

**Background Characterization and Discrimination in the
Final Analysis of the CDMS II Phase of the
Cryogenic Dark Matter Search**

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

Matthew Christopher Fritts

**IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy**

February, 2011

© Matthew Christopher Fritts 2011
ALL RIGHTS RESERVED

Acknowledgements

I'd like to thank a lot of the people, in order of appearance, who helped me on this path.

First my family, non-scientists all in the strictest sense at least, but all unflinching science fans. In particular my mother, from whom I get much of my sense of humor and other good qualities, and my father, whose love of reason and collection of science fiction and fact got me started in this direction.

My terrific high school physics teacher, Mr. Gish, helped convince me that physics is the way. I suppose other responsible parties include Jeff Howe and a guy named Robert Anton Wilson (believe it or not). Mr. Gish recommended the University of Missouri-Rolla where I met Timothy Gay who gave me my first physics job and who sometimes described elegant mathematical derivations as “sexy”.

Rob Toth has been calling me “Doctor” for nearly 20 years. I'd like to express my appreciation for that.

Thanks to Dick Schechtman and Larry Curtis who gave me another physics job after a decade-long break from the business. Monica deserves thanks for this happening as well, but I cover that in the Dedication. Thanks to others at the University of Toledo Department of Physics, for easing me back onto the graduate school path.

My fellow zeroth-years at Minnesota — Charlie Blackwell, Pearl Sandick, Andrew Thompson — helped me study for the written exam and gave me the impression that physics grad students are a cool and attractive bunch. I beat you to the written exam, but I'm not surprised that you all got doctorates before me. Thanks to a lot of others from my cohort, like Abe DeBenedetti, Beth Lusczek, Scott Thaller, and Jesse Woodroffe, for being significantly more than study partners. And thanks for being so young so I could do my cranky-old-man act. Also thanks to Lawren Ripton for hot

coffee and cool conversation.

Woods Halley is a good friend who got me out of some TA work and served as match-maker between me and a dilution refrigerator nicknamed Little Blue. I had never heard of dilution refrigerators before that. I want to thank several theorists for indulgence and advice (not to mention great teaching), including Keith Olive, Tony Gherghetta, Marco Peloso, Y.Z. Qian, and Serge Rudaz. Now I know that experiment is better, because nature grades your work, and I always liked getting grades. Also thanks for guidance in the early grad school years to Jim Kakalios and Yuichi Kubota.

Thanks to my advisor Prisca, for inviting me into the family (as has been said of the Mafia, no one ever really leaves CDMS) and for serving as a constant reminder that physics is fun and exciting. Also thanks for her significant contributions to the readability of this thesis. It was a sometimes painful process (for both of us I think) but the improvements were huge.

Thanks to the local CDMS team. The old guard — Long Duong, Angie Reisetter, Xinjie Qiu — all served as great mentors. The new guard — Oleg Kamaev, Scott Fallows, Roxanne Radpour, Tom Hofer, Hassan Chagani, Anthony Villano — have all been great to work with and hash out the details. Vuk Mandic spans both categories really, many thanks to him as well. And additional thanks to Scott for his sophisticated tastes . . . and on a related note thanks to Italian Spiderman for inspiration.

The CDMS collaboration has been a wonderful group to be a part of. I appreciate working with all of you, but let me single out just a few with whom I've greatly enjoyed working closely: Zeesh Ahmed, Lauren Hsu, Dave Moore, and Oleg again, with whom I hacked through the Runs 125–8 analysis; and Dan Bauer and Jim Beaty, constant guiding presences for Soudan shifts, whether physical or virtual.

Thanks as well to the many loyal and hardworking undergraduates who've been part of the Little Blue crew, many of whom have been instrumental in the “undergraduate-proofing” of the operations manual. Particular thanks to Ashley Rumney, who always cheered me up, Mark Williams, and Kirby Myers, who arranged therapeutic karaoke. Though he doesn't quite fit in this category I'll thank Andy Schofield here as well, with whom I've had many refrigerator and machine shop adventures.

Finally thanks to Simon, Billie, and Monax for their presence and unflagging friendship.

Dedication

To Monica, who gets the credit for all my success and happiness, except for the decision to be with her - that was mine.

Abstract

The Cryogenic Dark Matter Search (CDMS) is designed to detect Weakly-Interacting Massive Particles (WIMPs) in the Milky Way halo. The phase known as CDMS II was performed in the Soudan Underground Laboratory. The final set of CDMS II data, collected in 2007-8 and referred to as Runs 125–8, represents the largest exposure to date for the experiment.

We seek collisions between WIMPs and atomic nuclei in disk-shaped germanium and silicon detectors. A key design feature is to keep the rate of collisions from known particles producing WIMP-like signals very small. The largest category of such background is interactions with electrons in the detectors that occur very close to one of the faces of the detector. The next largest category is collisions between energetic neutrons that bypass the experimental shielding and nuclei in the detectors. Analytical efforts to discriminate these backgrounds and to estimate the rate at which such discrimination fails have been refined and improved throughout each phase of CDMS.

Next-generation detectors for future phases of CDMS require testing at cryogenic test facilities. One such facility was developed at the University of Minnesota in 2007 and has been used continuously since then to test detectors for the next phase of the experiment, known as SuperCDMS.

Contents

Acknowledgements	i
Dedication	iii
Abstract	iv
List of Tables	ix
List of Figures	x
1 Properties and Detection of Dark Matter	1
1.1 Known Dark Matter properties	2
1.2 Theorized Dark Matter properties	4
1.3 Direct detection principles	5
1.4 Terminology for direct detection experiments	11
2 Experimental Overview of CDMS II	15
2.1 CDMS detectors (ZIPs)	15
2.2 Shielding and Cosmic Ray Screening	20
2.3 Run History	21
2.4 Organizational Overview	24
3 Radiogenic Neutron Background	27
3.1 Origins	27
3.2 Identification and discrimination	28
3.3 Final background estimate	28

4	Cosmogenic Neutron Background	31
4.1	Origins	31
4.2	Identification	31
4.2.1	The muon veto cage	31
4.3	Discrimination criterion	37
4.4	Final Background Estimate	40
4.4.1	Monte Carlo simulations and observational benchmarks	40
4.4.2	Calculation of the background estimate	51
5	Bulk Electron Recoil Background	53
5.1	Origins	53
5.2	Identification	54
5.3	Discrimination criterion	56
5.4	Final background estimate	56
6	Surface Electron Recoil Background	59
6.1	Origins	60
6.1.1	Categories of surface interactions	61
6.1.2	Rate estimates for surface interactions in WIMP-search running	61
6.2	Identification	63
6.2.1	Phonon timing	64
6.2.2	Calibration data samples	64
6.2.3	Differences between calibration and WIMP-search surface interaction distributions	66
6.2.4	Phonon timing as a discrimination tool	68
6.2.5	Discriminator categories	70
6.2.6	Algorithms for extracting phonon pulse shape information	70
6.2.7	Position/energy correction of discriminators	75
6.3	Discrimination criterion	75
6.3.1	Combinations of discriminators	75
6.3.2	Estimating leakage and efficiency/exposure	77
6.3.3	Maximizing exposure on a ZIP-by-ZIP basis	80
6.3.4	Choosing a target leakage	82

6.3.5	Choosing a discriminator	83
6.4	Final background estimate	88
6.5	Additional comments	91
6.5.1	Lessons learned from discrimination studies	91
6.5.2	Evaluation of discrimination efforts	93
7	Results of the CDMS II Final Analysis	97
7.1	Unblinding the data	97
7.1.1	Preliminary expectations	97
7.1.2	Events satisfying WIMP selection criteria	98
7.2	Interpreting the results	102
7.2.1	Upper limit of WIMP-nucleon cross section	102
8	Testing SuperCDMS Detectors	107
8.1	Purpose of cryogenic tests	107
8.2	Design and construction of the testing facility	110
8.2.1	Preexisting equipment	110
8.2.2	New refrigerator hardware	112
8.2.3	Other modifications	114
8.3	Testing techniques	117
8.4	Summary of Little Blue test runs	121
	References	123
	Appendix A. Little Blue Operations Manual	127
A.1	Install payload	127
A.2	Install refrigerator in cryostat	131
A.3	Pump out fridge circulation lines.	134
A.4	Cool to 77 K.	137
A.5	Cool to 4 K.	139
A.6	Cool to 1.4 K.	142
A.7	Cool to base.	144
A.8	Monitoring and troubleshooting.	146

A.9 Warm up.	150
A.10 Cleaning mixture.	154
A.11 Liquid helium transfer.	158

List of Tables

2.1	Run history of CDMS II	23
4.1	Observed cosmogenic neutron interactions	44
4.2	Observed cosmogenic neutron events	44
6.1	Rates of surface interactions with WIMP-like signals	62
6.2	Fractions of surface interactions in various detector-face and energy categories	63
6.3	Detector-face ratios of surface interaction distributions	68
6.4	Expected number of surface interactions with WIMP-like signals in Runs 125–8	79
6.5	List of surface event discrimination criteria	87
6.6	Phonon pulse fit algorithm efficiencies	92
6.7	Details of detector-varied discrimination criteria	95
6.8	Leakage estimates for detector-varied discrimination criteria	95
7.1	Summary of expected backgrounds	98

List of Figures

1.1	Theoretical predictions of WIMP properties	6
2.1	Map of collision energy information in a ZIP	17
2.2	CDMS detector (ZIP)	19
2.3	Arrangement of ZIPs	20
2.4	Shielding configuration	22
3.1	Comparison of phonon rates from Runs 123 and 125	29
4.1	Muon veto cage	33
4.2	Photo of muon veto panels	34
4.3	Muon veto panel traces	35
4.4	Muon veto panel trace maxima	36
4.5	Surface interaction characteristics of cosmogenic neutron candidates . .	43
4.6	Cosmogenic neutron interaction recoil energies and ionization yields . .	45
4.7	Cosmic ray event with delayed interactions	46
4.8	Ionization yield distribution of delayed interactions	48
4.9	Delay time distribution of delayed subevents	49
4.10	Recoil energy distribution of delayed subevents	50
5.1	Recoil-type discrimination via ionization yield	55
6.1	Comparison of phonon pulses for bulk nuclear recoil vs. surface electron recoil	65
6.2	Energy distributions of surface electron recoils	67
6.3	Surface interaction discrimination via phonon pulse shape	69
6.4	Surface interaction discriminator categories	71
6.5	Rise time/fall time walk algorithm	73
6.6	Pulse fitting with PipeFitter	74

6.7	Discrimination based on a linear combination of discriminators	76
6.8	Discrimination including consistency criteria	78
6.9	Steps to calculate surface interaction leakage vs. discriminator threshold	81
6.10	Results of discrimination optimization	82
6.11	Experimental sensitivity as a function of target leakage	84
6.12	Primary phonon pulse delay with respect to ionization pulse	92
7.1	Phonon timing characteristics of WIMP candidates	100
7.2	Hypothetical upper limit for different target surface interaction leakages	101
7.3	WIMP-nucleon cross section upper limit	104
7.4	Cross section upper limit superimposed with theoretical predictions . . .	105
8.1	Schematic of ZIP phonon sensor	109
8.2	Resistance vs. temperature of a phonon sensor array	109
8.3	Photos of Janis CF-25 dilution refrigerator	111
8.4	New refrigerator hardware	113
8.5	Installation of new hardware	115
8.6	Superconducting transition of a phonon sensor array	118
8.7	Critical current measurements	119
A.1	Little Blue gas-handling system	135
A.2	Little Blue pressure-relief system	148
A.3	Little Blue mixture cleaning cycle	155

Chapter 1

Properties and Detection of Dark Matter

Astronomical and cosmological observations on a huge range of scales, from single galaxies up to the entire visible universe, indicate that the majority of the matter in the universe is not composed of any known particle. These observations can be summarized as follows:

- The orbital velocities of material outside of the visible disk of galaxies imply the existence of an invisible massive halo extending far past the disk [1].
- The total mass of clusters of galaxies as inferred by the velocities of the constituent galaxies [2], the temperature of the x-ray-emitting intergalactic gas [3], and gravitational lensing [4] far exceeds the mass of the visible matter.
- Cosmological measures of the matter density of the universe, utilizing such probes as the cosmic microwave background anisotropy and baryon acoustic oscillation studies, indicate that the total mass density is approximately five times that of the mass density of baryonic matter (i.e. matter whose mass is primarily protons and neutrons) [5].

In what follows I will usually use the expression *dark matter* as shorthand for the majority non-baryonic component. Technically this is inappropriate since in its broad sense the term includes subdominant components such as baryonic dark matter (massive

compact halo objects), relic neutrinos produced in the big bang, and possible contributions from multiple new particles.

1.1 Known Dark Matter properties

Dark matter is a fundamental part of the picture for astronomers and cosmologists and a very active area of study. From these studies a number of properties can be deduced or inferred, both microscopic (particle properties) and macroscopic (distributions).

First, dark matter is, indeed, *matter*. The particles have mass and interact gravitationally. So far, all that is known about dark matter (apart from inferences from non-observations) is derived from its gravitational effects.

Secondly, dark matter is *dark*. This term is meant both to evoke our lack of knowledge and, more specifically, to indicate that dark matter does not interact in the ways that make baryonic matter visible. Not only does dark matter not emit, absorb, or scatter light; it also travels through matter unimpeded, as dramatically demonstrated by gravitational lensing maps of dark matter distributions in e.g. the Bullet Cluster [6]. To astronomers it is sufficient to describe dark matter as entirely non-interacting (excluding gravity of course); this is a working assumption in large-scale models and simulations of dark matter. From a microscopic viewpoint, this property leads to a clear conclusion — dark matter does not interact by either of the two strongest interactions: the strong nuclear interaction and electromagnetism. This leaves open the question of whether it interacts by the weak nuclear interaction.

A third property of dark matter is that it is *cold*. This qualifier is related to dark matter's rôle in structure formation. The very small density fluctuations of the universe as seen in the cosmic microwave background must have grown to produce the large-scale structures seen today. This growth must be primarily due to dark matter particles that were moving at speeds much slower than the speed of light. If dark matter particles were moving at relativistic speeds at the time that structure formation was beginning, density fluctuations would tend to shrink rather than grow [7], and the large-scale structure seen today could never have developed. The term *cold* defines particles moving at non-relativistic speeds at the onset of structure formation. This property can be used to make some inferences about the fundamental properties of dark matter particles. For a

dark matter particle produced thermally (as all known particles are), and thus originally in thermal equilibrium with the universe, a lower limit on the particle's mass can be determined. This rules out the standard model neutrinos as candidates for dark matter particles.

Dark matter particles must also be stable. Of the long list of known particles, most don't exist long after being created — it is the tendency of fundamental particles to decay into lighter particles, unless some conservation law prevents it. While the characteristic lifetimes of unstable particles cover a vast range, the familiar decay processes produce lifetimes that are short on cosmological scales. Because they are abundant in the universe today, dark matter particles must be stable, or at least have an extremely long lifetime.

The body of knowledge on the macroscopic properties of dark matter — namely its density and velocity distributions — is growing at a steady rate. Of particular relevance to terrestrial direct detection experiments are the details of these distributions in our astronomical neighborhood. Through modeling of structure formation and measurements of velocities of nearby visible matter, it is known that the dark matter halo of the Milky Way galaxy has the following features:

- It is nearly spherically symmetric, with an approximate radial density distribution proportional to $1/r^2$.
- The velocity of its constituent particles has an approximate isotropic Maxwellian distribution.
- At the galactic radius of our solar system¹:
 - the dark matter mass density is approximately $0.3 \text{ GeV}/c^2/\text{cm}^3$.
 - the characteristic dark matter speed relative to Earth is approximately 220 km/s.
 - the galactic escape speed is approximately 544 km/s. (Whatever the specific dark matter velocity distribution, this represents a maximum cutoff.)

¹These values are typically used as standards by which interpretations of different direct detection experiments are compared; occasionally different values are used based on developments in the astronomy literature.

1.2 Theorized Dark Matter properties

A natural first step in trying to guess at the nature of dark matter particles is to draw inferences about how they were created. For known particles, we have a very good understanding of how the early conditions and subsequent evolution of the early universe led to their relic abundances. Indeed Big Bang Nucleosynthesis (BBN) is a major success story for the big bang model of cosmology. Our experimental knowledge of nuclear processes can be applied to a hot, dense universe whose evolution is governed by the laws of general relativity. In so doing the relic abundances of the various isotopes of the lightest elements can be predicted, and the results agree with astronomical observations (apart from lithium, which remains something of a mystery).

If dark matter particles are created in this well-understood way, we should be able to draw conclusions about their interaction properties through our knowledge of their relic density. For example, if we assume that dark matter particles interact at the characteristic strength of the weak interaction, their relic density itself can be used to set the scale of the mass of an individual dark matter particle. The result is on the order of $m_\chi \approx 100 \text{ GeV}/c^2$, i.e. about 100 times as massive as the proton. The result is intriguing because this is the same mass scale as the gauge bosons that carry the weak interaction.

The expression *Weakly-Interacting Massive Particle* (WIMP) is used to refer to the class of proposed dark matter particles that interact via the weak nuclear force and with mass on the order of $100 \text{ GeV}/c^2$ (within a couple of orders of magnitude). What is particularly compelling about WIMPs as dark matter candidates is that the existence of a stable WIMP is predicted by particle theory quite independently of dark matter considerations.

The Standard Model of particle physics was formulated in the 1970s. No experiment to date has definitively contradicted it, and no particle has been discovered that is not a part of it. Despite its great success, it is known to be incomplete and contains suspicious “unnatural” features that can be solved by theories beyond the standard model. While no collider experiment has yet achieved the energy scale to test such theories, many of them point to new physics at weak-scale energies, which should be achievable by the current generation of collider.

A compelling class of theories predicting WIMPs are those involving supersymmetry, a conceptually natural extension to the Standard Model that predicts the existence of a supersymmetric partner to every known particle. None of these superpartners has yet been detected (indicating that they should have high enough masses to have escape detectable production in collider experiments). Many theories with supersymmetry naturally predict a stable WIMP (usually the *neutralino*, a superpartner to the neutral electroweak gauge bosons).

There are many good reviews summarizing the relationship between supersymmetry and dark matter [8]. In terms of efforts to directly detect dark matter, however, the details are not of critical importance. Supersymmetry primarily serves to motivate a search for WIMP dark matter. Specific formulations of supersymmetric theories somewhat constrain parameters such as mass and interaction cross-section, producing hypothetical maps (figure 1.1) to bolster our courage as we explore unknown territory.

Thus cosmological arguments and independent theoretical predictions strongly motivate the hypothesis that dark matter is composed of WIMPs. There do exist, however, other dark matter candidates. Probably the next most attractive is the axion. Axions were also originally proposed as extensions to the Standard Model independent of dark matter considerations, and they satisfy all of the necessary properties described in the previous section. However they are not produced thermally, and their properties (and corresponding detection strategies) are significantly different. The axion dark matter hypothesis is not quite as compelling because axion theories don't automatically produce the correct dark matter abundance. While the Cryogenic Dark Matter Search (CDMS) is designed specifically to search for WIMPs, it has also been applied to indirect axion detection [13], a testament to the flexibility of the experiment and the creativity of the collaboration.

1.3 Direct detection principles

Dark matter in the Milky Way halo is constantly passing through the earth. Because ordinary matter interacts via the weak nuclear force as do WIMPs, WIMPs should in principle interact with ordinary matter. Calculating the details of such interactions sheds light on the feasibility of direct detection. A useful review of these details was

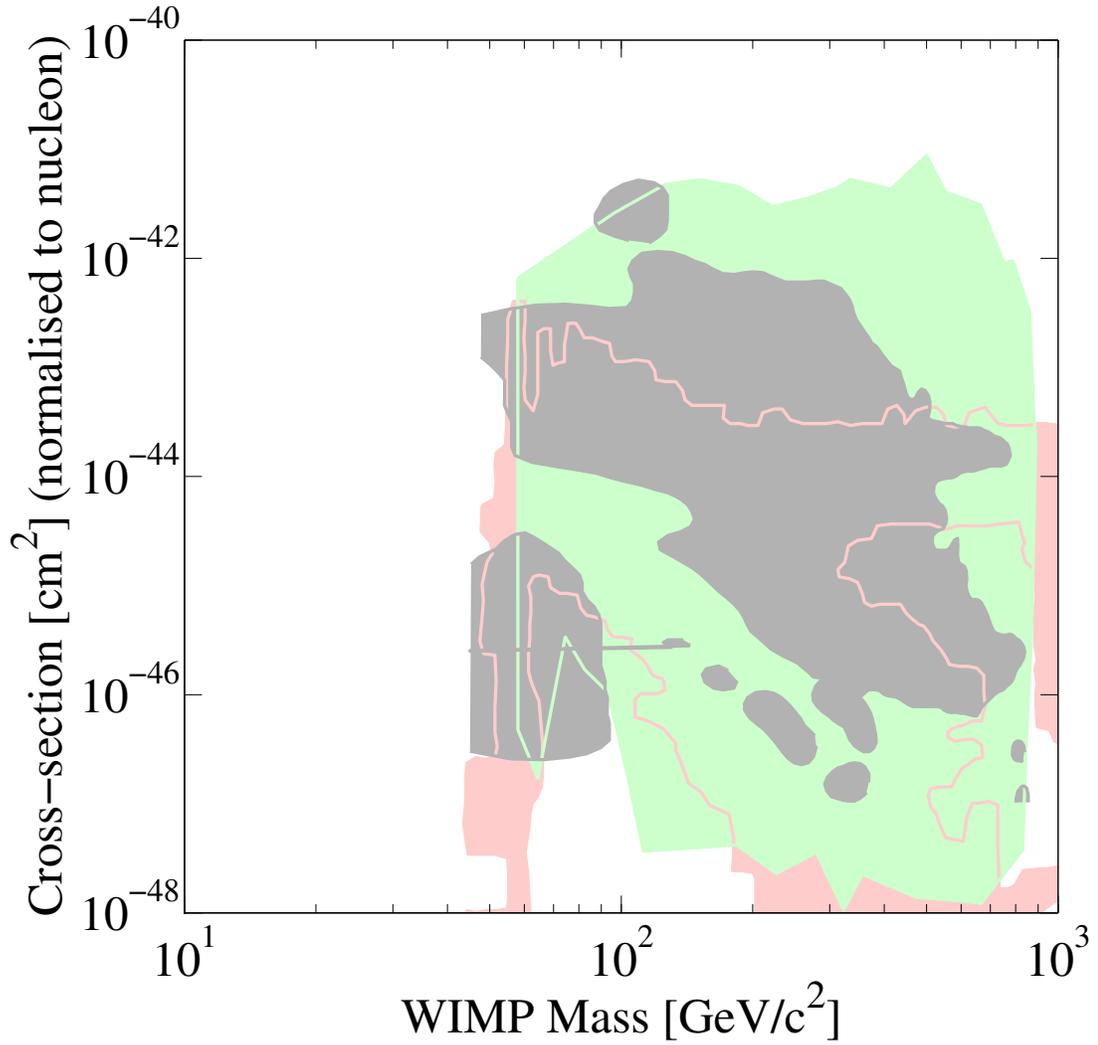


Figure 1.1: Predicted WIMP properties from supersymmetric theories utilizing experimental constraints. Grey: Constrained Minimal Supersymmetric Standard Model (cMSSM) [9]. Light green: Low-Energy Effective Supersymmetric Theory (LEEST) [10]. Pink: Minimal Supergravity (mSUGRA) [11]. (Figure courtesy [12])

provided by J. D. Lewin and P. F. Smith in 1996 [14] and has served to inform strategies for direct detection experiments ever since. What follows is a brief summary of the features most relevant to understanding the design principles and analytical interpretations of CDMS (and related experiments).

Direct detection of WIMPs is challenging for two reasons. First, the weak nuclear force is of very short range, with correspondingly low interaction rates. We can draw a comparison to neutrino detectors, which despite large detector masses have very small signal rates. However the neutrinos detected have nuclear-scale energies or greater, since they are typically created in reactors, accelerators, or astrophysical processes. An additional challenge for detecting halo WIMPs is due to their low kinetic energies. (A more apt comparison may be the notoriously difficult prospect of detecting the low-energy relic neutrinos produced in the big bang.)

Considering the galactic escape velocity, a halo WIMP of rest energy 100 GeV cannot have greater than 200 keV kinetic energy; characteristic kinetic energies are on the order of a few 10s of keV. This is rather small in terms of experimental particle physics, in the range of hard x-rays (and some 9 orders of magnitude smaller than collision energies at the Large Hadron Collider).

The most obvious way to detect halo WIMPs is to look for the recoil of an ordinary matter target after an elastic collision with a WIMP. The characteristic fraction of the WIMP's kinetic energy that is transferred in the collision, from simple kinetic arguments, is

$$\left\langle \frac{E_{recoil}}{E_{kinetic}} \right\rangle = \frac{4M_W M_T}{(M_W + M_T)^2} \quad (1.1)$$

with M_W and M_T the respective masses of the incident WIMP and the target particle. This is maximum when WIMP and target have the same mass, falling off the greater the mismatch in their masses. Therefore the signal is maximized by choosing a target with the same mass as a WIMP. Considering the range of possible WIMP masses (50 GeV/c² or greater) this suggests a heavy atomic nucleus as an appropriate target. Since the de Broglie wavelength of an incident WIMP is on the order of a fm, to first approximation a nucleus serves as a single target.

Even more important than maximizing the collision recoil energy is to maximize the probability of a collision, i.e. to maximize the collision cross section σ . It has been

shown [15] that in the relevant non-relativistic limit the interaction strength between WIMPs and nucleons can be characterized by four parameters: f_p , f_n , a_p , a_n . These appear in the general interaction Lagrangian expression

$$\mathcal{L}_{\chi N} \approx 4\chi^\dagger\chi \left(f_p\eta_p^\dagger\eta_p + f_n\eta_n^\dagger\eta_n \right) + 16\sqrt{2}G_F\chi^\dagger\frac{\vec{\sigma}}{2}\chi \cdot \left(a_p\eta_p^\dagger\frac{\vec{\sigma}}{2}\eta_p + a_n\eta_n^\dagger\frac{\vec{\sigma}}{2}\eta_n \right) \quad (1.2)$$

This represents the fundamental interactions between WIMPs and nucleons (protons and neutrons). In this expression, χ is the WIMP wave function, η_p and η_n are the proton and neutron Weyl spinors, and $\vec{\sigma}$ is the spin operator. The first term represents so-called spin-independent (SI) interactions, while the second indicates the possibility of coupling of the spins of the WIMP and the target, so-called spin-dependent (SD) interactions.

In the SI case, the collision cross section can be calculated from (1.2) to be

$$\sigma_0^{SI} = \frac{4}{\pi}\mu_{WT}^2 [Zf_p + (A-Z)f_n]^2 \quad (1.3)$$

where the subscripted 0 indicates that this holds in the extreme nonrelativistic limit. A and Z are the mass number and atomic number of the target nucleus. μ_{WT} is the reduced mass of the WIMP and the target nucleus,

$$\mu_{WT} = \frac{M_W M_T}{M_W + M_T} \quad (1.4)$$

All of this holds true for a generic WIMP. But we can simplify further if we assume $f_p \approx f_n$ as is true in supersymmetric models. We then find

$$\sigma_0^{SI} = \frac{4}{\pi} \left(\frac{M_W M_T}{M_W + M_T} \right)^2 (A f_{p,n})^2 \quad (1.5)$$

Considering that M_T is proportional to A , σ_0^{SI} increases very strongly with increasing A : at least as strongly as A^3 (A^4 in the limit $M_W \gg M_T$). If the cross section were directly proportional to the target mass number, the rate at which collisions occur would be insensitive to the choice of target nucleus for a given total detector mass. But since the cross section has a much stronger dependence on mass number, detectors with heavier atomic nuclei will have a large enhancement to their WIMP collision rate.

The enhancement due to large mass number is due in part to a coherence effect that is responsible for the squaring of the bracketed term in (1.3). This is related to the nucleus appearing as a single target to the WIMP. However, this only holds at the limit of zero momentum transfer; a correction is necessary as the recoil energy increases and internal structure of the nucleus becomes increasingly important. The net result is that the mass number enhancement is weakened, and the gains achieved at higher target mass taper off, particularly at the high end of the recoil energy spectrum.

The SD cross section can also be derived from (1.2). Some of the features of SD cross sections [16] are:

- The strong mass number enhancement is not present.
- Nuclei with large spins are favored.
- Nuclei with an odd number of protons or an odd number of neutrons are favored.

Thus to probe SD WIMP collisions, nuclear isotopes with non-zero spins are the preferred targets. In general more effort has been put into designing experiments for SI collisions, both because total cross sections are expected to be greater and because suitable target isotopes are more readily available.

Armed with the information above we can predict the expected interaction rate and recoil energy spectrum for SI collisions of halo WIMPs with target nuclei. The details depend on the two least-known parameters (WIMP mass and $f_{p,n}$, the latter usually expressed in terms of the SI cross section between a WIMP and a single target nucleon). However regardless of the specific values, the recoil energy spectrum will be approximately a decreasing exponential with characteristic energy in the tens of keV, and the overall rate will be small because WIMPs interact weakly. The WIMP signal will therefore face competition from familiar forms of radiation that interact much more strongly. Most of the effort in experimental design goes into dealing with potential background signals from collisions between known particles and the matter in the detector.

The two natural sources of background signals are radioactive matter in the vicinity of the experiment and cosmic rays. The relevant types of radiation produced by these sources are α , β , γ , and n. Unlike WIMPs, all of these particles can be blocked by

physical shielding; however, such shielding can never be 100% effective. Some energetic cosmic ray muons will penetrate even to the deepest sites. Shielding against backgrounds from radioactive materials can be nearly perfect, but only against radioactivity external to the shield. To reduce radioactivity internal to the shield — in the shield material itself, in the detector hardware, and in the detector mass — the most radiopure materials available are used. Invariably, however, there is enough residual radiation to produce background signals at a much higher rate than that of a potential WIMP signal.

One class of strategies to extract a WIMP signal in the presence of background exploits differences between the global features of the background signal and the expected WIMP signal. If the background spectrum is understood well enough, it can be subtracted from the total spectrum, leaving any residual WIMP signal. This strategy has been employed by CoGeNT [17]. Alternatively one can look for a feature of the WIMP signal that is not expected in any background signal. As mentioned above, the expected WIMP spectrum is essentially free of distinguishing features such as peaks. However an additional expected feature of a WIMP signal is a spectrum that changes over the course of a year due to the earth's velocity relative to the sun. This motion will alternately enhance and diminish the average relative velocity of the WIMP wind, which will in turn modify the WIMP signal spectrum and overall rate. A (WIMP + background) signal could be studied over time for signs of annual modulation, which would serve as a WIMP-signal signature (assuming an annual modulation in any background source can be ruled out). This strategy is used by DAMA/LIBRA [18].

In contrast to global approaches, a different strategy for extracting a WIMP signal is to exploit the differences between WIMP interactions and background particle interactions. The most prevalent background is typically photons from radioactivity inside the shielding. When a photon of the relevant energy collides with an atom, it collides with one of its electrons, rather than the nucleus. In a sense the atom's electrons serve as shielding for the true WIMP target: the nucleus of the atom. (WIMP-electron collisions are not impossible, but the cross section is much smaller and the deposited energy typically unmeasurable.) Generally the type of particle detectors used in WIMP search experiments produce a collision signal whose strength is proportional to the recoil energy. However most such detectors respond to electron recoils and nuclear recoils differently. If two independent signals are produced by a collision, the difference between

the two in terms of recoil-type response can be used to distinguish nuclear recoils from electron recoils. Such *two-channel detectors*² are employed by several collaborations, including CDMS. Detailed time resolution of the collision signal can produce pulse-shape information that also may be useful in distinguishing recoil type.

Even if electron recoils can be distinguished from nuclear recoils, there remain known radiation products that produce nuclear recoils and hence could produce a collision signal indistinguishable from a WIMP collision signal. Neutrons have sufficient penetrating power to present an irreducible background in most experiments. In particular neutrons produced by cosmic rays can appear at a significant rate even in experiments with extensive shielding.

Apart from physical shielding and recoil-type discrimination, another tool for removing background is *coincidence*. If multiple, physically distinct detector volumes exist, a strongly-interacting background particle has a significant probability to collide in multiple detectors, producing (nearly) simultaneous signals. In contrast a WIMP can only be expected to interact once if at all. Another way in which multiple simultaneous collision signals can occur is if the particle is produced in a process that produces multiple particles, as for example the showers of particles produce by cosmic ray muons interacting with matter.

1.4 Terminology for direct detection experiments

The following terms are used frequently in direct detection WIMP searches, which are by nature rare event searches. Usage can vary; the definitions given here are the sense in which I will use these terms in this thesis.

Event: information from all detectors in an experiment collected simultaneously over a restricted time interval. *Coincidence* is evidence of simultaneous signals in two separate detector volumes.

Live Time: the total amount of time in which the detectors are ready to detect a new signal. For a period of data taking this is real time minus the intervals over which

²One-channel detectors, such as those employed by DAMA and CoGeNT, also respond differently to nuclear recoils vs. electron recoils, but the difference cannot be exploited. Two energy calibrations are necessary, one for each recoil type.

events are being recorded. If periods of data taking are later rejected due to high noise or other problems, the live time is adjusted accordingly.

Raw Exposure: the live time multiplied by the detector mass, typically expressed in $\text{kg} \cdot \text{days}$ or $\text{kg} \cdot \text{years}$. This is the basis for characterizing the rate of the WIMP signal, which is the primary aim of a direct detection experiment. The greater the exposure, the more sensitive the experiment is to a low-rate signal.

Background: signals from ordinary radiation that satisfy all WIMP identification criteria. Different background categories exist for different types of ordinary radiation. For a given experimental exposure, background is expressed in units of *events*. For low-background experiments the expected background is often less than one (an expectation value). Typically the background of an experiment falls between two extremes:

- **zero-background:** a total background that is much less than one event for a given exposure. In a zero-background experiment the sensitivity is directly proportional to the exposure.
- **background-limited:** a total background that is proportional to the exposure, and which cannot be efficiently separated from the WIMP signal by global features. In a background-limited experiment the sensitivity cannot be increased by increasing the exposure.

Discrimination: identification of potential background that distinguishes it from a WIMP signal. Any events thus identified are eliminated as background.

Leakage: potential background that is not eliminated by some discrimination criterion. Leakage is expressed in units of *events*. The **leakage fraction** is the fraction of some set of potential background that is not eliminated by a discrimination criterion.

Efficiency: the fraction of potential WIMP signal events that satisfy one or more WIMP selection criteria, including background discrimination criteria. The **exposure** is the product of the raw exposure and the combined efficiency of all WIMP selection criteria.

Blind Analysis: a procedure to avoid unconscious bias in which the researchers finalize all WIMP selection criteria before examining any data that could contain an identifiable WIMP signal. Any conclusion drawn in a rare event search is much stronger if a blind analysis is performed. Data files from which potential WIMP signal events have been automatically removed are often referred to as “blinded.” “Unblinding” is the opening of a complete data set for analysis.

WIMP Candidate: an event satisfying all WIMP selection criteria. Any WIMP candidate is considered to be either background or a WIMP interaction. If the number WIMP candidates is significantly larger than the expected background this constitutes evidence of a WIMP signal.

Chapter 2

Experimental Overview of CDMS II

2.1 CDMS detectors (ZIPs)

The primary target element for CDMS (the Cryogenic Dark Matter Search) is germanium. Germanium nuclei make excellent targets for SI WIMP-collisions. At an atomic mass of $72.64 \text{ g} \cdot \text{mol}^{-1}$, the collision rate for WIMPs of mass $100 \text{ GeV}/c^2$ can be anywhere from 10^{-4} to 10^3 detectable collisions¹ per kg·year, considering the range of cross sections depicted in figure 1.1. (The higher two orders of magnitude have already been ruled out by experiment.)

Silicon (atomic weight 28.09) is a secondary target element for CDMS. The sensitivity of silicon to WIMP collisions is an order of magnitude less than germanium under the same assumptions listed above. It is even smaller for detectors of the same size, due to silicon's lower density. In the presence of a signal, an experiment with detectors made of different materials would have certain advantages: WIMP properties could be characterized in more detail, and it would be easier to rule out a neutron background because the ratio of signal rates in germanium vs. silicon is quite different for neutrons than for WIMPs. Because no signal has been detected yet, and because other techniques have proven effective at managing the neutron background (see chapters 3 and 4), silicon is

¹assuming a readily achievable energy threshold of 10 keV

essentially being phased out for future runs in favor of the more sensitive germanium.

The target atoms form a high purity semiconductor crystal which is held at a very low temperature using a ^3He - ^4He dilution refrigerator, at about 40 mK (hence the “cryogenic” part of the acronym). When a target recoils after a collision with an energetic particle, the recoil energy is deposited into the crystal in two parts: 1) valence electrons are excited into the conduction band producing electron/hole charge carrier pairs (*ionization*), and 2) high-frequency lattice vibrations in the crystal are produced (*primary phonons*). If a bias voltage is applied across the detector, the charge carriers will drift to the detector surfaces where they can be collected (the *ionization channel*). As they drift, the potential energy they have by virtue of the applied voltage is transmitted to the lattice via collisions (*Luke phonons* [19, 20]). Upon absorption into electrodes at the detector surface, the de-exciting charge carriers release their excitation energy in a third form of lattice vibration: *recombination phonons*. All three kinds of phonons are collected by phonon sensors on the detector surface forming a combined *phonon channel*. This necessitates operation at low temperatures to reduce the noise from thermal phonons.

The CDMS detectors, known as ZIPs (Z-dependent Ionization Phonon), are thus two-channel WIMP detectors. The way in which energy information from the collision is transmitted to the sensor channels is illustrated in figure 2.1. The true recoil energy can be recovered from the data by combining the information from the channels. The bias voltage energy as determined from the ionization channel information is subtracted from the total phonon energy, as illustrated in equation 2.1, in which N_q represents the number of charge carriers ionized:

$$E_{\text{recoil}} = E_{\text{phonon}} - N_q e V_{\text{bias}} \quad (2.1)$$

The proportion of the collision energy that is converted into ionization (rather than primary phonons) is different for electron recoils vs. nuclear recoils (about 3 times smaller for the latter). Thus recoil-type information is also easily extracted. Recoil type is characterized by a derived quantity referred to as *ionization yield*. First the ionization channel is calibrated using a radiation source that produces only electron-recoils. This “charge energy” gives the “electron-equivalent” recoil energy (units keVee),

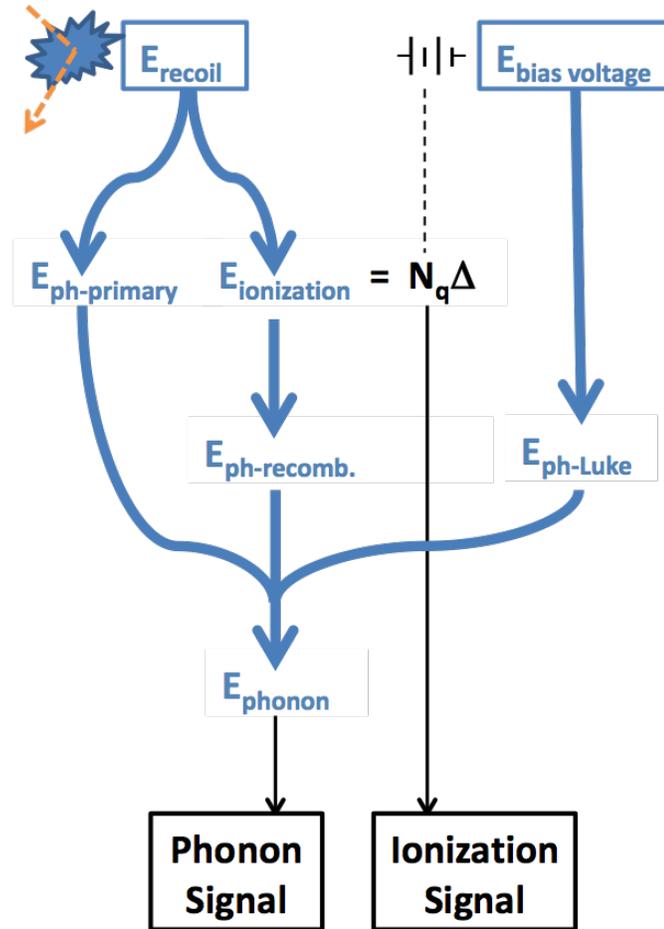


Figure 2.1: Schematic map of how energy information travels through a ZIP and is recovered in the sensor channels. The energy recorded in the phonon signal is the sum of the recoil energy and the potential energy of the charge carriers provided by the voltage biasing. This latter energy can be subtracted by means of information from the ionization signal. Recoil type information is recorded in the proportion of recoil energy that goes into ionization.

an accurate measure of the recoil energy for electron recoils but not for nuclear recoils.² This is illustrated in equation 2.2, in which Δ indicates the band gap energy and ϵ_{ee} represents the fraction of the recoil energy that is converted into ionization energy in the case of an electron recoil:

$$E_{\text{charge}(ee)} \equiv \frac{N_q \Delta}{\epsilon_{ee}} \quad (2.2)$$

Ionization yield is defined as the charge energy divided by true recoil energy (derived from information from both channels) as illustrated in equation 2.3:

$$Y \equiv \frac{E_{\text{charge}(ee)}}{E_{\text{recoil}}} \quad (2.3)$$

Ionization yield is thus 1 on average for electron-recoils, and around 0.3 for nuclear-recoils. This quantity is the basis for discrimination by recoil type.

Time resolution of the phonon channel pulse provides additional information relevant to the depth of the collision within the detector (which is the meaning of the “Z-dependent” portion of the ZIP acronym). This is important because if an electron recoil occurs very close to the detector surface, the ionization yield is often not properly characterized, so that an electron recoil may be misidentified. The phonons produced in surface interactions arrive at the phonon sensors characteristically faster than those produced by bulk interactions. This phenomenon is understood in terms of phonon propagation. The three types of phonons discussed above are produced with different characteristic frequencies, which affect the speed at which they propagate through the crystal. At low enough frequencies, the mean free path of a phonon is greater than the size of the detector; these ballistic phonons travel fastest, at the speed of sound. At higher frequencies scattering processes become important, slowing the phonon. The primary phonons produced by a particle interaction are mostly of the high frequency (slower) type. Over time, scattering processes cause anharmonic decay of the phonon, producing multiple lower frequency phonons. Thus high frequency phonons eventually down-convert to ballistic phonons.

When a particle collides very near either surface of a ZIP, rapid down-conversion

²Using only the ionization channel information in this way is equivalent to treating CDMS as though it were a one-channel experiment. One-channel experiments generally report energy spectra in keVee.

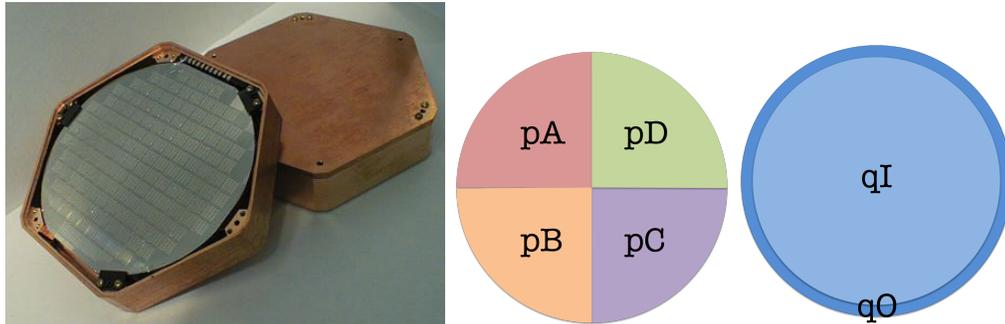


Figure 2.2: Photograph of ZIP in housing and schematic showing the layout of the four phonon (pA-pD) and two ionization (qI, qO) sensor arrays.

of the primary phonons occurs due to their interactions with the metal films on the surface. This is the main reason that surface interactions produce faster phonon signals. Luke phonons are important to phonon pulse timing as well. Unlike primary phonons, Luke phonons are generally ballistic. Since electron recoil interactions produce a larger fraction of Luke phonons than nuclear-interactions do, electron recoils generally produce faster pulses. The two phenomena combine to allow a means of discrimination between surface electron recoils and bulk nuclear recoils.

Figure 2.2 shows the layout of a ZIP. The detectors are 7.6 cm diameter, 1.0 cm thick disks with sensor arrays photolithographed onto each face. One face (the “phonon side”) is divided into four quadrants, each an independently-instrumented array of phonon sensors. Also on the phonon side is a single electrode grid for the ionization channel. The opposite face (the “charge side”) is divided into an inner circular region and an outer annular region, which contain independent electrode grids for two ionization sensors. This arrangement yields much information concerning collision location.

The ionization signals are amplified via Field-Effect Transistors (FETs). Each phonon quadrant consists of 1036 Quasiparticle-trap-assisted Electrothermal-feedback Transition-edge sensors (QETs): phonons are absorbed in superconducting aluminum fins and converted to quasiparticles which are trapped in superconducting tungsten films held at their critical temperature (approximately 80 mK). The feedback circuit response is measured via Superconducting Quantum Interference Devices (SQUIDs). Details of these sensors have been reported by Akerib et. al. [21].

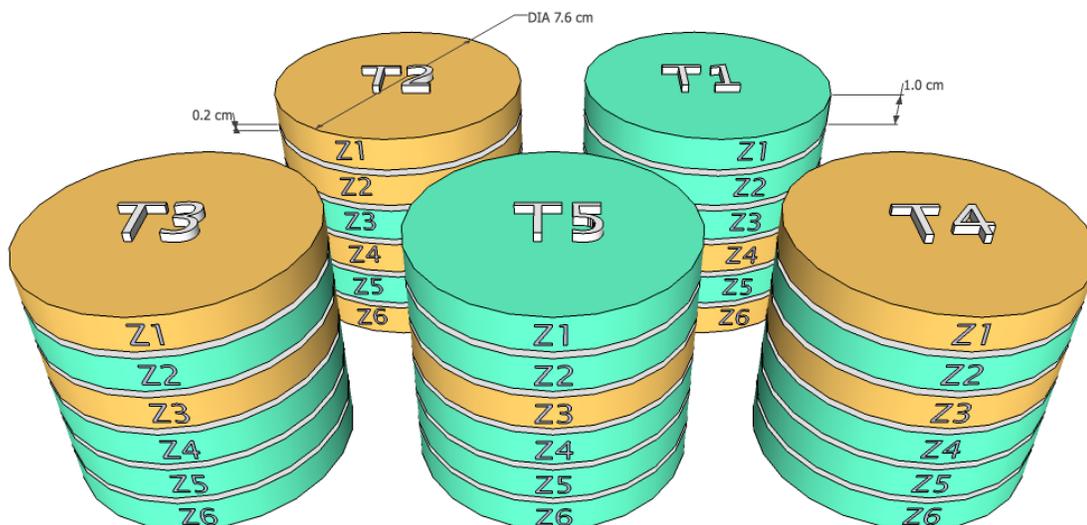


Figure 2.3: Physical arrangement of the ZIPs in the final CDMS II configuration. The detectors are labeled by tower (T1-T5) and position within tower (Z1-Z6) (e.g. “T2Z5”). ZIPs are color-coded by type in the diagram: beige for silicon, aqua for germanium.

The ZIPs are arranged in stacks of six inside a hexagonal copper housing above which are situated the cards containing the sensor circuit elements; this assembly is referred to as a “tower”. There are five towers and thus 30 ZIPs, 19 composed of germanium and 11 of silicon. Due to detector failures or poor performance some ZIPs are not used for WIMP search, but all 30 serve as coincidence detectors. The arrangement is shown in figure 2.3.

2.2 Shielding and Cosmic Ray Screening

The ZIP towers are inside the innermost of six nested copper cans, each connected to a different temperature stage of the cryogenic system by concentric copper tubes extending through the layers of shielding. The wires from the internal ZIP electronics are run through a second, similar set of tubes. Low radioactivity material is used for the cans and other internal hardware. The cans and housings provide an average of 3 cm of copper shielding, effectively shielding the ZIPs from alpha and beta radiation external to the cans. The space between the cans is under vacuum.

A 7.6 to 10.2 cm layer of polyethylene surrounds the copper cans, primarily intended to moderate any neutrons emitted by radioactive contamination in the surrounding lead shielding. Two layers of lead shielding serve to absorb external gamma radiation. The inner 4.4 cm thick layer is ancient lead, which greatly reduces its radioactivity³. The outer lead layer provides an additional 17.8 cm. An outer 40 to 40.6 cm thick polyethylene layer surrounding the lead serves to moderate neutrons from external radioactivity. The entire shielding volume outside of the copper cans is purged continuously with pure nitrogen gas, providing a radon-free environment.

Particles produced by cosmic ray muons can have sufficient energy to penetrate all these layers. For this reason the shielding is surrounded by muon detectors that serve as coincidence detectors so that cosmogenic backgrounds can be eliminated. This is referred to as the “muon veto cage”. Figure 2.4 shows the arrangement of the shielding and muon veto. The experiment is housed in a class-10,000 clean room in order to minimize radioactive contamination of the detectors and hardware during handling. This room is RF-shielded against external electromagnetic signals that could produce noise in the sensitive detectors.

From 2003 until the end of CDMS II in 2009, the experiment was located in the Soudan Underground Laboratory in northern Minnesota. The rock overburden is 713 meters (2090 meters water equivalent) thick, reducing the cosmic ray muon flux by a factor of 50,000.

2.3 Run History

Data-taking periods are defined by refrigerator runs during which the detectors are kept continuously at sub-kelvin temperatures. The ZIPs must be tuned and calibrated at the beginning of a new run. Typically, multiple runs are performed with a given configuration of detector towers and analyzed as a group. Table 2.1 summarizes the entire run history of CDMS II. Before the move to Soudan, early CDMS II runs were performed at the Stanford Underground Facility, a shallow site (10.6 m, 17 meters water equivalent) with a cosmic ray muon flux approximately 300 times greater than that at

³Freshly-mined lead contains trace amounts of the radioactive isotope ^{210}Pb with a half-life of 22 years; in lead taken from ancient sources such as shipwrecks, isolated from its natural environment for hundreds of years, this isotope is essentially absent.

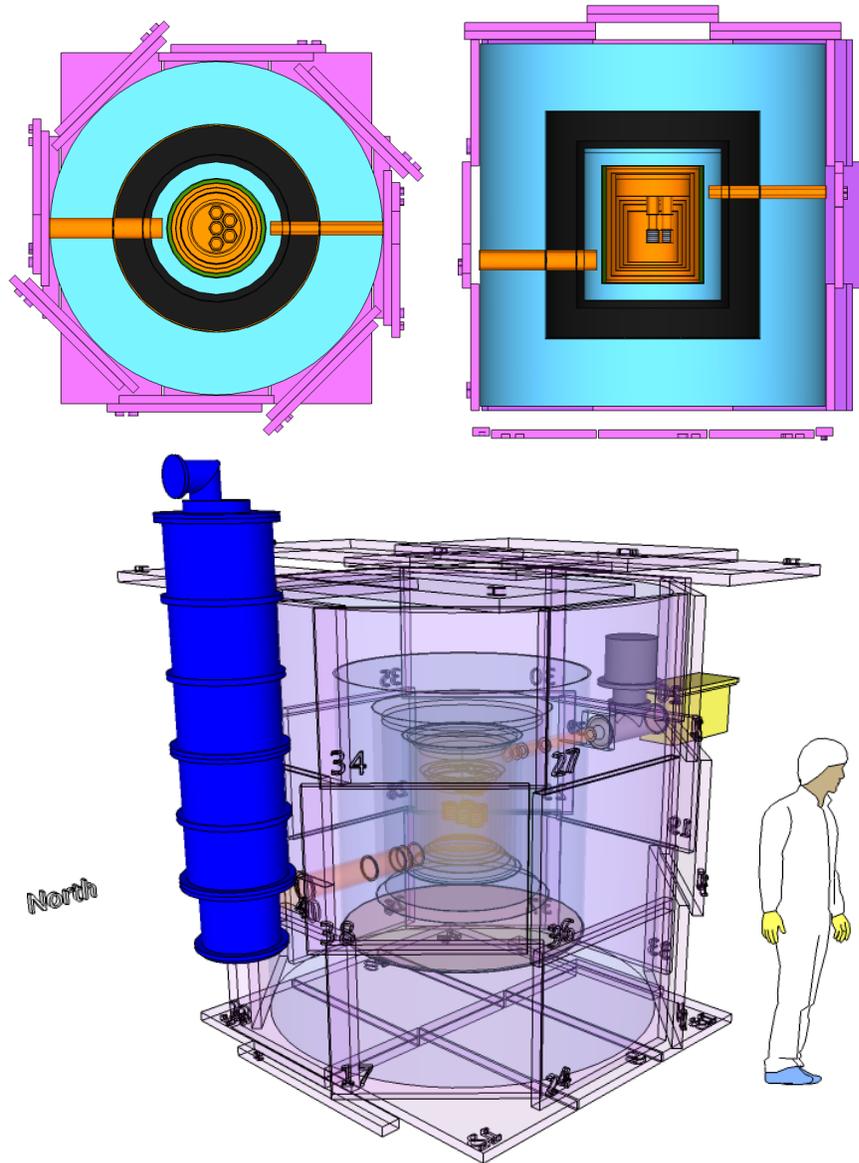


Figure 2.4: Schematic of shielding and muon veto cage of CDMS II. Cross sections from the top (upper left) and from the west (upper right) are shown depicting the muon detector panels (lavender), polyethylene shielding (blue), lead shielding (dark gray), and copper hardware (orange). The lower copper tubing is for cryogenics; the upper copper tubing is for electronics. Also shown is a partially transparent perspective view of the apparatus with a human figure for scale. In the foreground is the Dewar for the ^3He - ^4He dilution refrigerator (blue). In the background is the electronics box from which the ZIP signals are read (yellow).

Run Number	Dates	Location	ZIP Towers
20	Apr. 2000 – Dec. 2000	SUF	1
21	Aug. 2001 – July 2002	SUF	1
118	Oct. 2003 – Jan. 2004	Soudan	1
119	Mar. 2004 – Aug. 2004	Soudan	1,2
123	Aug. 2006 – Mar. 2007	Soudan	1-5
124	Apr. 2007 – July 2007	Soudan	1-5
125	July 2007 – Jan. 2008	Soudan	1-5
126	Jan. 2008 – April 2008	Soudan	1-5
127	May 2008 – Aug. 2008	Soudan	1-5
128	Aug. 2008 – Sep. 2008	Soudan	1-5

Table 2.1: Run history of CDMS II, including all runs in which WIMP-search data was taken. SUF indicates the Stanford Underground Facility. The table groups runs whose data were analyzed together.

Soudan.

This thesis describes the analysis of Runs 125–8, the second collection of data involving five ZIP towers.

2.4 Organizational Overview

The following rough outline describes the data collection and blind analysis procedure used in CDMS II for a given set of data. Refer to section 1.4 for definitions of terms used here.

1. Collect WIMP search data
2. Collect calibration data
3. For every background category
 - (a) develop a discrimination criterion
 - (b) estimate the leakage
4. Calculate the total efficiency
5. Calculate the total expected background (sum of the leakages)
6. “Unblind” the data
7. Compare the number of WIMP candidates to the expected background
8. Characterize the WIMP signal rate (by calculating an upper limit and/or a confidence interval)

This outline oversimplifies several aspects of the procedure. For example, apart from background discrimination there are numerous data quality and detector performance criteria that must be applied to the data. A discrimination criterion may dominate the background rate and the efficiency, so that steps 3, 4, and 5 are intertwined in a complicated way. This is the case with the *surface electron recoil* background, the subject of chapter 6. Another oversimplification is in step 8. Rather than simply characterizing the WIMP signal rate, the known and theorized properties of WIMPs discussed in sections 1.1 and 1.3 are utilized to characterize the WIMP/nucleon cross section. This allows for comparison of multiple direct detection experiments to one another and to theoretical predictions in plots similar to figure 1.1.

In the following four chapters I will discuss the four main categories of background. These are *neutrons from radioactivity*, *neutrons from cosmic rays*, *bulk electron recoils*,

and *surface electron recoils*. *Bulk electron recoils* and *surface electron recoils* correspond roughly but not precisely to gamma and beta radiation, respectively. In each chapter I will describe

1. the nature and origin(s) of the background
2. the means by which it can be identified
3. the discrimination criterion
4. the final background estimate.

In a simplified sense, the “IP” part of the ZIP acronym allows for discrimination of bulk electron recoils, via ionization yield. The “Z” part of the acronym is necessary for discrimination of surface electron recoils, which can defeat discrimination based on ionization yield. Information from a single ZIP cannot be used to discriminate neutrons; we must instead rely on shielding and coincidence.

In chapter 7 I will discuss the results of the unblinding of the data from Runs 125–8. Chapter 8 presents work I am doing toward the next phase of CDMS, known as SuperCDMS.

Chapter 3

Radiogenic Neutron Background

Neutrons produce nuclear recoil interactions in ZIPs and are sufficiently penetrating to interact in the bulk. Thus they can produce signals identical to those that a WIMP would produce. Neutrons differ from WIMPs in that they may collide with more than one nucleus in a detector, but the combined signal is still unlikely to be distinguishable from that of a single interaction.

The first line of defense against neutrons is the use of radiopure materials and shielding. This minimizes the rate of neutron interactions. The second strategy is identification and discrimination of neutrons via event coincidence, i.e. the simultaneous appearance of signals in multiple detectors (either ZIPs or detectors in the muon veto cage).

Much of the work of managing the neutron background is in the design and location of the experiment. The remaining tasks are to ensure that the methods of identifying neutrons via coincidence are robust and to understand the sources of neutrons well enough to produce estimates of the neutron backgrounds. For both radiogenic neutrons (this chapter) and cosmogenic neutrons (chapter 4) the estimate relies on Monte Carlo simulations supported by observational benchmarks, as will be described further.

3.1 Origins

Radiogenic neutrons are produced by decay processes of radioactive materials near the detectors. Uranium and thorium exist in the rock surrounding the experimental cavern

at the Soudan Underground Laboratory. Neutrons can be produced with kinetic energies up to 6 MeV by spontaneous fission or by secondary collisions with neutron-rich nuclei ((α, n) processes). However the polyethylene shielding surrounding the ZIPs serves as an efficient neutron moderator, reducing the energy of such neutrons to a level far below the experimental recoil energy threshold of 10 keV. There remains a small possibility that radiogenic neutrons originating external to the shielding could bypass the polyethylene through one of the two holes in the shield.

A larger potential source of radiogenic neutron background lies interior to the shielding, due to contamination by uranium and thorium in the experimental hardware. With no moderating polyethylene in their paths, these neutrons can produce nuclear recoils in the ZIPs within the WIMP-search energy range.

3.2 Identification and discrimination

Unlike the other background categories described in this thesis, no specific identification strategy exists for radiogenic neutrons. If a radiogenic neutron interacts in a single ZIP with the right recoil energy, it will be indistinguishable from a WIMP. ZIP-coincidence can be used to reject some radiogenic neutrons but not all of them. It is thus of great importance that the rate of radiogenic neutron events be small.

3.3 Final background estimate

To calculate the radiogenic neutron background, it is necessary to estimate the level of contamination of the inner hardware components. Most of the events recorded in CDMS II are produced by photons, as described in chapter 5. Data from Runs 123–4 were used to produce a detailed spectrum of electron recoils from these photons. Peaks in this spectrum correspond to gamma ray emissions from radioactive decay connected with the ^{238}U and ^{232}Th decay chains and from ^{60}Co and ^{40}K . Monte Carlo simulations were performed to produce separate spectra corresponding to contamination by these isotopes in the tower hardware and the surrounding copper cans. A χ^2 minimization fit of these simulated spectra to the observed spectrum has produced estimates of the contamination level of the various components that are in good agreement with independent estimates

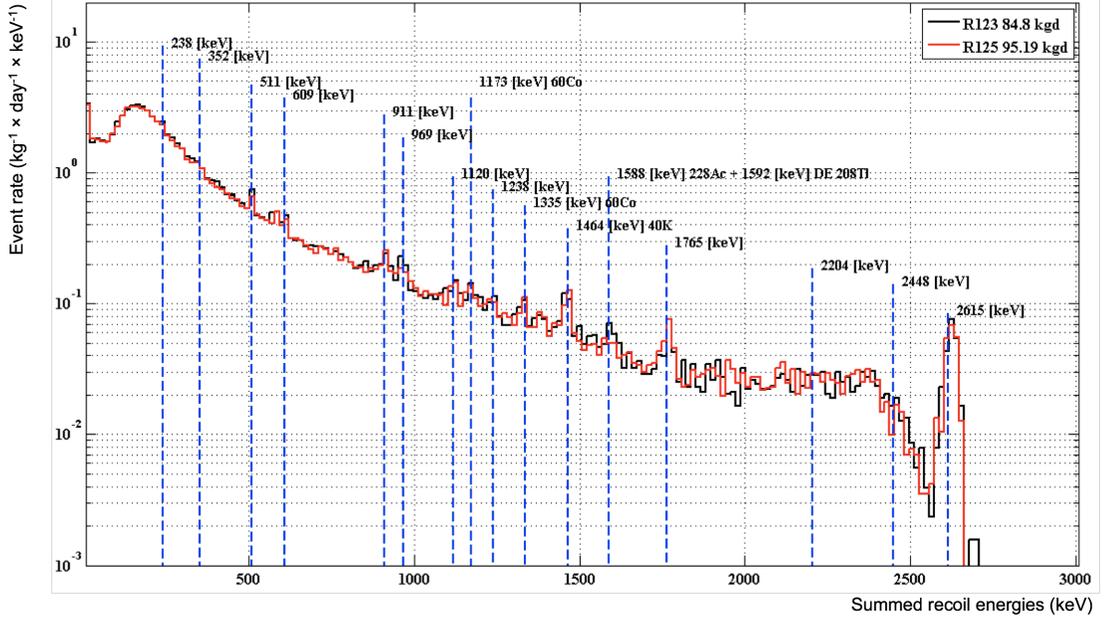


Figure 3.1: Photon energy spectra from Runs 123 (black) and 125 (red). Because high energy photons typically deposited energy in multiple ZIPs, the recoil energies of all germanium detectors with a signal in an event are summed. All lines seen in the spectrum are identifiable from radiation sources discussed in the text.

produced by high-purity germanium detector screening of the materials used [22].

Once the contamination levels have been estimated, a second simulation is used to calculate the rate at which radiogenic neutrons produce WIMP-like signals in WIMP-search detectors. This was done for the analysis of Runs 123–4. It is not likely that this rate would change significantly for Runs 125–8, which uses the same hardware. This is confirmed in figure 3.1, which compares photon rates from Runs 123 and 125.

The background estimate for Runs 125–8 is the product of the estimated background rate and the WIMP-search exposure. The exposure is calculated using all WIMP-search selection efficiencies, since any neutron rejected by WIMP-selection criteria would not constitute background. Uncertainties in the background calculation include the combined uncertainties of several contamination estimates. The resulting background estimate for radiogenic neutrons produced by contamination in the experimental hardware is $0.041^{+0.003}_{-0.017}$ for Runs 125–8.

An estimate of radiogenic neutron background from sources external to the shielding was calculated for the analysis of Runs 125–8. As mentioned previously neutrons from the cavern rock can bypass the polyethylene through either of the two holes in the shielding. Monte Carlo simulations were performed to estimate the rate at which neutrons produce WIMP-like signals for a neutron source located at a hole entrance. Data taken using a ^{252}Cf neutron emitter inserted into one of these holes at different depths were compared to the simulations to assess the degree of systematic error resulting from idealizations in the simulated hardware. Combined with calculations of the neutron flux from the Soudan cavern rock this yields a background estimate of 0.002 ± 0.002 . Thus the total radiogenic neutron background estimate for Runs 125–8 is $0.04^{+0.00}_{-0.02}$.

Chapter 4

Cosmogenic Neutron Background

The background from neutrons produced by cosmic rays (cosmogenic neutrons) differs from the radiogenic neutron background in two ways. First, the rate at which cosmogenic neutrons interact in ZIPs is much higher. Second, cosmogenic neutrons can be identified by the simultaneous detection of the cosmic ray muon or other particles in the cosmic ray shower.

4.1 Origins

Neutrons that may interact in the detectors can be generated in showers of particles produced by cosmic ray muons interacting with matter surrounding or within the experimental cavern. Muons reaching the depth of the Soudan laboratory have a median kinetic energy of approximately 10 GeV and a mean kinetic energy of approximately 200 GeV. They can produce neutrons within the shielding or hardware. Neutrons can also be produced external to the shielding with sufficient energy (at least 50 MeV [23]) to reach a ZIP and produce an interaction above the experimental threshold.

4.2 Identification

4.2.1 The muon veto cage

The depth of the Soudan Underground Laboratory reduces the cosmic ray muon flux by a factor of 50,000 relative to the surface. Nevertheless cosmogenic particles interact in

the ZIPs at a rate of about two per hour. Observations of data coincident with cosmic rays indicate that among these cosmogenic particles about six neutrons per kg-year produce WIMP-like signals in germanium ZIPs. Approximately 80% of these can be discriminated from WIMPs due to coincident signals in other ZIPs. If this were the only discrimination the background from cosmogenic neutrons would be about three times greater than the largest background category (surface electron recoils).

For this reason a cosmic ray detection system is built into the CDMS II design. The muon veto cage is an assembly of 40 particle detector panels surrounding the outermost layer of polyethylene shielding. The cage is 2.6 m high and between 2.3 m and 2.7 m wide at various levels. The six panels forming the top of the cage overlap to form a canopy with no gaps to particles coming from above. The side and bottom panels also overlap but some gaps exist. Most muons that reach the cavern have approximately vertical trajectories, with a median angle with respect to vertical of approximately 32° [23]. The shallowest angle forming a line-of-sight through a gap to the ZIPs is 67° , through the hole through which the detector wires are run. Only about 0.4% of cosmic ray muons have trajectories with angles greater than 67° . Details of the arrangement of panels are shown in figure 4.1 (and in the earlier figure 2.4).

The panels are each 5.1 cm thick and composed of Bicron BC-408 plastic scintillator. Attached to each panel is an acrylic light guide which directs the scintillation photons to a Hamamatsu R329-02 photo-multiplier tube (PMT). A photo is shown in figure 4.2. The panels are very sensitive to cosmic ray muons, which, as minimum ionizing particles, deposit approximately 10 MeV in a panel if they travel through the full thickness (more if they travel at an oblique angle). Other high energy charged particles in a muon-induced shower can deposit similar energies in the panels. Photons from environmental radioactivity also produce signals in the panels at a much higher rate, but with deposited energies less than 3 MeV.

The output of each PMT is digitized to 1 μs precision [24]. This information is recorded in two different ways. First, the time of any signal above a hardware threshold set for each panel is recorded in a history buffer. These thresholds are set just high enough to exclude most of the ambient radiogenic photon flux to keep the rate relatively low (around 10 Hz for each panel). The thresholds correspond to deposited energies of approximately 2 MeV. Second, for any event recorded in a WIMP-search data run, a full

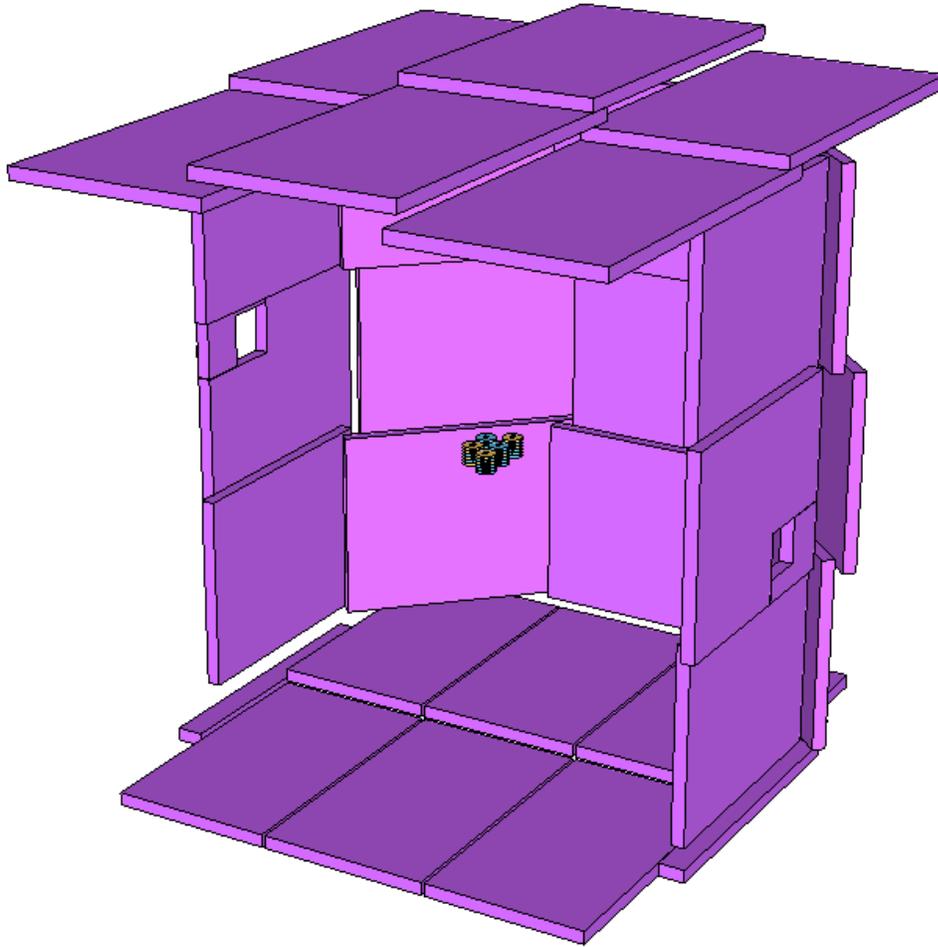


Figure 4.1: Schematic of panel placement of muon veto cage, with nine panels hidden to show interior detail including the position of the ZIPs. The image was constructed from a geometry file used for Monte Carlo simulations, which simplifies the two holes, one for electronics (upper) and one for cryogenics (lower). The actual holes are smaller and rounded to better fit the copper tubes.



Figure 4.2: Photo of a portion of the muon veto cage. Attached to the rectangular panels are triangular light guides. Also visible are the housings of the PMTs.

trace for each panel is recorded over a time window from -185 to $20 \mu\text{s}$ with respect to the event trigger. The trace is first integrated and conditioned so that pulse heights are representative of the deposited energy. Examples of such traces are shown in figure 4.3.

To collect cosmic ray data, CDMS II implements a veto trigger condition: if two panels simultaneously exceed their hardware thresholds, information from all detectors (veto cage panels and ZIPs) is recorded. This allows for the study of the performance of the veto cage and cosmic ray muon showers, whether or not interactions occur in the ZIPs. For Runs 125–8 these triggers occurred at an overall rate of 0.035 Hz (or 1 trigger every 29 seconds on average).

The simplest way to characterize a panel trace is by its maximum amplitude. Figure 4.4 shows the distribution of the highest peak in the traces for two of the muon veto panels: a top panel (Panel #4) and a side panel (Panel #16). The raw trace output is in volts; an approximate energy scale is also shown. At low energies the distribution decreases sharply with energy. This part of the distribution corresponds to the ambient radiogenic photon flux. No significant features are apparent in this region apart from selection artifacts of the muon veto trigger. Above a few MeV the distribution flattens, then increases sharply corresponding to the energies deposited by cosmic ray muons traversing the full thickness of the panel. For a top panel this “muon bulge” peaks at about 10 MeV as expected for muons traveling through the 5.1 cm thickness. Beyond

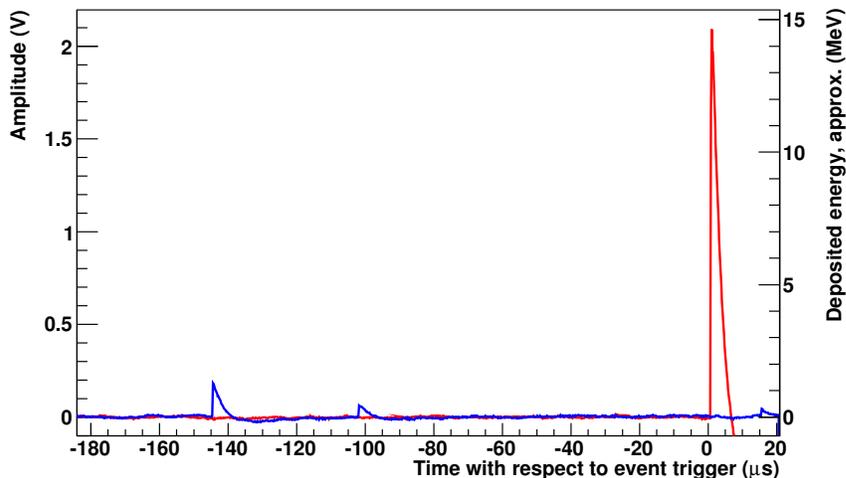


Figure 4.3: Traces recorded for two muon veto panels situated at the top of the cage. One trace (red) is characteristic of a cosmic ray muon. The other (blue) shows typical pulses from radiogenic photons. The energy can be estimated from the peak height with the approximate scale on the right of the figure.

the peak the distribution falls relatively slowly, corresponding to muons traveling at larger angles relative to vertical, or in some cases to multiple minimum-ionizing shower particles traveling through the panel simultaneously. For a side panel the muon bulge is stretched to higher energies, since often muons will traverse much of the vertical extent of the panel (76 cm). The plateau between the radiogenic photon region and the muon bulge is associated with muons that hit the edge of a panel in such a way that they do not pass through its entire thickness. For energies above the radiogenic photon region the distributions are matched well by those produced by Monte Carlo simulations.

After the completion of Runs 125–8 it was discovered that two failures had occurred in the discriminators that record information to the history buffer when a panel exceeds its hardware threshold. For one top panel (Panel #5) the discriminator was not operating at any time during Runs 125–8. For one side panel (Panel #22) the discriminator failed during Run 126 and was not operating during Runs 127 and 128. The trace

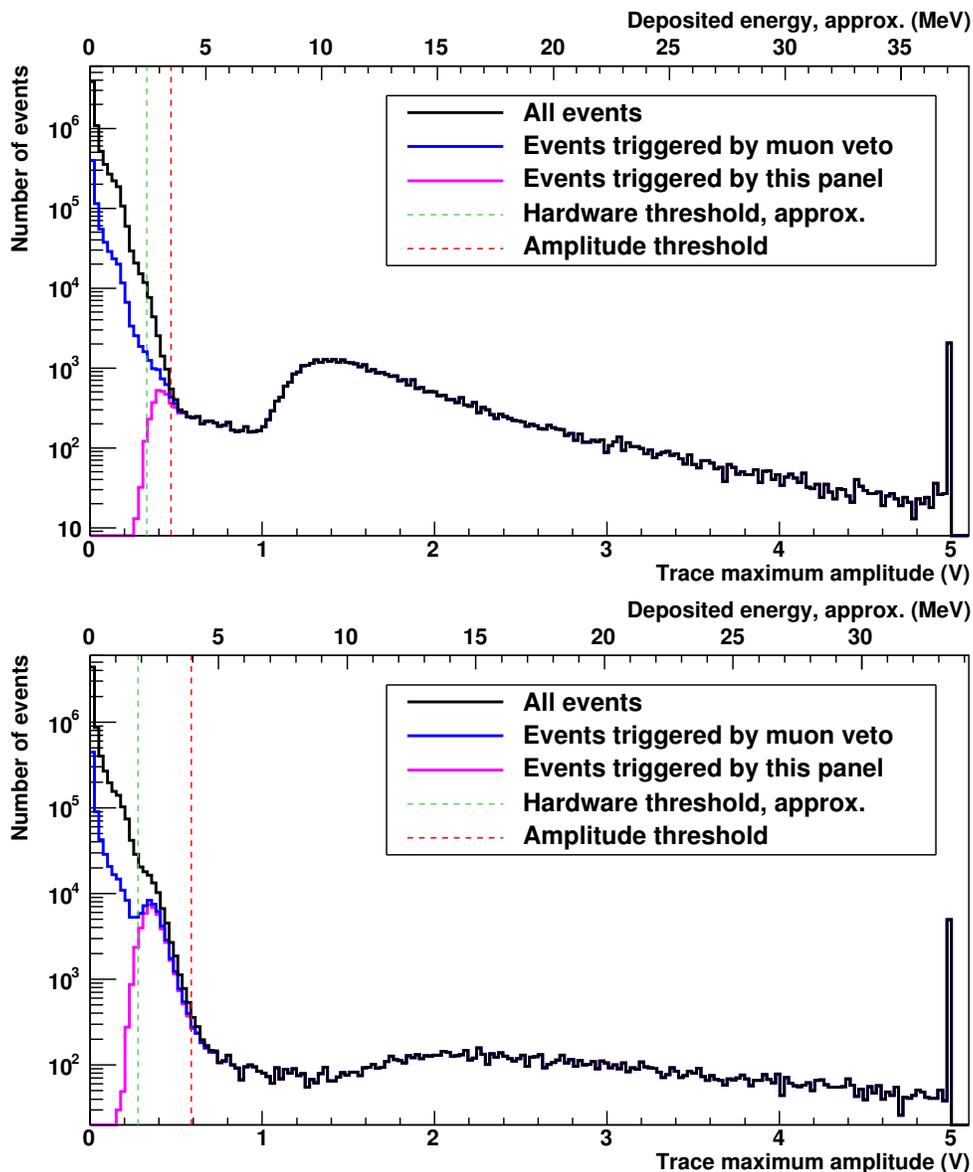


Figure 4.4: Distribution of the maximum value (highest peak) in a veto panel’s trace for a top panel (top figure) and a side panel (bottom figure). Features include the ambient radiogenic photon flux (roughly 0 to 3 MeV), the “muon bulge” (beginning at roughly 7 MeV), and a saturation effect at 5 V. The energy scale is approximated by matching the observed muon bulge to those produced by Monte Carlo simulations.

information of all panels was recorded properly for the entirety of Runs 125–8¹. Tracking the position of the “muon peak” in trace-maximum distributions over the course of Runs 125–8 revealed that the gains of the PMTs decrease over periods of several months in which they are continuously at high voltage, on the order of 1–2%. This small change in gain corresponds to more noticeable changes in the rate at which the discriminator records pulses to the history buffer, since each hardware threshold is set at a steeply-changing point in the radiogenic photon energy spectrum. For Runs 125–8 the average total rate of above-threshold signals in the veto cage panels was 407 Hz. If all of the discriminators had been operating, this rate would likely have been about 10 Hz greater.

4.3 Discrimination criterion

The information provided by the muon veto cage allows us to set criteria to reject interactions of cosmogenic neutrons that may otherwise be mistaken for WIMP signals. The simplest criterion would be to reject any signal in a WIMP-search detector if it were coincident with any activity in the veto cage. However it is unnecessary to reject signals coincident with radiogenic photons interacting in the panels, which are not associated with ZIP backgrounds. Indeed it would be undesirable to do so because it would greatly reduce the live time of the experiment. The hardware threshold set for each panel is not high enough to exclude all radiogenic photons.

The strategy for discriminating cosmogenic neutrons is to identify events which show some hint of a cosmic ray muon or muon-induced particle shower. Two separate discrimination criteria are used: one using information from the history buffer and one using information from the panel traces. I will refer to the former as the **hardware-threshold discrimination** and to the latter as the **trace-maximum discrimination**.

In figure 4.4 a vertical green line represents the approximate value of the hardware threshold. The times of signals with amplitudes above this threshold are recorded to the history buffer. A second, higher threshold (shown as a vertical red line) is set empirically for each panel using the distributions themselves, with the intent of

¹If there were a failure in the recording of trace information for any of the panels, it would have been noticed early by virtue of the automatic real-time data monitoring system.

rejecting all radiogenic photons while accepting all muon-like signals. This threshold is set conservatively to accept amplitudes in the “plateau” region of the distributions, identified as muons clipping a corner of a panel. If such a muon actually enters the cage, it is almost certain to produce a larger signal in another panel. Amplitude thresholds set for Runs 123–4 were found to be appropriate for distributions in Runs 125–8, suggesting that veto panel gains have not changed significantly. The average value of the amplitude thresholds is 0.58 V (≈ 3.8 MeV).

The trace-maximum discrimination rejects any WIMP-like signal if any muon veto panel’s trace maximum exceeds its amplitude threshold. It is of course possible for this maximum to occur at any point in the trace; hence the time window in which a muon-like signal can serve as a veto is -185 to 20 μs with respect to the time of the ZIP signal.

The intent of this criterion is to identify any muon that enters the veto cage. A cosmic ray muon that enters the veto cage will almost certainly exit it as well (since it necessarily had enough energy to travel through 713 m of rock). Thus almost all through-going muons will activate more than one panel. The efficiency with which a through-going muon is tagged by the veto cage has been measured using data from Runs 125–8. First muons passing through the overlap between the central upper veto panels were selected as a sample of muons known to have entered the cage. Excluded from this selection were any events in which any of the other top panels showed significant activity, thus excluding cosmogenic showers. The fraction of such events in which a muon-like signal was also seen in any of the side or bottom panels was $91.7 \pm 0.5\%$. This is taken to be the efficiency with which the bottom and side panels tag through-going muons (and is in good agreement with what one would estimate from the gaps between the bottom panels, which cover about 8% of the area). Similarly a selection of muon signals in either of the two central bottom panels with no simultaneous signal in the other bottom panels is used to estimate the efficiency of the top and side panels. The fraction of this selection of events in which a simultaneous muon signal is seen in any of the top or side panels is $99.07 \pm 0.05\%$. The overall tagging efficiency is estimated by calculating the probability of a tagging failure in *both* sets of panels (bottom+side and top+side), or $\epsilon_{\text{overall}} = 1 - (1 - \epsilon_{\text{b+s}}) \times (1 - \epsilon_{\text{t+s}})$. The resulting tagging efficiency of the veto cage for through-going muons is $99.923 \pm 0.006\%$.

The trace-maximum discrimination involves a generous definition of “simultaneous” (the entire time span of the trace) with a relatively strict definition of “muon-like energy” (exceeding the amplitude threshold). A second discrimination is defined using the history buffer information, which contains a record of signals exceeding the hardware threshold. If the history buffer contains a record of a signal in a panel within a $50 \mu\text{s}$ window before a signal in a ZIP, that signal is rejected as a WIMP candidate. This hardware-threshold discrimination has a stricter time window but a more generous energy selection.

The complete cosmogenic neutron discrimination criterion rejects any WIMP candidate if it fails either the trace-maximum discrimination or the hardware-threshold discrimination. To calculate the efficiency of the discrimination criterion, we make use of another category of recorded events in WIMP-search data. Interspersed among triggered events are events recorded at random times (so-called “random triggers”). The efficiency of the cosmogenic neutron discrimination criterion is taken to be the fraction of random-triggered events that are not rejected by the discrimination criterion. For Runs 125–8 this was found to be $97.88 \pm 0.01\%$. Most of the WIMP-search inefficiency introduced by the cosmogenic neutron discrimination criterion is due to the hardware-threshold discrimination, because the lower thresholds cause some events to be rejected accidentally due to radiogenic photon interactions in muon veto panels.

The $50 \mu\text{s}$ time window might seem much greater than necessary for identifying interactions in a ZIP coincident with a cosmic ray. A neutron with the minimum energy necessary to produce an above-threshold recoil moves at half the speed of light, and thus should arrive at a ZIP in less than a microsecond after it is produced. This is indeed what is observed in all the events at the Soudan Underground Laboratory in which an interaction in a ZIP is identified as a cosmogenic neutron. Only if a neutron were to follow a very complex path through the shielding could the delay between signals be significantly longer. In early runs at the shallow Stanford Underground Facility, where the shield and veto cage were very similar and the cosmic ray rate was much greater than at Soudan, some evidence of cosmogenic neutrons delayed as much as $40 \mu\text{s}$ after a veto panel signal was seen [25]. This motivated the $50 \mu\text{s}$ limit used in the hardware-threshold discrimination. It is entirely possible that the time limits in both parts of the discrimination criterion are more generous than necessary, though the WIMP-search

efficiency is great enough that there is little motivation to tighten them.

4.4 Final Background Estimate

4.4.1 Monte Carlo simulations and observational benchmarks

Monte Carlo simulations are employed to characterize the distribution and flux of cosmic ray muons that reach the experimental cavern. For the analysis of Runs 125–8 a new simulation was performed as a basis for estimating the cosmogenic neutron background. A simulated population of cosmic ray muons with energy and angular distributions appropriate for the Soudan Underground Laboratory was generated using information from a previous experiment at the same location (Soudan 2 [26]). A software toolkit designed to simulate the passage of particles through matter known as Geant4 [27] was used to simulate the passage of these muons through several meters of cavern rock, collecting data equivalent to 8.5 years of observation. Energy, momentum, and type information for the particles produced by these simulated interactions was stored to a file. Using this data, the interactions of these particles (and any secondary particles produced) with the CDMS II apparatus were simulated. This was repeated several times, producing data equivalent to 29 years of observation. The most important simulation data are the interactions in each of the detectors, the 40 muon veto panels and the 30 ZIPs. Information from previous simulations performed for CDMS II, including some using a different particle-physics software package (FLUKA [28]), were also used for the cosmogenic neutron background estimate.

Disagreement exists in the fraction of cosmic rays which produce neutrons as determined by simulations using the various software packages. The simulated fractions also disagree with the fraction observed in events identified as cosmogenic recorded during WIMP-search data-taking simulations. It is therefore key to study cosmic ray events to provide benchmarks to assist in the interpretation of Monte Carlo simulation results. The following is a list of some observations from Runs 125–8:

Cosmic ray rate in the muon veto cage

The simulations produce a set of events with no information about the event rate. Rather the rate at which cosmic rays interact with matter in the CDMS II apparatus is

determined from observed data. It is impractical to attempt to count every muon that hits a panel, since event recording is only triggered when two or more panels are hit simultaneously. Thus information regarding muons or showers that interact in only one panel is lost². Thus an appropriate benchmark is the rate of events in which at least two panels register a signal with a muon-like amplitude. A convenient energy threshold choice is 1 V (≈ 6.9 MeV) for every panel. This is higher than the thresholds used in the trace-maximum discrimination, thus rejecting most edge-hitting muons. This selection allows for normalizing the rate from a Monte Carlo simulation.

The failure of the Panel #5 discriminator must be taken into account to correctly estimate the cosmic ray rate in Runs 125–8, since muons that hit this panel and only one other would fail to trigger event recording. By considering the panels that are situated symmetrically to Panel #5 (the top corner panels) the number of events that failed to be recorded can be estimated by counting how many events register in one of those panels and only one other panel, and do not trigger a ZIP. In this way the cosmic ray rate in the veto cage for Runs 125–8 is calculated to be 15.54 ± 0.03 mHz, or one cosmic ray event every 64.4 ± 0.1 s.

Fraction of cosmic ray events with activity in ZIPs

As would be expected considering the geometry of the experimental apparatus, only a small fraction of cosmic ray events triggered by the veto cage show corresponding activity in ZIPs. This fraction can be calculated either by setting low energy thresholds on each ZIP, just high enough to reject noise, or by setting an energy threshold of 10 keV on each ZIP, which is the threshold used for WIMP search. In the former case, thresholds are set on each ZIP in accordance with its average noise level, with an average threshold of 1.8 keV³. This yields a fraction of 0.0219 ± 0.0003 . When the thresholds are set to a recoil energy of 10 keV for each ZIP, the corresponding fraction is 0.0194 ± 0.0002 . In both cases the definition used for a cosmic ray event is that at least one veto panel has a trace maximum greater than 1 V.

²A record exists in the history buffer but the thresholds are so low as to pollute this data with radiogenic photon interactions.

³With such low thresholds it is important to remove from consideration periods of data-taking in which the signal rate in any ZIP is unusually high due to some pathological behavior. For example in Runs 125–8 there were periods in which a slow vacuum leak led to the development of films of helium on some detectors, producing spurious low-amplitude signals. These are described in [22].

Nuclear recoil interactions in ZIPs coincident with cosmic rays

The criteria developed for identifying WIMP-signals (i.e. bulk nuclear recoils) are ideal for performing a search for cosmogenic neutron signals in ZIPs. Such a search was performed of the data from Runs 125–8, applying all WIMP-search criteria except muon veto coincidence, ZIP coincidence, and surface interaction discrimination. A further selection criterion was that the event be tagged as a cosmic ray event by failing the trace-maximum discrimination (thus guaranteeing that at least one panel showed a muon-like energy deposition).

The ZIPs included were the 14 germanium detectors designated for WIMP search and all 11 silicon detectors. The energy and ionization yield characterization of the silicon detectors is significantly less precise because they were not subjected to a procedure used to remove variations with interaction position. This procedure will be described in section 6.2.7. Consequently the selection of nuclear recoils via ionization yield is somewhat less reliable for the silicon ZIPs.

Ionization-yield selection produced a set of *neutron candidates*. Photons associated with cosmic ray showers could potentially induce surface electron recoils, which can have suppressed ionization yields and thus be mistaken for nuclear recoils. The discrimination of surface interactions is based on phonon timing. An analysis of the phonon timing characteristics of the neutron candidates was performed to check whether any might be surface electron recoils. The distribution is consistent with true bulk nuclear recoils, with only one of 14 neutron candidates having phonon timing that is also clearly consistent with a surface interaction. Because the distributions of phonon timing vary significantly with detector, it is easiest to illustrate this by calculating the percentile value of each candidate’s phonon timing in terms of the distributions in the corresponding detector; this is shown in figure 4.5. The resulting picture is in accord with the notion that most if not all of the interactions in this sample are true bulk nuclear recoils. This analysis was performed only for neutron candidates found in germanium detectors because no surface interaction discrimination was developed for silicon detectors.

Table 4.1 specifies the number of neutron interactions identified in each ZIP type and table 4.2 gives some details about the configurations of the events in which they appear. Three interactions satisfy all WIMP signal criteria except muon-veto coincidence and surface interaction discrimination (two of them fail the surface interaction

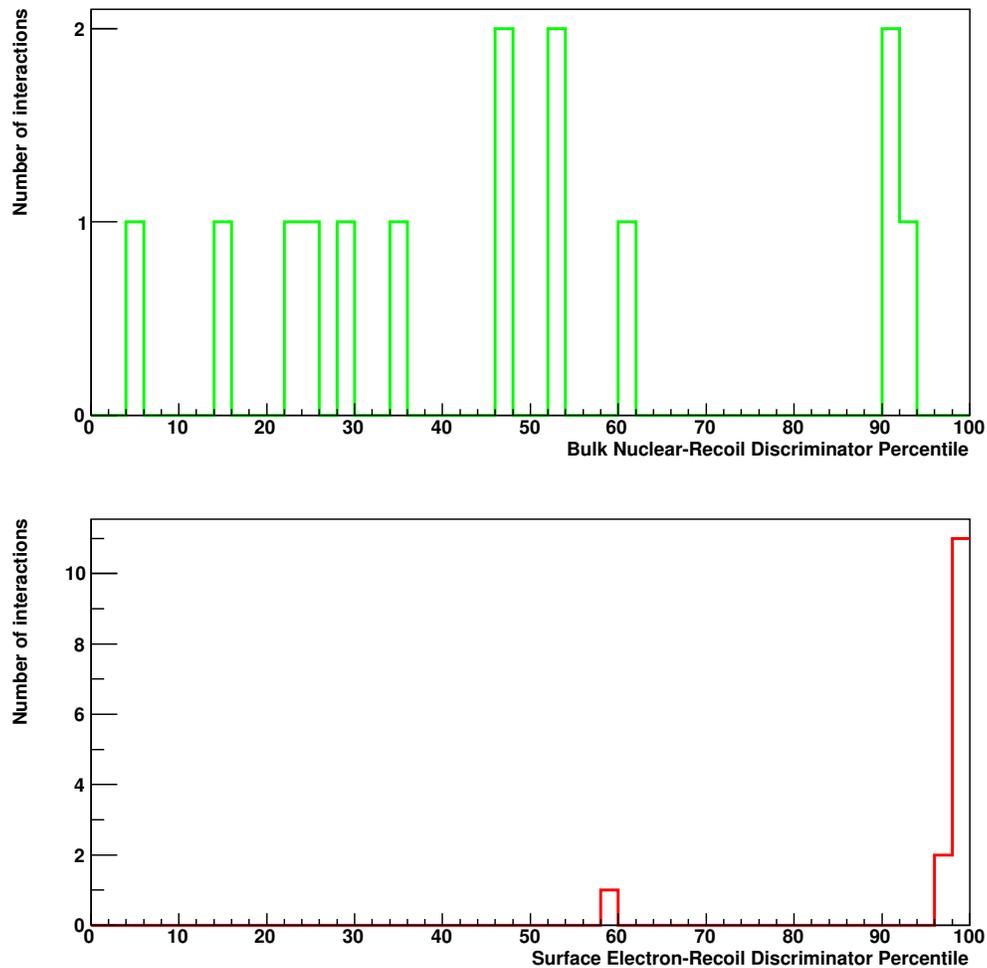


Figure 4.5: Distribution of phonon timing of the cosmogenic neutron candidates observed in germanium ZIP data in terms of their percentiles in the corresponding distributions for the bulk nuclear recoil sample (top) and the surface electron recoil sample (bottom). The phonon timing is characterized by a quantity referred to as the *discriminator*, defined in chapter 6.

	Ge ZIPs	Si ZIPs	All ZIPs
Neutron-candidate interactions	14	11	25

Table 4.1

	Total	No ZIP coincidence
Events with a neutron candidate in a Ge ZIP	14	3
Events with a neutron candidate in a Si ZIP	11	2
All events with a neutron candidate	22	5

Table 4.2

discrimination, consistent with its measured efficiency). Figure 4.6 displays the recoil energy and ionization yield distributions of the neutron candidate interactions.

Considering the exposure and selection efficiency corresponding to this search, the observed rate of cosmogenic neutron interactions in germanium detectors is 13_{-4}^{+5} per kg·year and the rate of neutron interactions with no ZIP-coincidence is 3_{-2}^{+3} per kg·year. The corresponding rates produced by the Geant4 simulation are 1.7 ± 0.1 per kg·year and 0.26 ± 0.04 per kg·year, respectively. Apparently the simulation severely underestimates the production of energetic neutrons by cosmic ray muons. However the fraction of interactions which are not coincident with other ZIP signals is the same, within statistical error, whether calculated from the observed counts or the simulated counts.

Indirect evidence of cosmogenic-neutron capture

Figure 4.7 is an example of all the traces from all 70 detectors for an event visualized in a single display. The event depicted is a cosmic ray event from Run 127 which includes an interaction identified as a neutron. In all the trace windows a vertical dashed cyan line represents the trigger time of the event. An unusual feature of this event is the presence of signals in some ZIPs that are delayed with respect to the trigger time by more than $100 \mu\text{s}$. Particles with sufficient energy to produce signals in ZIPs are not expected to be delayed more than a microsecond from other particles associated with the same production process. It is of course possible for an unrelated particle to accidentally appear within the same time window, but due to the low background rate this occurs in only about one in 10,000 random-triggered events. Since 5 of the 22 cosmogenic neutron events have delayed signals this is a significant effect.

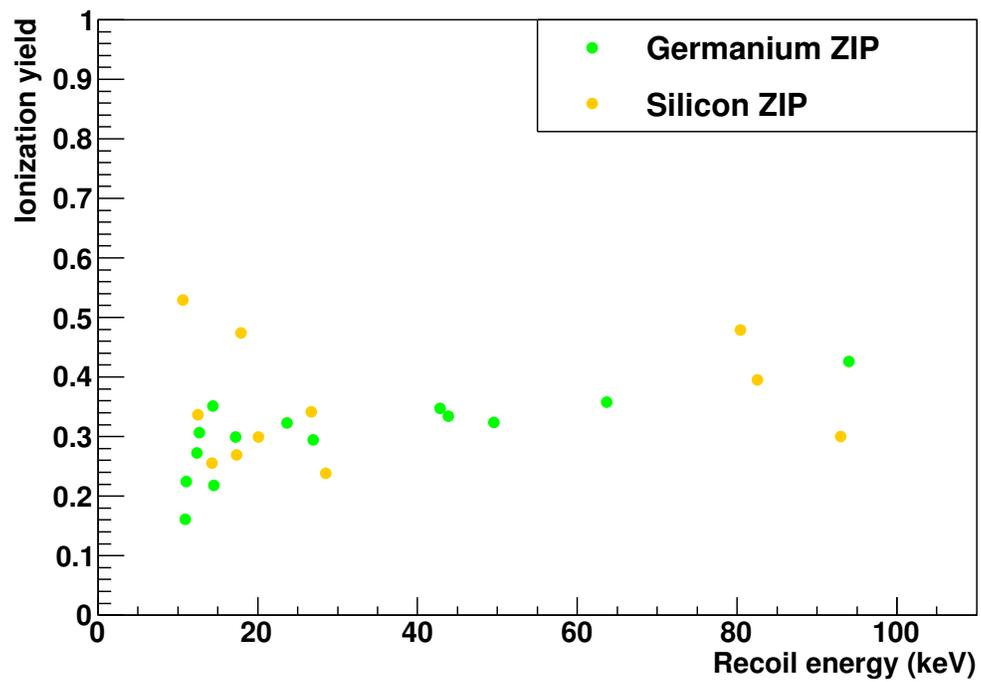
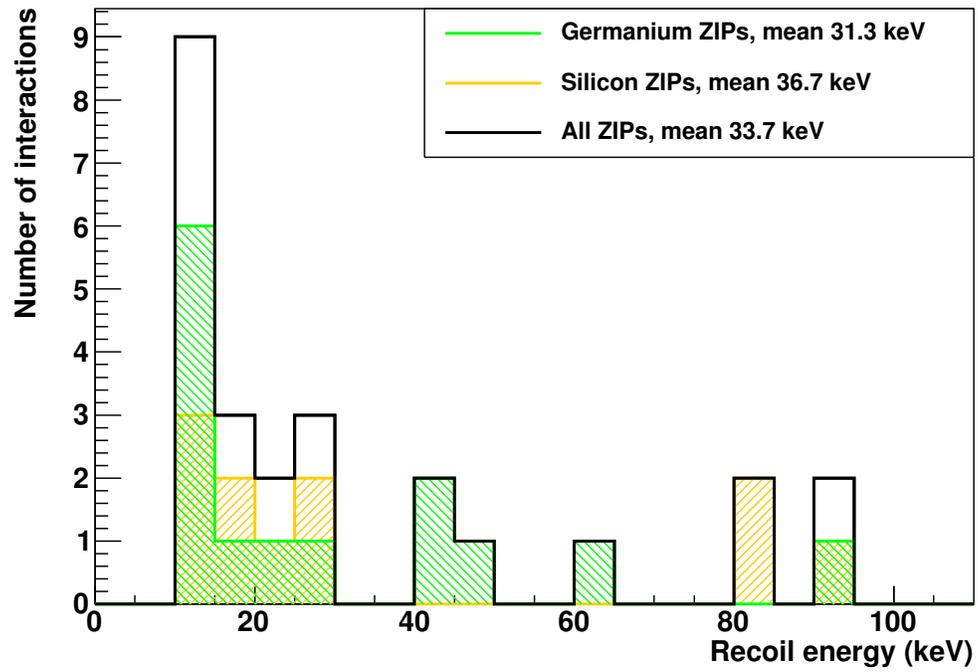


Figure 4.6: Recoil energy and ionization yield distributions of identified cosmogenic neutron interactions.

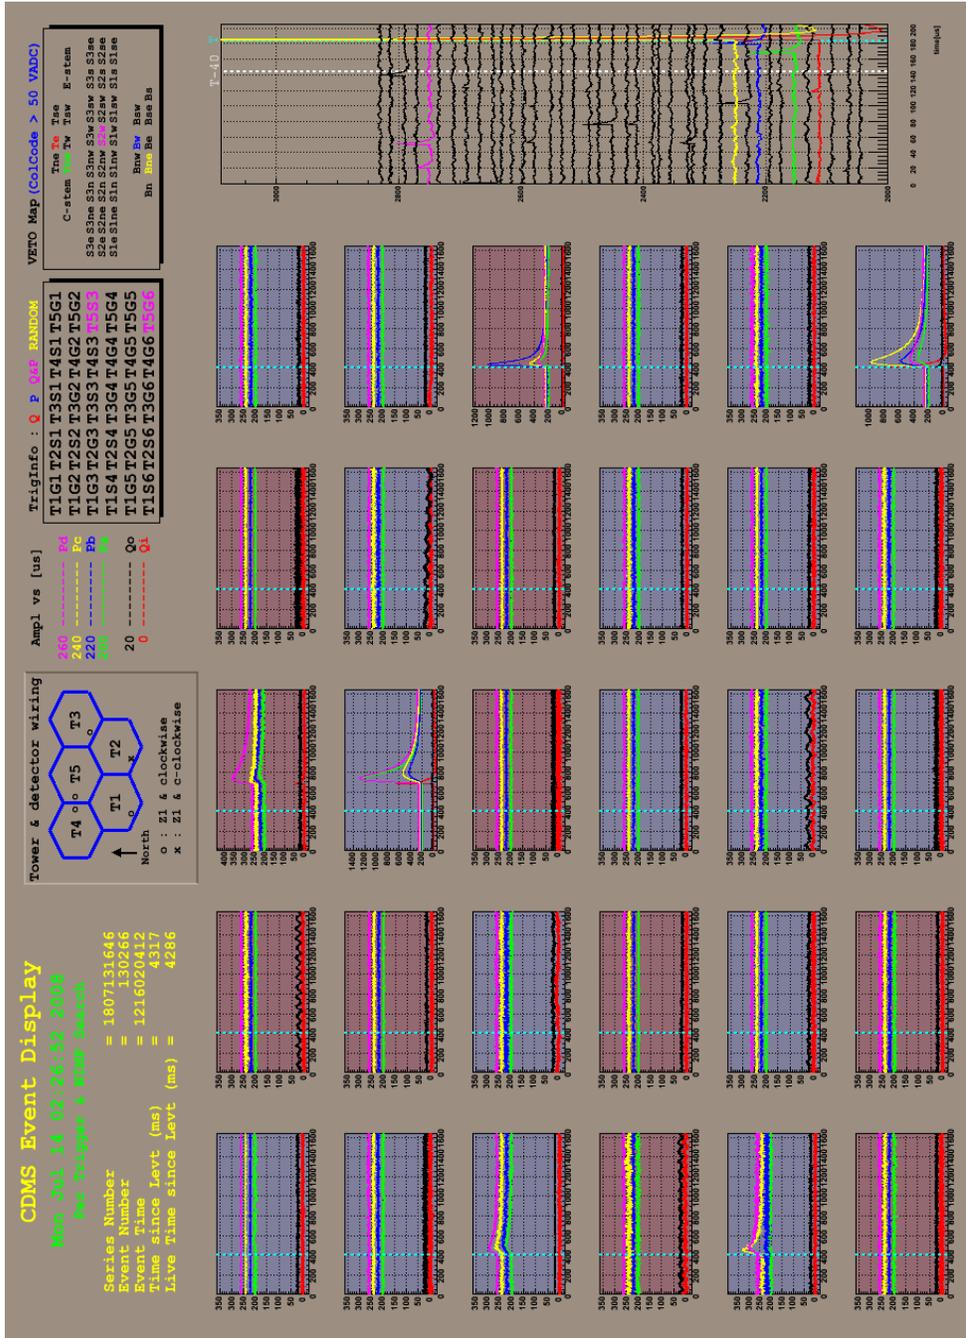


Figure 4.7: Display of detector traces from a cosmic ray event. The 30 ZIPs are represented in five columns, one for each tower. The 40 muon veto traces are shown together on the far right. The event trigger time is represented in all trace plots by a vertical dashed cyan line (mostly obscured in the muon veto plot by a strong veto pulse shown in yellow.) T5Z3 (5th column, 3rd row) shows an interaction identified as a neutron. This interaction and three others are simultaneous with the event trigger. T3Z1 and T3Z2 show signals occurring 272 μ s after the event trigger.

A search for such delayed interactions revealed that the rate is about seven times greater than chance for cosmic ray events. A likely interpretation is that these are photons released after neutron capture. The polyethylene shielding can reduce the energy of a neutron until it is thermalized. A thermal neutron is likely to eventually be captured by hydrogen, with a capture cross-section of 0.33 barns. Considering the speed of a thermal neutron this is consistent with a mean lifetime before capture of about 200 μs . Upon capture by a proton, an excited deuteron is produced which promptly releases a 2.23 MeV photon. If this photon interacts with a ZIP detector there will be a characteristic delay with respect to the cosmic ray processes that produced the parent neutron.

The data processing used in CDMS II was not designed to properly characterize the energies of pulses delayed by more than a few tens of microseconds with respect to the event trigger. Raw quantities such as pulse integrals or pulse maxima can be calibrated to produce rough energy reconstructions. This was done for the delayed interaction analysis. Energies from the ionization channel were constructed from the pulse maxima, with resolutions typically about 2% of the energy for energies between 0.5 and 3.0 MeV. Under the assumption that an interaction is an electron recoil the calibrated ionization energy can be taken to be the recoil energy, properly expressed in MeVee (electron-equivalent MeV).

Calibrating phonon signals at high energies is more difficult due to saturation effects. Calibrations were generated using integrated phonon pulses, with typical resolutions of 5–10% of the energy at a few MeV. The phonon energy was combined with the ionization energy to calculate the ionization yield, with a much poorer resolution than usual. In the selection of delayed interactions, high-radius interactions were not excluded, which will further distort ionization yield reconstruction. Figure 4.8 shows the distribution of ionization yields for delayed interactions seen in cosmic ray events. The distribution is roughly consistent with electron recoils, as would be expected from photons.

Photons with MeV energies will usually interact in more than one ZIP. An example of two delayed interactions in neighboring detectors is seen in figure 4.7. These two interactions occurred simultaneously. In some events multiple delayed signals are seen that are not simultaneous, suggesting that more than one delayed photon interacted in the detectors. I define a delayed *subevent* as a collection of simultaneous interactions

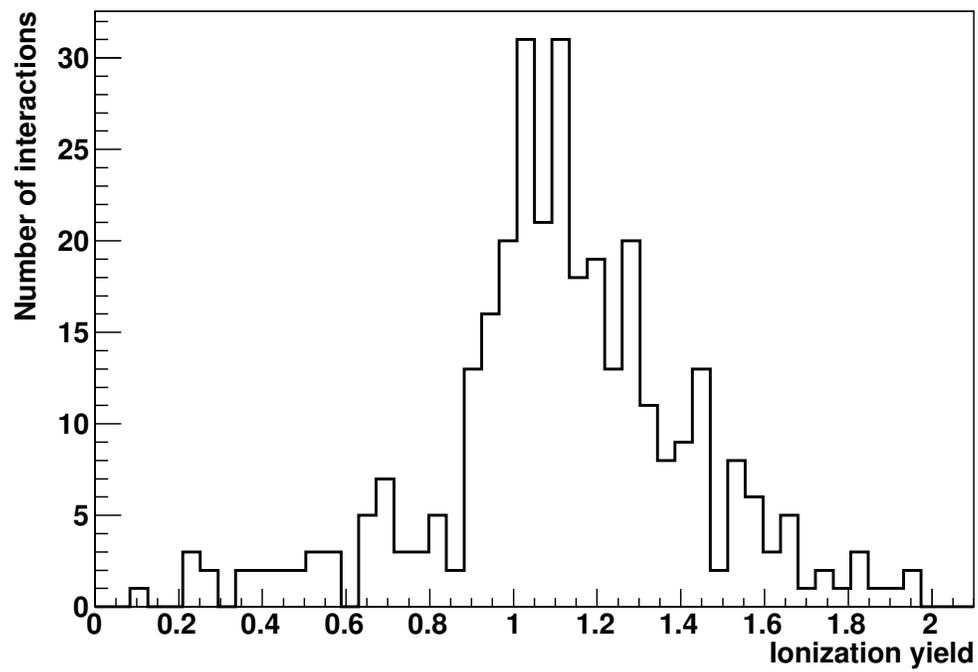


Figure 4.8: Ionization yield distribution of delayed interactions seen in cosmic ray events, as reconstructed from the calibrations described in the text. Not shown are the approximately 14% that fall outside of the range shown. These are clear cases of mis-reconstructions due to the phonon channel energy. The distribution is peaked near 1.0, consistent with electron recoils.

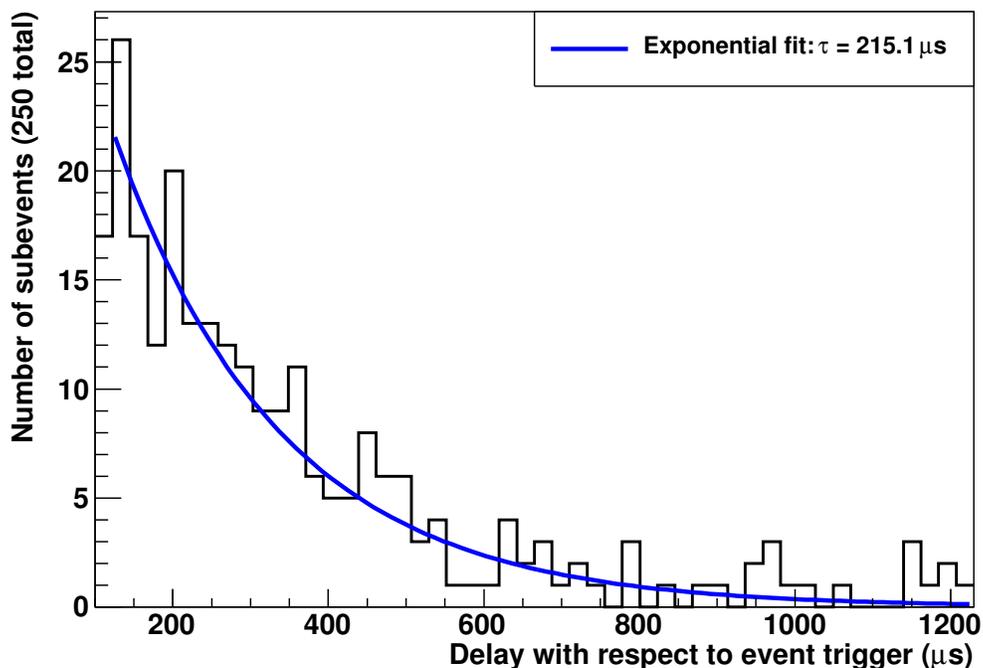


Figure 4.9: Distribution of delay times greater than $100 \mu\text{s}$ of the subevents found within cosmic ray events, with a fit showing the characteristic delay time.

(ionization signals which are within $5 \mu\text{s}$ of one another) which occur at least $100 \mu\text{s}$ later than the event trigger. 250 such delayed subevents were seen in cosmic ray events in the data from Runs 125–8. Figure 4.9 shows that the characteristic delay of these subevents is $215 \mu\text{s}$.

Figure 4.10 shows the ZIP-summed recoil energies of these subevents. The distribution can be represented approximately by a decreasing exponential up to an energy somewhat above 2 MeV. This is consistent with the notion of 2.2 MeV photons produced by neutron capture on hydrogen subsequently losing energy as they travel to a ZIP detector, depending on how much polyethylene they traverse. There also appears to be a population of photons with energies greater than 3 MeV. Photons produced by thermal neutron capture on nuclei in the copper cans surrounding the detectors could be responsible. The mean free path of a thermal neutron in polyethylene is around 40 cm,

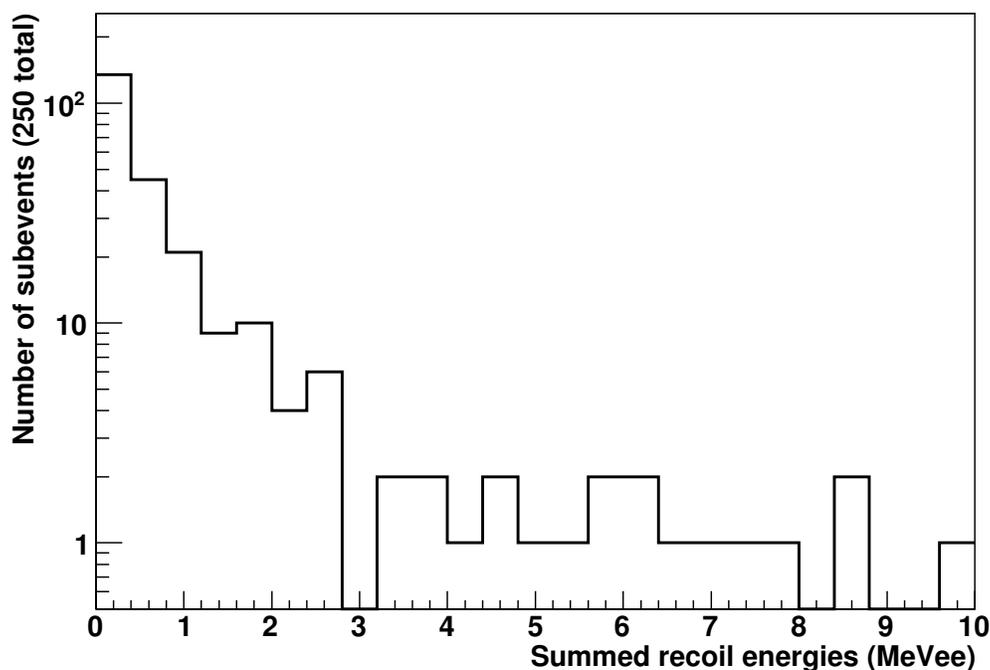


Figure 4.10: Summed recoil energy of the delayed subevents found within cosmic ray events. The energies are reconstructed from calibrations of the ionization signals for each ZIP. Hence the energies are accurate under the assumption that the interactions are electron recoils.

so it is not unlikely that some escape the polyethylene shielding after becoming thermalized. The thermal neutron capture cross-section on copper is 3.8 barns, and the number density of copper nuclei is about the same as the number density of hydrogen nuclei in polyethylene. Thus a thermal neutron is some ten times more likely to be captured by a copper nucleus than a hydrogen nucleus, if it manages to escape the polyethylene. The photons released by a copper nucleus after neutron capture have energies between 7 and 8 MeV. This could potentially explain the full spectrum of delayed subevents.

These neutron capture signals represent a larger sample of observed cosmogenic neutrons to compare with the results of Monte Carlo simulations. This could be a useful tool for assessing the apparent discrepancy in cosmogenic neutron production rates between simulation and observation.

4.4.2 Calculation of the background estimate

The cosmogenic background is defined as the number of cosmogenic neutrons expected to pass all WIMP selection criteria. The estimate is based on both simulation and observation. A neutron rate is determined from observed data. This is combined with simulation results to produce an estimate of the number of neutrons expected to escape discrimination (“leaked” neutrons).

There are several possible categories of observed neutron events that can be chosen to determine the neutron rate. There is a tradeoff between reducing statistical error by choosing a category with a high count and reducing systematic error by choosing a category for which observation and simulation should agree well. The category chosen was neutrons in germanium without ZIP coincidence. The count for this category is small, only three events (see table 4.2).

Basing the neutron rate on all germanium neutron events would have introduced a significant systematic error. This is because many ZIPs were rejected as WIMP-search detectors for Runs 125–8 due to poor performance, in contrast with the full complement of germanium detectors assumed to be working perfectly in the simulation. Of 19 germanium ZIPs, the number deemed good enough for WIMP search ranged from 14 during Run 125 down to 10 for Run 128. Therefore a significant number of events that include a neutron interaction in a germanium ZIP cannot be identified as such⁴.

To minimize systematic error, multiple data sets were used to calculate the neutron rate. Cosmic ray coincident data from Runs 123–4 and Runs 125–8 were combined, thus including all data taken in the 5-tower configuration. No neutrons without ZIP coincidence were observed in Runs 123–4, so the count is still three. The combined neutron search exposure is 844 kg-days. This is calculated using the efficiencies of all selection criteria used to identify cosmogenic neutrons.

To calculate the leakage rate, the neutron rate is multiplied by an event ratio calculated from the simulation results. This is the ratio of non-vetoed neutrons to vetoed neutrons for events without ZIP coincidence. Three different simulations, all using the

⁴In principle the simulation data could be reanalyzed to represent the configuration of well-functioning detectors that existed during Runs 125–8. Small counts in some key categories however made this impractical.

Soudan 5-tower configuration, were combined to produce this ratio. The individual ratios were $2/36$, $0/9$, and $33/561$, from the most recent Geant4 simulation, from an older Geant4 simulation, and from a FLUKA simulation, respectively. Note that these ratios are statistically consistent with one another. The resulting leakage rate is

$$\frac{3}{844 \text{ kg} \cdot \text{days}} \times \frac{2 + 0 + 33}{36 + 9 + 561} = 2.0 \times 10^{-4} \text{ per kg} \cdot \text{day}, \quad (4.1)$$

or 0.07 per kg · year.

Once the leakage rate is known, one need only multiply by the WIMP-search exposure for Runs 125–8, 194.1 kg·days, to calculate the background estimate. The final estimate of the cosmogenic neutron background is

$$2.0 \times 10^{-4} \times 194.1 = 0.04_{-0.02}^{+0.04}(\text{stat}) \quad (4.2)$$

Chapter 5

Bulk Electron Recoil Background

Neutrons are the only category of background that produce nuclear recoils in the bulk of a ZIP. Other background particles produce electron recoils, which are usually identifiable by their ionization yields. This identification can fail if the interaction occurs very close to one of the faces of a ZIP. These *surface electron recoils* are the subject of chapter 6. In this chapter I describe the yield-based discrimination of *bulk electron recoils*.

5.1 Origins

In standard low-background running (i.e. WIMP-search mode), the CDMS II experiment records information when triggered by any of the phonon signals exceeding a certain threshold above their noise. The greatest trigger rate is produced by photons produced by radioactive materials in the hardware proximate to the ZIPs, as described in section 3.3. In Runs 125–8 the rate of such photon triggers was 0.127 Hz (or 1 trigger every 7.9 seconds on average). The rate is 0.091 Hz (or 1 trigger every 11.0 seconds on average) for photons which deposit energies greater than 10 keV, which is the threshold used for the standard WIMP-search analysis. These photons usually interact in ZIPs by producing bulk electron recoils.

5.2 Identification

Bulk electron recoils are identified by means of their ionization yields. Each ZIP must be studied to characterize its response to recoil type. Interactions from calibration data are studied for this purpose. These calibration data-taking periods are interspersed among the WIMP-search data-taking periods so that any time variations in detector behavior can be tracked. If such variations are small, the procedure provides as unbiased a sample as possible.

The electron recoil calibration sample is generated using a ^{133}Ba gamma radiation source. The nuclear recoil calibration sample is generated using a ^{252}Cf neutron emitter. This source produces both photons and neutrons, so both interaction types are present in the data. Figure 5.1 shows an example of ionization yield vs. recoil energy for data from both sources, with high-yield interactions removed from the ^{252}Cf data to highlight the nuclear recoils. Plotted in this way, electron recoils and nuclear recoils form separate bands centered at yields of approximately 1.0 and 0.3, respectively.

The reconstructed ionization yield varies with the interaction's position in the X-Y plane of the ZIP. Discrimination is significantly improved by applying a correction procedure to the calculated ionization yield to remove its variations with X-Y position. The position of an interaction is determined by comparing the signals from the four phonon sensor quadrants. For every interaction in a large sample of bulk electron recoils, a correction factor is determined by calculating both a global average and a local average of the yield at its position. Any ionization yield can then be corrected by multiplying by the correction factor corresponding to the nearest interaction in the sample. The yield bands are significantly narrower after this procedure is applied.

The yield bands are characterized by fitting ionization yield distributions in several recoil energy bins to Gaussian functional forms. The resulting widths (σ values) and centers are then fit as functions of energy. The center of the electron recoil band is nearly constant with energy, as expected by the definition of ionization yield. The energy dependence of the nuclear-recoil band is in good agreement with semiconductor theory [29].

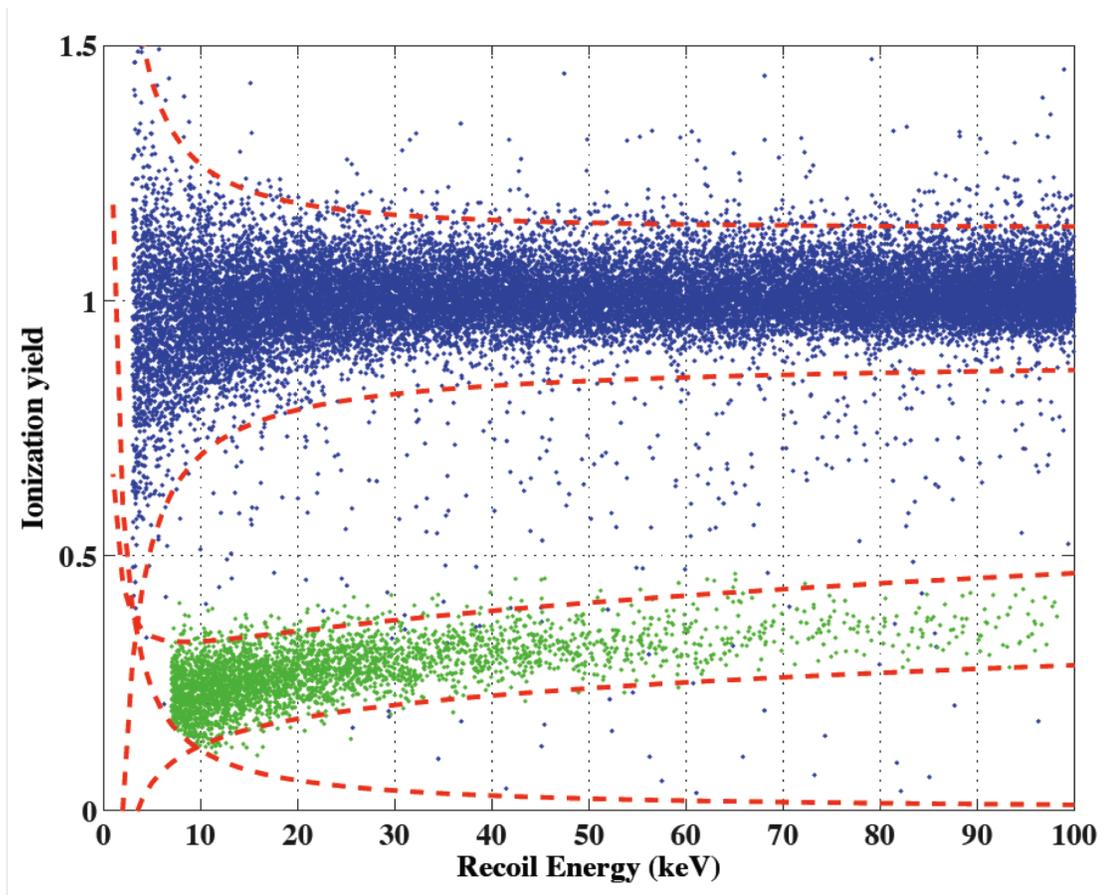


Figure 5.1: Ionization yield vs. recoil energy in a ZIP. Data are from two different radiation sources used for calibration purposes. Blue points are electron recoils from a ^{133}Ba gamma source. Green points are nuclear recoils from a ^{252}Cf neutron source. Four dashed red lines indicate band definitions used to identify the recoil type. An additional line (with asymptotes on both axes) delineates a threshold for the ionization signal.

5.3 Discrimination criterion

To discriminate the bulk electron recoil background, a WIMP candidate is required to satisfy all of the following:

- The ionization yield is within 2σ of the nuclear recoil band center. To the extent that the yield distribution is Gaussian, this should be 95% efficient.
- The recoil energy is above a threshold of 10 keV. This avoids the region in which the yield bands overlap.
- The ionization signal strength is above a noise threshold, defined as 4σ greater than the center of a Gaussian fit to noise signals. If no significant ionization signal is seen, yield reconstruction is not considered trustworthy. Since nuclear recoils produce smaller ionization signals this threshold is significant for low-energy nuclear recoils. In a yield vs. energy plot such as figure 5.1 this threshold appears as a sharply-falling curve with asymptotes on both axes.
- The ionization yield is more than 3σ below the electron recoil band center. For most ZIPs this is a redundant condition because the yield bands are well-separated above the recoil energy threshold.

The efficiency of this discrimination criterion varies with energy and with detector. Overall it is about 95% as expected within the recoil energy analysis range (10–100 keV). The efficiency decreases by a few percent at recoil energies below about 20 keV due to the effects of the ionization signal threshold and the requirement that the ionization yield be below the electron recoil band. The efficiency also decreases slightly as recoil energy increases to the maximum of the WIMP search analysis range. This is because the nuclear recoil band's deviation from Gaussianity is greater at higher energies.

5.4 Final background estimate

Because above-threshold bulk electron recoils are so readily discriminated, this background is extremely small. Using the observed detector resolutions, Monte Carlo simulations were performed to calculate the leakage fraction for each ZIP. Combining this

with the bulk electron recoil rate and live time for each ZIP, the total background is calculated to be less than 5×10^{-4} events for Runs 125–8. Were this the only category of background, CDMS II would qualify as a background-free experiment.

Chapter 6

Surface Electron Recoil Background

In figure 5.1 it can be seen that most electron recoils produced by photons from the ^{133}Ba calibration source fall into the electron recoil ionization yield band. However it is also apparent that there is a population of interactions with ionization yields below the electron recoil band. Those with yields within the nuclear recoil band cannot be discriminated from WIMPs using the criterion of the previous chapter.

In the recoil energy range of interest for WIMP search (10–100 keV), about 1 in 10,000 photon interactions in events without ZIP-coincidence have a measured ionization yield consistent with a nuclear recoil. Thus there is a 1 in 10,000 failure rate in the ability to discriminate photons in this energy range from WIMPs via yield and coincidence. Electron recoil interactions with yields significantly below the electron recoil band are identified as surface interactions.

A great deal of analysis effort goes into surface electron recoil discrimination. This is primarily because about 50% of potential WIMP candidates must be rejected in order to remove this background. Thus the discrimination criterion has a very strong effect on the WIMP sensitivity of the experiment. Surface interaction distributions can vary greatly by ZIP due both to details of the ZIP response and to the location of the radioactive contaminants. Development of surface interaction discrimination is made more difficult because of low statistics of surface interaction data samples, both

from calibration and from low-background (WIMP-search) data sets. Procedures for developing discrimination criteria have evolved during CDMS II, so that for Runs 125–8 they are rather complex.

To simplify the discussion to follow I define the following abbreviations and terms:

SER: Surface Electron Recoil

BNR: Bulk Nuclear Recoil

Discriminator: a quantity characterizing phonon timing derived from phonon pulse shape analysis. If this quantity takes on different distributions for SERs and BNRs, it can be used as a means of discrimination.

6.1 Origins

The mechanism for suppressed ionization yields in surface interactions is well understood. Charge carriers excited by a bulk interaction are eventually collected in the electrode of the opposite charge sign. However for interactions occurring within a layer very close to the surface, the nearby electrode will collect a significant number of same-sign charge carriers [30]. This suppresses the ionization yield signal. The “dead layer” in which this occurs has been measured to be approximately $10\ \mu\text{m}$ thick [31]. The physical definition of a *surface interaction* is one that occurs primarily in the dead layer.

A small fraction of photons interact in the dead layer with recoil energies above 10 keV. Photons can also induce surface interactions by producing energetic electrons as a secondary particles. These electrons interact with electrons in the crystal. An electron with less than 100 keV kinetic energy will penetrate less than $10\ \mu\text{m}$ into a germanium detector.

The rate of photon interactions seen in low-background running cannot entirely account for the rate of surface interactions seen. Electrons produced by beta decay of radioactive isotopes are another source of surface interactions. Since such electrons are easily blocked by the innermost layer of copper surrounding the detectors, the relevant radioactive contamination can only be on the inner surfaces of the ZIP housings and on the detector surfaces themselves. Studies of the energy spectra of surface interactions

and of interactions identified as alpha radiation show that the primary contamination is from ^{210}Pb , a radon daughter isotope [32]. Such contamination can be reduced by adopting clean procedures for handling detectors during production and testing. Improvements in these procedures have resulted in lower surface interaction rates seen in later detectors.

6.1.1 Categories of surface interactions

As with other types of background the ability to discriminate SERs varies with the recoil energy. The properties of SERs also vary by which face of the detector they occur on. These differences arise from ZIP design and fabrication.

Surface interactions are often accompanied by a simultaneous interaction in the adjacent ZIP. The second interaction can be caused by either a photon scattering between adjacent surfaces, or by a photon ejecting an electron from one surface which then interacts on the adjacent surface. This provides a means for identifying the face on which the surface interaction occurred. This identification tool is useful in characterizing surface interactions, both in calibration data and in WIMP-search data.

The two detector faces are labeled phonon-side and ionization-side based on the sensor locations. The phonon-side dead layer is somewhat thicker due to fabrication effects, so that phonon-side surface interactions more often have suppressed ionization yields. Phonon-side SERs also have faster phonon collection because some of the slow primary phonons are collected promptly. Thus phonon timing-based discrimination is less effective on ionization side SERs than on phonon-side SERs.

Surface interactions can also be categorized by whether they have ionization yields within the nuclear recoil band and whether they are coincident with other ZIP signals. The specific properties of non-coincident, nuclear-recoil-band SERs are important for discrimination, since these will escape other discrimination criteria.

6.1.2 Rate estimates for surface interactions in WIMP-search running

To estimate the rate at which SER interactions occur with WIMP-like signatures, data from previous five-tower runs (Runs 123–4) were analyzed. The least-biased sample consists of SERs which satisfy all WIMP selection criteria (except SER discrimination

Detector	count	live time (days)	rate	averaged rate
T1Z2	9	138.4	0.07 ± 0.02	0.053 ± 0.011
T1Z5	8	137.0	0.06 ± 0.02	0.053 ± 0.011
T2Z3	5	82.3	0.06 ± 0.03	0.053 ± 0.011
T2Z5	2	97.3	0.02 ± 0.01	0.053 ± 0.011
T3Z2	1	38.4	0.03 ± 0.03	0.027 ± 0.006
T3Z4	4	87.7	0.05 ± 0.02	0.027 ± 0.006
T3Z5	0	64.1	0	0.027 ± 0.006
T3Z6	6	69.0	0.09 ± 0.04	0.092 ± 0.021
T4Z2	4	97.6	0.04 ± 0.02	0.027 ± 0.006
T4Z4	3	153.4	0.02 ± 0.01	0.027 ± 0.006
T4Z5	5	137.5	0.04 ± 0.02	0.027 ± 0.006
T4Z6	13	137.7	0.09 ± 0.03	0.092 ± 0.021
T5Z4	3	147.6	0.02 ± 0.01	0.027 ± 0.006
T5Z5	4	153.2	0.03 ± 0.01	0.027 ± 0.006

Table 6.1

of course). Data from Runs 125–8 could not be used since it would violate the CDMS blind-analysis policy.

Table 6.1 shows rates of surface interactions with WIMP-like signals as determined from data from Runs 123–4. The table includes the 14 ZIPs designated as WIMP-search detectors for some or all of Runs 125–8. This information was used in the development of the SER discrimination criteria for Runs 125–8. Due to the very low counts on which these calculations are based, a decision was made to average the rates over certain categories of detectors expected to have similar rates. The ZIPs can be categorized by whether they belong to Towers 1–2 or Towers 3–5 — the latter category is known to have smaller contamination from radon daughter isotopes. They can also be categorized by whether they are situated at the top, bottom, or an intermediate position in the stack — the three categories have very different rates for surface interactions that escape coincidence tagging. The 14 WIMP-search ZIPs fall into three such categories: Tower 1–2 intermediate-position, Tower 3–5 intermediate-position, and Tower 3–5 bottom.

Information about how SERs are distributed was also taken from Runs 123–4 for use in discrimination development for Runs 125–8. The information in table 6.2 summarizes how surface interactions are distributed in energy and by detector face. The fractions of surface interactions that fall into various energy and face categories are shown. The

Face	Energy	T1–2 intermediate	T3–5 intermediate	T3–5 bottom
Ionization side	Low	0.306 ± 0.037	0.131 ± 0.020	0.511 ± 0.093
	Medium	0.060 ± 0.020	0.030 ± 0.011	0.090 ± 0.065
	High	0.093 ± 0.024	0.012 ± 0.008	0.210 ± 0.084
Phonon side	Low	0.191 ± 0.032	0.293 ± 0.026	0.037 ± 0.022
	Medium	0.115 ± 0.026	0.146 ± 0.021	0.031 ± 0.021
	High	0.235 ± 0.034	0.388 ± 0.028	0.120 ± 0.034

Table 6.2

recoil-energy categories are 10–20 keV, 20–30 keV, and 30–100 keV. The data used are surface interactions with ionization yields falling in the nuclear recoil band for which the ZIP face is identified due to coincidence in an adjacent detector. For the bottom ZIPs (T3Z6 and T4Z6), ionization-side SERs cannot be identified because there is no adjacent ZIP on that side. For these ZIPs, SERs not identified as phonon-side are assumed to be ionization-side.

If an SER is not ZIP-coincident the face on which it occurred cannot be identified. Applying the ratios in table 6.2 to non-coincident SERs relies then on the supposition that the face ratios are the same between face-identified and non-coincident SERs. This supposition is strengthened somewhat by the similarity between the energy spectra of nuclear-recoil-band SERs in these two categories.

The SER categories described in this section were specifically chosen to represent populations of SERs whose properties differ significantly. The purpose of this is to minimize systematic error in leakage estimates. To minimize statistical error, SER populations are grouped together when their properties appear consistent, as is the case for example in the detector groupings described above.

6.2 Identification

The phenomenological definition of an SER is an electron recoil whose ionization yield falls below the electron recoil band. Thus by definition these interactions are not distinguishable from BNRs based on ionization yield. SER discrimination is instead based on phonon timing as characterized by analyzing phonon pulse shapes.

6.2.1 Phonon timing

Phonons produced in surface interactions arrive at the phonon sensors characteristically faster than those produced by bulk interactions. Figure 6.1 shows representative examples of phonon pulses produced by an SER and a BNR. The trace from the phonon quadrant that receives the greatest energy is referred to as the *primary quadrant*. This quadrant almost always has the fastest pulse. Experience has shown that the primary quadrant contains almost all of the information that can be used for discrimination.

As will be described in the following sections, SER discrimination in CDMS II is essentially based on empirical information obtained by studying SERs (from calibration data and low-background data) and BNRs (from calibration data) on a detector-by-detector basis. Detector Monte Carlo simulations which explore the phonon physics in ZIPs are currently in development and may play a rôle in developing discrimination criteria in future phases of CDMS.

6.2.2 Calibration data samples

Calibration data are used to study the discrimination of SERs from BNRs. The SER data sample is taken from ^{133}Ba data. The surface interactions are identified by ionization yields more than 5σ below the (bulk) electron recoil band to avoid contamination of the sample with bulk interactions in the non-Gaussian tails of their distribution. Such contamination can be seen empirically by changes in the distribution of pulse shape discriminators when less-strict selection criteria are applied.

Because only a small fraction of particles from ^{133}Ba cause surface interactions, a large amount of calibration data must be collected in order to produce sufficient statistics to develop discrimination criteria. Two factors exacerbate the situation. The first is that the allowable leakage fraction for surface interactions is so small that the critical sub-population of the surface interaction sample lies in the tail of the distribution where statistics are smaller still. The second is that systematic differences exist between the surface interaction populations in the calibration sample and in WIMP-search data. This makes it important to divide the data into categories such as those shown in table 6.2. Naturally there is a tradeoff between amassing ^{133}Ba calibration data and maximizing the time and resources devoted to searching for WIMPs. The low statistics

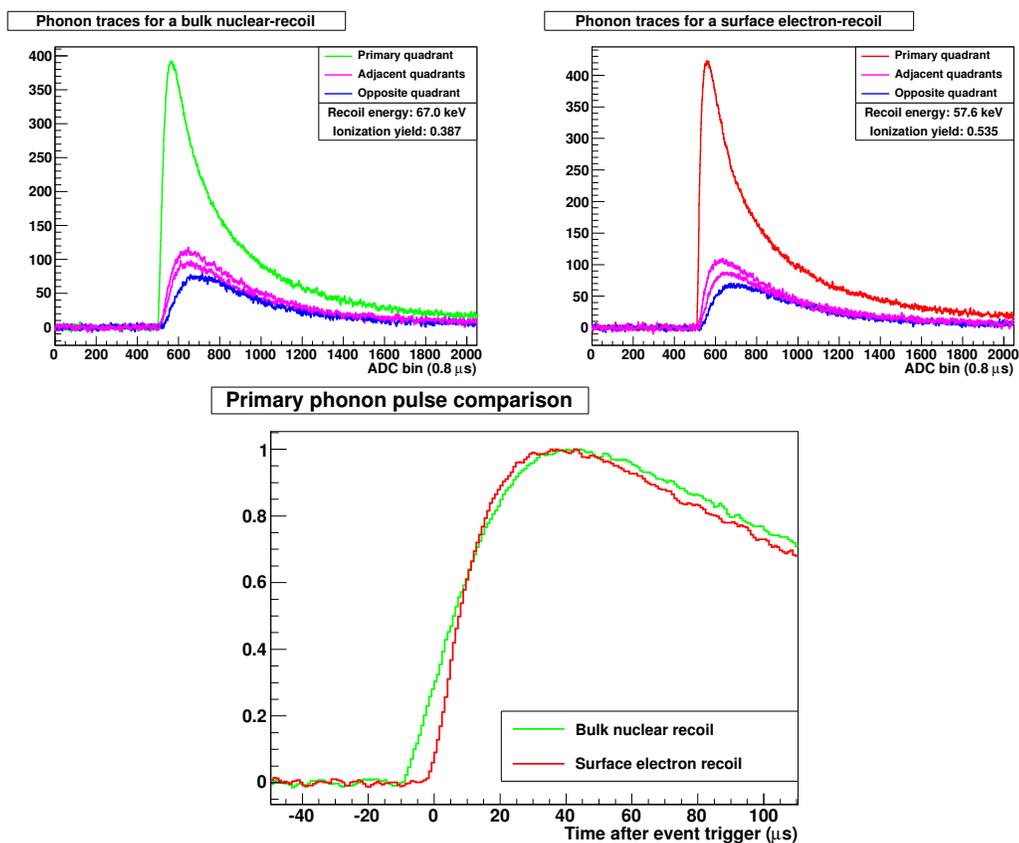


Figure 6.1: Comparison of pulse shapes for a typical SER vs. a typical BNR. All four raw phonon traces for a BNR (upper left) and an SER (upper right) are shown. The phonon quadrant registering the greatest energy is designated the “primary” quadrant, with “adjacent” and “opposite” quadrants defined with respect to the primary quadrant. A comparison of the normalized primary phonon traces (bottom) shows the characteristic faster rise for the surface interaction. The nuclear recoil is from ^{252}Cf calibration data and the electron recoil is from ^{133}Ba calibration data. In both cases the interaction type and location are identified by their ionization yields.

in the surface interaction sample are a limiting factor in developing and characterizing surface interaction discrimination.

The BNR data sample is taken from ^{252}Cf data. BNRs are selected by choosing ionization yields within 2σ of the nuclear recoil band centroid. The BNR samples are much larger than the SER samples.

6.2.3 Differences between calibration and WIMP-search surface interaction distributions

The energy spectra of SERs in calibration data and WIMP-search data differ as illustrated in figure 6.2. The most significant difference is at low energies, where there are fewer calibration interactions. This is because bulk electron recoils must be rejected when selecting the sample, and the ionization yield band for bulk electron recoils is wider at low energies (see figure 5.1). Any assessment of surface interaction discrimination that did not take these differences into account would be biased. The bias would be reduced if we considered only those calibration data whose ionization yields fall within the nuclear recoil band, but then the sample size would be reduced by an order of magnitude.

The relative number of ionization-side vs. phonon-side surface interactions also differs since the particles producing SERs originate from very different locations. For calibration data this is a single source exterior to the copper cans. For WIMP-search data the radioactive contaminants are located on the detector surfaces and on the inner walls of the copper towers. Table 6.3 shows the ratio of ionization-side interactions to phonon-side interactions. For each WIMP-search detector, ratios are shown for the calibration data, for WIMP-search data, and for WIMP-search data in the nuclear recoil ionization yield band (WS, NR yield). All information in the table is from surface interactions that are “tagged” by a coincident signal in a neighboring detector. For the bottom ZIPs, the ratio of non-phonon-side interactions to phonon-side interactions is shown (in italics). These two detectors are not included in the summed ratios shown in the bottom row.

Differences between the first column (all calibration data) and the last column (WIMP-search interactions with ionization yields falling in the nuclear recoil band) are important for assessing discrimination. Also of note is that the values in the second

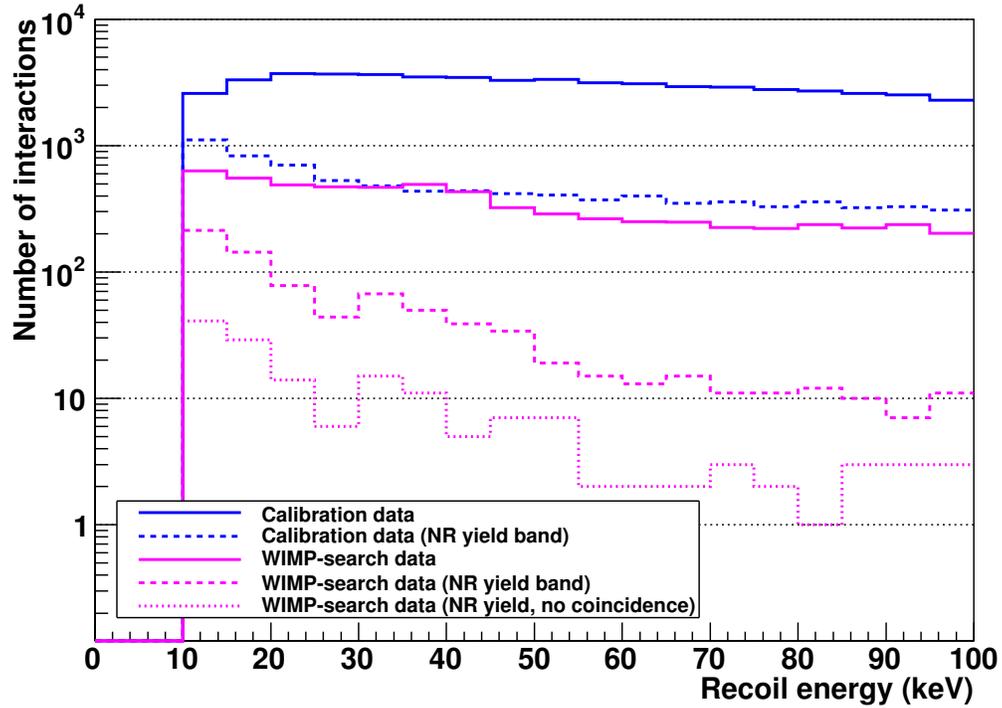


Figure 6.2: SER energy distributions. All data are from Runs 125–8, and each distribution curve shown combines data from all 14 WIMP-search detectors. The most important distributions are the top curve, representing the full calibration data sample from a ^{133}Ba gamma radiation source, and the bottom curve, representing interactions satisfying all WIMP-signal criteria apart from surface interaction discrimination. (In principle the latter distribution could include WIMPs and neutrons, but in retrospect we know these possible contributions to be very small.)

Detector	Calibration	WIMP-search	WS, NR yield
T1Z2	1.44	1.33	0.50
T1Z5	1.11	1.79	0.47
T2Z3	0.45	1.56	1.22
T2Z5	0.91	2.10	1.59
T3Z2	0.90	0.41	0.21
T3Z4	0.47	0.84	0.12
T3Z5	0.64	2.39	0.71
T3Z6	<i>4.71</i>	<i>2.53</i>	<i>2.39</i>
T4Z2	0.47	1.39	0.74
T4Z4	0.30	0.85	0.19
T4Z5	0.39	0.95	0.32
T4Z6	<i>0.94</i>	<i>1.25</i>	<i>0.31</i>
T5Z4	0.66	0.68	0.11
T5Z5	0.32	0.79	0.04
All inner	0.53	1.20	0.44

Table 6.3: Ratio of ionization-side to phonon-side SERs

column are invariably larger than the values in the last column. This demonstrates that ionization-side SERs are less likely to have ionization yields low enough to fall into the nuclear recoil band.

6.2.4 Phonon timing as a discrimination tool

While phonon pulse shape information (phonon timing) provides SER discrimination, it is by no means as powerful as the ionization yield discrimination between bulk electron recoils and nuclear recoils, since phonon pulse shapes for bulk nuclear recoils vary over a wide range. The problem is illustrated in figure 6.3. In comparison with ionization-yield discrimination (figure 5.1) the pulse shape discriminator bands of the two event categories are strongly overlapping. This is why SER discrimination typically has an efficiency of only about 50%.

A convenient feature of SER discriminators is that they tend to vary little with recoil energy over the 10-100 keV analysis range, as can be seen for the discriminator shown in figure 6.3. If there were a strong energy dependence, it would be much more difficult to characterize the discrimination since data samples are so small. In contrast to bulk electron recoil discrimination, in which energy-dependent yield bands are constructed,

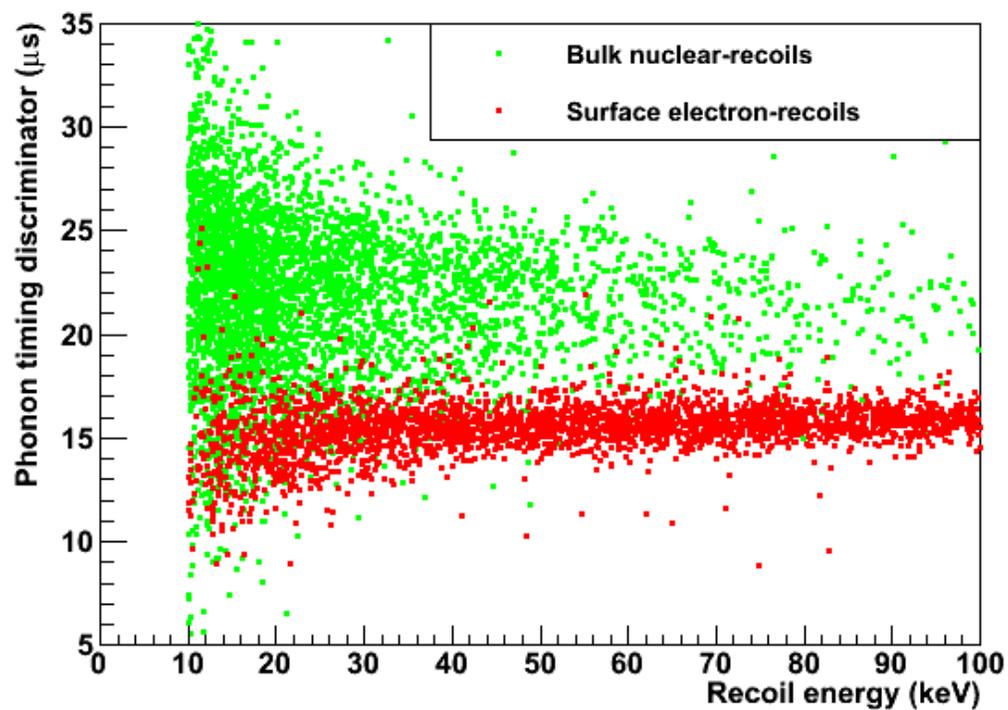


Figure 6.3: Primary phonon pulse timing discriminator vs. recoil energy. The shape of the rising edge of the primary phonon pulse has been characterized by a single quantity with the purpose of discriminating SERs from bulk nuclear recoils. (Shown here is the discriminator ultimately chosen for the analysis of Runs 125–8.)

SER discrimination is based on a simple discriminator threshold (with some elaborations that will be discussed in this chapter).

It is the central value of the discriminator distribution that tends to be constant in energy; the width of the distribution increases with decreasing energy, as the signal-to-noise ratio decreases. Thus SER leakage fractions are strongly energy dependent.

6.2.5 Discriminator categories

Several types of quantities have been used to characterize the discriminating features of the primary phonon pulse shape. The two which provide the most discriminating power are:

Rise time: the time interval in which the pulse rises from one percentage of its maximum height to another.

Delay with respect to the ionization signal: the time interval between the onset of the fast ionization pulse and a point along the rising edge of the phonon pulse. Unless the point chosen is the “0% time” (the instant the pulse begins to rise above the baseline), this quantity will necessarily also include rise time information. In the remainder of this chapter I will refer to these discriminators as simply **delay** discriminators.

Figure 6.4 illustrates these discriminator categories.

6.2.6 Algorithms for extracting phonon pulse shape information

To calculate these discriminators, a number of different methods for extracting pulse shape information from the raw pulses have been developed. Those that proved most useful for discrimination are the following:

“Walk” algorithms

One way to find the times at which a certain percentage of the maximum height is reached is to find the pulse maximum, then “walk” back one data point at a time until the percentage in question is reached. For interaction energies in the lower part of the analysis range, the noise in the signal would make this procedure unreliable if

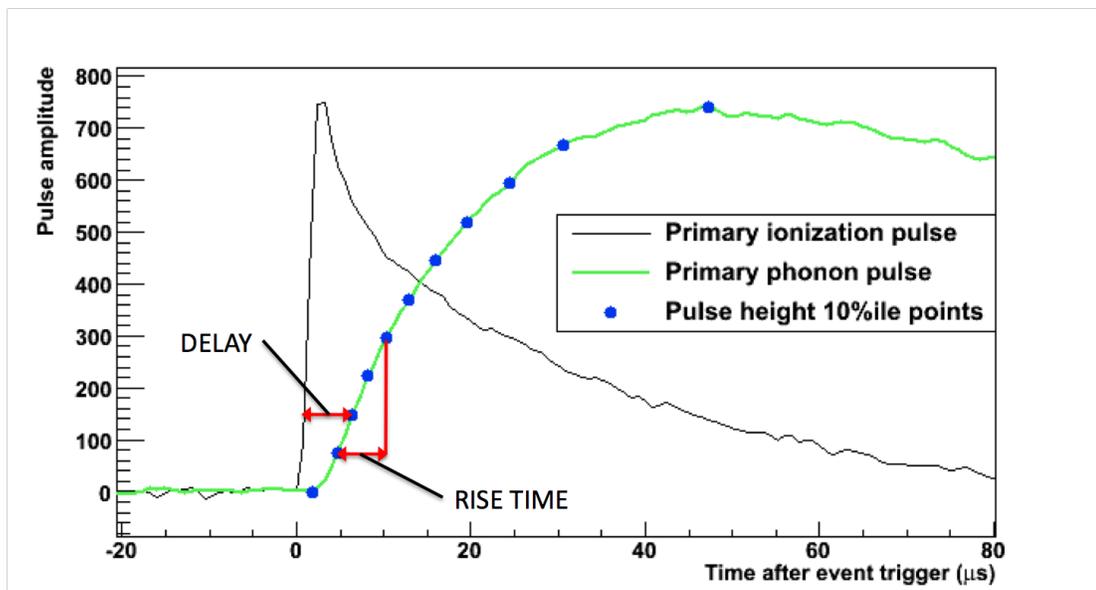


Figure 6.4: Schematic illustration of two categories of SER discriminator. In this case the *delay* is characterized as the time interval between the onset of the ionization pulse and the 20% time of the phonon pulse. The *rise time* depicted is the time interval in which the phonon pulse rises from 10% to 40% of its maximum height.

performed on a raw pulse. Therefore the method that has become standard for CDMS II germanium ZIPs is to first apply a low-pass Butterworth filter with a 50 kHz cutoff, then perform the walk algorithm on the smoothed pulse. Figure 6.5 shows an example of the results of this procedure.

For Runs 125–8 a second walk algorithm was also implemented using a filter cutoff frequency that varied with the signal-to-noise ratio. The intent was to apply stronger smoothing to small pulses and weaker smoothing to large pulses to optimize the tradeoff between loss of pulse shape information and misrepresentations introduced by noise fluctuations.¹

The discriminator quantities produced by this variable-frequency version tend to have a strong energy dependence at low energies. Apparently low-energy pulses are “over-smoothed” by the algorithm, such that their rise times are exaggerated. Despite this failing, the quantities produced by this algorithm were expected to be less susceptible to noise in low-energy pulses than those from the constant-frequency cutoff version, so a correction procedure was applied to produce discriminators without energy dependence.

Thus there are two different walk algorithms that were found useful for discrimination. For brevity, in this thesis I will refer to these two versions of the walk algorithm as:

CF-walk: the walk algorithm using a constant cutoff frequency for the Butterworth filter

VF-walk: the walk algorithm using a variable cutoff frequency for the Butterworth filter, along with an energy correction for the derived discriminators.

PipeFitter functional form fit

An alternative method to extract the pulse shape information is to fit a functional form to the trace, then calculate the desired quantities analytically from the fit function. An algorithm referred to as *PipeFitter* was used in Runs 123–8 for this purpose.

¹The adjustments to the filtering are particularly important for characterizing the timing of phonon pulses from the non-primary quadrants because they are smaller. This information is key to determining the location of an interaction in a plane parallel to the detector faces, but not directly useful for identifying surface interactions.

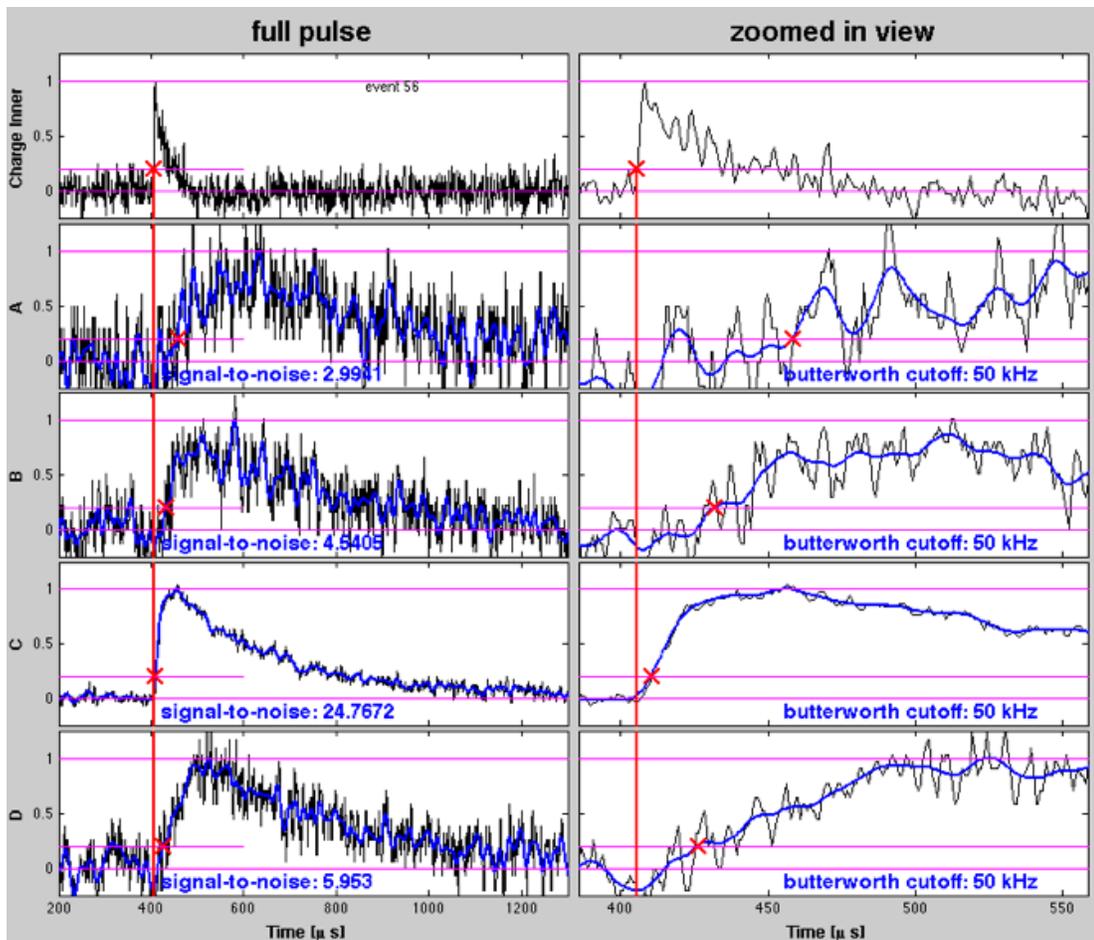


Figure 6.5: Results of the CF-walk algorithm on the four phonon traces (and the inner ionization trace, “Charge Inner”) of a sample interaction. The red crosses show the 20% point on the rising edge as determined by the algorithm.

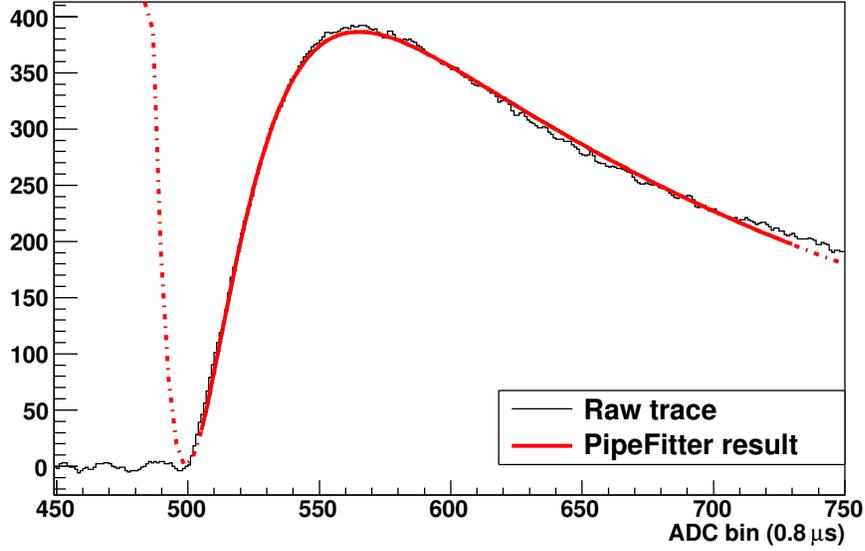


Figure 6.6: Results of the PipeFitter algorithm on a primary phonon pulse. Also shown is an extrapolation of the fit function outside of the fit range.

In the PipeFitter algorithm the empirically-chosen function

$$f(t_{\text{ADC}}) = a \left[1 - \exp\left(-\frac{t_{\text{ADC}} - t_1}{\tau}\right) \right] \left[\exp\left(-\frac{t_{\text{ADC}} - t_1}{\kappa}\right) - b \exp\left(-\frac{t_{\text{ADC}} - t_1}{\tau}\right) \right] \quad (6.1)$$

with free parameters a , b , τ , κ , and t_1 is fit to the phonon pulse over a range beginning at the 5% point on the rising edge and ending approximately 100 ADC bins ($80 \mu\text{s}$) after the maximum. This fit is applied to phonon pulses for all interactions with energies within the analysis range of 10–100 keV (although at the lowest energies the fit range is extended to include the entire pulse).² Figure 6.6 shows an example of the results of the PipeFitter algorithm for a sample primary phonon pulse.

²PipeFitter as originally written was also intended to provide an alternative energy reconstruction. However in Runs 125–8 only information from equation 6.1 was used, and only for the purpose of SER discrimination. What is described here is only a portion of the full PipeFitter algorithm, which fits different functional forms to pulses of different sizes.

6.2.7 Position/energy correction of discriminators

A correction procedure is used to remove variations with X-Y position from every candidate discriminator. For the primary phonon pulse these pulse shape variations are largely due to variations in the fraction of fast phonons that are collected in the primary quadrant. Irregularities in the fabrication of QET phonon sensor arrays can also cause position dependence. In order to maximize the effectiveness of surface event discrimination, it is necessary to separate this X-Y dependence from the Z dependence.

Unlike the correction procedure used for ionization yield, the correction for SER discriminators also removes residual energy variations. This is achieved by simply including energy with the position quantities to define the metric which describes the “location” of an interaction. The local average value of the discriminator is thus calculated over interactions that are close to one another in both X-Y position and energy.

Position/energy correction greatly increases discrimination power by tightening the distributions of rise time and delay discriminators, both for BNRs and SERs.

6.3 Discrimination criterion

6.3.1 Combinations of discriminators

There are a number of ways to base a discrimination criterion on two or more discriminator quantities. The simplest is to use a linear combination of the quantities as a discriminator and choose a threshold for it (see Figure 6.7). In the process of finding an optimal discrimination criterion (described in section 6.3.5) it was found that criteria based on linear combinations of two quantities were more powerful than those based on one. Smaller gains were seen when a third quantity was added.

In concert with the simple threshold criterion for a two-quantity linear combination discriminator (generally chosen by considering a maximum acceptable SER leakage), further requirements were added to reject interactions whose pulse shape characteristics lay far from the main distribution of BNRs. As shown in figure 6.8 a region in the discriminator space is defined which encloses the bulk of the BNR distribution. The boundaries of this region are the main discrimination threshold and three “consistency” boundaries. These appear in the figure as a line parallel to the main threshold line and

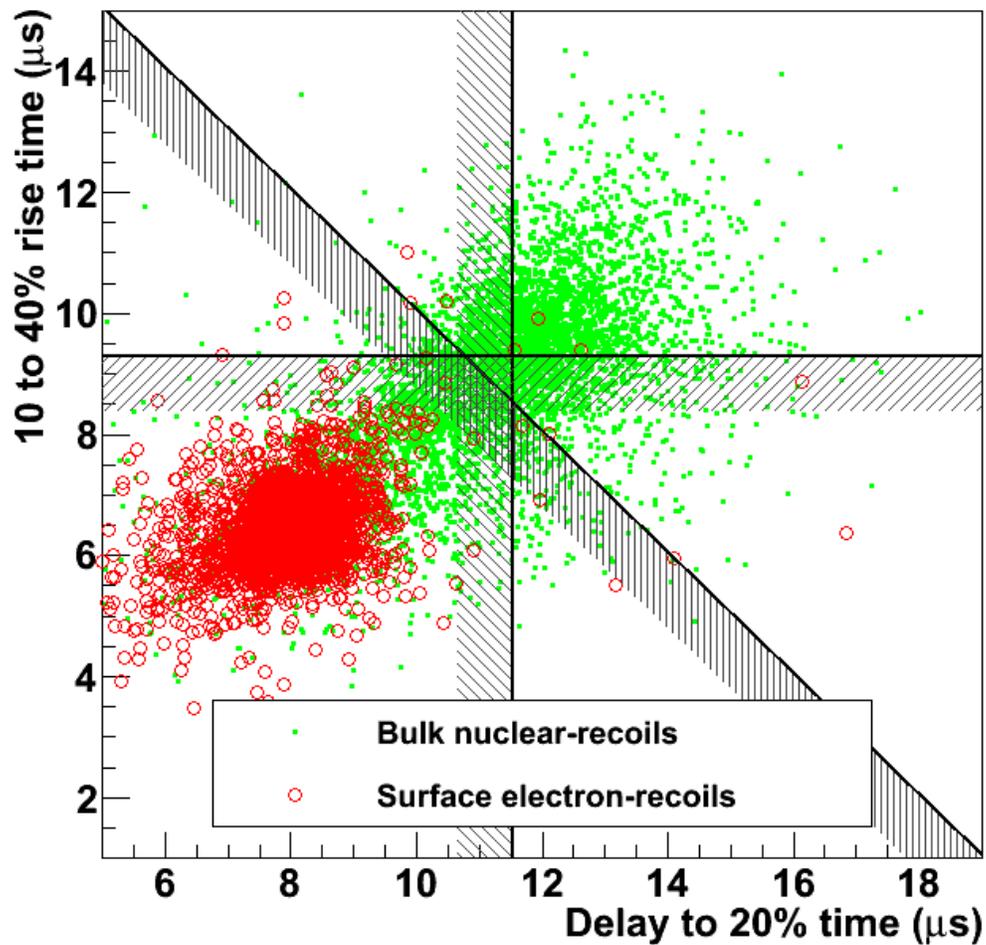


Figure 6.7: Three different discrimination criteria illustrated by three lines, with excluded regions highlighted. In each case the criterion is set to accept one in 300 SERs. The criteria are based on a delay quantity (vertical), a rise time quantity (horizontal), and the sum of the two (diagonal). The efficiencies of each criterion are 46%, 43%, and 57%, respectively. The discriminators are derived from the PipeFitter algorithm and the detector depicted is T1Z2.

two lines with slope opposite to the main threshold line. Each of these boundaries was set to accept 99.5% of BNRs. Thus without great loss to efficiency, interactions were removed which, due to their presence in the “tails” of the distribution, might be suspected of having poorly-reconstructed phonon timing characteristics. By applying the consistency criteria before calculating the threshold for the primary criterion, it is usually the case that the overall efficiency is improved, because certain surface interactions which “leak” into the bulk interaction region are already removed, and the primary threshold can be set to accept more bulk interactions.

6.3.2 Estimating leakage and efficiency/exposure

The SER discrimination criterion is set in an attempt to optimize the sensitivity of the experiment. To do so it is necessary to estimate the leakage and the efficiency resulting from the criterion.

For any discriminator, and for each ZIP, the SER expected leakage estimate in Runs 125–8 is

$$n(t) = N \sum_{e,f} s_{e,f} b_{e,f}(t) \quad (6.2)$$

where $n(t)$ is the number of surface electron recoils expected to “leak” into the WIMP-signal region expressed as a function of the discriminator threshold t . The subscripts e and f indicate specific surface interaction categories for recoil energy and for detector face. $s_{e,f}$ are the fractions of surface interactions expected in each category as shown in table 6.2. The $b_{e,f}(t)$ values are the leakage fractions of SERs in the various categories expressed as a function of the discriminator threshold. These are calculated using the calibration data. Because of the small data samples, the tails of the discriminator distributions in each category are fit to a functional form.

Finally, N is the number of SERs with WIMP-like signals expected in Runs 125–8. This is calculated as the product of the rate from table 6.1 and the live time for Runs 125–8. The results are shown in table 6.4.

I wrote a version of this procedure in the data analysis framework ROOT [33]. In this version the tails of the discriminator distributions are fit to a generalized Pareto

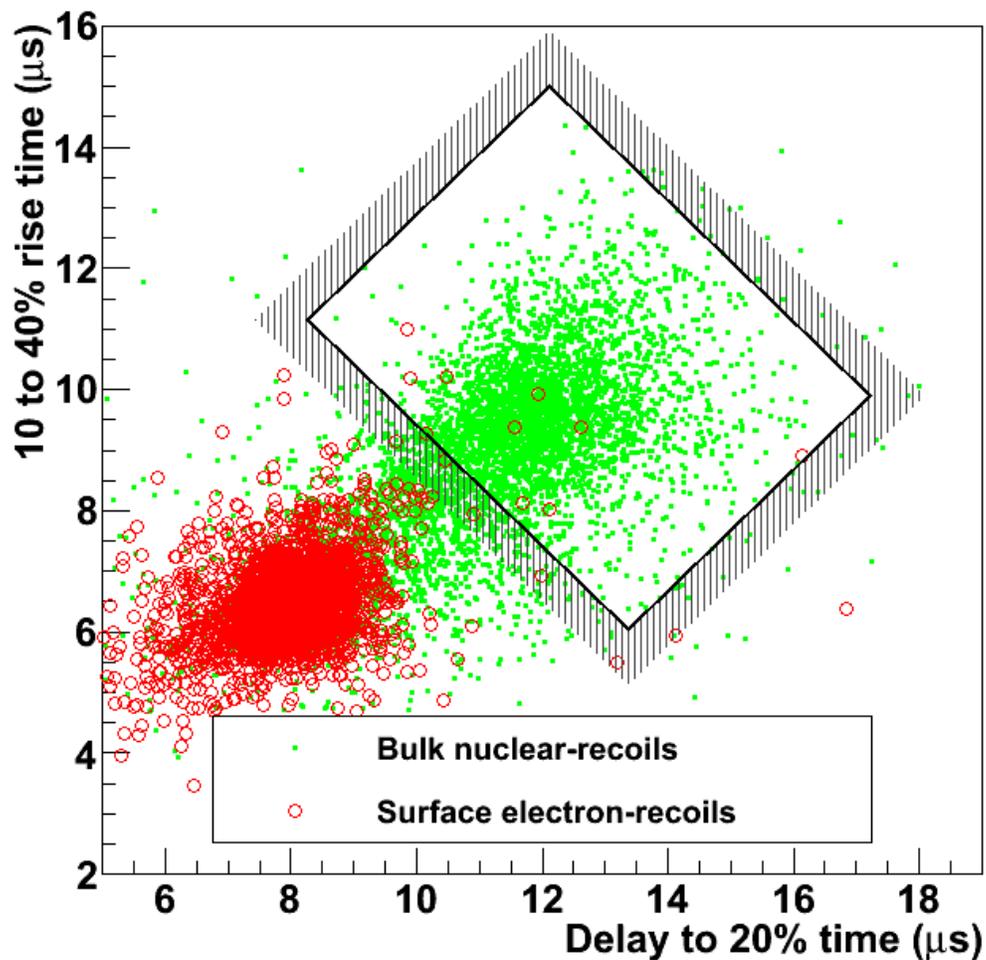


Figure 6.8: A discrimination criterion combining a threshold on a linear combination discriminator (the sum of a delay and a rise time) with “consistency” criteria designed to eliminate outlying interactions. The same quantities and data from the same detector are represented here as in figure 6.7. The efficiency of the discrimination criterion represented in this figure is 62%, in contrast to 57% when only the threshold criterion is used (as shown in the previous figure). In both cases the allowed leakage fraction is one in 300.

Detector	averaged rate	live time (days)	count
T1Z2	0.053±0.011	218.2	11.5±4.7
T1Z5	0.053±0.011	163.4	8.6±2.0
T2Z3	0.053±0.011	159.6	8.4±2.0
T2Z5	0.053±0.011	182.5	9.6±2.2
T3Z2	0.027±0.006	81.0	2.2±0.5
T3Z4	0.027±0.006	236.6	6.5±1.3
T3Z5	0.027±0.006	147.6	4.0±0.8
T3Z6	0.092±0.021	227.4	20.9±4.8
T4Z2	0.027±0.006	198.2	5.4±1.1
T4Z4	0.027±0.006	229.6	6.3±1.3
T4Z5	0.027±0.006	210.6	5.7±1.2
T4Z6	0.092±0.021	224.9	20.7±4.7
T5Z4	0.027±0.006	224.2	6.1±1.2
T5Z5	0.027±0.006	223.1	6.1±1.2

Table 6.4: Expected number of SERs with WIMP-like signals in Runs 125–8

function of form

$$L(t) = A \left(1 + \frac{C}{B}(t - t_0) \right)^{\frac{1}{C}-1} \quad (6.3)$$

with t representing the discriminator threshold, free parameters A , B , and C , and fixed parameter t_0 . In the limit $C \rightarrow 0$, this reduces to

$$L(t) = A \exp\left(\frac{t - t_0}{B}\right). \quad (6.4)$$

The initial value used for C is 0.01, so that the function is nearly exponential, but the fit is allowed to “relax” into a power-law function, up to a power as high as -1 . This allows for a possibly more conservative estimate than an exponential fit would provide in cases where the fit must be extrapolated beyond the data.

An example of the results of the leakage estimation procedure for a particular detector and discriminator is shown in figure 6.9. In this example the leakage estimate is dominated by ionization-side surface interactions in the lowest energy category. Also shown in the figure is the final threshold chosen for this detector and discriminator by the optimization procedure. Note that this value is greater than the discriminator value for the furthest outlier in any of the SER categories, necessitating that the fit functions

extend past the fit range.

The WIMP exposure can also be calculated as a function of discriminator threshold. The exposure can be taken to be the product of the live time shown in table 6.4 and the efficiency calculated using BNR calibration data. An accurate exposure calculation would also incorporate the efficiencies of all other WIMP selection criteria. Additionally, since these efficiencies are energy-dependent in general, the true efficiency is an average weighted by the expected WIMP spectrum. This spectrum depends on the WIMP mass; a 60 GeV/ c^2 WIMP is used as a standard. The energy dependence of the SER discrimination efficiency is small, so for the purposes of discrimination criteria development it can be calculated as a simple acceptance fraction. The weighted efficiency calculation is used for all other WIMP-selection criteria.

Once the SER leakage and WIMP exposure are both known as a function of the discriminator threshold, the exposure can be expressed as a function of the leakage for each ZIP. In the ROOT-code version of the procedure, this curve is stored as a set of points, one for every 0.001 leaked interactions. As can be seen in figure 6.10 this resolution is sufficient for encoding the shape of the curve.

6.3.3 Maximizing exposure on a ZIP-by-ZIP basis

The discrimination criterion is set by first choosing a target value for the total leakage. The exposure vs. leakage information is then used to maximize the exposure. A leakage setting can be chosen for each detector such that the sum equals the target total leakage. After an initial arbitrary choice for these settings, an optimization algorithm is applied to find the greatest total exposure by adjusting the leakage setting on a detector-by-detector basis, while keeping the total leakage constant. In the ROOT-code version, this is accomplished by shifting the leakage setting of one detector up by 0.001 while simultaneously shifting the leakage setting of another detector down 0.001. Whenever such a pair of shifts increases the total exposure, it is implemented and another exposure-increasing pair of shifts is sought. When no further exposure increase can be found with any possible pair of shifts, the algorithm ceases. Figure 6.10 illustrates the final result of this optimization procedure for a particular discriminator combination. It can be seen that, at the points on the exposure vs. leakage curves corresponding to their final leakage settings, the derivatives of the curves are approximately equal to one another.

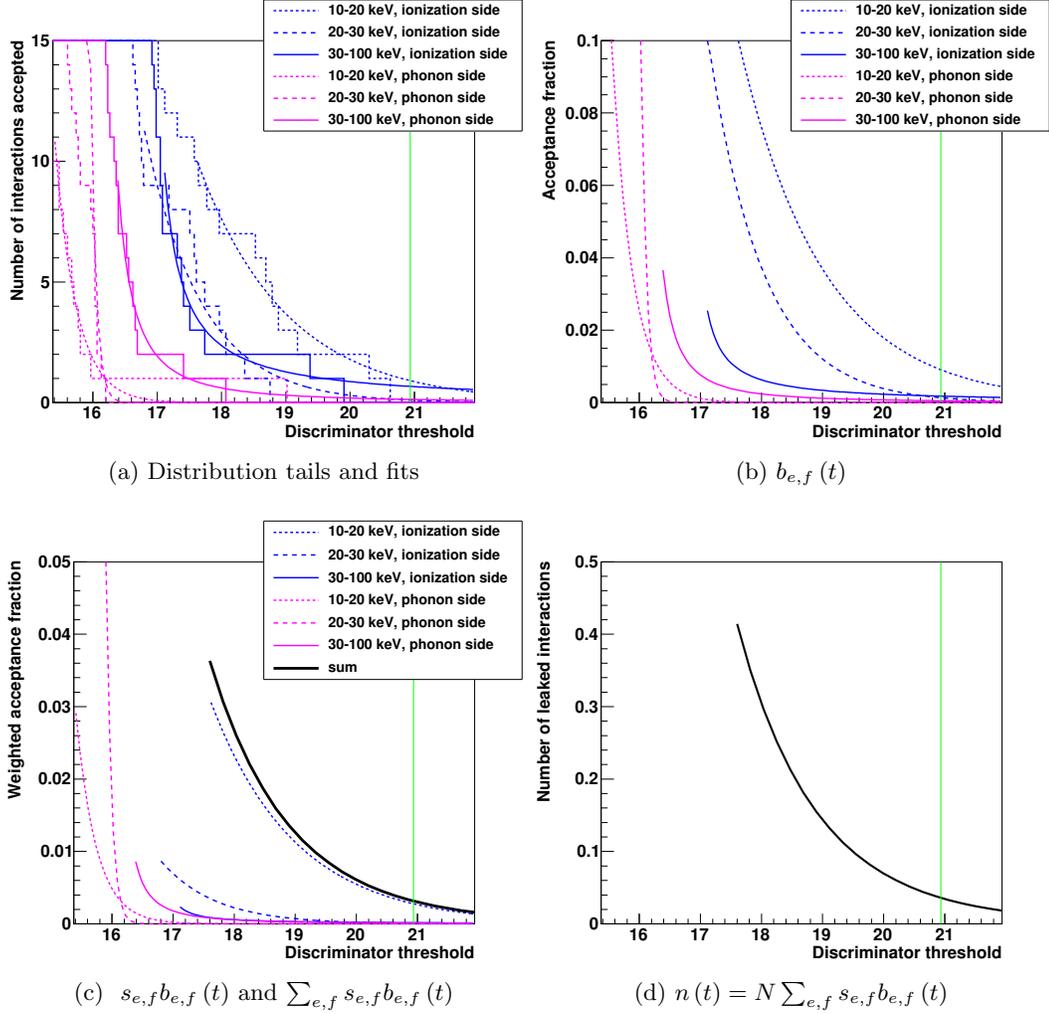


Figure 6.9: Steps in calculating the estimated SER leakage of one detector as a function of discriminator threshold, labeled according to equation 6.2.

- Fits to tails of cumulative distributions of calibration SERs in six categories.
- Fraction of SERs above discriminator threshold in each category.
- Leakage fractions weighted by fraction of SERs in each category, and their sum.
- Estimated SER leakage for Runs 125–8.

The vertical green line indicates the threshold set by the final optimization. The discriminator is a sum of CF-walk-derived rise time and delay and the detector depicted is T1Z2.

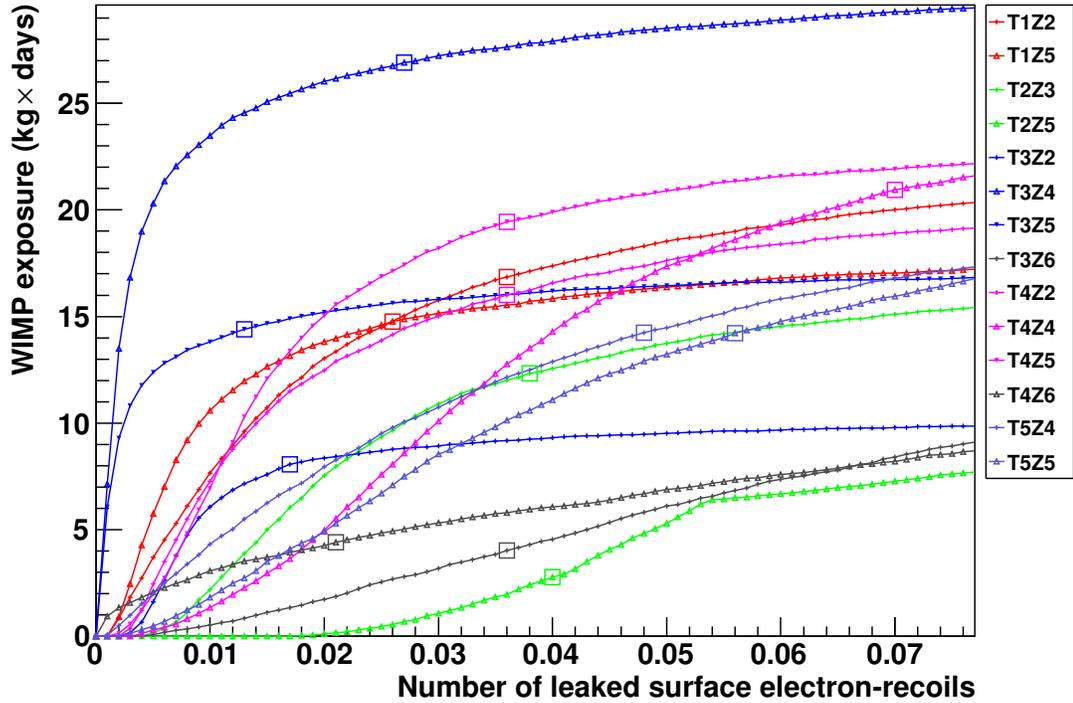


Figure 6.10: Exposure vs. leakage curves for each WIMP-search detector. The final leakage settings as a result of optimization are indicated by squares. The discriminator is a sum of CF-walk-derived rise time and delay.

This is the final stable condition of the algorithm, since in this state a marginal increase in one leakage combined with a marginal decrease in another will produce no net change in total exposure.

6.3.4 Choosing a target leakage

The total WIMP-exposure cannot be the only consideration for choosing the target leakage since leakage necessarily increases with exposure. If the leakage is set very low, the exposure will be unacceptably small; if it is set too large, the presence of SERs in the WIMP signal region will weaken any conclusion that can be drawn from the result.

If treated as a simple counting experiment³, one can calculate an upper limit (to a given confidence level) on the WIMP interaction rate based on the number of interactions seen and the total WIMP exposure using the Poisson distribution. Given a certain expected leakage one can also calculate the probability for any number of background events using Poisson statistics. Under the assumption that these are the only interactions in the signal region, the average upper limit value can be calculated. Hence knowledge of the exposure as a function of leakage allows one to determine the average upper limit that can be set as a function of leakage. This is a measure of the sensitivity of the experiment. For Runs 125–8, a typical high-performance discriminator yielded an average upper limit vs. leakage curve with a broad minimum occurring in the range 0.5 to 0.8 leaked interactions. This is illustrated in figure 6.11. Since systematic effects might be expected to lead to somewhat larger final estimates of the leakage, a conservative choice for target leakage of 0.5 interactions was adopted.

6.3.5 Choosing a discriminator

Once a target leakage is chosen, one can compare the effectiveness of any two discriminators by using the maximized exposure as a figure of merit. To make a proper assessment, it is necessary to have some idea of the uncertainty in this calculation.

The statistical error in the the total leakage estimate can be calculated considering the contributions from all the terms in equation 6.2, including the errors from the functional form fits of the SER distribution tails. Monte Carlo simulations of the procedure were performed on randomly generated samples of surface interactions with probability distributions determined by the fits on the SER distribution tails. These simulations produce the same size errors, confirming the appropriateness of the functional forms. For a total leakage of 0.50 interactions the statistical error is 0.06.

The estimate of the total WIMP exposure for a given set of discrimination criteria is much more precise, since the calculation is based on highly populated portions of the BNR discriminator distributions rather than the tails of the distributions. However to the extent that the exposure is used as a figure of merit to judge the effectiveness

³The actual technique used for setting an upper limit on the WIMP-nucleus interaction rate, the Optimum Interval Method [34], is more sophisticated, taking into account the recoil energy of each interaction in the signal region and the expected WIMP spectrum.

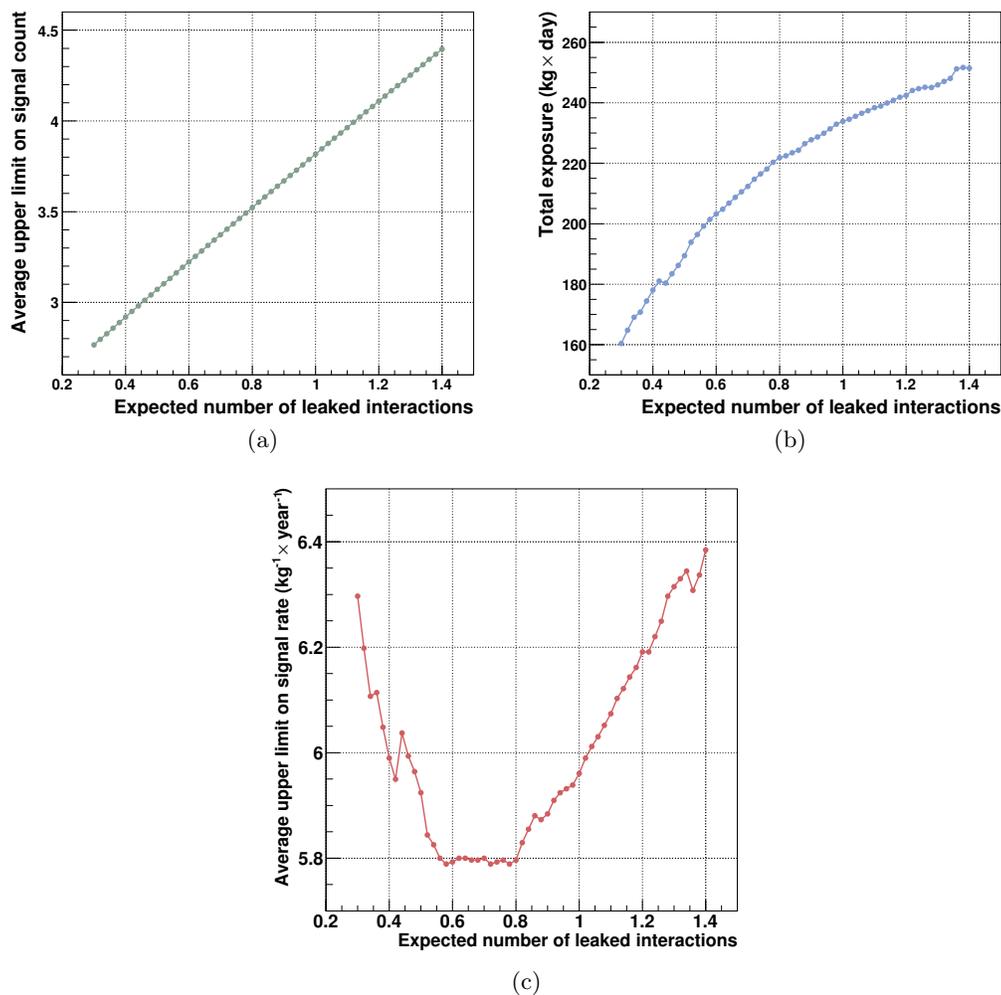


Figure 6.11: Determining a target leakage to optimize experimental sensitivity. The discriminator used is a sum of CF-walk-derived rise time and delay.

(a) Average upper limit (at 90% confidence level) that can be set on the signal count for an experiment with a given expected background count (in the absence of any actual signal) based on Poisson statistics.

(b) Total exposure in kg \cdot days for Runs 125–8 as a function of target leakage after maximizing exposure on a ZIP-by-ZIP basis.

(c) Average upper limit (90% CL) that can be set on the signal rate vs. target leakage, as calculated by dividing the values in the first plot by those in the second.

of the discrimination, it is appropriate to assign an error to it in accordance with the error in the leakage. From an examination of the total exposure vs. total leakage curve, it is determined that for a leakage of 0.50 interactions the corresponding error in the exposure is approximately 4%. If two different discriminators produce exposures that differ by less than this error, the difference cannot be considered significant.

In the previous analysis of Runs 123–4, the discriminator combination used was the sum of the 10-40% rise time and the delay to 20% as derived from the CF-walk algorithm. For Runs 125–8 it was decided that this combination would be used for the final discrimination unless a significant improvement was found with a different combination. This default combination is certainly the most well-studied, and is based on a phonon pulse shape algorithm that is arguably the simplest. I will refer to this combination as the **default discriminator combination** and to the individual discriminators as the **default discriminators**.

Combinations of discriminators from a list of 17 were compared using maximized exposure as a figure of merit. The target leakage was set to 0.5 events. A comprehensive scan over 19 different relative weightings was performed for all pairs of discriminators in the list using the ROOT-code version of the optimization procedure. More than 5000 linear combinations were tested. The relative weightings were expressed as angles from 0 to 90 degrees in 5-degree steps.

The best linear combinations included one rise time and one delay. They tended to have roughly equal weightings between the two discriminators, within a factor of two. When the rise time was based on the 10% to 40% interval, the optimal result tended to assign a smaller weight to the delay (which was based on either the 10% time or the 20% time). In cases where the rise time was based on the 0% to 20% interval the rise time was more likely to be assigned the smaller weight.

Because the weighting angles included 0 and 90 degrees, the linear combination scan implicitly included single-discriminator cases, none of which performed as well as the two-discriminator cases. An exploration of three-discriminator linear combinations was also performed. This was done with the same code used to test two-discriminator combinations with only minor modifications. Instead of using two discriminators from the list as input, one discriminator from the list and one two-discriminator combination were used as input. Three representative powerful two-discriminator combinations

were combined with each of the other discriminators in the list. However, no three-discriminator combination showed an improvement in exposure greater than 2%, so that the improvements cannot be considered significant.

For all the cases mentioned thus far, the same linear combination was used for every WIMP-search ZIP. However, detectors vary enough in their properties that it is likely that the most powerful linear combination differs with detector. This possibility was explored in several ways. First, the weighting angle was varied on a detector-by-detector basis using the default discriminators. The optimal weighting angles varied widely, from 0 degrees for T2Z5 (meaning the discrimination was based entirely on the rise time) to 70 degrees for T5Z4. The improvement in exposure compared to the best globally-applied combination was 5%. In another exploration, the discriminator combinations were varied by detector, chosen from a list of discriminator pairs and using the weightings that had been found to work best when applied globally. This resulted in an improvement of 4% over the the best globally-applied combination. Finally, an attempt to find the optimal discriminator pair *and* weighting for each detector produced an improvement of 8% over the best globally-applied combination. In general, however, the more complicated the set of discrimination criteria, the more likely that it is “tuned” to the specific details of the tails of the calibration SER distributions. In such cases the leakage as determined from the calibration sample will be significantly underestimated.

Table 6.5 is a list of six sets of linear combinations that were considered as the top candidates. The first four use the same linear combination for each detector. Each of these includes one 10-40% rise time and one delay to 20%. In set #5 the default discriminators are used but the weightings are varied by ZIP. In set #6 the discriminators and the weightings are varied. The final column in the table shows the total exposure for each set. All were developed with a target total leakage of 0.5 interactions.

The sets that vary the discriminator by detector (#5 and #6) were ruled out for use as the final SER discrimination. The increase in exposure was not considered large enough to outweigh their disadvantages, including the possibility that they were over-tuned to the sample. Since the first four criteria sets in the list produced very similar total exposures, the final choice was to use the default discriminator combination (#1) for SER discrimination.

The exposure using this discrimination is 194.1 kg · days. This differs from that

#	Discriminator combination	Algorithm	Total exposure in $\text{kg} \times \text{days}$
1	rise time + delay	CF-walk	188.5
2	rise time + $0.75 \cdot \text{delay}$	CF-walk	190.3
3	rise time + delay	VF-walk	191.8
4	rise time + delay	PipeFitter	194.6
5	$W_{\text{RT}} \cdot \text{rise time} + W_{\text{D}} \cdot \text{delay}$	CF-walk	198.2
6	$W_{\text{RT}} \cdot \text{rise time} + W_{\text{D}} \cdot \text{delay}$	various	203.9

Table 6.5: List of SER discriminator sets

shown in table 6.5 because it was calculated using a second version of the procedure for estimating leakage and optimizing for exposure. This version, written in the MATLAB programming language, makes more precise calculations at several stages, including the calculation of the WIMP exposure. Individual thresholds produced by the MATLAB version can differ significantly from those produced by the ROOT version due to details in the fitting algorithms. However differences in exposures calculated by the two versions are small (approximately 2%) and uniform, such that comparisons between discriminators agree very closely.

6.4 Final background estimate

After the final discrimination criterion was chosen, a leakage estimate for Runs 125–8 was calculated using three independent samples of SERs. Each sample consists of interactions with yields below the bulk electron recoil band and which are not coincident with cosmic rays. They differ by the data set (WIMP-search or ^{133}Ba calibration), whether the yields fall in the nuclear recoil band, and whether interactions with ZIP coincidence are considered. The samples are as follows:

1. WIMP-search data in the nuclear recoil band with ZIP coincidence
2. WIMP-search data not in the nuclear recoil band
3. all ^{133}Ba SER data.

The first two samples comprise all WIMP-search SER data that was not embargoed by the blinding procedure. The blinded data are WIMP-search interactions in the nuclear recoil band without ZIP coincidence. This includes SERs and possibly WIMPs (or leaked neutrons). Since these SERs are the target of discrimination I will refer to them as *target SERs*. The SER samples are listed in order of increasing size and decreasing similarity to the target SERs.

Sample 1 is only about four times as large as the expected number of target SERs. However it is expected to be very similar to the target SERs in terms of energy and face distributions. The discrimination criterion can be applied directly to this sample to produce a leakage fraction for each ZIP. Sample 2 has face and energy distributions that differ significantly from those of the target SERs. To account for these differences the leakage estimate is calculated in the same way as it is for the calibration sample as detailed in section 6.3.2 using equation 6.2. Sample 3 is the one used to develop the discrimination criterion. The leakage estimate for this sample is calculated as part of the optimization algorithm.

Of the 14 WIMP-search ZIPs, 12 are situated at intermediate positions in their towers, so that there are adjacent ZIPs above and below them. The other two are at the bottom of their towers. For these bottom ZIPs, Sample 1 is not a good representation of the target SERs. By requiring ZIP coincidence, almost all ionization-side SERs are rejected from this sample. Most of these rejected SERs become target SERs. Thus the

face ratios of Sample 1 SERs and target SERs are very different. For this reason Sample 1 is only used to estimate leakage in intermediate-position detectors.

The expected numbers of target SERs as listed in table 6.4 were calculated using the rates observed in Runs 123–4. A semi-independent estimate can be obtained using the ratio of non-coincident to coincident nuclear-recoil-band SERs observed in Runs 123–4 and multiplying this by the number of coincident nuclear-recoil-band SERs in Runs 125–8. The two estimates were found to be consistent.

The first attempt to estimate leakages using Samples 1 and 2 did not use the three energy categories referred to in table 6.2. Only two energy categories (10–30 keV and 30–100 keV) were used to account for the energy distribution differences. Evidence indicated that this binning was insufficient to account for said differences. For intermediate-position detectors the leakage estimates using Samples 1 and 2 were 1.4 ± 0.5 and 0.3 ± 0.1 , respectively. The large discrepancy, along with an examination of the energy distributions, motivated the change to the three-energy-category scheme. The estimates from the three samples are therefore not completely independent. This was included in the systematic error calculation for the final estimate.

When the final energy binning scheme was applied, the leakage estimates for intermediate-position detectors using samples 1, 2, and 3 were all consistent, at 0.51 ± 0.27 , 0.29 ± 0.17 , and 0.45 ± 0.09 , respectively. Estimates based on samples 2 and 3 including all 14 ZIPs were also consistent, at 0.82 ± 0.59 and 0.51 ± 0.14 , respectively. A combination of estimates from all samples produced a total leakage estimate of 0.52 ± 0.12 . All errors stated here are statistical.

Studies performed after the embargoed data set for Runs 125–8 was unblinded revealed that a further adjustment to the leakage estimate was necessary. Calculating the time of the ionization signal is important for discrimination, since delay discriminators are defined as the time interval between the ionization pulse and the phonon pulse. In general the time of ionization pulses can be calculated very precisely since they rise much faster than phonon pulses. However for recoil energies near the lower threshold the pulses can be quite small. This is particularly true for SERs, since they have low ionization yields. For these small pulses noise becomes important, and the ionization pulse time can be miscalculated by a few μs . Since a much larger fraction of target SERs have

energies near the threshold than SERs in Samples 1 and 2 (see figure 6.2), the systematic error between leakage estimates is enhanced. A final leakage calculation accounted for this effect, employing a Bayesian estimation scheme as described in [32]. The final surface interaction leakage estimate for Runs 125–8 was $0.82^{+0.12}_{-0.10}(\text{stat})^{+0.20}_{-0.19}(\text{sys})$.

6.5 Additional comments

The description of the analytical efforts to produce SER discrimination described in the preceding sections do not do justice to their full complexity. In this final section I provide a few additional details of several aspects of the analysis, including useful new knowledge concerning discrimination.

6.5.1 Lessons learned from discrimination studies

A second functional form fit referred to as *WedgeFit* was applied to data from Runs 125–8. It was designed to extract information from the earliest part of the pulse, a region not directly sampled by either the walk algorithms or PipeFitter. The functional form is piecewise-defined as a constant (to fit the pre-pulse baseline) up to a time t_0 , and a quadratic at times greater than t_0 :

$$f(t_{\text{ADC}}) = \begin{cases} 0 & \text{if } t_{\text{ADC}} \leq t_0, \\ a(t_{\text{ADC}}^2 - t_0^2) + b(t_{\text{ADC}} - t_0) & \text{if } t_{\text{ADC}} > t_0. \end{cases} \quad (6.5)$$

The functional form has only three free parameters: t_0 , a , and b , and is fit over a range beginning before the onset of the pulse and ending somewhat after the 20% time.

WedgeFit is the only algorithm that provides a reliable measure of the the time that the pulse first begins to rise above the baseline. A delay discriminator measured to this 0% time incorporates no information about the rise time of the pulse. It was interesting to find that this quantity has virtually no discriminatory power (see figure 6.12). Apparently some phonons arrive promptly regardless of interaction type or location. This indicates that the discriminatory power of 20%-time delay quantities actually lies in the the shape of the pulse between the 0% and 20% times.

There is evidence that a small fraction of the primary phonons from a particle interaction in the detector are ballistic (and therefore fast), and that this effect is actually enhanced in nuclear recoils [35]. If so, it is reasonable to expect that the earliest phonons from a BNR arrives just as promptly as those from an SER. Only information from later in the pulse will provide discrimination, since the majority of primary phonons from the BNR will propagate slowly. This explains the observations illustrated in figure 6.12.

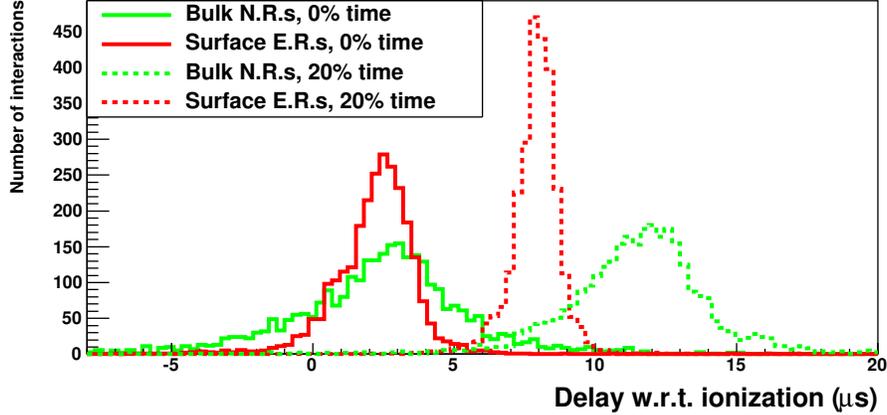


Figure 6.12: Time intervals (“delays”) between the onset of the ionization signal and the primary phonon pulse, for BNRs and SERs. The delays shown are to the 0% time and 20% time of the primary phonon pulse. The former is characterized by the WedgeFit algorithm and the latter by the CF-walk algorithm. Note that the ability to discriminate the recoil types is not present at the pulse onset.

Recoil energy range	PipeFitter	WedgeFit
5–25 keV	95.0%	80.0%
25–55 keV	98.2%	90.5%
55–125 keV	96.2%	96.9%

Table 6.6

For both the WedgeFit and PipeFitter algorithms, rejection criteria were developed based on failure of the χ^2 -minimization and outlying χ^2 values for the fit. Table 6.6 lists the BNR selection efficiencies of these criteria in three different recoil energy ranges. For both fit algorithms the efficiency is lowest at low energies, where noise can make a trace difficult to fit. For PipeFitter it is also somewhat smaller at high energies, where pulse shapes begin to deviate due to nonlinearities in phonon sensor response.

WedgeFit-derived rise time discriminators were explored for 0%-to-10% and 0%-to-20% intervals. These may have competed better with other discriminators were it not for the fact that fit quality criteria reduced their efficiencies. Much less time was spent on WedgeFit development than on PipeFitter development. It is entirely possible that the algorithm could be modified to make it more stable, thus improving its efficiency.

Another phonon pulse shape characterization tool used in Runs 125–8 is an algorithm to find the maximum and minimum of the differential pulse. The differential pulse is constructed by calculating the difference in pulse amplitude from one ADC bin to the next. It is smoothed by averaging each bin with its four nearest neighbors, then the maximum and minimum of this smoothed pulse are located. Quantities representing the magnitude and time of the most quickly changing parts of the pulse on the rising edge and falling edge are thus obtained.

The only discriminators deriving from the maximum/minimum differential algorithm that were studied in detail were time intervals calculated between the 0% time and the time of the maximum pulse differential. The potential of this algorithm could have been studied further. The amplitude of the maximum differential, appropriately normalized, could serve as a sort of inverse rise time, and might have produced unique discriminatory power.

While a “maximum differential” discriminator was not explored, some discriminators were studied which do not fit into either “rise time” or “delay” categories. These include derivatives of the pulse form calculated at various points on the rising edge using the PipeFitter or WedgeFit functions. These discriminators contain information very similar to rise times. Another type of discriminator that showed some promise was based on the pulse width, defined as the interval between the 80% time on the rising part of the pulse and the 80% time on the falling part. Usually information from the falling edge has no discriminatory power. It is possible that the effectiveness of the pulse width discriminator is related to partial pulse saturation effects: a faster-rising pulse is also a taller pulse and thus more likely to be affected by any nonlinearities in phonon sensor response. In the end we confirmed that the shape of the early part of the primary phonon pulse contains the best information for SER discrimination.

6.5.2 Evaluation of discrimination efforts

The analysis of Runs 125–8 was the first in which discrimination thresholds were set for each detector so as to optimize total WIMP-exposure. In previous analyses thresholds were set to produce equal leakage fractions in each ZIP. The optimization significantly improved the sensitivity of the experiment.

Certainly if detectors were identical in all respects there would be no advantage to

the optimization procedure. However significant differences exist in detector properties, in their positions with respect to other detectors, and in their exposure to contaminants. The increase in total exposure afforded by the optimization procedure is approximately 15% with respect to the previous method. In fact most of the advantage is achieved by assigning the detectors equal *numbers* of expected leaked surface interactions. This is a simpler method of weighting the threshold selection for each detector by its overall contribution to the experimental sensitivity. In principle the optimization procedure can entirely remove a detector from the analysis, if its contribution to leakage outweighs its contribution to WIMP exposure. This sometimes occurred during discrimination-criteria development in cases when the target total leakage was set below 0.5 interactions.

The criteria that vary the discriminator by detector (#5 and #6 in table 6.5) represent more extreme attempts to optimize WIMP exposure by accounting for differences between detectors. Details of these criteria are provided in table 6.7. The possibility that these criteria were over-tuned to the calibration sample, such that their leakages were underestimated during the optimization procedure, can be tested by calculating the leakages on independent samples. Table 6.8 shows the ratio of the leakage estimate calculated using an independent sample to the leakage estimate using the calibration sample, for the complex discrimination criteria #5 and #6, with the default discrimination criterion #1 included for comparison. There is a suggestion that indeed the independent leakage estimates are somewhat higher for the complex criteria, although in all cases the leakage estimates agree within the statistical errors. Apart from concerns about underestimating the leakage, these complex discrimination criteria are somewhat unsatisfying since it is not understood why different discriminators work better with different detectors. A thorough study of how phonon pulse shapes vary by detector, particularly in the case of surface interactions, might have better motivated such complex discrimination criteria.

Discriminators produced by the PipeFitter algorithm performed at least as well as the default discriminator. The pattern in table 6.7 suggests that PipeFitter discriminators perform better than walk-algorithm discriminators for detectors in Towers 1–2. Due to differences in fabrication the detectors in Towers 1–2 have on average three times greater phonon pulse amplitudes at a given recoil energy than do detectors in Towers

ZIP	DC #5		DC #6			
	W_{RT}	W_D	W_{RT}	rise time	W_D	delay
T1Z2	1	0.577	1	PF	0.466	PF
T1Z5	0.7	1	1	PF	0.268	PF
T2Z3	0.839	1	1	PF	0.176	PF(10)
T2Z5	1	0	1	CF-walk	0.364	PF
T3Z2	0.466	1	0.364	VF-walk	1	VF-walk
T3Z4	1	0.268	1	CF-walk	0.268	CF-walk
T3Z5	0.7	1	1	CF-walk	0.466	PF
T3Z6	1	0.268	1	CF-walk	0.268	CF-walk
T4Z2	0.7	1	0.577	CF-walk	1	VF-walk
T4Z4	1	0.268	1	CF-walk	0.466	VF-walk
T4Z5	1	0.466	1	CF-walk	0.466	VF-walk
T4Z6	0.839	1	0.839	CF-walk	1	CF-walk
T5Z4	0.364	1	0.268	PF	1	PF
T5Z5	0.839	1	0.466	VF-walk	1	VF-walk

Table 6.7: Weighting factors used for discriminator combination set #5 and weighting factors and algorithms used for discrimination combination set #6 as listed in table 6.5. In DC #5 the default rise time and delay discriminators are used for each ZIP. In DC #6 a 10-40% risetime and a delay to 20% are used for each ZIP except for T2Z3, for which a delay to 10% is used.

	Set #1	Set #5	Set #6
Int. detectors, calibration sample	0.45 ± 0.09	0.45 ± 0.11	0.46 ± 0.08
Int. detectors, sample 1	0.51 ± 0.27	0.85 ± 0.37	0.66 ± 0.31
Ratio	1.1 ± 0.7	1.9 ± 0.9	1.4 ± 0.7
All detectors, calibration sample	0.51 ± 0.14	0.50 ± 0.14	0.50 ± 0.15
All detectors, sample 2	0.82 ± 0.59	0.89 ± 0.59	1.00 ± 0.61
Ratio	1.6 ± 1.2	1.8 ± 1.3	2.0 ± 1.4

Table 6.8: Leakage estimates for various sets of discrimination criteria calculated from calibration data samples and independent samples from WIMP-search data sets. Also shown are the ratios of the independent estimates to the calibration estimates. The discrimination criteria set #1 (the final set) and sets #5 and #6 (the sets using different discriminator combinations for each detector) are described in table 6.5. The data samples are described in section 6.4.

3–5, with correspondingly better signal-to-noise ratios. Evidently the PipeFitter functional form fit produces better discrimination for such detectors.

Because PipeFitter was originally written to also provide a reconstruction of recoil energy, it is not fully optimized for the purpose of SER discrimination. A more limited fit range which better emphasized the early part of the pulse may improve the discrimination power. It is also possible that the functional form could be simplified.

In the analysis of Runs 125–8 we have used all the surface interaction data at our disposal to calculate a leakage estimate. In the effort to reduce systematic errors, statistical errors have increased. It is probable that we have reached the limit of how precisely the leakage can be estimated. Even if the calibration data sample were greatly increased in size, the small numbers of SERs in the WIMP-search data represent a significant limiting factor, since these numbers are necessary to characterize the rates of various categories of SERs that could constitute background.

In CDMS II the limiting factor for WIMP-selection efficiency is the discrimination of SER interactions. One strategy to mitigate this loss of efficiency is to reduce the contamination that produces such interactions; however greater efforts than those now implemented are impractical. A more important strategy that is being implemented for future phases of CDMS is to use new detectors designed for the ability to identify surface interactions much more definitively, thus yielding highly efficient discrimination.

Chapter 7

Results of the CDMS II Final Analysis

In the CDMS II analysis procedure, the embargoed WIMP-search data set is unblinded after all WIMP-selection criteria are finalized, including the discrimination criteria described in the previous chapters. The key elements for interpreting the results are the total estimated leakage and the total WIMP exposure.

7.1 Unblinding the data

7.1.1 Preliminary expectations

The result of unblinding is a set of interactions satisfying all WIMP-selection criteria (WIMP candidates). The number of WIMP candidates must be significantly greater than the expected background to be considered evidence of a WIMP signal.

Table 7.1 lists the background estimates from the previous four chapters. The total background estimate for Runs 125–8 is 0.9 ± 0.2 . It is dominated by the SER background, which is ten times greater than the combined neutron background.

Given 0.9 expected background interactions, the Poisson distribution can be used to judge how many WIMP candidates constitute evidence of a signal. If one requires only 90% confidence, three candidates would be enough, since the probability of three or more background events is only 6%. Four candidates would constitute a stronger hint,

Category	Expected Background
Radiogenic neutrons	0.04 +0.00/−0.02 (stat)
Cosmogenic neutrons	0.04 +0.04/−0.02 (stat)
Bulk electron recoils	$< 5 \times 10^{-4}$
Surface electron recoils	0.82 +0.12/−0.10 (stat) +0.20/−0.19 (sys)
Total	0.9 ± 0.2

Table 7.1

since the probability of a background fluctuation then drops to 1.3%; at five candidates the confidence level exceeds 99%, and so on. These probabilities can be converted to the equivalent for a Gaussian distribution. In these terms, the significance is greater than “1 σ ”, “2 σ ”, “3 σ ”, or “4 σ ” for three, four, six, or eight WIMP candidates, respectively.

In the analysis of Runs 123–4, zero WIMP candidates were found. From this result an upper limit was set on the WIMP interaction rate. The WIMP exposure for Runs 125–8 is only 60% greater than that of Runs 123–4, so if a large number of WIMP candidates were seen in Runs 125–8 it would tend to contradict the null result of the earlier analysis. We can pose the question: what is the probability of seeing zero WIMP candidates in Runs 123–4 given a rate estimated using the results of Runs 125–8? If three candidates were seen in Runs 125–8 this probability would be 27%. With five candidates it falls to 8%; at eight candidates it falls to 1.2%.

Thus there is a rather narrow range of possible outcomes of the unblinding that would be considered evidence of a signal while not being inconsistent with earlier results.

7.1.2 Events satisfying WIMP selection criteria

The embargoed data set from Runs 125–8 was unblinded on 5 November 2009. Two WIMP candidates were revealed. Since this result is consistent with the expected background, it does not represent evidence of a WIMP signal.

The recoil energies of the WIMP candidate interactions were 12.3 keV and 15.5 keV. All characteristics of the WIMP candidate events are in accord with WIMP selection criteria: no ZIP coincidence, no muon coincidence, no pathological detector behavior, no unusual run conditions, etc. The candidates were scrutinized to confirm that the selection criteria worked as intended. During this period of study it was found that

the algorithm for determining the time of the ionization signal was not ideal¹, which prompted the adjustment to the SER background estimate.

In the standard CDMS analysis no assumption is made about the identity of an individual WIMP candidate — it could be background (neutron or SER) or a WIMP. Given the lack of evidence for a WIMP signal, the most reasonable guess for either candidate is that it is an SER. In figure 7.1 the WIMP candidates are shown in the context of SER discrimination for their corresponding detectors. Although each candidate has appropriate characteristics for a BNR (as it must), each also appears as though it could be an outlier of the SER population. If either candidate were a BNR it certainly could have appeared in the wide region of the BNR distribution that lies far from the SER distribution.

The alternative SER discrimination criteria (#2–#6 in table 6.5) were applied to the WIMP search data of Runs 125–8. If the discrimination criterion had been #3, #5, or #6 instead of #1, there would have been only one WIMP candidate. The other would have failed SER discrimination. The candidate that would have failed was seen in ZIP detector T3Z4. For this detector the same discriminator is used in #5 and #6 (see table 6.7), which uses the default discriminators but with a much weaker weight assigned to the delay discriminator. This candidate would also have been rejected by the default discriminator if a corrected ionization time reconstruction had been used. With one fewer WIMP candidate, the Poisson upper limit on the rate would have been 37% lower.

The effect of different target leakages for developing the SER discrimination with the default discriminator combination was also explored. The target leakage used was 0.5. If it had been 0.32, only one WIMP candidate would have been seen; at 0.30 or lower, there would have been zero candidates. If the target leakage had been 0.78, a third candidate would have been revealed. Figure 7.2 shows what the Poisson upper limit on the WIMP interaction rate would have been for SER target leakages from 0.3 to 1.4. This figure can be compared to the predicted sensitivity plot in figure 6.11. Note however that the final leakage estimate is greater than the target leakage due to the

¹One candidate would have been rejected by the SER discrimination if this algorithm were adjusted to better characterize the time of the ionization signal. However a proper interpretation of this would require a complete reprocessing and reanalysis of the data. Such a reanalysis is currently being pursued, but it will of course not be considered a blind analysis.

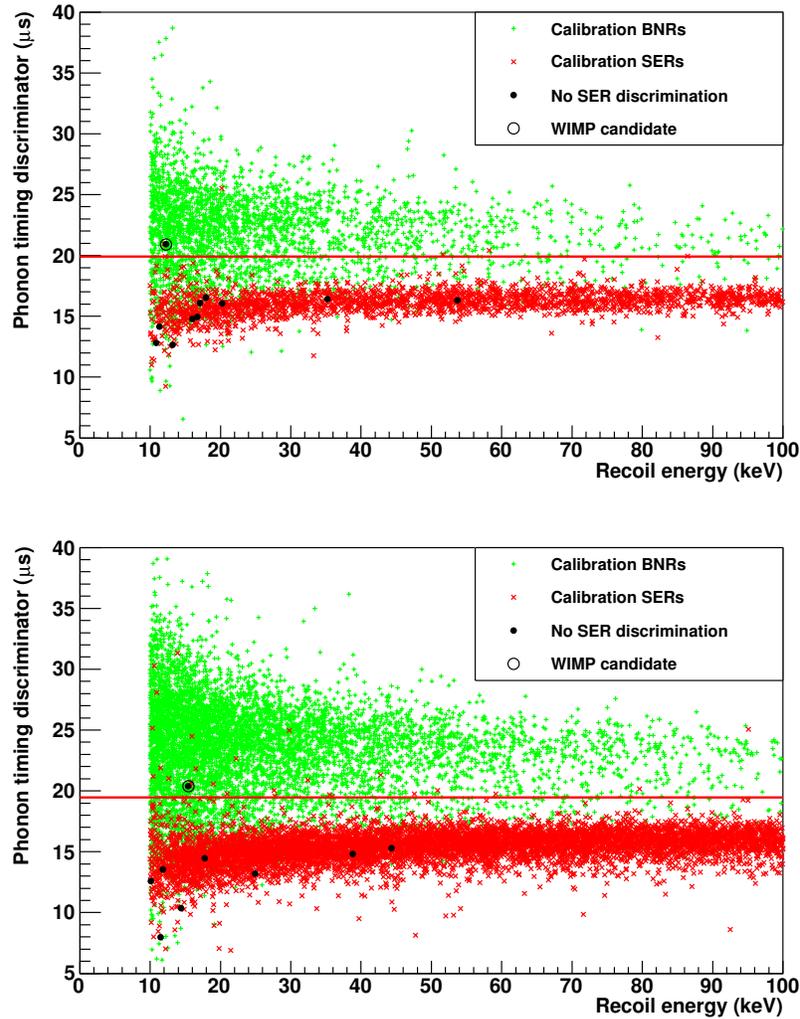


Figure 7.1: SER discriminator values of the two WIMP candidates vs. recoil energy plotted along with the calibration SER and BNR populations for the corresponding ZIP detectors (top T1Z5, bottom T3Z4). Also shown are the interactions that satisfied all WIMP selection criteria *except* SER discrimination. The horizontal line indicates the threshold of the SER discrimination for the detector.

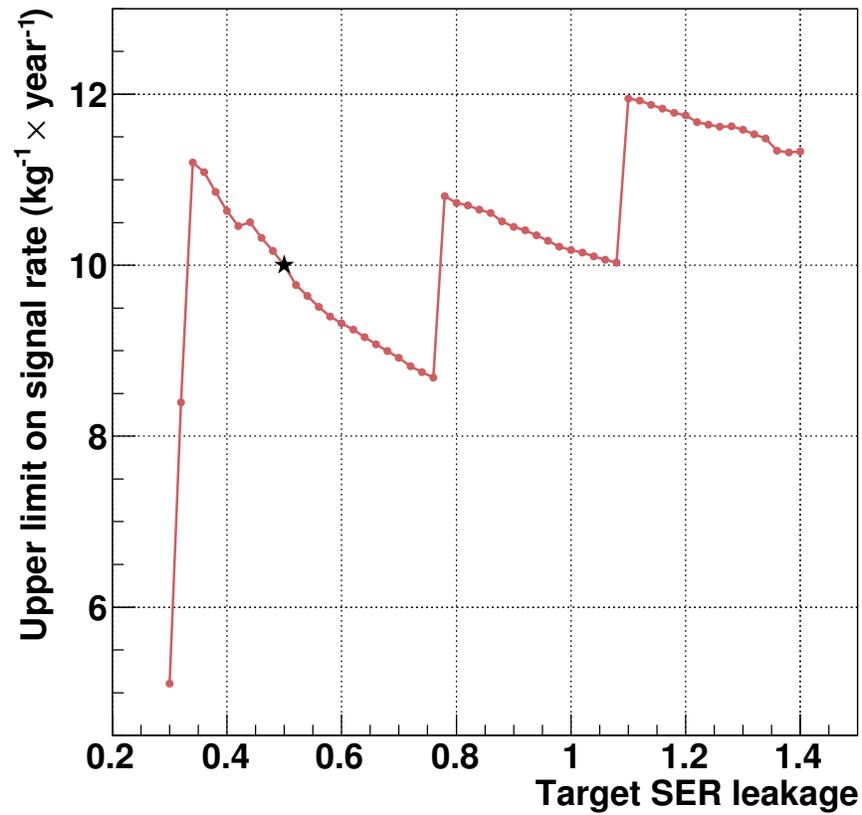


Figure 7.2: Upper limit that could have been set on the WIMP interaction rate with the Runs 125–8 data for different SER target leakages using the default discriminator. Sharp increases occur each time an additional WIMP candidate appears in the data. The black star indicates the target leakage actually used. Limits are calculate for a WIMP mass of $60 \text{ GeV}/c^2$ using Poisson statistics.

correction concerning the ionization signal time.

The above scenarios are of course hypothetical: the SER discrimination criterion was fixed before the unblinding. If one were to perform a non-blind reanalysis, resetting the SER discrimination criterion based on this a posteriori knowledge, a statistical penalty would be necessary to properly calculate the upper limit.

7.2 Interpreting the results

The exposure of Runs 125–8 is 194.1 kg·days for a WIMP mass of $60 \text{ GeV}/c^2$. Treated as a simple counting experiment, we can use this result to set an upper limit of 10.0 interactions per kg·year on the WIMP interaction rate in germanium for recoil energies between 10 and 100 keV. Combined with the Runs 123–4 result (zero candidates, 121 kg·days exposure) the upper limit is 6.2 interactions per kg·year.

This calculation depends on the WIMP mass because the expected WIMP recoil energy spectrum depends on the mass, and the exposure depends on the energy-weighted average efficiency. An upper limit on the spin-independent (SI) WIMP-nucleon cross section can also be calculated using the considerations discussed in section 1.3. However the CDMS analysis procedure uses a more sophisticated method than Poisson statistics to calculate this limit, a method that makes use of the recoil energies of the WIMP candidates. The choice of statistical technique is made before unblinding as part of the CDMS blind analysis policy.

7.2.1 Upper limit of WIMP-nucleon cross section

The technique used for setting an upper limit on the WIMP-nucleon cross section is called the Optimum Interval Method [34]. This method is designed to set an upper limit on the rate of a process with a known spectral shape, in the presence of background events that may have a different spectrum. The principle is to consider all energy intervals between pairs of candidates and make a statistical judgement of how likely it would be, for a given rate, to see as few events in that interval as are observed. The algorithm finds the interval from which the strongest limit can be inferred. In this way a stronger upper limit can be set than one based on Poisson statistics if the background spectrum differs significantly from the expected WIMP spectrum. In the case of zero

WIMP candidates, the method is equivalent to a Poisson statistical treatment. The method requires no knowledge concerning the background. An upper limit can be calculated regardless of whether there is evidence of a signal.

The upper limit can be calculated for every possible WIMP mass. Figure 7.3 shows the results of this for Runs 125–8, along with the combined results of previous runs, and the combined result of all CDMS II runs. The sharp features in the Runs 125–8 curve are caused by shifts in the optimum interval chosen by the algorithm. This interval depends on the expected WIMP spectrum which in turn depends on the WIMP mass.

It is noteworthy that the combined limit including Runs 125–8 is not as strong as the combined limit not including Runs 125–8 for a certain range of WIMP masses. This is a consequence of the relatively unlucky occurrence of two candidate interactions. The two WIMP candidates have recoil energies near the lower end of the analysis range, which makes it more difficult to rule out a light mass WIMP. The lack of high energy candidates allows for a strong limit on higher mass WIMPs.

Also shown in Figure 7.3 is an upper limit curve produced by another direct detection experiment known as XENON100 [36]. XENON100 is similar to CDMS in that it employs a two-channel strategy to discriminate electron recoils from nuclear recoils. However the target is xenon, which has a peak sensitivity to a lower WIMP mass than does germanium. Thus the two experiments are complementary. CDMS II currently leads the world in upper limits for WIMP masses greater than approximately $85 \text{ GeV}/c^2$, with XENON100 setting the strongest limits for masses in the $10\text{--}85 \text{ GeV}/c^2$ range.

Figure 7.4 shows the CDMS II and XENON100 upper limit curves superimposed on the theoretical predictions shown in figure 1.1 in chapter 1. Direct detection experiments have begun to constrain these theories. However an increase in sensitivity of more than two orders of magnitude will be necessary to probe the full parameter space depicted in this figure. CDMS along with other collaborations is planning future phases with greatly increased total detector mass. If backgrounds continue to be managed as effectively as they were for CDMS II, the sensitivity will increase proportionately.

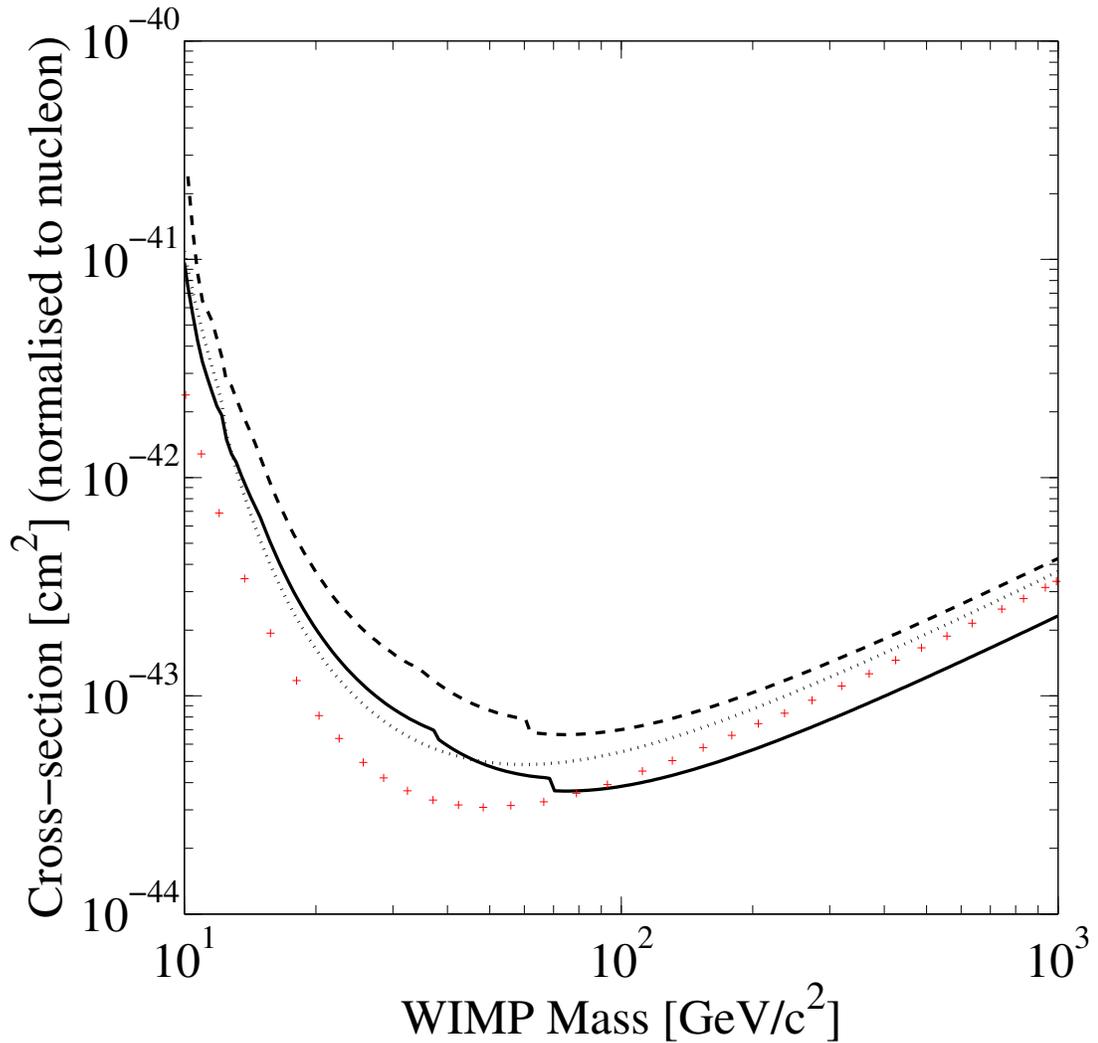


Figure 7.3: Upper limit on the SI WIMP-nucleon cross section set by the blind analysis results of CDMS II. Dashed black: using only data from Runs 125–8. Dotted black: using combined data from Runs 118–123. Solid black: using all CDMS II data. Also shown (red pluses) are the results of a non-blind analysis from another direct detection experiment, XENON100 [36]. (Figure courtesy [12])

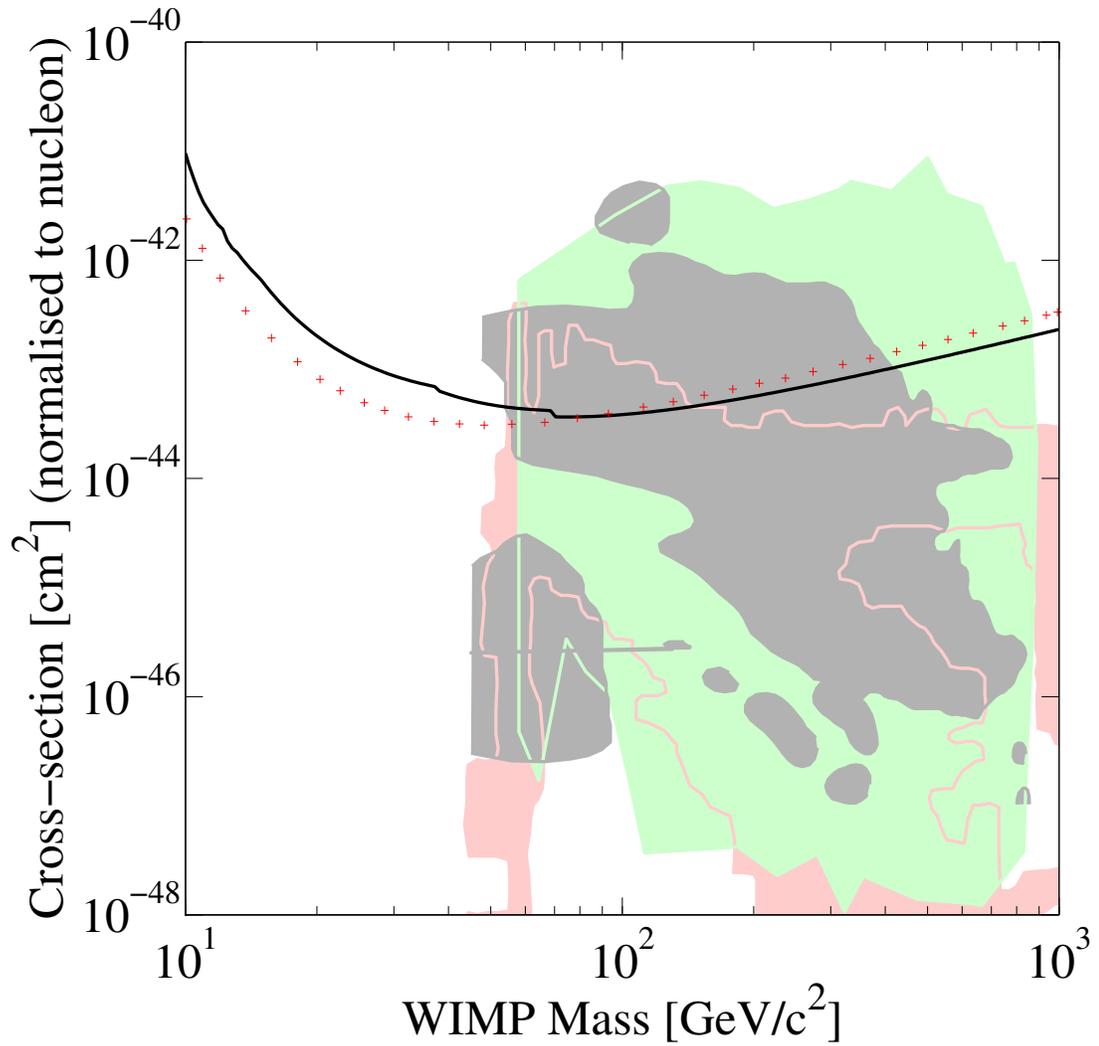


Figure 7.4: CDMS II (solid black) and XENON100 (red pluses) WIMP-nucleon cross section upper limits and theoretical predictions. See figure 1.1 for a description of the supersymmetric theories. (Figure courtesy [12])

Chapter 8

Testing SuperCDMS Detectors

The phase of CDMS that follows CDMS II is known as SuperCDMS. SuperCDMS incorporates detector designs that are upgrades of the ZIPs used in CDMS II, although many features remain: cylindrical germanium crystals with aluminum and tungsten transition-edge phonon sensor arrays and aluminum grid ionization sensors photolithographed onto the faces.

CDMS detectors undergo extensive testing at aboveground facilities before deployment at underground experimental sites. Tests are performed to evaluate design and fabrication changes and to check individual detectors for fabrication errors. These tests range from simple electrical measurements to full characterization of their detection performance.

In 2007 I developed a detector test facility at the University of Minnesota by adapting a dilution refrigerator with which I had previous experience. Nicknamed “Little Blue”, this facility fills an intermediate functionality: it is able to cool detectors below the superconducting transition of the phonon sensors so that measurements can be made at operational temperatures, but it is not designed to incorporate the electronics necessary to run the detectors in detection mode.

8.1 Purpose of cryogenic tests

The tungsten transition-edge sensor arrays of ZIPs are designed to have low transition temperatures to minimize noise. The transition temperatures must also be high enough

above the ambient crystal temperature so that the sensors can relax to their equilibrium states after receiving a signal. Typical critical temperatures of the phonon sensor arrays after photolithography are near 140 mK. In a later fabrication step iron ion implantation lowers these temperatures to a typical target value of approximately 80 mK.

In early ZIPs uneven tungsten deposition resulted in a critical temperature gradient along the Y-axis of the detector surface. This is an undesirable condition, particularly if the gradient is significant within a single array, because it may be impossible to simultaneously bias all of the individual sensors to their superconducting transitions. Such critical temperature gradients can be removed by a graded ion implantation.

This process of tuning the critical temperatures required pre-implantation and post-implantation measurements of the critical temperatures for early ZIPs, which naturally require a refrigerator capable of base temperatures below 100 mK. Improvements in fabrication have resulted in more uniform results, reducing the necessity of such measurements.

Some fabrication errors that can result in a nonfunctional phonon sensor array are only detectable with cryogenic testing. Figure 8.1 shows the basic layout of a transition edge sensor: a tungsten strip with aluminum fins on a semiconductor substrate. A full sensor array includes many such sensors connected in parallel. Figure 8.2 shows the measured resistances of a good sensor array as a function of temperature. Below 1.2 K there is a sharp drop in resistance as the aluminum becomes superconducting. Between 1.0 K and 0.1 K the array resistance is a few tenths of an Ohm. This is the normal resistance of the tungsten array. At a temperature of about 85 mK the tungsten transition appears. Below this temperature a sensitive measurement will show resistances of only a few $\text{m}\Omega$ (the so-called *parasitic resistance*), due mainly to contact resistance of the wirebonds that form the external connection to the aluminum on the detector surface.

An error in photolithography could lead to aluminum shorting a tungsten sensor if it forms a bridge between aluminum fins. If such an error is not apparent by visual inspection, it can only be detected by a cryogenic test. The signature of an aluminum short is a sensor array resistance that drops to a minimum value (less than 0.1Ω) at temperatures below 1.2 K, with no measurable tungsten transition at lower temperatures. Because the sensors are connected in parallel a single shorted sensor will short

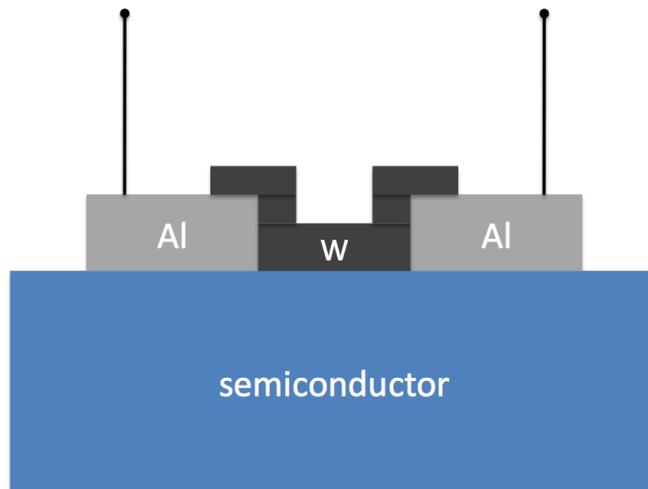


Figure 8.1: Schematic of a transition-edge phonon sensor on the surface of a ZIP. The tungsten sensor and aluminum fins on the substrate are depicted. The two lines indicate the external leads. An actual phonon sensor array is many such sensors connected in parallel.

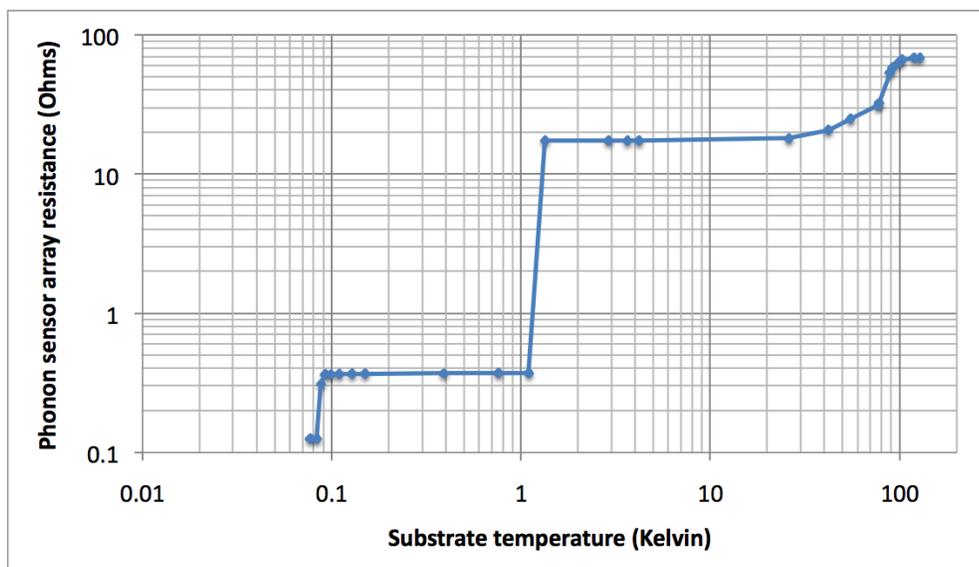


Figure 8.2: Measured resistance of a phonon sensor array as a function of temperature. Notable are the aluminum superconducting transition at 1.2 K and the tungsten superconducting transition at 0.085 K. (At the lowest temperatures the measuring instrument reads its minimum value — the actual resistance is closer to 10 m Ω .)

the entire array.

A number of other cryogenic tests can ensure that a phonon sensor array is healthy. A broken circuit due to a break in an aluminum connector path may not be measurable at room temperature because the semiconductor substrate is sufficiently conductive to complete the connection. Only at low temperatures is the substrate resistance high enough to reveal such an open circuit. The substrate resistance itself must be great enough at operating temperatures to prevent current leakage between arrays. This can be measured using the leads of two different sensor arrays. Unusual values of the normal-state resistance, the parasitic resistance, or the critical current (the minimum current required to drive a sensor out of its superconducting state) are also indicators of possible detector pathologies.

8.2 Design and construction of the testing facility

Several modifications and additions to existing equipment were necessary to construct a facility suitable for testing CDMS ZIPs. This was achieved at low cost by using as much of the existing hardware as possible. The refrigerator and associated equipment are shared with a condensed matter research group, so no permanent modifications were implemented that would hinder those experiments. The efforts to develop the Little Blue facility have been advantageous for all users because procedures have been standardized and codified and because certain infrastructure improvements have made the system more stable and reliable. A complete manual for cryogenic operations is included in this thesis as appendix A.

8.2.1 Preexisting equipment

The refrigerator used in the Little Blue facility is a Janis Research Co., Inc. model CF-25 dilution refrigerator. It has a rated cooling power of $35 \mu\text{W}$ at 100 mK and a measured cooling power of $60 \mu\text{W}$. The lowest base temperature achieved was 42 mK. Figure 8.3 includes photographs of the CF-25 unit. The refrigerator fits inside a 58.5-inch-deep Dewar with a maximum inner diameter of 6.56 inches. During normal operation this Dewar is filled with liquid helium. The Dewar is housed in a steel rack. Attached to the rack is a hanger for the refrigerator for removal and installation of experimental

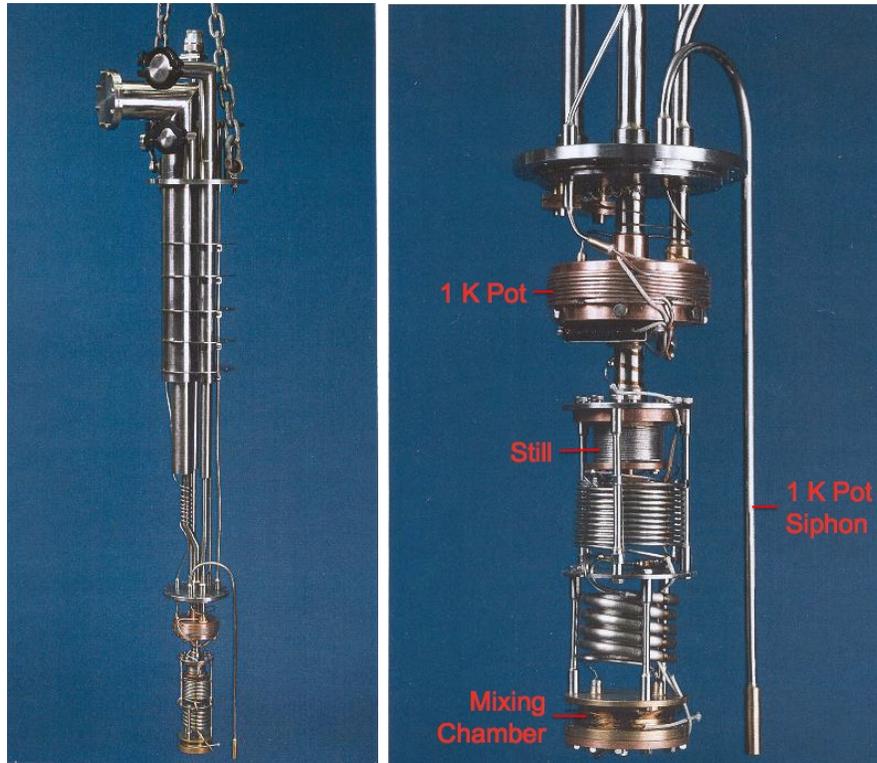


Figure 8.3: Photos of dilution refrigerator. On the left (full view) the pumping ports for the vacuum can, 1 K pot, and mixture circulation can be seen at the top of the unit. On the right (close-up) key refrigerator components are labeled. A good reference for the principles of dilution refrigeration is [37].

hardware. A pulley system allows the fridge to be raised and lowered into the Dewar or onto the hanger.

The refrigerator was originally used for superfluid helium studies [38, 39]. Mark Williams worked on this research as a graduate student and built a gas handling system that allows for storage, cleaning, and circulation of the ^3He - ^4He mixture. Details of this system are described in appendix A.

A vacuum can was built to accommodate the experimental cells used for superfluid helium studies. A 2.1-inch-high copper bracket attaches to the bottom of the mixing chamber. This bracket is used to mount the experimental hardware. A copper tube attaches to the 1 K pot to provide a 1.4 K shield for the refrigerator components and

experimental hardware to block 4.2 K thermal radiation from the vacuum can walls. This shield is open at the bottom and thus does not provide full coverage.

Two bundles of wires extend from the cold stages of the refrigerator up to vacuum feedthrough connectors at the top of the refrigerator. One bundle is original to the refrigerator and connects to three built-in RuO₂ thermometers mounted onto different temperature stages and two resistive heaters mounted to the still and mixing chamber. Twelve additional wires from this bundle are available to attach to experimental hardware. A second bundle of 28 wires arranged in twisted pairs provides for additional connections. The wires in the second bundle are 0.006-gauge Manganin, an alloy often used for cryogenic wiring due to its nearly constant resistivity over a large range of temperatures. The resistance of these wires is 70 Ω each. The wires of the first bundle are also Manganin with a resistance of 90 Ω each. Wires from each bundle are heat-sunk at three temperature stages. This arrangement allows for ten 4-wire connections to the experimental hardware and thermometers.

Two RuO₂ thermometers mounted onto copper plates are available to attach to the experimental hardware. These were used in superfluid helium experiments but their calibration information was lost. Breakout boxes for the refrigerator wiring were built for the specific needs of the superfluid helium studies.

Resistances of thermometers and other experimental components can be measured with a BTi Model 1000 conductance bridge. This instrument has one channel for 4-wire constant-voltage measurements with excitation settings from 10 to 300 μV . The minimum resistance measurable is 0.25 Ω .

8.2.2 New refrigerator hardware

The inner diameter of the vacuum can used for helium experiments ranges between 3.5 and 3.63 inches. The housing for a 7.6 cm ZIP is 3.6 inches wide and thus too large to fit into the preexisting vacuum can. A new vacuum can, shield, and mounting hardware were designed and built to accommodate the refrigerator, Dewar, and ZIP housings. Figure 8.4 is a schematic of the new hardware.

The new stainless steel vacuum can has two pieces. The upper piece is more narrow and connects to the refrigerator flange. The lower piece has an inner diameter

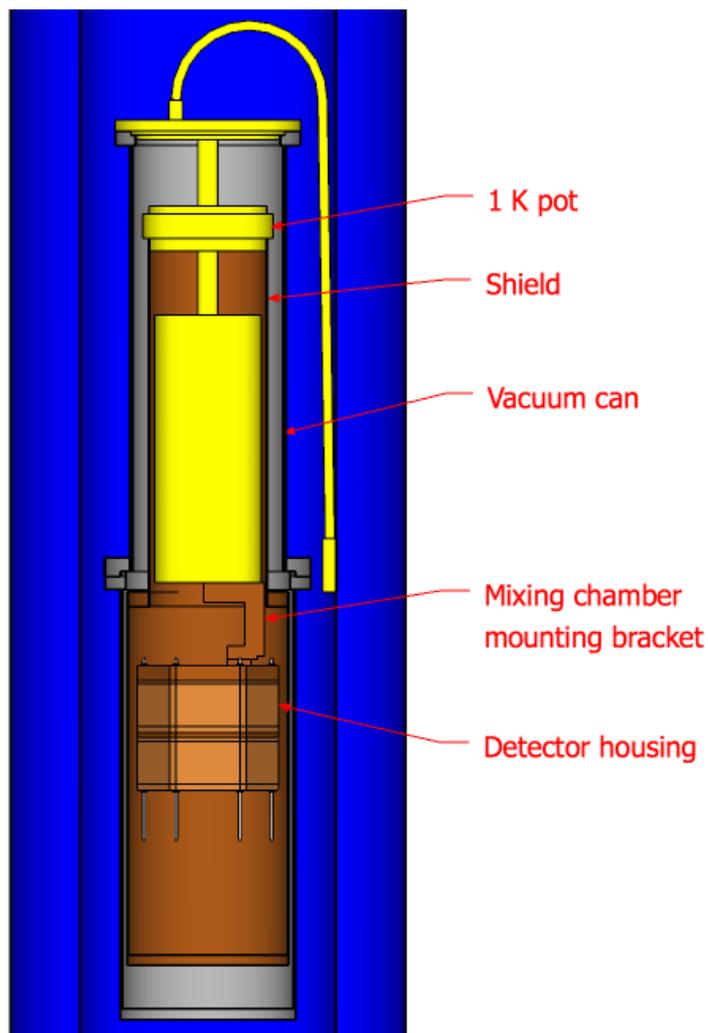


Figure 8.4: Schematic of the preexisting and new hardware. The upper piece of the new vacuum can (light gray) fits the existing refrigerator flange (yellow). The lower piece is wide enough to accommodate ZIP housings and a shield (brown) attached to the 1 K pot. It is narrow enough to fit in the Dewar (blue) and not interfere with the 1 K pot siphon (yellow).

of 4.22 inches at the opening and 4.38 inches throughout most of its depth. This allows clearance for a shield between the detector housings and the vacuum can walls. Clearance must also be allowed for the wiring connections on the sides of the detector housings. The position of the 1 K pot siphon is a limiting factor on the width of the lower piece. Another limiting factor is the necessary clearance for a tube used to remove liquid nitrogen from the Dewar that must be able to reach past the vacuum can to the bottom of the Dewar. The flanges are designed for indium o-ring vacuum seals.

A copper shield was built to attach to the 1 K pot and to provide complete coverage for the lower temperature stages of the refrigerator and the experimental payload. The shield has three pieces. The inner diameters of the upper and lower pieces are 3.00 inches and 4.09 inches, respectively. A third ring-shaped piece is an adapter for the upper and lower pieces. The shield pieces, vacuum can pieces, and experimental payload must be installed in a particular order, as shown in figure 8.5.

A hexagonal copper mounting plate was built to attach to the mixing chamber mounting bracket. Two hexagonal copper separating plates were also built. This allows for 1–3 detector housings to be arranged in a stack to attach to the mixing chamber. The mounting plate and separators have milled depressions to allow clearance for the heads of small screws used to attach the lids of the detector housings. Thus the lids need not be removed from the housings, reducing the possibility of contamination by dust containing radioactive isotopes. Holes in the plate and separators match holes in the housings and lids so that threaded rods with nuts and washers can be used to build the stack. The arrangement allows for good thermal contact between the housings and the plates.

8.2.3 Other modifications

The original wiring connectors used inside of the vacuum can were manufactured by Microtech. For the original bundle, connectors are only used at the point at which the wires connect to the experimental hardware. For the added bundle there are two additional connector buses attached at the 1 K pot and the mixing chamber. The Microtech connectors are miniature connectors with a 0.062 inch spacing between the pins. The connectors have a “solder cup” design for attaching the wires.

Early Little Blue runs were plagued with broken connections, often occurring only

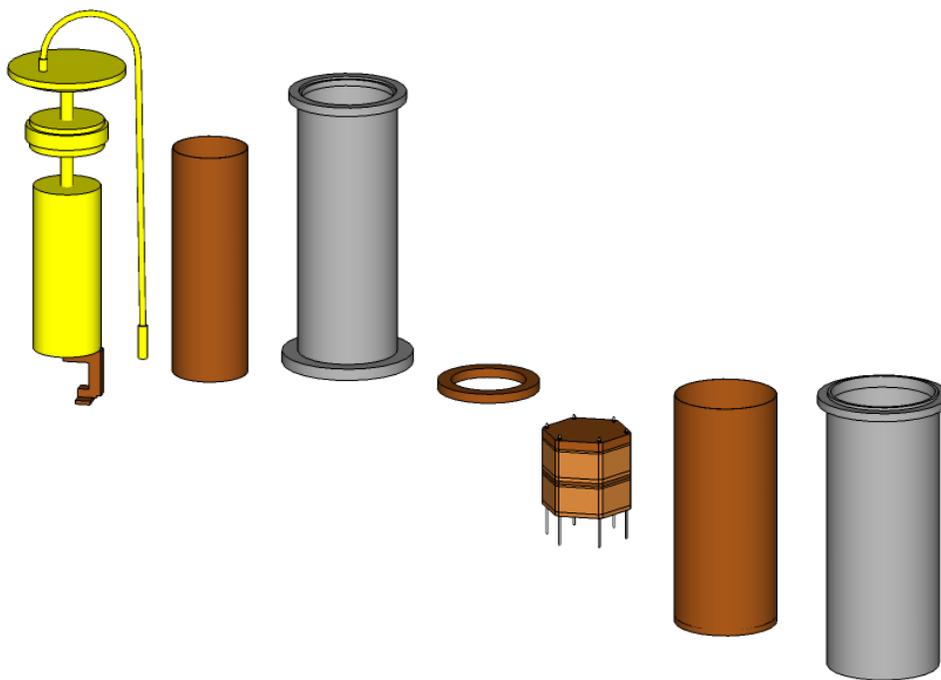


Figure 8.5: Components of new hardware. The components must be installed in the order shown: upper shield, upper vacuum can, shield adapter, detector stack, lower shield, lower vacuum can. Removal is in the reverse order.

after cooling, presumably due to thermal contraction. Eventually all of the Microtech connectors were replaced with larger connectors (0.100 inch pin spacing) manufactured by Samtec. Wires attach to these connectors by wrapping the ends around posts and soldering them. These more secure connections have proved quite stable and reliable.

A new breakout box was built for the wires from the second wire bundle. This breakout box was built with a switchboard design for flexibility of use. Any combination of wires can be connected to any of ten positions of a switchable 4-wire output for convenient connection to a bridge, multimeter, or other instrument.

Ruthenium oxide thermometers are common in cryogenics due to the strong dependence of their resistance to temperature in the mK range and their ability to retain their properties with thermal cycling. They require individual calibration, and commercially available RuO₂ thermometers are much more expensive if purchased calibrated. In early Little Blue runs the temperature calibrations being used were called into question due to disagreements in certain transition temperatures compared to other test facility results. Careful recalibrations were performed using a borrowed professionally calibrated thermometer mounted alongside the two portable thermometers. These calibrations have a precision within approximately 2 mK in the 100 mK range and show good agreement with the results of other test facilities.

Two features are shared with a newer CDMS test facility (nicknamed K100) now located in the same room as Little Blue. One is a HEPA-filtered clean room for assembly of the detector stack and storage of the stack hardware. The other is a LakeShore Model 370 AC resistance bridge. This instrument complements the BTi conductance bridge. It makes constant-current measurements rather than constant-voltage measurements and has a wider range of excitation settings and measurable resistances. During K100 runs this instrument is not available for use with Little Blue.

A clean tent was built for installing the detector stack in the refrigerator. It is composed of clear vinyl walls and ceiling attached to an aluminum frame and hung from the ceiling during installation and removal. A HEPA-filtered air flow allows for a dust-free environment to protect detectors and housings from contamination.

8.3 Testing techniques

Various methods have been developed to measure detector properties:

Resistances

4-wire resistance measurements are typical in cryogenic experiments to cancel large lead resistances. One simple method is to run a known current through the element to be measured using one pair of leads, then measure the voltage across the element using the other pair of leads. The bridge instruments described above measure resistances in a conceptually similar way, although these are AC rather than DC measurements.

A key issue for measuring resistances in cryogenic experiments is Joule heating. Heating on the scale of a μW or more could alter the base temperature of the experiment; heating on a smaller scale could alter the temperature of the element being measured so that it may no longer match the measured temperature.

Experience has shown that 4-wire measurements with a Keithley model 2000 multimeter cause excessive Joule heating in typical measurements below approximately 1 K. At lower temperatures a bridge can be used. The availability of two bridges, one that makes constant-voltage measurements and another that makes constant-current measurements, is useful when results are unstable. For example, a superconductor at its transition edge is stable under constant-voltage measurements but unstable under constant-current measurements. A third, DC technique is also useful when measurements are ambiguous. This involves taking voltage vs. current curves, with the current supplied by a 12 V DC power supply in series with a 1 M Ω ballast resistor and the voltage measured with the Keithley 2000 multimeter.

If four wires are not available for a measurement, 3-wire or 2-wire measurements are still possible. The meter resistances will then be a sum of the component resistance and one or two lead resistances. We gained much experience interpreting such measurements in early runs when wiring connections were lost.

Typical resistance measurements include normal-state resistances of tungsten sensors ($\sim 0.5\Omega$), substrate resistances ($\sim 1\text{M}\Omega$ at base temperature), and parasitic resistances ($\sim 10\text{m}\Omega$). Parasitic resistance measurements require the sensitivity of the LakeShore bridge.

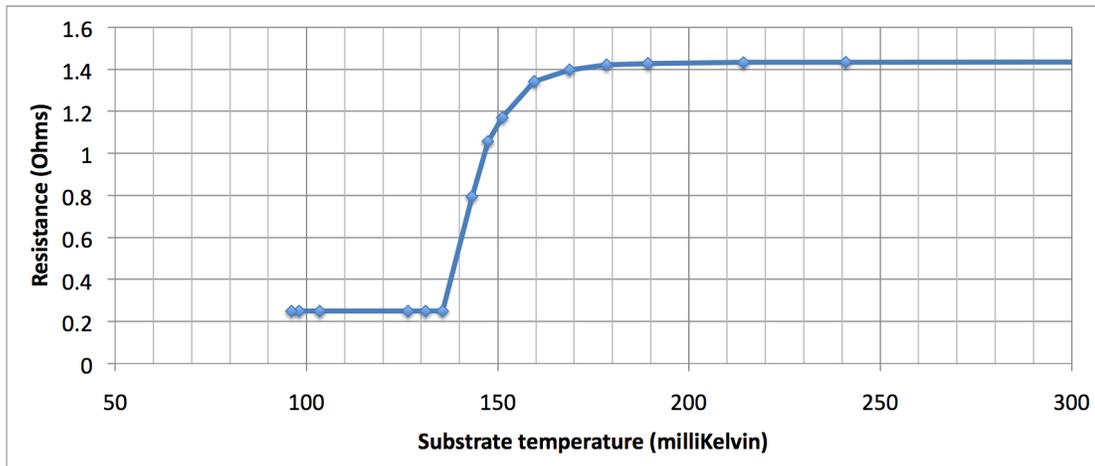


Figure 8.6: Superconducting transition of a phonon sensor array as measured by the BTi bridge with a $30 \mu\text{V}$ excitation. The minimum resistance measurable by the BTi bridge is 0.25Ω .

Critical temperature

The critical temperature of the tungsten sensors can be determined by taking resistance vs. temperature measurements. Figure 8.6 is an example. With detailed information it is also possible to characterize the width of the transition.

In figure 8.6 the resistances are 4-wire measurements taken with the BTi bridge. Below 135 mK it reads its minimum value of 0.25Ω . If a 3-wire measurement were performed the curve would be similar but with a pedestal of 70Ω due to the lead resistance. On some occasions when only the BTi bridge was available this has actually proven useful for measuring the lower resistance part of the curve.

If the width of the transition is not of interest, critical temperatures can be measured very quickly by monitoring the resistance during a temperature sweep. When the transition is seen the temperature can be noted. When the approximate temperature has been found, a more precise measurement can be made with slower temperature sweeps. If agreement is found between heating and cooling sweeps one can rule out any effects due to lags in heating or cooling.

The critical temperature depends on the excitation of the measurement, so that the

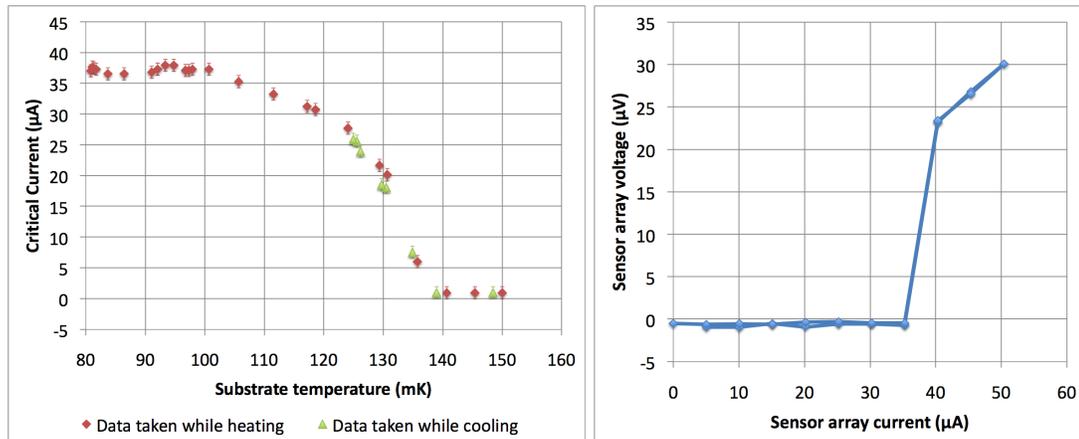


Figure 8.7: Critical current of a phonon sensor array. The critical current vs. temperature curve (left) is generated by taking voltage vs. current curves such as the one shown (right, taken at 87 mK). The voltage vs. current curve shows a very clear transition from superconducting to ohmic behavior.

excitation should be reported along with the critical temperature. True critical temperature can be considered to be the limiting case of zero excitation. Very low excitations however result in noisier measurements so a convenient standard is often chosen. At high enough excitation currents the tungsten can be driven out of its superconducting state regardless of the temperature as described below.

Critical current

The critical current for a phonon sensor array is the maximum current at which the array is in a superconducting state. The critical current depends on temperature and drops to zero above the true critical temperature. The most direct way to measure the critical current is to collect DC voltage vs. current data as described above for resistance measurements. The transition from zero resistance to ohmic behavior is quite clear (see figure 8.7). Such measurements are taken over a range of temperatures to produce critical current vs. temperature data. The individual voltage vs. current curves are taken both while current is increasing while it is decreasing as a check for hysteresis effects such as heating/cooling lags. For similar reasons the critical current vs. temperature curves are generated both while heating and while cooling.

Generating critical current information in this way is time-consuming because each data point involves a current sweep, and because manual switching is required between multiple sensor arrays. Partial automation of the process using the LabVIEW programming language for data acquisition speeds the process somewhat. A faster way to acquire a rough estimate of critical current is to simply increase the excitation setting on a bridge until a non-zero resistance is seen. More precise results are achieved by sweeping the temperature at each excitation setting until the transition is seen. Only a few data points can be gathered this way because consecutive excitation settings on the bridges are separated by a factor of three. These AC measurements may not be consistent with DC measurements.

Temperature sweeps

After achieving base temperature, the temperature of the detector stack can be controlled by use of the mixing chamber heater. This is done manually with a 6-volt lantern battery in series with a decade resistance box. Very steady temperature control can be achieved this way. For temperatures around 100 mK an appropriate current is approximately 0.4 mA. The heater resistance is 130 Ω , so this corresponds to a heating power of approximately 20 μW .

The BTi and LakeShore bridges include temperature control functionality but to date we have not used them, relying on the more direct method described above.

More sophisticated testing techniques can certainly be imagined. For example the critical current measurements could probably be done more quickly using a function generator and oscilloscope with LabVIEW interfaces, both of which are already available. However the speed would still be limited by manual switching between the different elements to be measured. Electronic switching would require an additional material investment. So far the existing techniques have proven perfectly adequate for the measurements performed with Little Blue. Typically all measurements are completed within a day of reaching base temperature.

8.4 Summary of Little Blue test runs

The Little Blue test facility has performed 16 test runs to date measuring the properties of 16 different detectors. Some have been measured multiple times; for example in two cases where an aluminum short circuit was found, a repair was attempted and the detector retested to determine its success. Apart from aluminum shorts, tests have revealed problems such as low substrate resistances, high normal-state resistances, and high parasitic resistances. In other cases test results revealed no pathologies, and the detector was sent to other facilities for full characterization. Little Blue has also been employed for measurements on sample wafers to test specific elements of new fabrication strategies.

The detectors tested include two new designs: the mZIP and the iZIP. Both are 2.5 times thicker than CDMS II ZIPs, and both have modified phonon sensor arrays for improved X–Y position reconstruction. The iZIP has twice as many phonon arrays and ionization grids and employs a different strategy for identifying surface interactions: interleaved ionization electrodes of alternating voltage on each face produce symmetric signals for bulk interactions and asymmetric signals for surface interactions. Six of the detectors tested with Little Blue are currently in use for SuperCDMS at the Soudan Underground Laboratory.

References

- [1] V. C. Rubin, W. K. J. Ford, and N. Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 ($R = 4\text{kpc}$) to UGC 2885 ($R = 122\text{ kpc}$). *Astrophys. J.*, 238:471–487, June 1980.
- [2] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110, 1933.
- [3] A. Cavaliere and R. Fusco-Femiano. X-rays from hot plasma in clusters of galaxies. *Astron. Astrophys.*, 49:137–144, 1976.
- [4] E. Jullo *et al.* Cosmological Constraints from Strong Gravitational Lensing in Clusters of Galaxies. *Science*, 329(5994):924–927, 2010.
- [5] M. Kowalski *et al.* Improved Cosmological Constraints from New, Old, and Combined Supernova Data Sets. *Astrophys. J.*, 686:749–778, October 2008.
- [6] D. Clowe *et al.* A direct empirical proof of the existence of dark matter. *Astrophys. J.*, 648:L109–L113, 2006.
- [7] G. R. Blumenthal, S. M. Faber, J. R. Primack, and M. J. Rees. Formation of galaxies and large-scale structure with cold dark matter. *Nature*, 311:517–525, 1984.
- [8] G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric dark matter. *Physics Reports*, 267(5-6):195 – 373, 1996.
- [9] L. Roszkowski, R. R. de Austri, and R. Trotta. Implications for the Constrained MSSM from a new prediction for $b \rightarrow s\gamma$. *Journal of High Energy Physics*, 2007(07):075, 2007.

- [10] J. Ellis, K. A. Olive, Y. Santoso, and V. C. Spanos. Update on the direct detection of supersymmetric dark matter. *Phys. Rev. D*, 71(9):095007, May 2005.
- [11] E. A. Baltz and P. Gondolo. Markov chain Monte Carlo exploration of minimal supergravity with implications for dark matter. *Journal of High Energy Physics*, 2004(10):052, 2004.
- [12] R. Gaitskell, V. Mandic, and J. Filippini. DMTools Limit Plot Generator. <http://dmttools.brown.edu:8080>.
- [13] Z. Ahmed *et al.* (CDMS). Search for Axions with the CDMS Experiment. *Phys. Rev. Lett.*, 103:141802, 2009.
- [14] J. D. Lewin and P. F. Smith. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. *Astroparticle Physics*, 6:87–112, December 1996.
- [15] A. Kurylov and M. Kamionkowski. Generalized analysis of the direct weakly interacting massive particle searches. *Phys. Rev. D*, 69(6):063503, Mar 2004.
- [16] M. T. Ressell *et al.* Nuclear shell model calculations of neutralino-nucleus cross sections for ^{29}Si and ^{73}Ge . *Phys. Rev. D*, 48(12):5519–5535, Dec 1993.
- [17] C. E. Aalseth *et al.* (CoGeNT). Results from a Search for Light-Mass Dark Matter with a P-type Point Contact Germanium Detector. 2010. [arXiv:1002.4703](https://arxiv.org/abs/1002.4703).
- [18] R. Bernabei *et al.* First results from DAMA/LIBRA and the combined results with DAMA/NaI. *The European Physical Journal C - Particles and Fields*, 56:333–355, 2008.
- [19] B. S. Neganov and V. N. Trofimov. Possibility of producing a bulky supersensitive thermal detector at a temperature close to absolute zero. *J. Exp. Theor. Phys.*, 28:328, 1978.
- [20] P. N. Luke. Voltage-assisted calorimetric ionization detector. *J. Appl. Phys.*, 64:6858, 1988.

- [21] D. S. Akerib *et al.* (CDMS Collaboration). Exclusion limits on the WIMP-nucleon cross section from the first run of the Cryogenic Dark Matter Search in the Soudan Underground Laboratory. *Phys. Rev. D*, 72(5):052009, Sep 2005.
- [22] T. Bruch. *A Search for Weakly Interacting Particles with the Cryogenic Dark Matter Search Experiment*. Ph.D. thesis, Universität Zürich, 2010.
- [23] A. Reisetter. *Results from the Two-Tower Run of the Cryogenic Dark Matter Search*. Ph.D. thesis, University of Minnesota, 2005.
- [24] J. Sander. *Results from the Cryogenic Dark Matter Search Using a Chi Squared Analysis*. Ph.D. thesis, University of California, Berkeley, 2007.
- [25] S. Golwala. *Exclusion Limits on the WIMP-Nucleon Elastic-Scattering Cross Section from the Cryogenic Dark Matter Search*. Ph.D. thesis, University of California, Berkeley, 2000.
- [26] W. W. M. Allison *et al.* Measurement of the atmospheric neutrino flavour composition in Soudan 2. *Phys.Lett.B*, 391:491, 1997.
- [27] J. Allison *et al.* Geant4 developments and applications. *IEEE Transactions on Nuclear Science*, 53(1):270–278, Feb 2006.
- [28] A. Ferrari, P. R. Sala, A. Fass, and J. Ranft. *FLUKA: A multi-particle transport code (program version 2005)*. CERN, Geneva, 2005.
- [29] J. Lindhard, M. S. V. Nielsen, and P. Thomsen. Integral equations governing radiation effects. *J. Appl. Phys. Mat. Fys. Medd. Dan. Vid. Selsk.*, 33(10):1–42, 1963.
- [30] T. Shutt *et al.* A solution to the dead-layer problem in ionization and phonon-based dark matter detectors. *Nucl. Instrum. Meth.*, A444:340–344, 2000.
- [31] V. Mandic *et al.* Study of the dead layer in germanium for the CDMS detectors. *Nucl. Instrum. Meth.*, A520(1-3):171 – 174, 2004.
- [32] J. Filippini. *A Search for WIMP Dark Matter Using the First Five-Tower Run of the Cryogenic Dark Matter Search*. Ph.D. thesis, University of California, Berkeley, 2008.

- [33] R. Brun and F. Rademakers. ROOT - An Object Oriented Data Analysis Framework. *Nucl. Inst. & Meth. in Phys. Res. A*, 389:81–86, 1997. See also <http://root.cern.ch/>.
- [34] S. Yellin. Finding an upper limit in the presence of an unknown background. *Phys. Rev. D*, 66(3):032005, Aug 2002.
- [35] A. T. Lee *et al.* Measurements of the ballistic-phonon component resulting from nuclear and electron recoils in crystalline silicon. *Phys. Rev. B*, 54(5):3244–3256, Aug 1996.
- [36] E. Aprile *et al.* First Dark Matter Results from the XENON100 Experiment. *Phys. Rev. Lett.*, 105:131302, 2010.
- [37] F. Pobell. *Matter and Methods at Low Temperatures*. Springer, 2nd edition, 2002.
- [38] M. Williams. *Creation and Detection of Low Temperature Helium Atom Beams and Suspension of Superfluid Helium for the Study of Bose-Einstein Condensation in Superfluid Helium*. Ph.D. thesis, University of Minnesota, 1998.
- [39] K. Lidke. *Transmission Experiments with Superfluid Helium*. Ph.D. thesis, University of Minnesota, 2002.

Appendix A

Little Blue Operations Manual

Instructions for operation of the Janis refrigerator used for Little Blue were included in the theses of Mark Williams [38] and Keith Lidke [39]. These proved quite useful for my own training on the equipment. This operations manual is intended to be significantly more thorough. The reason for this is that Little Blue is now a multi-user facility. Both research groups currently employing it utilize temporary undergraduate workers. The instructions below are meant to be appropriate for an operator who has had a couple of weeks of training.

A.1 Install payload

At this point the fridge is removed from the cryostat and it is assumed that all vacuum can and 1K-shield pieces are removed. IVC stands for Inner Vacuum Can. MC stands for Mixing Chamber.

THESE INSTRUCTIONS ARE GENERIC. The order in which parts are assembled, plus any additional steps (e.g. installing feedthroughs, cell pumpout, and clean-handling procedures) are specific to the experiment (e.g. helium experiments or CDMS).

1. Check wiring before attaching 1K-shield to 1K Pot.
 - (a) Attach both wiring cables to top of fridge.
 - (b) Use ohmmeter to check continuity of relevant wires from the fridge to the breakout boxes in the rack.

- (c) Use conductance bridge and LabVIEW software to check for correct readings from built-in thermometers (1K Pot, Still, MC) and other thermometers (e.g. THIN, THICK, CASE).
 - (d) Wiring repairs may be necessary at the plugs attached to the 1K Pot, or the plugs attached to the MC, or any connections below that. (Of course it is also possible that repairs are necessary at the top of the fridge, in the cables, or at the breakout boxes.)
2. Attach 1K-shield to 1K Pot.
- (a) Polish copper 1K-shield with a Scotch Brite pad if necessary.
 - (b) Prepare surfaces of contact between 1K Pot and 1K-shield by polishing them with a Scotch Brite pad.
 - (c) Attach 1K-shield tightly to 1K Pot with appropriate screws. The heads must be short enough to allow clearance for the IVC. The 1K-shield must be aligned properly to allow clearance between it and all fridge and payload elements below the 1K Pot.
3. Attach experimental payload to MC bracket.
- (a) Prepare surfaces of contact between MC bracket and experimental payload by polishing them with a Scotch Brite pad.
 - (b) Attach payload to MC bracket tightly with correct screws (3 of them). Note: dont over-tighten brass screws, as they are relatively delicate, and they will become tighter as the experiment is cooled.
4. Attach and secure wiring to experiment.
- (a) Connect fridge wires to experiment as appropriate.
 - (b) Secure all wiring to experiment so that no wires will touch the 1K-shield or IVC. There are several methods:
 - i. Dental floss. Tie down wiring with floss, cut off excess. Clean, but can be difficult.
 - ii. Scotch or masking tape. Easier than floss perhaps but can be messy (adhesive on wires and experiment), adhesive may outgas (although weve never had a problem with pumpout), and tape may lose adhesion under vacuum or when cold. Can secure tape with floss.
 - iii. Paper. White paper can be taped down around experiment. This avoids putting tape directly on wires or experiment.
 - iv. Kapton film and dental floss. Lengths of dental floss can be woven through slits in a sheet of Kapton film of the appropriate size to wrap around the experiment. This is the prefer method if the experiment must be kept in a clean environment.

5. Attach Inner Vacuum Can (IVC).
 - (a) Form an indium o-ring from indium wire on the vacuum can o-ring groove. Make sure it fits snugly and overlaps itself. Both o-ring surfaces should first be cleaned of indium and wiped with methanol using a cotton swab.
 - (b) Attach the can with the appropriate screws, but do not tighten yet. Insert each screw so that the gap is uniform, a millimeter or two.
 - (c) Tighten the screws in a star pattern. Continue to tighten them until all are as tight as can comfortably be using a hex key.
 - (d) Optional: Leak-check inner vacuum can. The indium seals are pretty reliable, so this can be skipped – there is an opportunity to leak-check later (after cooling to 77K) also. If you do leak-check, check the connections and electrical feedthroughs at the top of the fridge, and then work down to the vacuum can seals.

6. Pump out IVC.
 - (a) Rough pumping: Attach metal hose from the IVC flange at the top of the fridge to the wall gas-handling system (WGHS). Pump out IVC through the large black IVC valve and the WGHS using the rough pump in the pump alley. Monitor pressure using the MKS gauge and the thermocouple gauge on the rack.
 - (b) When pressure stops dropping, close black valve and GHS valve, and attach a metal hose from the IVC flange to the turbo pump station.
 - (c) Turn on turbo station rough pump using power strip switch. Open valve on pump station to pump out hose leading to IVC flange. Pump for about 5 minutes.
 - (d) Open IVC valve at top of fridge.
 - (e) (If desired, rough pump for about 10 minutes or so, then close the valves and remove the hose, skip ahead to the Install section ??, then finish turbo pumping after installing the fridge.)
 - (f) Turn on turbo pump by plugging it into the power strip, waiting for the green Power light on the controller, then plugging in and removing the 25-pin starter plug to the controller. The yellow Start light should come on, and you will begin to hear the whine of the turbo pump increase in pitch.
 - (g) Turn on the cold cathode pressure gauge on the pump station. If it is off the top of the scale, turn it back off so as not to foul the filament.
 - (h) Pump for at least a few hours. Target pressure may vary with experiment but should at least be below 10^{-4} Torr. (It will never be below 10^{-6} Torr.)

7. Shutting down the turbo pump:

- (a) Turn off cold cathode gauge.
- (b) Close valve to pump.
- (c) Turn off turbo pump by unplugging the controller.
- (d) Wait about 15 minutes for the turbo pump to spin down.
- (e) Turn off rough pump by switching off power strip.

Quick notes on order of steps:

- CDMS

1. Check wiring.
2. Install upper 1K shield.
3. Install upper IVC.
4. Install 1K shield ring.
5. Set up purge skirt.
6. Assemble detector stack in cold desk.
7. Install detector stack and connect wiring.
8. Install lower 1K shield
9. Install lower IVC.
10. Pump out IVC.
11. Check wiring.

- HELIUM EXPERIMENTS

1. Payload
2. Shield
3. Upper IVC
4. Feedthroughs
5. Lower IVC

For details on the helium suspension experiment see Keith Lidke's thesis [39].

A.2 Install refrigerator in cryostat

At this point the fridge is removed from the cryostat, the experiment has been installed, and the Inner Vacuum Can (IVC) is sealed. Cryostat refers to the blue refrigerator Dewar.

Installing the refrigerator in the cryostat is a two-person job. The two-person steps are indicated with a slash: first person / second person.

1. Disconnect everything from top of fridge.
 - (a) Disconnect pumping hose from IVC and any other tubing.
 - (b) Disconnect wiring connectors.
 - (c) Clear loose things on the top of the fridge rack out of the way.
 - (d) Check that nothing is at the bottom of the cryostat.
 - (e) Check that the cryostat o-ring is in place.
 - (f) Remove the black stopper from the still pumping line at the top of the fridge rack.
2. Connect rope to top of fridge.
 - (a) Loosen left rope on the pillar by the fridge gas-handling system rack so that the rope connector can be lowered to the top of the fridge. The right rope should be secured on the wall hook.
 - (b) Secure eyelets to the three threaded holes in the fridge top flange.
 - (c) Connect hooks from the rope connector to the fridge eyelets. Guide the ropes to the eyelets so that they are as vertical as possible.
3. Raise fridge with pulley system / guide fridge.
 - (a) The guider should stand on a stool and keep the fridge from banging against the fridge rack.
 - (b) The puller should pull the fridge up using the pulley system with a firm grip on the rope.
 - (c) Pull the fridge up until it clears the top of the fridge rack.
4. Position fridge above cryostat with pulley system / guide fridge.
 - (a) The puller should secure the lifting rope firmly on the wall hook, and then loosen the rope from the right hook.
 - (b) Use the rope from the right hook to move the fridge horizontally above the cryostat. The guider should judge when it is centered above the hole.

- (c) Secure the rope on the right hook.
5. Lower fridge into cryostat / guide fridge.
- (a) Loosen the lifting rope from the left hook and hold firmly.
 - (b) The puller should lower the fridge slowly with the rope while the guider guides the fridge into the cryostat.
 - (c) The guider should take care when the 1K Pot siphon reaches the cryostat — it may be necessary to hold it in against the IVC. Also be careful not to knock the cryostat o-ring out of place. As the fridge is lowered further it may be necessary to guide it back and forth a little so that the siphon can clear the weld joint part-way down the cryostat.
 - (d) When the fridge flange is 1 cm above the cryostat top the puller should pause while the guider rotates the fridge so that the large still pumping port lines up with the still pumping hose.
 - (e) As the fridge is lowered the rest of the way the guider should line up the holes in the fridge flange with the threaded holes on top of the cryostat.
 - (f) After the fridge is in place, disconnect the hooks from the fridge flange eyelets.
 - (g) Raise the pulley system well out of the way above the fridge and secure the rope to the right hook.
 - (h) Remove the eyelets from the fridge flange.
 - (i) Insert the six screws connecting the fridge flange to the cryostat, but leave them loose at this point.
6. Connect tubing.
- (a) Still pumping line: This is the largest pumping line on top of the fridge. Find the hex key that is in the box on top of the fridge rack — it has been cut short for better clearance. Remove the cover from the fridge pump port. Check that the o-ring is in place in the fridge flange. Pull the flexible hose forward and twist it as necessary so that its flange lines up with the fridge flange; it might be easier to have a second person for this. The screw holes are not aligned perfectly so it is necessary to insert the screws in a certain order. First insert a screw into the hole labeled 1, FORWARD at the arrow. Do not tighten all the way — leave at least 1 mm. Next insert a screw backward into the hole labeled 2, BACKWARD far enough to hold the flange in place. You may need to wrestle a bit with the hose at this stage. Next insert the other six screws forward — don't tighten fully. Then remove the screw that was inserted backwards and insert it forward through the same hole. Finally tighten all the screws in a star pattern.

- (b) Still return line: this is the small port clockwise from the still pumping line and situated low (close to the fridge top) with a KF-10 Quick Flange connection. Remove the clamp and the blank from the port. Attach the flexible hose with the proper end flange; this normally hangs down in front of the back wall, and leads to the fridge gas-handling system. Check that the VCO connection of this hose to the gas-handling system is tight.
 - (c) 1K Pot pumping line: this is the high port counterclockwise from the still pumping line with a KF-40 Quick Flange connection. The flexible hose that connects to this line is normally resting on top of the large still pumping line hose and leads to the large gray pumping line that extends to the rear corner of the lab. Remove the clamps and blanks from the port and the hose. Attach the hose to the port.
 - (d) Top flange: finally tighten the screws attaching the top flange of the fridge to the cryostat. Tighten them in a star pattern.
 - (e) Store the clamps, blanks, hex keys etc. in the box on top of the fridge rack.
7. Connect wiring cables.
- (a) The two wiring cables can now be connected to the top of the fridge. The connections are different so it should be obvious which goes where. Attach and tighten carefully — these are vacuum feedthroughs and if too much stress is put on the pins the vacuum seal will be broken. One of these connectors tightens with a threaded nut; the threads are damaged so it is difficult to tighten properly. Do your best to screw it on to a point where the connection feels tight.
 - (b) This is a good opportunity to do a room-temperature wiring check to check the integrity of the electrical connections at the wiring breakout box.

A.3 Pump out fridge circulation lines.

At this point the fridge has been installed in the cryostat and all plumbing connections are connected. The air must be pumped out of the still lines before LN₂ cooling. GHS refers to the Refrigerator Gas-Handling System, most of which is located on a stand north of the refrigerator. Refer to Figure A.1.

WGHS refers to the tubing mounted on the wall, which allows access to the rough pumps in the Pump Alley and the pressure gauges on the electronics rack (TC1, TC2, MKS1, MKS2). To use the WGHS there are various ports to connect to, near valves W1, W2, W3, W6, W7, and W15. Valve W8 leads to the rough pump. Check which valves are closed. Valves W12 and W13 should generally always be left open, since they lead to pressure gauges MKS1 and TC1.

At this point it is assumed that the mixture is stored in Dump 1 (and in the tubing in front of and behind the pump). Before proceeding be sure that this is the case and check that all valves are closed. G4 should be reading more than 20" up from the 30" mark.

Note on Gauges: the mechanical gauges on the GHS display pressures above atmospheric in PSI and below 1 atm in "Hg of vacuum — thus zero pressure is approx. 30 "Hg. In this manual, the convention is to express pressures below 1 atm in "absolute" inches, reading up from the minimum. Thus for example 2" refers to a dial reading of 28"Hg.

1. Check system before pumping.
 - (a) Check that all GHS valves are closed.
 - (b) Check that G4 (Dump #1) reads ≥ 20 ". This shows the mix is stored there.
 - (c) Check that G1 reads 0". This helps to confirm that mix has not been left elsewhere in the GHS.
 - (d) Check that the 77K Trap is at room temperature.
2. Open fridge to GHS.
 - (a) Open valves V4, V11, and V10. G1 will read close to 1 atm.
3. Pump through GHS.
 - (a) Attach a rough pump to the port (near V6) on the GHS with a metal hose. This can be either the pump on the Turbo pump station, or a rough pump

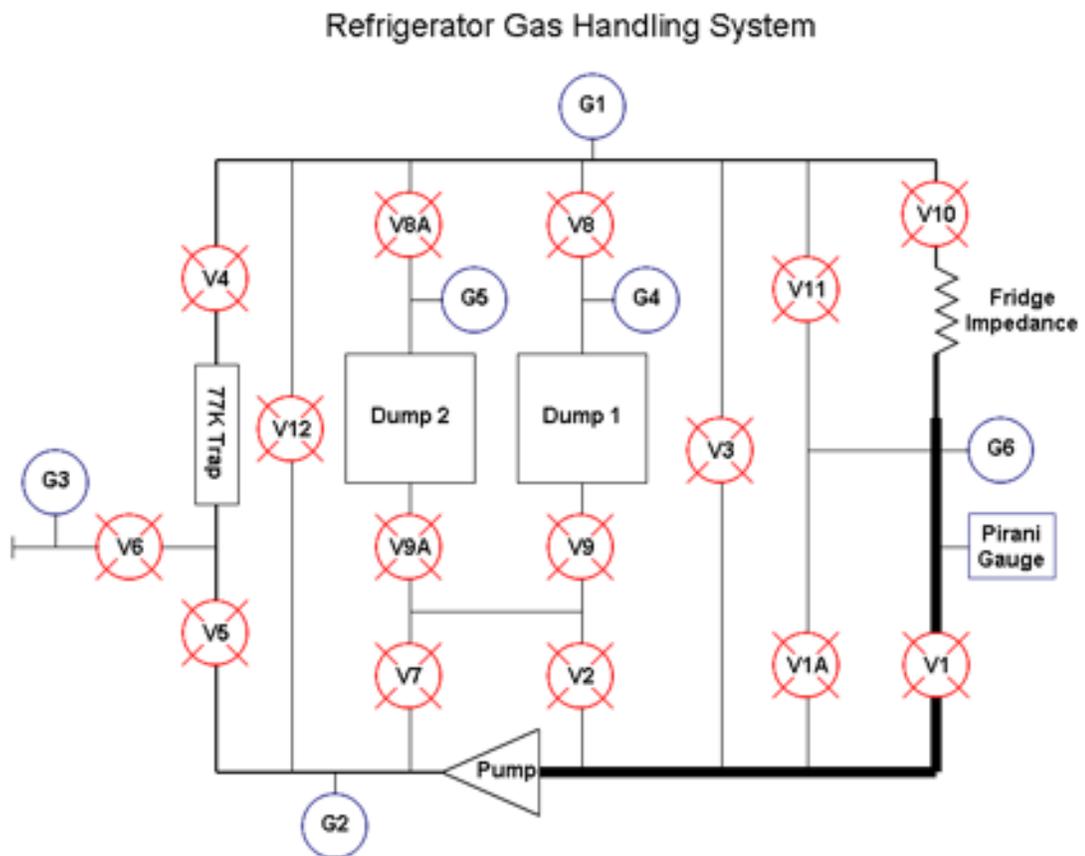


Figure A.1: Little Blue gas-handling system (GHS) showing manual valves (red circles) and mechanical pressure gauges (blue circles). The outer perimeter represents the circulation path during ordinary operation.

in Pump Alley via the WGHS. For turbo pump instructions refer back to section A.1: Install Payload step 6c to step 6g.

- (b) Pump out line leading to port. The pressure on G3 should drop.
 - (c) Open V6. The pressures on G1 and G6 should begin to fall.
 - (d) When the pressure falls below 2 Torr (less than 1 hour) the digital display on the Pirani gauge will start to register. (This pressure can also be monitored with the same LabVIEW VI that monitors temperature.)
 - (e) Wait at least until the pressure on the Pirani gauge (still pressure) is below 400 mTorr. This takes about 2 hours.
4. Cryopump fridge by cooling cold trap.
- (a) Close V6 and remove pump line from GHS.
 - (b) Insert the 77K Trap into its 5 liter Dewar behind the GHS. Feed the rubber hose from the fill funnel into the Dewar, and plug the hole with a piece of newspaper.
 - (c) Fill another 5 liter Dewar with liquid nitrogen. Pour this into the fill funnel attached to the GHS rack to fill the 77K Trap Dewar.
 - (d) The pressure on the Pirani gauge should drop further. Wait at least until this is below 300 mTorr before proceeding to liquid nitrogen transfer (should be less than 1 hour).
5. Clean mixture.
- (a) After the fridge lines have been pumped out, the GHS is free to clean the mixture. This can be done in parallel with the next two sections (A.4: Cool to 77 K and A.5: Cool to 4K). Follow instructions in section A.10: Cleaning Mixture.

A.4 Cool to 77 K.

Liquid N_2 is transferred into the cryostat using the round 25 liter Dewar with tubing designed to pressurize the Dewar with N_2 gas and push LN_2 into the cryostat.

The 1K Pot pumping line has a mechanical pressure gauge. Near this is a manual valve labeled “1K Pot Pressurization”. At the top of the fridge is a black knurled knob labeled “1K POT NEEDLE VALVE”.

WGHS refers to the wall gas-handling system (see section A.3 for description). The pressure can be read from the MKS digital gauge, MKS1 (which occasionally must be re-zeroed) on the electronics rack and also the thermocouple gauge dial, TC1, at the top left of the rack (which is sensitive to lower pressures but is not calibrated in pressure units). Tubes above the WGHS lead to the He and N_2 gas cylinders in the corner of the lab (green knobs = N_2 , blue knobs = He). These can be fed to ports on the WGHS.

Before proceeding with this step, the fridge lines must be evacuated. Follow the procedure in section A.3: Pump Out Fridge Circulation Lines.

1. Pressurize 1K Pot with helium. (The purpose of this step is to prevent ice or frozen air from entering the 1K Pot siphon during LN_2 cooling.)
 - (a) Close the 1K Pot Needle Valve completely but not too tight.
 - (b) Open top valves on He and N_2 cylinders in back corner of room.
 - (c) Run a metal hose from the He supply to the 1K Pot Pressurization port. Open the blue valve from the He supply.
 - (d) Open the 1K Pot Pressurization valve enough so that the gauge reads a few PSI.
 - (e) It will probably be necessary to leave this valve open to keep the 1K Pot pressurized during LN_2 cooling.
2. Put transfer gas in IVC.
 - (a) Attach a metal hose from the WGHS (near W6) to the IVC port.
 - (b) Pump out the line up to the IVC valve for a few minutes (until the thermocouple gauge TC1 reads low pressure/high current).
 - (c) Close valve W8 on the WGHS to isolate it from the pump.
 - (d) Open the IVC valve.
 - (e) Open a green valve from the N_2 gas supply to purge the hose, and then attach hose to a port on the WGHS.

- (f) Let 1 atm of N_2 into the IVC through the WGHS and pump it back out. Repeat this flush once more.
 - (g) Let 30 Torr of N_2 into the IVC as a transfer gas. Leave the IVC open to the WGHS to monitor the pressure with MKS1.
3. Transfer liquid N_2 into cryostat.
- (a) Fill the round 25 liter Dewar with liquid N_2 .
 - (b) Attach o-ring adapter to the fill port (labeled "CRYOSTAT FILL") at the top of the fridge, if it's not already attached. Insert the long L-shaped transfer tube into the fridge fill port. This is kept on the rack attached to the square pillar. Insert until it touches the bottom of the cryostat (you will need to rotate it back and forth as you insert it).
 - (c) Position the 25 liter Dewar on a stool in front of the fridge.
 - (d) Insert the tubing apparatus in the 25 liter Dewar without sealing the stopper.
 - (e) Connect the rubber hose extending from the end of the tubing apparatus to the end of the cryostat transfer tube.
 - (f) Insert the stopper into the 25 liter Dewar and attach the spring hooks to the Dewar handles.
 - (g) Connect a hose from the N_2 supply to the end of the tube near the opening of the 25 liter Dewar.
 - (h) Open the green valve to pressurize the 25 liter Dewar. Liquid N_2 should start to flow into the cryostat through the rubber hose which will freeze stiff.
 - (i) Record IVC pressure (MKS1), Still pressure (digital Pirani gauge on GHS), and fridge temperature at intervals during the transfer. IVC pressure should fall by about $1/2$ in an hour.
 - (j) The transfer will take about 1 hour, depending somewhat on the gas pressure. The easiest way to determine that the transfer is finished is that the rubber hose will go limp. At this time close the green valve from the N_2 gas and disconnect the spring hooks on the 25 liter Dewar and lift up the stopper to release the pressure.
 - (k) Monitor temperature and allow at least 6 hours for fridge to cool. The 1K Pot temperature should be reading below 90 K at this time. The IVC pressure should be 9 Torr or less.
 - (l) About 12 hours after the LN_2 transfer the temperature will begin to rise again. If you won't be in the lab at that time, you should transfer about 5 additional liters before leaving for the day, using the same method as the original transfer.

A.5 Cool to 4 K.

At this point liquid nitrogen has been in the cryostat for at least 6 hours, there is N₂ transfer gas in the IVC, and the 1K Pot pumping line is pressurized to a few PSI with He gas. If the cryostat has been left overnight, the temperature should be 77 K unless there was not enough LN₂ in the cryostat, in which case it may have begun to warm up again. If the fridge thermometers are reading much above 100 K, you should transfer more LN₂ and allow to cool before proceeding.

1. Pump out and leak check IVC.
 - (a) Open W8 on the WGHS to pump out IVC to less than 1 Torr.
 - (b) Close IVC valve and WGHS valve W6 and remove hose from IVC.
 - (c) Attach Alcatel leak detector to IVC with metal hose. (Leak-detection is optional; if skipped, be sure to pump out IVC thoroughly.)
 - (d) Pump out line to IVC with leak detector or turbo pump.
 - (e) Open IVC valve to leak detector or turbo pump. Record background signal in log book if using LD. 10^{-9} to 10^{-10} mbar·l/s is normal.
2. Remove LN₂.
 - (a) Make sure the L-shaped cryostat transfer tube is pushed down all the way to the bottom of the cryostat.
 - (b) Attach rubber hose from the cryostat transfer tube to the 25 liter Dewar (if not still attached). Insert stopper in Dewar, lock down with spring hooks, and remove gas-feed hose from 25 liter Dewar tubing apparatus.
 - (c) Attach a rubber He gas supply hose to the port at the top of the fridge labeled "CRYOSTAT EXHAUST".
 - (d) Open the blue valve from the He gas supply to pressurize the cryostat. LN₂ should start to flow through the hose into the 25 liter Dewar.
 - (e) Monitor leak detector signal while removing LN₂. Signal should not increase much.
 - (f) Removal time is typically 30-60 minutes. Hose should go mostly limp. (Make sure that the He gas cylinder isn't empty.) When this occurs, close blue valve to He gas supply and disconnect hose from fridge exhaust port.
 - (g) Close IVC valve, remove and turn off leak detector.
3. Fill IVC with ³He exchange gas.

- (a) The ^3He exchange gas supply is located in an orange bottle on the wall next to the WGHS. Check that the black valve on the bottle and the black valve on the regulator are closed.
 - (b) Connect a metal hose from port W6 on the WGHS to the IVC port. Pump out this line through the WGHS.
 - (c) Make sure that valve W9 is closed. (This used to lead to a second rough pump, but currently we only have one working pump. The upper left portion of the WGHS is actually a separate system with its own pressure gauges, TC2 and MKS2.)
 - (d) Attach a metal hose from port W2 to port W4 (this connects the two systems). Pump out the lines leading to the ^3He bottle regulator through valves W2 and W4.
 - (e) After the lines to the IVC and to the regulator are thoroughly pumped out, close pump valve W8. Open the IVC valve.
 - (f) Open the big black valve on the ^3He bottle, and then close it. This lets out just enough gas to fill the regulator tubing.
 - (g) Open the silver regulator valve (clockwise), and then close it (counterclockwise).
 - (h) Open the smaller black valve attached to the regulator, and then close it. Make a note of WGHS pressure (MKS1 gauge) in Torr and write it in the lab book.
4. Stop pressurizing 1K Pot.
 - (a) Close the 1K Pot Pressurization valve at the top of the fridge.
 - (b) Close the blue valve from the He gas supply and remove the hose from the 1K Pot Pressurization port.
 - (c) Open the 1K Pot Needle Valve a few full turns.
 5. Do initial LHe transfer.
 - (a) Remove the hose from the L-shaped cryostat transfer tube (when it is warm enough to do so). Remove the cryostat transfer tube from the cryostat by completely unscrewing the o-ring adapter. You may need to allow the tube to warm up so that the o-ring adapter can be removed and put back onto the cryostat fill port.
 - (b) Put the tubing apparatus and the 25 L Dewar away. (The remaining LN_2 can be poured into 5 L Dewars for later use in the cold trap.)
 - (c) Follow instructions in section A.11: Liquid Helium Transfer and follow the specific instructions for Initial Transfers.

6. Pump out IVC.
 - (a) Open valve W8 on the WGHS and pump out IVC for about 5 minutes.
 - (b) Close IVC valve and valve W6 and remove hose from IVC port.
 - (c) Attach a metal hose from the turbo pump station to the IVC.
 - (d) Turn on the rough pump on the turbo pump station by switching on the power strip. Open valve at turbo station and pump out line leading to IVC for 5 or 10 minutes.
 - (e) Open IVC port valve.
 - (f) Turn on turbo pump (refer back to section A.1: Install Payload step 6c to step 6g).
 - (g) After turbo spins up turn on cold cathode gauge.
 - (h) Pump for at least 1 hour before cooling below 4 K.
7. Clean the $^3\text{He}/^4\text{He}$ mixture.
 - (a) If not yet finished, follow the instructions in section A.10: Cleaning Mixture.
8. Dump mixture into fridge.
 - (a) Check that all GHS valves are closed except perhaps V4, and that the 77K Trap is immersed in LN_2 .
 - (b) Open V5, V7, V9, and V4 (if not already open) on the GHS.
 - (c) Slowly open V11 about $1/4$ turn, just enough for the reading on the digital Pirani gauge (still pressure) to change. It will quickly go off-scale.
 - (d) Wait until the pressure on G6 (still pressure) is the same as that on G4 (dump pressure), and then open V10.
 - (e) Leave the mixture to freeze out any remaining contaminants onto the fridge lines overnight before beginning circulation. You can proceed to section A.6: Cool to 1.4 K, but not to section A.7: Cool to Base.

A.6 Cool to 1.4 K.

At this point there is liquid helium in the cryostat (at least 10 inches above the 1K Pot siphon), the IVC has been turbo-pumped out, and the $^3\text{He}/^4\text{He}$ mixture has been dumped into the fridge. The 1K Pot Needle Valve is a black knurled knob on a rod extending from the top of the fridge. There are three 1K Pot Pumping Valves in parallel situated between the fridge and the back wall. These will be referred to as small (micrometer-type knob), medium (red handle), and large (orange handle). The 1K Pot pump is sometimes referred to as the 5-inch pump in reference to the 5-inch diameter pumping line.

1. Start the 1K Pot pump.
 - (a) Close all valves to the 5-inch pumping line. This line extends from the fridge area to the northeast corner of the lab and along the north wall. Usually one valve is left open, near the gas cylinders in the northwest corner, to vent the pump.
 - (b) Check that the valve from the 5-inch pump exhaust to the He recovery system is closed. This is located in the northeast corner of the lab, behind the K100 fridge, on a vertical copper pipe near the big gray pipes.
 - (c) Get key E1 located in the desk and enter room S10 down the hall. The 1K Pot pump is down the stairs and to the right, the green pump closest to the room entrance. Take a flashlight to view the oil level in the pump.
 - (d) Turn on the pump using the green START button. Sometimes on startup it draws too much current and trips the breaker. The breaker box and switch numbers are written on a box above the pump. The breakers are located in a room at the bottom of the stairs.
 - (e) When the pump is running check the sight glass to be sure the oil level is correct. The level should be about halfway up the sight glass.
2. Close the 1K Pot Needle Valve.
 - (a) Close the 1K Pot Needle Valve on top of the fridge completely but do not over-tighten.
 - (b) Check that the valve to the 1K Pot Pressurization valve at the top of the fridge is closed.
3. Open the 1K Pot pumping valves in sequence.
 - (a) Open the small 1K Pot pumping valve 5 rotations and wait 5 minutes. The 1K Pot pressure should fall to about 6" (gauge near 1K Pot port at top of cryostat) and the 1K Pot temperature should fall to about 2.5 K.

- (b) Open the medium valve 2 rotations and wait 5 minutes. The pressure should drop to about 3" and the temperature to about 2 K.
 - (c) Open the large valve 2 rotations. The 1K Pot temperature should drop to less than 1.5 K in about 5 minutes.
4. Open the 1K Pot pump exhaust to the recovery system.
- (a) In the northeast corner of the lab open the valve from the 5-inch pump exhaust to the He recovery system. (If the 1K Pot is being re-started during a run, this step will not be necessary.)

A.7 Cool to base.

Cooling to base involves circulating the $^3\text{He}/^4\text{He}$ mixture through the refrigerator lines using the Alcatel circulation pump in the Pump Alley. A common cause of fridge failure is the formation of an impedance blockage due to snow from contaminant gases. The mixture should be left cooling in the fridge lines overnight before beginning this section in the hope that this snow will stabilize on the walls of the circulation lines.

At this point it is assumed that the 1K Pot is at 1.4 K. The following GHS valves are assumed open: V4, V5, V7, V9, V10, V11, so that the mixture is open to the 77K Trap and the fridge. The mixture has been cleaned (section A.10), the 77K Trap Dewar is full of LN_2 , and the Alcatel circulation pump is running.

1. Pump mixture into circulation lines.
 - (a) Close V7.
 - (b) Open V2 about $1/4$ turn. Wait 10 minutes.
 - (c) Open V2 completely. Wait 10 minutes.
 - (d) Close V2 and V11.
2. Prepare log book.
 - (a) Make a table to record time, amount valve opened, pressure at gauge G1, Still temperature and pressure, and base temperature. This is useful for monitoring changes in fridge performance.
3. Begin circulation.
 - (a) Double-check that V2 and V11 are closed.
 - (b) Open V1A about $1/3$ of a turn.
 - (c) Record data in table.
 - (d) Wait a few minutes to see if the G1 pressure is increasing.
 - (e) If the G1 pressure is stable or decreases, open V1A completely.
 - (f) Record data in table.
 - (g) If the G1 pressure is stable or decreases, begin opening V1 in stages. This is the large valve on the thick circulation tubing located by the wall to the right and behind the refrigerator (it is the only valve in the GHS which is not located on the GHS rack). This valve has a certain amount of free rotation before it begins to open. In the first step open V1 $1/12$ of a rotation.

- (h) Iterate recording data, waiting for G1 to stabilize or decrease, and opening V1 in steps, up to $1/4$, $1/2$, 1, 2, 4 rotations, etc., until V1 is completely open. The pressure on G1 will typically not exceed 5" above vacuum. It should not be allowed to exceed 20" above vacuum. If this occurs, consult section A.8: Monitoring and Troubleshooting.
- (i) Consult data from previous runs to compare still temperature and pressure, G1 pressure, base temperature, cooling times etc.

A.8 Monitoring and troubleshooting.

Both the 1K Pot and the $^3\text{He}/^4\text{He}$ circulation line operate by pumping fluid through a constriction, or impedance. In both cases failure is most often due to blockage of the impedance by frozen material.

Another common cause of failure to cool is leakage of helium into the IVC, either from the cryostat bath, or for helium experiments possibly from the experimental cell or helium fill line.

Other heat leaks such as wires to the payload touching warmer surfaces or insufficient thermal radiation shielding can severely increase base temperatures though they usually don't make refrigeration impossible.

1. Monitoring Dilution Refrigerator Performance

The most common failure is impedance blockage. The following are typical signs of an impedance blockage:

- (a) Slowly but steadily increasing base temperature is an early sign. (This assumes that the refrigerator was working well at one point.) If the system is not stressed significantly, the fridge should be able to keep a steady temperature.
- (b) Too-low, or a drop in, still pressure. The proper operating value of the still pressure is generally no less than about 35 mTorr. This will be quite stable if the system is not stressed and the fridge is operating well. A drop below about 20 mTorr is the clearest indicator of an impedance blockage. Proceed to step 4a: Suspending Circulation.
- (c) High pressure on G1. This should also be steady, a few inches above minimum. If it climbs above 15" (about 0.5 atm) the circulation should probably be shut off — certainly if it approaches 1 atm. Proceed to step 4a: Suspending Circulation.

2. Monitoring 1K Pot Performance

The most common cause of blockage in the 1K Pot is pumping on it when the 1K Pot siphon is not immersed in the helium bath, e.g. if the bath is allowed to boil off too low before transfer. This will cause frozen air snow to be pulled in and block flow. Signs are too-low 1K Pot pressure and rising 1K Pot temperature.

3. Automatic System Emergency Responses

(a) Alcatel Circulation Pump Automatic Shutoff

Two conditions can cause the circulation pump to shut off

- i. Interruption in flow of cooling water. If this is the cause it should be possible to fix it and restart circulation again.
- ii. High pressure at pump exhaust (8 PSI is the threshold). This commonly occurs if fridge failure happens during a time when the fridge is not being monitored.

(b) **High-Pressure Relief Valves in the GHS and 1K Pot Pump Line**

If cold fluids are allowed to warm and expand in enclosed spaces there is always the danger of bursting the vessel in which they are confined. In the GHS are pressure-relief valves that allow mixture to escape into Dump #2 if the pressure exceeds 15 PSI. (Normally Dump #2 should be kept empty for this reason.) There is also an 80 PSI pressure-relief valve on Dump #2 that opens to the room, and a similar valve on the 1K Pot pumping line. See figure A.2.

4. Handling Emergencies

(a) **Suspending Circulation**

This should be done if G1 is too high (more than 20" above minimum) or if base temperature or still temperature is 1K or higher or if an impedance blockage has occurred.

You must stop circulation and open the fridge up to the dump so that the mixture does not expand in the fridge lines.

- i. Close V1 and V1A on the GHS.
- ii. Open V11, V7, and V9.
- iii. If you wish to try to circulate again, follow procedure in section A.7: Cool to Base.

(b) **If 1K Pot Stops Working, or if LHe Bath Level Falls Too Low**

Signs of 1K Pot failure include low 1K Pot pressure and rising temperature. The 1K Pot will certainly fail if the bath level falls too low. If there is a delay in getting liquid helium, and you know that you will miss a transfer, follow this procedure to shut down 1K Pot.

- i. First shut down circulation as in step 4a above.
- ii. Close all three 1K Pot pumping valves.
- iii. Open the 1K Pot needle valve a few full turns. 1K pot pressure should quickly reach 1 atm and 1K Pot temperature will go to 4 K.

5. Preventative Measures Checklist

- (a) LHe bath should not be allowed to get too low. About 8" above the siphon should be the minimum. Boiloff rate is approximately 1"/hour so plan accordingly. If transfers are done to 32" or higher you should be able to do

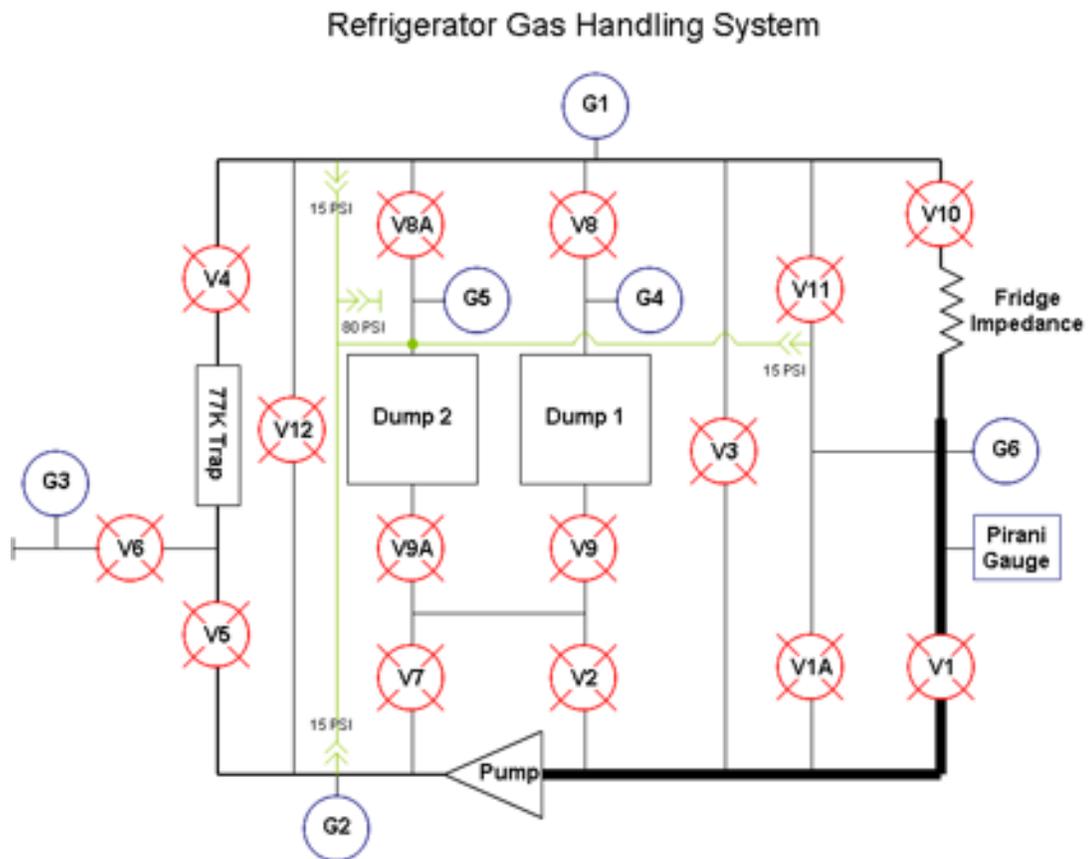


Figure A.2: Little Blue gas-handling system as in figure A.2 with locations of pressure-relief valves shown.

transfers once every 24 hours. Doing transfers at the same time every day and to about the same height is best.

- (b) 77K Trap should be filled daily.
- (c) 77K Trap should probably be cleaned periodically. Many runs are short, less than a week at base temperature, so this is often not a major concern. But weekly cleanings are probably a good idea to prevent buildup of contaminants that could cause a blockage. To clean the trap when the fridge is running:
 - i. Close V1 and V1A.
 - ii. Close V10. Still pressure will rise. Keep an eye on G6 — don't allow it exceed 10" above minimum.
 - iii. Close V5.
 - iv. Open V7, V9, and V3. Wait 10 minutes.
 - v. Close V3.
 - vi. Remove 77K Trap from its Dewar, move Dewar out of the way, and hang trap on its hanger on the GHS.
 - vii. Heat trap with heat gun pointing about 2" away from trap for 10 minutes.
 - viii. Turn off heat gun and wait another 5 minutes.
 - ix. Record G1 pressure in lab book.
 - x. Attach a metal hose from a rough pump (WGHS or turbo station rough pump) to the GHS port (near V6) and pump out the line.
 - xi. Open V6 to pump out the trap contents for 5 minutes.
 - xii. Close V6 and remove hose.
 - xiii. Return trap to its Dewar.
 - xiv. You can now cool to base again following the procedure in section A.7: Cool to Base. (Read the introductory paragraphs to get the GHS in the correct initial state.)

A.9 Warm up.

When a run is finished proper procedures must be followed for an orderly warm-up and removal of payload. The most important considerations are recovering the $^3\text{He}/^4\text{He}$ mixture and making sure that no condensed gases are allowed to expand in sealed volumes.

At this point it is assumed that the fridge is in ordinary circulation mode: GHS valves V10, V1, V1A, V5, and V4 are open and the Alcatel circulation pump is on.

1. Stop circulation and pump mixture out of 77K Trap.
 - (a) Close V1 and V1A on the GHS.
 - (b) Close V5 and V10.
 - (c) Open V7 and V9 (if not already open).
 - (d) Open V3 and wait about 5 minutes.
 - (e) Close V4 and V3.
 - (f) Remove 77K Trap from its Dewar and hang it from its hooks to let it warm to room temp.
2. Open fridge lines to dump.
 - (a) Open V12, V10, and V11.
3. Stop 1K Pot.
 - (a) Close all three 1K Pot Pumping Valves.
 - (b) Open the 1K Pot Needle Valve a few full turns. 1K pot pressure should quickly reach 1 atm and 1K Pot temperature will go to 4 K.
 - (c) Close valve connecting the 1K Pot pump exhaust line to the recovery system. This is located in the northeast corner of the lab, behind the K100 fridge, on a vertical copper pipe near the big gray pipes.
 - (d) Turn off the 1K Pot pump. This can be done from inside the lab by pressing the STOP button on the switchbox labeled "5 in".
 - (e) Open the valve to vent the 5-inch pumping line. This is a small red valve in the northwest corner of the lab near the gas cylinders (verify that its on the same 5-inch pumping line that leads to the fridge). Open it just enough to hear a hiss of air.
4. Pump mixture from fridge to dump.

- (a) If at all possible wait until base temperature and still temperature are close to 4 K (at least above 1 K). Otherwise pumping mixture to the dump will proceed very slowly. This will typically be overnight if you are beginning when the fridge is at base temperature. It is best to proceed while the fridge is still below 5 K, although warmer temperatures are acceptable.
- (b) Close V12.
- (c) Open V1A. Wait until G6 reads low (less than 2 above minimum).
- (d) Open V1 $\frac{1}{4}$ turn.
- (e) Wait until all the mixture is in Dump #1. This can be determined by watching G4 — make sure it reaches as high as it was before the run (around 21" above minimum, though this may vary since we sometimes adjust the mixture.) Still pressure should fall to about 10 mTorr and base and fridge temperatures should be about 4 K. (Still Pressure may oscillate somewhat as the mixture is removed.)
Note: if the fridge is still cold enough this can take a very long time. If necessary it can be left pumping to the dump all night, although it is probably better to restore the GHS to its state at the end of step 2, then wait and try pumping later. If desired, the mixing chamber and still can be heated with as much as 10 mA to speed recovery.
- (f) When all mixture is recovered to Dump #1, close V1, V1A, V7, and V9. If possible this should be done well before fridge temperatures reach 20 K so that any hydrogen that might be condensed in the fridge lines is not pumped to the dump.
- (g) Turn off Alcatel circulation pump and close the two valves for the cooling water to it.

5. Fast Warming Technique

This technique can be skipped if you are not in a hurry. If you do skip it, warm-up can take several days.

- (a) Wait until the LHe level in the cryostat is too low to measure with the thumper (usually 2 days after the last LHe transfer).
- (b) Attach a hose from one of the He supply lines to the WGHS, at W3 for example. Be sure to bleed the line before attaching it.
- (c) From another port on the WGHS, W2 for example, attach a metal hose. Attach to the port on the front of the cryostat jacket. (Make sure that the valve on the front of the jacket is closed.) It may be necessary to connect two hoses.
- (d) Use the WGHS pump (valve W8) to pump out this line to 0.0 Torr on the MKS1 gauge (greater than 50 μ A on the TC1 gauge).

- (e) Let between 15-30 Torr of He gas into the WGHS. Open the valve on the front of the cryostat and slowly let the gas into the cryostat jacket.
 - (f) Record the time and how much gas was let into the jacket. Make sure that the total amount of gas let into the jacket does not exceed 1 atm (760 Torr).
 - (g) You should notice the temperature start to rise. Also after a minute or so you should notice that the gauge on the helium recovery system is rotating faster. After a few minutes the temperature will fall again and the pressure will get quite low (cryopumping).
 - (h) When the TC1 gauge gets back to $\geq 50 \mu\text{A}$ and the gauge on the helium recovery system slows down close the valve on the front of the cryostat jacket.
 - (i) You can now repeat steps 5e through 5h until the pressure on the digital gauge no longer falls to zero. This is an indication that all the LHe in the bath has boiled away. At this point the temperature will continue to increase. It should warm to 300 K overnight.
6. At 300 K, let N_2 gas into fridge lines.
- (a) Wait until fridge temperatures reach 300 K. Thermometer calibrations are poor at this temperature so it is not easy to judge exact temperature; it is best to have the LabVIEW VI monitor temperature continuously and wait until it has been flat for at least several hours.
 - (b) Make a reminder to pump out the cryostat jacket with the turbo pump before the next run.
 - (c) Attach a metal hose from WGHS port W2 to the GHS port near V6.
 - (d) Open W2 and W8 to pump out line to GHS port. 77K TRAP MUST BE WARM!
 - (e) Open V6 and V4 on the GHS.
 - (f) Close WGHS valve W8 (pump valve).
 - (g) Open the N_2 cylinder in the back of the room. Attach a rubber hose from a N_2 gas supply line to WGHS port W1.
 - (h) Watching the WGHS pressure on MKS1, let N_2 into the WGHS through the port until the pressure is around 740-760 Torr. G1 and G6 should read about 1 atm.
 - (i) Close V10 and V11 on the GHS.
 - (j) Open the WGHS valve W8 (pump valve). Pump until WGHS pressure on the TC1 gauge is greater than $50 \mu\text{A}$.
 - (k) Close V6 and valve W2 on the WGHS and remove the hose from the GHS.
7. Remove fridge from cryostat.

- (a) Close valve to Helium recovery system on the wall near the WGHS and remove rubber hose from cryostat bath exhaust port.
 - (b) Disconnect 1K Pot pumping line from 1K Pot port on fridge. Cover both flanges with KF-flange blanks. Rest the 1K Pot pumping hose on top of the large fridge pumping line.
 - (c) Disconnect small fridge return line from fridge return port. Cover flange on port with a KF-flange blank. Let the metal hose hang behind the fridge rack.
 - (d) Remove the screws to disconnect the large still pumping line from the still pumping port. Use the short hex key in the box on top of the fridge rack. Push the pumping line out of the way. Cover the still pumping port by screwing the brass cover onto it. It is OK to use just a few screws and not tighten them much. Cover the pumping line with the black rubber stopper (this will be easier after the fridge has been removed).
 - (e) Disconnect the wiring cables from the connectors at the top of the fridge. Allow these to hang behind the fridge rack.
 - (f) If anything is attached to the IVC port, close the IVC valve, remove the hose, and cover the IVC port flange with a plastic KF-flange cover.
 - (g) Remove the screws from the fridge flange which attach it to the cryostat and place them in the box on top of the fridge rack.
 - (h) Remove fridge using rope and pulley system and lower it onto the fridge rack hanger using essentially the same procedure as in section A.2 steps 2 through 5. Be careful that the o-ring at the top of the cryostat does not catch on the fridge as it is rising causing it to fall into the cryostat.
8. Open IVC and remove payload.
- (a) Let gas into IVC. If the fridge is room temperature and if the payload is not sensitive to exposure to air, simply remove the cover from the IVC port and open the valve. Otherwise, fill the IVC with a little more than 1 atm (760 Torr) of clean N₂ gas. (Remember to open the valve gradually if the IVC contents are fragile.)
 - (b) Remove all but two screws from the lower IVC piece. Loosen the remaining two screws so that the heads are about 1/4 inch above the flange. Insert two other screws into the threaded push-holes in the flange and tighten these until the lower IVC piece comes loose. Then remove it entirely.
 - (c) Remove top IVC piece in a similar way.
 - (d) Remove copper 1K shield.
 - (e) Disconnect wiring from payload.
 - (f) Remove payload.

A.10 Cleaning mixture.

These instructions are for cleaning the $^3\text{He}/^4\text{He}$ mixture before a run. Over time some air is expected to contaminate the mixture. This must be removed from the mixture before it is introduced to the cold refrigerator lines. It is removed by cycling it through the 77K Trap.

This procedure is typically performed after section A.3 Pump Out Fridge Circulation Lines, but it can also be performed when the fridge is not connected to the GHS.

It is very important to understand what you are doing in each step of this procedure so that you do not accidentally lose any $^3\text{He}/^4\text{He}$ mixture. Note that between runs the mixture is generally stored in Dump #1 but is also in the pump lines behind and in front of the Alcatel circulation pump.

GHS refers to the Refrigerator Gas-Handling System, most of which is located on a stand north of the refrigerator. (WGHS refers to the tubing mounted on the wall, which allows access to the rough pumps in the Pump Alley.)

1. Put GHS into proper initial state.
 - (a) Check that G1 shows vacuum (minimum). If it doesn't, find out why!
 - (b) Close V10 and V11 if not already.
 - (c) Open V4 if not already. G1 should not increase.
 - (d) Check that all other GHS valves are closed except V4.
 - (e) Put the 77K Trap into its Dewar and fill Dewar with LN_2 , if not already. (See section refsec:pumpoutfridge step 4 for instructions.)
 - (f) Open V5 and V8. This lets mixture into the trap. G1 should increase.
2. Start the Alcatel circulation pump if it's not already running. This is the orange pump near the door in Pump Alley.
 - (a) Locate the two copper water pipes leading from the right side of the pump. Follow these back to the cooling water manifold and open the two valves to flow cooling water through the pump.
 - (b) Press the green START button on the wall above the pump (or outside the door on the wall). Note: pump will not start if cooling water is not flowing.

Refrigerator Gas Handling System

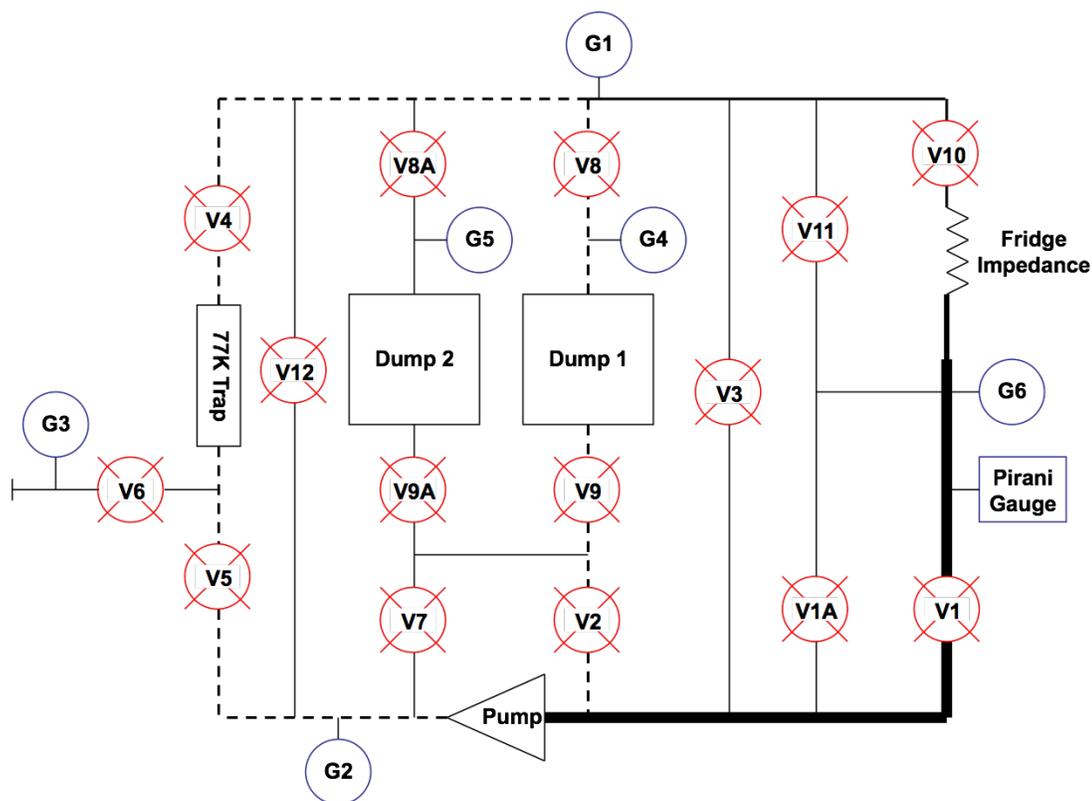


Figure A.3: Little Blue gas-handling system (GHS). The dashed lines indicate the cleaning cycle circulation path. Circulation is regulated with valve V2.

3. Circulate mixture through 77K Trap.
 - (a) Open V9.
 - (b) Open V2 a little more than $\frac{1}{4}$ turn. You will be able to feel a pulsation in one of the flexible metal hoses leading to the GHS on the right. You should also see frost build up on one of the tubes from the 77K Trap.
 - (c) Note the start time of the cleaning in the log book.
 - (d) Monitor the pressures on G1 and G2. If G2 exceeds G1 by 5" or more, proceed immediately to step 4.
 - (e) Let mixture circulate for 1 hour. (Sometimes 30 or 45 minutes is done instead. Occasionally when the mixture is particularly dirty for some reason, an overnight cleaning at a very low circulation rate is done.)
 - (f) Note duration of cleaning in log book.
4. Pump mixture out of trap back to the dump.
 - (a) Close in order V2, V8, and V5.
 - (b) Open V7 and V3.
 - (c) Wait 10 minutes.
 - (d) Close V3, V7, and V9. G1 should be at minimum.
5. Warm up 77K Trap.
 - (a) Lift 77K Trap out of its Dewar.
 - (b) Move Dewar out of the way and hang trap on the hooks on the GHS rack.
 - (c) Position heat gun on wooden platform so that it is pointing at the bottom of the trap a distance of about 2 inches.
 - (d) Turn on heat gun and leave on for 10 minutes.
 - (e) Turn off heat gun and wait an additional 5 minutes.
 - (f) Make a note of the pressure on G1 in the log book.
6. Pump out contaminants from 77K Trap.
 - (a) Attach a flexible metal hose from the GHS port near V6 to either the WGHS or the turbo station pump.
 - (b) Pump out the line to the GHS port for a few minutes. (G3 should go down, although at the time of this writing the needle has a large offset.)
 - (c) Open V6.
 - (d) Wait 5 minutes. G1 should fall to minimum.

- (e) Close V6.
 - (f) Close valve to WGHS or to the turbo station pump and remove hose from GHS port.
7. Return 77K Trap to Dewar.
- (a) Take 77K Trap off hooks, move its Dewar into place and lower trap into Dewar.
 - (b) Refill Dewar with LN₂.
8. Repeat as necessary.
- (a) The pressure on G1 from the step 5f should be less than 2". If it is not, return to step 1 and repeat the procedure.
 - (b) If the amount of contaminants removed does not decrease with each cleaning, troubleshooting is necessary. It is likely in that case that there is a leak somewhere in the system.

A.11 Liquid helium transfer.

These instructions are for Daily Transfers. [Modified instructions for Initial Transfer are in brackets.]

Throughout these instructions, “Dewar” refers to the LHe storage Dewar, and “cryostat” refers to the blue refrigerator bath Dewar.

Room locations are given from the perspective of someone entering the room from the door at the bottom of the stairs (south end).

If you are using a new Dewar, move Dewar onto dolly by locking dolly, tilting Dewar, and maneuvering onto dolly. Unscrew the cap from the top of the Dewar and screw on the o-ring adapter. Attach fill tube onto the side gas port (blue valve). [Initial Transfer will probably require a new Dewar — it requires at least 50 liters of LHe.]

(Note: currently we are not set up to leave the Dewar attached to the recovery system between transfers. These instructions assume that.)

1. Position Dewar dolly by fridge.
 - (a) Roll dolly to the right side of the fridge platform, such that the corner of the dolly is 6 inches past the front corner of the fridge platform.
 - (b) Lock the dolly.
2. Raise Dewar, and position Dewar.
 - (a) Turn crank on dolly until top port of Dewar is approximately level with top port of cryostat.
 - (b) Position the Dewar until the horizontal distance between the Dewar port and the cryostat port is 31 inches. (This is the width of the transfer tube at the top. Positioning the Dewar properly prevents stresses on the transfer tube.)
3. Measure cryostat level with thumper. [Initial Transfer: cryostat is empty, so skip this step.]
 - (a) Locate thumper, usually found hanging from the wire rack on the wall near the fridge or on the fridge rack.
 - (b) Remove stopper from top port of cryostat.
 - (c) With your finger over the rubber at the top of the thumper, lower the bottom of the thumper into the cryostat until you hear the helium start to boil.
 - (d) Hold it there for a few seconds, then lower and raise the thumper until you feel a sharp change in the frequency and/or amplitude of the vibrations at the top of the thumper.

- (e) Narrow down the level until you can repeatedly move the thumper up and down a few mm and feel a sharp change in the vibration.
 - (f) Mark the level on the thumper by attaching a clip to the tube at the level of the cryostat port.
 - (g) Measure the length to the nearest .25 inch from the point on the thumper marked "SIPHON" to the bottom of the clip.
 - (h) Record this level on the Liquid Helium Transfer Log hanging on fridge rack.
4. Attach lengthener tube to transfer tube
- (a) Locate the transfer tube, usually found hanging from the wire rack on the wall near the fridge.
 - (b) Locate the lengthener tubes (3 different lengths) usually found on the wall by the WGHS.
 - (c) Based on the last measured Dewar depth, choose a lengthener tube that will extend at least 3 below LHe level in Dewar when the transfer tube is inserted (Allow 4 inches or so for the top of the transfer tube to extend above the Dewar port.) If it is a new Dewar, the shortest tube should be fine. [Initial Transfer: use the medium-length tube.]
 - (d) Attach lengthener tube tightly to threaded end of transfer tube. (You don't want it to fall off into the Dewar!)
5. Insert transfer tube. Note time.
- (a) [Initial Transfer: Attach the cryostat exhaust to the He recovery system. First disconnect the rubber exhaust hose from the recovery tubing near the wall. Then attach the other end to the cryostat exhaust.]
 - (b) Locate the o-ring adapter for the top port of the cryostat, usually left attached to the transfer tube. Remove it, being careful not to drop the o-ring and metal o-ring holder ring. [Initial Transfer: it should already be on the cryostat fill port.]
 - (c) Remove the stopper from the cryostat port and attach the o-ring adapter. Don't tighten completely.
 - (d) Open the top port of the Dewar and close the pressure relief valve.
 - (e) Lower the long end of the transfer tube into Dewar, inserting far enough that the brass section is inserted into the o-ring and it is in far enough to be below the liquid level. Make sure short end is not too close to you. Hold it here until a bluish plume comes out the short end. [Initial Transfer: you can skip this step and just insert the transfer tube simultaneously into the Dewar and the cryostat.]

- (f) Raise the transfer tube and re-lower it, short end into cryostat, until it is about 4 inches above either port. Note the time.
 - (g) Tighten the o-ring adapters on both the Dewar and the cryostat.
 - (h) [Initial Transfer: Attach the hose leading from the cryostat exhaust to the He recovery system port on the wall and open the valve to the recovery system.]
 - (i) Record transfer start time on Liquid Helium Transfer Log.
6. Pressurize Dewar.
- (a) Open the helium gas cylinder at the left rear corner of the room (up the stairs).
 - (b) Locate a rubber hose extending from one of the blue valves near the wall behind the fridge on the lab He supply.
 - (c) Open the blue valve and flush helium from the hose for about 5 seconds, then pinch off the hose.
 - (d) Open the side port on the Dewar a little, just enough to flush helium from it, then attach the hose to the adapter.
 - (e) Remove the pincher from the hose to pressurize the Dewar. Open Dewar side port valve completely.
7. Determine that transfer is finished. Note time.
- (a) Whether a transfer is complete can be determined by one of several methods:
 - i. By recovery gas flow speed. The gas meter near the south door will show an increasing rate as the transfer continues. When the rate reaches 1 cubic foot in 5-6 seconds, the transfer should be nearly complete.
 - ii. By how stiff or frosty the exhaust hose from the top of the cryostat is. If it is stiff all the way to the copper recovery tubing, it should be within a few minutes of complete.
 - iii. By time. Helium transfers normally take about 10 to 15 minutes. This time might depend on the level in the Dewar, the setting of the regulator valve on the helium gas cylinder, etc. Consult log for previous times. [Initial Transfers take approximately one hour (no less than 45 minutes).]
 - iv. [Initial Transfer: ignore above techniques. Monitor the base temperature. When it reaches 4.2 K it will suddenly level off. Wait 15 minutes (20-25 minutes for the older IVC used for helium experiments) after this occurs — transfer is finished at that time.]
 - (b) Note the time and record it on the Liquid Helium Transfer Log.

8. Stop pressurizing Dewar. Release pressure.
 - (a) Close the side port to the Dewar, the blue valve for the rubber hose, and the helium gas cylinder.
 - (b) Pinch off hose and remove from side port of Dewar.
 - (c) Open the pressure relief valve on the Dewar.
9. Remove transfer tube.
 - (a) Loosen the o-ring adapter at the top port of the Dewar.
 - (b) Completely unscrew the o-ring adapter at the top of the cryostat. (The cryostat o-ring will get stuck on the transfer tube - leave it on the transfer tube as you remove it.)
 - (c) Put on gloves and remove the transfer tube. (You can temporarily hang it on the fridge rack, but it should be stored on the wire rack on the wall.)
10. Measure cryostat level and Dewar level.
 - (a) Use the thumper to measure the cryostat level using the same procedure as in step 3. This should be at least 33 inches to last 24 hours.
 - (b) Use the thumper to measure the Dewar level, except you should measure the distance from the bottom of the clip to the bottom end of the thumper. (This is the distance below the Dewar port to the LHe level.) Close the top port valve on the Dewar.
 - (c) Record both measurements on Liquid Helium Transfer Log.
11. Insert rubber stopper in fridge transfer port.
 - (a) After the release of gas from the top of the fridge has quieted down a little, insert rubber stopper in fridge transfer port.
12. Return Dewar to standby location.
 - (a) Lower Dewar dolly platform using crank.
 - (b) Unlock Dewar dolly.
 - (c) Roll Dewar dolly to standby location.
 - (d) Check that pressure relief valve is open.
13. Fill 77K Trap Dewar.
 - (a) Pour LN₂ from the Dewar next to the GHS rack into the funnel attached to the GHS rack. When the LN₂ in the 77K Trap Dewar starts to overflow, stop pouring.

14. Double Checks

- (a) Make sure the pressure relief valve is open on Dewar (at least one valve should always be open whenever Dewar is left for more than 10 minutes or so).
- (b) Make sure helium gas cylinder is closed (top valve on cylinder).

15. [Initial Transfer: Go back to section A.5 step 6.]