now been fixed in firmware. The effect is included in the overall efficiencies in this chapter.

5.2.2 Expected Background Calculation

Background estimates, or calculation of the number of background events expected to pass all analysis cuts, were also done while the data was still masked. Any gamma background event would be called a surface event, so the expected gamma background is 0 events. Using data and monte carlo simulations, the expected neutron background is 0.05 events in Ge and 0.05 events in Si; the errors in these estimates are on the order of the estimates themselves.

The expected background from surface events is dominant. This calculation requires two input variables per detector. The first is the *passage fraction*, the fraction of low-yield beta events that pass all other selection cuts that are expected to also pass the timing cut. The original intention was to use betas from the Ba data to evaluate the passage fractions. The Ba data was divided into two halves with approximately equal numbers of events, with the expectation that one half would be used to define analysis cuts, and the other half would be used to calculate expected backgrounds. As discussed in Chapter 4, however, because of the different position and energy distributions of the betas in the calibration data and WIMP-search data, there is a systematic difference in the timing parameters between the two data sets. Therefore, WIMP search data are used to calculate passage fractions. This leads to an increase in statistical error because there are fewer beta events to use, roughly 20% of the statistics level of the Ba data. However, this approach reduces the systematic uncertainty which would have been more difficult to evaluate using the Ba data. The passage fractions from these data sets are shown in Table 4.1.

The second input required for the surface event background calculation is the actual number of low-yield events in each detector that pass all cuts but the timing cut. The final calculation is simply the product of this number and the passage fraction. There are several methods of estimating this number without unmasking the WIMP-search data, and these were used while we were deciding at what values to set the timing cuts. However, the final background calculation is not based on estimates, but an actual counting. Accordingly, the unmasking of the WIMP search data in the signal region took place in two stages. In the first, all analysis cuts except the timing cuts were applied to the data. Table 5.1 shows the distribution of events that survived this first

stage of cut application across detectors. Only events above threshold (10 keV in Ge, 10 keV or 20 keV in Si) and in the two sigma nuclear recoil band are considered to have survived and were used to calculate expected backgrounds. In total, there were 15 such events in the Ge detectors and 17 such events in the Si detectors. Detectors T1Z1, T1Z6, and T2Z1 were not included in this analysis.

The expected background was calculated in two different ways, using two different passage fractions, and results from both calculations are summarized in Table 5.2. The first method used the passage fraction of all WIMP-search multiples in the energy range 10-100 keV with ionization yields between 0.1 and the bottom of the five sigma electron recoil band. Most of these events are close to the electron recoil band and therefore far from our signal region. Each detector's number of events passing all prior cuts, including event selection, fiducial volume, and signal region cuts except the timing cut, was multiplied by the passage fraction for that detector, and then the expected backgrounds were summed over detectors. This method of estimation predicted 0.40 events in Ge and 0.89 events in Si. The Si estimate is heavily dominated by the expected background in T2Z6. As the only end detector used in the analysis (the detectors at the top of each tower and T1Z6 are excluded), this detector has a higher rate of single scatter events than the others in the analysis because of the lack of a detector below it. The expected number of background events in this detector alone is 0.7. The Ge estimate is dominated by T2Z5, at 0.29 expected events. This dominance is less due to the high number of events than to the large passage fraction. This detector's timing parameter distribution shifts positively from the calibration data to the WIMP-search data and consequently its passage fraction is higher than expected.

Table 5.1: Events passing all analysis cuts but surface event rejection (timing) cuts and expected backgrounds, by detector. The two different beta populations used for the two methods of background calculation are described in the text.

Detector	Events	Expected Bkgd Method 1	Expected Bkgd Method 2
T1Z2 (Ge)	0	0	0
T1Z3 (Ge)	1	0	0
T1Z4 (Si)	7	0.099	0
T1Z5 (Ge)	3	0.020	0
T2Z2 (Si)	2	0.033	0
T2Z3 (Ge)	5	0.093	0
T2Z4 (Si)	1	0	0
T2Z5 (Ge)	6	0.29	0.032
T2Z6 (Si)	15	0.76	0

Table 5.2: Expected beta background and statistical error in Ge and Si using all WIMP-search beta multiples, only nuclear recoil band WIMP-search beta multiples, and the weighted average of the two.

	Using all WIMP-search beta multiples	Using nuclear recoil WIMP-search beta multiples	Using weighted average
Ge	0.40 ± 0.3	0.18 ± 1.1	0.37 ± 0.15
Si	0.89 ± 1.0	0.00 ± 2.6	0.43 ± 0.86

The second method of calculating expected beta backgrounds used the passage fraction of only the multiple scatter WIMP-search events that lie in the nuclear recoil band, the signal region itself. There are far fewer of these kinds of events, but they are closer to the expected behavior of our background. Again, statistical error for this estimate is higher, but the systematic error, far more difficult to calculate, is smaller. These passage fractions were multiplied by the same numbers of events after all other cuts. This method predicts 0.18 events in Ge and 0.0 events in Si after the primary timing cut. The only background is predicted to be in T2Z5.

The published background estimate is a weighted average of these two methods of calculation. The dominant systematic error is taken to be large enough to include both estimates in 1 sigma. The result is a prediction of $0.37 \pm 0.15(\text{statistical}) \pm 0.20(\text{systematic})$ in Ge, and $0.43 \pm 0.86(\text{statistical}) \pm 0.49(\text{systematic})$ in Si. The statistical errors are calculated using the binomial error function. The systematic error is dominated by the difference between the two background estimates. Additionally, it accounts for the systematic errors in the inner electrode cut and the cut that rejects events that occur in the questionable region in T2Z5. Several checks were performed regarding this background estimate's consistency with alternate methods of prediction. In both the Ge and Si cases, the $1\text{-}\sigma$ systematic error quoted is consistent with the background prediction from using the ^{133}Ba data as well as the prediction using the multiple events in the WIMP-search data and simply the ratio of the number of singles to multiples.

It should be noted that the discovery potential of this analysis is such that three events would have to survive the cuts before any significance could be assigned to the result; one or two events would be consistent with expected background.

5.3 Results

5.3.1 Unmasking

The second step of unmasking the WIMP-search data occurred at 12am Pacific Daylight Time on April 1, 2005. The surface event rejection cut was applied to the events in the signal region. The results are shown in Fig. 5.8 and Fig. 5.9. One event passes all cuts in T2Z5. There is another

event in T2Z5 that passes the timing cut and lies just below the two sigma nuclear recoil band. This event is not in the signal region, but if the nuclear recoil band had been adjusted downward to account for the differences between neutrons and WIMPs, as described in Chapter 4, this event would most likely have been in the signal region.




Figure 5.8: Events passing all cuts except timing cuts (dots) and all cuts including primary timing cut (x's) for Tower 1. Vertical axis is ionization yield and horizontal axis is recoil energy (keV). Missing detectors were not analyzed in this final stage due to known performance problems.

After unblinding, the 10.488 keV event in T2Z5 that passes all cuts was found to occur in a run during which T2Z5 had reduced ionization collection. This data set has some peculiar characteristics that allowed it to pass our data selection routines, even while its behavior was truly abnormal. Fig. 5.10 is an illustration of yield vs time in the this particular data set, taken on June 26, beginning at 2:50 pm. Prior to this dataset, operators at the mine had been attempting to use a ^{60}Co source to produce new lines to use for calibration of the ionization channels. However, the rate of this source was so high as to make avoiding pile-up, even in individual detectors, impossible. This source was abandoned as a secondary gamma analysis, but the heavy bombardment seems to




Figure 5.9: Events passing all cuts except timing cuts (dots) and all cuts including primary timing cut (x's) for Tower 2. Vertical axis is ionization yield and horizontal axis is recoil energy (keV). Missing detectors were not analyzed in this final stage due to known performance problems.

have affected T2Z5 adversely. As shown in the plot, the electron recoil band was wider than normal and quickly began to droop as time wore on. By the time this event occurred, the electron recoil band overlapped with the nuclear recoil band, presumably due to poor neutralization of some kind in the detector. As always, four hours into the run, the DAQ automatically stopped data taking, grounded the detectors, flashed all LEDs, waited for the icebox to cool, and resumed data taking. For the last few hours of the data set, the detector behaved as expected.




Figure 5.10: Yield vs time for the data set 140616_1450, detector T2Z5. All events are shown; single events are Xs. Large, dark Xs are events that passed the timing cut with energies between 7–10 keV. Large, lighter Xs pass are above 10 keV and pass all cuts but the timing cut; the additional circle indicated that the event passes the timing cut. The circled event is the 10.488 keV event that passes all cuts. It occurred in the middle of a period when the gamma band in this detector had clearly drooped away from 1. After the LED flash, the detector recovered.

This event is clearly not a candidate or even surface background event. The malady suffered by T2Z5 in this data run seems to have been very rare. Other detectors did not show the wear of the Co calibration, and only a few detectors in other datasets, all within a week of this dataset, show peculiar characteristics, and none of them as badly as T2Z5 in this dataset. However, the data selection will be changed in the future to better remove this kind of data series. The relatively well behaved periods at the beginning and the end of the run seem to have kept low yield event rates from appearing too far from average in the KS tests.

Despite the unfortunate oversight of allowing this questionable dataset into the main body of

data, the damage is minimal. The inclusion of a background event so close to the phonon energy threshold in the Ge does little to affect the calculated limit. It is included in all limit plots shown.

5.3.2 WIMP Search Limits

Upper limits on the WIMP-nucleon cross section were calculated for a range of WIMP masses. The limits are calculated upper limits on the WIMP-nucleon cross section for a range of WIMP masses, assuming standard halo parameters and spin independent interactions. These upper limits are frequentist, and should be interpreted as meaning that the probability of obtaining the CDMS result is less than 10% for cross-section/WIMP-mass theories below the limit. The limits are calculated without background subtraction using the optimal interval method described in [56]. This method has been used for data taken in Soudan, where there are few background events. One of its advantages over methods used previously is that it does not require assumptions to be made about the nature of the background. It depends only on the total exposure (kg-d, taking efficiencies into account) and the numerical expected background.

A probability calculation is done for each cross-section and WIMP-mass combination, fixing the WIMP mass and then finding the upper limit cross-section for that mass. For each combination, the expected signal is calculated from the proposed cross-section and known exposure: the expected number of events is called μ . The energy region (WIMP-mass interval) in which the largest number of events is expected but none are observed is defined: it is the “maximum gap” and its width is x . Then the probability of measuring a signal in this region is calculated. The probability is C_0 , where

$$C_0(x, \mu) = \sum_{k=0}^m \frac{(kx - \mu)^k e^{-kx}}{k!} \left(1 + \frac{k}{\mu - kx}\right) \quad (5.1)$$

and m is the greatest integer $\leq \mu/x$. For the run 119 case, in which one event at the edge of the expected signal region passed all cuts, the maximum gap is large: it is only marginally smaller than it would have been if no event had passed all cuts. The proposed cross-section and WIMP mass are excluded if C_0 is more than 90%, and the upper limit published is at the 90% value.

Fig. 5.11 shows two resulting limits from the Ge data. The first is the result from the analysis of R119 Ge data in the energy range of 10-100 keV. The second, lower curve is calculated using the same physical assumptions, but combines the analysis of R119 with the previous run at Soudan, and additionally includes the 7-10 keV bin for the R119 analysis. This figure also shows the spin independent limits using the Si. Standard halo assumptions are still used.

Fig. 5.12 shows the limits calculated from Si, also using 7–10 keV and combining the analysis with the previous run. These data set new limits down to a WIMP mass of 6 GeV. Even in the higher WIMP mass range, the Si data from CDMS are approaching the sensitivities of other experiments using heavy elements.

The Ge limits, being far lower than the Si, are the primary result of the CDMS experiment. The second limit curve for Ge is more than a factor of 10 lower than any experiment’s published result

for most of the plotted WIMP mass range. It excludes theoretical parameter space for WIMPs never before excluded, including the theoretical models of CMSSM that have been beyond reach in the past. The Si limit, while not as groundbreaking, is the first publication of a limit using Si, and the first time this limit was low enough to start to compare to the Ge results.

There is a theoretical possibility that the WIMP-nucleon interactions are spin-dependent, and therefore the cross section is only non-zero for unpaired nucleons. The DAMA allowed regions for these interactions are less easily ruled out by experiment because, although the cross sections are orders of magnitude greater, the rates are not proportional to A^2 . Fig. 5.13 shows the DAMA allowed region and limits obtained by other experiments, as well as the new limits from CDMS, both from Si and Ge. In the pure WIMP-neutron coupling, CDMS sets the most restrictive limit over most WIMP masses. For pure WIMP-proton coupling, indirect experiments, such as that conducted at Super-Kamiokande, still set the most restrictive limits. These allowed regions and limits are still significantly higher in cross section than theories, which are not shown on the plot because their scale is so much lower than the region of experimental interest.

5.4 Future Prospects

Five towers of detectors are now installed in the icebox at Soudan, with a total of 18 Ge and 12 Si ZIPs. If these detectors are operated, as anticipated, for roughly a year, another factor of 10 over our current limits will be gained in the Ge spin independent interactions. Several changes for that run have been made, including better grounding of the detector face during fabrication, additional dilution refrigerator plumbing, and improvements in DAQ readout and throughput. The most significant hardware change has been the addition of a CryoCooler at the 4K stage in the icebox to handle the additional heat load from 5 towers of detectors. Also, 30 new computers have been added to the analysis farm on the surface at Soudan to keep up with the expected high event rate from Ba calibrations of 30 detectors.

Beyond this simple extension of current techniques, the CDMS collaboration has proposed a SuperCDMS experiment, involving significantly different detectors, larger masses, and a deeper site [58]. The detectors would be 1 inch (rather than 1 cm) thick, to reduce the surface area to volume, and therefore surface event rate. Each face would have interlaced phonon and ionization sensors, improving symmetry and collection efficiency. This work, still in development and simulation stage, would significantly improve CDMS' sensitivity, rivaling that of proposed liquid Xe experiments.

A key question is whether future CDMS work may be background limited. It is clear that even at the Soudan site, expected neutron backgrounds will not be more than a very few events for several years, and an additional external veto would extend the life of the Soudan setup. A move to a deeper site is also proposed to lower that expected background even further. The dominant background most likely will always be surface events. The three new towers in the Soudan icebox have been carefully guarded from Rn exposure, and the background rates in these detectors in the new run will tell us a great deal about the source of our beta backgrounds. However, analysis

techniques used even in run 119 have proven to be very useful in rejecting these events at the analysis level, and they can be improved to be far more strict than they were in run 119. These two facts, along with design features of the new ZIP, point toward CDMS being able to keep surface event background at a reasonably low level for quite some time.

Run 119 has proven its usefulness as a step in advancing operational and analysis techniques for CDMS in the Soudan Laboratory. The lessons learned are even now being applied in further running, and great improvements in sensitivity are very near.




Figure 5.11: Spin-independent limits from R119 (dashed) and combined advance R119 and R118 (solid), compared to other limits and theories. Standard halo assumptions are used.




Figure 5.12: Spin-independent limits from R119 Si data. Two curves from Fig. 5.11 are included: the combined R118 and R119 Ge limit and the ZEPLIN-I limit, both of which are vertical at a WIMP mass of 10 GeV. Standard halo assumptions are used.




Figure 5.13: Spin dependent allowed regions and limits for (a) pure WIMP-neutron coupling and (b) pure WIMP-proton coupling. Results in black are new CDMS results, Ge in solid and Si in dot-dash. The dashed curve is combined Ge R118 and R119 analysis, and the grey (neutron plot only) is from the last CDMS run at the shallow SUF site. In red is the allowed DAMA region, and (b) includes limits from CRESST (blue crosses), SIMPLE (green circles), and SuperKamiokande (dark green stars). Figure from [57], produced by [41].

Appendix A

Characterization of Surface Contamination

Numerous tests were performed at Case Western Reserve University and the University of Minnesota to determine the source of the surface contamination on CDMS ZIP, and these tests continue in an ongoing study. Thin wafers, created in the same batch alongside ZIPs at the Stanford Nanofabrication Facility, are mailed to both institutions. There they are tested for trace amounts of beta emitters that may be causing the surface event background in CDMS data. The description of tests and results in this appendix will focus mainly on the UMN work, but will compare with Case results when appropriate.

A.1 Description of Tests

The only test performed by both institutions is a Rutherford Backscattering (RBS) test, in which the sample is mounted on a target that is then bombarded with high energy ions. The RBS detector measures the energies of the ions backscattering off the target sample.

The UMN accelerator is an MAS 1700 pelletron tandem ion accelerator, and it can produce He beams with energy ranges of 0.2-5.1 MeV [59]. The Case machine, an NEC 5SDH tandem pelletron, is very similar and has the same energy range for He ions [60]. The standard energy used for much of the CDMS analysis is 2 MeV, but tests for Carbon used a machine setting of 4.32 MeV. The UMN detector system, an Analytical Endstation RBS 400 from Phi Evans, includes detectors for RBS as well as others, such as the PIXE detector discussed later. The accelerator is the main component of the Ion Beam Analysis laboratory at the UMN Characterization Facility. The accelerator, detector, and associated analysis software are all user facilities and provide training for users. This provision extends analysis flexibility and keeps costs down as well. The tests done using the Ion Beam Analysis facility are not destructive.

For a given element, the RBS spectrum gives depth information; deeper buried atoms give lower energies to the backscattered ions. With known elements on the surface, then, RBS can provide a detailed depth profile. There can be some confusion of unknown elements, however, if it is also not known whether they are on the surface or buried. The software program used to do analysis at the IBA facility is a DOS-based program called HYPRA. Fig. A.1 is a typical RBS spectrum from HYPRA using data from CDMS wafers, showing large-scale structure. In the first sample, one with only am-Si on the surface, there is only a Si edge and a small W peak at the surface W channel of 605 (corresponds to an energy of 1.8 MeV; 3 keV/bin). This trace W is always found in the am-Si deposition step. The next wafer has had Al deposited on top of the Si. The Si edge has moved to the left (lower energy), and an Al peak near that edge is there. The W peak, more diffuse, has also moved left. The third wafer has W deposited on top of the Al. Everything has again moved leftward, and there is new, very large W peak at channel 605. In this sample, the smaller W was too small to see on this scale. Contamination peaks are very small and lie between the Si edge and the W peak.

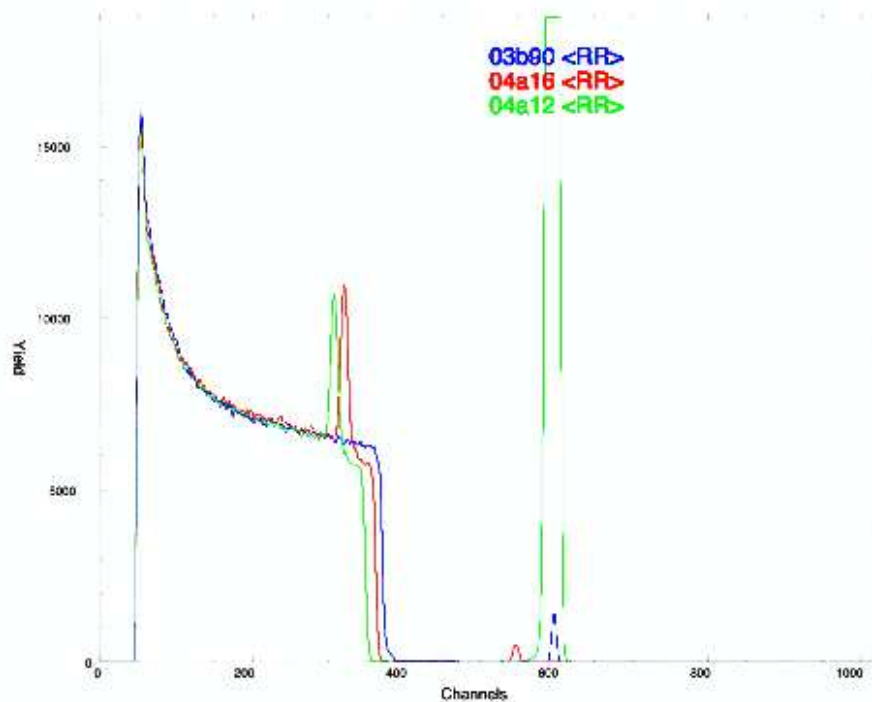


Figure A.1: RBS spectrum of 3 CDMS wafers: 03B90 has only am-Si, 04a16 has am-Si and Al, and 04a12 has am-Si, Al, and W. See text for interpretation.

HYPRA also allows the user to model layer profiles and compare to data, as well as measure

abundances in the peaks, as long as the element is known.

Alongside the RBS test, a Particle-Induced X-ray Emission (PIXE) detector also records data during IBA runs, both at UMN and at Case. This detector simply records energies of gamma rays emitted from the inner shells of elements in the sample. It is useful in sorting out which elements are present in the sample, because it does not suffer from the interference between buried and surface elements in RBS. Therefore, PIXE can be used to positively identify an element and then its depth and abundance can be measured in RBS. A typical PIXE spectrum from HYPRA, obtained using a CDMS wafer, is shown in Fig. A.2. Note the peak at twice the Al x-ray emission energy, labeled Al res. This peak is just below the energies for X-ray emission in K and Ar, and therefore caused confusion with these two elements. However, it is not present for wafers without Al. The Fe contamination is common on CDMS wafers. Also, there is a contamination peak present in all samples run on the UMN machine that has not yet been identified.




Figure A.2: PIXE spectrum of 04A12, which has am-Si, Al, and W. See text for interpretation.

Other tests run at UMN and Case include Auger electron spectroscopy. The machine at the UMN Characterization Facility is a Physical Electronics model 595 scanning auger microprobe (SAM) [61]. The instrument fires 3 or 5 keV electrons at the target and records the energies of

electrons, specifically auger electrons emitted from low energy levels. The auger electron measurement itself is not destructive; however, it is only sensitive to the first 30 Å of the sample. To obtain information from deeper elements, the surface of the sample is sputtered with a differentially pumped ion gun at an ion energy off 3 keV.

Fig. A.3 shows typical CDMS results from Auger analysis. While this technique was very useful in quantifying the atmospheric contamination on the wafers themselves (N, C, O), it was not sensitive to the heavier contamination like the beta-emitter candidates without tweaking. It was not used beyond the characterization of atmospheric contamination.

Another test that is just being explored at UMN is Inductively-Coupled Plasma - Mass Spectrometry (ICPMS). Two machines are operated by the MN geology department and have been help in identifying contaminants. One, which is more available than the other but has lower resolution, is a Thermal Electron PQExCell ICPMA quadropole unit. The other is a high-resolution Finnigan magnetic sector ICPMS. Samples were run through the first machine to identify mass ranges of interest and then scanned in the high-resolution machine. This method is destructive in that the sample is dissolved in acid and then then transformed into a plasma using Ar gas.

One test performed at Case but not at UMN is Secondary Ion Mass Spectroscopy (SIMS). In this test, the sample is bombarded with ions and secondary ion emission of surface material occurs. Mass spectroscopy of the emitted material is then performed. It is a destructive technique because the surface materials are actually released. It is the most sensitive test of all those performed for CDMS, and works well with tests such as RBS to confirm its findings.

A.1.1 Results

A combination of RBS and PIXE studies on various wafers have successfully identified several features of surface contamination; however, no sufficiently abundant beta emitters were discovered. Surface Carbon on samples studies at UMN was measured at between 4 and 6×10^{15} atoms/cm², which is 2 to 3 monolayers. This contamination is truly on the surface of the detector and is most likely due to exposure to air; the wafer samples were not stored under vacuum. Iron contamination was found on some samples but not others, with no apparent correlation to fabrication steps. However, Fe is not a beta emitter and therefore does not pose a problem at its low level of contamination.

For some time, there was uncertainty as to whether the RBS tests had detected K contamination on the wafers. Fig. A.4 shows the RBS peak thought to be K. It is sufficiently wide to be indistinguishable from the elements on either side of it, Ar and Ca. Eventually, extended runs and analysis of PIXE data obtained from them, also shown in Fig. A.4, revealed that the contamination element is actually Ar, the noble gas used to fill the chamber during RF etching during fabrication. Contamination with Ar is not surprising and not problematic, as it is not a beta emitter. Therefore, any beta emitters contaminating the surface of the detectors introduced during fabrication are most likely beyond the sensitivity of the RBS and PIXE tests.

The SIMS tests at Case have identified ³⁹K on the wafers at less than a monolayer abundance.





Figure A.3: (above) Auger scan of surface of CDMS wafer 03A36, which has W, Al, and am-Si. (below) Depth profile obtained using Auger and sputtering. From [61].




Figure A.4: (above) Zoomed-in RBS scan of sample 03B90 showing a clear contamination peak around the K energy, but also near the Ca and Ar energies. (b) PIXE spectrum of a long run (180 μ C) on 04A36, a sample with only W deposited on its surface (no Al). The contamination is identifiable as Ar.

However, because the SIMS test uses Ar as well in bombarding the surface of the wafers, it is not sensitive to ^{40}K , the beta emitting isotope. The question of whether there is beta-emitting K on the surface of the detectors is therefore still an open one, although it is limited to a level at which it cannot be the only beta contaminant.

The ICPMS tests at UMN have helped to identify two other candidates for beta contaminants. One is Bi, identified by the low-resolution machine. Work continues to explore this possibility, but the high-resolution machine could not resolve Bi from the Pb present in the acid used to dissolve the sample. The other identified possibility is Antimony. Natural Sb occurs in all samples, but tests with the high-resolution ICPMS identified the presence of ^{125}Sb , a beta-emitter, in all samples, including one that underwent no actual depositions or etchings. Fig. A.5 shows the plot from the high-resolution mass spectrometer with a definitive ^{125}Sb peak. The mass number is present at the level of $\sim 2 \times 10^{13}$ atoms/cm², to within a factor of two, in each sample, but is unclear what fraction of this amount is Sb and what fraction is the stable daughter of the beta decay. Work continues to try to identify this contaminant.




Figure A.5: High-resolution ICPMS scan around Sb masses of sample 04A16, a sample with am-Si and Al deposited. Natural Sb masses of 121 and 123 are clearly visible, as is the peak of the radioactive 125 isotope.

Appendix B

Details of Cosmogenics Simulations

Neutrons mimic WIMPs in the ZIP detectors because they cause nuclear recoils, so understanding and characterizing the neutron background for the CDMS experiment is important. It will become more so as the exposure increases in future data runs. The neutron rate at Soudan is low, however (1 neutron event has been observed in the course of two runs), and data is not useful in studying the characteristics of neutrons at the depth of the Soudan Underground Laboratory. Several simulations have been developed to study the neutron background and calculate the expected neutron rates, veto fractions, and spectra.

The simulations are based on Geant4 and FLUKA. The UMN simulation described here is based entirely on Geant4. The code used is release 4.5.p01 from the Geant collaboration at CERN [48]. The code is being updated to version 4.7, but all results shown here are using code from the earlier release. To begin an event, a muon is randomly generated along a $64\text{ m} \times 68\text{ m}$ plane 10m above the experimental cavern with energy and theta direction chosen from the theoretical distributions in [49]. The muon is propagated through Soudan Greenstone, the composition of which is found in [50], using the physics list QGSC provided by Geant4. The muon and its associated shower particles enter a $4\text{ m} \times 8\text{ m} \times 4\text{ m}$ cavern.

Fig. B.1 shows the energy distribution of the simulated cosmogenic muons as they enter the experimental cavern. The average muon energy is 225 GeV. Fig. B.2 shows the angular distribution of the same muons. These muons are accompanied by several different kinds of particles. Fig. B.3–B.5 show the energy spectra of neutrons, pions, and gammas in the showers associated with the muons as they enter the experimental cavern.

The neutrons are, of course, the background of interest, and neutrons well separated from their parent muons are potentially the most dangerous background. Not all neutrons enter the cavern accompanied by a muon; 21% are not accompanied by muons into the cavern. Of all neutrons




Figure B.1: Muon kinetic energy spectrum in the Soudan experimental cavern. Average energy is 225 GeV.




Figure B.2: Muon \cos (azimuthal angle) distribution in the Soudan experimental cavern. A value of -1 is straight down. Average value is -0.81.




Figure B.3: Neutron kinetic energy spectrum in the Soudan experimental cavern. Average energy is 41 MeV.




Figure B.4: Pion kinetic energy spectrum in the Soudan experimental cavern. Average energy is 2.4 GeV.

that enter the cavern, 35.8% enter within 1 m of a muon. Considering only those neutrons that are accompanied by muons, all those accompanied by a muon, 45.39% enter within 1 m of that muon. Fig. B.6 shows the distribution of distances of neutrons to their associated muons as the particles enter the cavern. Only the 79% of neutrons associated with muons are included in the plot.

In the simulation, the cavern is filled with air, and the CDMS shield and veto is centered at (0, 2 m, -0.4 m). It is closer to one wall than others to reflect the geometry of the actual CDMS setup. The scintillator panels surrounding the shield are sensitive detectors in the simulation, and any particles depositing energy there are recorded. Also, any neutron that intersects these volumes is recorded, whether or not it deposits energy, so the veto panels also serve as neutron counters. A typical scintillator panel energy deposition spectrum is shown in Fig. B.7. It can be compared to Fig. 4.5, the veto spectrum from R119 data. The simulation does not include ambient gammas in the cavern, but otherwise the two spectra agree well.

The use of the veto panels as neutron counters and scintillator allows calculation of the veto fraction of all neutrons hitting the veto. Of all neutrons incident on the veto, $74.7\% \pm 0.9\%$ were vetoed. This number is consistent with the estimates from other simulations using FLUKA and significantly higher than the rough estimate of 40% that we had used before performing the simulations. The increased veto fraction arises largely from the inclusion of shower particles. Only 31% of the neutrons entering the panels are vetoed by muons. Most are well separated from the muons that originated the shower. However, the accompanying shower particles, including pions and gammas, are also effective in vetoing neutrons that enter the CDMS shield.

Further study is being conducted which includes analysis of the particles that penetrate the CDMS shield and reach the icebox, where the ZIP detectors are housed. Current statistics levels, the equivalent of 88 live days of running at Soudan, are low enough that no neutron that was not vetoed has penetrated to the icebox. This agrees with expectations from data. Because muons that travel through the shield produce most of the neutrons that intersect the icebox, and these are all vetoed, over 99% of the neutrons in the simulation that reach the icebox have been vetoed. The quoted number for this fraction in the text of Chapter 4 is 95%, but this is a very conservative estimate. Work continues to quantify at what exposure the neutron background at Soudan may cease to be negligible. The simulations continue to be improved. It may be that additional vetoes will be useful. However, it is unlikely that CDMS will detect unvetoes neutrons in the near future.




Figure B.5: Gamma kinetic energy spectrum in the Soudan experimental cavern. Average energy is 35 MeV.




Figure B.6: Distance (m) between neutrons and muons that enter the cavern in the same shower event.




Figure B.7: Simulated spectrum of energy (MeV) deposited in events in veto panel 1, a panel on the top of the CDMS shield.

References

- [1] Zwicky, F. *Helv. Phys. Acta* **6**, 110 (1933).
- [2] Sofue, Y and Rubin, V. *Annual Review of Astronomy and Astrophysics* **39**, 137(2001).
- [3] Chang, Clarence. Doctoral Thesis, Stanford University (2004).
- [4] Weinberg, D., et al. *Astrophysical Journal* **490**, 564 (1997).
- [5] Sofue, Y et al. *Astrophysical Journal* **523** 136 (1999).
- [6] Sirko, E et al (SDDS), *Astronomical Journal* **127**, 914 (2004).
- [7] C. Alfonso et al., *Astron. and Astroph.* **400**, 951 (2003).
- [8] C. Alcock et al., *Astroph. J.* **542**, 281 (2000).
- [9] Sheldon, ES, et al (SDDS), *Astronomical Journal* **127**, 2544 (2004).
- [10] Knop, RA et al (SCP), *Astrophysical Journal* **598**, 102 (2003).
- [11] Gascon, J. arXiv:astro-ph/0504241, 2005.
- [12] Gaitskel, R, *Annual Review of Nuclear and Particle Science* **54**, 315 (2004).
- [13] Tonry, JL et al (High-Z Supernova Search Team), *Astrophysical Journal* **594**, 1 (2003)
- [14] Peebles, PJE, *Astrophysical Journal* **263**, L1 (1982).
- [15] Bennett, CL, et al (COBE), *Astrophysical Journal* **464**, L1 (1996)
- [16] Lang, A.E. et al (Boomerang), *Physical Review D* **63** 042001 (20001).
- [17] Spergel, DL et al (WMAP), *Astrophysical Journal Supplement.* **148**, 175 (2003).
- [18] Coutu, S, et al, *Astroparticle Physics* **11**, 429 (1999).
- [19] Haino, S, et al, *Phys. Letters B* **594**, 35 (2004).
- [20] Wright, EL, *New Astronomy Reviews* **47**, 87 (2003).
- [21] Hartman, RC, et al. *Astrophysical Journal Supplement* **123**, 79 (1999).
- [22] Tegmark, M et al (SDDS), *Physical Review D* **69**, 103501 (2004).

- [23] Lewin, JD and Smith, PF, *Astroparticle Physics* **6**, 87 (1996).
- [24] Linstrom, U. hep-th/0204016 (2002).
- [25] Mandic, Vuk. Doctoral Thesis, University of California, Berkeley (2004).
- [26] Mandic, Vuk. Private Communication, 16 Jan 2002.
- [27] Duffy, L. et al (ADMX) astro-ph/0505237 (May 2005).
- [28] de Freitas Pacheco, JA and Peirani, S. *Gravitation and Cosmology*, special issue **Cosmion** **2004**, 2004.
- [29] Bernabei, R., et al. *Physics Letters B* **424**, 195 (1998).
- [30] Bernabei, R., et al. *International Journal of Modern Physics D* **13**, 2127 (2004).
- [31] Abusaidi, R., et al. *Physics Review Letters* **84**, 5699 (2000).
- [32] Benoit, A., et al. *Physics Letters B* **513**, 15 (2001).
- [33] Kudryavtsev, V.A., et al. astro-ph/0406126 (2004).
- [34] Angloher, G., et al. *Astroparticle Physics* **23**, 325 (2005).
- [35] Akerib D., et al. *Physical Review Letters* **93**, 93 (2004).
- [36] Schnee, R., for CDMS collaboration. April APS meeting presentation (2005).
- [37] Sanglard, V. for EDELWEISS collaboration. proceedings of IDM2004, astro-ph/0411629 (2004).
- [38] Alner, G. J., et al. (UK Dark Matter Collaboration). *Astroparticle Physics* **23**, 444 (2005).
- [39] Aprile, E., et al. astro-ph/0503621 (2005).
- [40] Bertone, G, Hooper, C. and Silk, J. *Physics Reports* **405**, 279 (2005).
- [41] Gaitskell, R. and Mandic, V. http://dendera.berkeley.edu/plotter/entryform_cdmsprivate.html (2005).
- [42] Golwala, Sunil. Doctoral Thesis, University of California, Berkeley (2001).
- [43] Lindhard, J. *Mat. Fys. Medd. Dan. Vid. Selsk* **33** (1963).
- [44] Mahapatra, R. CDMS R119 note 107, 050213.
- [45] Qiu, Xinjie. UMN CDMS backgrounds internal note 18, 050712.
- [46] Sander, J. CDMS R119 note 81, 050125.
- [47] Thompson, J. P. talk give at CDMS collaboration meeting, July 2004.
- [48] Geant4 homepage: <http://geant4.web.cern.ch/geant4>.
- [49] Cassidy, G.L., Keuffel, J.W., and Thompson, J.A. *Physics Review D* **7**, 2022 (1973).

- [50] Eichblatt, S. CDMS internal note 9701011, 27 Jan 1997.
- [51] Yellin, S., and Hennings-Yeomans, R. Work presented at CDMS monte carlo meeting 15 Oct 2004.
- [52] Mahapatra, Rupak. CDMS R119 note 155, 050313.
- [53] Sander, Joel. CDMS R119 note 191, 050321.
- [54] Ogburn, R. Walter. CDMS R119 note 171, 050316.
- [55] Duong, L. private communication, 9 Feb 2005.
- [56] Yellin, Steven. Physics Review D **66**, 032005 (2002).
- [57] Fillippini, J, et al. "First Limits on Spin-Dependent WIMP-Nucleon Interactions from CDMS". astro-ph(050000).
- [58] CDMS collaboration. Super-CDMS proposal submitted to NSF and DOE. Oct 2004.
- [59] Details on instrumentation at the UMN Characterization Facility.
http://www.charfac.umn.edu/CF_Brochure2001.pdf
- [60] <http://vulcan2.case.edu/groups/Ernst/nec.html>
- [61] Thomas, J.H. III. report to P. Cushman October 2003.