UNIVERSITY of CALIFORNIA
Santa Barbara

Results from the Cryogenic Dark Matter Search Using a Chi Squared Analysis

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Physics

by

Joel Sander

Committee in charge:

Professor Harry Nelson, Chair
Professor Claudio Campagnari
Professor Mark Srednicki

December 2007

The dissertation of Joel Sander is approved:

Claudio Campagnari

Mark Srednicki

Chair: Harry Nelson

December 2007

Results from the Cryogenic Dark Matter Search Using a Chi Squared Analysis

Joel Sander

To my mom and dad

Acknowledgements

All great accomplishments are the product of the hard work of many people. While this thesis is hardly one of the seven wonders of the world, it certainly would never have been completed without the help and support of many people a few of whom are mentioned here.

I'd like to thank my family: Dad, Mom, Jan, Laura, Elena, and Tim. You've been there for me since before I learned to spell 'cat'. Thanks yet again! Verne, Joanne, Tim, Mary, and Garrett Gish, you also deserve a huge thanks! You've been family to me during this entire grad school thing. Thanks also to friends from homegroup, basketball, and church/SBCC that encouraged me and helped me keep my sanity. Marie, thanks for your love and patience.

Special thanks to my advisor Harry Nelson for patiently teaching me something of what it means to be a scientist and a leader. Your ability to reduce complex ideas to their simple core is special; I'll be forever trying to copy it. You were a far better advisor than I had hoped to have. Ray Bunker and Rupak Mahapatra, you both deserve enormous thanks for answer my many dumb questions. Working with you was a joy; I'm lucky to have you as friends. Thanks also to Sam Burke for teaching me so much as we worked together on the trigger electronics.

Many others in CDMS also deserve mention. I especially wish to thank Richard Schnee for all his gentle instruction and constant attention to my analysis, Long Duong for making Soudan a pleasure, Dan Bauer for his constant

efforts for and coordination of Soudan, Walter Ogburn for teaching me Settlers, and Angela Reisetter for her unflagging efforts in somehow obtaining a result and writing a paper with me. Many other people have labored with me at Soudan and elsewhere: Dan Akerib, Mike Attisha, Jim Beaty, Paul Brink, Blas Cabrara, Dan Callahan, Debbie Ceder, Clarence Chang, Rodney Choate, Jodi Cooley-Sekula, Mike Crisler, Mike Dragowsky, Ron Ferril, Dave Hale, Don Holmgren, Bryan Johnson, Wayne Johnson, Sharmila Kamat, Susanne Kyre, Bruce Lambin, Samuel Leclercq, Lap-Yan Leung, Vuk Mandic, Nader Mirabolfathi, Robert Nelson, Stan Orr, Chris Savage, Rich Schmitt, Dennis Seitz, Bruno Serfass, Gensheng Wang, and others. Without your prodigious efforts, this thesis would not exist. Thank you!

Last, I want to thank Jesus Christ who gave me perspective and so much more throughout.

Curriculum Vitæ

Joel Sander

1996–1998	Undergraduate Researcher, Department of Phyisics and Astronomy, Butler University, Indianapolis, IN
1998	B.S. *cum laude* in Physics, Butler University, Indianapolis, IN
1998–1999	Teaching Assistant, Department of Physics, University of California, Santa Barbara, CA
1999–2007	Graduate Student Researcher, Department of Physics, University of California, Santa Barbara, CA

Selected Publications

CDMS Collaboration (D. S. Akerib *et al.*), *Limits on spin-dependant WIMP-nucleon interactions from the Cryogenic Dark Matter Search.* Phys. Rev. D **73** 011102 (2006); astro-ph/0509269.

CDMS Collaboration (D. S. Akerib *et al.*), *Limits on spin-independant WIMP-nucleon interactions from the two-tower run of the Cryogenic Dark Matter Search.* Phys. Rev. Lett. **96** 011302 (2006); astro-ph/0509259.

CDMS Collaboration (D. S. Akerib *et al.*), *Exclusion limits on the WIMP-nucleon cross section from the first run of the Cryogenic Dark*

Matter Search in the Soudan Underground Laboratory. Phys. Rev. **D72** 052009 (2005); astro-ph/0507190.

CDMS Collaboration (D. S. Akerib *et al.*), *First results from the Cryogenic Dark Matter Search in the Soudan Underground Laboratory.* Phys. Rev. Lett. **93** 211301 (2004); astro-ph/0405033.

CDMS Collaboration (D. S. Akerib *et al.*), *New results from the Cryogenic Dark Matter Search.* Phys. Rev. **D68** 082002 (2003); astro-ph/0306001.

D. Abrams *et al.*, *Exclusion limits on the WIMP-nucleon cross section from the Cryogenic Dark Matter Search*, Phys. Rev. **D66** 122003 (2003).

Abstract

Results from the Cryogenic Dark Matter Search Using a Chi Squared Analysis

by

Joel Sander

Most of the mass-energy density of the universe remains undetected and is only understood through its affects on visible, baryonic matter. The visible, baryonic matter accounts for only about half of a percent of the universe's total mass-energy budget, while the remainder of the mass-energy of the universe remains dark or undetected. About a quarter of the dark mass-energy density of the universe is comprised of massive particles that do not interact via the strong or electromagnetic forces. If these particles interact via the weak force, they are termed weakly interacting massive particles or WIMPs, and their interactions with baryonic matter could be detectable.

The CDMS II experiment attempts to detect WIMP interactions in the Soudan Underground Laboratory using germanium detectors and silicon detectors. A WIMP can interact a with detector nuclei causing the nuclei to recoil. A nuclear recoil is distinguished from background electron recoils by comparing the deposited ionization and phonon energies. Electron recoils occurring near detector surfaces are more difficult to reject.

This thesis describes the results of a χ^2 analysis designed to reject events

occurring near detector surfaces. Because no WIMP signal was observed, separate limits using the germanium and silicon detectors are set on the WIMP cross section under standard astrophysical assumptions.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This chapter begins by giving a nutshell history of some of the reasons for looking for the dark matter candidate, weakly interacting massive particles or WIMPS, which are the particles that the Cryogenic Dark Matter Search (CDMS) seeks to detect. Next the standard parameterization that is used to describe the universe is outlined to provide a foundation of meaning for the results to be described later in this chapter. After the foundational math has been described, some key tools that were used to find much of the evidence for dark matter are related. The bulk of this chapter is then devoted to describing the various methods by which we understand the universe and the behavior of the dark matter.

1.1 A Short History

The first evidence that much of the matter in the universe is dark, or invisible to the eye, came from Fritz Zwicky's 1933 observations [1] of the Coma cluster. He measured the radial velocities of 8 galaxies gravitationally bound to the Coma cluster, and found that these galaxies moved with velocities indicative of a gravitational force about 400 times greater than the expected potential from the mass inferred from the cluster brightness. If he had a more accurate value of the Hubble constant, the estimated cluster mass discrepancy would have been reduced to a factor of 50. Zwicky believed that the difference was due to unobserved matter and was the first to use the term "dunkle kalte Materie" or cold dark matter. Zwicky did not use the term in the modern sense of nonrelativistic matter that doesn't interact with light. (The first use of the term in the modern sense was by Bond et al. [2] in 1983.)

Little additional progress was made until 1970 when Rubin [3] showed that the rotational velocity away from the center of the spiral galaxy Andromeda surprisingly remained constant as a function of radius. From observing the distribution of visible mass, one would expect the rotational velocity to decrease as the distance from the galactic center increased. Rubin's observations are therefore not well explained by the distribution of visible mass. Within spiral galaxies, the rotational velocity of stars and gas, particularly at high radii, about the galactic center is still not well explained by currently observed galactic mass and Newtonian gravity. In the 73 years since Zwicky's discovery,

the dark matter puzzle has only deepened.

Historically, deviations from the expected motions of astrophysical bodies have been explained by the presence of unobserved or dark matter or by refinements gravitational theory. The explanation of planetary dynamics within the solar system has required both the existence of unobserved masses as well as refinements to Newtonian gravity. Deviations from the expected motion of the planet Uranus led J. C. Adams and U. J. Le Verrier to independently predict the position of a new unobserved mass. On September 23rd 1846, J. C. Galle observed the predicted mass, Neptune, a day after receiving Le Verrier's prediction. Le Verrier's 1859 observation of the unexpectedly large precession of Mercury's perihelion was, however, explained by a refinement to Newtonian gravity, general relativity, rather than by missing mass.

Two proposed explanations for the observed deviations of bodies in clusters and galaxies are dark matter and modified Newtonian gravity. For the dark matter explanation to work, dark matter must be distributed differently than the visible matter. Self-gravitating dark matter which doesn't interact would be naturally distributed in a spherical, nearly isothermal halo, which is the simplest feasible distribution. Measurements of the cosmic microwave background (CMB) and large scale structure (LSS) indicate that a large fraction of the matter of the universe is currently unobserved and nonbaryonic, potentially providing a natural solution to the puzzle. Another potential explanation is that Newtonian gravity is for some unknown reason altered below

a critical acceleration, $a_0 \approx 1 \times 10^{-10} m/s^2$, giving an effective acceleration of $(a_0 G_N M/R^2)^{1/2}$ [4].

A candidate for the particle that comprises the dark matter naturally arises from many extensions to the Standard Model of particle physics, particularly from the concept of supersymmetry. If supersymmetry, a postulated symmetry between half integer and integer spin particles, is restored just above the weak scale, the lightest neutral supersymmetric particle can easily be a massive and long-lived particle that has interactions of a magnitude comparable to the weak scale. Relic weakly interactive massive particles (WIMPs) from the big bang with the properties required by the CMB and LSS measurements could naturally provide the dark matter density needed to explain Zwicky's puzzle.

This thesis presents some of the current motivation for WIMP searches; the detailed configuration setup of an experiment for direct detection of dark matter, the Cryogenic Dark Matter Search (CDMS); the analysis of CDMS data; the resulting current best limit on spin-independent WIMP interactions; and a projection of future CDMS WIMP sensitivity.

1.2 Big Bang Parameterization

Dark matter is believed to have played a significant role in the evolution of the universe. It is instructive to begin with an overview of the cosmology of the early universe. The description below follows a path similar to standard text books on cosmology such as [5] [6]. The Particle Data Group's review of Big-

Bang cosmology [7] is also a particularly useful reference. The assumption that the universe is homogeneous and isotropic leads to the Friedmann-Robertson-Walker metric,

$$ds^2 = dt^2 - a^2(t)\left(\frac{dr^2}{1-kr^2} + r^2(d\theta^2 + sin^2\theta d\phi^2)\right) \tag{1.2.1}$$

where a(t) is the scale factor between comoving and laboratory coordinates and k is a constant which sets the spatial curvature. k = 0 (-1 or 1) for a flat (open or closed) geometry. The scale factor, $a(t)$, is defined by

$$d_{lab} = a(t)d_{co} \tag{1.2.2}$$

where d_{lab} are the laboratory coordinates and d_{co} are the comoving coordinates. For this metric, the time-time component of the Einstein field equations gives the Friedmann equation,

$$\left(\frac{\dot{a}}{a}\right)^2 + \frac{k}{a^2} = \frac{8\pi G_N}{3}\rho_m + \frac{\Lambda}{3} \tag{1.2.3}$$

where ρ_m is the mass energy density and Λ is the cosmological constant or vacuum energy. A cute, non-rigorous derivation [8] of the Friedmann equation can be made using the assumption of Newtonian gravity. The expansion rate of the universe is given by the Hubble constant

$$H = \frac{\dot{a}}{a}. \tag{1.2.4}$$

The Hubble constant has a present day value of $H_0 = 100\ h^{-1}$ km s^{-1} Mpc^{-1} where h is the scaled Hubble parameter. The scaled Hubble parameter was

introduced as a convenient method of expressing the Hubble constant when its exact value was uncertain. Because many measurements depend on the value of the Hubble constant, it is common to quote a value in terms of the scaled Hubble parameter. The current best value for the Hubble parameter was obtained by WMAP's best fit to the cosmic microwave background experiments WMAP, ACBAR, and CBI; the 2dF galaxy redshift survey; and the Lyman α forest data: [9]

$$h = 0.71^{+0.04}_{-0.03}. \tag{1.2.5}$$

Defining the vacuum energy density to be

$$\rho_\Lambda = \frac{\Lambda}{8\pi G_N}, \tag{1.2.6}$$

and the total energy density, $\rho = \rho_m + \rho_\Lambda$, the energy density of a flat (k = 0) universe, called the critical density, is given by

$$\rho_c = \frac{3H^2}{8\pi G_N}. \tag{1.2.7}$$

The current day critical density has a value of $\sim 10h^2\,\mathrm{GeV\ m^{-3}}$ or approximately 5 hydrogen atoms per cubic meter. It is useful to define the energy density of a given component relative to the critical density as Ω_i. The Friedmann equation can be rewritten as

$$1 = \Omega_k + \Omega_m + \Omega_\Lambda \tag{1.2.8}$$

where Ω_k is defined as

$$\Omega_k = -k/(aH)^2. \tag{1.2.9}$$

Much of the effort in cosmology is directed at understanding Eqn. 1.2.8 to understand whether the universe is flat, open, or closed; to understand the various mass densities; and to understand the eventual fate of the universe: eternal expansion, re-collapse, or something in between. Efforts to understand the various critical densities provide the framework on which the dark matter is understood.

The first law of thermodynamics equates the change in internal energy, $d(\rho a^3)$, with the loss of energy due to the decreased pressure as the volume of the universe expands, $-pd(a^3)$. Upon integration this relation can be rearranged to give

$$\rho \propto a^{-3(1+\omega)} \tag{1.2.10}$$

where the pressure is assumed to be related to the density by

$$p = \omega\rho. \tag{1.2.11}$$

ω is assumed to be time independent and has different values for radiation ($\omega = 1/3$), non-relativistic matter ($\omega = 0$), and the vacuum energy ($\omega = -1$). The current measured value of ω for the vacuum energy density is $-0.97^{+0.07}_{-0.09}$ [10]. Applying Eqn. 1.2.10 to the Friedmann equation for an early, radiation-dominated or mass-dominated universe, one finds that $a \propto t^{2/3(1+\omega)}$. For a radiation dominated universe, $a \propto t^{1/2}$ and $\mathrm{H} = 1/2t$. Later, when non-relativistic matter dominates, one finds that $a \propto t^{2/3}$ and $\mathrm{H} = 2/3t$. The dark matter dominates the later evolution of the universe. Understanding the evolution of the universe provides insight into the nature of the dark matter.

1.3 Cosmological Tools

As light takes time to travel to an observer on Earth, observation of distant objects provides information about the early universe which is essential for understanding how the universe evolved to its present state. Three basic tools used to measure distances are the redshift, the distance modulus, and the Tully-Fisher relation. The use of the redshift is so fundamental that distance to an object is often quoted in terms the object's redshift.

1.3.1 Redshifts

Understand the evolution of the universe requires observations at different epochs. Light from a distant object is observed at a frequency, ν_{obs}, that is lower than the emitted frequency, ν_{em}. Redshifting of the frequency of emitted light, ν_{em}, to a lower observed frequency, ν_{obs}, on large distance scales is due to the expansion of the universe. If the frequency of the emitted light is known, the redshift can be used to determine distance to the emitter. The redshift can be defined to be

$$z \equiv \frac{\nu_{em} - \nu_{obs}}{\nu_{obs}} \simeq \frac{\Delta v}{c} \tag{1.3.1}$$

where Δv is the relative velocity between the emitting object and the observer. The approximation used, $\frac{\lambda_{obs}}{\lambda_{em}} \simeq 1 + \beta$, in the second half of Eqn. 1.3.1 is valid when the expansion velocity is nonrelativistic. From the current value of the Hubble constant, it is clear that the expansion of the universe in the current

epoch is nonrelativistic and the approximation is valid. For redshifts due to the expansion of the universe, $1 + z = a_0 \ / \ a$ where a_0 is the present day value of a. Using Eqn. 1.2.2, the Hubble Law can be written

$$v = \dot{d}_{lab} = \frac{\dot{a}}{a} d_{lab} = H d_{lab}. \tag{1.3.2}$$

Objects further away are receding from us faster than closer objects.

1.3.2 Distance Modulus

A second tool for determining distance is the distance modulus. Distances are often quoted using the distance modulus, μ_B, rather than the distance. The distance modulus is defined to be the difference between the apparent and absolute magnitudes of an object, and can easily be derived. The absolute magnitude of an object is defined to be

$$M \equiv -2.5 log(L) + constant \tag{1.3.3}$$

where L is the luminosity. The flux from a distant star is

$$f = L/(4\pi d_L^2) \tag{1.3.4}$$

where d_L is the distance to the star and L is the star's luminosity. The apparent magnitude is defined to be

$$m \equiv -2.5 log(f) + constant = M + 5 log(d_L) + constant. \tag{1.3.5}$$

The distance modulus is

$$\mu_B \equiv m - M = 5 log(d_L) + C. \tag{1.3.6}$$

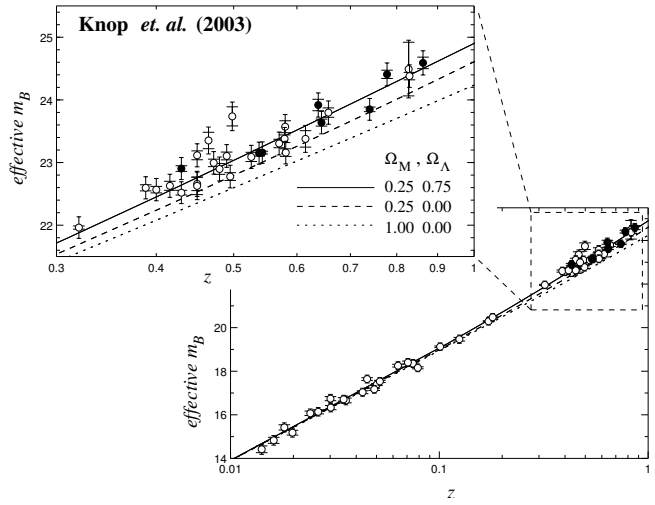

Figure 1.1: Hubble diagram of the effective distance modulus, m_B, vs. log of the supernovae redshift. Filled circles represent supernovae observed by the Hubble Space Telescope (HST). The larger error bars include the 0.17 magnitude intrinsic dispersion in addition to the measurement uncertainties. Figure from [11].

The constant, C, is defined to be the distance at which the apparent and absolute magnitudes are equal. A typical distance is 10pc which gives $C = 25$ for distances given in Mpc. Figure 1.1 illustrates the relationship between distance and redshift by showing a corrected or effective distance modulus versus redshift for supernovae. The distance modulus is usually corrected for Milky Way and host galaxy extinction and redshifting of light out of the observed band.

1.3.3 Tully-Fisher Relation

One method for measuring the distance to a galaxy is through observation of the width of its galactic absorption lines. While an absorption line from an individual star in the galaxy will be either red-shifted or blue-shifted depending on whether the star is moving away from or towards the observer, the observed galactic spectral absorption line will be from the entire galaxy. The velocity dispersion of stars in the galaxy will widen the line. For an isothermal distribution, which is a reasonable approximation at larger radii, the rotational velocity is related to the velocity dispersion

$$v_c = 2^{1/2}\sigma \tag{1.3.7}$$

where σ is the one dimensional line-of-site velocity dispersion. More massive galaxies are more luminous and have a higher rotational velocity. The Tully-Fisher relation [12] for spiral galaxies relates the rotational velocity to the luminosity

$$v_c = 220\frac{L}{L_*}^{0.22}\frac{km}{s} \tag{1.3.8}$$

in the infrared 2.2 micron K-band. L_* is a characteristic galaxy luminosity and is $10^{10}\,h^2\,L_\odot$ in the visual band where $L_\odot$ is the solar luminosity. An L_* galaxy is a bright galaxy and has a luminosity that is roughly similar to that of the Milky Way. The Tully-Fisher relation describes the motion of stars primarily under the gravitation influence of dark matter.

1.4 Early Evidence for Dark Matter

To describe the observed structure of the universe and the resulting evidence for dark matter, we will begin by describing the early universe and trace its early history. Each short description will include or be followed by a discussion of the role of dark matter in the evolution of the universe. The dark matter has a profound effect on the universe from early times through the present time, providing numerous indirect evidences of the existence of the undetected dark matter.

1.4.1 Primordial Nucleosynthesis

Prediction

The hot primordial universe synthesizes the light elements from hydrogen. Observation of light-element fractional abundances can be compared with the prediction from primordial nucleosynthesis. The comparison yields an estimate of baryonic matter density, Ω_b. Two useful references on the big-bang nucleosynthesis are [5] and [7].

In the early universe, the equation for a particle's number density is given by the Boltzmann equation

$$\frac{dn}{dt} = -3Hn - < \sigma v > (n^2 - (n^{eq})^2) \tag{1.4.1}$$

where $< \sigma v >$ is the thermally averaged annihilation cross section times the velocity and n^{eq} is the thermal equilibrium number density. For elastically-

scattering massive particles, the equilibrium number density is given from the Maxwell-Boltmann distribution as

$$n^{eq} = g \left(\frac{mT}{2\pi} \right)^{3/2} e^{-m/T} \tag{1.4.2}$$

where m is the particle mass, T is temperature, and g is the number of degrees of freedom available to the particle. g is temperature dependent and falls with the expansion of the universe. See Jungman et al. for further information on g including a plot of g for temperatures between 100 MeV and 1 TeV [13]. When the temperature falls such that $T \approx m$, the particle leaves equilibrium and is considered to be frozen-out. Since massive particles freeze out earlier, WIMPs would be expected to freeze out before baryonic matter allowing structure to begin to form earlier. Section 1.4.2 explains why the early freeze-out of WIMPs is important for explaining the formation of observed structure.

The WIMP density is not required to remain constant after freeze-out. WIMP interactions subsequent to freeze-out can reduce their density. Many supersymmetric models require a stable lightest particle that is its own antiparticle for which a WIMP would be a natural candidate. The WIMP density would then decrease due to WIMP-WIMP annihilations. Section 1.4.2 will describe the role that primordial WIMPs play in the formation of large scale structure.

Comparison of light-element abundance ratios with the ratios expected from primordial nucleosynthesis gives an estimate of the baryonic mass of the uni-

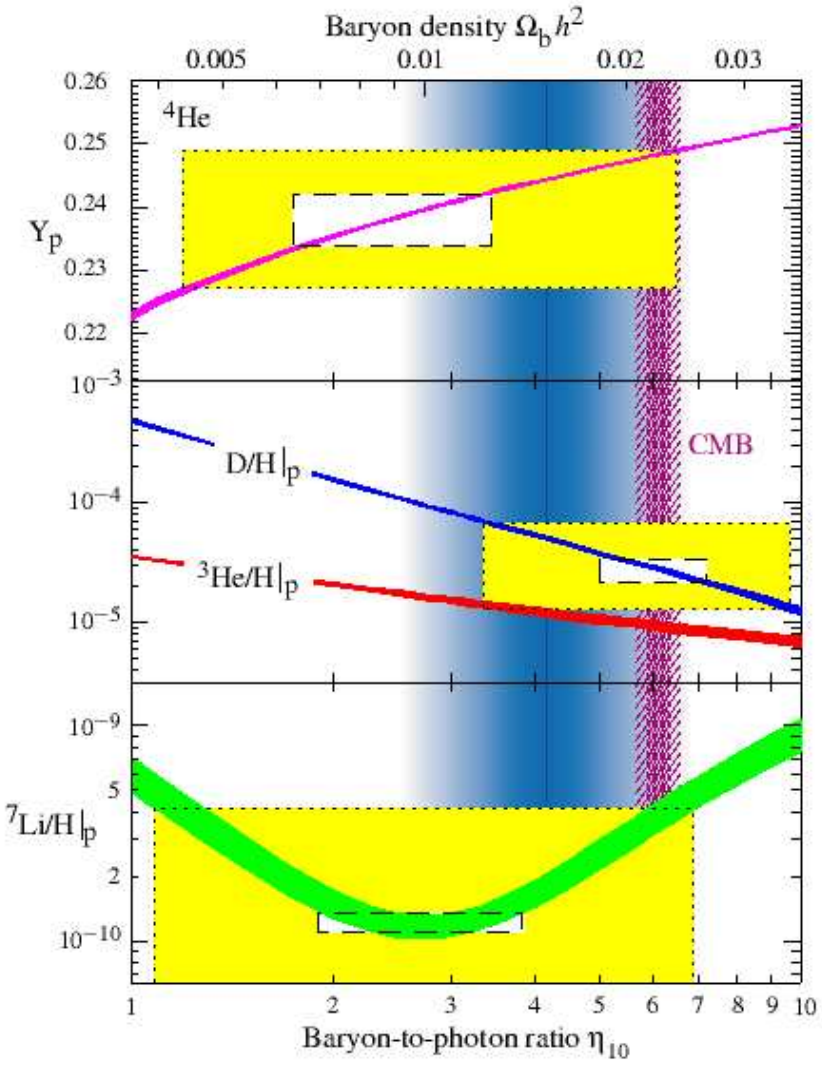

Figure 1.2: The abundances of the light elements as predicted by the standard model of big-bang nucleosynthesis. Smaller boxes indicate the 2σ statistical errors on the observed light element abundances while larger boxes indicate the 2σ statistical and systematic errors added in quadrature. The narrow vertical band gives the cosmic microwave background measure of the cosmic baryon density, η. Figure from [7].

verse. From Eqn. 1.4.2, the ratio of neutrons to protons is

$$\frac{n}{p} \simeq e^{-(m_n - m_p)/T}. \tag{1.4.3}$$

When the temperature falls below the mass difference of ∼1MeV, the neutron-proton ratio begins to freeze into place with a value of approximately 1/6. As the universe continues to cool, ambient photons no longer have sufficient energy to break nuclei apart and the light elements hydrogen, deuterium, helium, and lithium are formed. Free neutrons, which have a half-life of 887s, can decay after freeze-out, reducing the neutron-proton fraction for light element synthesis to approximately 1/7.

The expected ^{4}He abundance can be easily estimated. Nearly all free neutrons combine with protons to form the lightest stable element ^{4}He. The ^{4}He abundance is typically quoted in terms of the primordial mass fraction, Y_p. Y_p can be estimated by a simple counting argument

$$Y_p = \frac{4n_{He}}{n_n + n_p} = \frac{2(n/p)}{1 + n/p} \simeq \frac{1}{4} \tag{1.4.4}$$

where n_{He}, n_n, and n_p are the ^{4}He, neutron, and proton number densities and a neutron proton ratio of 1/7 is used.

Production of light elements depends on baryon density which is usually expressed relative to the photon density

$$\eta \equiv \frac{n_B}{n_\gamma}. \tag{1.4.5}$$

where n_γ is obtained from the CMB temperature. For convenience, η_{10} is

defined to equal $10^{10}\eta$. Ω_b is related to η by

$$\Omega_b \simeq \frac{\eta_{10} h^{-2}}{274} \tag{1.4.6}$$

After accounting for statistical and systematic error as well as theoretical uncertainties, the values of η, predicted by the various abundances of the light elements, are in reasonable agreement within a 2σ confidence region of $3.4 \leq \eta_{10} \leq 6.9$. See Fig. 1.2 for the preferred values of η and therefore Ω_b from the abundance of light elements. Light element abundances are usually quoted relative to the hydrogen abundance. Clearly, Ω_b is only a small fraction of Ω_{total}.

Synthesis of light elements begins with deuterium. Deuterium production will be balanced by photon-dissociation until the ratio of photons above deuterium's binding energy of $\Delta_D = 2.23\,\mathrm{MeV}$ per baryon, falls below 1 at temperatures below $\simeq 0.1\,\mathrm{MeV}$. The ratio at freeze-out is given by $\eta^{-1} e^{-\Delta_D/T}$.

Observation

Light elemental abundances are observed at times far after production, complicating attempts to compare observation with prediction. Stars synthesize the light elements into heavier elements, so observations of light elements abundances are made in high redshift systems in which low amounts of heavy elements such as carbon, nitrogen, oxygen, and iron have been observed.

An estimate of Ω_b from big-bang nucleosynthesis expected value for η using Eqn. 1.4.6 is $\Omega_b = (0.012 - 0.025)h^{-2}$. WMAP's measurement of Ω_b to be

0.024 ± 0.001 [19], which translates to $\eta_{10} = 6.58 \pm 0.27$, is in reasonable agreement with the above value of η from big-bang nucleosynthesis. Lower values of η_{10} down to 5.2 can be obtained, if one varies the shape of the CMB power spectrum [20] [21].

Observation of ^{4}He in clouds of ionized gas within metal poor systems such as dwarf galaxies can be used to constrain Y_p. Small amounts of ^{4}He and heavy elements: carbon, nitrogen, and oxygen are observed in metal poor systems [17]. Since stellar production of ^{4}He is correlated with heavy element production, an extrapolation to no heavy elements can be made. Upon extrapolation, an estimate of $Y_p = 0.238 \pm 0.002(stat) \pm 0.005(sys)$ [18] is obtained.

Measurements of deuterium are particularly interesting since there is no known astrophysical production mechanism for deuterium [14]. The most precise measurements of the deuterium abundance fraction, from high-resolution spectra of low-metallicity quasar absorption systems, give a deuterium to hydrogen ratio of $(2.78 \pm 0.29\ (\text{stat})) \times 10^{-5}$ [15] [16]. The spread in measured deuterium abundance fractions is larger than expected from statistical error indicating significant systematic error.

Baryonic matter is only a small contributor to the mass-energy density budget given by the Friedmann Eqn. 1.2.8. However, because WIMPs freeze-out earlier, they would be expected to contribute significantly more to the mass-energy budget than baryons.

1.4.2 Large Scale Structure

The growth of density fluctuations

One of the challenges facing cosmology is to explain how the highly-homogeneous, early universe developed the structure observed today. Larger structures, such as clusters, are believed to have grown from mergers of older, smaller structures. Measurements of the cosmic microwave background show temperature fluctuations, $\delta T/T$, of a few 10^{-5} at the time of recombination which is about five orders of magnitude below the current day level. Radiation and baryonic matter alone are insufficient to provide sufficient growth in the density fluctuations from the time of recombination to the current time. The problem is solved by assuming that most of the mass is weakly interacting, since weakly interacting particles decouple earlier allowing time for larger density fluctuations to develop. Figure 1.3 shows a schematic view of how the density fluctuations grow in comoving coordinates when the presence of dark matter is included or excluded.

To understand how structure develops, it is useful to begin by studying how a small density fluctuation in the early universe develops following [6] or [23]. A perturbation, $h_{\alpha\beta}$, to the standard Friedmann space time model is assumed, such that

$$ds^2 = dt^2 + (\delta_{\alpha\beta} - h_{\alpha\beta})dx^\alpha dy^\beta. \tag{1.4.7}$$

If the matter acts as an ideal fluid with density, ρ, and pressure, p, the stress-

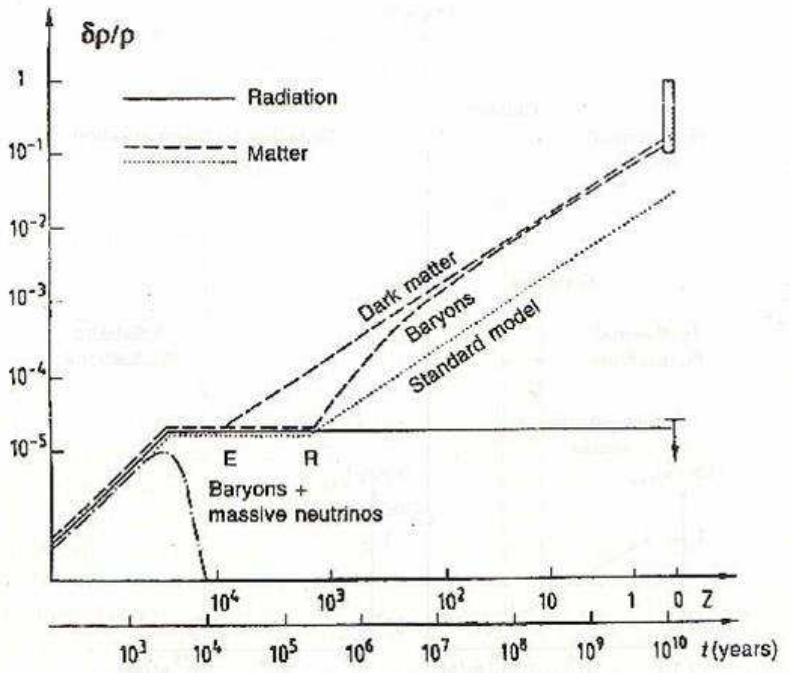

Figure 1.3: Growth of density fluctuations in comoving coordinates as a function of time. The figure shows how density fluctuations could grow under three different scenarios. The right side of the plot, at $z = 0$ shows the current density fluctuations. The density fluctuations in the early universe are adiabatic, primordial fluctuations which grow until they contain the horizon mass. The solid line shows the history of density fluctuations for a radiation-dominated universe. The dotted line shows the growth of baryonic density fluctuations which only begin to grow after recombination, R. The density fluctuations for weakly interacting dark matter, which begin to grow at an earlier time, are shown by the upper dashed line. The lower dashed curve shows the growth of baryonic density fluctuations under the influence of the density fluctuations from the weakly interacting dark matter. The dot-dashed curve illustrates how neutrino dark matter would wash out the structure. Figure from [22].

energy tensor is given by

$$T^{ij} = (\rho + p)u^i u^j - g^{ij}p \tag{1.4.8}$$

where $u^k = (1, 0, 0, 0)$ is the four-velocity in the comoving coordinates. The pressure is assumed to be related to the density as in Eqn. 1.2.11. The time-time component of the Einstein gravitational field equation yields an equation describing the evolution of a density perturbation, $\delta = (\rho - \rho_b)/\rho_b$:

$$\ddot{\delta} + 2\frac{\dot{a}}{a}\dot{\delta} = 4\pi G_N \rho_b (1+\omega)(1+3\omega)\delta. \tag{1.4.9}$$

where ρ_b is the mean background mass density. For a flat $\Omega = 1$ early universe, the average mass energy density is approximately equal to the critical density given by Eqn. 1.2.7. From Section 1.2, $a \propto t^{2/(3+3\omega)}$ where ω equals 0 for non-relativistic matter and 1 for relativistic matter. Equation 1.4.9 becomes

$$\ddot{\delta} + \frac{4t^{-1}}{3(1+\omega)}\dot{\delta} - \frac{2}{3}\frac{1+3\omega}{1+\omega}t^{-2}\delta = 0. \tag{1.4.10}$$

Late baryons

The question of whether the observed fluctuations in the cosmic microwave background can be explained by matter density fluctuations can now be addressed. Equation 1.4.10 suggests a solution of the form $\delta \propto t^n$. The two solutions are $n = -1, \frac{2}{3}\frac{1+3\omega}{1+\omega}$. From Section 1.2, the solution for growing density fluctuations is $\delta \propto t^{2/3} \propto a$ for non-relativistic matter and $\delta \propto t \propto a^2$ for relativistic matter. Decoupling of photons and baryons occurs around $z \sim 10^3$,

allowing baryonic structure to begin developing. If non-relativistic matter that coupled to light is the dominant matter component in the universe, one would expect current day density fluctuations to be of the order $10^{-5} \times 10^3 = 10^{-2}$, far below the observed value. If the dominant matter component of the universe was relativistic and coupled to light, the density fluctuations at decoupling should have grown to be of the order 10^{-4} to 10^{-3}. These rough calculations illustrate how difficult it is to explain the growth of structure since the early universe with visible matter alone.

Weak-scale formation

Matter that interacts with a weaker cross section would decouple earlier and have more time to produce density fluctuations. A natural question is if the dominant component of the matter interacted with a weak-scale cross section would it be expected to produce the observed mass density fluctuations. A massive, weakly-interacting particle will fall out of equilibrium when $m \approx T$ from Eqn. 1.4.2. The number density at freeze-out, n^f, becomes

$$n^f \approx g \left(\frac{m_w^2}{2\pi} \right)^{3/2} \tag{1.4.11}$$

where g is the number of available degrees of freedom. For temperatures between a few fractions of a GeV and 1TeV, g ranges from approximately 60 to 110 [13]. The relic density is inversely proportional to the cross section, so

$$\sigma_w^{-1} \propto \frac{n_w^0 m_w}{\rho_c} = \frac{n_w^f m_w}{\rho_c} \left(\frac{a_w^f}{a_0} \right)^3 . \tag{1.4.12}$$

where n_w^0 and m_w are the current number density and mass of the weak-scale particles and ρ_c is the critical density, and n_w^f is the number density at freeze-out. The weak force is approximately 10^9 times weaker than the electromagnetic force. Taking the ratio of weak-scale and baryonic cross sections gives

$$a_w^f = 10^{-3/2} \left(\frac{m_b}{m_w} \right)^{1/6}. \tag{1.4.13}$$

For particles with masses within several orders of magnitude of a proton this gives $a_w^f \sim 10^{-4} - 10^{-5}$. Density fluctuations should have grown by a factor of $\sim 10^4 - 10^5$ since weakly-interacting, massive particles are expected to have decoupled which is in agreement with observation.

A rough WIMP mass estimate

A crude estimate of the mass of a relic weakly-interacting particle, that is the dominant form of matter in the universe, can be made. The current mass density, ρ_0, is $0.3\rho_c$. Using Eqn. 1.4.11 for the WIMP number density at freeze-out, one finds that

$$m_w = \frac{\rho_0}{n_0} = \frac{\rho_0}{n^f \epsilon} \left(\frac{a_w^f}{a^0} \right)^3 \approx (2\pi)^{3/8} \left(\frac{0.3\rho_c}{\epsilon g} \right)^{1/4} \left(\frac{a_0}{a_w^f} \right)^{3/4} \tag{1.4.14}$$

where ϵ is the WIMP survival fraction since freeze-out. Assuming that the number of WIMPs since freeze-out remains constant, the mass of a WIMP is then estimated to be of order 10^3 GeV.

1.4.3 Cosmic Microwave Background

The early universe is believed to be homogeneous and isotropic. When photons and baryons decoupled at $z \sim 10^3$, the photons were free to travel (mostly) without interaction. Observed photons from the cosmic microwave background (CMB) incident from any direction have traveled an equal distance since decoupling and originate on a surface of a sphere centered on the observer. This surface is termed the surface of last scattering. Radiation from the CMB is well described by a blackbody spectrum with a temperature of 2.725±0.001K [24]. At the time of decoupling, the temperature was much higher, but the expansion of the universe redshifted the photons. The photon density, Ω_γ is then well known and can be considered to be fixed when the effect of varying the densities Ω_m, Ω_b, and Ω_Λ is considered later in this section.

Initial Fluctuations

While the early universe at the time of decoupling was nearly homogeneous and isotropic, it was not quite perfectly so. There will be small nearly Gaussian fluctuations that the self-gravitation of weakly interacting dark matter will grow. The small fluctuations lead to potential wells on the surface of last scattering causing small variations in observed temperature from one part of the sky to the next.

Spherical Harmonics - Lowest Multiples

Traditionally the anisotropies on the surface of last scattering are expressed in terms of an expansion of spherical harmonics.

$$T(\theta, \phi) = \sum_{l,m} a_{l,m} Y_{l,m}(\theta, \phi). \tag{1.4.15}$$

Power versus the log of the multiple, l, is traditionally plotted. The power at each l is proportional to $(2l + 1)\mathrm{C}_l$ where $\mathrm{C}_l \equiv < |a_{lm}|^2 >$ and is also proportional to the temperature variation squared, $(\Delta T/T)^2$. The angular scale is related to the multiple, l, by the rough relationship $\theta \sim 2/l$. The zeroth spherical harmonic, $a_{0,0}$, to which more recent difference mapping experiments like WMAP are insensitive, give the mean temperature of the CMB blackbody spectrum. The first spherical harmonic, $l = 1$, is dominated by a Doppler shift due to the motion of the solar system [25].

Spherical Harmonics - Higher Multiples

Variations in the observed higher multiples of the observed CMB are interpreted as being 1) a result of primary fluctuations at the time of decoupling and 2) a result of secondary fluctuations from processes that affect the light as it propagates to an observer. This discussion will begin by outlining how some of the physical parameters that are relevant to a dark matter search affect the observed power at the higher multiples. The theoretical calculation of these affects on the power spectrum is not trivial and will only be briefly outlined

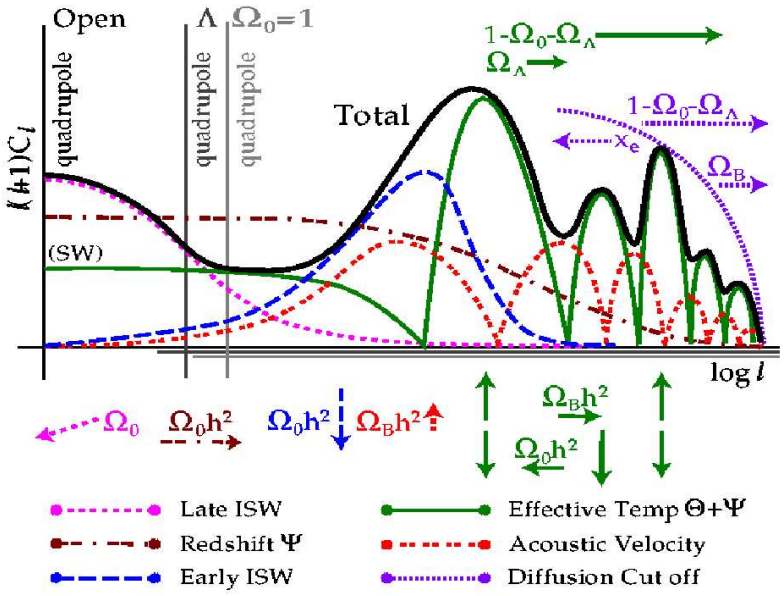

Figure 1.4: A nearly-standard expected CMB power spectrum. The plot is not quite standard in that the x-axis is logarithmic. The effect of the integrated Sachs-Wolfe effect (ISW), the acoustic peaks, and the damping of higher multiples are described in the text. Figure from [27].

below. A fuller treatment can be found in Wayne Hu's thesis [26]. Figure 1.4 illustrates many of the affects of the physical parameters on the CMB power spectrum, and is a useful diagram for much of what is described below.

The initial condition for the photon-baryon fluid is assumed to be adiabatic with the fluid in an anticompressed (coolest) state. The first acoustic peak corresponds to the mode in which the fluid compressed once before recombination. The second peak corresponds to the mode in which the fluid oscillated once back to the uncompressed state, etc. The position of the first peak is highly

sensitive to the curvature of the universe. Increasing Ω_k in Eqn. 1.2.8 an open geometry, makes the observed peaks appear at smaller angles or higher multiples. The location of the peaks is also slightly sensitive to Ω_Λ. The vacuum energy increases the expansion rate of the universe, decreasing the comoving distance to the surface of last scattering. The anisotropies then appear at slightly larger angular scales.

Acoustic peaks at higher order multiples around $l = 10^3$ are damped, because the surface of last scattering has some thickness. The temperature of the observed photons from a point on the surface is the mean of the temperature of photons from the near to far regions of the surface. The averaging washes out the smaller scale detail as shown in Fig. 1.4.

The baryon density affects the acoustic peaks. Increased baryon density pulls the photon-baryon fluid deeper into the potential wells shifting the baseline of the oscillation to higher ΔT and slightly increasing the amplitude of the oscillation. Since the power at a given multiple is proportional to the temperature variation squared, the acoustic oscillations are sensitive only to fluctuations from the mean temperature and increasing Ω_b enhances the compression peaks and diminishes the anticompression peaks. The suppression of the second peak relative to the first and third peaks allows Ω_b to be measured.

The third peak is sensitive to the ratio of the dark matter density, Ω_m - Ω_b, to the radiation density Ω_γ. In the very early universe when the smallest scale modes begin oscillating, the radiation density dominates the matter density. In

the absence of dark matter, the gravitational potentials form from fluctuations in the radiation density and decay at decoupling. The decay of the potential leads to much larger oscillations particularly at small scales. The presence of dark matter damps all the acoustic peaks, but increases the ratio of the higher peaks to the first peak.

Secondary Effects

As photons from the CMB travels, intervening processes can change the energy of the photons changing the CMB power spectrum. Three such effects are the integrated Sachs-Wolfe (ISW) effect, the Sunyaev-Zeldovich (SZ) effect, and foregrounds. The ISW effect is due to the gravitational potentials, through which the photons pass on their way to an observer, not being constant. This is true for both the initial potential well as well as for additional potential wells that a photon may pass through. The ISW effect manifests itself in two ways: the early ISW effect and the late ISW effect. The early ISW effect is due to the photon density shortly after decoupling being small but not quite negligible resulting in a small decay in the potential wells. This slightly redshifts the higher energy photons from the overdense regions and broadens the low multiple shoulder of the lowest acoustic peak. The late ISW effect occurs if $\Omega_\Lambda > 0$ or if $\Omega_k \neq 0$. Since the vacuum energy and curvature of the universe effect large scales, the late ISW effect increases the power at large angles. Interpreting observations of the CMB are made more difficult by the presence of foregrounds

including dust and radio point sources.

CMB photons can inverse Compton scatter off of gas changing the energy spectrum of the CMB. The SZ effect, which is more significant when the CMB photons pass through clusters, takes two forms. First, clusters are not stationary objects. Photons scattering with the cluster gas will be blue shifted in the direction of the cluster velocity. This is called the kinematic SZ effect. Second, inverse Compton scatters with the hot cluster gas changes the CMB energy spectrum by energizing low energy photons. Clusters subtend small areas of the sky, so the SZ effect is limited primarily to high precision measurements of small scales.

Significant amounts of microwave radiation are emitted by the Milky Way and extragalactic sources [28]. Microwaves emitted by foreground sources typically have a different frequency spectrum than the CMB. The effect of foregrounds can be treated by observing portions of the sky away from the Milky Way and by making observations in multiple frequency bands.

CMB Results

Figure 1.5 shows the CMB power spectra as measured by several recent experiments. Using a 6 parameter fit to the WMAP data alone where Ω_{tot} is

assumed to be 1, the following best fit values are obtained [29]:

$$\begin{aligned} \Omega_m h^2 &= 0.127^{+0.007}_{-0.013} \\ \Omega_b h^2 &= 0.0223^{+0.0007}_{-0.0009} \\ h &= 0.73^{+0.03}_{-0.03} \end{aligned} \tag{1.4.16}$$

Notice that a significant component of nonbaryonic dark matter is required by the fit. Including data from other CMB experiments as well as supernova data described in Section 1.4.4 further tighten the constraints. Combining the WMAP data with data from the Supernova Legacy Survey constrains Ω_k to be $-0.015^{+0.020}_{-0.016}$ [29], consistent with a flat universe.

1.4.4 Type Ia Supernovae

If one can find astronomical objects of known luminosity, or standard candles, with known lines whose redshift can be measured to give the distance, then the distance modulus can be used to test if the flux observed from a standard candle is as expected. Using standard candles at a variety of redshifts allows one to make a measurement of the cosmology of the universe that is complementary to the CMB measurement and is from a more recent time period.

Supernovae are divided into two classes based on their observed lines. Type I supernovae do not have hydrogen Balmer lines while type II do. Type Ia

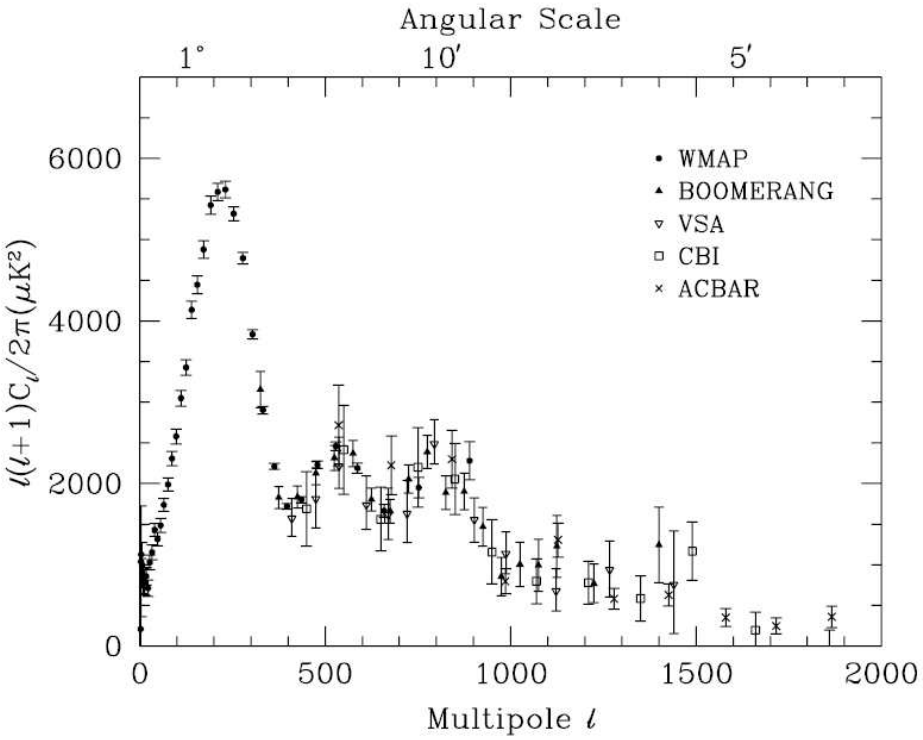

Figure 1.5: The CMB power spectra as measured by several recent experiments: WMAP [29], BOOMERANG [30], VSA [31], CBI [32], and ACBAR [33]. Figure from [27].

supernovae, which can be recognized by a silicon line that is present at 615nm, have a very different formation history than type Ib, Ic, and type II supernovae which occur when massive stars exhaust their hydrogen supply. Type Ia supernovae are formed from white dwarfs stars. White dwarfs are small stars that have exhausted their hydrogen supply and are supported against gravitational collapse by the electrons' Fermi degeneracy pressure. In the presence of a companion star, they can accrete matter until their mass rises to the Chandrasekhar limit of about $1.4M_{\odot}$. Above the Chandrasekhar limit, sudden gravitational collapse leads to fusion of carbon and oxygen forming silicon. Because white dwarfs explode into supernovae with similar masses and chemical composition, their explosions are expected to be fairly homogeneous. Type Ia supernovae are observed to be highly similar and are considered to be (nearly) standard candles. Figure 1.6 shows type Ia supernovae luminosity as a function of time.

As a standard candle, the observed flux depends on the distance to the supernova and on cosmology experienced by the light traveling to a detector. The flux is related to the supernova's luminosity by

$$F = \frac{L}{4\pi d^2(1+z)^2} \tag{1.4.17}$$

where d is the proper distance to the supernova. The redshift of a supernova can be obtained by measuring the redshift of known spectral lines. Measurement of the redshift of the host galaxy can also be used to obtain or cross-check a supernova's redshift. The distance is usually quoted in terms of the distance

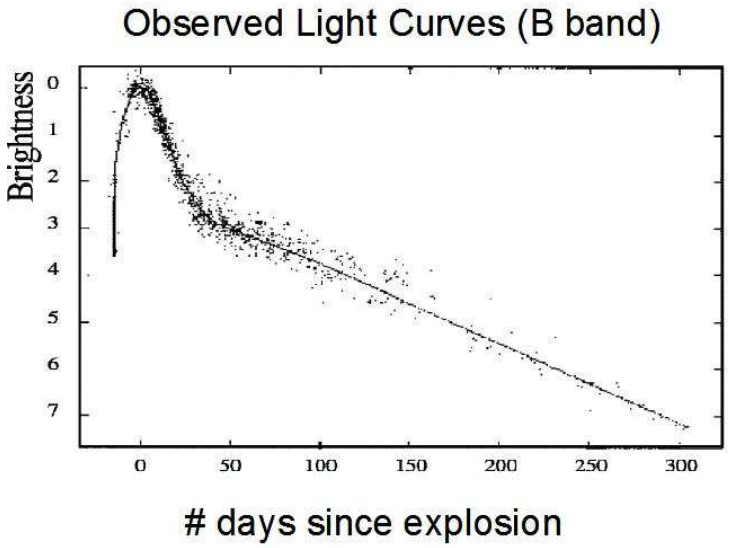

Figure 1.6: Supernovae as standard candles. The figure shows the observed absolute luminosity of supernovae light curves in the B band as a function of time. The supernovae follow a similar light curve with a dispersion of 0.4 magnitude. Figure from [34].

modulus. See Section 1.3.2 for a description of the distance modulus.

In order for measurements of supernovae to provide an interesting cosmological probe, corrections to the observed supernovae's B-band luminosities are required. As supernovae are redshifted, light emitted in the B-band begins to be shifted into the R or I-band. The observed light spectra must be corrected for redshift or K-corrected. After K-correction, supernovae are still not quite ideal standard candles; the magnitude, fall-time, and color of a supernova are correlated, causing a 0.4 magnitude dispersion in supernovae luminosity. More luminous supernovae release their energy more slowly than less luminous su-

pernovae. Accurate color templates from nearby supernovae can be used to correct the correlations. One of the current difficulties in correcting supernovae observations is that complete templates do not exist for all colors. Correlation corrections reduce the dispersion to a magnitude of 0.15.

The measured distance modulus depends on Ω_m, Ω_Λ, and z. Measurements of supernovae out to high redshift are required to break the degeneracy between Ω_m and Ω_Λ. High redshift supernovae are more sensitive probes of different cosmological models, since they have had a greater exposure to the cosmology. Figure 1.7 shows the results of the most recent measurement from the first year of data from the Supernovae Legacy Survey (SNLS). The SNLS best fit result for a flat ΛCDM model is Ω_m is 0.263 ± 0.042(statistical) $\pm$ 0.032(systematic) [35]. Another recent measurement by the Supernova Cosmology Project measured Ω_m to be $0.25^{+0.07}_{-0.06}$(stat) $\pm$ 0.04(identified sys) [36].

1.5 Evidence From Clusters and Galaxies

Additional evidence for dark matter comes from present day observations of mass-to-light ratios, clusters of galaxies, and the rotation curves of galaxies. Some of the earliest and strongest evidence for dark matter comes from observations of clusters of galaxies and the rotation curves of spiral galaxies. Evidence for dark matter within our own spiral galaxy is particularly important for any experiment attempting to directly detect dark matter.

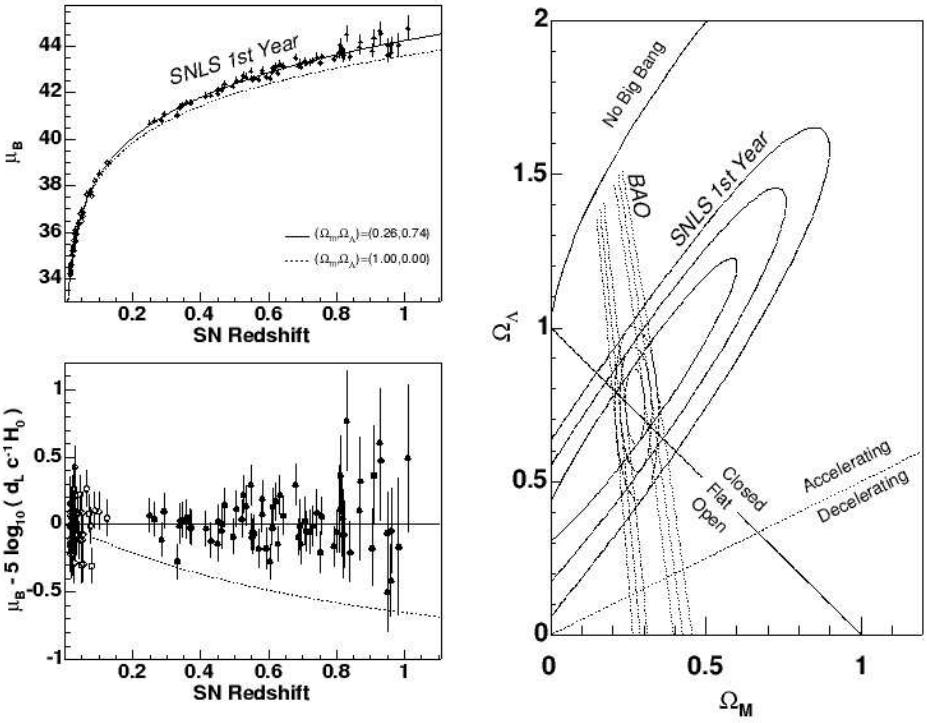

Figure 1.7: Results from the first year of the 5-year Supernovae Legacy Survey (SNLS).The plots on the left show the distance modulus as a function of supernova redshift. The solid line shows a standard $\Omega_m = 0.26$, $\Omega_\Lambda = 0.74$ cosmology, and the dashed line shows a $\Omega_m = 1.0$, $\Omega_\Lambda = 0.0$ cosmology. On the right is shown the 1σ, 2σ, and 3σ confidence contours for the SNLS result for Ω_Λ vs Ω_m. The BAO result is a baryon acoustic oscillations measurement from the Sloan Digital Sky Survey (SDSS). Figure from [35].

1.5.1 Mass-to-Light Ratio

Following reference [6], observed light and a mass-to-light ratio can be used to estimate of the mass density of the universe which can then be compared with the critical density. The galaxy luminosity function, ϕ, is defined to be the mean number of galaxies with luminosity between L and $L + dL$ per unit volume

$$dn = \phi(\frac{L}{L_*})\frac{dL}{L_*}. \tag{1.5.1}$$

L_* is the luminosity of a characteristic galaxy. If the luminosity density of the universe can be found, the mass-to-light ratio can be used to estimate the mass density of the universe. The luminosity function can be approximated by a broken power law or by the more elegant form [37] [38]

$$\phi(y) = \phi_* y^\alpha e^{-y}. \tag{1.5.2}$$

Here ϕ_* and α are constants and $\alpha = -1.28 \pm 0.03$ [39] as measured by the 2dF galaxy redshift survey. By measuring M_*, ϕ_* can be found to be $0.0161e^{\pm 0.8}h^3 Mpc^{-3}$. Integrating the galaxy luminosity function, a luminosity density, j, can be found

$$j = \int_0^\infty L dn = (1+\alpha)!\phi_* L_*. \tag{1.5.3}$$

Most of the luminosity of the universe is due to galaxies with luminosity near L_*. Multiplying the luminosity density by the mass-to-light ratio from a typical L_* galaxy gives an estimate of the mass density of the universe. A

System	$\frac{M_i}{L_i}\frac{L_\odot}{M_\odot}$	Ω_i
Solar neighborhood	5	$0.003h^{-1}$
Rotation curve of NGC 3198	$> 28h$	> 0.017
Cores of elliptical galaxies	12h	0.007
M87 X-ray halo	> 750	$> 0.46h^{-1}$
Groups of galaxies	260h	0.16
Clusters of galaxies	400h	0.25

Table 1.1: Selected mass-to-light ratios relative to the local mass-to-light ratio. $M_\odot$ and $L_\odot$ are the mass and luminosity of the sun. Ω_i shows the mass density of each type of object relative to the critical density. Data from [41].

typical L_* mass-to-light ratio in solar units is $12e^{\pm 0.2}h$ [40]. The mass-density becomes

$$\rho_* = 0.05h^2 GeV m^{-3} \tag{1.5.4}$$

and the density relative to the critical density is $\Omega_* = 0.005$. Since the observed luminosity is dominated by the visible portion of the galaxies, the above density can be taken to be an estimate of the baryonic energy density of the universe.

If most of the mass were baryonic, one might naively expect the mass-to-light ratio of all systems to be similar. Most galaxies are less luminous than an L_* galaxy, so most of the systems are composed of galaxies that are typically less luminosity than an L_* galaxy. These systems have higher mass-to-light ratios which is evidence that these systems contain significant dark mass. The mass-to-light ratio for a number of objects relative to the solar mass-to-light ratio is shown in Table 1.1.

1.5.2 Clusters

Clusters are groups of galaxies, ranging from "poor groups" of tens of galaxies to "rich clusters" with thousands of galaxies. Cluster galaxies are typically old ellipticals and are some of the largest gravitationally-bound systems in the universe. A number of techniques exist for measuring the mass of an entire cluster. The mass of the constituent galaxies is estimated from the cluster luminosity. The mass inferred from the cluster luminosity is then compared with the measured mass of the entire cluster. The mass of the entire cluster can be measured by use of the virial theorem, x-rays, or gravitational lensing. Techniques that measure the mass of the entire cluster give higher mass estimates than expected from cluster luminosity. The additional dark mass could be interpreted as being due to the higher dark matter content. Section 1.4.2 points out that the visible structure appears to form within potential wells of dark matter. It is expected that the deepest wells would attract additional luminous matter. One would also expect the ratio of dark mass to luminous mass to change. If however all cluster mass had the same source, it would be more challenging to explain the higher mass-to-light ratio.

Virial

The virial theorem provides a useful method for estimating the mass of a cluster that has relaxed into gravitational equilibrium. The viral theorem cannot be applied to clusters that have undergone a recent merger or show

other signs of inequilibrium. Clusters typically grow from mergers with other clusters and groups. Unfortunately, most clusters exhibit substructure which is evidence for not having achieved dynamic equilibrium. The virial theorem will overestimate the mass of a cluster that is out of equilibrium. Assuming that a cluster is virialized within some radius, r_{vir}, the mass of a cluster within the virial radius, M_{vir} is given by

$$M_{vir} = \frac{2\sigma^2 r_{vir}}{G_N} \tag{1.5.5}$$

where σ is the velocity dispersion within the virial radius. In practice, the line of sight component of the velocity distribution is measured for galaxies within a cluster. After converting galaxy velocities into galactocentric velocities, the main cluster body is identified by finding the most significant peak in the velocity distribution. Galaxies outside of the peak are rejected as non-cluster members.

The cluster mass estimate typically has a fairly large error, $\sim$40%. One difficulty is estimating the proper virial radius. Velocity anisotropies of galaxy orbits in the central region of a cluster can have a strong effect on the measured line of sight velocity dispersion, but have little effect on the velocity dispersion calculated using the entire cluster[42]. For nearby clusters, the integral velocity dispersion profile for a given radius can be strongly increasing or decreasing in the central cluster regions. At very high radii, the velocity dispersion profile of some clusters shows a strong increase due to the presence of nearby systems. A virial radius is chosen for which the velocity dispersion profile is fairly constant

with radius. The cluster mass is then given by the estimated mass within the virial radius plus a correction, typically 20%, for the cluster not being fully contained with the virial radius. The error on the mass estimate ranges inversely with cluster mass from ∼75% for less massive clusters with mass $< 5 \times 10^{13} M_{\odot}$ to ∼30% for massive clusters with mass $> 5 \times 10^{14} M_{\odot}$.

Since clusters are so large, a reasonable assumption is that the cluster energy density is similar to the energy density of the universe. The mass-density of the universe can then be estimated by

$$\Omega_m = \frac{M}{L}\frac{\rho_L}{\rho_c} \tag{1.5.6}$$

where ρ_L is the cluster luminosity density. A typical cluster value for Ω_m is given in Table 1.1. The mass of the cluster's constituent galaxies can be estimated from the cluster luminosity. Figure 1.8 shows the cluster luminosity versus mass from the Two-degree Field Galaxy Redshift Survey. It is clear that the mass-to-light ratio for a cluster is much higher than one would expect from the visible light of the constituent galaxies.

X-Ray

A cluster's mass distribution can also be probed by observing the temperature of X-rays coming from the gas within the cluster. Observation of the X-ray temperature provides an excellent cross-check of the mass estimate from the virial theorem, since the luminosity from the cluster gas can be observed beyond the virial radius. X-ray data can provide evidence for whether a cluster

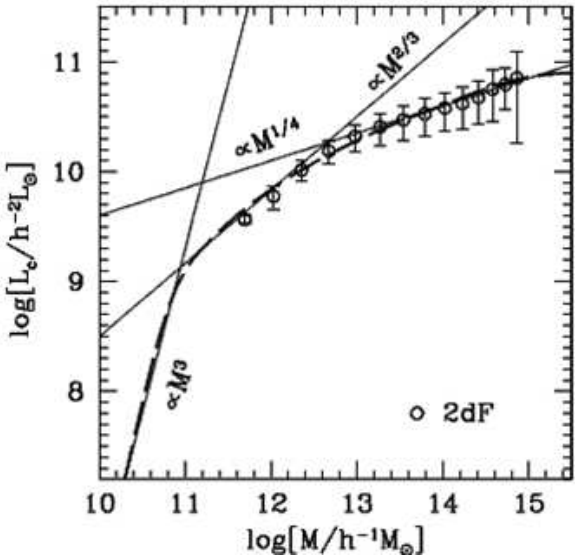

Figure 1.8: The mean central galaxy luminosity, L_c, as a function of system mass, M, for a range of system masses. The data points are from the Two-degree Field Galaxy Redshift Survey, 2dF. The cluster mass-to-luminosity ratio is shown to not be constant for all halo masses and to increase with system mass. Figure from [43].

is in equilibrium [44]. Figure 1.9 shows the temperature of X-rays coming from the Coma cluster. In equilibrium, the gas pressure of a spherical system is balanced by the gravitational force

$$\frac{dp}{dr} = -G_N \frac{M(r)\rho(r)}{r^2}. \tag{1.5.7}$$

Assuming an ideal gas, the pressure equals $nk_BT = \rho k_BT/\mu m_p$ where n is the number density, m_p is the proton mass, and μ is the average molecular weight of ≈ 0.6. Equation 1.5.7 then becomes

$$M(r) = -\frac{k_BTr^2}{G_N\mu m_p}\left(\frac{dlnn}{dr} + \frac{dlnT}{dr}\right). \tag{1.5.8}$$

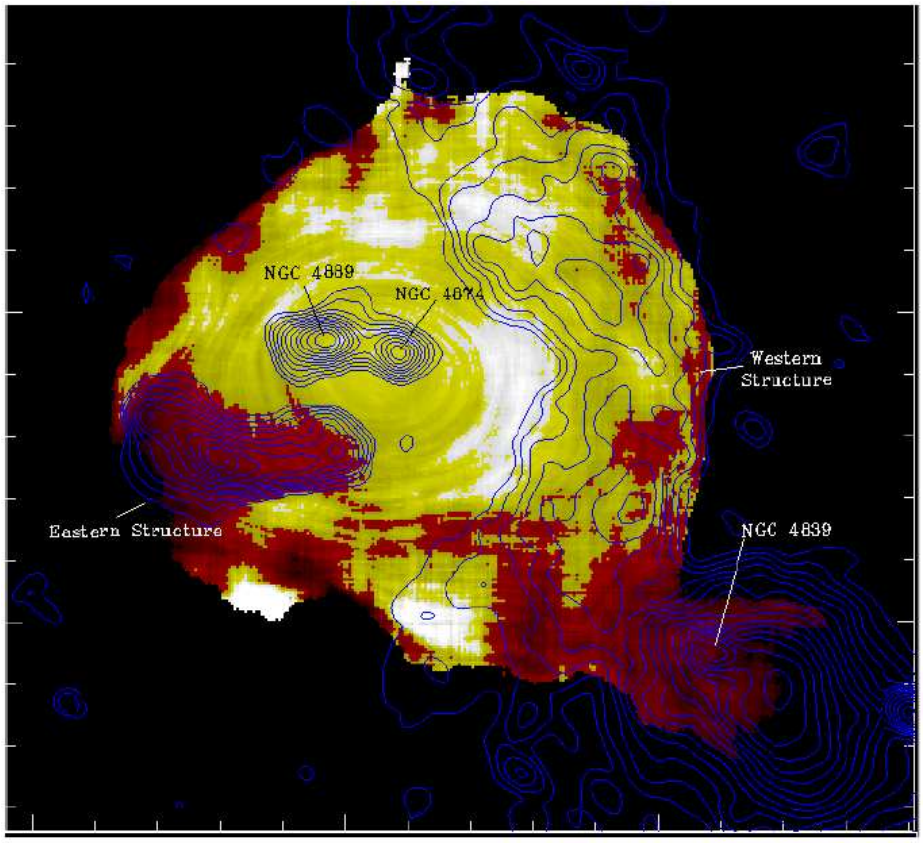

Figure 1.9: The hardness ratio of the Coma cluster as observed by XMM-Newton. The hardness ratio is obtained by dividing the image in the 2-5keV band by the image in the 0.5-2keV band. Assuming an isothermal spectrum, red regions correspond to regions of the cluster with temperatures below 8keV, yellow regions have temperatures above 8keV, and white regions have temperatures above 10keV. The contours give the significance of the residuals. The image shows the sum of all exposures and has been background and vignetting corrected. Dark blue regions correspond to count rates of $0.016\,\mathrm{cnts}\,/\,\mathrm{s}\,/\,\mathrm{arcmin}^2$ while white regions correspond to count rates of $\geq 0.288\,\mathrm{cnts}\,/\,\mathrm{s}\,/\,\mathrm{arcmin}^2$. The vertical axis ranges from 27.1 to 28.6 degrees and the horizontal axis ranges from 195.4 degrees on the left to 194.1 degrees on the right. Figure from [45].

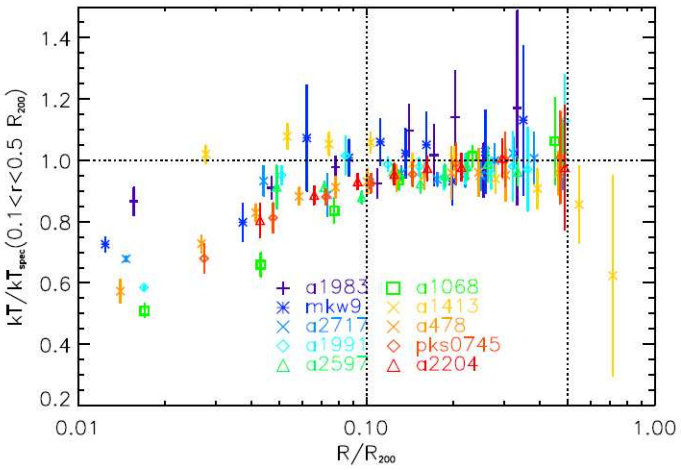

Figure 1.10: The temperature profile of 10 clusters as a function of radius. R_{200} is the radius at which the density falls to 200 times the critical density. The temperature profiles have been normalized to the spectroscopic temperature in $0.1 < r < 0.5\ R_{200}$. Figure from [46].

Figure 1.10 shows that the temperature of cluster gas is roughly constant outside the cluster core. Assuming that the gas density at away from the core is in an isothermal distribution with $n \propto r^{-2}$ the mass contained within a given radius is

$$M(r) \approx 7 \times 10^{13} \left(\frac{kT}{1keV}\right) \left(\frac{r}{1Mpc}\right) M_{\odot}. \tag{1.5.9}$$

The typical observed cluster temperature ranges from 1keV in poor groups to 15keV in rich clusters. Much of the rich Coma cluster emits X-rays in the 8-10keV range and above [45] giving an estimated cluster mass within a $1.8h^{-1}Mpc$ virial radius of $1.4 \times 10^{15}h^{-1}M_{\odot}$ in agreement with the mass estimated by the virial theorem of $\approx 2 \times 10^{15}h^{-1}M_{\odot}$.

Gravitational Lensing

When light from distant galaxies pass through a gravitational potential, the light is bent resulting in the observed image being distorted. The gravitational potential altering the path of the light is referred to as a gravitational lense. Studying the distorted light images can provide information on the gravitational potential of the lense. Gravitational lenses can range in mass from microlensing by brown dwarfs to weakly or strongly lensing rich clusters. Strong lensing by a cluster, which occurs when the observed sources of light are highly distorted, occurs infrequently and requires significant effort to interpret due to its non-linearity. Weak lensing of a cluster, which is more common, occurs when the lensed light sources are slightly distorted, or sheared and cannot be individually determined to have been lensed.

Weak lensing of a cluster is used to determine the mass of the cluster without requiring the cluster to be in equilibrium. Lensing instead requires that a cluster lense many galaxies. Each of the galaxies lensed by a cluster has its own ellipticity that is oriented in a random direction. The lensing cluster shears the light from the lensed galaxy slightly changing the ellipticity. No useful information can be obtained from a single lensed galaxy. The lensing shear can be estimated by averaging over many lensed galaxies. Most clusters lense an area of around 1/5 of a square degree. The population density of sufficiently bright distant blue galaxies available for a cluster to lense is typically 100 per square arcminute [47], however typically less than half of the source galaxies

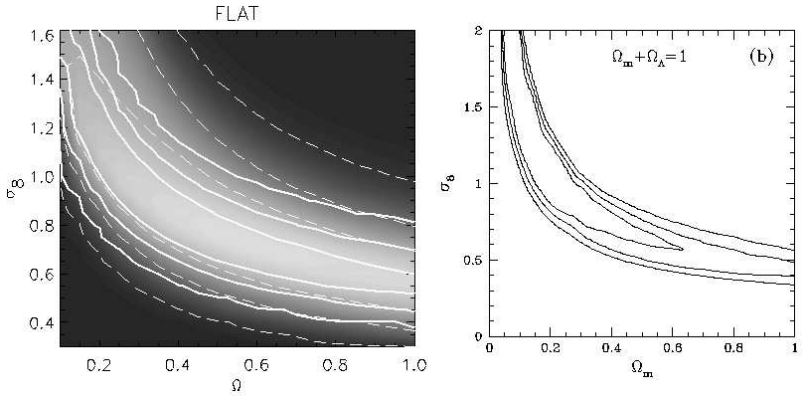

Figure 1.11: Constraints on Ω_m and σ_8 from weak lensing.Left: The VIRMOS-DESCART survey [48]. The dashed and solid lines show the 68%, 95%, and 99.9% confidence regions under different assumptions of source galaxy redshift and marginalization. Right: the RCS survey [49] showing the 1, 2, and 3σ confidence regions. A flat universe has been assumed by both groups. Figure from [50].

are usable.

Measurements of the shear are sensitive to $\sigma_8\Omega_m^\alpha$ where σ_8 is the amplitude of fluctuations of the matter density inside a sphere with radius $8h^{-1}Mpc$ and alpha is determined by the depth of the survey. One recent weak lensing result at the 95% confidence level on the mass-density of the universe is $\sigma_8\Omega_m^{0.52} = 0.46^{+0.05}_{-0.07}$ [49] under the assumption of a ΛCDM model. The exact power of Ω_m depends on the depth of the survey. See Fig 1.11 for plots showing weak lensing results for σ_8 versus Ω_m. The value of σ_8 found by other groups often exceeds the two sigma error for σ_8 quoted above, indicating that the systematic errors are sometimes underestimated. A source of observational

systematics is atmospheric distortion which blurs the image, causing it to appear more circular. Nevertheless, weak lensing provides an important tool that is independent of cluster luminosity of estimating cluster mass. Peter Schneider review [50], *Weak Gravitational Lensing*, is a good source for further information on gravitational lensing.

The Bullet Cluster

The galaxy cluster 1E0657-56, also known as the bullet cluster, is one of the most interesting clusters in the universe and provides unique, direct evidence for the existence of dark matter. The bullet cluster is composed of two merging clusters: a main cluster and a subcluster that appears to have recently exited the core of the main cluster at high speed. The cluster is uniquely ideal for observation because 1) the collision is occurring within the plane of the sky, 2) it is only 3.4 billion light years away, and 3) the subcluster is moving away from the main cluster at the extremely high rate of 4500^{+1100}_{-800} km/s producing a strong, $M = 3.0 \pm 0.4$ [51], gas shock wave behind the subcluster.

Most of the baryonic mass resides in gas which emits light in the form of X-rays. The remainder of the baryonic mass resides in the galaxy's stars. Because the galaxies interact only gravitationally, their path is little affected by the collision. The gas also interacts electromagnetically causing it to slow down and fall behind the galaxies. The bullet cluster has been observed with optical and X-ray telescopes. Weakly interacting dark matter would be expected to

be affected by the collision in much the same way as the stars. The cluster mass distribution has also been probed via weak gravitational lensing. If the cluster mass is dominated by baryonic matter, the lensing should map the gas. But if the mass is dominated by weakly interacting dark matter, the lensing should map the dark matter. Figure 1.12 shows the weak lensing mass contours overlayed on the optical and X-ray images. The right panel shows that the cluster gas is not located in the same location as the cluster mass, while the left-hand panel shows that the mass is distributed similarly to the visible galaxies.

Explanations of the rotation curves of spiral galaxies and cluster observations that replace dark matter with modified gravity have difficulty explaining the separation between the subcluster's mass center and the gas bullet. The separation has been argued to be 2σ direct evidence against modified gravity [52].

In addition to providing direct observational evidence for dark matter, the bullet cluster can be used to set a limit on the self-interaction cross section of dark matter using at least two different methods. The derived limits are only strong enough to limit strongly interacting dark matter and not weakly interacting dark matter. Both methods use a scattering depth to characterize the typical number of dark matter scatters as the subcluster passed through the core of the main cluster

$$\tau_s = \frac{\sigma}{m}\Sigma_s. \tag{1.5.10}$$

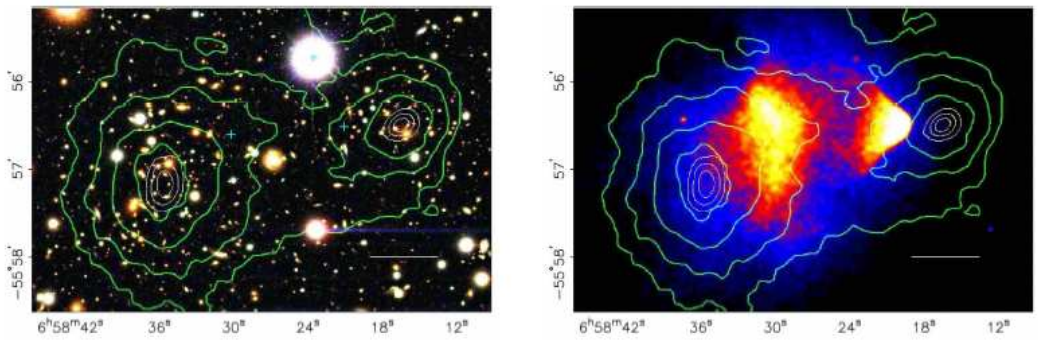

Figure 1.12: The bullet cluster in the optical, X-ray, and weakly lensed. The subcluster is on the right-hand side of each panel and the main cluster is on the left. The left panel shows an overlay of weak lensing mass contours on the optical image of the bullet cluster. Each contour line indicates a $1\,\sigma$ confidence interval on the mass centers of the clusters. The blue +s show the centers of the hot gas. The right panel shows an overlay of the weak lensing mass contours on top of the X-ray image. The hot gas in the subcluster exhibits a bowed shock front. The gas centers is separated from the mass centers by $8\,\sigma$. Figure from [53].

Here σ is the cross section for self interaction, and Σ_s is the dark matter surface mass density. The two methods set the following limit

$$\frac{\sigma}{m} < \zeta \times 10^{-23}\ cm^2\ GeV^{-1} \tag{1.5.11}$$

where ζ equals 1 or 0.2 for the two methods [54].

1.5.3 Rotation Curves of Galaxies

The evidence described above for dark matter only points to the existence of far more nonbaryonic matter in the universe than baryonic. Evidence for dark matter from spiral galaxies is particularly important, for searches for direct

dark matter interactions within a laboratory, as it could provide evidence for local dark matter. Some of the strongest evidence for dark matter comes from looking at the rotational curve, which is the rotational velocity vs. distance from the galactic center, of spiral galaxies.

The dominant component of the visible matter in a spiral galaxy is in the shape of a disk whose surface brightness can be modeled as a function of galactic radius r as

$$I(r) = I_0 e^{-r/\alpha} \tag{1.5.12}$$

which indicates that most of the visible mass is can be found at low radii. Here α is the disk scale length. For giant spirals like the Milky Way, the disk scale length is approximately $4\mathrm{h}^{-1}$ kpc. At small radii, such as a couple kpc in the Milky Way, the galactic bulge becomes significant. Neutral hydrogen gas can be found near the disk out to $\sim$ 20 kpc [55]. If only visible matter exists, the rotation velocity away from the center of the galaxy should fall as $r^{-1/2}$, since most visible mass is found near the galactic center. Observations of spiral galaxies show that away from the center of the galaxy, the rotation velocity becomes roughly constant. The mass density required to explain the rotation curve for a disk or halo can be calculated by equating the gravitational acceleration with the centripetal acceleration.

$$\frac{G_N M(r)}{r^2} = \frac{v_c^2}{r}. \tag{1.5.13}$$

The enclosed mass, M, must be proportional to the radius. Isothermal, spher-

ically distributed particles under hydrostatic equilibrium can be described by

$$\sigma^2 \frac{\partial \rho(r)}{\partial r} = -\frac{G_N M(r)}{r^2} \rho(r) \tag{1.5.14}$$

where the velocity dispersion, σ, is assumed to be independent or radius or isothermal and the pressure equals σ times ρ. At high radius, this gives $\rho \propto r^{-2}$.

Universal Rotation Curves

Persic, Salucci, and Stel [56] showed that a galaxy's rotation curve can be predicted by its luminosity. The defined a universal rotation curve as

$$V_{URC} = (V_d^2 + V_h^2)^{\frac{1}{2}} \tag{1.5.15}$$

where V_d^2 and V_h^2 give the contributions from the disk and halo. V_d^2 and V_h^2 are only functions of the luminosity and radius. The universal rotation curves are normalized to equal the observed rotational velocity at the optical radius. The optical radius, R_{opt} is defined to be the radius containing 83% of the total integrated light.

The parameterization of the dark halo is usually written as

$$\rho(r) = \frac{\rho_0}{(r/R)^\gamma [1 + (r/R)^\alpha]^(\beta - \gamma)}. \tag{1.5.16}$$

and can fit the more widely used halo models. The values of the parameters for a modified isothermal halo are $\alpha = 2$, $\beta = 2$, $\gamma = 0$, and R = 3.5 kpc. For a Navarro, Frenk, and White (NFW) [57] halo, the parameters are $\alpha = 1$, $\beta = 3$, $\gamma = 1$, and R = 20 kpc. The additional term in the denominator keeps

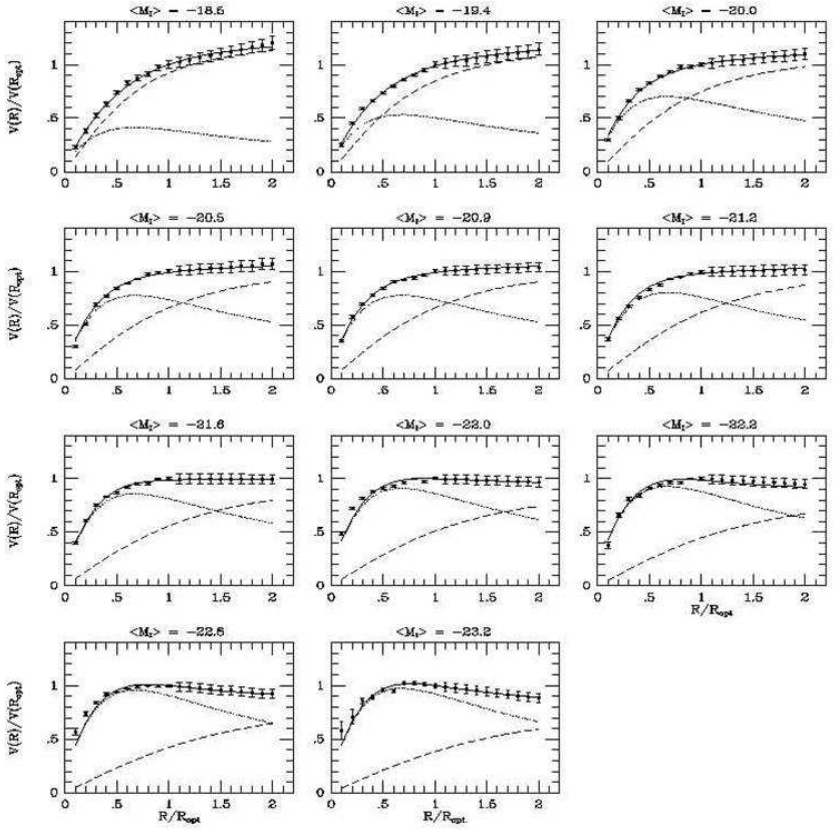

Figure 1.13: Best two-component fits to the universal rotation curve. The dotted line shows the fit to the disk and the dashed line shows the halo. The dominant source of error comes from the low radii bins where inclusion of other visible matter components, such as the bulge, would improve the fit. Figure from [56].

the density finite as the radius approaches 0 which allows the visible matter to remain the dominant component at low radii.

Particles, still distributed in an approximately isothermal distribution, must have cross sections much less than visible matter. Visible matter interacts sufficiently to have collapsed from an initially spherical distribution to the observed disk-like distribution. Particles with interaction strength at the weak-scale or below would be expected to have had sufficiently few interactions to still be in a distribution similar to an isothermal, spherical distribution. They will not have been able to collapse down to a disk like the visible matter.

The Milky Way

Because any experiment looking for dark matter will be located in the local galaxy, evidence for dark matter within the Milky Way is of particular interest. The mass distribution of matter within the solar radius of 8kpc has been well studied [58] [59] [60]. Dark matter is not required to describe the region of the galaxy within the solar radius. Figure 1.14 shows the rotation curve of the inner portion of the Milky Way.

Observationally, it is more difficult to measure the rotation curve of the Milky Way for radii higher than that of the Solar circle where the dark matter component of the mass becomes dominant in a typical spiral galaxy. Recent measurement of globular clusters [62], field blue horizontal branch stars [63], red giants [64], and satellite galaxies [65] allow for the rotation curve to be

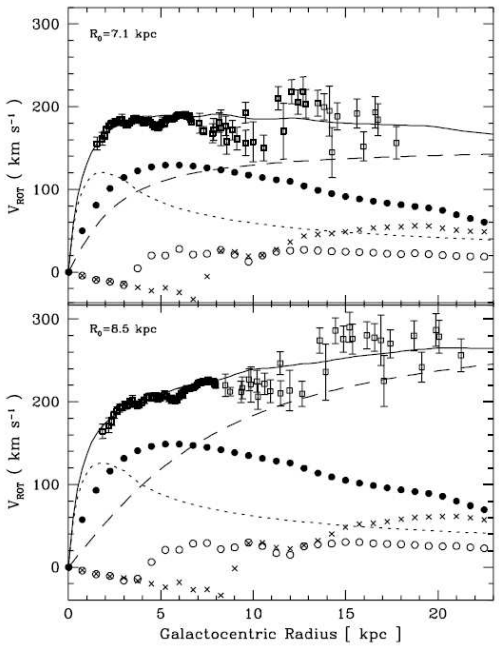

Figure 1.14: The rotation curve for the Milky Way. The components contributing to the rotation curve (solid line) are shown. The contribution of the bulge is shown by the dotted line, the disk by filled circles, the H_1 layer by crosses, the H_2 by open circles, and the dark halo by the dashed line. The upper figure is for a solar radius of 7.1 kpc and a rotational velocity of 185 km s^{-1}. The lower figure is for the solar radius of 8.5 kpc and a rotational velocity of 220 km s^{-1}. Figure from [61].

measured out to higher radii. Figure 1.15 shows that the rotation curve of the Milky Way falls more slowly with increasing radius than would be expected without the presence of a spherical dark matter halo.

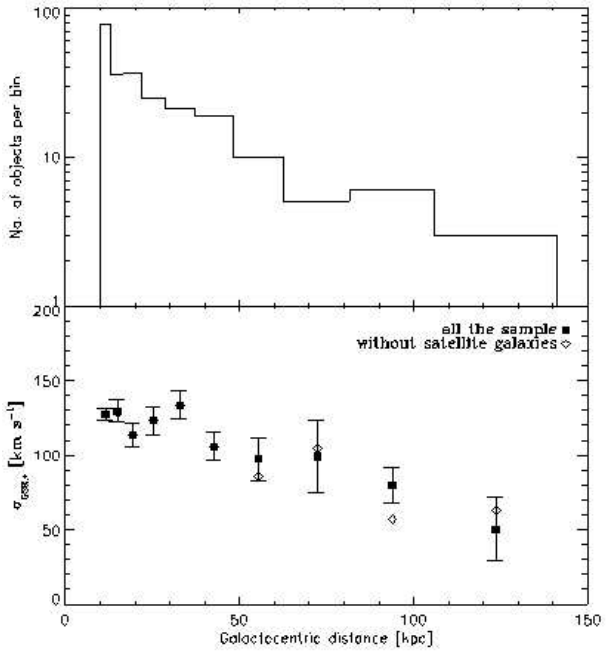

Figure 1.15: The velocity dispersion of the Milky Way at high radii. The upper panel gives the number of observed objects in each bin. The lower panel gives the radial velocity dispersion. The squares (diamonds) give the dispersion profile when the satellite galaxies are included (excluded) in the sample. Figure from [66].

It is useful to ask what is the expected WIMP energy density in the local region of the galaxy. Several estimates of these quantities exist. Bahcall *et al.* [67] estimated $\rho_0 = 0.34$ GeV cm^{-3}, Caldwell and Ostriker [68] found a slightly lower value of $\rho_0 = 0.23$ GeV cm^{-3}, and Turner [69] found an acceptable range

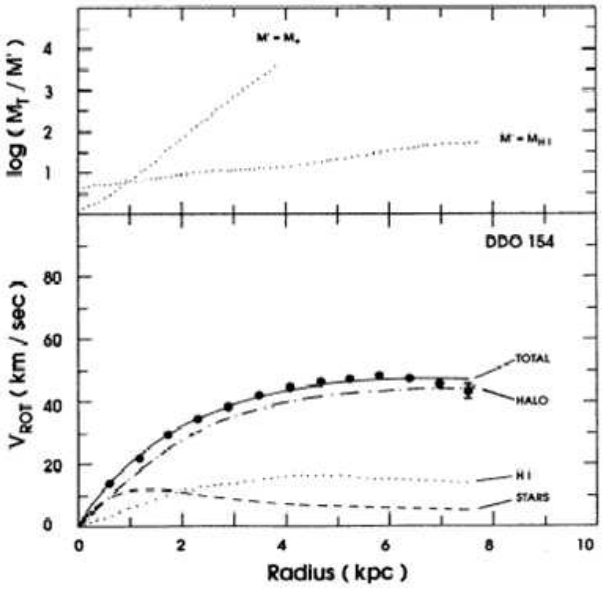

Figure 1.16: Rotation curve of the dwarf irregular galaxy DD0154. The top panel shows the ratio of the total mass to the stellar mass, M_* and to the HI mass, M_{HI} as a function of radius. The lower panel shows the rotation curve of the galaxy as well as the contribution to the rotation curve from the stars, gas, and dark halo. The error bars that are not visible are smaller than the symbol. Figure from [70].

of $rho_0 = 0.3 - 0.6$ GeV cm^{-3}. Throughout this thesis, a local WIMP energy density of 0.3 GeV cm^{-3} will be assumed.

Dwarf Galaxies

Dwarf galaxies can provide exceptionally strong evidence for dark matter. Their low visible matter content usually requires a mass-to-light ratio above that of spiral or elliptical galaxies in order to explain their rotation curves.

Mass-to-light ratios of up to 100 are observed. Dwarf galaxies are often found as satellite galaxies around larger galaxies such as the Milky Way. If they are true satellites rather than debris from mergers, dwarf galaxies become useful as high-radii tracers of the galactic mass. See Fig. 1.15 for a velocity dispersion measurement of a large galaxy, the Milky Way, using satellite galaxies.

Dwarf galaxies are split into two categories: dwarf irregulars, dIrr, and dwarf spheroidals, dSph. Dwarf irregular galaxies are dominated by dark matter and typically have non-negligible dark matter content even within the stellar core. The dark matter content of dIrr's follows the HI gas. Dwarf irregular rotation curves can be explained by multiplying the observed gas density by a factor of about 10. Figure 1.16 shows the rotation curve for the dIrr, DD0154. The evidence for dark matter within dwarf spheroidals is weak since measurement of mass density is degenerate with velocity anisotropies making it difficult to determine their dark matter content. The rotation curve of the dSph's cannot be measured out to high radii due to their visible matter content.

Chapter 2

Direct Detection of WIMPs

2.1 Introduction

The previous chapter described the evidence for the existence of nonbaryonic, weakly-interacting, massive dark matter in the universe and within our galaxy. While astrophysical observations can provide interesting constraints on the dark matter, they can not answer some fundamental questions such as what are the mass, spin, and cross section on a known target nucleus of a dark matter particle.

Observation of dark matter interactions within a terrestrial laboratory can provide answers to some of these questions. WIMPs interact with nuclei via the weak force, leading to nuclear recoils. The principle experimental technique, known as direct detection, is to try to detect the deposited energy. Information on the WIMP kinetic energy distribution can also be obtained from the distribution of energy depositions.

Detectors capable of reconstructing the direction of the recoiling nucleus from a WIMP interaction have been designed. Reconstruction of recoil directions within a detector provides an additional major challenge, but would provide additional means for characterizing the local dark matter velocity distribution.

This chapter addresses some of the basics of direct detection of WIMPs: what is the WIMP density within the local region of the Milky Way, how WIMPs are expected to scatter the nuclei in a detector, what interaction rate and energy recoil spectra would be expected, and what strategies and techniques are useful for suppressing the background below that of a weak signal level.

2.2 The Local Halo

Experiments attempting to directly detect dark matter scatters in earth bound detectors depend on knowledge of the density and velocity distribution of the dark matter particles in the local region of the galaxy as well as the velocity of the earth. In section 1.5.3 the local WIMP energy density was given within a factor of a few to be $0.3\,\mathrm{GeV\ cm^{-3}}$. For a WIMP with a mass of $100\,\mathrm{GeV}\ c^{-2}$, this translates to a few WIMPs per liter. Direct, but not very stringent, limits on the density of the dark matter within the solar system have been calculated. Estimates of the possible effect of dark matter on the orbit of Uranus [73][74] and the precession of the perihelion of the asteroid Icarus [75]

result in an upper limit of $6 \times 10^7\,\mathrm{GeV\,cm^{-3}}$, $2{\times}10^8$ weaker than the density expected from galactic considerations. More recently, detailed analysis of the precession of the perihelion of Earth and Mars, while accounting for smaller effects such as 301 large asteroids, leads to about a 300 times more stringent limit of $2 \times 10^5\,\mathrm{GeV\ cm^{-3}}$ [76]. While interesting, limits derived from the motion of solar objects are not sensitive to the dark matter density inferred from galactic rotation curves.

A reasonable approximation for the velocity distribution of the local WIMP halo is a Maxwellian distribution is given by

$$f(v)d^3v = \frac{1}{v_0^3\pi^{3/2}}e^{-v^2/v_0^2}d^3v \tag{2.2.1}$$

where v_0 is the velocity dispersion. Use of this distribution is supported by the central limit theorem which says that many statistical distributions become Gaussian or normal as a limiting case. The velocity dispersion, v_0, is typically assumed to be similar to the local rotational velocity, v_c.

Measurements of the proper motion of Cepheids by the Hipparcos satellite were used to estimate the value of the local rotational velocity, v_c, to be $(27.2 \pm 0.9)(\mathrm{R_0/kpc})\mathrm{km\ s^{-1}}$ [77] where $\mathrm{R_0}$ is the solar radius. Interferometry measurements of the proper motion of the SgrA* radio source at the center of the galaxy have been used to find another estimate of v_c to be $(30.1 \pm 0.9)(\mathrm{R_0/kpc})\mathrm{km\ s^{-1}}$ [79] [78]. A compilation of the different measurements of the solar radius yield $\mathrm{R_0} = 8.0 \pm 0.5$ kpc [80]. The local rotational velocity deduced from Cepheids is (218 ± 15) km s^{-1} and (241 ± 17) km s^{-1} from SgrA*. I will take the value

of the rotational velocity $v_c = 220$ km s^{-1} and assume a velocity dispersion of $v_0 = v_c$.

The velocity distribution of the halo should be truncated at the escape velocity of the galaxy, v_{esc}. The escape velocity has been estimated to be 580 km s^{-1} $< v_{esc} \leq$ 625 km s^{-1} [71] and $v_{esc} >$ 475 km s^{-1} [72]. I will assume an escape velocity of 600 km s^{-1}.

To understand WIMP scatters in a terrestrial detector, one can make the simplifying assumption that the WIMPs are stationary in the galactic frame. Understanding the velocity of the earth, $\mathbf{v_E}$, as it moves through the (on average) stationary WIMPs is important. The velocity of an earth-bound detector in the galactic rest frame is measured is three steps. The velocity of the earth about the sun, $\mathbf{u_E}$, the velocity of the sun relative to the local standard of rest (LSR), $\mathbf{u_s}$, and the local disk's rotational velocity, $\mathbf{u_r}$, are all measured. The LSR frame is defined by the mean velocity of nearby stars relative to their rotational velocity. Then, ignoring the diurnal rotation of the earth, the velocity of the earth will be

$$\mathbf{v_E} = \mathbf{u_r} + \mathbf{u_s} + \mathbf{u_E}. \tag{2.2.2}$$

Velocities will be quoted in Galactic coordinates in which the x-coordinate points in the direction of the Galactic center, the y-coordinate points in the direction of galactic rotation, and the z-coordinate points in the direction of the north galactic pole. The north galactic pole is defined to be a point within the constellation Coma Berenices that is perpendicular to and above the galactic

plane. The local galactic rotational velocity is $\mathbf{u_r} = (0,\ v_c,\ 0)$ km s^{-1}. A typical value quoted for the velocity of the sun relative to nearby stars is $\mathbf{u_s} =$ (9, 12, 7) km s^{-1} [81] which has statistical errors in each velocity component of about a few km s^{-1}. A more recent measurement using data from the Hipparcos satellite gives a velocity of (10.0, 5.2, 7.2) km s^{-1} [82] [78] with statistical errors on each component of about 0.5 km s^{-1}. Thus, omitting u_E for simplicity,

$$|\mathbf{v_E}| = \sqrt{10.0^2 + (220 + 5.2)^2 + 7.2^2}\,\mathrm{km\,s^{-1}} = 226\,\mathrm{km\,s^{-1}}. \tag{2.2.3}$$

The earth has an orbital speed about the sun of 30 km s^{-1}. The angle between the axis of the earth's orbit and the velocity of the sun is 30.7^o [83]. Incorporating the Earth's velocity about the sun,

$$|\mathbf{v_E}| \approx v_0 \left[1.03 + 0.07\cos\left(\frac{2\pi(t - t_m)}{1\,\mathrm{year}}\right)\right], \tag{2.2.4}$$

where t_m is the time relative to January first that the earth's velocity gives the maximum scattering rate. The scattering rate has an annual modulation as the Earth's orbit moves its velocity vector into and out of alignment with the sun's velocity. Under the simplifying assumptions of assuming an isotropic Maxwellian halo and including only the component of the Earth's motion in the galactic plane, the Earth's orbit achieves maximal alignment with the sun's motion on the 153rd day of the year, or June 2nd. More complicated assumptions result in a 15% to 20% change in the expected maximum scattering rate and a somewhat earlier point of maximum scattering rate [84].

2.2.1 WIMP Scatters

Given the above estimate of the density and velocity distribution of WIMPs in the local halo as well as the movement of the earth, the scalar cross section, the expected rate, and the energy spectrum of WIMP scatters on a target nucleus can be calculated. It is useful to begin with a discussion on the kinematics of two body scattering. Schematic 2.1 illustrates WIMP scattering in the lab frame. In the center of momentum frame, the WIMP has an initial (final) momentum of $\vec{p}_{\chi i}$ ($\vec{p}_{\chi f}$).

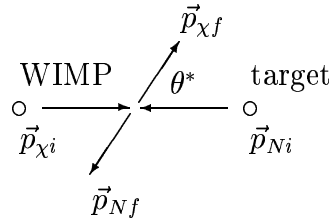

Schematic 2.1: WIMP scattering in the center of momentum frame. A WIMP with momentum, $\vec{p}_{\chi i}$, incident on a target nucleus with momentum, $\vec{p}_{Ni} = -\vec{p}_{\chi i}$. The final momentum of the WIMP (target) is $\vec{p}_{\chi f}$ ($\vec{p}_{Nf} = -\vec{p}_{\chi f}$).

In this frame, the initial momentum of the reduced mass, $\vec{p}_i$, is $\mu_r \vec{v}_i$ and the final momentum of the reduced mass, $\vec{p}_f$, is $\mu_r \vec{v}_f$. The magnitude of the initial and final momentums, p, equals $\mu_r v$ where v is $\mu_r |\vec{v}_1 - \vec{v}_2|$. The square of the momentum transferred to the reduced mass is

$$q^2 = (\mid \vec{p}_f - \vec{p}_i \mid)^2 \, , = 2\mu_r^2 v^2 (1 - \cos\theta^*) \tag{2.2.5}$$

where θ^* is the scattering angle and is isotropic in the center of momentum frame. Since the momentum transfer is invariant, the recoil energy of the

target nucleus is

$$E_r = \frac{q^2}{2m_N} = \frac{\mu_r^2 v^2}{m_N}(1 - \cos\theta^*). \tag{2.2.6}$$

To obtain the recoil energy spectrum, Eqn. 2.2.6 must be convoluted with the WIMP velocity distribution given by Eqn. 2.2.1.

Detectors have an energy threshold, E_{Th}, below which they are unable to observe energy depositions. The smallest possible WIMP velocity, resulting in a nuclear recoil that is not below threshold, occurs when the maximum momentum is transferred or when $\cos\theta^* = -1$. The minimum velocity can be expressed as

$$v_{min} = \left(\frac{E_{Th} m_N}{2\mu_r^2}\right)^{1/2}. \tag{2.2.7}$$

Coherent Scattering

Until now, discussion of WIMP scattering has been restricted to the momentum transfered to a target nucleus and the resulting recoil energy of that nucleus. An experimentalist has no control over the energy spectrum of incident WIMPs, but can control the target material. The choice of target nuclei has a major impact on the type of WIMP scatters allowed and the scattering rate in addition to impacting the recoil energy spectra.

WIMPs scatter off of nucleons in the nucleus. Two different types of scatters are possible. The simplest interaction, spin-independent, occurs when a WIMP scatters off all the nucleons in a nucleus. WIMPs could also scatter via a spin-dependent interaction. For a spin-dependent interaction, the incident

WIMP must have spin, and the target nucleus must also have net spin from an unpaired nucleon. The unpaired nucleon and the WIMP interact. I will assume the simpler case of spin-independent interactions. The spin-independent interaction will be enhanced by a coherence factor of A^2 for spin-independent interactions if qR is ≤ 1 where R is the nuclear radius. The radius of the target nucleus is approximately given by the formula [85]

$$R_N \approx (0.89A^{1/3} + 0.30)\text{fm}. \tag{2.2.8}$$

For higher momentum transfers and higher A nuclei, the spin-independent begins to lose its coherence. The loss of coherence as a function of A and q is described by a Form Factor in Section 2.2.1.

Cross Sections and Form Factors

The cross section for WIMP scattering can be calculated from Fermi's golden rule

$$\frac{d\nu}{dt} = 2\pi \mid \mathcal{M} \mid^2 \frac{dN}{dE}. \tag{2.2.9}$$

Here $d\nu/dt$ is the transition rate, $\mathcal{M}$ is the transition matrix element, and dN/dE is the density of final states. The initial and final state wave functions, normalized to a unit volume, of the reduced mass are given by $\psi_{i,f}(\vec{r}) = e^{-i\vec{p}_{i,f}\cdot\vec{r}}$, and the potential, $\phi(\vec{r})$, is $\rho(\vec{r})Af$, where f is the coupling coefficient of a WIMP

to a nucleon. The transition matrix element can be calculated

$$\begin{aligned} \mathcal{M} &= \int d^3r\, \psi_f^*(\vec{r})\, \phi(\vec{r})\, \psi_i(\vec{r}) \\ &= Af \int d^3 r\, e^{-i(\tilde{p}_i - \tilde{p}_f)\cdot\tilde{r}} \rho(\tilde{r}) \\ &= A f\, F(q), \end{aligned} \tag{2.2.10}$$

where q is the momentum transferred to the nucleus in the lab frame in which the nucleus was initially at rest. If protons and neutrons have different coupling coefficients, the transition matrix element has two parts and can be written as

$$\mathcal{M} = [Zf_p + (A - Z)f_n]F(q). \tag{2.2.11}$$

At this point, a brief detour from the calculation of the cross section will be made to discuss the form factor, $F(q)$. Lewin and Smith's paper is a good source for information on the form factor [81]. The form factor expresses the effect of the structure of the nucleus on the cross section as the momentum transfer increases and the interaction is not quite coherent. When the momentum transfer becomes large enough that the wavelength becomes similar to the size of the nucleus, the structure of the nucleus becomes important. The form factor is given by the Fourier transform of wavefunction of the target nucleus.

$$F(q) = \int \rho(r) e^{-i\vec{q}\cdot\vec{r}} d^3r \tag{2.2.12}$$

The form factor is normalized to be unity when no momentum is transfered.

To find the form factor, one must approximate the nucleon density. The simplest approximation is to assume point-like nucleons. The density can be

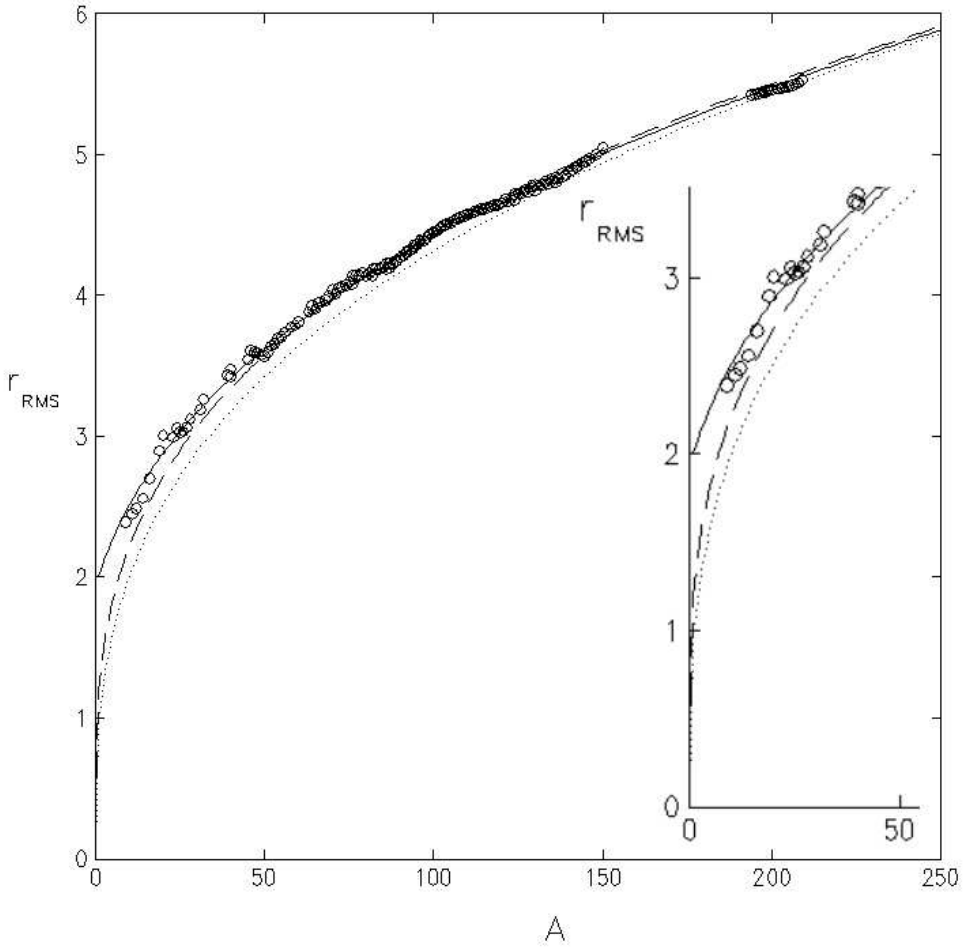

Figure 2.1: Fits to nuclear radii, r_{rms}, as a function of A. The o's represent data from muon scattering. The solid line shows the least squares fit, the dotted shows Engle's parameterization [86] of Helm's form, and the dashed shows a fit by Eder [85]. The inset region is a zoomed in version of the relevant low A portion of the figure. Figure from [81].

written as $\rho(r) = \sum_i \delta^3(r_i)$ where the sum is over the nucleons. A more reasonable form for the density that gives an analytic form factor for elastic scattering was found by Helm [87] from scattering of 187 MeV electrons off of several nuclear targets including ^{28}Si. The Helm form factor approximates the nucleus as a core with nearly constant density and a surface with a thickness, s.

$$\rho(r) = \int d^3r' \rho_0(r') \rho_1(r - r'). \tag{2.2.13}$$

ρ_0 defines the radius of the target nuclei and is given by a constant inside a sphere of radius, $R \simeq 1.2A^{1/3}$fm, and is zero outside the sphere. ρ_1 defines the surface thickness of $\simeq$ 1fm by $\rho_1(r) \propto e^{-r^2/2s^2}$. ρ_0 and ρ_1 are normalized such that $\int d^3r' \rho_{0,1}(r') = 1$. Scattering at the energies tested by Helm depends primarily on the surface nucleon density [88].

The Helm form factor can be written

$$F(q) = \frac{3j_1(qR)}{qR} e^{-(qs)^2/2}, \tag{2.2.14}$$

where j_1 is a spherical Bessel function of the first kind. A least squares fit [81] to more recent muon scattering data [89] gives values for the nuclear radius and skin thickness.

$$\begin{aligned} R^2 &= c^2 + \frac{7}{3}\pi^2 a^2 - 5s^2 \qquad (2.2.15) \\ c &\simeq 1.23A^{1/3} - 0.60\text{fm} \\ a &\simeq 0.52\text{fm} \\ s &\simeq 0.9\text{fm} \end{aligned}$$

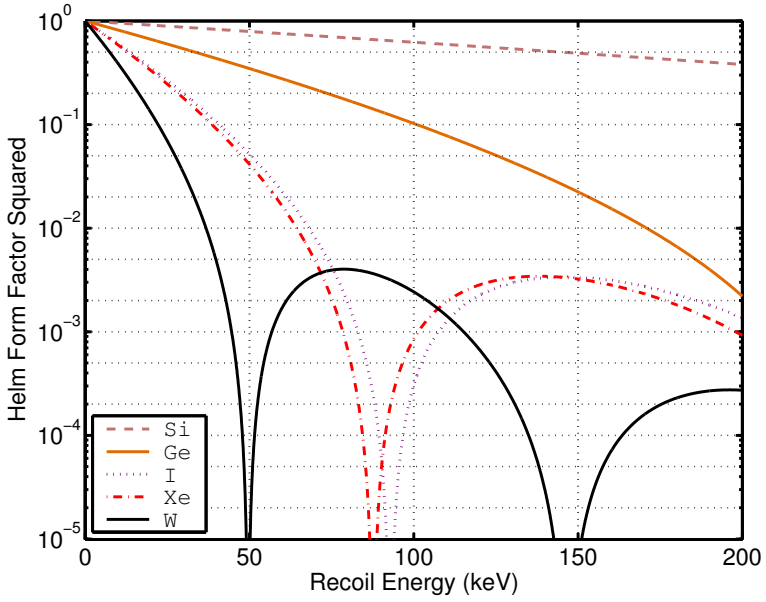

Figure 2.2: The Helm form factor squared for several typical target nuclei as a function of recoil energy. A value of 1, or no form factor suppression, is equivalent to the A^2 scaling of point-like nucleons. Figures 2.3 and 2.4 show the typical recoil energy of various target nuclei for sample WIMP masses.

Figure 2.1 compares various fit nuclear radii as a function of atomic mass, A. Helm's value for the nuclear radius provides a reasonable approximation for A greater than 10, while the least squares fit gives a reasonable approximation to the nuclear radius down to low A. The approximations to the nuclear radius and skin factor given by least squares fit are reasonably accurate as illustrated by Fig. 2.1 and will be used in this thesis.

The form factor will reduce the cross section differently for different target nuclei. Figure 2.2 shows the effect of the Helm form factor as a function of recoil energy for several typical target nuclei. Due to loss of coherence, higher recoil energies and more massive nuclei have a smaller form factor and hence lower cross section.

The next step in the calculation of the cross section is to estimate the density of final states. Using $dE_\mu = \mu_r^{-1} p dp$, one finds

$$\frac{dN}{dE_r} = \frac{d^3p}{(2\pi)^3 dE_r} = \frac{\mu^2 v d\Omega}{(2\pi)^3}. \tag{2.2.16}$$

Bringing the density of final energy states together with the transition matrix element, the golden rule can now be used to calculate the differential cross section. The transition rate equals the cross section times the relative velocity per unit volume, so

$$\frac{d\sigma}{d\Omega} = \frac{1}{v}\frac{d\nu}{dt}\frac{1}{d\Omega} = \frac{(\mu_r f A F)^2}{(2\pi)^2}. \tag{2.2.17}$$

Using Eqn. 2.2.5 to relate $d\Omega$ and dq^2, one finds

$$\frac{d\sigma}{d|q^2|} = \frac{(fAF)^2}{4\pi v^2}. \tag{2.2.18}$$

The differential cross section is frequently described in terms of a cross section for zero momentum transfer, σ_0, which is defined to be

$$\sigma_0 = \int_0^{4\mu_r^2 v^2} \frac{d\sigma(q=0)}{d|q^2|} d|q^2| = \frac{(\mu_r f A)^2}{\pi}. \tag{2.2.19}$$

The differential cross section then becomes

$$\frac{d\sigma}{d|q^2|} = \frac{\sigma_0 F^2(E_r)}{4\mu^2 v^2}. \tag{2.2.20}$$

Several target nuclei are used in WIMP searches making it useful to obtain a simple method for intercomparing results. Typically this is done by quoting a cross section normalized to a single nucleon, σ_1. The conversion between differential cross section of a given target nuclei and a single nucleon is given by

$$\frac{d\sigma}{d|q^2|} = \frac{\mu_r^2 A^2 F^2(E_r)}{\mu_1^2} \frac{d\sigma_1}{d|q^2|} \tag{2.2.21}$$

where μ_1 is the reduced mass of a single nucleon and incident WIMP.

Scattering Rate

The rate of scatters per unit detector mass is

$$R = \frac{n_\chi \sigma \bar{v}}{m_N}, \tag{2.2.22}$$

where $n_\chi = \rho_\chi \ / \ m_\chi$ is the WIMP number density, m_N is the mass of the detector's nucleus, and $\bar{v}$ is the mean incident velocity of WIMP's on the detector. This simple expression for the rate neglects two factors. Detectors can only observe scatters that deposit energies above a threshold, E_{Th}, and the differential

cross section depends on the velocity of the incident WIMP through a form factor, $F(q)$ (see Eqn. 2.2.14). The differential rate per unit detector mass is

$$dR = \frac{\rho_0}{m_\chi m_N} v f(v) \frac{d\sigma}{d|q\ |^2} d|q\ |^2\ d^3v, \tag{2.2.23}$$

where $f(v)$ is the velocity distribution of WIMPs on the detector, ρ_0 is the local WIMP energy density, m_N is the mass of a target nucleus, and q is the momentum transfered. Using Eqn. 2.2.18 for the differential cross section and noting that $dE_r = d|q|^2/2m_N$, the differential rate can be rewritten in terms of the recoil energy as

$$\frac{dR}{dE_r} = \frac{\sigma_0 \rho_0}{2 m_\chi \mu_r^2} F^2(E_r) \frac{f(v)}{v} d^3v. \tag{2.2.24}$$

Assuming the Maxwellian velocity distribution given in Eqn. 2.2.1 and also assuming that the inclusion of velocities above the escape velocity introduces a negligible increase in rate, one can integrate over the velocity distribution. The differential rate is found to be

$$\frac{dR}{dE_r} = \frac{\sigma_0 \rho_0}{m_\chi \mu_r^2} \frac{F^2(E_r)}{v_0 \pi^{1/2}} e^{-E_r m_N / 2\mu_r^2 v_0^2}. \tag{2.2.25}$$

If the form factor is ignored, the differential rate can be integrated to find a rate of

$$R(E_r) = \frac{2}{\pi^{1/2}} \frac{\sigma_0 \rho_0 v_0}{m_\chi m_N} e^{-E_r m_N / 2\mu_r^2 v_0^2}. \tag{2.2.26}$$

A more accurate expression for the differential scattering rate can be calculated if the velocity of the Earth, v_E, and the WIMP escape velocity of $v_{esc} \approx$

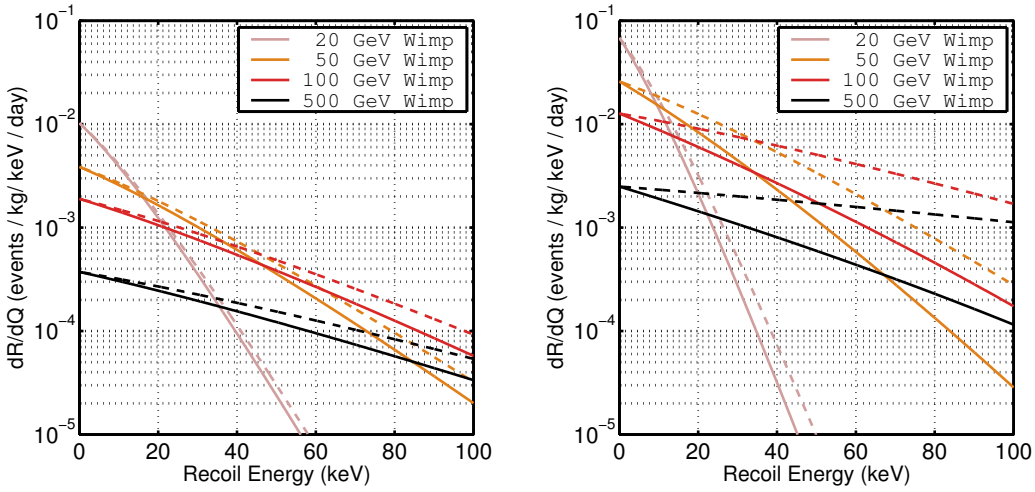

Figure 2.3: The differential scattering rate on silicon (left) and germanium (right) for selected WIMP masses. The cross section for a WIMP scatter on a neutron or a proton is assumed to be identical and equal to $1 \times 10^{-42}\mathrm{cm}^2$. The standard value of ρ_0 and a time averaged value of the velocity of the Earth from Eqn. 2.2.4 were used. The dashed lines ignore the effect of the form factor while the solid lines include the effect of the Helm form factor.

600 km/s are not neglected. The velocity of the Earth is given by Eqn2.2.4. The differential scattering rate becomes

$$\begin{aligned} \frac{dR}{dE_r} &= \frac{\sigma_0 \rho_0}{2 m_\chi \mu_r^2} \frac{F^2(E_r)}{v_0^3 \pi^{3/2}} \int v^{-1} e^{-(v+v_E)^2/v_0^2} d^3v \\ &= \frac{\sigma_0 \rho_0 F^2(E_r)}{4 v_E m_\chi \mu_r^2} \left[\mathrm{erf}\left(\frac{\mathrm{v_{min}} + \mathrm{v_E}}{\mathrm{v_0}} \right) - \mathrm{erf}\left(\frac{\mathrm{v_{min}} - \mathrm{v_E}}{\mathrm{v_0}} \right) \right] \end{aligned} \quad (2.2.27)$$

where v_{min}, given by Eqn. 2.2.7, is the minimum incident WIMP velocity that induces a nuclear recoil above threshold.

The rate can be significantly impacted by the mass of the target nucleus and the mass of the incident WIMP. Figure 2.3 shows the differential scattering rate

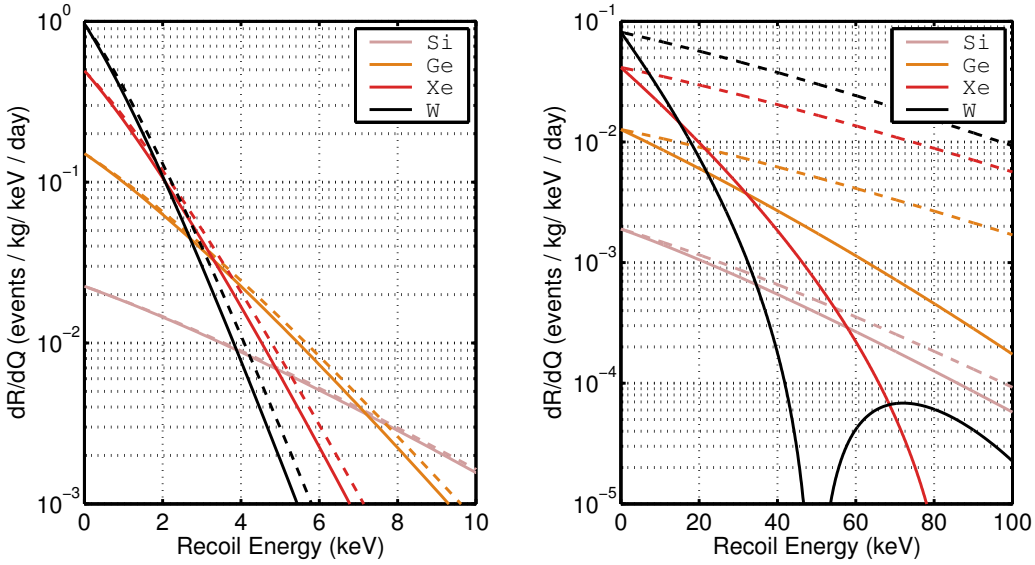

Figure 2.4: The differential scattering rate on typical targets for WIMPs with masses of 10 GeV and 100 GeV. The left panel shows the rate for 10 GeV WIMPs and the right panel shows the rate for 100 GeV WIMPs. The cross section for a WIMP scatter on a neutron or a proton is assumed to be identical and equal to $1 \times 10^{-42}\text{cm}^2$. Scattering rates for iodine are similar to xenon. The standard value of ρ_0 and a time averaged value of the velocity of the Earth from Eqn. 2.2.4 were used. The dashed lines ignore the effect of the form factor while the solid lines include the effect of the Helm form factor.

for WIMPs with various masses incident on germanium. Figure 2.4 illustrates the effect of changing the target material by showing the differential scattering rate for 10 GeV and 100 GeV WIMPs incident on several typical targets. For both figures, the WIMP cross section on a neutron or proton is assumed to be equal to $1 \times 10^{-42}\text{cm}^2$. These figures also illustrate the effect of the form factor, particularly for higher mass WIMPs and target nuclei, as well as indicating the typical low recoil energy induced by WIMP scatters. The exponential fall in

the WIMP interaction rate as well as the low energy of the induced nuclear recoils makes sensitivity to low energies crucial for attempting to observe WIMP scatters.

The rate of events observed is expected to be quite small. A convient equation for describing the expected rate in units of events per kg·day is provided by Lewin and Smith [81] under the simplifying assumptions $\mathbf{v_E} = 0$ and $v_{esc} = 0$.

$$R = \frac{3.61 \times 10^{-4}}{m_N m_\chi} \left(\frac{\sigma_0}{10^{-42}\mathrm{cm}^2}\right) \left(\frac{\rho_\chi}{0.3\,\mathrm{GeV\,c^{-2}\,cm^{-3}}}\right) \left(\frac{v_0}{220\mathrm{km\,s^{-1}}}\right). \quad (2.2.28)$$

where m_χ is the WIMP mass and ρ_χ is the local WIMP energy density. Using Eq. 2.2.21 and ignoring the suppression of the rate due to the form factor, one can expect $\sim$0.5 events per kilogram of germanium per day for a 100 GeV WIMP with a cross section on a single nucleon of $10^{-42}\,\mathrm{cm}^2$.

2.3 Experimental Strategies

Electromagnetic Backgrounds

The primary challenges for direct searches for WIMP recoils are the low rate of WIMP scatters, the much greater ambient background rate, and the low typical recoil energy of the events. To develop a feeling for the ambient background rate, one can calculate the number of ^{40}K decays per second in a 16 ounce pint of beer. Beer typically has a radioactivity due to ^{40}K of 390 pCi per kilogram [90] or about 7 decays per second. Placing a pint of beer next to

a typical CDMS detector of mass 0.25 kilograms will result in an observed rate in excess of 10^5 events per day. A typical person contains six to seven hundred times more ^{40}K than a pint of beer [90]. Clearly, methods for dealing with the ambient background are crucial for any experiment hoping to detect WIMP scatters.

Several strategies for dealing with this challenge are used in current experiments. Five basic strategies for increasing sensitivity to WIMPs are 1) to increase the detector mass to increase the signal event rate, 2) use shielding to reduce the frequency of background events, 3) to use low-radioactive background materials surrounding the detectors, 4) to employ background rejection techniques, and 5) to exploit expected variations of the signal rate over time. The DAMA experiment uses large detector mass and exploits an expected annual variation in the rate of signal to overcome a large background as is briefly described at the end of this section.

Increasing the target mass results in limited benefits for an experiment with significant background as the sensitivity only improves as the square root of the quantity mass times detector livetime. It is difficult to obtain significant improvement in sensitivity for a background limited experiment by increasing detector mass without accompanying background reduction.

Shielding against Electromagnetic Backgrounds

Because the electromagnetic background is so large, every experiment attempting to directly detect dark matter employs shielding to reduce the rate of electromagnetic events. Shielding materials are chosen for their density and purity. High density materials are used for their better shielding of gammas. High purity materials are also common as radioactive isotopes in the shielding can contribute to the electromagnetic background. Lead and copper are typical choices. Experiments searching for WIMPs are conducted underground to avoid backgrounds induced by cosmic rays. Experiments operating underground also need to keep the radon-contaminated air away from their detector.

Rejecting Electromagnetic Backgrounds

Techniques for rejecting background typically exploit the fundamentally different nature of nuclear recoils induced by WIMP scatters versus background electron scatters due to incident gammas and electrons. Nuclear recoils travel very short distances, typically a few tens of nm, and are fairly inefficient at producing ionization. They can deposit most of their energy in phonons or scintillation instead. Electron recoils travel much larger distances and deposit more of their energy in ionization. Comparison of the phonon or scintillation energy with the ionization energy allows background electron scatters to be rejected. The CDMS experiment employs detectors that are sensitive to both phonons and ionization.

Neutron Backgrounds

Experiments capable of rejecting background electron scatters become sensitive to the smaller background of nuclear recoils induced by neutrons. Background neutrons are more pernicious than the electromagnetic background. Their lack of electric charge allows them to bypass electrons and interact with the nuclei. The resulting neutron-induced nuclear recoils are indistinguishable on an event by event basis from WIMP induced nuclear recoils.

Neutrons are produced by nuclear decays in materials surrounding the detectors and by nuclear spallation caused by cosmic muons. Lower energy neutrons, produced by nuclear decays in the surrounding rock, can be moderated by hydrogenic shielding down to energies that induce nuclear recoils that are below threshold. Most neutrons with sufficient energy to penetrate the shielding and induce a nuclear recoil above threshold are caused by the muonic component of the cosmic ray background. A 1 MeV neutron will typically induce a nuclear recoil with recoil energy in the few tens of keV. Nuclear recoils induced by neutrons that are much lower in energy will usually be below threshold.

The flux of muons, and hence of dangerous neutrons, can be greatly reduced by placing an experiment underground at depths of order thousands of meters water equivalent. Figure 2.5 shows the muon flux as a function of depth in meters water equivalent. Neutrons produced by muons inside the shielding can be rejected by placing an active muon veto around the detector shielding. Some neutrons scatters can be rejected if multiple detectors are employed as neutrons

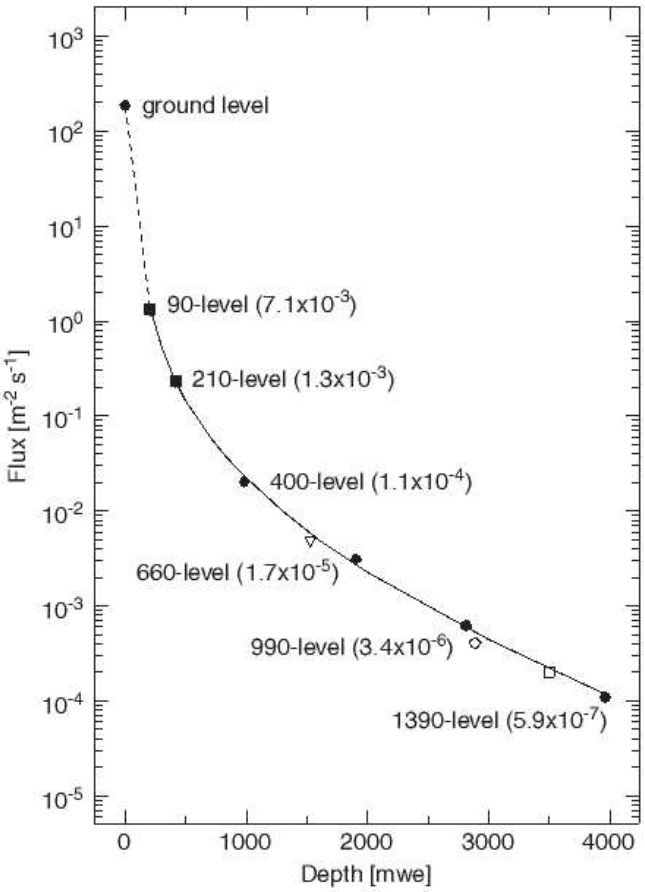

Figure 2.5: Muon intensity as a function of depth in meters water equivalent (mwe). Full circles are from the Pyhäsalmi mine in Finland [91], full squares show an earlier measurement with a similar detector system, open circles from [92], and open triangles from [93]. The uncertainties are quoted as being smaller than the size of the symbol when not shown. The curve is a χ^2 fit to a function of the form $F(x) = A(x_0/x)^{\eta}\exp(-x/x_0)$. Figure from [91].

can self veto by scattering in multiple detectors. WIMP scatter rates are far too low for WIMP scatters to simultaneously occur in multiple detectors.

Signal Modulation

As was shown in Section 2.2, the motion of the earth will affect the energy spectrum and direction of nuclear recoils from WIMP scatters. As was shown in Section 2.2, when the earth's motion reaches maximal alignment with the motion of Sun, the earth's velocity in the galactic frame will increase. This results in: 1) the WIMP flux on the earth being increased, and 2) the energy spectrum of the nuclear recoils being altered, enhancing the fraction of higher energy, above threshold recoils. Similarly when the earth's motion is maximally antialigned with the Sun, the WIMP flux will decrease, and the energy spectrum of the nuclear recoils will be altered decreasing the fraction of above threshold recoils. The ensuing annual modulation in a WIMP signal can be exploited to increase WIMP sensitivity in the presence of significant background. After accounting for the WIMP velocity distribution, the modulation in integrated rate is about 5% from Table 1 of [81], less than expected from the variation of the velocity of the earth alone given by Eqn. 2.2.4. From Table 1 of [81], including the effect of the altered energy spectrum, increases the modulation of the rate of higher energy recoils to about 17%.

Chapter 3

The CDMS II Experiment at Soudan

3.1 Overview

The CDMS experiment consists of shielding from the overburden of rock in the Soudan mine, a dilution refrigerator, an active veto, passive man-made shielding, a μ-metal shield, copper refrigeration cans, germanium and silicon WIMP detectors, and a data acquisition system. I will briefly describe it from the outside in.

Outermost to the CDMS experiment is the natural muon shielding provided by the overburden of rock at the Soudan mine. Within the mine, an active muon veto surrounds layers of passive shielding consisting of polyethylene to reduce neutrons and lead to reduce gamma-rays. The veto and passive shielding are housed within a clean room. The CDMS detectors, whose ability to reject electromagnetic backgrounds is one of the keys to the impressive

WIMP sensitivity achieved by CDMS, are located within the passive shielding. The detectors are cooled to their operational temperature of $\sim$50 mK by a dilution refrigerator. Signals from the detectors and active muon veto are processed by the cold and detector electronics of the data acquisition system (DAQ), which will be described in some detail. Particular attention will be paid to the trigger system and the Java based software used to control and monitor the experiment.

3.2 Shielding and Veto

The CDMS shielding blocks most cosmic ray muons that could produce neutrons in the rock surrounding the experiment, reduces backgrounds from gamma-rays, and moderates neutrons produced outside the passive shielding down to energies that will not produce nuclear recoils that are above the energy threshold. The rock overburden in the Soudan mine attenuates most muons. The muon veto system identifies muons that could produce neutrons within the passive shielding. The muon veto also rejects some neutrons produced outside the passive shielding, because those neutrons are typically accompanied by a muon or other charged particles to which the muon veto is sensitive.

3.2.1 Soudan Underground Laboratory

The CDMS experiment resides within the Soudan Underground Laboratory (SUL) which is located in Northeastern Minnesota. SUL is located in the 27th level of the Soudan Underground Mine at a depth of 2341 feet (713 meters) [94] below the surface. Figure 3.1 shows the entrance to the Soudan Mine. SUL is a retired iron mine located in Soudan, a small town in northern Minnesota. Opened in 1884, it was the oldest and deepest iron mine in the state. At one point, the iron mines in northern Minnesota were the leading supplier of iron in the United States. In December of 1962 the Soudan Mine was closed for business operations. The Soudan Mine was registered with the National Register of Historic Places in 1966[95].

Since then, the mine has been operated by the Minnesota Department of Natural Resources (DNR) who maintain the mine in its original state to preserve its historical significance.

The DNR allowed the construction of the Soudan Underground Laboratory in the 1980s. The SUL is leased by the University of Minnesota [96] from the DNR. The first experiment to utilize the SUL was the now defunct Soudan2 experiment. Currently the CDMS and MINOS experiments reside in the SUL. Figure 3.2 shows a schematic view of the Soudan Underground Laboratory.

Figure 3.3 shows a view of the portion of the SUL in which CDMS resides. Figure 3.4 shows a schematic view of how the CDMS experiment is laid out within the Soudan 2 cavern.

Figure 3.1: A view of the entrance to the Soudan Underground Mine on a typical winter day. The two gates in the center of the photo with metal bars in the shape of an "X" and yellow signs are the entry point for the "cage" or "boat". The cage is the name of the elevator compartment that one enters to be lowered down into the mine. It is particularly important not to miss the boat up to the surface at the end of the work day!

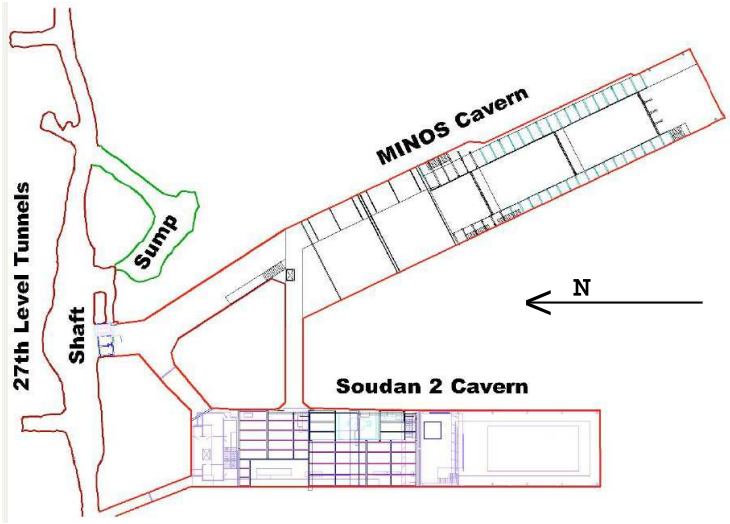

Figure 3.2: A layout of the SUL on the 27th level of the Soudan mine. The Soudan 2 cavern has a length of 71.6 m and a width of 13.9 m. The MINOS cavern has a length of 82.3 m and a height of 13.6 m. The width of the MINOS cavern is 15.9 for the northern 37.7 m of the cavern and 13.8 m for the remainder of the cavern. The CDMS experiment resides in the Soudan 2 cavern just to the south of the east-west passage to the MINOS cavern.

Figure 3.3: A view of CDMS in the Soudan Underground Laboratory. The picture is taken on top of the mezzanine looking south. To the left, the top of the RF-shielded clean room housing the detectors can be seen protruding through the mezzanine. Immediately behind the clean room is the DAQ room. Through the tunnel, the top of which can be seen to the far left, is the chamber housing the MINOS experiment. Behind the white wall is where the Soudan2 experiment was located before it was removed. After this picture was taken, a ping pong table, which was particularly useful for sanity preservation, was located on the far right hand side of the mezzanine.

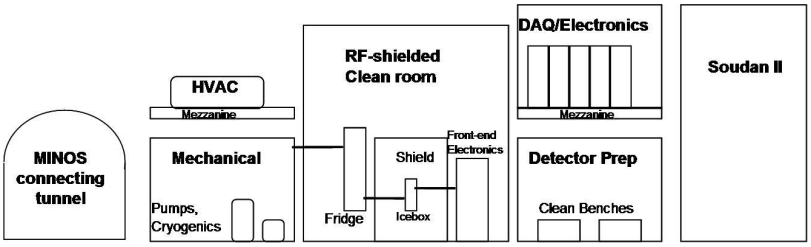

Figure 3.4: A schematic view of the CDMS layout at in the Soudan 2 cavern when looking towards the east. The Soudan 2 cavern has a length of 71.6 m and a height of 11.2 m. North is to the left. The mechanical section is also called the cryopad. The cryopad is a raised area where the pumps, controls, and monitoring computers for the dilution refrigerator reside. The RF-shielded clean room houses the veto, passive shielding, and detectors. Detector pulses travel from the detectors through the front end boards to the DAQ room.

The University of Minnesota employs a mine crew to maintain the facility, and this crew has provided helpful assistance numerous times. While the ambient air temperature in the mine is a nippy 50 ^{o}F (10 ^{o}C), the air entering the SUL is heated to approximately 68 ^{o}F (20 ^{o}C).

The rock overburden provides 2090 meters water equivalent (mwe) of shielding from cosmic radiation, reducing the muon flux by a factor of 5×10^4 [97]. Figure 2.5 illustrates how the muon flux decreases with increased depth. The depth of SUL has decreased the number of neutrons observed in approximately 70 kg-d of exposure from $\sim$50 at a depth of $\sim$15 mwe (the depth of the Stanford Underground Facility) to less than 1 neutron.

The surrounding rock in the Soudan mine contains ^{238}U, ^{232}Th, and ^{40}K. The concentration of these isotopes in the surrounding rock was measured to

be 0.17±0.06, 0.89±0.20, and 0.79±0.04 parts per million (ppm) respectively [98]. The walls of the mine are coated with 2.5 cm thick shotcrete and have been measured to have the following concentrations 0.78±0.06, 3.21±0.20, and 0.91±0.04 ppm respectively [98].

^{222}Rn from the Uranium/Thorium decay chain can plate-out on detector surfaces. The radon content in the air of the SUL ranges between 5.5 pCi/l in the winter to 19 pCi/l in the summer [99]. ^{222}Rn has a half-life of 3.8 days. Since radon decays produce hard-to-reject electrons, the air volume near the detectors is purged with 3 to 4 week-old, radon-free air.

3.2.2 Veto

The veto system is intended to reject muons that can interact with nuclei in the passive shielding and produce neutrons which can fake WIMP events in the detectors. The CDMS detectors and passive shielding are surrounded by an active muon veto. The muon veto is composed of 40 BICRON BC-408 plastic scintillator paddles. The paddles are arranged such that adjacent panels have a slight overlap to minimize the possibility of muons entering the passively shielded region undetected. Figure 3.5 shows a cartoon view of how the veto panels are arranged while Fig. 3.6 and Fig. 3.7 show schematic views of the veto panels and passive shielding from the top and side. Acrylic light guides direct the scintillation photons in the panels to attached photo multiplier tubes (PMTs). The passive shielding and veto were designed and fabricated at UCSB

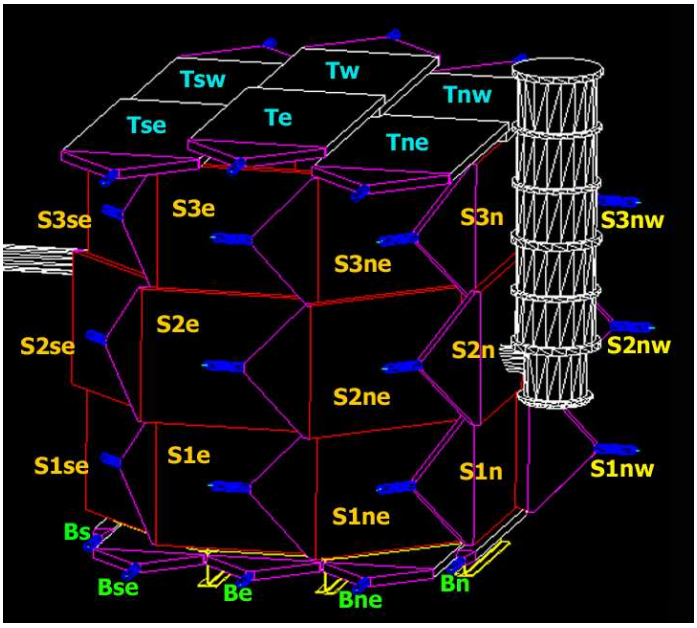

Figure 3.5: A schematic view of the CDMS veto looking from above to the west south west. North is to the right or towards the dilution refrigerator. Veto panels are labeled by their position. T for top, B for bottom, and S for side. North, south, east, and west are indicated by n s, e, and w. The veto is about 2.6 m tall and 2.6 m in diameter. Figure by Susanne Kyre and Raymond Bunker.

by Ray Bunker, Dan Callahan, Dave Hale, Susanne Kyre, and Steve Yellin. The PMTs are Hamamatsu R329-02. They are operated with voltages from 1.2 kV to 1.3 kV, such that muons produce, after amplification by the veto electronics, approximately 100 pC of collected charge. Power to the 40 veto panels is supplied by a single CAEN SY2527 Universal Multichannel Power Supply System.

Veto panels surround the detector volume and are labeled by their position.

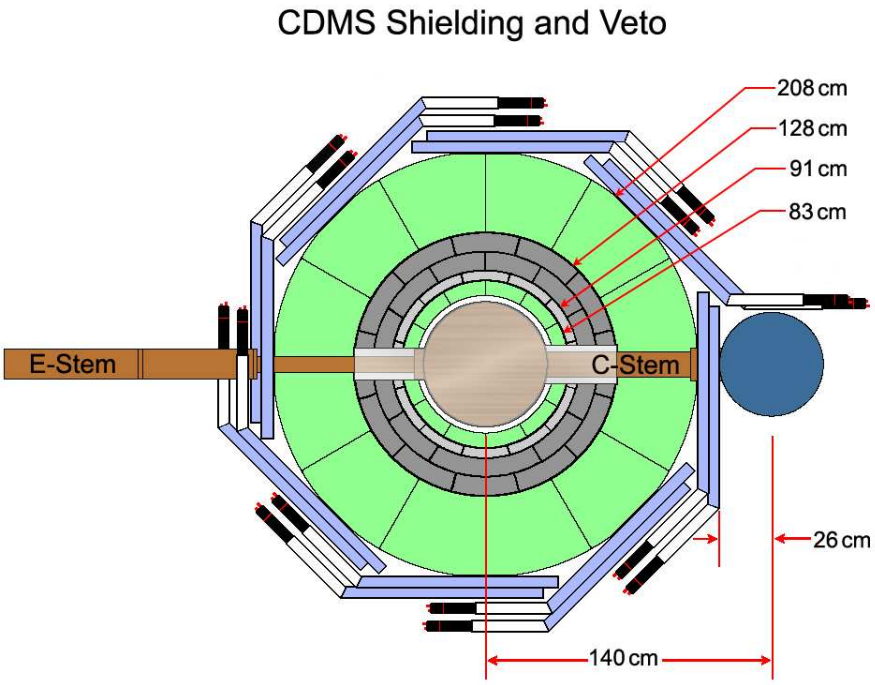

Figure 3.6: Schematic view of the top of the veto and passive shielding. The side veto panels (light blue) surround the outer poly (green). Within the outer poly, the three circles of gray pieces are the lead bricks. Within the lead bricks is a layer of inner poly (green). The innermost circle (light metal colored) represents the mu metal shield. The E-stem and C-stem (brown) pass between veto panels and through the passive shielding and connect to the mu metal shield. The E-stem connects to the E-box out of which the detector signals pass and the C-stem connects to the dilution refrigerator (dark blue), located to the right or north of the passive shielding and veto, and provides the thermal contact between the cold layers and the fridge. Figure from Susanne Kyre. Modified for clarity by Raymond Bunker and Marie Burggraf.

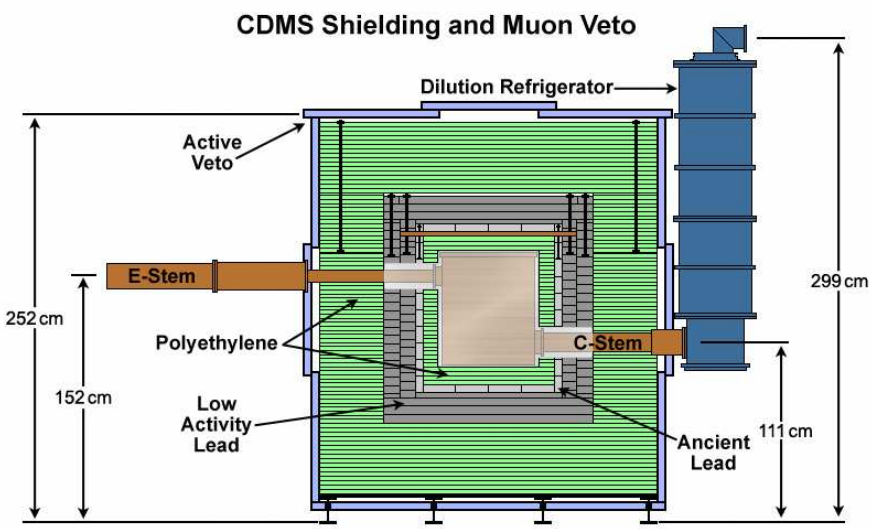

Figure 3.7: Schematic view of the side of the veto and passive shielding. The poly shielding is shown in green, and the lead shielding is shown in gray. The veto panels are shown in light blue, and the four steel I-beams, on which the passive shielding rests, are shown in black. The E-stem and C-stem are shown in brown, and the dilution refrigerator is shown in dark blue. Figure from Susanne Kyre. Modified for clarity by Raymond Bunker and Marie Burggraf.

T for top, B for bottom, and S for side. Figure 3.5 shows a schematic view of the veto with the visible veto panels appropriately labeled. The position of a veto panel is further specified by its geographic location. For example, Bsw for the bottom panel on the southwest side. Likewise, the top veto panel on the north east is labeled Tne. There are three rows of side panels. The side panel on top north east is labeled S3ne. The 6 top veto panels are 78.7 cm×137.2 cm×5.1 cm, the 6 larger bottom veto panels are 62.9 cm×109.2 cm×5.1 cm, and the 2 bottom crack panels, Bs and Bn, are 10.2 cm×109.2 cm×5.1 cm. The side panels,

except for those abutting the cold stem (c-stem) and electronics-stem (c-stem) on the north and south, are 76.2 cm×96.5 cm×5.1 cm. The panels abutting the stems have a curved region cut out of them to allow the stems to pass through with minimal gaps.

Distinguishing Muons from Compton-scattered Electrons

While the veto was designed to reject incident muons, most events in the veto are due to the Compton scattering of γ-rays that originate from Uranium/Thorium decays in the surrounding rock. The veto paddle thickness of 5 cm was chosen so that typical muons will deposit more energy than typical low-energy Compton events. All the energy from low-energy Compton electrons, typically less than 2 MeV, is absorbed by the scintillator paddles. A minimum ionizing muon typically deposits 2.0 MeV g^{-1} cm^2 or about 10 MeV in a 5 cm thick veto paddle with a density of 1.03 grams per cm^3.

A hardware veto threshold is set to allow muons to be rejected. The hardware threshold level is set as far below the level needed to reject typical muons as is possible without allowing too many Compton-scattered electrons to be over the veto threshold. If too many Compton-scattered electrons are over the threshold, then too high a fraction of potential WIMP interactions will be rejected. I measured the rate of Compton events above threshold in the combined 40 veto paddles to be $R_c = 603 \pm 2$ Hz [100], using the time difference between a trigger and the veto event immediately preceding the trigger. Figure 3.8 shows

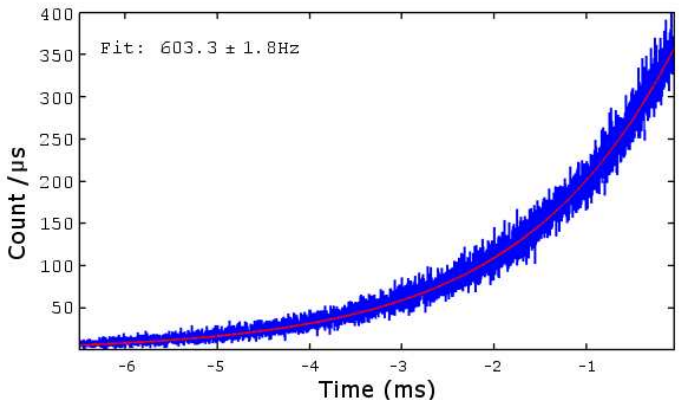

Figure 3.8: A histogram of the time from the veto event immediately preceding the trigger to the time of the trigger. the red curve shows the best fit veto rate to the histogram. The best fit veto rate is 603±2 Hz.

a histogram of the veto events preceding the global triggers and the fit to these events.

Figure 3.9 shows the distribution of muons and Compton-scattered electrons in a sample veto panel, Tw. After passing through a veto amplifier described in Section 3.4.3, the 10 MeV deposited by typical muon translates to about 100 pC. Typical Compton events yield than 30 pC. Muons that clip the edge of a veto panel yields less than 100 pC and are the population of events in Fig. 3.9 between the muon and Compton peaks.

Muon Rejection

Neutrons produced by muon interactions can cause nuclear recoils and must be addressed. These interactions can occur in two different locations and are

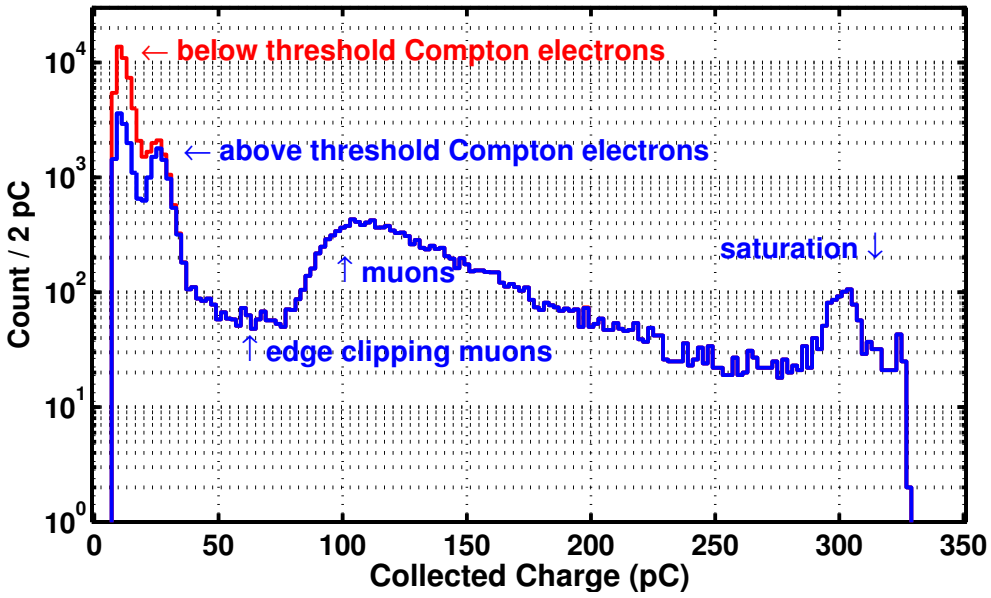

Figure 3.9: In situ energy spectra from one of the top veto panels, Tw. The darker blue line shows the energy spectra of all events in the WIMP-search data run. The lighter red line shows the spectra for events for which at least 2 veto panels were over threshold. The veto threshold is around 20 pC. Normally incident muons yield about 100 pC of charge. The valley between the muon peak and the gamma peaks is populated by muons that clip the edge of the veto panel.

dealt with by different means. 1) Events due to neutrons produced in the passive shielding can be rejected by vetoing events in the dark matter detectors that are coincident with a muon in the veto. These neutrons are the primary source of background nuclear recoils. The efficiency of the veto system for rejecting these events is discussed in the veto efficiency section to follow. 2) The secondary source of background nuclear recoils is neutrons produced outside the veto. These events can sometimes be rejected because of associated charged particles incident on the muon veto.

The veto has some ability to reject muon-induced neutrons produced outside the veto. These neutrons are typically accompanied by a muon or shower of particles that can be tagged by the veto. When incident on a veto panel, external neutrons and their accompanying particle(s) or showers can deposit energy above a veto threshold. Simulations indicate that at least 70% of external neutrons are accompanied by charged particles, including muons, and are rejected by the veto [101]. (Simulations subsequent to the completion of this analysis indicate that 93% of external neutrons are rejected by the veto [102].)

Veto Efficiency

The veto efficiency is estimated using three steps. First, a sample of muons entering the veto are tagged. The fraction of these muons observed as they exit the veto is used to approximate the efficiency of the veto for rejecting exiting muons. Second, a sample of muons are tagged exiting the veto. The fraction

of these muons observed entering the veto is used to approximate the efficiency of the veto for rejecting muons as they enter. Third, missing a muon requires missing it in both top and bottom portions. The veto efficiency will then be

$$\epsilon_{veto} = 1 - (1 - \epsilon_{in}) \times (1 - \epsilon_{out}) \tag{3.2.1}$$

where ϵ_{in} (ϵ_{out}) is the efficiency for tagging muons as they enter (exit) the veto.

Since the veto panels extend beyond the passively shielded region, some care must be taken in choosing samples of tagged muons to prevent underestimation of the veto efficiency. Muons passing through regions of veto panels that extend beyond the passively shielded region might not enter the shielded volume and thus have no chance to strike a second panel. See, for, example the Tse veto panel in Fig. 3.5. These muons would be mischaracterized as veto inefficiency.

Another class of events that must be excluded from a measurement of the veto efficiency are shower events. Showers are events that strike 8 or more veto panels. Because a typical shower event is observed in many veto panels, inclusion of shower events can cause an overestimate of the veto efficiency. Shower events are rejected by requiring the veto panels near the panel with the tagged muon to be unhit or consistent with Compton electrons by requiring that their deposited charge be below 15 pC.

To measure the efficiency of the veto for rejecting the muons as they entered, a sample of muons tagged exiting the passively shielded volume was selected by using muons in the bottom paddles, Be and Bw. A sample of 16273 events containing pulse heights in the bottom panels consistent with a muon was

selected. Shower events were excluded from the sample by requiring no activity in the remaining 6 bottom panels. No perfect selection of exiting muons can be defined, since all bottom and side veto panels extend beyond the shielded volume. The Be and Bw panels extend beyond the shielded volume by the least amount (about 1 cm or less than 1% of the 110 cm long panels). Tagged muons are considered rejected by the veto if one of the top or side veto panels is above threshold. 16244 of the selected muons were vetoed as the entered the veto. From these numbers, the efficiency for rejecting muons as they enter the veto is 99.70+0.08-0.10% (95% confidence level).

The efficiency is expected to be underestimated by less than 1%. About 1% of the muons that strike the Be and Bw panels will miss all other panels entirely, because the bottom panels extend beyond the dimensions of the adjacent panels as can be seen in Fig. 3.5. The effect leading to the underestimate is minimized by the placement of the top veto panels which extend further beyond the dimensions of the other panels than do the bottom veto panels. Some of the muons hitting regions of Be and Bw that extend beyond the dimensions of the adjacent panels will also pass through one of the top veto panels. Thus the efficiency is expected to be underestimated by less than 1%.

Muons passing through the overlap region between the Te and Tw veto panels, were used as a sample of muons tagged entering the passively shielded volume. The smaller sample was used because the top and side veto panels all extend beyond the shielded region. Only portions of the veto panels that do not

extend past the shielded region can be used to tag entering muons. The overlap region of these two panels is entirely above the shielded region and is directly above the detectors. Of the 893 muons tagged entering the shielded volume, 830 were vetoed as they exited. At the 95% confidence level, the efficiency for rejecting muons exiting the veto is 92.95+1.59-1.88%.

There is an expected inefficiency for rejecting exiting muons due to the presence of four steel I-beams on which the passive shielding rests. The four I-beams, shown in Fig. 3.7, cause 2.3 cm wide gaps in the bottom veto panels. The two central I-beams have the same 208 cm length as the diameter of the outer poly shielding. The two outer I-beams have a length of 86 cm. The fraction of bottom area not covered by the veto is 4%. The larger measured inefficiency can be explained by the vertical nature of the muons and the 78.7 cm wide top veto panels being wider than the bottom veto panels. Because the top panels are wider, the Te and Tw panels reside directly above 2 I-beams. The 2 central I-beams occupy about 5.8% of the area directly below the overlap region. A similar fraction of the highly vertical muons passing through the overlap region would be expected to pass through the I-beams and miss the bottom veto panels entirely.

From Eqn. 3.2.1, the veto efficiency for tagging through-going muons is 99.98±0.02%, reducing the rate of nuclear recoils from neutrons produced within the shielding by over three orders of magnitude. The background from neutrons produced by muon interactions within the passive shielding can be

neglected.

Muons identified hitting one of the towers of detectors during the data run on which this thesis is based can also be used to check the efficiency of the veto. The expected rate of muons on the towers is about 1 per day, so a measurement using these muons will suffer from low statistics. All of the 72 muons (34 in tower 1 and 38 in tower 2) that hit at least 5 consecutive detectors were vetoed and had deposited muon-like energy in at least 1 veto panel. The lower limit on veto efficiency is 95% (at the 95% confidence level).

3.2.3 Passive Shielding

In addition to the surrounding rock, the CDMS shielding is also composed of man-made passive shielding. The passive shielding has 3 parts: an outer layer of polyethylene (poly), two layers of lead, and an inner layer of polyethylene as can be in Fig. 3.6 and Fig. 3.7. The poly is used to moderate neutrons down to energies that induce nuclear recoils that are below the detectors' thresholds. The lead attenuates the external gamma-ray flux. Figure 3.10 and Fig. 3.11 show the shielding as it was being installed.

The outer polyethylene layer is composed of a number of one inch thick slabs that are designed to stack together around the outer lead brick layer. The outer poly layer is 40 cm thick on the sides and 40.6 cm thick on top and bottom.

The lead reduces the ambient gamma background. The lead shielding is

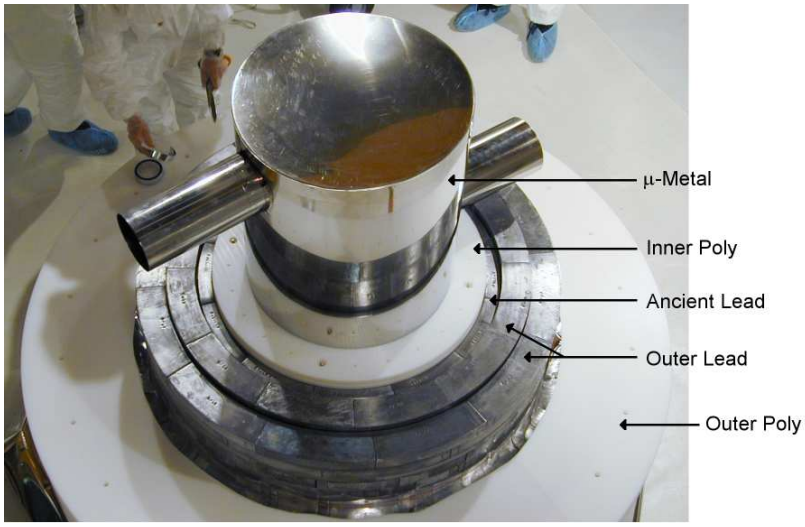

Figure 3.10: A view of the passive shielding as is it is being assembled at Soudan. The outer poly, 3 brick layers of lead, and the inner poly can be seen. The mu metal shielding, in which the detectors reside, is also shown.

Figure 3.11: A view of the top of the lead and poly shield as they are being lowered into place. On the left, Dave Hale guides the lead lid into place. On the right, Dave Hale motions the poly lid into place.

composed of inner and outer layers of lead bricks. Each layer is shaped in a cylinder surrounding the detectors. The lead bricks are arranged to minimize the number of lines of site through which a gamma might penetrate. The outer layer is composed of low-activity lead bricks from the Doe Run mine. The outer layer has a thickness of 17.8 cm (4.4 cm) on top and bottom and an outer thickness of 17.8 cm (4.4 cm) on the sides. The inner lead brick layer is also cylindrically shaped and is composed of ancient lead from a shipwreck off the coast of France. ^{210}Pb, which has a half-life of 23 years, is no longer present in ancient lead. The ancient lead is used to attenuate backgrounds due to radioactive isotopes with half lives of order one hundred years or less in the Doe Run lead shielding. The inner layer has a thickness of 4.4 cm on top, on bottom, and on the sides. The layer of lead composing the top of the cylinder rests upon a supporting 2.5 cm thick copper layer.

The inner layer of polyethylene shielding is inside the lead layers and provides further neutron moderation. The inner poly is 10 cm thick on the side, 7.6 cm thick on top, and 10.2 cm thick on bottom. Simulation indicates that in the absence of the inner poly, neutrons that penetrated the lead would be likely to scatter repeatedly within the lead, since each scatter would only slightly reduce the neutrons energy. Neutrons would then have multiple opportunities to interact and leave significant energy in the detectors. The inner poly, with its high density of hydrogen and carbon, is far more effective at moderating the energy of the neutrons.

Within the inner poly layer, a can of "mu metal" can [103] was placed to reduce the external magnetic field that might effect the performance of part of the cold electronics, the SQUIDs. The mu metal can has a height of 70.7 cm and a diameter of 61.4 cm. The detectors reside within the mu metal shield. Testing of the mu metal shield with a gauss probe at UCSB showed that it reduced the background magnetic field by about two orders of magnitude.

3.3 Detectors

A CDMS detector, or Z(depth)-sensitive Ionization and Phonon detector (ZIP), is a cylindrical crystal composed of germanium or silicon. The detectors are so named because they have sensitivity to the depth in the detector at which an event occurs in addition to sensitivity to the x-y location of the event. Sensitivity to the depth of an event is particularly important because CDMS's main background is from electromagnetic events on the surface of the detectors. Each crystal has a height of 1 cm and a radius of 3.81 cm. A germanium detector has a mass of 0.25 kg and a silicon detector has a mass of 0.10 kg. The detectors are cooled to an operational temperature of 50 mK by a dilution refrigerator. The detectors are operated in units of 6 detectors; each unit of 6 detectors is said to be a tower. The recent data run at Soudan from March 25th 2004 through August 8th 2004, on which this thesis is based, utilized two towers of detectors. Figure 3.12 shows one of the towers just before it was installed. Figure 3.13 gives a schematic side view of how the detectors

Figure 3.12: Dan Bauer (right) and Vuk Mandic (left) look at one of the towers. The 6 detectors in the tower are in the 1.2 cm tall copper houses at the bottom of the copper structure. The long, vertical, copper extender cards carry signals between the detectors and the cold electronics which are mounted like a crown on top.

are installed and connected to the cold electronics described in section 3.4.1. Figure 3.14 shows the two towers installed and connected to the electronics just before the detector volume is closed in preparation for the beginning of the data run. A detector naming scheme listing the tower number followed by the detector's slot, counting from the top, will be used. In this naming scheme, the first detector in tower one is termed T1Z1. The following detectors are germanium: T1Z1, T1Z2, T1Z3, T1Z5, T2Z3, and T2Z5. The remaining six

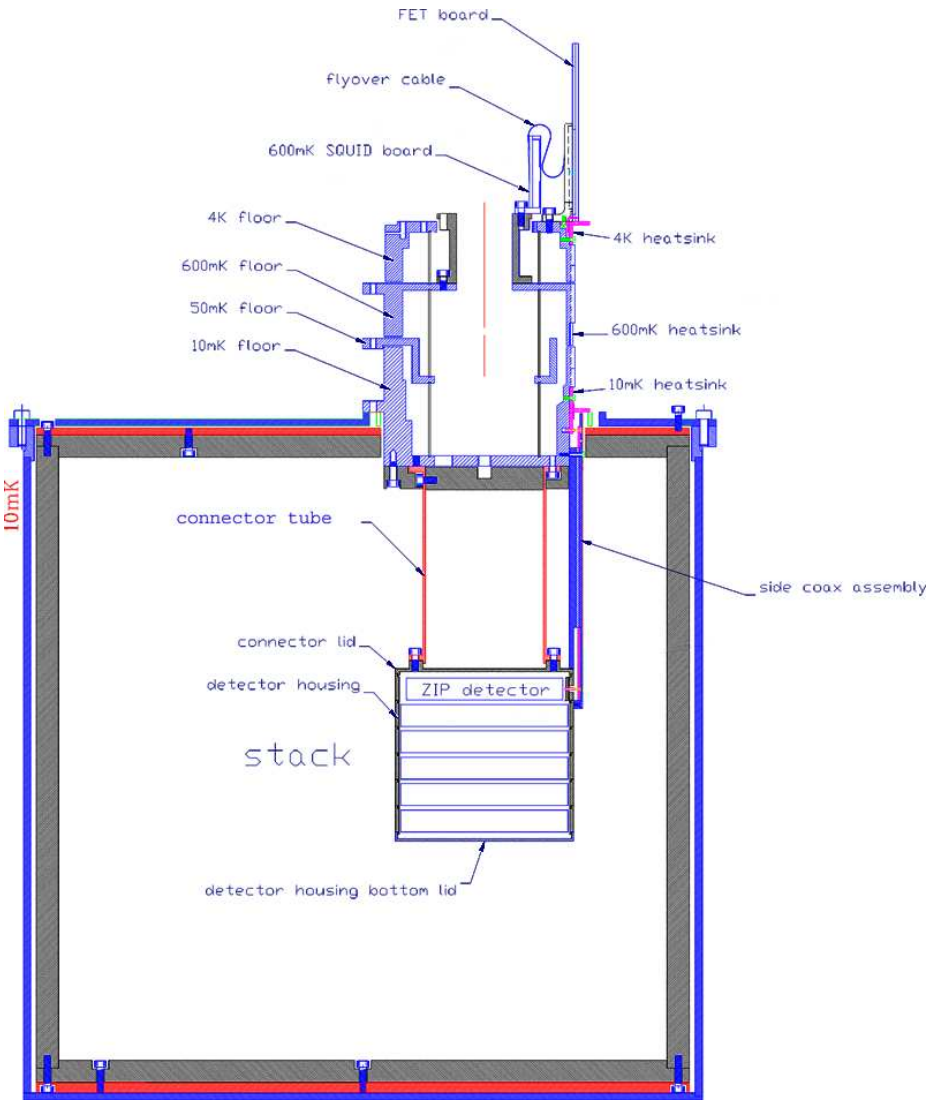

Figure 3.13: A schematic view of the side of a single tower installation. A sample SQUET card, described in section 3.4.1, can be seen at the top. Below the SQUET cards are the heat sinks. A connector tube connects the detector housing with the 10 mK sink. Extender cards, one of which is shown, surround the connector tube and carry signals between detectors and the SQUET cards.

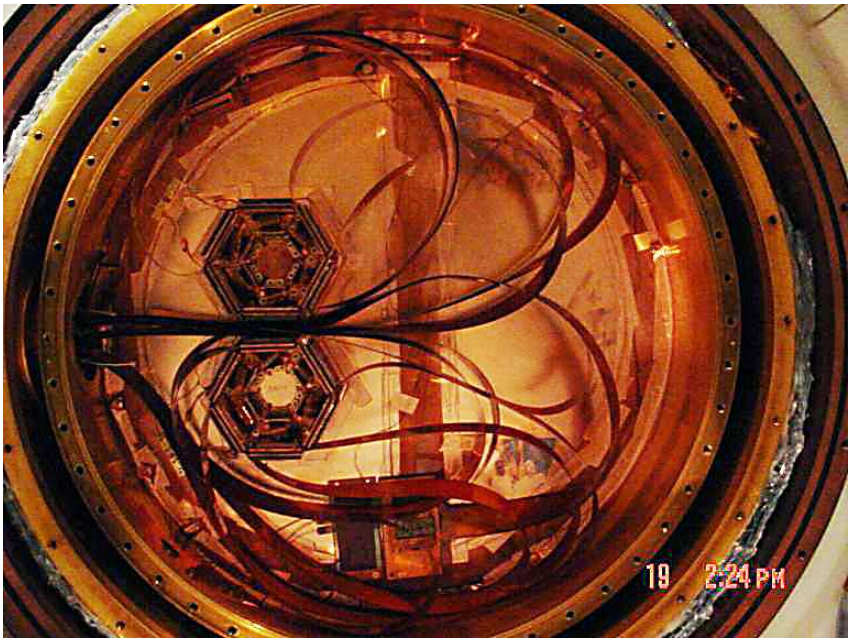

Figure 3.14: A picture looking down at the installed towers. Tower 1 is on top and tower 2 on bottom. The 6 SQUET cards, one for each detector, have been installed are are the "crown" on top of the towers. The long lines attached to the SQUET cards running out the E-stem on the left are the strip lines which carry the detector signals to the detector electronics.

detectors are silicon. Occasionally the detectors will be referred to by number where Z1 is the first detector in the first tower and Z12 is the last detector in the second tower. The theses of Clarence Chang [104], Vuk Mandic [105], and Gensheng Wang [106] are good sources for detailed descriptions on the design and performance of the ZIPs.

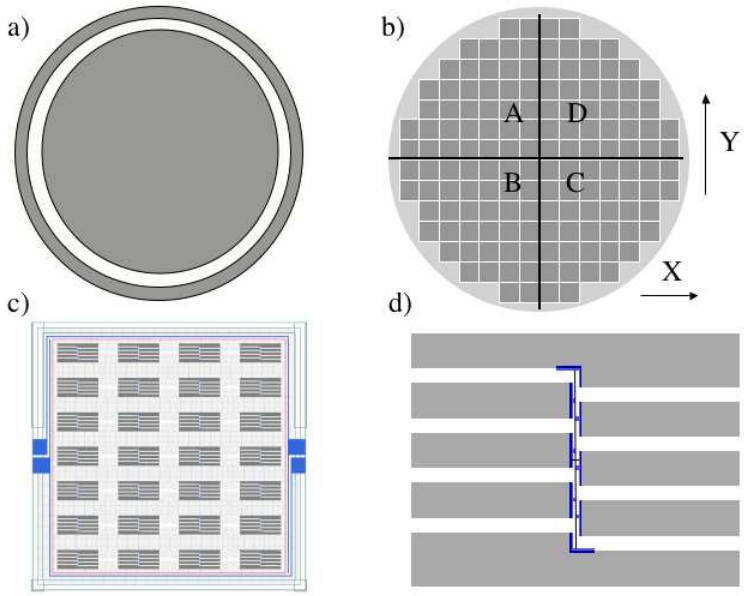

Figure 3.15: Schematic of the ZIP ionization electrodes and phonon sensors. Part a) shows the ionization electrodes on the under-side of a detector. Part b) shows the four phonon quadrants. Each phonon channel has 37 photolithographic sensor dies. Part c) shows one sensor die. Each die contains 28 QET phonon sensors. Each athermal phonon sensor contains 8 Al phonon-collector fins attached to a 1 μm thick T transition edge sensor. Part d) shows a zoom of 8 aluminum collector fins transition edge sensor.

3.3.1 Ionization Electrodes

The detectors are instrumented with inner and outer ionization electrodes on the lower side of the detector, with the top side of the detector acting as a ground. Part a) of Fig 3.15 shows the electrodes on the bottom side the a detector. The inner electrode is disk-shaped with a radius of 34.4 mm and centered on the center of the detector. A 1±0.1 mm trench separates the inner electrode from the outer electrode. The outer electrode is ring-shaped with an inner radius of 35.4 mm and an outer radius of 37.4 mm. The outer electrode then has a width of 2.0±0.4 mm. The distance from the edge of the crystal to the outer edge of the trench is 2.7±0.2 mm. The electrodes are used to define a fiducial volume. If events in the inner half the the trench are assumed to deposit their energy in the inner electrode, then the fraction of the detector contained within the fiducial volume can be estimated to be 83.9%. Under the more conservative assumption, that events within the trench deposit part of their energy in the outer electrode, the fraction of the detector contained within the fiducial volume is estimated to be 81.5%.

The germanium ZIPs are operated with a -3 V bias and the silicon ZIPs are operated with a -4 V bias. Ionizing particles free electrons creating electron-hole pairs. The negative bias voltage drifts the freed electrons towards the phonon sensors on the top surface and the holes towards the ionization electrodes on the lower surface. The drifting charges interact with the nuclei in the detector producing Luke-Neganov phonons[107][108]. Luke-Neganov phonons are low

energy, ballistic phonons that have a combined energy equal to the work done by drifting the charges across the crystal[109]. Since the Luke-Neganov phonons are ballistic they tend to arrive at the phonon sensors more quickly than the primary phonons from events within the bulk of the detector events (bulk events).

3.3.2 Phonon Sensors

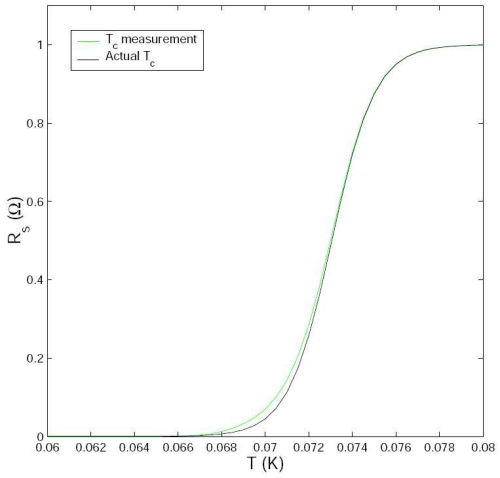

Figure 3.16: Comparison of the resistance verses temperature for a phonon sensor near the critical temperature. The light green is the measured curve for the combined QETs in a single phonon channel while the black is the transition curve of a single QET. Operating multiple QETs in parallel reduces the sharpness of the transition curve.

Four phonon sensors are located on the top surface of the detector. The sensors are sensitive to phonons produced by interactions within the detector. Part b) of Fig. 3.15 is a schematic of the detector showing the layout of the four

phonon channels. Each phonon channel had 37 photolithographic sensor dies. Each die contains 28 Quasi-particle assisted Electrothermal feedback Transition edge sensors (QETs). Each QET contains 8 aluminum phonon-collector fins attached to a 1 μm thick tungsten transition edge sensor. The transition edge sensors within a phonon quadrent are wired in parallel and are operated at their critical temperature. The critical temperature is the temperature at which a material rapidly transitions between normal and superconducting. The tungsten, which has had iron nuclei implanted to adjust and make uniform its critical temperature, has a critical temperature of about 60-80 mK. At the critical temperature very small temperature changes in the tungsten result in a large change in resistance. Low tungsten heat capacity, and hence low mass, is important to achieve significant temperature change for a small energy deposition. Figure 3.16 shows the resistance change as a function of temperature for the QETs in a phonon channel.

The detectors, which are operating at 50 mK, act as a thermal bath for the QETs. The tungsten is operated under a bias current. The bias current is chosen so that power dissipated by the bias current warms QETs to their critical temperature. For the phonon sensors to be operationally stable, a method for making the transition edge of the tungsten maintain a point of stable equilibrium is required. Stability can be accomplished if the QETs are voltage biased rather than current biased. When voltage biased, the power dissipated is the voltage squared divided by the resistance. As the energy deposited by

an event warms the QETs, the resistance increases and the power decreases. Similarly, QETs are protected from cooling below their critical temperature because the power dissipated increases as the resistance decreases and the point of equilibrium becomes naturally stable.

At 50 mK, the aluminum collector fins are superconducting. Phonons from a particle interaction break apart Cooper pairs in the aluminum. The quasi-particles, or broken Cooper pairs, diffuse into the tungsten. Since tungsten has a lower band gap than aluminum, the quasi-particles are trapped in the tungsten and warm the tungsten. Since the tungsten is held at its critical temperature, the small change in temperature results in a large change in resistance. The resulting change in bias current is measured by an array of Superconducting Quantum Interference Devices (SQUIDs). SQUIDs are discussed below in the cold electronics section.

3.3.3 Sensitivity to Event Position

Phonons propagate through a crystal at speeds that allow the event position to be reconstructed. The x-y coordinates of an event can be independently triangulated from the arrival times and amplitudes of the phonon pulses in the sensors located over the four quadrants of a crystal. Figure 3.17 shows the two methods, using the relative start times and amplitudes of the phonon pulses from the four quadrants, for identifying the x-y coordinates of an event.

The physics of the phonon interactions in a ZIP allow one to infer the depth

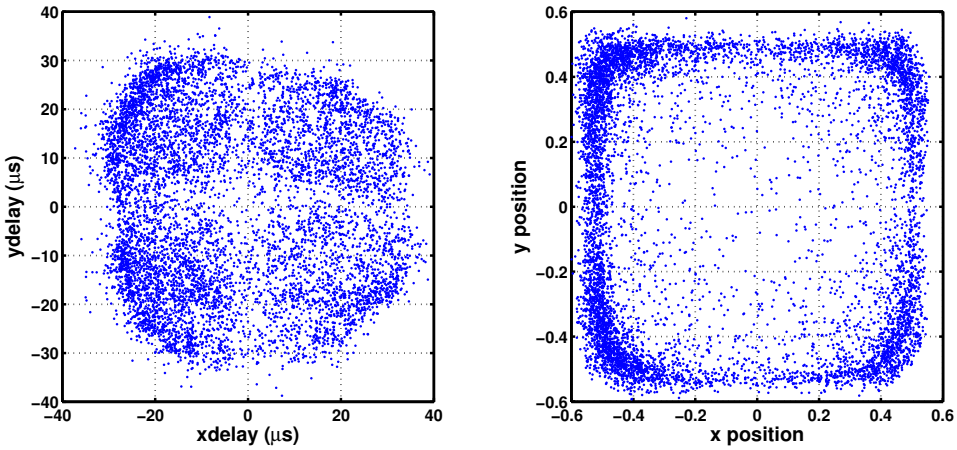

Figure 3.17: Delay and box position plots for T1Z2. The plot on the left shows the x-y position of an event as inferred from the relative start times of the pulses from the four phonon quadrants. The plot on the right shows the x-y position as inferred from the amplitude of the pulses from the four phonon quadrants.

of an event which allows background events on the surfaces of the detectors to be rejected. Electron and gamma events near surfaces of the ZIPs are tricky to reject as some of the ionization energy does not arrive in the electrodes causing the events to appear more like neutrons. Simulations have shown that the detector surfaces have a thickness of approximately 20 nm. Events deeper in the surface tend to lose less ionization energy than shallower events. Both the timing of the four phonon signals and the energy distribution among the four phonon channels provide useful information for rejecting surface events. Primary phonons produced by a particle interaction are typically high energy and do not propagate far before interacting and downconverting to lower en-

ergy. At sufficiently low energy, the phonons become ballistic and propagate more quickly. Phonons produced by event interactions near a detector surface interact with the metals on the surface causing the phonons to become ballistic more quickly. The phonons resulting from surface events arrive in the QETs more quickly causing the phonon pulses to arrive sooner, or start more promptly relative to the instantaneous ionization signal. The phonon pulses also have have quicker risetimes than for events within the bulk of a detector. As Fig. 3.15 illustrates, the phonon channels cover most of the surface area on the top side of a detector. Events occurring on the top surface will tend to deposit an increased fraction of their energy in the phonon quadrent in which they occurred.

3.4 Data Acquisition Hardware

The data acquisition (DAQ) is the system of hardware and software used to control, monitor, and obtain data from the detectors and veto photomultiplier tubes (PMTs). This Section is organized to describe the DAQ in the following order:

- The hardware used to control and obtain data from the detectors is described beginning with the hardware nearest the detectors, the cold electronics, in Section 3.4.1. The description then follows the path of the detector signals from the cold electronics on to the detector electronics

that are described in Section 3.4.2.

- The veto hardware is described in Section 3.4.3 by following the path of the veto signals from the PMTs through the hardware. A trigger on events depositing muon-like energy in multiple veto panels is also described.

- The software used to control and monitor the experiment and the software used to read out and package events is described in Section 3.5.

3.4.1 Cold Electronics

The term cold electronics refers to the SQUETs (the acronym is defined shortly) cards, which operate at cryogenic temperatures, and the striplines. A SQUET card is so named because it houses SQUIDs and field effect transistors (FETs). Phonon and ionization signals from a single detector first pass to a SQUET card; the phonon signals are amplified by the SQUIDs, and the ionization signals are amplified by the FETs. A SQUET card is shown in Fig. 3.18. The 6 SQUET cards for a tower of detectors are mounted to the top of a tower as is shown in Fig. 3.13. A SQUET card is shown in Fig. 3.18. SQUET cards are further described in [105].

A stripline carries a detector's phonon and ionization signals from a SQUET card out through the E-stem to the E-box. The E-stem and E-box are show in Fig. 3.19. Each of the 12 striplines are approximately 3 m long. On one side,

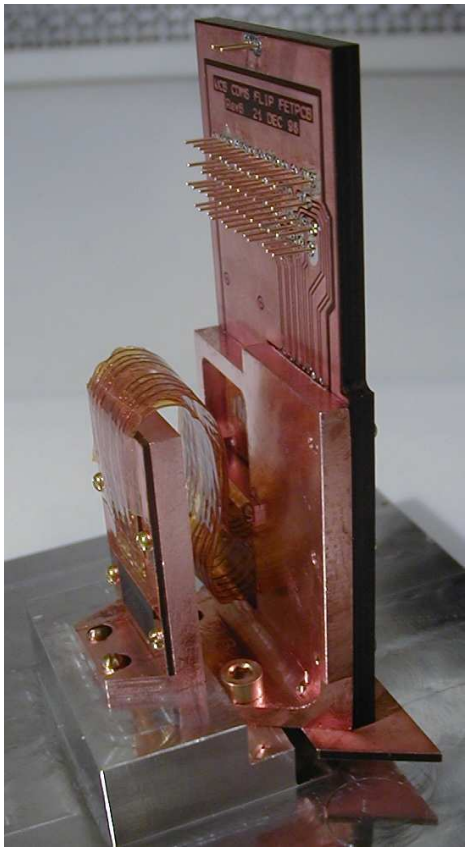

Figure 3.18: A picture of a SQUET card. A SQUET card is used to amplify the ionization and phonon signals from a detector. The ionization signals are amplified by field effect transistors (FETs) which are located in the enclosed copper area. The FETs operate at 4 K. The phonon signals are amplified by the arrays of SQUIDS. The SQUID arrays are located on the smaller part of the card to the left and operate at 600 mK. A SQUET card is mounted to the top of a tower of detectors. A stripline, which connects to the pins at the top of the card, carries the ionization and phonon signals out to the detector electronics.

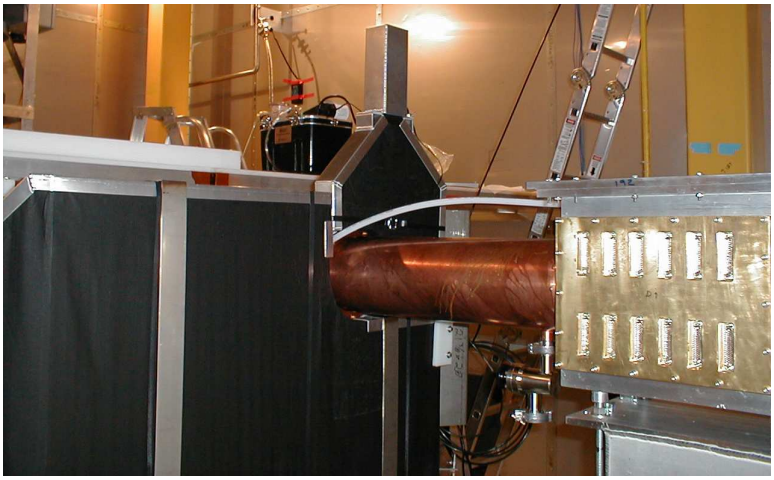

Figure 3.19: A picture the E-stem and E-box before the veto assembly was completed. The strip lines connected to the inside of the visible E-box connectors. Cables connecting the outside of the E-box connectors to the FEBs were not installed at the time of this picture.

a stripline connects to the pins at the top of a SQUET card, as are visible in Fig. 3.18. On the other side, a stripline connects to the inside of one of the connectors on the E-box. Cables run from the E-box to the front end boards (FEBs) and the detector electronics.

3.4.2 Detector Electronics

This section describes the functionality of the detector electronics by following how the detector electronics process an event in a detector, and it describes the physical setup of the detector electronics. The detector electronics is composed of 8 different types of devices which take signals from the detectors and

perform several tasks. These tasks include controlling of detectors, filtering of detector pulses, triggering on events, recording event information, and monitoring the performance of detectors and detector electronics. The 8 different devices that perform these tasks are: the front end boards (FEBs), receiver trigger filter boards (RTF boards), trigger conditioner boards (TCBs), trigger logic board (TCB), digitizers, a scaler, a history buffer, and slow ADCs. Figure 3.20 illustrates how detector information flows between the 8 devices. The functionality and the physical specifications of these 8 devices will be examined sequentially with particular attention to the TCBs and TLB which were designed and tested at UCSB.

- A single FEB controls a detector and SQUET card, and it amplifies the 4 phonon and 2 ionization detector signals from a SQUET card. A FEB controls all settings for a single detector. The settings controlled include detector bias voltage, SQUID bias and lockpoints, QET biases, LED flashing, and initial heating of the FETs on startup. A FEB amplifies a detector's signals and passes them to a RTF board.

  The 12 FEBs, which reside in the clean room about 1 m east of the E-box, receive detector signals from the E-box, amplify them, and pass the signals on to the RTF boards. Signals from a SQUET card proceed via a Belden type 9995 shielded differential twisted-pair cable from the E-box to a FEB. A FEB then amplifies each detector signal by a user-specified factor. After amplification, a detector's phonon and ionization

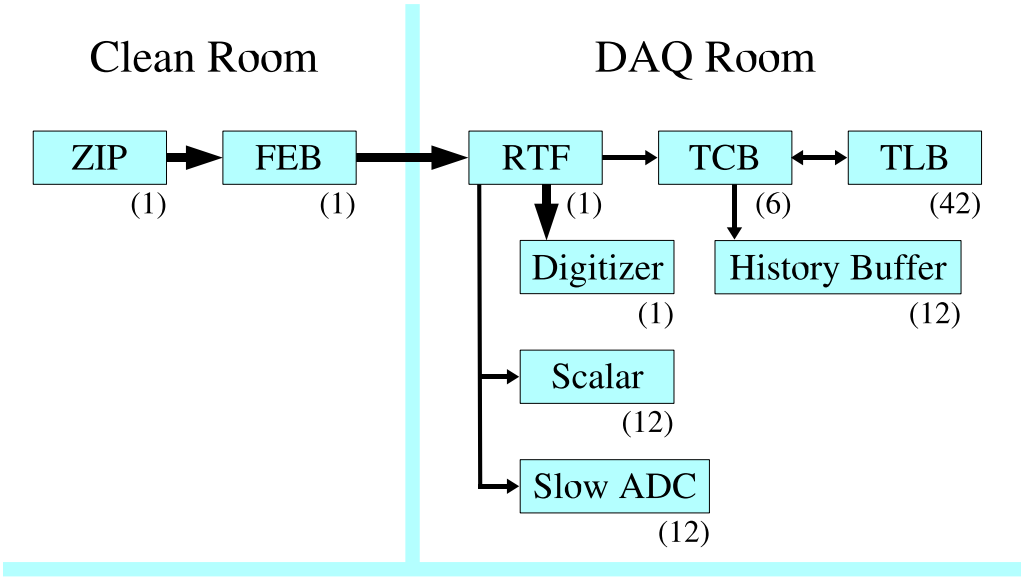

Figure 3.20: A schematic view of the detector electronics. The left (right) side shows the portion of the detector electronics that exists within the clean room (DAQ room). The thick (thin) lines denote the travel of signal pulses (thresholds, rates, and trigger records). The arrows show the direction in which the data travels. The numbers in parentheses indicate the maximum number of detectors whose signals are processed by a single board. For example, a single history buffer card records the event history of up to 12 detectors.

signals pass from a FEB to a RTF board on Belden type 9734 cables. A cable carrying a detector's signals departing from a FEB terminates at a patch panel on the southern wall of the clean room. Another Belden cable continues carrying a detector's signals from the patch panel through a hole in the northern wall of the DAQ room to a RTF board. The Belden 9734 cables are shielded differential twisted-pair cables.

The FEBs reside within a Fermilab-modified Weiner UEV-6021 9U crate.

The FEBs were designed at Fermilab [110] [111]. The crate containing the FEBs is powered by a Fermilab-designed power supply [112].

- The four purposes of a single RTF board are to compare a detector's signals with user-specified thresholds in order to determine whether an above-threshold event occurred, to filter a detector's signals, to pass the filtered signals to a digitizer, and to provide monitoring information to the appropriate devices. To determine whether an above-threshold phonon (ionization) event has occurred, a RTF board makes an analog sum of the pulses from the 4 phonon (2 ionization) channels and compares the summed pulse with user-specified high and low thresholds. The summed phonon pulse is also compared to a user-specified ultra-low threshold. When a threshold has been met, a RTF board generates a ~1 μs-wide logical-level signal. The logical-level signals from a RTF board are routed to a TCB, which is housed in the same crate, through the RTF crate backplane.

  A detector's signals are passed to a digitizer after the detector's RTF board filters the signal and introduces an offset. A 0.5 MHz low-pass filter is utilized. Offsets of -2 V are introduced to utilize the full dynamic range of a digitizer. The filtered signals are passed via a Fermilab-made cable to a digitizer. The end of a cable connecting to a RTF board has a AMP 226298-6 6-pair flat coaxial connector; the end that connects to a digitizer breaks out into 6 LEMO connectors.

A RTF board provides monitoring information to a scaler and slow ADCs. A copy of the logical-levels are sent to a scaler via a 34-pin ribbon cable that connects to the back of the crate housing a RTF board. The ribbon cable carries all logical-levels for a single tower's 6 RTF boards. A RTF board's thresholds and offsets are sent to a slow ADC via 2 96-pin ribbon cables that connect to the back of the crate housing the RTF boards.

The 12 RTF boards are housed inside a Fermilab-modified Weiner UEV-6021 9U crate. The upper half of Fig. 3.21 shows the RTF boards in the 9U crate. The 6 RTF boards that handle the Tower 1 detectors occupy slots 2 through 7 on the left hand side of the crate. The RTF board occupying slot 2 handles T1Z1, etc. The 6 RTF boards that handle Tower 2 detectors occupy slots 14 through 19 on the right hand side of the crate. The RTF board occupying slot 14 handles T2Z1, etc. The RTF boards were designed at Fermilab and are further described here [111]. The crate housing the RTF boards is powered by an identical Fermilab-designed power supply as the crate housing the FEBs [112].

The more complex detector electronics used to trigger the experiment, the TCBs and the TLB, will be described after the simpler digitizers, scalers, and slow ADCs are described.

- The detector pulses are digitized by a Struck SIS3301 8-channel, 14-bit analog to digital converter (ADC) or digitizer. The 4 phonon and 2

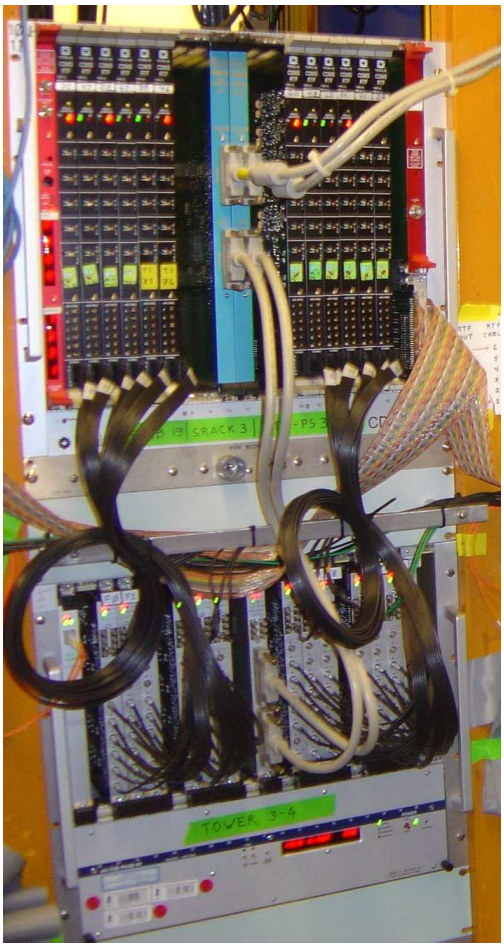

Figure 3.21: A picture of the RTF crate (upper) containing the RTF boards and TCBs and the VME crate (lower) containing the digitizers and history buffer. The 2 blue cards in the center of the upper crate are the TCBs. The 6 black cards to the left of the TCBs and the 6 black cards to the right of the TCBs are the RTF boards. Black ribbon cables carry the phonon and ionization pulses from the RTF boards to the digitizers in the VME crate. The lower (upper) SCSI cables from the TCBs head to the history buffer in the VME crate (head to the TLB which is not shown). The left-most card in the VME crate is the VME-to-PCI bridge.

ionization signals were fed into a single digitizer leaving 2 unused channels in each digitizer. This simple arrangement yielded a one to one mapping between detectors and digitizers. The digitizers operate at a minimum sample rate of 65 MHz which is far faster than it is necessary to sample detector pulses. The digitizers use an averaging algorithm to increase the effective resolution to 14 bits resulting in a 1.25 MHz sampling rate. 512 pretrigger samples and 1536 post trigger samples are obtained for a recorded time window of 1.64 ms.

The 12 digitizers reside in a Weiner VME 6023/611 crate. The lower half of Fig. 3.21 shows the Weiner VME 6023/611 crate containing the digitizers. The digitizers in the VME crate are read out via a SBS model 620 VME-to-PCI bridge.

- A scaler is used to monitor the rate of above threshold interactions in the 12 detectors. The scaler records changes in the logical-levels from the 12 RTF boards, indicating when the detector phonon (ionization) hi, low, and ultra low (hi and low) thresholds are exceeded.

  A Joerger VS64 scaler resides in a Weiner VME 6023/611 crate and accepts 2 34-pin ribbon cables. Each ribbon cable carries the 30 logical-levels from the RTF boards for a single tower of detectors from the backplane of a RTF crate to a scaler. The scaler is read out via a SBS model 620 VME-to-PCI bridge.

- A slow ADCs is used to monitor RTF threshold levels and SQUID offsets. Each of the 2 inputs on a slow ADC record the following information for 3 detectors: changes in the 3 phonon threshold levels for each detector (hi, low, and ultra-low), changes in the 2 ionization threshold levels for each detector (hi and low), and changes in each detector's 4 phonon SQUID offsets. Each tower's threshold levels and SQUID offsets are monitored by a single slow VMIC3128 slow ADC. Each VMIC3128 slow ADC has 2 32-channel differential inputs which are scanned by a 14-bit ADC.

  Two VMIC3128 slow ADCs reside in the same Weiner VME 6023/611 crate as the scaler, and each slow ADC accepts 2 96-pin ribbon cables. Each ribbon cable carries the 15 (9 phonon and 6 ionization) differential threshold levels as well as 12 differential SQUID offsets for 3 detectors. The slow ADCs are read out via a SBS model 620 VME-to-PCI bridge.

- Because of its simpler description, the history buffer will be described before the TCBs from which it receives detector trigger records.

  The history buffer keeps a digital record as a function of time of when detector phonon and ionization thresholds are exceeded. The history buffer is operated with a sample rate of 1 MHz. The history buffer typically has a pretrigger record containing all times phonon and ionization thresholds were exceeded since the last event. The minimum pretrigger time window is about 5 ms. The history buffer's posttrigger time window is about 20 ms.

A Struck SIS3400 VME board functions as the history buffer. The history buffer resides within the same VME crate as the digitizers and is visible in the lower portion of Fig. 3.21. A SIS3400 board accepts 2 SCSI-2 inputs. A SCSI-2 cable carries the logical levels for each tower from the TCBs to the history buffer.

The TCBs and TLBs were designed and tested at UCSB. Therefore, they will be described in more detail in the two immediately following subsections. The TCBs and TLBs were designed by Sam Burke to provide flexible experimental triggering. Sam Burke and I were responsible for insuring that the TCBs and TLBs functioned properly. The left (right) hand side of Fig. 3.22 shows a TCB (TLB) in our test setup in the lab at UCSB. Additional information on the design of the TCBs and TLBs can be found here [113].

The TCBs and TLB were designed to provide flexible triggering options to accommodate future, unforeseen triggering schemes. The flexibility is accomplished by having field programmable gate arrays (FPGAs) handle detector trigger logic. An FPGA's logic is loaded from an erasable programmable read-only memory chip (EPROM). TCB and TLB detector trigger logic can by changed by inserting EPROMs with updated logic.

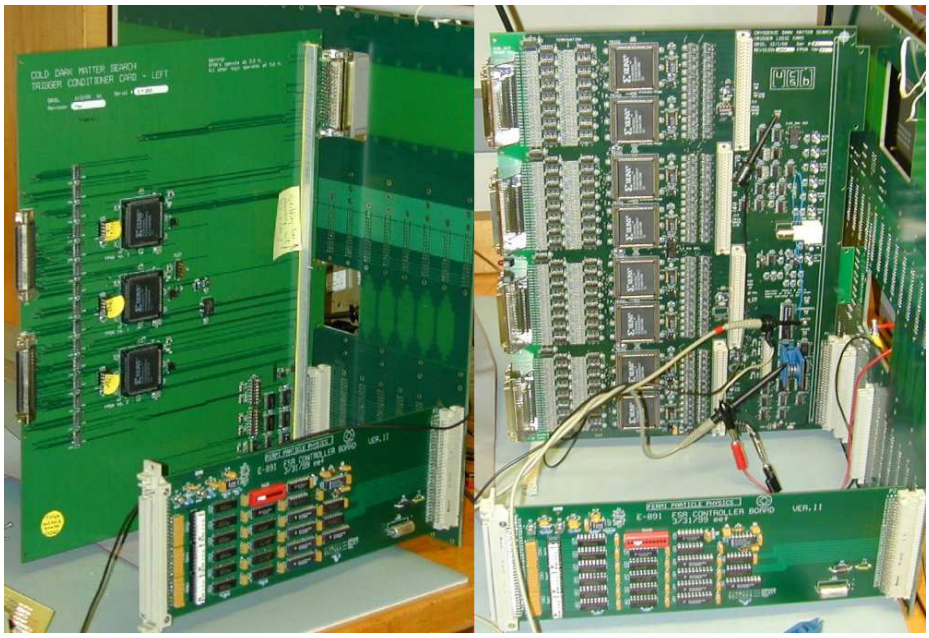

Figure 3.22: A picture of a TCB (left) and TLB (right) as they were being tested at UCSB. The large, square, black chips on the TCB and TLB are FPGAs that handle that contain the board logic. To the right of the FPGAs are EPROMs containing the logic which is loaded into the FPGAs on powerup. The upper (lower) SCSI-2 connector visible on the front of the TCB carries the logical-levels from the TCB to the TLB (history buffer). The SCSI-2 connectors on the front of the TLB receive the logical-levels from up to 7 TCBs. The connector on lower back (right) end of the boards is where the boards receive commands from software through the controller card. A Fermilab-designed controller card can be seen in the first, right-most slot in front of the TCB and TLB.

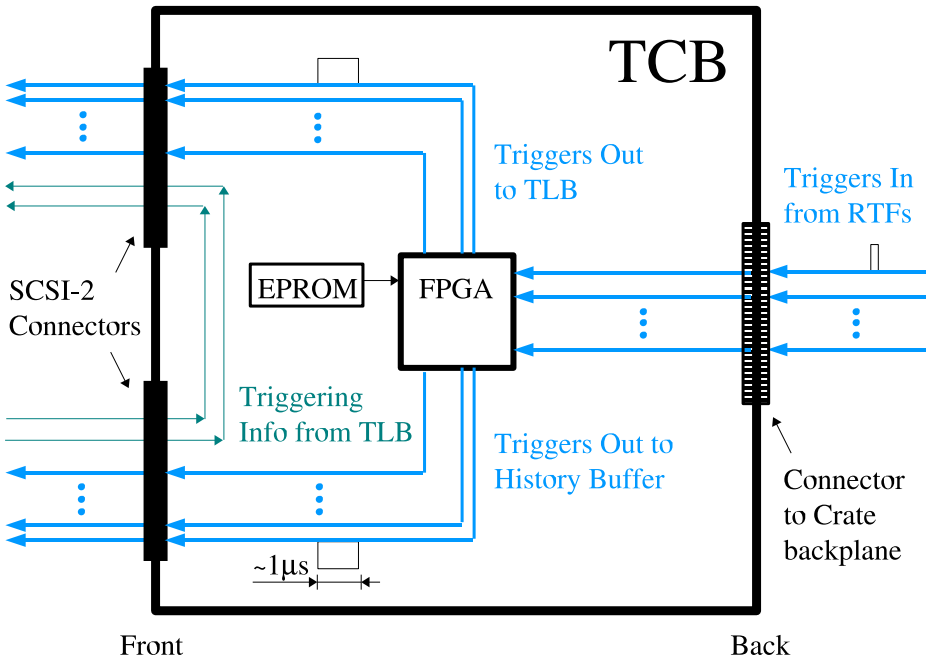

Figure 3.23: A schematic view of a TCB. The thick, light blue lines show how the logical-levels from the RTF boards passes through a TCB to the TLB and history buffer. Logical-levels from the RTF boards are stretched to a width of 700 to 900 ns. FPGA logic is loaded from the EPROM upon powerup. In the event that the rate of above-threshold events in a particular detector channel is unreasonable, the FPGA can disable passing that channel's logical-levels to the history buffer. Two lines on each TCB are used as a conduit to pass triggering information from the TLB to the history buffer.

TCBs

A TCB serves 4 purposes: it creates a differential copy of the logical-levels received from an RTF board, it sends a differential copy of the logical-levels to a history buffer, it sends a differential copy of the logical-levels to the TLB, and it routes TLB specific triggering information (global trigger, trigger enable, random trigger, veto pulser trigger, and veto multiplicity trigger described in Section 3.4.3) from the TLB to the history buffer. (The global trigger, trigger enable, random trigger, and veto pulser trigger are described in the TLB description to follow.) Figure 3.23 shows a schematic view of how a TCB works. A TCB takes a logical-level from an RTF board and creates a 700 ns to 900 ns-wide differential TTL logical-level. The differential logical-levels are sent to the history buffer and the TLB.

TCBs are designed in two flavors, left or right, to accommodate the different pin mappings from the RTF boards. The 6 RTF boards in slots 2 through 7 have a different mapping through the crate backplane to the TCB in slot 10 than the 6 RTF boards in slots 14 through 19 have through the backplane to the TCB in slot 11. A left (right) TCB is designed to have the mapping needed to function properly in slot 10 (11). Tower 1 (2) detectors were handled by a left (right) TCB. The upper portion of Fig. 3.21 shows the 2 TCBs in the 9U crate containing the RTF boards. Further information on the design of the TCBs can be found here [113].

The differential logical-levels proceed from a TCB to the history buffer and

TLB on SCSI-2 cables.

TLBs

The primary function of the TLB is to determine when to create a global trigger, the signal for the DAQ software (described in Section 3.5) to retrieve and save event information from the detector electronics and veto electronics (described in Section 3.4.3). Figure 3.24 shows a schematic view of how the TLB works. A TLB forms a global trigger using the TTL logical-levels from up to 7 TCBs as well as 3 external TTL logical-levels from the random trigger line, the veto multiplicity trigger line, and the veto pulser trigger line. The TLB creates a global trigger based on the OR of the logical-levels from the TCBs and the 3 external logical-levels. The external logical-levels pass directly into the OR, while the TLB has 4 different triggering modes for each detector's phonon and ionization logical-levels. The 4 modes are: off, high, low but not high, and low (default). The TLB was operated with all the detector phonon (ionization) channels in the low (off) trigger mode. Thus, the TLB created a global trigger whenever any detector's phonon low threshold logical-level transitioned from low to high, or when any of the 3 external logical-levels transitioned from low to high.

Only 2 of the external triggers, the random and veto multiplicity triggers, occur while taking WIMP-search data or calibration data. The random trigger is used to obtain information on detector phonon and ionization channels noise

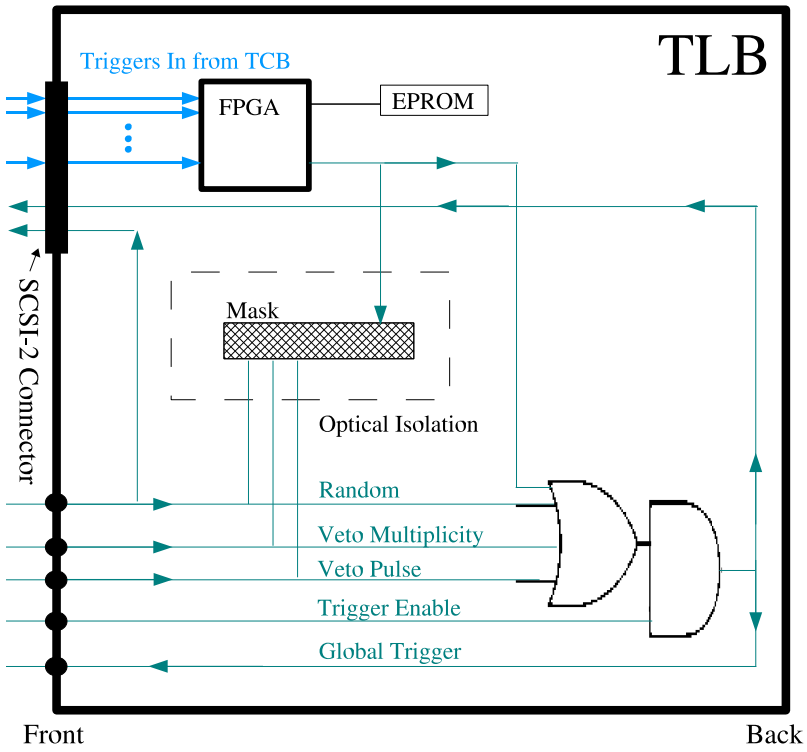

Figure 3.24: A schematic view of a TLB. The thick, light blue lines show the detector logical-levels from a TCB arriving on a SCSI-2 cable. The light green lines show the triggering information. FPGA logic, used to form triggers from the detector logical-levels, is loaded from the EPROM upon powerup. Two lines on the SCSI-2 connector to each TCB are used to feed back triggering information through each TCB to the history buffer. The mask, which was not used, carries a record of what caused the trigger.

profiles. The veto multiplicity trigger is used to trigger on interesting veto events and is described in Section 3.4.3.

After a global trigger has been created, no additional global triggers can occur until after the TLB has been reset by toggling its trigger enable line. The enable line is toggled by pulling the enable line low then high.

The TLB feeds the trigger enable logical-level, global trigger logical-level, and the 3 external triggers' logical-levels through the SCSI-2 cables to the TCBs to be routed to the history buffer. The trigger enable logical-level, the veto pulser trigger logical-level, and the veto multiplicity trigger logical-level are fed through left TCBs to the history buffer from the odd numbered SCSI-2 TCB connectors (1, 3, 5, and 7) on the TLB. The global trigger logical level and the random trigger logical level are fed through right TCBs to the history buffer from the even numbered SCSI-2 TCB connectors (2, 4, 6, and 8) on the TLB. The left (right) TCB, which is connected to the first (second) SCSI-2 connector, passes these logical-levels to the history buffer.

The trigger logic board was designed to be flexible and has other features which were not used. The logic controlling the detector phonon and ionization trigger options is contained in EPROMs which load the logic into FPGAs on powerup. Each EPROM contains the triggering logic for 1 tower which is loaded into 1 FPGA. New triggering modes can be introduced by changing the EPROMs. The upgrade for the next WIMP-search run included replacing the phonon high trigger mode with a new trigger mode. Another unused feature

of the TLB is the ability to cancel a user selected fraction of veto coincident triggers. The TLB also creates an unused record of which detectors and towers were hit. The record contains a copy of the random and veto trigger bits as well as an 8-bit ordinal number that is incremented after each event. The upgrade for the next WIMP-search run includes utilizing this record.

Control Hardware

Two additional hardware devices, GPIB boxes and controller cards, are used by the DAQ software (described in Section 3.5) to control the FEBs, RTF boards, TCBs, and TLB. Two custom GPIB boxes, which were built at Berkeley, control these boards. Each box is identified by a unique GPIB number. Commands to one of the GPIB boxes are routed to the 9U crate located in the clean room containing the FEBs. Commands to the other box are routed to all the 9U crates in the DAQ room which contain the RTF boards, the TCBs, and the TLB.

Each of the Fermilab-modified 9U crates has a controller card in the first, right-most slot that receives the GPIB commands from its GPIB box. Each controller card has a subrack number that is user selected to be unique. A controller card parses each received GPIB command for its subrack number to determine if that command was intended for the controller card's crate. GPIB commands for a particular crate are passed by the controller card along that crate's backplane to the waiting boards. Each GPIB command contains the

crate slot number of the board for which the command was issued.

3.4.3 Veto Electronics

I'll begin describing the veto electronics by following the veto signals from the veto panels to the digitizers, history buffer, and veto scaler. A schematic of the veto layout is shown if Fig. 3.25. Events in the 40 veto panels are detected out by Hammamatsu R329-02 photomultiplier tubes (PMTs). The PMTs are powered by a CAEN SY2527 high voltage supply located in the DAQ room. Figure 3.26 shows a picture of the veto electronics, including the CAEN power supply, that was taken before the veto trigger was added. The high voltage supply is controlled via the ethernet cable that runs down the right side of the figure. Signals from the PMTs pass up to the DAQ room on long 60 ft coaxial cables to a patch panel.

The remainder of the veto electronics resides in the DAQ room. The coaxial cables terminate in a patch panel. From the patch panel the signals proceed on LEMO cables to integrate and stretch amplifiers. These amplifiers, which were designed at UCSB by Raymond Bunker, make two copies of the PMT signals. One copy is passed to 8 channel LeCroy 623B discriminators, while the other is integrated, conditioned, and driven out to a set of Joerger VTR812 5 MHz digitizers. The veto pulses are integrated because the integral of the veto pulse provides a better estimate of the event energy than the pulse height. This is because reflections within the veto panels result in double peaked pulses.

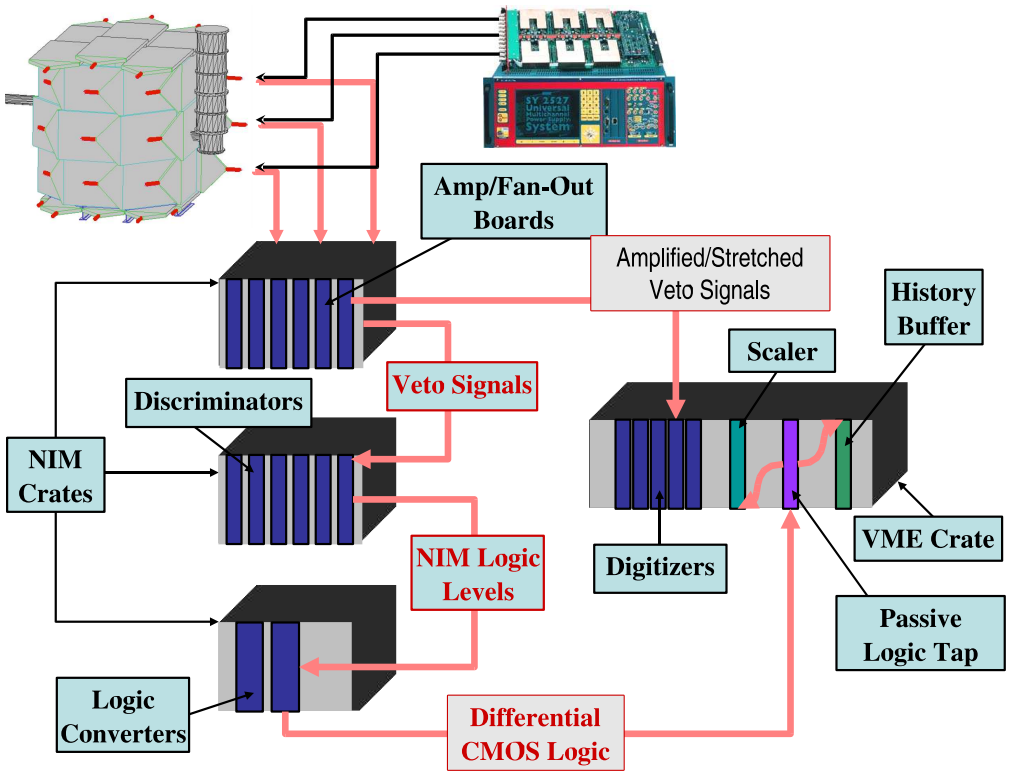

Figure 3.25: A schematic view of the veto electronics. Modified from a schematic made by Raymond Bunker.

Now I'll trace the path of the veto pulses from the discriminators to the veto history buffer and veto scaler. Discriminator thresholds are set so that muons that just clip the edge of a panel are generally above threshold whereas most of the ambient gamma background is below threshold. This translates to the minimum possible threshold of -30 mV. NIM levels from the discriminators pass to 2 logic converter boards designed by Sam Burke. The converter boards act as a giant OR for the discriminator levels. The discriminator levels for 20 veto panels are passed to the first converter board. The resulting OR of the first converter board along with the discriminator levels for the other 20 veto panels are passed to the second converter board. The resulting OR of the second converter board defines whether an event is veto coincident. The converter boards also convey a record of veto panel hits via SCSI cables to a Struck SIS3400 veto history buffer through a passive logic tap. The veto history buffer records veto hits with 1 μs precision. The logic tap was designed at UCSB by Sam Burke. The logic tap passes the creates a second copy of the veto panel hit record which is routed to a Struck SIS3800 veto scaler via SCSI cables. The veto scaler is used to keep track of each veto panel's rate of above threshold interactions.

A veto multiplicity trigger was added to record events that deposit muon-like energy in multiple veto panels. These events occur far less frequently than ambient, over-threshold gammas (about 1 per hour versus 600 Hz) and are valuable for a couple reasons. First, these events can be used to quantify the veto's

Figure 3.26: A picture of the rack housing the veto electronics before the veto trigger was added. The high voltage power supply resides at the top of the rack and directly below it is the patch panel. Below the patch panel are the NIM crates which hold the veto amplifiers, discriminators, and logic conditioners.

efficiency for tagging muons. Second, they can be used to develop understanding the rate of muons and showers. Knowing the muon and shower rates is useful as a normalization for simulations of the external neutron background.

The veto multiplicity trigger is designed to trigger on events hitting more than one veto panel. This is accomplished by feeding a second output from the discriminators is fed into a LeCroy 380A multiplicity module. The multiplicity module, whose output is fed to the trigger logic board, is set at a multiplicity of 2. The output of the multiplicity unit is used to trigger the experiment on any event hitting at least 2 veto panels. 26.5% of all global triggers in WIMP-search data were due to the veto multiplicity trigger of which 28.2% were due to muons or showers yielding 106k recorded muons or showers.

3.5 Data Acquisition Software

The system of DAQ computers and software is used to configure and obtain data from the detector and veto hardware which were described in Section 3.4.2 and Section 3.4.3. Entirely new software was written for controlling the CDMS experiment at Soudan. The new DAQ software is composed of servers, GUIs (graphical user interfaces), configuration files, and an event builder. A server is a software program used to provide a specific service to client software that can be running on a different computer. The client software can be another server or a GUI. A GUI is a software program that utilizes a graphical interface to allow a human to communicate with a server. To help facilitate remote control, the

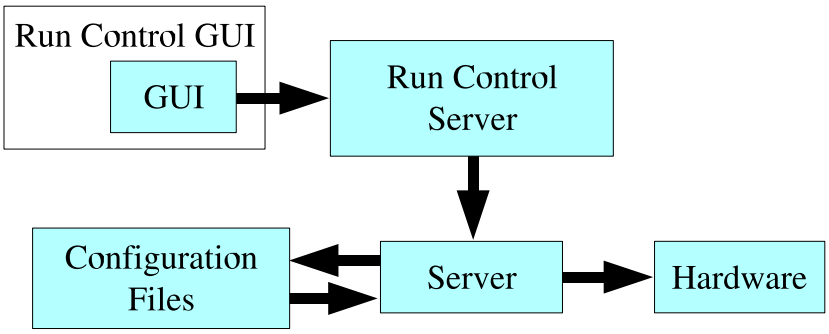

Figure 3.27: A figure depicting the GUI-server relationship used to control the experiment. GUIs used to control the experiment are only accessible from the run control GUI. A command issued on a GUI is conveyed to the relevant server via the run control server. The relevant server also obtains experimental run settings from configuration files. The relevant server issues the given command to the intended hardware device. The flow of data is omitted from this figure.

DAQ software utilizes a GUI-server relationship. This relationship is depicted by Fig. 3.27.

The top level GUI in CDMS is called the run control GUI. The run control GUI enables a user to configure, start, and stop data runs. A user can also "select down" from the run control GUI to access other GUIs that enable the configuration of the specific elements of detectors. User choices may be saved in configuration files.

A configuration file is a human-readable text file that tells a server how to configure a hardware device in a predetermined way. The contents of a configuration file is specified by a use of a GUI.

The event builder is a set of software programs, running on multiple com-

puters, that are used to retrieve event data and write it to disk. The event builder retrieves detector and veto data rapidly to insure little time is lost due to event readout.

3.5.1 DAQ Software Requirements

The software for the CDMS-II data acquisition (DAQ) had multiple design specifications that were intended to lead to a DAQ that was a significant improvement over the Labview-based DAQ used at the shallow Stanford site. While the new DAQ software is complex, it is designed so that a typical user will be able to obtain data at Soudan without developing more than a superficial knowledge of the DAQ software.

- The new DAQ was designed to accommodate a higher data rate. The Labview DAQ was limited to an event rate of about 1 to 2 Hz while the new DAQ can take data at an event rate of 30 Hz.

- The new DAQ should be easily scalable, because the number of detectors has and will continue to be increased. This will be achieved by designing the DAQ software to be modular.

- Since the experiment is situated underground and is usually directly accessible, for only 10 hours a day, it is essential for the experiment to be operable from offsite locations. The remoteness of Soudan, Minnesota also makes it desirable for members of the CDMS collaboration who are

not on shift at the experiment to be able contribute by remotely monitoring the experiment. Members of the collaboration use several different operating systems making it desirable for monitoring graphical user interfaces (GUIs) to run properly on all operating systems.

- It is also desirable for the DAQ to be written using modern object-oriented languages: Java and C++. CORBA was to be used to provide communication channels between the various DAQ servers that may be operated across different computer operating systems.

3.5.2 DAQ Software Implementation

The implementation of the DAQ software provides four primary functionalities: slow control of the experiment, monitoring of the experiment, fast event data collection, and inter-server communication.

These functionalities are achieved by using a modular software design coupled with a GUI-server framework. Each server was designed for a single, specific task. For example, one server is designed to control all custom electronics that are connected via GPIB (General Purpose Interface Bus), a second server is designed to control the CAEN high voltage supply, a third server is designed to provide data on the dilution refrigerator's performance, etc. Because the software is modular, the servers are able to operate on different computers dedicated to specific tasks. Hardware configuration settings are provided to the servers by means of configuration files. Table 3.1 shows which servers operate

Server/Process	Computer
Run Control Server (1)	`control`
GPIB server (2) CAEN server Configuration Manager (3)	`monitor`
Fridge Data Server (4) Monitoring Data Server (5)	`datasrv`
Fast Analysis Server (6) Event Builder Core	`builder`
Event Builder Nodes	`tower1` `vetocrate`
Fridge Monitoring Server	`Intellution`*

Table 3.1: A list of DAQ servers and computers illustrating the relationship between the DAQ servers (and the Event Builder) and the DAQ computers. The Intellution computer, indicated by a star, is not part of the DAQ. The numbers in parenthesis are used in Table 3.2 to indicated to which of the servers the GUIs communicate.

DAQ GUIs	
Slow Control	Monitoring
Run Control GUI (1)	Fridge Historical GUI (4)
Run Configuration GUI* (3)	Fridge Intellution GUI (4)
FEB GUI* (2)	Fridge Monitoring GUI* (5)
RTF GUI* (2)	Scalar GUI* (5)
Bake GUI* (2)	Veto GUI* (5)
	Noise Viewer GUI* (6)

Table 3.2: List of GUIs used to control or monitor the experiment. GUIs denoted with a * are only accessable from the run control GUI. The numbers in parenthesis indicate which of the servers listed in Table 3.1 may be communicated with by using these GUIs.

on which computers and Table 3.2 lists the main GUIs used to control and monitor the experiment by connecting the user to the various servers.

Slow Control

The DAQ slow control component allows the user to configure the detector and veto hardware. Slow control is handled by 4 servers and 5 slow control GUIs. These GUIs are listed in Table 3.2. The servers used to control the experiment are the run control server, the GPIB server, the CAEN server, and the configuration manager. After a server's functionality is described, any configuration files used by that server will be mentioned followed by a description of all GUIs that communicate with that particular server.

The run control server is designed to contain the primary logic for the experiment, by orchestrating the actions of the other servers and routing user commands issued on the slow control GUIs to the appropriate servers. The run control server also enforces security by verifying the user names and IP addresses of operators accessing the run control GUI. The run control server reads in the top level configuration file to obtain experimental run settings.

The run control GUI, from which all other GUIs controlling the servers are accessed, is the primary GUI for user control and monitoring of the experiment. Figure 3.28 shows a picture of the run control GUI. The run control GUI communicates directly with the run control server. The run control GUI allows the user to change the state of a data run, view recent error messages, access

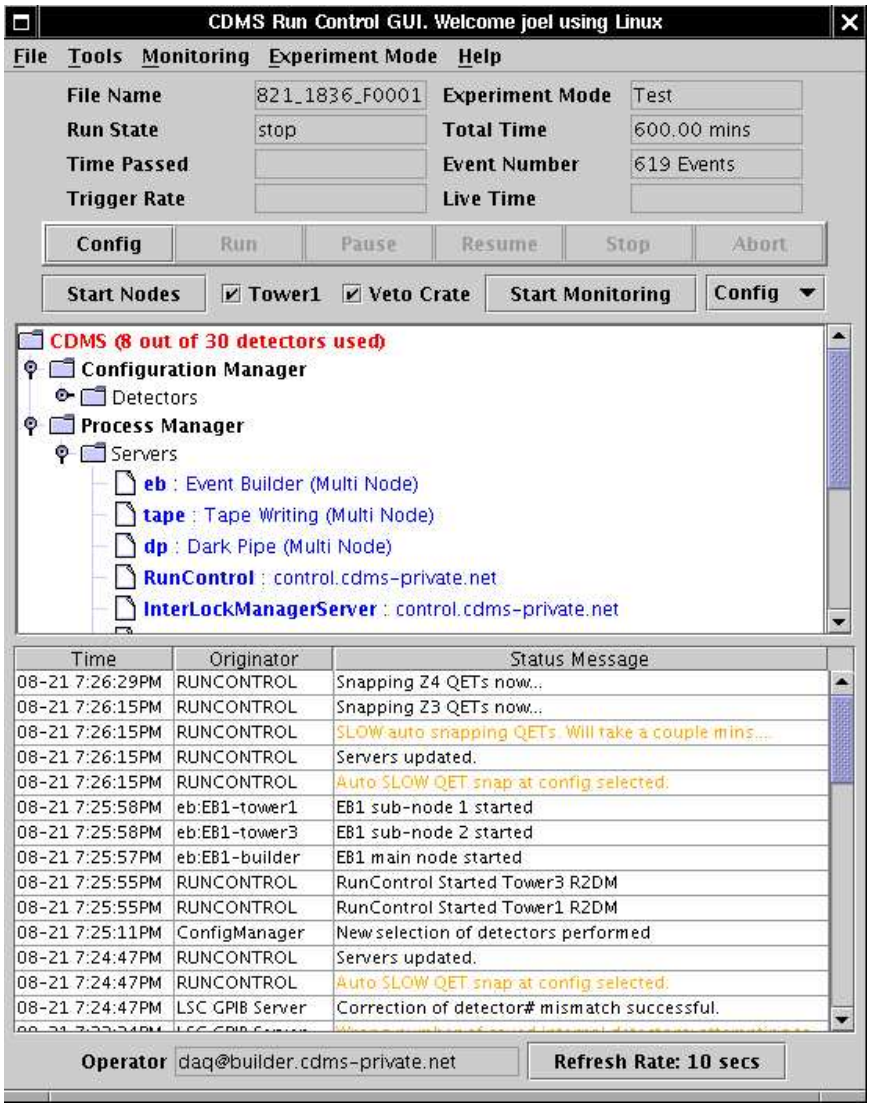

Figure 3.28: A picture of the run control GUI. The name of the last data file, experiment mode, event number, live time, etc are displayed at the top of the GUI. A users can use the buttons at the top of the GUI to change the experimental run mode. Just below the buttons are settings to restart the event builder nodes on `tower1` and veto crate. The central portion of the GUI lists the detectors being operated as well as the servers that are running. The lower portion of the GUI lists the recent errors and allows the user to specify a refresh rate for the GUI.

the GUIs that are connected to the other servers, as well as access the GUIs used to monitor the experiment. The run control GUI also allows a user to issue warnings on the error channel.

The GPIB server is responsible for executing any user command to change a setting on any device that is connected to the GPIB (FEBs, RTF boards, TCBs, and TLB which are described in Section 3.4.2). The GPIB server accomplishes this in three steps. First, it verifies that the current experiment state is one for which it is valid to modify hardware settings. This guarantees that the experimental configuration is not modified during a data run. Second, it formats the user command into the GPIB protocol. Third, it executes the command on the proper GPIB device.

When WIMP-search data or calibration data is to be obtained, the GPIB server is responsible for preparing the GPIB devices for each new experimental run state. To prepare the GPIB devices for a new experimental run state, the GPIB server reads in the top level configuration file and the FEB, RTF, TCB, and TLB configuration files that were written by the configuration manager. The GPIB server then sends all commands in the hardware configuration files to the appropriate GPIB devices. When the experimental run state is changed, the GPIB server's responsibilities include updating detector configuration, setting trigger thresholds, and setting the detector trigger modes.

The FEB and RTF GUIs allow the user to change settings on any FEB or RTF board. Figure 3.29 shows a picture of the FEB GUI. The FEB GUI allows

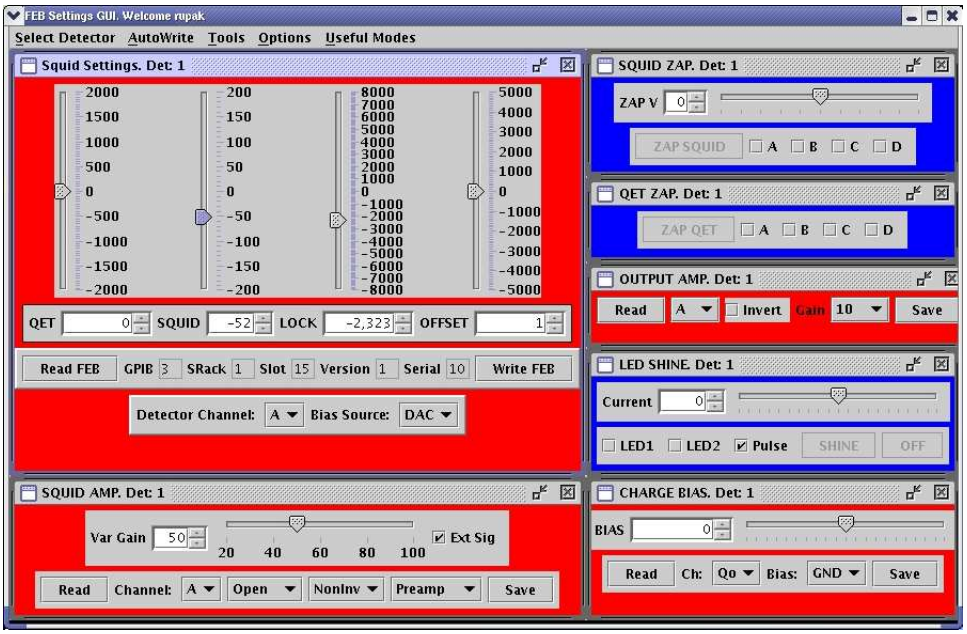

Figure 3.29: A picture of the FEB GUI. The left hand side of the GUI allows the user to set or read back SQUID settings for the 4 phonon channels. The 2 upper right portions of the GUI allow the user to ZAP the SQUIDS or QETs. The central right portions of the GUI controls the gain for the phonon and ionization channels as well as flashing of the LEDs. the LED flashing is better controlled by the bake GUI. The lower right portion of the GUI allows the user to set the bias voltage for the inner and outer ionization channels.

Figure 3.30: A picture of the bake GUI. The bake GUI allows the user to specify the LED flashing parameters used to neutralize the detectors.

the user to modify detector settings. The RTF GUI provides the user interface for controlling detector trigger levels and offsets. A user can save the current FEB or RTF settings into FEB and RTF settings files. These settings files can be reloaded into the GUIs to restore a useful detector state or be loaded into the run configuration GUI to specify experimental settings.

The bake GUI allows the user to specify the parameters used to neutralize the detectors. Figure 3.30 shows a picture of the bake GUI. The bake GUI is designed to give the user flexibility in neutralizing the detectors. This is accomplished by allowing each detector to have unique LED flashing parameters and by allowing any combination of detectors to be neutralized.

The CAEN server sets and monitors the CAEN high voltage supply's high voltages for the 40 PMTs in the veto system. It monitors the CAEN voltages and currents to make sure that they remain within a user prescribed range. The CAEN server provides monitoring information on the CAEN voltages and

currents and provides warning in the case of an incorrect voltage or current.

The configuration manager is responsible for writing all configuration files. When a user updates the experimental configuration, the configuration manager translates those updated settings into the appropriate GPIB format corresponding to each GPIB hardware device. The updated settings are written to the appropriate hardware device's configuration file.

The run configuration GUI allows the user to specify which FEB and RTF settings files are used for a given hardware configuration. It also allows the user to specify each detector's phonon and ionization trigger modes. Once the user has selected the run settings, the user can issue the command to generate new configuration files. The specified run settings are then sent to the configuration manager to create the new configuration files.

Monitoring

It is important to monitor the experiment to be sure that the detectors and veto panels behave properly as WIMP-search and calibration data are obtained. Monitoring information is used to stop a data run with an obvious problem. Monitoring information is also used to remove suspect data using the bad events cut described in Section 4.3.1.

Several servers and GUIs are dedicated to monitoring the experiment. The servers keep track of the hardware, electronic signal noise, fridge, and data quality and provide monitoring information on them to users via GUIs and

the monitoring channel. The four monitoring servers are the monitoring data server, the fast analysis server, the fridge monitor, and the fridge data server. Data quality is monitored by viewing plots produced by the event display which is described later in the section.

The monitoring data server keeps track of the detector trigger rates, the detector offsets, the detector thresholds, the fridge variables, the veto rates, and the CAEN voltages and currents. The information tracked by the monitoring data server is available for viewing by a user from the veto GUI, scaler GUI, and fridge monitoring GUI.

The fast analysis server is used to process the first data file in any given run and calculates the power spectral density of each phonon and ionization channel. The first datafile consists of 500 events which are random triggers. A Fourier transform is performed on the digitized trace of the phonon and ionization channels for these random trigged events. The results, which are made available to all users via the noise viewer GUI, allow each channel's noise level for a given dataset to be compared with prior datasets. Figure 3.31 shows a picture of the noise viewer GUI. The fast analysis server and noise viewer GUI together provide feedback on the quality of the data being taken within a few minutes of a data run starting.

The fridge monitoring software, existing in three parts, is used to allow users, onsite or offsite, to monitor the state of the cryogenic system. The three parts of the fridge monitoring software are: the fridge monitor server, the fridge

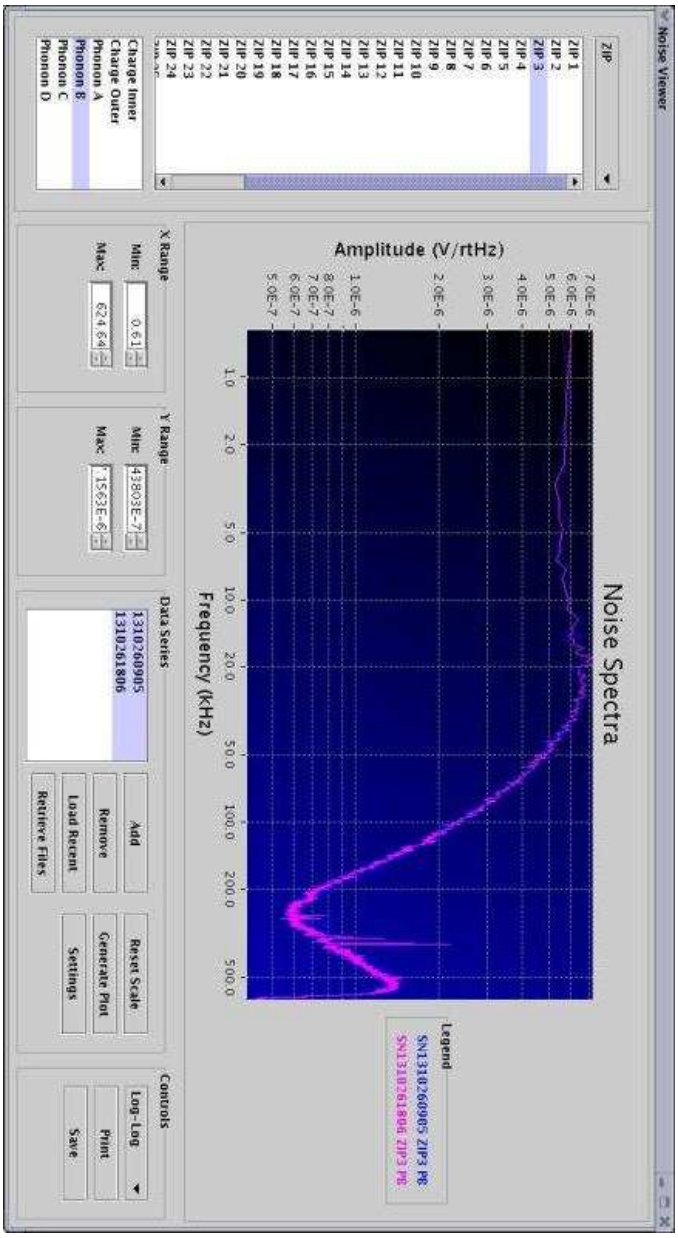

Figure 3.31: A picture of the noise viewer GUI. The noise viewer GUI allows the user to compare each detector channel's noise level with that of prior data runs.

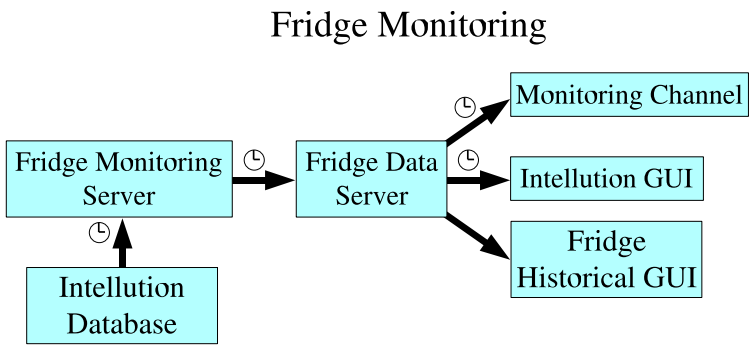

Figure 3.32: A schematic view of the fridge monitoring software. The fridge monitoring server operates on the `Intellution` computer and obtains the fridge data from the Intellution database. The fridge data server runs on `datasrv` and retrieves the fridge data from the fridge monitoring server and saves it to file. The fridge data server serves out the fridge data to the monitoring channel, the intellution GUI, and the fridge historical GUI. The clock symbol above an arrow indicates that program to the right of the arrow retrieves data from the program to the left of the arrow every minute.

data server, and the 2 GUIs used to view the fridge data: the "intellution GUI" and the "historical GUI". Figure 3.32 illustrates the flow of fridge monitoring information between the servers and GUIs.

The fridge monitoring software was designed with the goal of reliably providing a highly useful method for many simultaneous users to obtain information on the performance of the fridge without purchasing additional Intellution licenses. To accomplish this goal, the fridge monitoring software had to satisfy two criteria.

- The fridge monitoring information should be available from a stable com-

puter. The fridge data was made available for remote users from a robust DAQ computer rather than the unreliable Intellution computer that controls the fridge. The `Intellution` computer is the computer on which the propriety Intellution software is used to control the fridge. The fridge data server, which serves out the data to users, runs on the DAQ computer, `datasrv`. The amount of fridge data to be retrieved has changed several times. Currently 153 fridge variables are retrieved from Intellution.

- Fridge experts should be able to view the fridge data in the two ways they find most useful. It is valuable to be able to study how a few fridge variables change over time, and it is valuable to be able to view a snapshot of the current fridge variables. Two GUIs were designed to cover both uses. The "historical GUI" was designed to allow experts to see trends in the fridge data. The "intellution GUI" was designed to allow users to easily view the most recent fridge data.

The fridge monitoring server obtains the fridge data from the fridge Intellution database. This server must run on the Intellution computer in order to have access to the fridge via the Intellution software. The fridge monitoring server obtains the value of all fridge variables every minute from the Intellution software's database by means of an executable file provided by Intellution. The fridge monitoring server bundles these fridge values together and makes them available upon request for the fridge data server.

The fridge data server, which runs on the DAQ computer, `datasrv`, retrieves the fridge values from the fridge monitoring server every minute and stores them to file. Upon request, the fridge data server provides the requested fridge data for the time window and variables requested. Two GUIs, the "intellution GUI" and the "historical GUI" provide data to users.

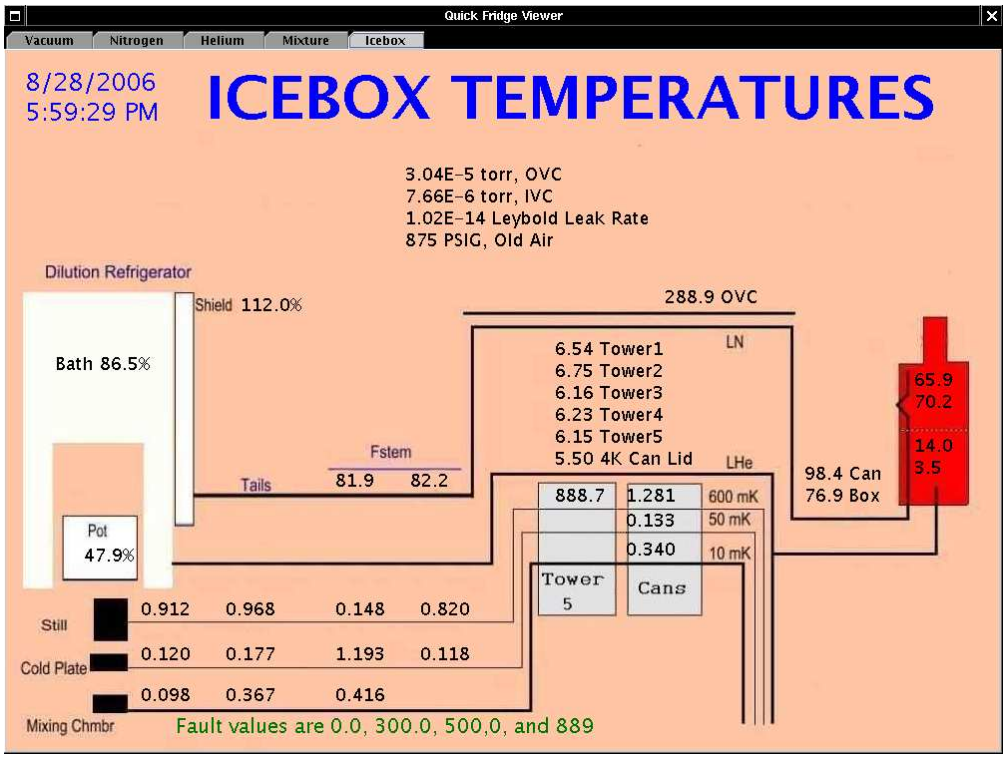

Figure 3.33: A picture of the fridge intellution GUI. The intellution GUI shows a snapshot of the current fridge values. To make it easier for the user to interpret the fridge data, the snapshot display is patterned after the Intellution display.

The intellution GUI allows a user to see a snapshot of the most recent

fridge values by retrieving them from the fridge data server and displaying them. Figure 3.33 shows a picture of the intellution GUI. The intellution GUI automatically updates its display every minute.

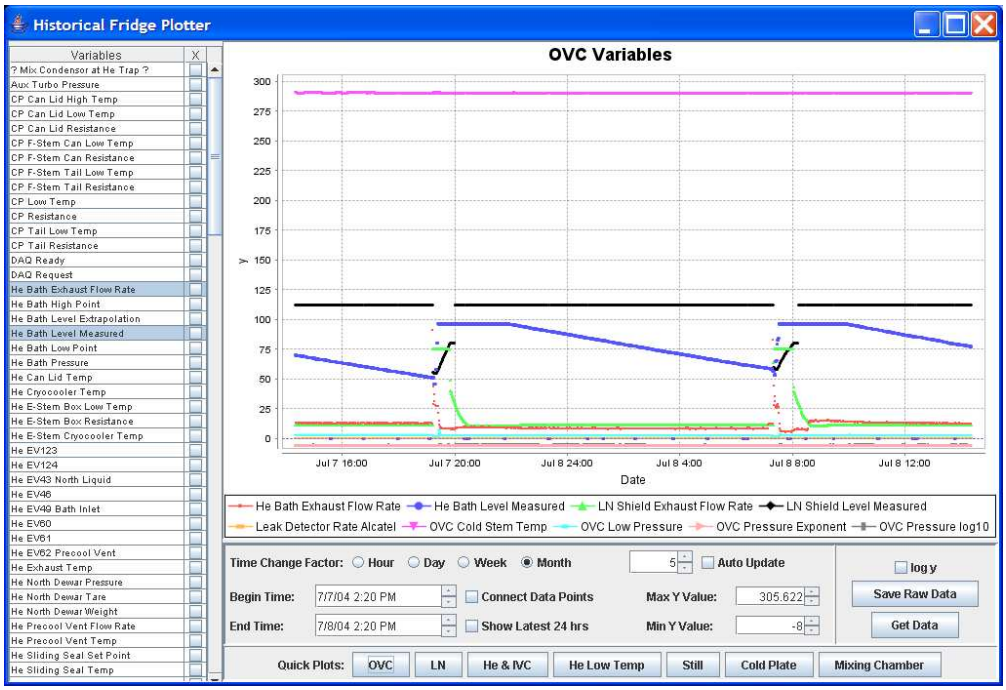

Figure 3.34: A picture of the fridge historical GUI. The historical GUI shows a record of the selected fridge variables for the requested time window.

The fridge historical GUI allows a user to view the value of a user-specified set of fridge variables over a user specified time window. Figure 3.34 shows a picture of the historical GUI. The fridge historical GUI, which I wrote, allows any member of the experiment to view trends in the fridge data.

The event display shows standard plots of the data that are useful for mon-

itoring the quality of the detector data. The event display updates the data quality plots in real time and posts a copy of the data quality plots on `datasrv` for offsite viewing. Figure 3.35 and Fig. 3.36 show portions of the event display. The plots show information on the distribution of events in a data series including detector event rates, histograms of the energy distribution in each ionization and phonon sensor, yield plots, and position plots.

The event display also gives a fuller picture of the most recent event. The picture includes a record of which detectors and veto panels were hit as well as showing the traces of all detector phonon and ionization channels. During a WIMP-search data run, all events are shown in the plots. Only a fraction of the events in a calibration data run are shown because of the high data rate. The event display, which is written in C++, runs on `builder`.

Fast Event Data Collection

The event builder retrieves, formats, and saves event data from the detector electronics and veto electronics. The event builder is a modified version of the Run II DATA Manager (R2DM) which is a general purpose DAQ developed at Fermilab for the CDF experiment [117]. It allows for distributed, parallel data readout and processing. Because the event builder was written in C++, the JAVA-based slow control servers communicate with it via the CORBA channels described at the end of this section.

The event builder is composed of a core and two nodes. The core of the

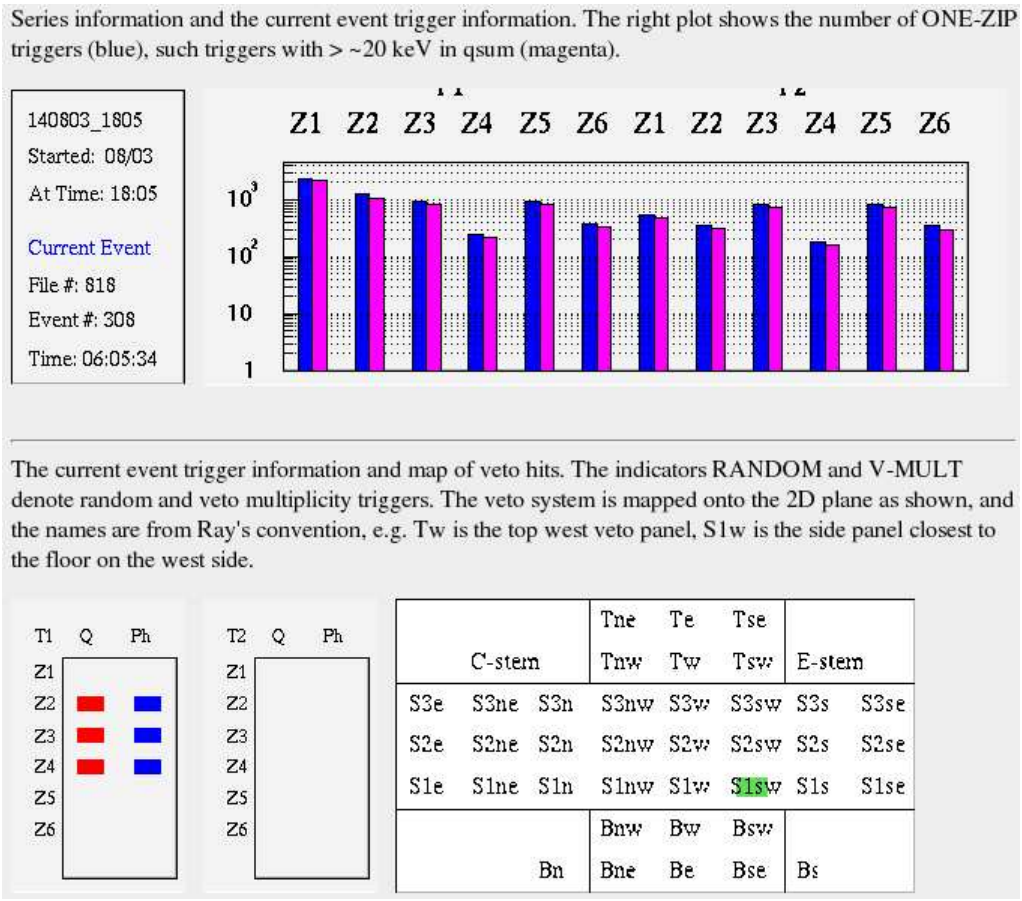

Figure 3.35: A picture of the first portion of the event display. The event display contains far more information than can be displayed in a single figure.

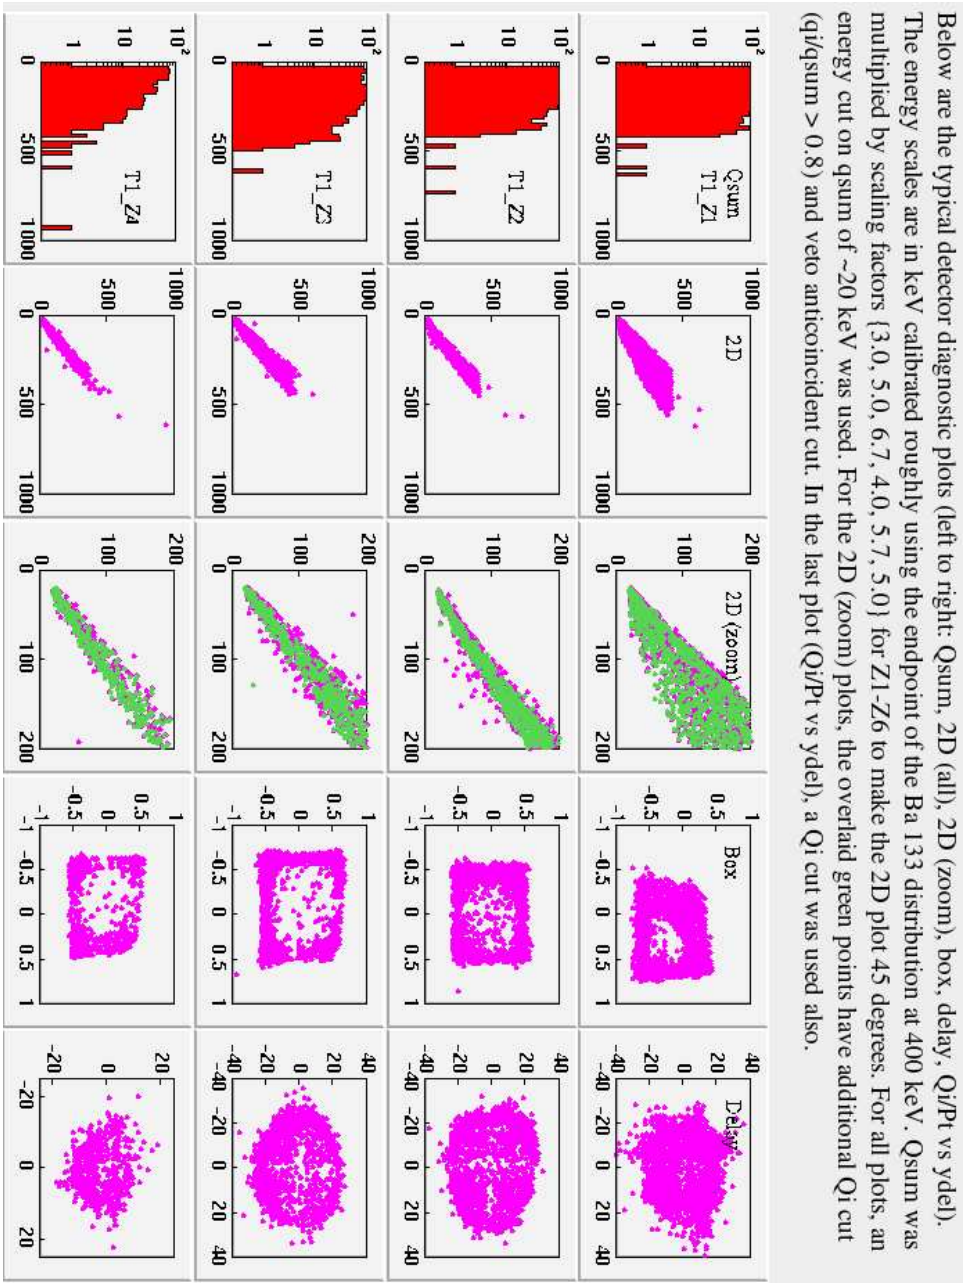

Figure 3.36: A second picture of a portion of the event display. The event display contains far more information than can be displayed in a single figure.

event builder runs on `builder` and is responsible for communicating with the slow control servers, packaging and saving the event data, and coordinating data readout with preparing the hardware for resumed data running. The two event builder nodes, which run on `tower1` and `vetocrate` are responsible for passing data from the digitizers and history buffer to the event builder core.

The event builder is able to quickly obtain event data to minimize the livetime lost due to event readout. The event builder nodes require ~20 ms to retrieve an event's data from the hardware, and the event builder core requires ~5 ms to concatenate the two data fractions from the nodes and to write a full event to disk. Because of the rapid hardware readout and parallel structure, the event builder is able to take data at a rate of ~30 Hz.

The event data is passed from the event builder to the data processing computers located on the surface by means of a Perl script. These computers are responsible for preparing the data for analysis.

Inter-server Communication

Because the DAQ software is modular, split between six computers, and written in multiple languages, it is critical for the various pieces of software to be able to communicate effectively and securely. The DAQ was designed to have two methods of inter-communication: Remote Method Invocation (RMI) and Common Object Request Broker Architecture (CORBA).

RMI, a JAVA specific protocol that allows one JAVA-based program (server

or GUI) to call a pre-approved method on a remote JAVA-based server, is used for two purposes: 1) it is used by all GUIs to communicate with the servers, and 2) it is used by JAVA-based servers to issue commands on other JAVA-based servers. For example, a command from a user operating the FEB GUI to bias a detector would be issued by the FEB GUI to the run control server via RMI; the run control server would then issue the same command to the GPIB server via RMI.

RMI commands from GUIs are typically issued to the run control server, because the computer on which the run control server operates (`control`) has access to the public network and therefore is guaranteed to be reachable by the GUI. The run control server then relays the command to the appropriate server. For example, commands issued by a GUI for a GPIB device are routed through the run control server to the GPIB server which runs on `monitor`.

Only authorized users are allowed to control the experiment. Security is controlled by files that specify allowed user names and IPs that can control the experiment. Typically only the daq account from the DAQ computers and a single computer on the surface at Soudan are allowed to control the experiment. At times, remote experts have been granted temporary control.

CORBA is used to create channels of communication that programs written in different languages can utilize. CORBA allows the event `builder`, which was written in C++, to communicate with the JAVA-based slow control of the experiment. CORBA is not easily used by remote users because it needs

significant installation effort, so the GUIs are restricted from using CORBA.

CORBA was used to create three channels of communication. The three CORBA channels used are: an error channel, a monitoring channel, and an interlock channel. All servers use the error channel and the interlock channel. The monitoring data server, the fridge data server, and the CAEN server use the monitoring channel.

- The error channel is used to convey errors ranging in severity level to the servers and the user. A data run is aborted if a sufficiently dangerous error occurs.

- The interlock channel is used to convey the run state so that servers and the event builder are apprised of updated run conditions.

- The monitoring channel is used to convey monitoring information. Information passed on the monitoring channel include detector trigger rates, trigger thresholds, phonon offsets, fridge data, veto rates, and the CAEN voltages and currents. If any of these monitored quantities deviates significantly from normal, automated corrections are attempted.

Chapter 4

Defining a WIMP Search

This chapter describes an analysis of the data obtained with the apparatus described in Chapter 3 to search for WIMPs. The goal of the analysis is to detect WIMP interactions in the detectors or, in the absence of a WIMP interactions, to set a limit on the WIMP mass and interaction cross section with nucleons.

4.1 The Data

During the data run, three different types of data were taken: WIMP-search data, gamma calibration data, and nuclear recoil calibration data. The calibration data was used to define cuts, calculate cut efficiencies, and estimate backgrounds. The data run was composed of many datasets. A dataset is composed of the data taken during a period of time after the detectors have been prepared to be sensitive to WIMP signals. Typical datasets last from two

hours to half a day and end when the detectors are turned off. The labeling of a datasets gives the date and time when the dataset began: 1YMMDD_HHTT where Y is the last digit of the year, MM is the month, DD is the day, HH is the hour, and TT is the minute.

The WIMP-search data consists of 1.42 million events split between 257 WIMP-search datasets. A typical dataset lasted for about seven hours. Most candidate events in the WIMP-search data are due to the ambient gamma background, to instances where the detector noise was above the threshold, or in coincidence with multiple veto panels; only a tiny portion might be due to WIMP interactions.

Gamma calibration data was taken using a ^{133}Ba source to illuminate the detectors. Typical gammas from calibration data were due to the source rather than from the ambient gamma background. About half of the eight million gamma calibration events (3.85 million) were used to define data cuts and were termed "open". The remaining gamma calibration events (4.17 million) were used to estimate the efficiency of the cuts for passing signal events and the inefficiency for rejecting background, or leakage. These events were termed "closed". There are 77 gamma calibration datasets that were taken throughout the data run. Gamma calibration datasets were typically split between open and closed by placing alternating groups of 50k events (for a duration of about 2 hours) in the open and closed data. (First 50k events to open and the second 50k to closed, etc.)

Nuclear recoil calibration data was taken using a ^{252}Cf source of neutrons and gammas. About 210 thousand events were obtained. Most of the events in the nuclear recoil calibration data are electron recoils due to gammas from the ^{252}Cf source. A fraction of the events are due to nuclear recoils caused by neutrons and are used to define the region or band in which nuclear recoils tend to fall. There were six nuclear recoil calibration datasets taken on four different days in the data run.

4.2 Analysis Overview

This analysis is composed of five parts.

1. The analysis begins with determination of the exposure. Exposure is defined to be the livetime during which the detectors were prepared to observe a WIMP induced nuclear recoil times the detectors' fiducial mass. The time during which the detectors were sensitive to WIMPs is called the livetime. There are portions of time during which some or all detectors were not behaving properly. These times are removed offline, resulting in a reduced livetime.

   Portions of each detector are not reliably sensitive to interactions and are excluded as described in Section 4.8. The most reliable portion of each detector is called the fiducial volume.

2. In order to facilitate a fair analysis, the analysis is blinded to a potential

signal in three ways.

(a) Data cuts were defined independent of potential signal events by using side bands and open calibration data. The background rejection cuts are defined before potential signal events are examined.

(b) Two versions of the WIMP-search data were created: a masked version and a complete version. The masked version has all potential WIMP-search events masked or removed to prevent a data analyzer from accidental awareness to potential signal events. Cuts where the definition utilized WIMP-search sidebands were defined using masked WIMP-search data.

(c) The analyzers agreed not to look at any potential signal events until all the cuts were defined and the background estimate had also been finalized. In the event that the masked version was imperfect, the integrity of the blinding would survive.

Unfortunately, the integrity of the masked version was compromised. The algorithm implementing the masking was flawed, and a few potential signal events were present in the masked version of the data. Figure 4.1 shows the intended and actual masking for a sample detector. However, the analyzers had mutually agreed to not look at the signal region. These potential signal events did not influence how the development of the cuts or the background estimate. The mutual, intentional blindness of the

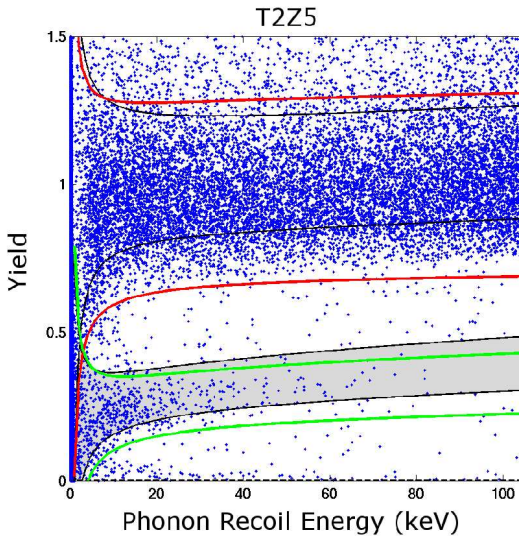

Figure 4.1: The intended and actual masked regions of T2Z5. The actual masked region is the shaded region. The intended region is the 3σ nuclear recoil band where the 2σ nuclear recoil band is shown in green. Blue datapoints are events from gamma and neutron calibration data. Modified figure from R. W. Ogburn.

analyzers preserved the blindness of the analysis.

3. The efficiency for signal events to fall within the fiducial volume and to survive each background-rejection cut is calculated as a function of nuclear recoil energy. Each cut's efficiency is measured using the appropriate dataset. Some efficiency calculations were cross checked by using multiple methods for estimating the efficiency or by using both calibration data and WIMP-search sidebands.

4. The expected backgrounds from surface events and neutrons are esti-

mated as well as potential systematic biases in the analysis.

5. The result is obtained by unblinding potential signal events. If no signal is observed, the result is scrutinized and a limit on the cross section of incident WIMPs is set under the assumption of a standard dark matter halo. If a potential signal is observed, the result must be carefully scrutinized before a possibly exciting claim can be made.

4.3 Livetime Estimation

To determine the exposure, one begins by calculating the livetime. If all detectors worked properly during all the times that they were prepared to listen for WIMPs, then the exposure would be given by

$$E_i = N_i \times M_i \times \tau_L \tag{4.3.1}$$

where E_1 (E_2) is the germanium (silicon) exposure, N_1 (N_2) is the total number of germanium (silicon) detectors, and M_1 (M_2) is the fiducial mass of a germanium (silicon) detector. Before removing periods of bad detector response, the experimental livetime, τ_L, is 77.4 days. This livetime is the sum of the livetimes for each event as calculated from data recorded by the event builder. A livetime,computed for each event, lasts from the time that the trigger became enabled after processing the previous event to the time the event in question triggered.

In practice, all detectors did not work properly at all times; as a result some data was of low quality and, therefore, excluded from the analysis. Data quality was compromised in three ways. First, some detectors consistently provided low quality data. All data from these detectors were excluded from the analysis. Second, there are periods of time when the entire experiment did not function properly. These malfunctions affected the data taken by all detectors, and data from these time periods are removed offline using the "bad events" cut. Malfunctions affecting the entire experiment decreased the livetime of all detectors equally. Third, some malfunctions affected a single detector, and the resulting compromised data was also removed offline by criteria incorporated in the same bad events cut. Malfunctions affecting individual detectors resulted in each detector having a unique livetime. After applying the bad events cut, the exposure becomes

$$E_i = \sum_j (M_{ij} \times \tau_{Lij}) \tag{4.3.2}$$

where j is an index that steps over the good germanium or silicon detectors, and τ_{Lij} is the jth germanium or silicon detector's livetime. Rejecting bad events and detectors was done to insure that potential signal events have data of excellent quality.

Even when data from a detector was deemed to be unsuitable for WIMP detection, that data was still useful for rejecting events hitting multiple detectors.

4.3.1 Livetime and the Bad Events Cut

The bad events cut removes portions of the data from all detectors or a single detectors based on the following reasons. Events are removed from all detectors for first two reasons and from a single detector for the remaining reasons.

1. Events occurring during failures such as a power outage, cryogen transfer, or software error configuring the detectors are rejected for all detectors. The computer on which the event builder runs is on a UPS and would continue to take flawed data during a power outage.

2. Datasets with few events are rejected, because the low numbers of events make the Kolmogorov-Smirnov (KS) tests unreliable. The purpose of a KS test is described in the third item of this list.

3. If the data from one detector's dataset failed one of several KS tests, all data from the failing detector's dataset is rejected. A KS test is used to determine if the data for a given detector and dataset is statistically consistent with a template dataset. The template dataset is one of the earlier datasets that was hand selected for its high quality data.

4. Several detectors had poor signal to noise for a few datasets. These datasets are rejected for the given detectors.

5. Some gamma calibration datasets are rejected for all detectors since they

had exceptional rates of low yield events.

6. A few events had an unexpected data processing error code and are rejected for all detectors.

We will now discuss each of the seven reasons for event rejection.

ExperimentWide Failures

A number of experiment-wide failures occurred. Data taken during a failure is removed by the bad events cut. Some of the failures that occurred were a power dip due to a thunder storm, a pump trip, a cryogen transfer, and a GPIB conflict preparing the detectors to be run. The power dip occurred during the 140611_1708 dataset. Several electronics crates lost power as did some of the dilution refrigerator pumps. While the event rate was unaffected, the relevant portion of the dataset is rejected. The squids lost lock due to a pump trip during the 140407_0913 dataset which may have triggered a spike in the event rate that was observed near the end of the run. The relevant portion of this dataset is rejected. A cryogen transfer occurred after the first 4500 events in the 140508_1035 dataset. In addition it the difficulty in trusting data taken while cryogens are being dumped into the fridge, the liquid nitrogen level meter, which is operational only during a transfer, is known to induce 19 kHz noise on the detector channels, substantially increasing the trigger rate. Only the first 4500 events in this dataset are accepted. The 140729_0903 dataset has two problems. First, the record of which detector triggered was corrupted. Second,

the on-shift operator noted that a conflict had occurred while configuring the detectors. This entire dataset is rejected.

Low Stat Datasets

For the Kolmogorov-Smirnov (KS) tests to be reliable, a sufficient number of events are needed in each detector. We wanted a typical detector to have at least 10 to 20 non-noise events in a particular dataset. For 12 detectors, a minimum of a couple hundred non-noise events are needed. The first 500 events in a dataset are due to random triggers from the event builder and are not useful for the KS tests. About half the subsequent events are due to above threshold noise in a detector and are also not useful for the KS tests. Datasets with less than 1000 events were rejected for all detectors.

Kolmogorov-Smirnov Tests

It is important to distinguish a dataset with bad data from one with good data. Good datasets should have similar distributions of events. Kolmogorov-Smirnov (KS) tests are sensitive to the shape and position of the distribution of events and can be used to see if events from two datasets are consistent with being from the same probability distribution. One good dataset was chosen as a template. All other datasets were compared with the template dataset. One dimensional and two dimensional KS tests are used to test the quality of the signals from the four phonon and two ionization channels, the phonon

delays and energy partitions, yields, and energy distribution. The KS test returns a likelihood that a particular dataset is consistent with the template dataset. Figure 4.2 shows a couple two dimensional KS tests for a couple sample detectors in a dataset.

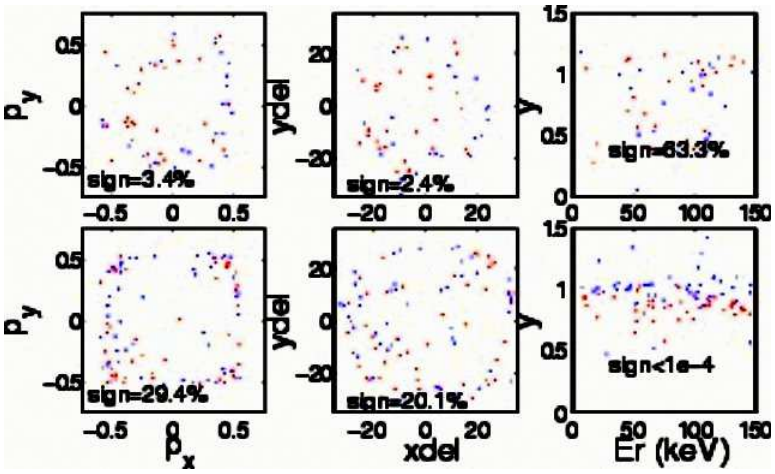

Figure 4.2: Sample two dimensional KS tests from a dataset for two detectors. Blue datapoint are from the template dataset. Each row shows KS tests for a single detector. The two left-hand plots show a comparison of the dataset's partition plots, the center plots show a comparison of the delay plots, and the right plots compare the yield plots. The likelihood that events in the tested dataset are from the same probability distribution is show on each plot as "sign". The lower (upper) detector fails (passes) the KS yield-recoil energy test.

Datasets can be divided into three classes: 1) datasets for which all detectors were consistent with the template dataset, 2) datasets for which many detectors have channels that are inconsistent with the template dataset, and 3) datasets with a few inconsistent detector channels. The first class of datasets

are retained and the second are rejected. The rest of this section will deal with the third class. Typically the data for a failing detector or two are excluded from the analysis, while the data for the rest of the detectors are retained. Note that events in rejected detectors are still used to veto events in multiple detectors. Vuk Mandic's thesis [105] contains additional information on the KS tests.

A WIMP-search dataset, 140505_1739, with a summed ionization energy KS test below 0.1 and a two-dimensional yield-recoil energy KS test also below 0.1 were rejected. Three gamma calibration datasets, 140513_0919 - D; 140525_1338 - B, D; and 140608_0932 - A, D, had T1Z6 phonon channel(s) that ceased functioning properly during the data run. Data from the affected detectors and datasets were rejected.

Some datasets did not have a sufficient number of events in a particular detector causing the KS tests to fail. Since the affected detectors and datasets can not be tested, they are rejected. Table 4.1 lists the affected datasets and detectors.

Poor Signal to Noise

Some calibration datasets had a fairly large pulse in a detector or two in the random triggered events used to define the baseline phonon noise leading to reduced signal to noise. The datasets were reprocessed without the offending pulse. After reprocessing the signal-to-noise of most datasets was restored. The

Dataset	Detector(s)
140712_2157	T2Z5
140713_2155	T2Z2
140715_1901	T1Z1, T1Z2
140720_1406	T1Z1, T2Z5
140720_1518	T2Z5
140720_2022	T1Z1
140726_1821	T1Z4
140728_1808	T1Z1
140712_1015	T1Z4

Table 4.1: List of datasets and detectors with insufficient statistics causing the KS tests to fail.

few remaining datasets are small and were rejected for the affected detectors. The rejected datasets and detectors are listed in table 4.2.

Poor Calibration Data

The high event rate of calibration causes the detectors to lose neutralization more quickly while taking calibration data. Partially neutralized detectors can have a higher rate of low yield events, because some bulk events have reduced ionization. The timing of the phonon pulses for the additional low yield events are similar to other bulk events are not suitable for defining a surface event rejection cut. Four gamma calibration datasets had unusually high rates of low yield events. The 140414_0757, 140615_1252, 140726_2123, and 140728_2226 datasets were rejected. The shorter neutron calibration datasets are not expected to suffer from bulk, low yield, electromagnetic events.

Dataset	Detector(s)
140712_1015	T1Z4
140712_2157	T2Z5
140713_2155	T2Z2
140715_1901	T1Z1, T1Z2
140720_1406	T1Z1, T2Z5
140720_1518	T2Z5
140720_2022	T1Z1
140726_1821	T1Z4
140728_1808	T1Z1

Table 4.2: List of datasets and detectors removed because of increased noise on phonon channels leading to reduced signal-to-noise.

Missing Event Records and Unusual Error Codes

Some events have problems with their event record. The processing of the event to create the event record was unable to create all parts of the event record for a few events. A few events are missing a record of a global trigger and rejected. One event is missing its record of veto hits; the 20 events surrounding this event are also rejected. The processing also tagged the event record for 11 random triggered events with an unusual error code. These events are rejected.

Excluded Detectors

Three detectors, a germanium detector, T1Z1, and two silicon detectors, T1Z6 and T2Z1, were not included in the analysis. T1Z1 was the first detector fabricated and had a high critical temperature variation across the face of the crystal leading to uneven detector response. In prior data runs, the uneven re-

Ge Detector	Livetime (days)
T1Z2	74.56
T1Z3	74.57
T1Z5	74.62
T2Z3	74.56
T2Z5	74.58
μ_{Ge}	74.58
Si Detector	**Livetime (days)**
T1Z4	74.51
T2Z2	74.56
T2Z4	74.60
T2Z6	74.61
μ_{Si}	74.57

Table 4.3: List of the livetime surviving the bad events cut for each detector.

sponse was observed to substantially increase the number of background events. Events in T2Z1 were observed to have poor phonon signal-to-noise in two of the four phonon quadrants in both earlier test data and in *in situ* calibration data. T1Z6 was excluded because it was known to be contaminated with ^{14}C from a source before the data run began. ^{14}C emits electrons that don't penetrate the surface layer of the crystal, increasing the number of surface events. Data from these three detectors were used to reject events occurring in multiple detectors. The data from the remaining five germanium and four silicon detectors were used in this analysis.

Resulting Livetime and WIMP Sensitivity

The removal of events and datasets decreases the livetime of each detector. The surviving livetimes for the 9 detectors included in the analysis are listed in table 4.3. The mean livetime for the germanium and silicon detectors is also given. The bad events cut reduces the livetime of a typical detector by about 3 days. The livetime is only effected by the bad events cut. All other cuts decrease the efficiency of the analysis rather than the livetime.

As was shown in Section 2.2.1, the WIMP recoil energy spectra falls exponentially with recoil energy. To calculated the expected WIMP sensitivity, the energy dependent efficiency of each subsequent cut must be estimated. The estimated efficiencies are used to form a combined efficiency for the germanium and silicon detectors. Equation 4.3.3 shows how in the absence of background the WIMP sensitivity is affected by the combined efficiency.

$$S_i \propto E_i \int \frac{dR_i}{dQ} \epsilon_i(Q) dQ \tag{4.3.3}$$

where S_1 (S_2) is the sensitivity of the combined germanium (silicon) detectors, Q is the recoil energy, dR_1/dQ (dR_2/dQ) is the differential rate in germanium (silicon), and ϵ_1 (ϵ_2) is the combined germanium (silicon) efficiency of this analysis.

4.4 Good Events Cut

The bad events cut rejected events coming from obviously flawed events, detectors, or periods of time. Potential WIMP candidate events must satisfy a higher standard. WIMP candidates are also required to not be abnormal candidates in four respects: the fit to the ionization pulse should not be bad [118], deviations in the prepulse baselines should not be overly large [119], no phonon channels should saturate [120], and phonon channels should not have reconstructed energies that were well below that of typical noise events [121]. The good events cut rejects events that are abnormal in any of these four respects. We now look at these four parts of the good events cut.

1. Ionization pulses are required to have the expected pulse shape. They are fit to a template. The goodness of the fit is given by the chi-squared of the fit. Events with large chi-squared values have badly fit ionization pulses. Higher energy events normally have larger chi-squared fits, so an energy dependent cut is defined to reject events with abnormally high chi-squared values.

2. The signal baseline is required to be stable. The signal traces for the $410\,\mu s$ before an event are saved and are used to determine if an event occurs on a stable baseline. This pre-event portion of the signal trace is referred to as the pre-pulse baseline. The $410\,\mu s$ pre-pulse baselines of the four phonon channels are used to give an indication of the environment

in which a particular event occurred. Abnormally large deviations from the normal standard deviation of the pre-pulse baseline indicate unusual detector behavior. From 5 to 10% of events during the data run are due to event builder generated triggers at random times during the run. These "noise triggers" are used to define normal phonon channel behavior. Events with pre-pulse baselines whose standard deviation is more than 5σ from the norm are rejected.

3. Events may not saturate a phonon channel. This analysis does not extend to recoil energies above 100 keV. All events saturating any of the phonon channels are rejected.

4. Events may not have negative energy. A few events exhibit significant negative energy for a single phonon channel. Apparent negative energy can result when another event occurs just before the event being analyzed. The earlier event causes undershoot in the phonon pulses. Events occurring on the undershoot can have poorly fit phonon pulses. Events with recorded phonon energy that is 6σ below the mean energy of noise events are removed. This cut was motivated by two non-candidate events in T1Z6 which had low (10.7keV and 114keV) recoil energies despite saturating at least one phonon channel for at least one digitizer bin, because other phonon channels had significant negative phonon energy. Figure 4.3 shows how the cut on negative phonon energy is defined for two sample phonon channel in T1Z5. This cut actually has little effect. While 2.8% of

open gamma calibration events have negative phonon energy in a phonon channel, all of them are rejected by other portions of the good events cut: the ionization threshold cut described in Section 4.6, or the inner electrode cut described in Section 4.8. Among WIMP-search events, 0.06% have negative phonon energies.

The efficiency of the good events cut is measured using masked WIMP-search data and is cross checked using gamma calibration data. (Usage of a sideband of WIMP-search data was not necessary, because the fraction of potential signal events in the masked data is far too small to affect the measurement.) Neutron calibration data is not used because of low statistics. The efficiency is measured by seeing what fraction of masked WIMP-search events pass the good events cut in each bin of recoil energy. The efficiency of the good events cut for all detectors is above 97% for recoil energies above analysis threshold of 7keV. The efficiency measured from gamma calibration data is similar but slightly lower. Because of the higher event rate in calibration data, more events have badly fit ionization pulses and bad standard deviations of the pre-pulse baseline. The efficiency of the cut was measured in bins from 5 to 7 keV, 7 to 10 keV, 10 to 20 keV, 20 to 30 keV, 30 to 50 keV, and 50 to 100 keV. The efficiency was then fit by the following functional form: efficiency $= C_1 + C_2/(E_r - C_3)^{C_4}$ [122]. Typical values for the fit coefficients are $C_1 = 1$, $C_2 =$ -20, $C_3 =$ -10, and $C_4 = 5$. Figure 4.4 shows the efficiency of the good events cut as a function of recoil energy for all detectors.

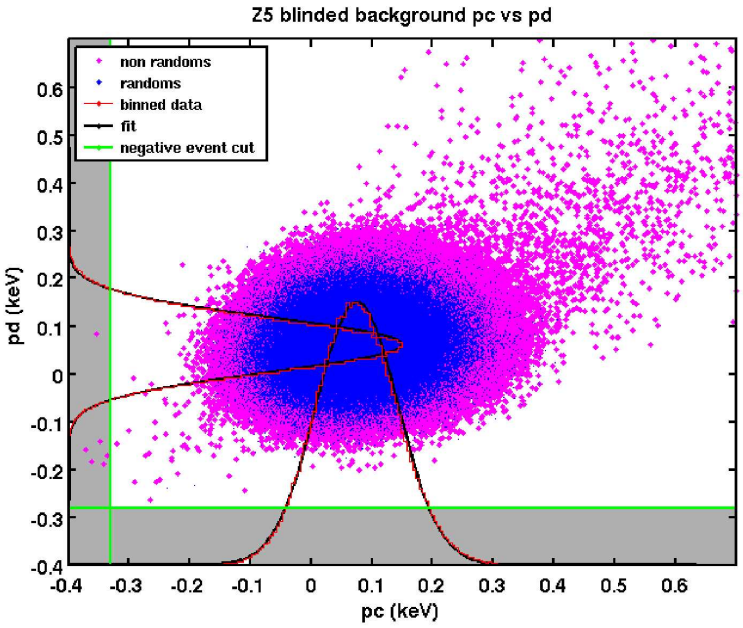

Figure 4.3: A 'comet' plot for two sample phonon channels in the sample detector T1Z5 showing how events with 6σ negative phonon energy in a phonon channel are rejected. The plot is made using WIMP-search data from which potential wimp candidates were excluded. The horizontal axis gives the energy in phonon channel C and the vertical axis gives the phonon energy in phonon channel D. Blue data points are noise events, magenta data points are events triggered by one of the 12 detectors, the red lines give the histograms of the noise triggers, and the black curves give the fits to the histograms. The light green horizontal and vertical lines indicate the -6σ cut. The gray shaded areas are excluded.

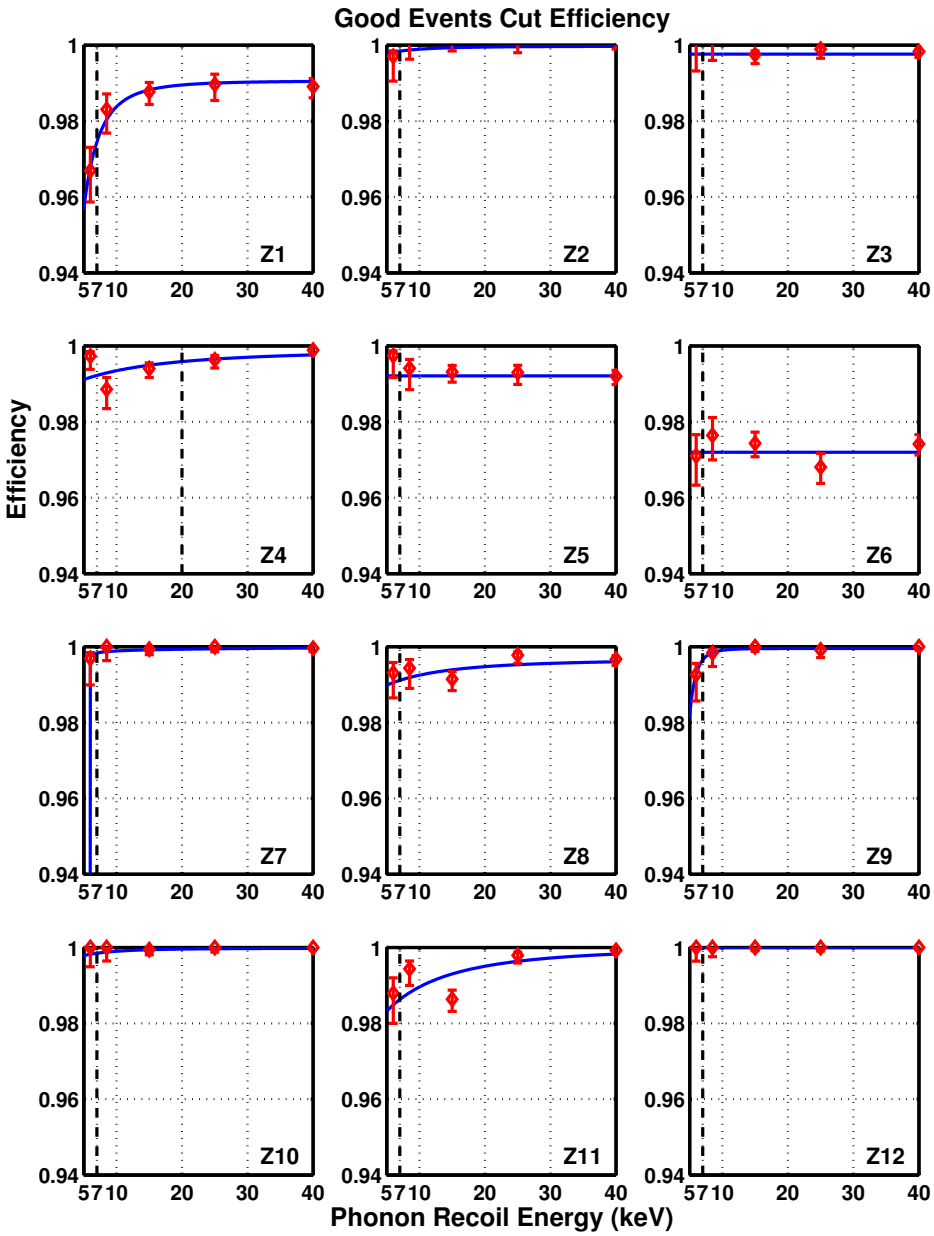

Figure 4.4: The efficiency of the good events selection cut as a function of recoil energy for all twelve detectors as estimated from gamma calibration data. The red data points show the measured good events selection efficiencies with 1σ error bars. The blue curve shows the fit efficiency, and the black dashed line indicates the analysis threshold. The efficiency remains flat above 40keV.

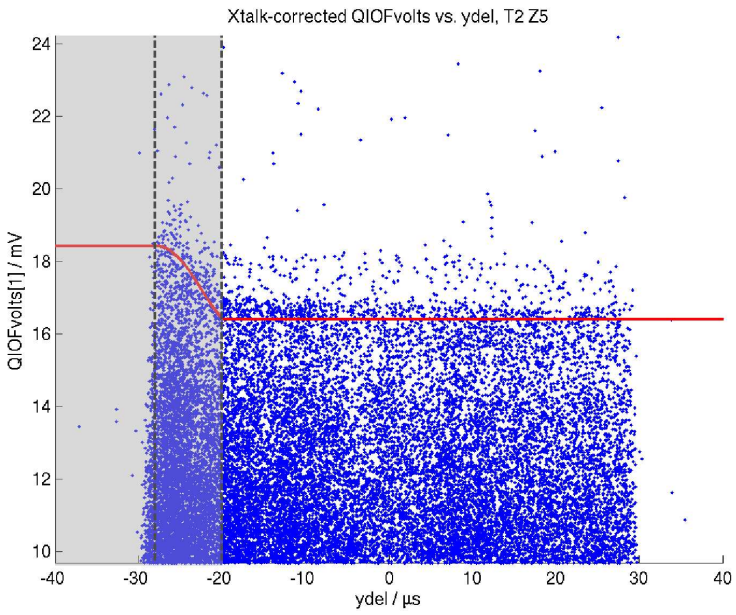

Figure 4.5: The discrete change in recorded inner electrode energy as a function of y-delay in T2Z5. The cliff cut is used to remove the events with y-delay less than -20μs. The shaded area is removed.

4.4.1 The Cliff Cut

A region of a single detector, T2Z5, exhibits unusual detector response [123], and events from the poorly functioning region are removed. T2Z5 suffers from a discrete change in ionization collection for events, with positions as inferred from phonon timing, with y-delay less than -20μs. Figure 4.5 shows the discrete change in raw recorded inner electrode energy as a function of y-delay. Events in T2Z5 with y-delay less than -20μs are abnormal and are rejected by the cliff cut. Figure 4.6 shows the efficiency of this cut to be 0.753±0.026 from

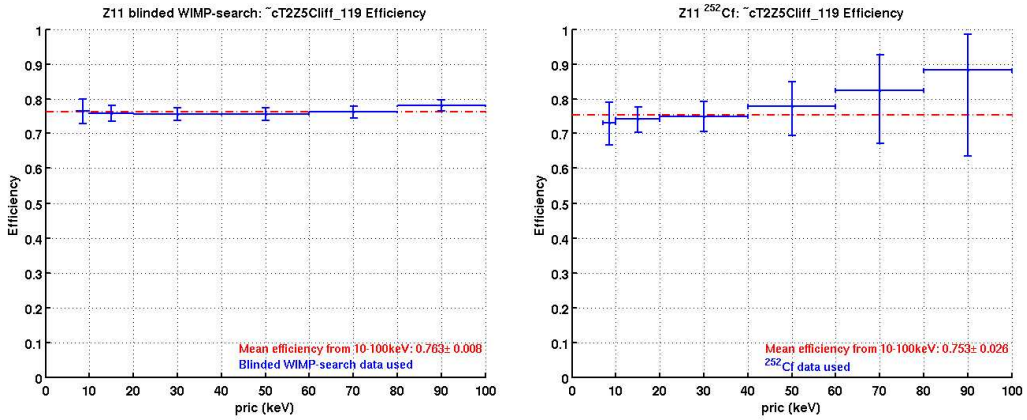

Figure 4.6: The efficiency of the cut designed to remove events from the region of the T2Z5 detector with abnormally high recorded ionization energy. The horizontal axis give the phonon recoil energy. Left: the efficiency calculated using WIMP-search data excluding potential candidate events. Right: the efficiency calculated from neutron calibration data. The estimates are in good agreement and consistent with a flat efficiency. The lower efficiency of 0.753 is used. The shown error bars given the 95% confidence limits.

neutron calibration data which is consistent with the efficiency of 0.763 ± 0.008 measured from masked WIMP-search data. Both efficiencies are flat in energy. The lower efficiency estimate from calibration data is used.

4.5 Search Regime and Trigger Efficiency

This section describes how the range of recoil energies used in this analysis were chosen as well as describing a measurement of the efficiency for an event to trigger the experiment.

A nuclear recoil search region of 7 keV to 100 keV was used for all detectors

except the silicon detector T1Z4. The upper threshold is chosen because the scatter rate above 100 keV for even massive WIMPs of 500 GeV is much less than the rate below 100 keV, while the expected background really does not fall rapidly. Figure 2.3 shows the differential scatter rate as a function of recoil energy.

The low energy threshold was chosen so that few electron recoil events would leak into the signal region because of misestimation of the ionization energy due to white noise on the ionization pulse [124]. Ionization pulses of various energies were simulated on a background of white noise. The simulated pulses were then fit using the same chi-squared minimization as is used for data. Larger pulses are harder to misfit. The fraction of events for which the fit amplitude was less than half the true amplitude was calculated. A phonon recoil energy threshold was chosen so that less than 0.1 of these misfit events in the entire WIMP-search data would be above threshold.

The low energy threshold was raised for T1Z4 to 20 keV since open gamma calibration data indicated worse rejection power for surface events at low recoil energy in this detector than in the other silicon detectors.

Some events above the 7 keV phonon recoil energy threshold fail to trigger the experiment, resulting in lost efficiency. The trigger efficiency is the fraction of events with a given recoil energy to which the experiment is sensitive. For an unbiased sample, the efficiency is calculated for a particular detector using events that triggered any other detector. The trigger efficiency for a given recoil

energy is the fraction of events that also triggered on this particular detector. The trigger efficiency is calculated using masked WIMP-search data and gamma calibration data [125]. The efficiencies measured using masked WIMP-search data tend to be slightly lower and are used as they are more conservative. Figure 4.7 shows the trigger efficiency as a function of phonon recoil energy for nuclear recoils for all 12 detectors. The two germanium detectors in tower 2, T2Z3 and T2Z5, have the worst trigger efficiencies resulting in significant lost efficiency for low energy nuclear recoils.

The measurement is made slightly more complex because the experiment triggers on total phonon energy, while the offline analysis is based on the phonon recoil energy. The total phonon energy is the phonon recoil energy plus the ionization energy due to the Luke effect. Typical electron recoils have a total phonon energy that is twice the phonon recoil energy (a yield of 1), while typical nuclear recoils have total phonon energy that is 30% greater than the recoil energy (a yield of 0.3). The trigger efficiency of a 10 keV electron recoil will be different than the trigger efficiency of a 10 keV nuclear recoil. Since the trigger efficiency is defined using electron recoils, the nuclear recoil equivalent phonon recoil energy must be calculated. This is done by converting each electron recoil to an equivalent nuclear recoil by calculating what the recoil energy would have been for a hypothetical nuclear recoil with identical total phonon energy. The hypothetical nuclear recoil is assumed to have an recoil energy dependent yield given by the centroid of nuclear recoil band.

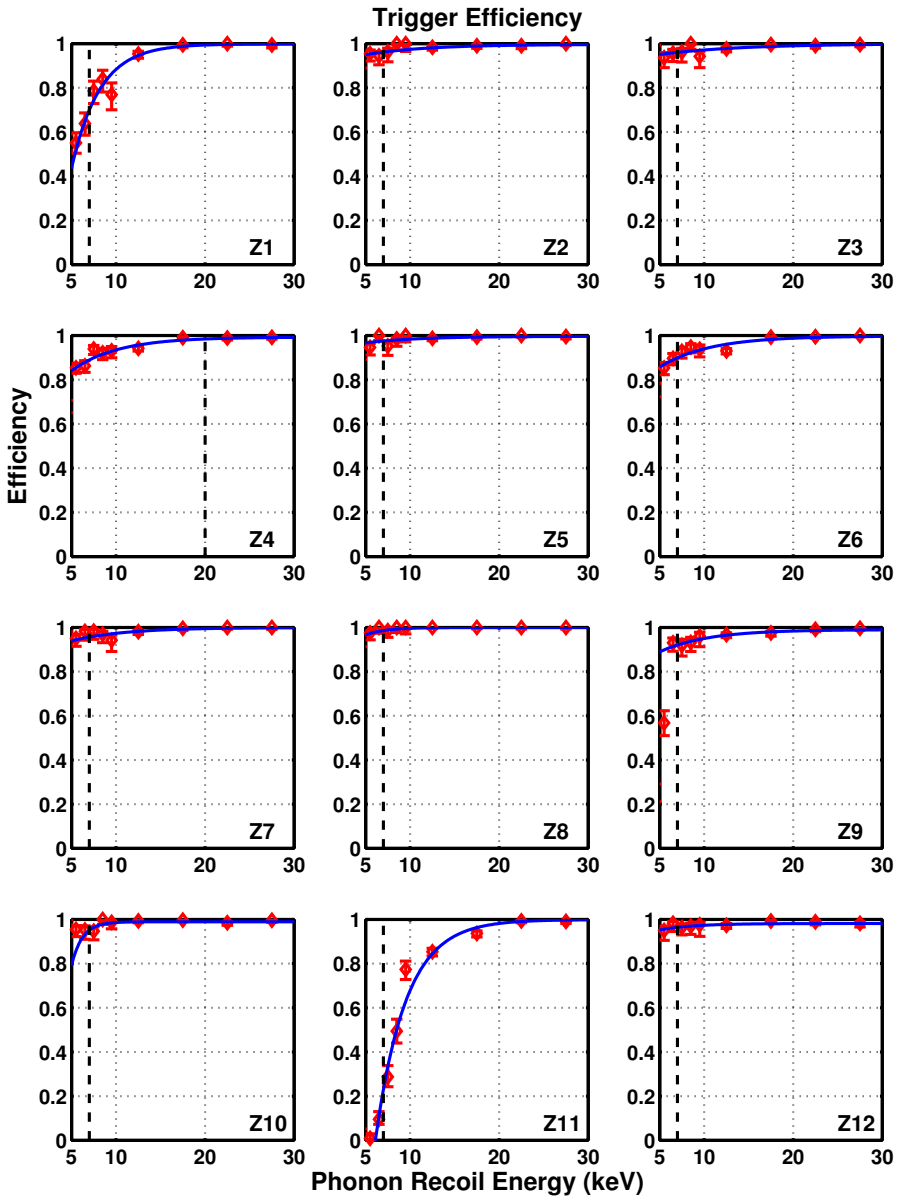

Figure 4.7: The trigger efficiencies for all twelve detectors as estimated from WIMP-search data. The red data points show the measured trigger efficiencies with 1σ error bars. The blue curve shows the fit efficiency, and the black dashed line indicates the analysis threshold. Trigger efficiencies for recoil energies below 6 keV were not used in the fitting. The efficiency remains flat above 30 keV.

The efficiencies were measured using 1 keV bins from 0 to 10 keV, using 5 keV bins to 30 keV, and a 70 keV bin extending to 100 keV. One would expect the trigger efficiency to be fit by an error function. Only the turn-off portion of the trigger efficiency lies above the analysis threshold, allowing the fit to be well approximated by the following functional form: efficiency $= C_1 + C_2/(E_r - C_3)^{C_4}$ [122]. Efficiencies for recoil energies below 6 keV were not used in the fit as including them had an adverse affect on the fit. Typical values for the fit coefficients are C_1 = -0.1, C_2 = 0.1, C_3 = 3, and C_4 = 1.

4.6 Ionization Threshold

Observing a potential dark matter signal with clarity by both phonon and ionization channels signal is important for obtaining a reliable signal, therefore it is desirable to reject events in which the signal in either the ionization or the phonon sensors is not significantly above the noise. A ionization threshold cut is set to guarantee that any signal event deposits significant energy in an ionization sensor in addition to being above the phonon analysis threshold [126]. A typical detector has about 80 good events in the masked R119 WIMP-search data, passing the fiducial volume cut described in Section 4.8, that have ionization energies consistent with noise. Since any potential signal event must have a clear ionization signal, these events must be excluded. An ionization threshold cut was set such that the expected number of these events surviving the cut was less than 0.01. Since these events have ionization signals consistent with

Detector	Number of Events	3.85σ Threshold (keV)
T1Z1	156	1.326
T1Z2	70	1.027
T1Z3	46	1.345
T1Z4	175	1.772
T1Z5	62	1.086
T1Z6	69	1.941
T2Z1	58	1.392
T2Z2	60	1.390
T2Z3	83	0.993
T2Z4	115	1.553
T2Z5	51	1.251
T2Z6	40	1.914

Table 4.4: List of number of events with phonon recoil energy between 5keV and 100keV passing a inner ionization electrode cut, ionization χ^2 cut, and phonon prepulse baseline consistent with noise cut. The second column contains the number of events passing these cuts for each detector and the third column contains the resulting ionization threshold cut.

noise, a cut of 3.85σ above the mean inner ionization electrode noise was set assuming a normal distribution. Table 4.4 gives the number of events without an ionization signal in each detector and the resulting ionization thresholds.

To estimate the efficiency of the ionization threshold cut, ionization pulses of various energies were simulated on a background of white noise and then fit using the same chi-squared minimization as is used for data. The energy of the fitted ionization pulse occasionally underestimates (but never overestimates) the energy of the true pulse by more than is expected from a normal distribution. The right hand portion of Fig. 4.8 illustrates how the efficiency of

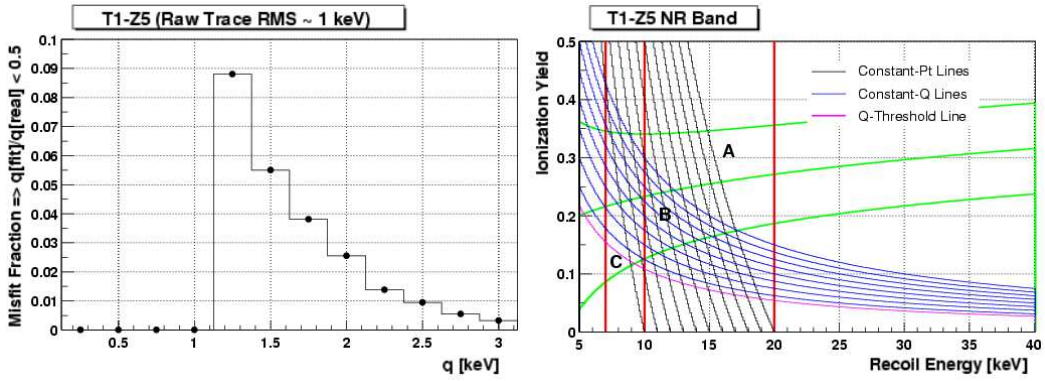

Figure 4.8: Sample plots showing the method by which the ionization threshold cut efficiencies were estimated. Left: the fraction of events with the inner ionization pulse misfit such that the reconstructed energy was less than half that of the true energy of the simulated pulse as a function of energy. Right: schematic showing the three regions relevant for calculating the efficiency of the ionization threshold cut. The outer green curves show the nuclear recoil band or the signal region. The inner green curve shows the centroid of the nuclear recoil band. The vertical red lines show the recoil energy bins in which the efficiency was estimated. The blue curves and (lowest) magenta curve show lines of constant ionization. The magenta curve gives the ionization threshold and the upper blue curve is for 3keV. The nearly vertical black curves are lines of constant total phonon energy. Events in region C are below the ionization threshold and contribute solely to the inefficiency. Events in region A, above the upper blue curve, have sufficient energy for misfitting of the ionization pulse to no longer contribute to the inefficiency. Events in region B will have varying degrees of efficiency. Simulation and plots by Long Duong.

the ionization threshold cut for nuclear recoils was estimated. The region for which ionization pulses are sufficiently large that misfitting of the ionization pulse is unimportant (above 3 keV) is labeled "A". The efficiency of the ionization threshold cut in this region is 100%. Region "B" is the region in which misestimation of the ionization energy could drop an event out of the nuclear recoil band signal region. When the ionization energy is underestimated, the total phonon energy remains constant. Events with misestimated ionization energy will follow the lines of total phonon energy. Potential signal events are assumed to have a Gaussian distribution about the centroid of the nuclear recoil band. The efficiency of region "B" is the fraction of potential signal events that would be misfit but still remain within the nuclear recoil band. Region "C" is the region in which ionization pulses are below threshold. All event in region "C" are rejected. The efficiency was estimated in bins from 7 to 10 keV, 10 to 20 keV, and 20 to 100 keV. The estimated efficiency was fit using a least squares fit to the following functional form: efficiency $= C_1 + C_2/(E_r - C_3)^{C_4}$ [122]. Typical values for the fit coefficients are $C_1 = 1$, $C_2 = -1$, $C_3 = 6$, and $C_4 = 2$. The efficiency resulting from the functional fit for all detectors is shown in Fig. 4.9.

4.7 Single Detector Requirement

Do to their very low cross section, WIMPs rarely scatter in multiple detectors, whereas electromagnetic and neutron backgrounds can do so. Events

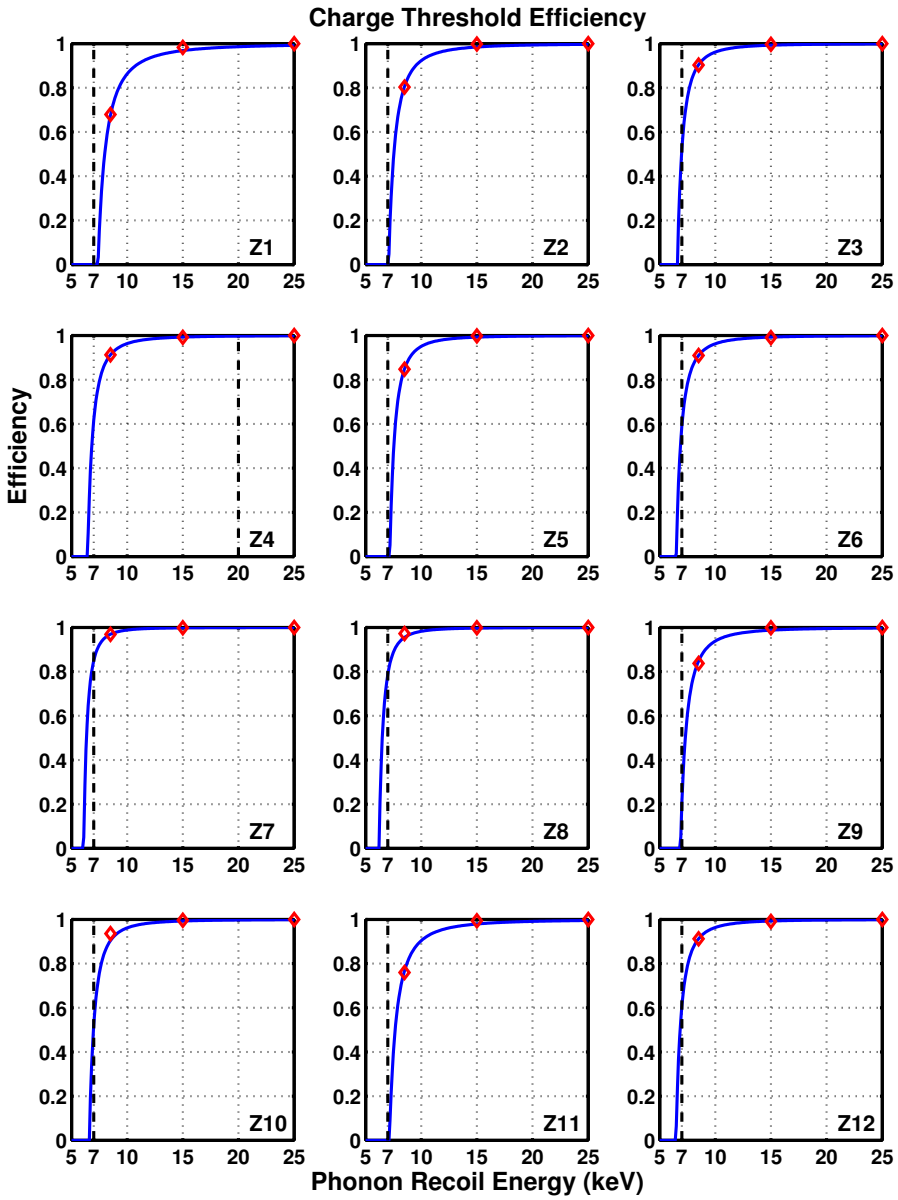

Figure 4.9: The ionization threshold cut efficiency trigger efficiencies for all twelve detectors. The efficiency remains flat above 20keV. The black dashed line indicates the analysis threshold. The red diamonds indicate the efficiency from testing for misfitted simulated ionization pulses in each energy bin as described in the text.

Tower1	T1Z2	T1Z2	T1Z3	T1Z4	T1Z5	T1Z6
Mean (keV)	0.8643	0.2269	0.2496	0.4831	0.2740	0.3888
Std (keV)	0.5880	0.1232	0.1253	0.3302	0.1355	0.2869
Efficiency (%)	99.965	99.976	99.978	99.965	99.978	99.977
Tower2	T2Z2	T2Z2	T2Z3	T2Z4	T2Z5	T2Z6
Mean (keV)	0.2728	0.1257	0.3588	0.0905	0.3137	0.0918
Std (keV)	0.1485	0.0670	0.1724	0.0566	0.1611	0.0656
Efficiency (%)	99.976	99.979	99.984	99.978	99.978	99.979

Table 4.5: List of detector phonon recoil energy means and standard deviations as well as the efficiency of the singles cut.

containing interactions in multiple detectors are rejected [127]. A detector is considered to have an interaction if its phonon recoil energy is 6σ above its mean noise value. Table 4.5 gives the phonon mean and standard deviations for the detectors. A cut requiring a single detector to have been hit will reject few WIMPs since rarely will a true single scatter in a detector be misidentified as a multiple scatter due to unusual noise in another detector.

The efficiency of the singles cut is measured using the "noise triggers" from masked WIMP-search data that passed the bad events cut described in Section 4.3. The efficiency is the fraction of events in a particular detector for which no other detector had an interaction. The measured efficiency of the singles cut is greater than 99.97% in all detectors [128]. Requiring events with single detector hits has a high signal efficiency. Table 4.5 gives the efficiency of the single events cut for each detector. More noise events in T2Z3 were outside the allowed $\pm 6\sigma$ range decreasing the measured singles efficiency of the other

detectors.

4.8 Inner Electrode Cut

Events occurring near the outer, vertical surfaces of a detector, shown in Fig. 3.15, can have uncollected ionization energy. These events can be rejected by excluding events for which the collected ionization energy is not contained within the inner ionization electrode. This rejection is accomplished by selecting only events for which the energy in the outer electrode is consistent with noise [129]. It was shown in section 3.3, based on the geometry of the electrodes, that a cut requiring all the ionization energy to be deposited in the inner electrode would be expect to have an efficiency of about 82%.

Because neutrons uniformly illuminate the detectors, nuclear recoils from neutron calibration data can be used to measure the achieved efficiency of the inner electrode cut. A reasonable estimate can be made by measuring the fraction of nuclear recoils from neutron calibration data that pass the inner electrode cut. This uncorrected efficiency is shown in Figure 4.12 for a sample detector.

In practice, the measured efficiency is expected to be biased low for two reasons.

- Unlike WIMPs, neutrons can multiple scatter within a given detector. Neutrons that scatter within the region under the inner electrode and

scatter within the region covered by the outer electrode will fail the inner electrode cut. This decreases the measured efficiency and biases a measure of the inner electrode efficiency low. This bias is not corrected for.

- Electromagnetic events near a detector surface can mimic nuclear recoils due to their lost ionization. Most events in neutron calibration data are from electromagnetic interactions, so any sample of nuclear recoils will be contaminated by surface events. A fraction of the events used to measure the efficiency of the inner electrode cut will be surface events. A disproportionate fraction of surface events are distributed under the outer electrode biasing the efficiency estimate low. One can attempt to correct for this bias by using gamma calibration data to estimate the number of electromagnetic events that mimic nuclear recoils. Two different corrections will be described below.

Two similar calculations of the inner electrode efficiency were made using slightly different techniques to attempt to correct for the bias due to electromagnetic events near a detector surface. Using data from the inner electrode alone will not reveal energy depositions under the outer electrode. To avoid bias, it is important to use yield and phonon recoil energy measures that use both the inner and outer electrode energies because these measures use both inner and outer electrode energies. Electromagnetic events in the nuclear recoil band contained (not contained) within the inner electrode raise (lower) an es-

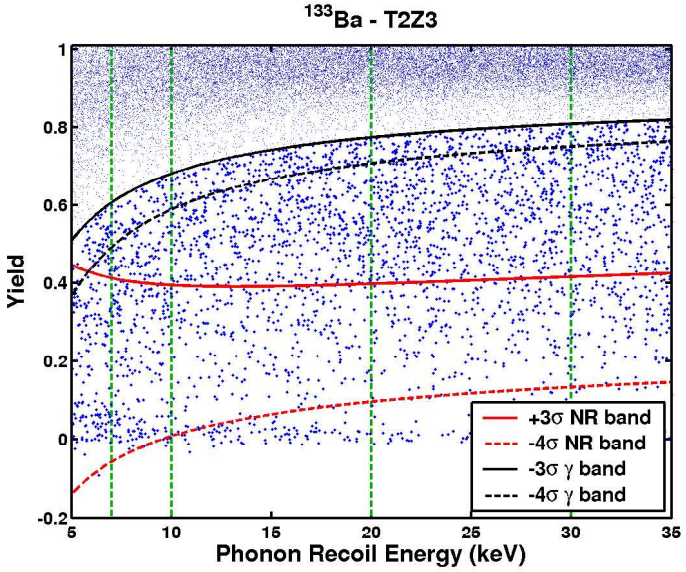

Figure 4.10: Defining correction ratios in a sample detector. The blue dots are events from gamma calibration data. Events below the -3σ gamma band (solid black) and above the 3σ nuclear recoil band (solid red) are used to estimate correction ratios into the +3σ to -4σ nuclear recoil band (dashed red). The -4σ (black dashed) gamma band is shown for reference. Energy bins are denoted by vertical lines (dashed green).

timate of the inner electrode efficiency. Both correction techniques use gamma calibration data to estimate the numbers of electromagnetic events contained and not contained within the inner electrode in the nuclear recoil band for neutron calibration data.

The first method was devised by Rupak Mahapatra and I [130]. Correction ratios were defined to be the ratio of the number of events for which the yield had drooped below the 3σ gamma band, but not into the 3σ nuclear recoil band, to the number of events that had fallen into the nuclear recoil band in

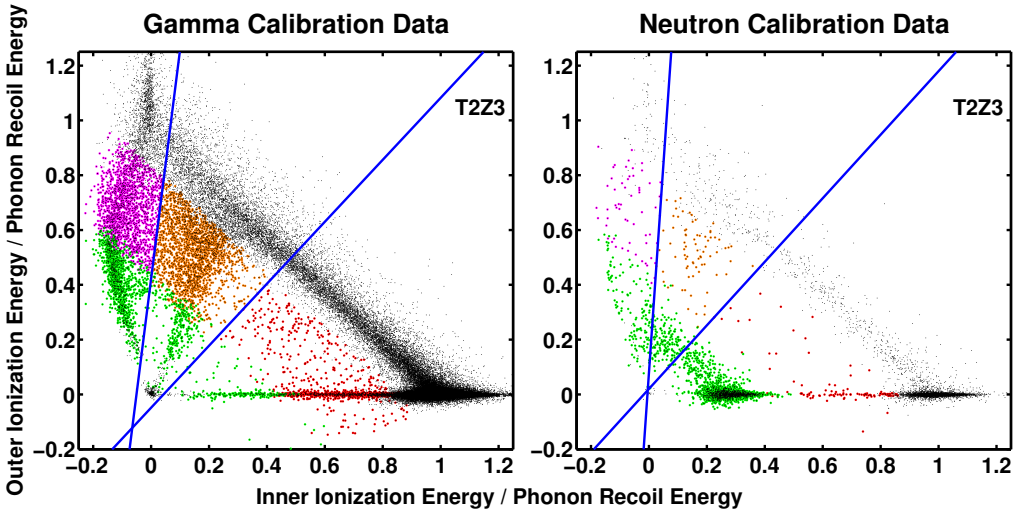

Figure 4.11: Defining and applying correction ratios in a sample detector. The three regions in each plot are (upper left) events occurring under the outer electrode, (upper center, between the blue lines) events occurring between the electrodes, and (right) events occurring under the inner electrode. The black dots are good events with phonon recoil above 10 keV. Events drooping below the gamma band, but not into the nuclear recoil band, are highlighted in magenta (upper left), orange (upper center), and red (right). Events in the nuclear recoil band are highlighted green. The ratio of drooping events to nuclear recoil in each region of gamma calibration data is used to define correction ratios. The correction ratios are normalized by the number of drooping events in neutron calibration data to subtract drooping events from the three regions of the nuclear recoil band.

the gamma calibration data. A wider, -4σ to +3σ, definition of the nuclear recoil band was used. Correction rations were defined for events clearly in qouter, events clearly in qinner, and shared events for each of the 6 energy bins. Figure 4.11 illustrates how the correction ratios were defined for the three classes of deposited ionization energy. The energy bins used were 7 to 10 keV, 10 to 20 keV, 20 to 30 keV, 30 to 40 keV, 40 to 60 keV, and 60 to 100 keV. Figure 4.10 shows the populations of events used to define the correction ratios for the lowest energy bins for a sample detector. The correction ratios were multiplied by the number of events between the gamma and nuclear recoil band in the neutron calibration data to estimate the number of events that needed to be removed from the efficiency estimate. Events under the outer electrode and under the shared region (inner electrode) fail (pass) the inner electrode cut. Removal of drooping events from these regions increase (decrease) the efficiency estimate. Some of the low energy bins did not contain enough events to use and their correction ratios were set to zero. This is the case for the 7-10keV bin in the silicon detectors.

The second method was calculated by R. W. Ogburn [131] and is similar but is normalized to the number of gammas contained within the 3σ gamma band. Figure 4.10 shows the gammas contained within the 3σ gamma band for a sample detector. The gamma calibration events were broken down into the same 6 energy bins. The energy bins were further divided, again using outer ionization electrode yield versus inner ionization electrode yield plane,

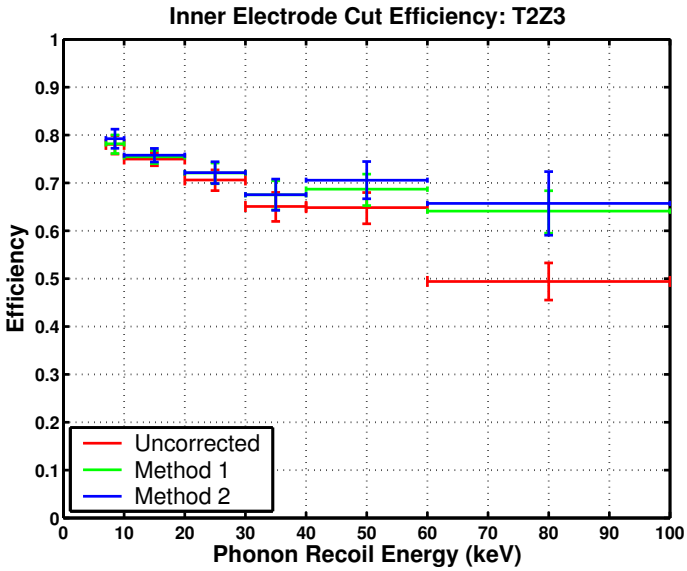

Figure 4.12: Comparison of the different methods for estimating the efficiency of the inner electrode cut for a sample detector, T2Z3. 1σ errors are shown.

into 6 regions of varying outer electrode yield versus inner electrode yield. The correction ratios were calculated using the total number of gammas in each energy bin relative to the number of events in the same $+3\sigma$ to -4σ nuclear recoil band. Figure 4.12 shows the efficiency of the inner electrode cut before and after application of both corrections for a sample detector.

The two measures of the inner electrode cut efficiency were fit using a least squares fit to the following functional form: efficiency $= C_1 + C_2/(E_r - C_3)^{C_4}$ [122]. Typical values for the fit coefficients are $C_1 = -1$, $C_2 = 2$, $C_3 = -6$, and $C_4 = 0.1$. The efficiency resulting from the functional fit for all detectors is shown in Fig. 4.13. To be conservative, the lower fit efficiency from the two methods is used.

4.9 Veto Anti-Coincidence Cut

The importance of rejecting events that are veto coincident was discussed in Section 3.2.2. These events are rejected by the veto anti-coincidence cut. The veto anti-coincidence cut has two parts.

First, detector events that are near in time to veto events of above-threshold pulseheight are rejected. Since the detector phonon pulses are significantly slower than the veto pulses, this is accomplished by rejecting events with veto activity shortly preceding the detector trigger. Figure 4.14 shows a histogram of the delay of (typically high energy) detector phonon triggers relative to the time of the incident muon. The typical phonon trigger delay increases

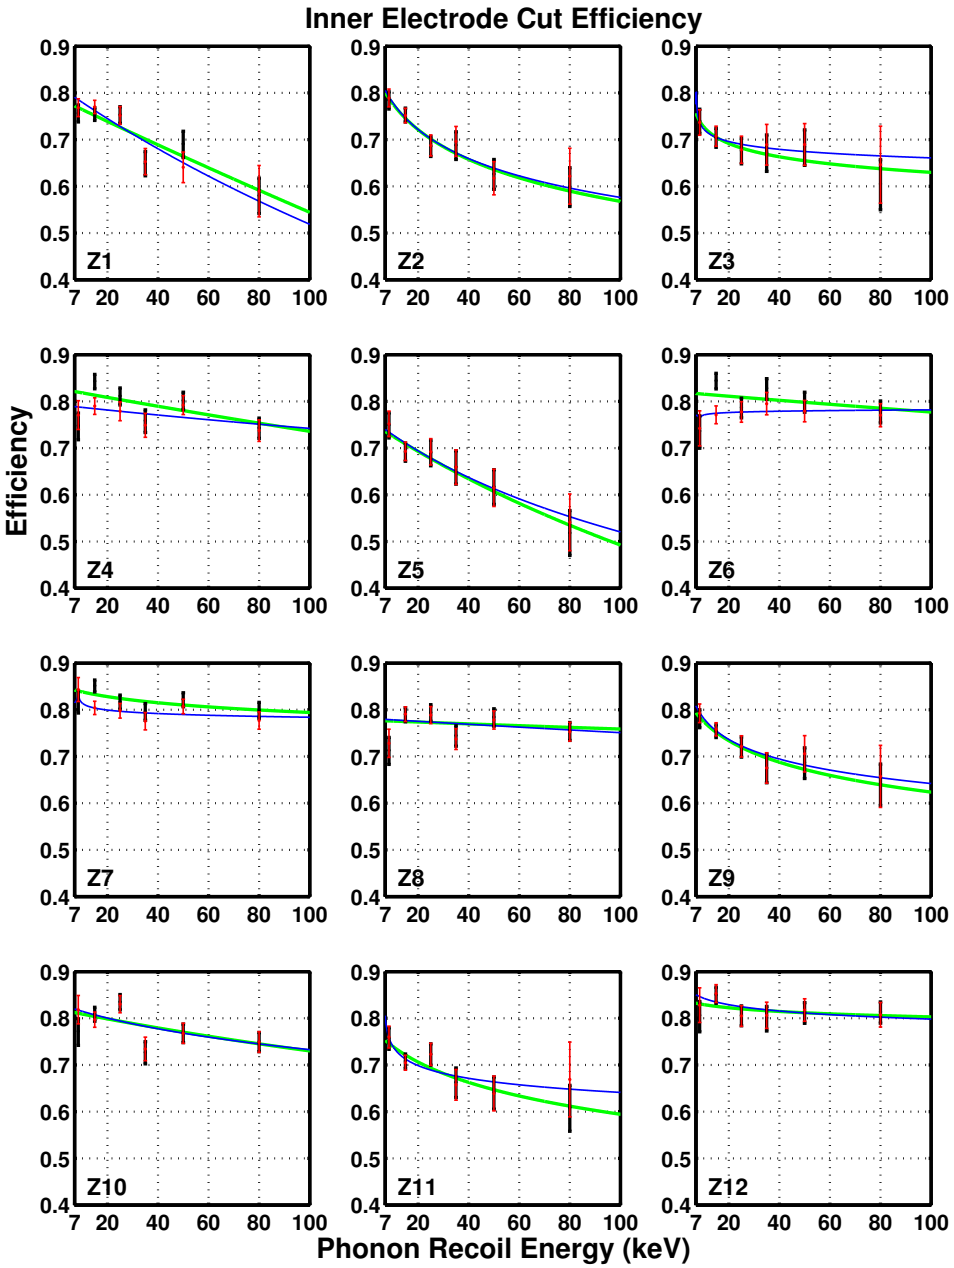

Figure 4.13: The inner electrode cut efficiencies for all twelve detectors. The black (red) data points show the efficiency from method 1 (method 2) with 1σ error bars. The green curve (blue curve) shows the best fit efficiency for method 1 (method 2).

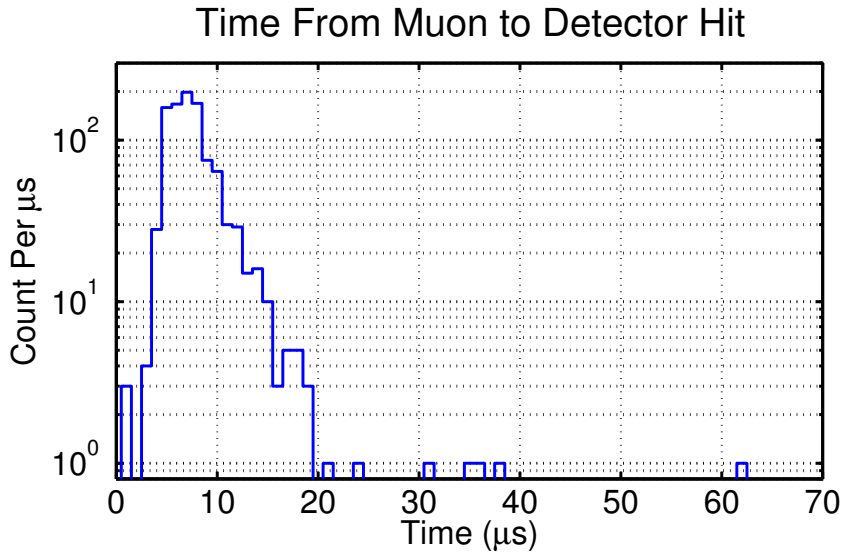

Figure 4.14: Histogram of the time delay from a muon incident on the veto to the first detector trigger. Shown delays are for events with phonon recoil energies above 2 keV in the 9 detectors used in the analysis. Of the events histogramed, 75% are above 100 keV while only 11% are below 10 keV.

as the deposited energy decreases. The results from an earlier run at the shallow Stanford Underground Facility, which had a larger muon flux, rejected events with veto activity in the 40 μs preceding the detector trigger [132]. To be conservative, events with veto activity in the 50 μs preceding the detector trigger are rejected.

Second, rare events can have a veto panel with a large, muon-like, energy deposition well before the detector trigger. Any such event within the entire 185 μs veto trace preceding the detector trigger is rejected.

The veto anti-coincidence cut causes a loss in efficiency for accepting WIMP

interactions. Section 3.2.2 gave a calculation of the rate of veto activity to be 603±2Hz. The probability of an accidental veto event is then 1 - exp(-$50\times10^{-6}\,\mathrm{s}\times603\,\mathrm{Hz}$) or 0.0303±0.0001. The efficiency of the veto anti-coincidence cut is then 0.9697±0.0001. The efficiency of the other cuts are not determined to this degree of precision, so the efficiency is rounded down to 0.969.

4.10 Nuclear Recoil Band Cut

Nuclear recoils typically deposit a higher fraction of their energy in phonons then in ionization than electromagnetic events and hence have a lower yield [133][134][135][136]. Nuclear recoils typically have a yield of about 0.3 that rises slowly with recoil energy. For a variety of reasons, including detector resolution, areas of poor ionization collection, and internal multiple scatters, the yield of nuclear recoils have some scatter about the mean yield for a given recoil energy. The signal region, termed the nuclear recoil band, containing all nuclear recoils within 2σ of the mean yield for a given energy is defined.

Nuclear recoils are characterized *in situ* using calibration data taken with a neutron source, ^{252}Cf. ^{252}Cf also emits gammas. Of the 215,000 good events obtained with the neutron source, 24,000 were good neutron calibration events with phonon recoil energies between 7 and 100 keV. These nuclear recoils were used to define nuclear recoil bands in the yield-recoil energy plane containing nuclear recoils within $\pm2\sigma$ of a mean nuclear recoil with similar recoil energy.

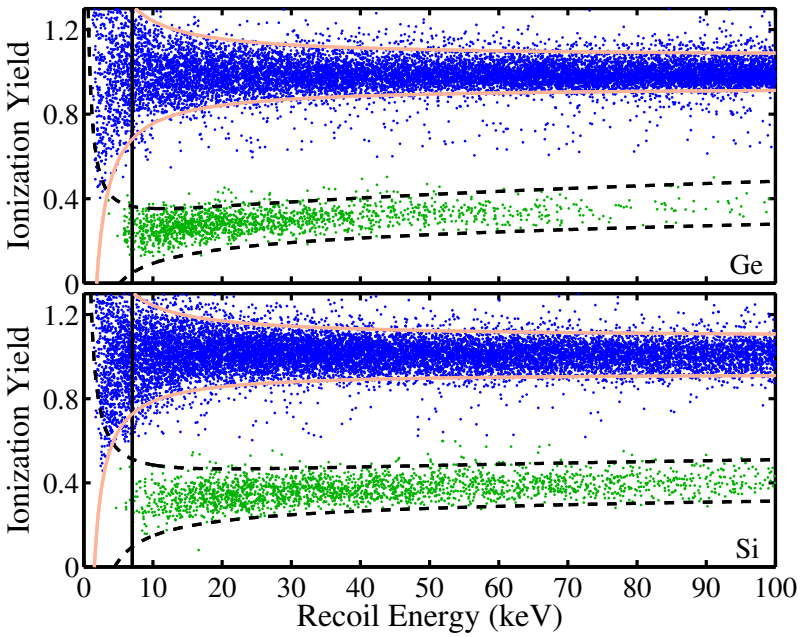

Figure 4.15: Defining the nuclear recoil bands from neutron calibration data. Nuclear recoils are shown in dark green and electromagnetic events in blue. The vertical, solid black line indicates the 7keV threshold (20keV in T1Z4). The dashed, black lines (solid, peach lines) show the 2σ nuclear recoil bands (2σ gamma bands) in T2Z5 for germanium and T2Z6 for silicon. Events in other detectors are rescaled to these detectors.

Figure 4.15 shows the neutron calibration data used to define the nuclear recoil bands. Because of the small number of events in the germanium detectors above 60keV, the mean of the nuclear recoil band at high energy was set by combining the events from the germanium detectors. Events in T1Z3 were not combined with the other germanium detectors since T1Z3's nuclear recoil band behaves somewhat differently.

Because the nuclear recoil region is designed to contain nuclear recoils within

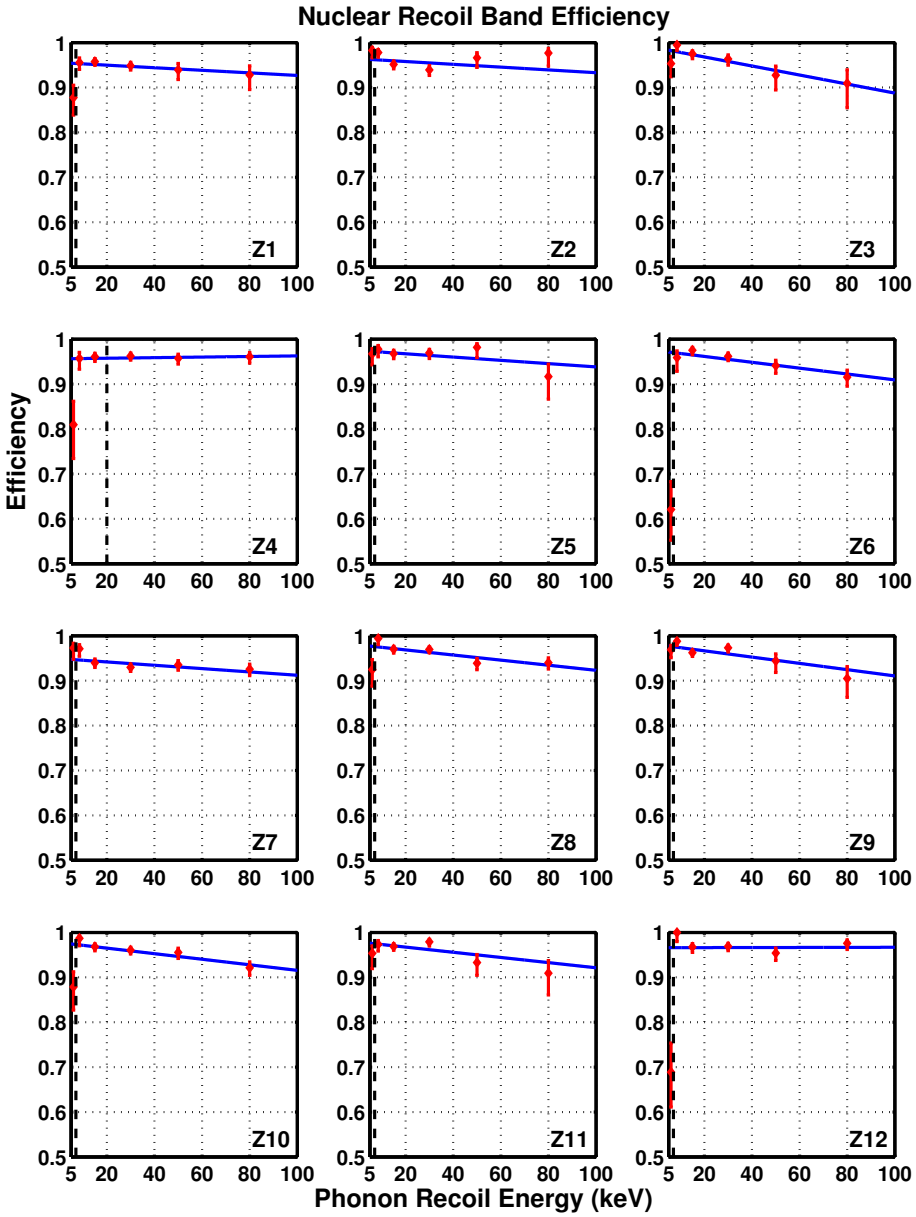

Figure 4.16: The 2σ nuclear recoil band efficiency trigger efficiencies for all twelve detectors. The red data points show the measured nuclear recoil band efficiencies with 1σ error bars. The 2σ nuclear recoil band efficiency is expected to be 0.9545. When the fit efficiency is above this level, it is assumed to be 0.9545. The black, dashed, vertical line shows the analysis threshold.

$\pm 2\sigma$ of a mean nuclear recoil, the nuclear recoil band efficiencies are expected to be about 95%. The nuclear recoil band efficiencies are fit for each detector by assuming an efficiency that is linear with energy with an offset. To be conservative, an efficiency of 0.9545 is used when the fit efficiency is greater than 0.9545. Figure 4.16 shows the fit nuclear recoil band efficiency for each detector.

4.11 Phonon Timing Outlier Cut

A crucial upcoming step in the analysis is to define a cut to reject events near detector surfaces. The surface event rejection cut will utilize the timing of the four phonon pulses for each event. As was pointed out in section 3.3, the timing of the four phonon pulses is used to reconstruct the x-y position of each event as in Fig. 3.17. An event whose reconstructed position based on phonon delay fell significantly outside that of the normal detector position reconstruction would have a suspect value for its surface event rejection cut and hence could not be considered a reliable candidate event. These outlying, phonon-delay position events are rejected by a phonon delay outlier cut. The values of the phonon delay outlier cut are given in Table 4.6. Note that the phonon propagation velocity is about twice as high in silicon as it is in germanium leading to the phonon delay plots for silicon detectors to have a radius that is about half the size as that of the germanium detectors. The phonon delay outlier cut is therefore about twice as tight in the silicon detectors.

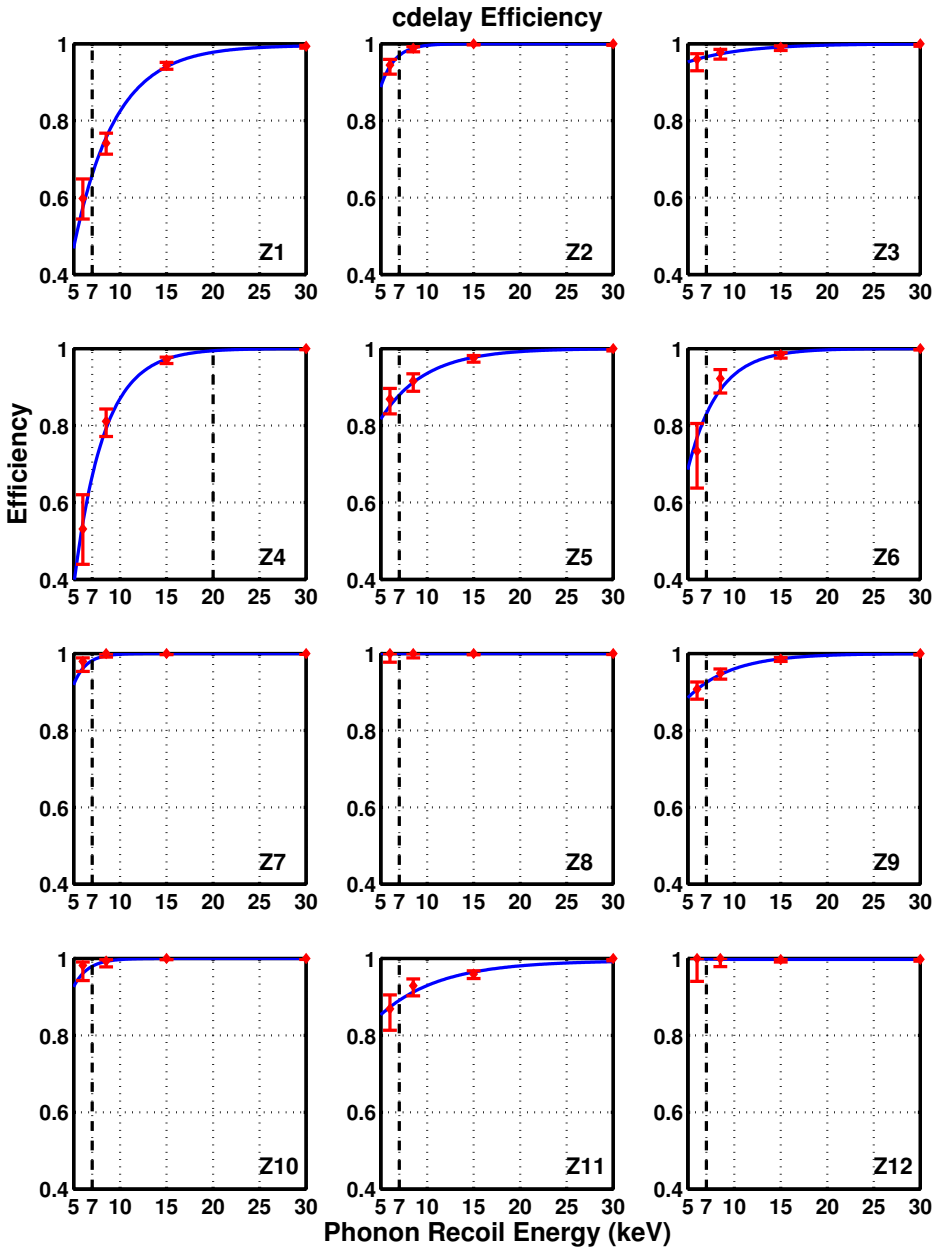

Figure 4.17: The efficiency of the phonon timing outlier cut for all 12 detectors. The red data points show the measured nuclear recoil band efficiencies with 1σ error bars. The blue curves shows the best fit efficiencies. The vertical, dashed line shows the low recoil energy threshold.

Tower1	T1Z2	T1Z2	T1Z3	T1Z4	T1Z5	T1Z6
Min (μs)	-70	-50	-50	-35	-40	-30
Max (μs)	70	50	50	35	40	30
Tower2	T2Z2	T2Z2	T2Z3	T2Z4	T2Z5	T2Z6
Min (μs)	-20	-20	-40	-20	-45	-20
Min (μs)	20	20	40	20	45	20

Table 4.6: List of detector phonon delay outlier cut values in μs.

The efficiency of the phonon timing outlier cut is estimated from neutron calibration data using good nuclear recoils passing the ionization threshold cut, the inner electrode cut, the T2Z5 cliff cut, and the veto anti-coincidence cut. The efficiency is nearly 100% above 20keV in all detectors and is found to fall rapidly at low phonon recoil energy and is fit by the following functional form: efficiency $= C_1 * e^{-C_2*(E_r - C_3)} + C_4$ [122]. Typical values for the fit coefficients are C_1 = -0.2, $C_2 = 0.6$, $C_3 = 3$, and $C_4 = 1$. Figure 4.17 shows the fit efficiency for each detector.

4.12 Surface Event Rejection Cut

As mentioned before, events near detector surfaces can suffer from incomplete ionization energy collection causing their yield to decrease from the gamma band towards the nuclear recoil band, as shown in Fig. 4.18. Events within about 5 microns of a detector surface suffer incomplete ionization collection. Events within $1\,\mu$m of the surface suffer nearly complete loss of their

ionization collection; events deeper than $\simeq$20 μm suffer only negligable loss of ionization collection [137].

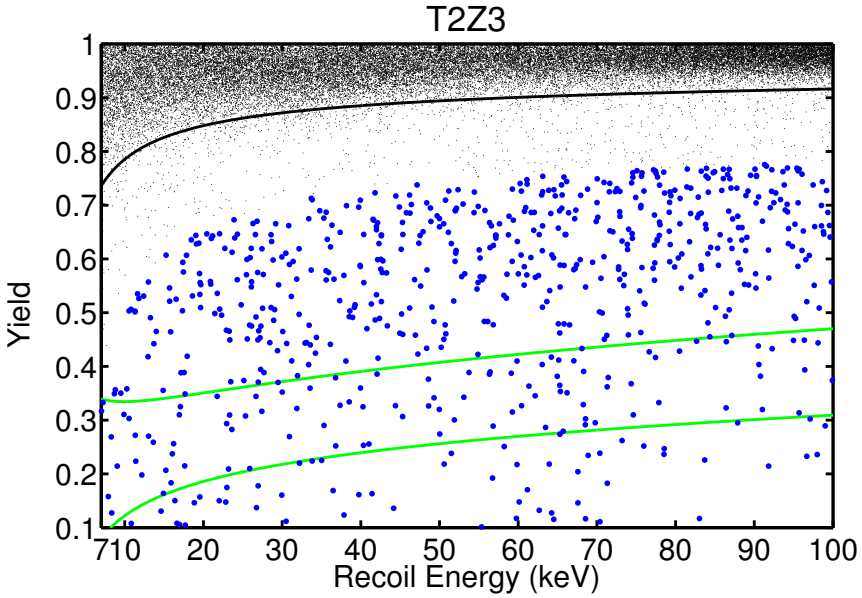

Figure 4.18: Yield-based surface event rejection in a sample detector, T2Z3, for open gamma calibration data. The green curves (black curve) show the bounds of the nuclear recoil band (lower bound of the gamma band). Events with yield more than 5σ below the mean gamma yield are assumed to be surface events and are denoted by blue dots; all other events are shown by small black dots.

Surface events occur on all surfaces of a detector. The application of prior cuts, excluding the nuclear recoil band cut, reject many of these surface events. For example, the inner electrode cut, described in Section 4.8, rejects events near the vertical, outer, detector surfaces. Events near the inner regions of the horizontal detector faces must still be rejected. As Fig. 4.18 shows, most of these low yield or surface events are rejected because they fall outside the

nuclear recoil band. The surface event rejection cut is intended to reject surface events that are inside the band.

4.12.1 Classification of Surface Events

There are three classes of background from surface events:

1. Ambient gammas that Compton-scatter off electrons in a detector and sometimes scatter in the first few microns of a detector. An electron that is scattered this near to the surface can be ejected from the detector toward the adjacent detector and is called an ejectron. Ejectrons frequently deposit energy in the adjacent detector.

2. Electrons Compton-scattered near the detector surface that do not escape the detector.

3. Electrons from beta-decay of radioactive contaminants on detector surfaces. Decays from ambient radon in the atmosphere implant ^{210}Pb less than a micron into the surface.

These three classes do not contribute surface events equally to calibration data and WIMP-search data. Most of the surface events in calibration data are from the first two classes of surface events because of the high rate of ambient gammas. Many of the surface events in WIMP-search data, however, are from the third class of surface events.

Calibration data is taken for a much shorter duration than WIMP-search data is. Gamma calibration data was taken for a total of about 4 to 5 days of livetime while neutron calibration data was taken for 2.6 hours of livetime. In contrast, recall that WIMP-search data was taken for a livetime of about 75 days. Because of the low duration of the calibration data, decay of radioactive isotopes on the detector surfaces contributes a much smaller fraction of the surface events in calibration data than in WIMP-search data.

4.12.2 Defining a Sample of Surface Events

We wish to obtain a sample of surface events for the purpose of defining and testing the surface event rejection cut. A sample of surface events with primarily ejectrons can be obtained by illuminating the detectors with a gamma source. Figure 4.18 shows events from open gamma calibration data for a sample detector; surface events are highlighted by large blue dots. This sample of surface events is used to define a surface event rejection cut using variables other than yield.

It is beneficial to have a sample with a large number of surface events to use in defining the cut. A larger sample decreases fluctuations in the surface event distributions. An obvious choice for a sample of surface events might be surface events within the nuclear recoil band of gamma calibration data. Since the number of surface electron recoils within the nuclear recoil band is rather limited even in gamma calibration data, it is desirable to define a wider

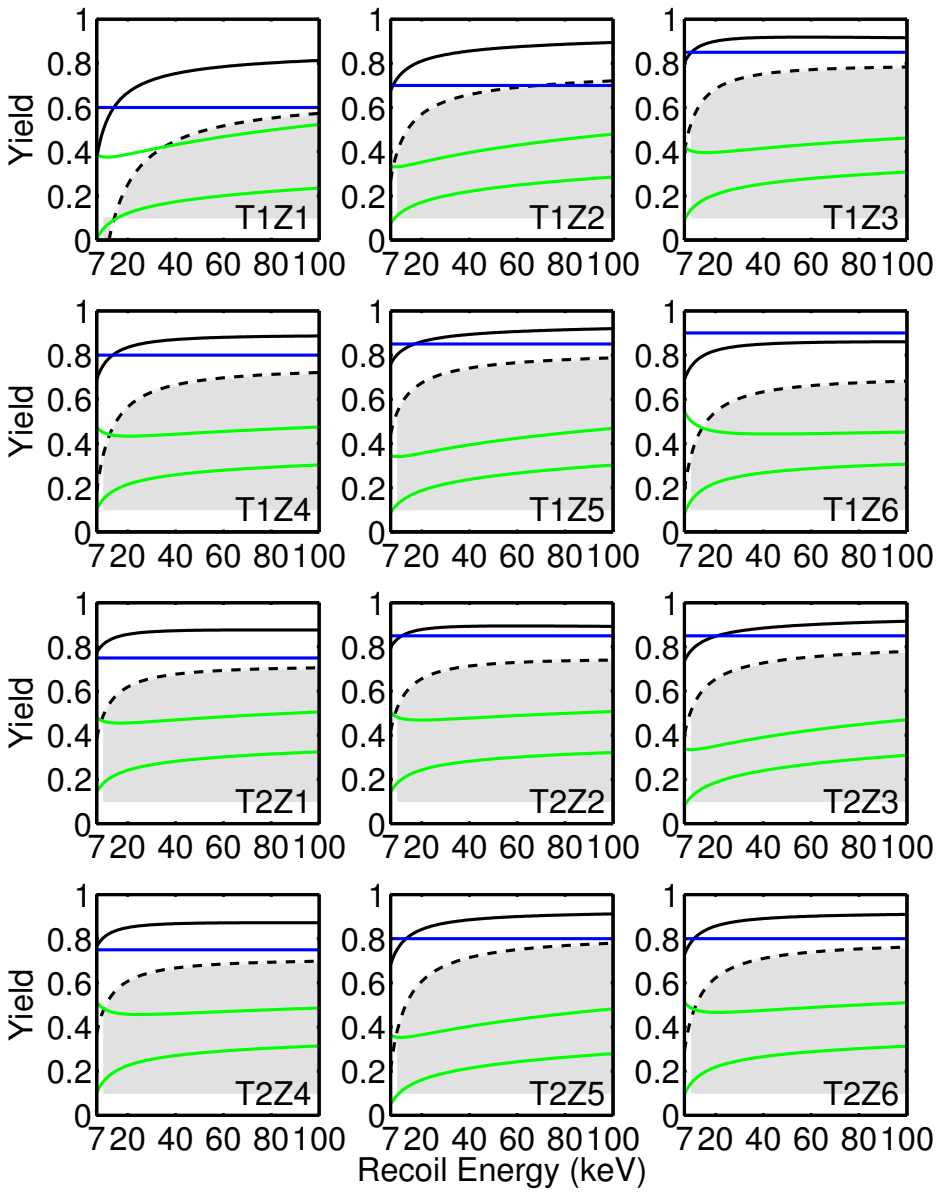

Figure 4.19: Figure showing the "wide beta" region of each detector in light gray. The solid (dashed) black curve shows the lower end of the 2σ (5σ) gamma band. The solid blue line shows the requirement that the yield be less than the value given in Table 4.7. The green curves show the nuclear recoil band.

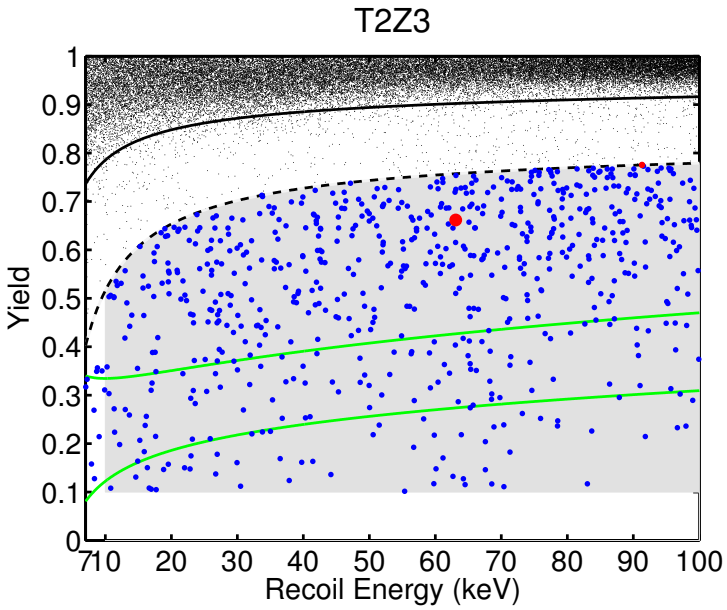

Figure 4.20: Surface events from a sample detector during open gamma calibration data. The gray shaded area is the "wide beta" region. The green curves show the nuclear recoil band. The solid (dashed) black curve shows the lower end of the 2σ (5σ) gamma band. Black (blue) dots denote bulk (surface) electron recoils. The large (small) red dots denote events that passed the primary (less restrictive secondary) cut.

selection region.

A "wide beta" region is used to define and test the surface event rejection cut. Surface events inside the "wide beta" region are assumed to have properties similar to surface events within the 2σ nuclear recoil band. The "wide beta" region was defined by Angela Reisetter [138] to contain events with yield more than 5σ below mean gamma yield, with yields greater than 0.1, and with recoil energies between 10 (20 in T1Z4) and 100 keV. Some higher energy electron

Tower	Z1	Z2	Z3	Z4	Z5	Z6
T1	0.60	0.70	0.85	0.80	0.85	0.90
T2	0.75	0.85	0.85	0.75	0.80	0.80

Table 4.7: List yield values for the "wide beta" region. Events in the wide beta region were required to have a yield below that listed in this table. Only events with recoil energies above 69 keV in T1Z2 are affected by this requirement.

recoils in the detector bulk had yields below the 5σ gamma band. To guarantee that higher energy electron recoils in the detector bulk did not contaminate the sample of surface electron recoils, each detector's yield was additionally required to be less than the value listed in Table 4.7. Figure 4.19 shows the "wide beta" region for each detector. Notice that in practice only T1Z2's "wide beta" region is affected by the requirement that its yield be less than the value given in Table 4.7. Figure 4.20 shows the surface events within the "wide beta" region of open gamma calibration data for a sample detector. Surface events within the "wide beta" region of open gamma calibration data will be referred to as WBO events.

4.12.3 Tools for Rejecting Surface Events

Most surface events are rejected by the cuts described earlier in this chapter. The remaining surface events can be rejected by a surface event rejection cut which exploits differences between the phonon signals for surface events and the phonon signals for events within the bulk of a detector (bulk events).

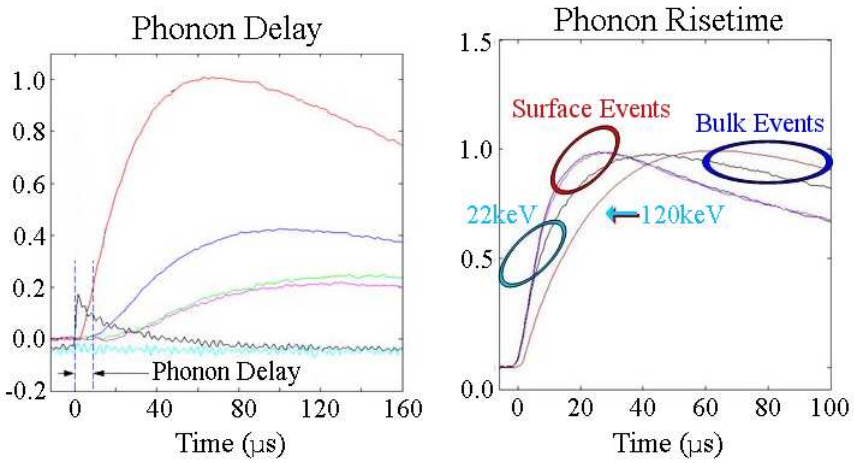

Figure 4.21: Sample phonon delays and risetimes. The plot on the left shows the phonon delay of the four phonon channels relative to the ionization pulse for a sample event. The plot on the right shows the risetime for sample phonon pulses. The risetime of bulk events is greater, but decreases with decreasing energy. These surface events are from calibration data taken using a beta source. The surface events shown in were obtained by shining a collimated beta source on a detector [139]. Most events pass through the collimator and occur in a single location on the detector. Events in other locations in the detector are bulk events.

The phonon pulses from each of the detector's four quadrants are different for surface events than for nuclear recoils in the bulk in three respects. Three surface event rejection parameters are defined to exploit each of these differences. The three surface event rejection parameters are:

1. Phonon delay, x_1 - phonon pulses from surface events arrive sooner relative to the ionization pulse than do the phonon pulses for nuclear recoils as illustrated by the left-hand panel of Fig. 4.21. The phonon delay is defined to be the time from the start of the ionization pulse to the time at which the largest phonon pulse rises to 20% of its maximum amplitude. Typical phonon delays are 2 to 15 μs. Pulses from surface events typically arrive sooner because interactions of the initially high frequency (slower) phonons with the surface metal cause quicker downconversion to lower frequency (faster) phonons [97] [140].

2. Phonon Risetime, x_2 - the time for a phonon pulse to rise for a surface event is shorter than that for a nuclear recoil. The risetime is defined to be the time for the pulse to rise from 10% to 40% of its maximum amplitude. Typical risetimes are 8 to 11 μs. The right-hand panel of Fig. 4.21 shows the delay of the four phonon pulses for a sample event. Pulses from surface events typically rise faster because interactions of the initially high frequency (slower) phonons with the surface metal cause quicker downconversion to lower frequency (faster) phonons [97] [140].

3. Energy partition, x_3 - the ratio of energy deposited in the local quadrant to the energy deposited in the diagonally opposite quadrant. Surface events, particularly at low energy, have a weak tendency to deposit a higher fraction of their energy within the local phonon quadrent. A large fraction of the surface events occuring on the phonon sensor side of the detector will typically be incident under a single phonon sensor, thus increasing the energy partition.

Figure 4.22 and Fig. 4.23 show the distribution of the surface event rejection parameters for WBO surface events and nuclear recoils from the 2σ nuclear recoil band of neutron calibration data for a sample detector.

We define a surface event rejection cut that utilizes these three different surface event rejection parameters to efficiently reject surface events while allowing a high fraction of nuclear recoils to survive. In the past, a simple two dimensional cut using the first two surface event rejection parameters was used [97] [141]. All the surface event rejection parameters are correlated, so it would be preferable to design a new surface event rejection cut accounted for the correlations. Section 4.12.4 will describe the definition of this cut using the WBO surface events.

4.12.4 A Chi Squared Analysis

An analysis is performed using two χ^2 deviations: 1) the χ^2 deviation for each event from the nuclear recoil or signal hypothesis (χ^2_n) and 2) the χ^2 devi-

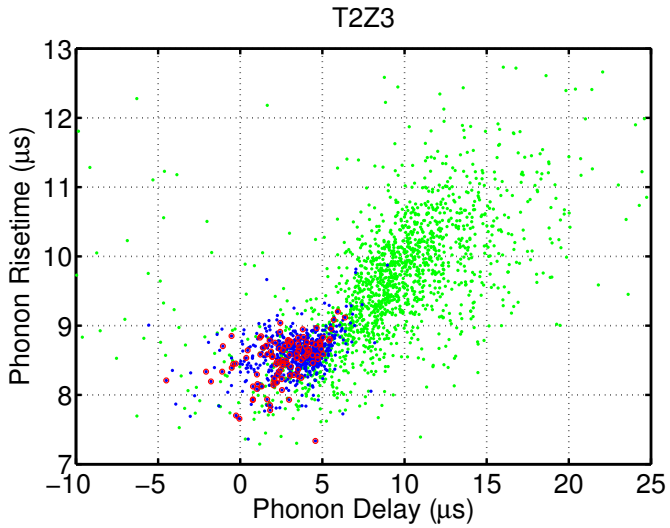

Figure 4.22: A plot of two surface event rejection parameters, phonon delay and risetime used to reject surface events in T2Z3. WBO surface events are plotted in blue. WBO surface events within the nuclear recoil band are indicated by red circles. Nuclear recoils from neutron calibration data are plotted in green.

ation from the surface background hypothesis (χ^2_b). A χ^2 analysis is typically used with Gaussian distributions. While the distributions of surface events and nuclear recoils are not precisely Gaussian, a χ^2 analysis can still be effective for separating different hypotheses.

We calculate the χ^2 values as follows. Nuclear recoils (and surface events) have distributions of the 3 rejection parameters described in Section 4.12.3. The distributions of these 3 parameters (the phonon risetimes, phonon delays, and

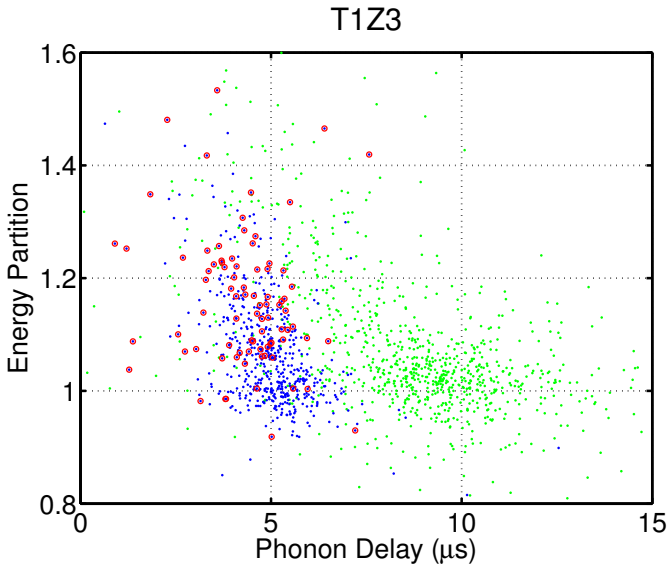

Figure 4.23: A plot of two surface event rejection parameters, phonon delay and energy partition, used to reject surface events in T1Z3. WBO surface events are plotted in blue. WBO surface events within the nuclear recoil band are indicated by red circles. Nuclear recoils from neutron calibration data are plotted in green.

energy partitions) are each characterized by a mean and standard deviation:

$$\begin{aligned} \mu_{ij} &= N^{-1}\sum_{k=1}^{N} x_{ijk} \\ \sigma_{ij} &= \sqrt{\frac{1}{N-1}\sum_{k=1}^{N}(x_{ijk}-\mu_{ij})^2} \end{aligned} \tag{4.12.1}$$

Here x_{ijk} is the value of the jth surface event rejection parameter for the kth event. The index i takes a value of n for the signal hypothesis or a value of b for the background hypothesis. μ_{nj} and σ_{nj} are defined using nuclear recoils

from neutron calibration data, and μ_{bj} and σ_{bj} are defined using WBO event. Note that μ_{ij} and σ_{ij} neglect any variation with recoil energy.

The element of the covariance matrix containing the pth row and qth column element expresses the covariance between the pth and qth variables and is defined by:

$$C_{ipq} = \frac{\Sigma_{k=1}^{N} \left(x_{ipk} - \mu_{ip}\right)\left(x_{iqk} - \mu_{iq}\right)}{N-1}. \tag{4.12.2}$$

The C_{npq} (C_{bpq}) elements are defined using neutron calibration data (open gamma calibration data). Because the correlation between variable p and variable q must be the same as the correlation between variable q and variable p, the correlation matrix is symmetric.

The χ^2 for the kth event, or the 'distance' of this event from the mean of the signal ($i = n$) or background ($i = b$) distributions, is calculated as:

$$\chi^2_{ik} = \begin{pmatrix} x_{i1k} - \mu_{i1} \\ x_{i2k} - \mu_{i2} \\ x_{i3k} - \mu_{i3} \end{pmatrix}^T \mathbf{C_i^{-1}} \begin{pmatrix} x_{i1k} - \mu_{i1} \\ x_{i2k} - \mu_{i2} \\ x_{i3k} - \mu_{i3} \end{pmatrix}. \tag{4.12.3}$$

The distributions of χ^2_b and χ^2_n are shown for a sample germanium detector in Fig. 4.24 and for a sample silicon detector in Fig. 4.25.

If for the kth event, χ^2_{bk} is less (greater) than χ^2_{nk}, then the event is more likely to be a surface event (nuclear recoil). Figure 4.26 shows a plot of χ^2_b versus χ^2_n for sample germanium and silicon detectors.

Two requirements on the values of χ^2_n and χ^2_b are imposed for an event to be considered a potential signal event. These two requirements are:

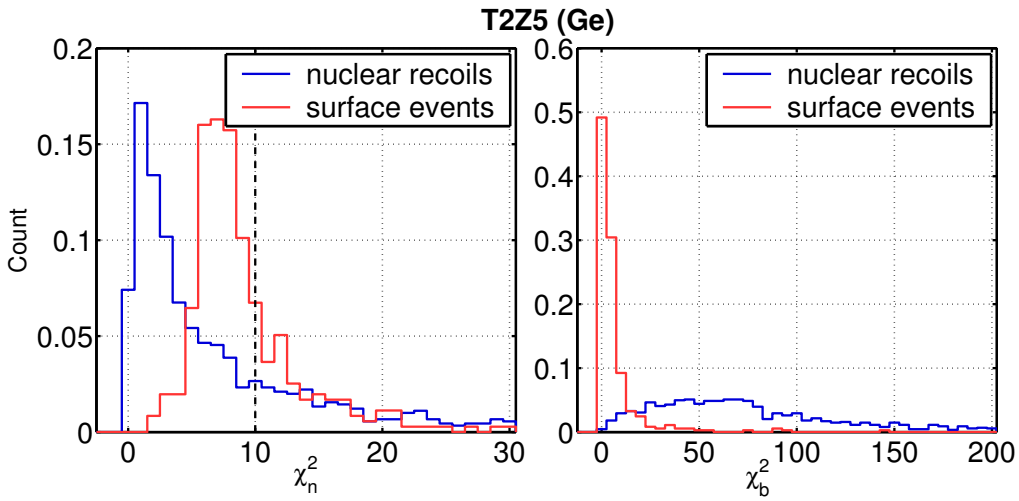

Figure 4.24: Histograms showing the distribution of χ^2_b and χ^2_n for a sample germanium detector. The histograms were made used nuclear recoils from neutron calibration data and WBO surface events. The black vertical dashed line in the left hand plot denotes the "outlier cut". Events with χ^2_n greater than 10 are not typical of nuclear recoils and are rejected. The area of each distribution is normalized to unity.

- A potential signal event must have a value of χ^2_n that is consistent with that of a nuclear recoil. This requirement is referred to as the "outlier cut". It requires χ^2_n to be less than 10, in order to reject outlier events that do not behave as typical nuclear recoils. If the surface event rejection parameters follow a Gaussian distribution, the "outlier cut" should pass 90% of nuclear recoils (for this 3 paramater χ^2). The distribution of χ^2_n values for surface events and nuclear recoils in not Gaussian, and the fraction of nuclear recoils (surface events) in the germanium detectors passing (rejected by) the "outlier cut" is 0.67±0.01 (0.34±0.01). The

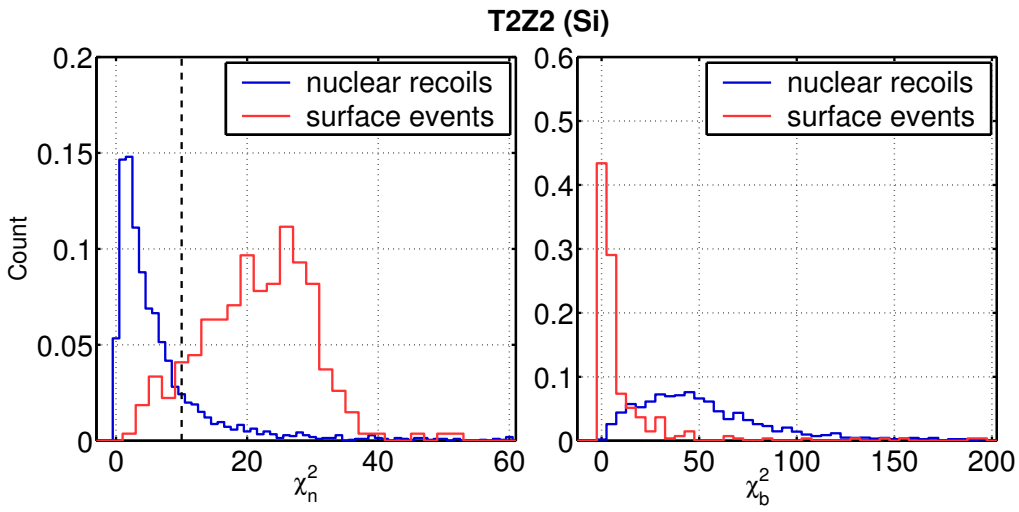

Figure 4.25: Histograms showing the distribution of χ^2_b and χ^2_n for a sample silicon detector. The histograms were made used nuclear recoils from neutron calibration data and WBO surface events. The black vertical dashed line in the left hand plot denotes the "outlier cut". Events with χ^2_n greater than 10 are not typical of nuclear recoils and are rejected. The area of each distribution is normalized to unity.

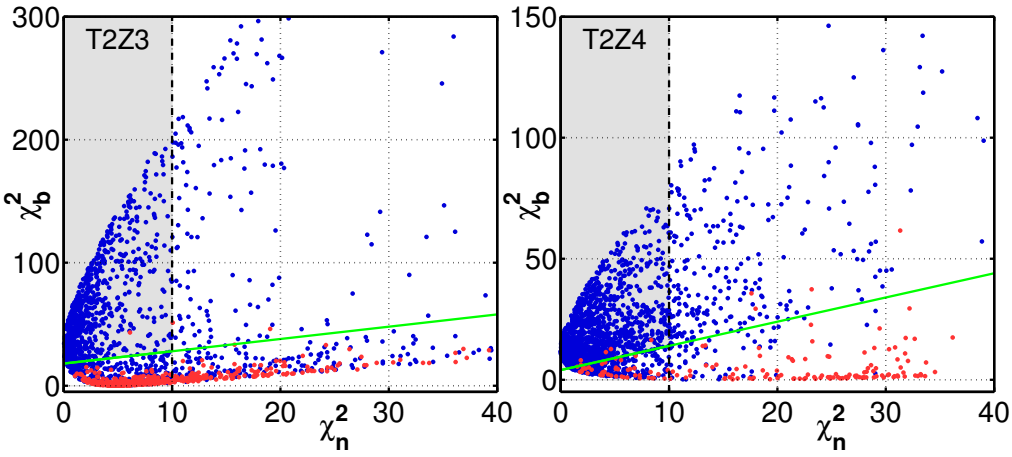

Figure 4.26: A distribution of surface events (light red) and nuclear recoils (dark blue) from calibration data in the χ^2_b versus χ^2_n plane. The left (right) hand pane shows a sample germanium (silicon) detector, T2Z3 (T2Z4). The vertical, dashed black line shows the "outlier cut" described in Section 4.12.6, and the diagonal solid green line shows the surface event rejection cut, $\chi^2_b - \chi^2_n > 18$ (4) for T2Z3 (T2Z4). The signal area is shaded light gray.

fraction of nuclear recoils (surface events) in the silicon detectors passing (rejected by) the "outlier cut" is 0.79±0.01 (0.56±0.02).

- For remaining events, an additional parameter is used to distinguish potential signal events from surface events. This parameter is related to the relative probability of an event being a nuclear recoil or surface event. The χ^2_n (χ^2_b) of a given event is the 'distance' of the event from the mean nuclear recoil (surface event) distribution. The parameter, d_{χ^2},

$$d_{\chi^2} < \ln\left(\frac{e^{-\chi^2_n}}{e^{-\chi^2_b}}\right) = \chi^2_b - \chi^2_n \qquad (4.12.4)$$

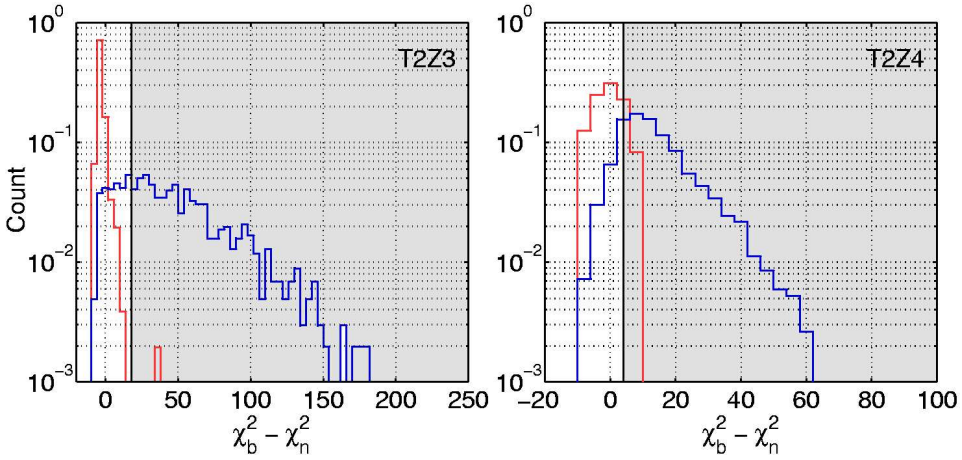

Figure 4.27: A histogram of $\chi^2_b - \chi^2_n$ for WBO surface events (light red) and nuclear recoils (dark blue) from neutron calibration data. The left (right) hand pane shows a sample germanium (silicon) detector, T2Z3 (T2Z4). Events failing the "outlier cut" described in Section 4.12.6 are not shown. The shaded area is the signal region. The area of each distribution is normalized to unity.

where $e^{-\chi^2_n}$ $(e^{-\chi^2_b})$ is proportional to the probability of a particular event being a nuclear recoil (surface event), characterizes whether an event is more like surface background or nuclear recoil. The value of d_{χ^2} should be < 0 for surface events, and > 0 for nuclear recoils. The selection of the particular cut on d_{χ^2} values is described in Section 4.12.8.

Figure 4.27 shows a histogram of the distribution of $\chi^2_b - \chi^2_n$ for surface events and nuclear recoils for sample germanium and silicon detectors. All good germanium detectors had similar distributions of $\chi^2_b - \chi^2_n$, allowing all them to have the same cut on d_{χ^2}. The variation of the populations in silicon is far greater, so each detector was given its own value of d_{χ^2}.

4.12.5 Tools for Estimating the Expected Background

We make two estimates of the number of surface-background events, B, in the 2σ nuclear recoil band signal region, for a given value of d_{χ^2}. The actual signal region for this analysis was blinded, so these estimates entail extrapolation from data samples other than the WIMP-search data for this run.

The first estimate utilizes WBO surface events, $N^o_{wba}(d_{\chi^2})$, with a difference in chi-squared value $\chi^2_b - \chi^2_n$ that exceeds a minimum value d_{χ^2}. We estimate $B_1(d_{\chi^2})$, the number of surface events in the signal region which pass the same cut on the difference in chi-squared value as:

$$B_1(d_{\chi^2}) = S_1 \times N^o_{wba}(d_{\chi^2}), \tag{4.12.5}$$

where S_1 is a scale factor. We assume that the scale factor S_1 is independent of d_{χ^2}, and estimate S_1 as:

$$S_1 = \frac{R_E \times N^p_{2\sigma s}}{N^o_{wba}} = 0.0066^{+0.0032}_{-0.0023} \tag{4.12.6}$$

where $R_E = 2.0 \pm 0.1$ is the efficiency corrected exposure of the current WIMP-search run relative to that of the previous run, $N^p_{2\sigma s} = 8$ is the number of single-scatter events, prior to surface rejection cuts, in the 2σ nuclear recoil band signal region of the *previous* WIMP-search run, and $N^o_{wba} = 2435$ is the number of WBO surface events for the *current* run, also prior to surface rejection cuts. $N^p_{2\sigma s}$ and other key event statistics of the *previous* WIMP-search run are summarized in Table 4.8.

ZIP	Wide Beta All (N^p_{wba})	3σ All ($N^p_{3\sigma a}$)	3σ Singles ($N^p_{3\sigma s}$)	2σ All ($N^p_{2\sigma a}$)	2σ Singles ($N^p_{2\sigma s}$)
T1Z2	65	18	3	13	1
T1Z3	112	47	5	30	3
T1Z5	131	30	9	18	4
Good Ge	308	95	17	61	8
T1Z1	38	86	51	48	30
T1Z4	89	36	11	24	8
T1Z6	417	274	184	193	129

Table 4.8: List of the number of events in the prior WIMP-search run within the "wide beta" region, within the 2 and 3σ nuclear recoil bands, and the number of single scatters within the 2 and 3σ nuclear recoil bands. The final row gives the numbers for the combined good germanium detectors (T1Z2, T1Z3, and T1Z5).

In Section 4.12.8, we will use S_1 to choose the number of WBO surface events, and hence the value of d_{χ^2}, that are allowed to pass the surface event rejection cut.

Other scale factors can be similarly estimated. In Section 4.14.1, a similar scale factor will be used to estimate the surface event background from the "wide beta" region of closed gamma calibration data. Our second estimate of the number of surface-background events will utilize a similar scale factor, S_2.

We use our second estimate to study the behavior of our expected results as a function of d_{χ^2} in the illustrative example in the next section. This estimate utilizes multiple scatter surface events from the "wide beta" region of the current WIMP-search run, $N_{wbm}(d_{\chi^2})$, with a difference in chi-squared value

$\chi_b^2 - \chi_n^2$ which exceeds a minimum value d_{χ^2}. We estimate $B_2(d_{\chi^2})$, the number of surface events in the signal region which pass the same cut on the difference in chi-squared value as:

$$B_2(d_{\chi^2}) = S_2 \times N_{wbm}(d_{\chi^2}), \tag{4.12.7}$$

where S_2 is a scale factor. We again assume that the scale factor S_2 is independent of d_{χ^2}, and estimate S_2 as:

$$S_2 = \frac{N^p_{2\sigma m}}{N^p_{wbm}} \times \frac{N^p_{wbs}}{N^p_{wbm}} = (4.9 \pm 1.0) \times 10^{-2}, \tag{4.12.8}$$

where the values of $N^p_{2\sigma m}$, N^p_{wbm}, and N^p_{wbs} are given in Table 4.9 and the error is statistical. The value of S_2 is estimated using surface events from the "wide beta" region of the *previous* WIMP-search run. Although, multiple scatter surface events and single scatter surface events within the "wide beta" region of the current WIMP-search run were not generally blinded, some subtle issues made use of the surface events from the current WIMP-search data difficult.

We use $B_2(d_{\chi^2})$ to study the variation of our sensitivity and discovery potential as a function of d_{χ^2} in the next section.

4.12.6 Strategies for Choosing a Value of d_{χ^2}

The strength with which we reject surface events, that is, the particular choice of the minimum value d_{χ^2} of $\chi_b^2 - \chi_n^2$ that an event may have and still be considered a nuclear recoil, is the result of a compromise between the desire for high efficiency in retaining true nuclear recoils and the desire for a low

ZIPs	Wide Beta	Wide Beta	Wide Beta	2σ
	All (N^p_{wba})	Mult. (N^p_{wbm})	Singles (N^p_{wbs})	Mult. ($N^p_{2\sigma m}$)
Good Ge	308	250	58	53

Table 4.9: Table containing the total number of events in the prior WIMP-search run at Soudan within the "wide beta" region, the number of multiple scatters within the "wide beta" region, and the number of single scatters within the "wide beta region". The total number of scatters within the 2σ nuclear recoil band is also given. The good germanium detectors are T1Z2, T1Z3, and T1Z5.

background of surface events. When there is no cut at all on $\chi^2_b - \chi^2_n$, which is described by $d_{\chi^2} = -\infty$, the efficiency for true nuclear recoils to pass this cut is 100%, but many surface events will also pass the cut, which makes the detection of a nuclear recoil signal difficult. As the value of d_{χ^2} is raised, surface events will first be preferentially rejected and true nuclear recoils events retained, improving the ability to distinguish nuclear recoil events. If the value of d_{χ^2} is raised too high, however, the efficiency to retain nuclear recoils is again reduced, making it hard to see a signal of nuclear recoils.

In this section we describe studies of the choice of d_{χ^2}. In actual fact, our choice of d_{χ^2} was made under time pressure, because a deadline had been chosen by which all surface event rejection cuts were to be defined, and the cut was to be tested on the closed gamma calibration data. Our choice of d_{χ^2} turned out to be a bit higher (resulting in a stronger rejection of surface background) than our studies here would indicate was optimum. Ideally, we would have done these studies before choosing d_{χ^2}, so we describe the studies now. Our post-

unblinding studies described here indicate that our choice of d_{χ^2} gave nearly optimum sensitivity.

Two approaches for choosing d_{χ^2} are maximizing sensitivity, and maximizing discovery potential. By sensitivity we mean the typical upper limit on the cross-section on WIMP-nucleon interactions, where the smaller the upper limit, the higher the sensitivity. Prior to unblinding, the collaboration had agreed to quote an upper limit. By discovery potential, we mean the cross-section where a convincing case for a signal could be asserted, where the smaller the cross-section the better the discovery potential. In general, the choice of d_{χ^2} that achieves the best sensitivity will be lower (allowing more surface events into the signal region) than the choice of d_{χ^2} that achieves the best discovery potential.

To understand sensitivity and discovery potential, it is helpful to consider a simplified, illustrative example. First we consider the sensitivity. The result of our WIMP search will be observation of N_o candidate nuclear recoils, with detection efficiency ϵ_{nr} in detectors of total mass m during a time interval with livetime τ_L. The N_o candidates are used to deduce an upper limit on the number of WIMP-induced nuclear recoils. Note that prior to unblinding the signal region, our best estimate for N_o is B, the number of surface events we estimate by extrapolation to be in the signal region.

The number of WIMP-induced nuclear recoils, N_i, is related to N_o by the

efficiency of our selection criteria, ϵ_{nr}.

$$N_o = \epsilon_{nr} \times N_i. \tag{4.12.9}$$

The efficiency ϵ_{nr} can be factored into a piece that does not depend on the surface-event rejection cut d_{χ^2} (ϵ_{ot}), and a part that does ($\epsilon_{\chi^2}(d_{\chi^2})$):

$$\epsilon_{nr} = \epsilon_{ot} \times \epsilon_{\chi^2}(d_{\chi^2}). \tag{4.12.10}$$

Figure 4.28 shows how ϵ_{nr} varies with d_{χ^2}.

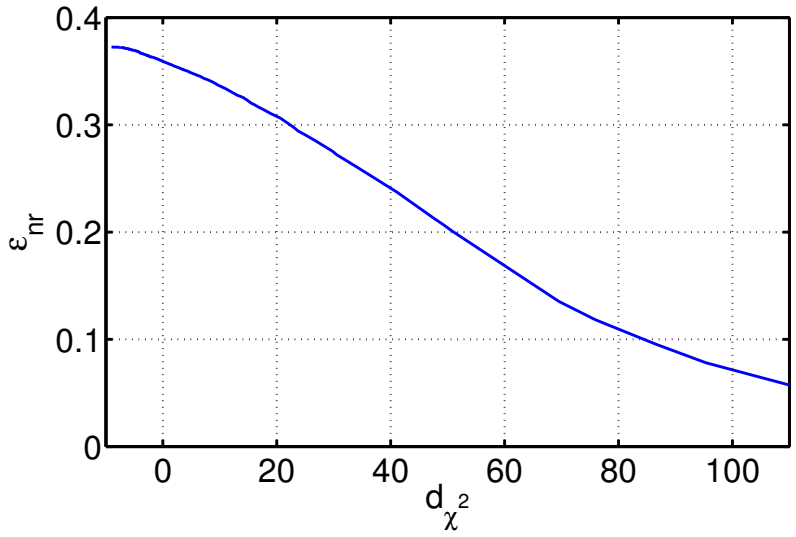

Figure 4.28: Figure showing ϵ_{nr} versus d_{χ^2}. The surface events used have already passed the "outlier" cut.

The typical 90% confidence level upper limit on the number of observed WIMP-induced nuclear recoils, $\bar{N}_o^{90}$, depends upon whether background is subtracted or not:

1. If background is not subtracted, the frequentist estimate is calculated in three steps. First, the 90% confidence upper limit when k events are observed, $N_o^{90}(k)$, is calculated. Second, the Poisson probability distribution is used to estimate the probability of observing k events given an expected background, B. Third, the probabilities are used to calculate a weighted average of the upper limits, $N_o^{90}(k)$. $\bar{N}_o^{90}(B)$ is then given by

$$\bar{N}_o^{90} = \sum_{k=0}^{\infty} \frac{e^{-B}\, B^k}{k!} \times N_o^{90}(k). \tag{4.12.11}$$

   Values for $N_o^{90}(k)$ area given in [149]. In the detailed limit procedure we ultimately use in Section 4.16, background is not subtracted.

2. If background is subtracted, the Feldman-Cousins technique is used to determine $\bar{N}_o^{90}$; the value so obtained is smaller than that obtained without background subtraction.

Figure 4.29 shows how $\bar{N}_o^{90}$ varies with B if background is subtracted and if background is not subtracted.

In either case, $\bar{N}_o^{90}$ is a function of the estimated surface background B, which itself is a function of d_{χ^2}. Figure 4.30 shows how B varies with d_{χ^2}. To make the dependencies explicit, we will use the notation $\bar{N}_o^{90}(B(d_{\chi^2}))$.

Given the value of $\bar{N}_o^{90}(B(d_{\chi^2}))$, we can proceed to the upper limit on the WIMP-nucleon cross section, σ in a few steps. The intrinsic quantity that experimentalists generally use to quantify direct dark matter experiments is the rate per kilogram, R. The typical upper limit (90% confidence level) on R

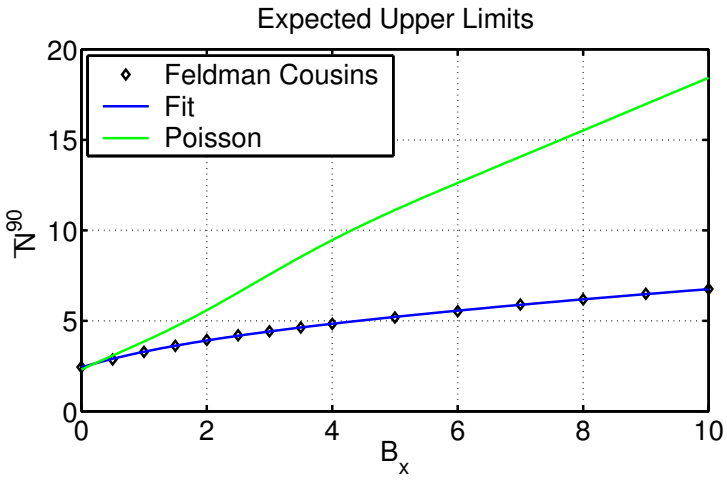

Figure 4.29: Figure showing how the mean 90% confidence upper limit on the number of observed signal events varies with expected background assuming background subtraction. The blue curve shows how the upper limit varies with no background subtraction. The black diamonds are datapoints assuming background subtraction from the paper by Feldman and Cousins [150]. The green line is a fit to the data points indicated by the black diamonds.

is

$$\bar{R}^{90} = \frac{\bar{N}_o^{90}(B(d_{\chi^2}))}{m\tau_L\epsilon_{nr}(d_{\chi^2})}. \tag{4.12.12}$$

To connect to more fundamental quantities, we must convert from rate of interaction per kilogram to rate of interaction per target nucleus. Assuming the atomic weight of the target is A, the number of target nuclei per kilogram, κ, is:

$$\kappa = \frac{10^3 N_A}{A}. \tag{4.12.13}$$

The connection to the WIMP-nucleon interaction cross section follows from the observation that the rate of interaction per target atom is also equal to

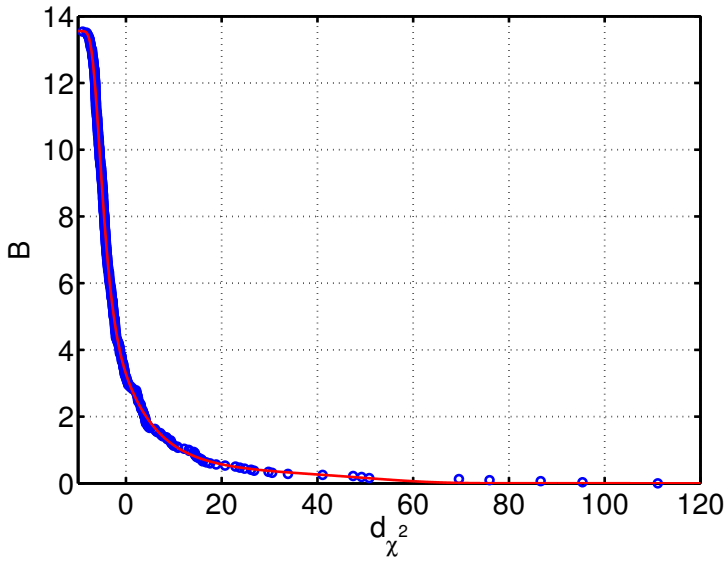

Figure 4.30: Figure showing how B changes as d_{χ^2} is varied. The red curve is a fit to the data.

the product of the galactic flux of WIMPs, ϕ_w, with the interaction cross section between WIMPs and a target nucleus, σ_t. We neglect for this simplified discussion the variation of the flux of WIMPs with initial WIMP velocity, as well as variations caused by the variety of energies of the nuclear recoil. The cross-section per target nucleus is proportional to the cross section per target nucleon: $\sigma_t = T\sigma$, where the coefficient T was described earlier in Section 2.2.1. Putting all the factors together, the limit on the WIMP-nucleon cross section is:

$$\sigma_t = \left(\frac{A}{10^3 N_A \phi_w T}\right) \times \left(\frac{1}{\tau_L m \epsilon_{ot}} \times \frac{\bar{N}_o^{90}(B(d_{\chi^2}))}{\epsilon_{\chi^2}(d_{\chi^2})}\right). \tag{4.12.14}$$

The first term in Equation 4.12.14 is a constant that does not depend on the

N	2	3	4	5	6	7	8
3σ	0.075	0.271	0.564	0.930	1.35	1.82	2.32
4σ	0.011	0.074	0.205	0.402	0.655	0.956	1.30

Table 4.10: List of the maximum allowed background level to obtain a 3σ hint of a signal with N observed events. Maximum allowed background levels are given for to obtain 4σ confidence as well.

choice of d_{χ^2}. The second term is $\bar{R}^{90}$. The best sensitivity is obtained by variation of d_{χ^2} to obtain the minimum value of $\bar{R}^{90}$. Figure 4.31 shows how $\bar{R}^{90}$ varies with B. The value of B that yields the best sensitivity is labeled $B_{\downarrow}$. $B_{\downarrow} = 0.49$ events when no background subtraction is used, and $B_{\downarrow} = 0.55$ events when background subtraction is used.

As can be observed from Fig. 4.31, employing or not employing background subtraction has little affect on the sensitivity as long as the the expected background is near $B_{x\downarrow}$. The optimum sensitivity when no background subtraction is employed is about 7% less than the optimum when background subtraction is employed.

Evaluation of the discovery potential follows a logical thread distinct from that of the evaluation of sensitivity. The first consideration is, 'What is the minimum number N_o of events that could be reasonably accepted as a discovery of nuclear recoils from WIMPs?' We take this minimum as three events, but we also evaluate the discovery potential for four events.

Given the number of events considered convincing, the remaining issue is to cut just hard enough on the surface variable, $\chi_b^2 - \chi_n^2$ to insure that the

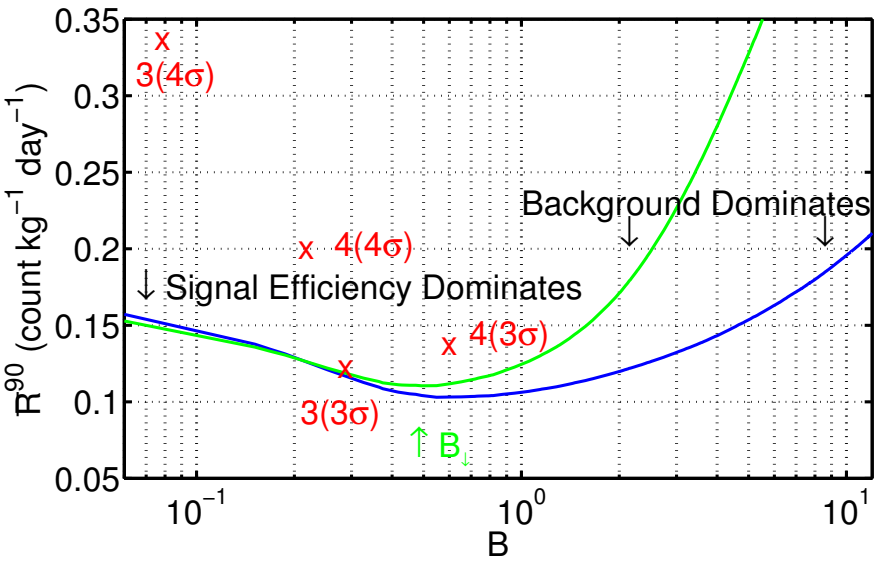

Figure 4.31: Figure showing how $\bar{R}^{90}$ varies with B for this WIMP-search run. The solid blue curve (green curve) shows how $\bar{R}^{90}$ varies with the B without (with) background subtraction. Without (with) background subtraction, the point of greatest sensitivity, $B_{\downarrow}$, occurs at $B = 0.49$ (0.55) events. The best value of $\bar{R}^{90}$ without background subtraction is about 6% worse than the best value with background subtraction. For reference in Section 4.12.7, the red x's show the mean values of R for the maximum levels of B allowed (0.271 events and 0.564 events) to observe a 3σ hint of a signal with 3 or 4 observed events.

probability that the background fluctuated to N_o or more events is sufficiently small. We take that probability to be 2.7 parts in 1000, corresponding to 3 standard deviations (3σ). The values of $B(d_{\chi^2})$ for a variety of N_o are given in Table 4.10. Table 4.10 also gives the values of $B(d_{\chi^2})$ for the same values of N_o when the probability is taken to be 63 parts in 10^6 (4σ).

The discovery potential cross-section is then:

$$\sigma_d = \left(\frac{A}{N_A \phi T}\right) \times \left(\frac{1}{\tau_L m \epsilon_{ot}} \times \frac{N_o}{\epsilon_{\chi^2}(d_{\chi^2})}\right) \tag{4.12.15}$$

where $\epsilon_{\chi^2}(d_{\chi^2})$ is the efficiency at the value of d_{χ^2} that gives the value of $B(d_{\chi^2})$ in Table 4.10, and where the second term is the interaction rate, R, and can be compared to that in Eqn. 4.12.14. In general, the discovery potential cross section will exceed that of optimum sensitivity. The discovery potential interaction rate for 3 and 4 events is shown in Fig. 4.31.

Figure 4.31 shows that the sensitivity varies only slowly near $B_{x\downarrow}$. Any surface event rejection cut with B near $B_{x\downarrow}$ will have near optimum sensitivity. A surface event rejection cut with 3σ discovery potential for 3 events has only about a 10% lower sensitivity than the optimum sensitivity. Therefore, we decided to choose a surface event rejection cut that gave 3σ discovery potential: we decided to choose a surface event rejection cut with discovery potential for 3 observed events and hence require the expected background to be ≤ 0.271 events.

In addition to the expected surface event background, there is the expected neutron background. To maintain discovery potential, the surface event background and the neutron background must together be ≤ 0.271 events. While under the pressure of time, we estimated that the neutron background would be ≈ 0.1 events. The remaining expected background, B^β, is due to surface events. In the next section, we will choose a surface event rejection cut with $B^\beta \leq 0.17$ events.

4.12.7 Choosing a value of B^β

We defined a surface event rejection cut under the pressure of time with the goal of allowing ≤ 0.17 surface events. The actual definition of the cut resulted in $B^\beta \approx 0.15$ events. While we had not yet made Fig. 4.31, we observed that the signal efficiency changed little as the cut was further tightened. Therefore, we decided to define a second, stricter cut with $B^\beta \approx 0.08$ events, initially to allow a systematic test after unblinding the data. Before unblinding, we decided to make the stricter cut our primary cut for subjective reasons.

4.12.8 Calculating the Values of d_{χ^2} for Germanium

We now wish to obtain the values of d_{χ^2} that will result in the desired levels of expected surface event background, B^β, for the looser and stricter cuts. To do so, we must first calculate the number of WBO surface events, L_o, passing the looser (stricter) cut that will result in the desired level of B^β. Using the equation, Eqn. 4.12.6, that was derived for this purpose

$$\begin{aligned} L_o(d_{\chi^2}) &= \frac{B^\beta}{S_1} - 1 \\ &= \frac{0.15\,(0.08)}{0.0066^{+0.0032}_{-0.0023}} - 1 = 22^{+11}_{-8}(11^{+6}_{-4}) \end{aligned} \tag{4.12.16}$$

for the looser (stricter) cut.

The -1 in Eqn. 4.12.16 is a correction for statistical bias derived by Bernard Sadoulet. This factor is significant for small numbers of accepted surface events.

The values of L_o that yield the desired value of B^β for the looser (stricter)

Detector	T1Z4	T1Z6	T2Z1	T2Z2	T2Z4	T2Z6
d_{χ^2}	12	10	3	1	4	17

Table 4.11: List of silicon detector d_{χ^2} values.

surface event rejection cut are 22 (11) WBO events. The corresponding values of d_{χ^2} are 11(18).

4.12.9 Calculating the Values of d_{χ^2} for Silicon

For each silicon detector, we chose a value of d_{χ^2} by eye for two reasons: we were under time pressure to finish defining our surface event rejection cut before testing the cut on the closed gamma calibration data, and at the time we were defining the cut the silicon detectors were viewed as relatively unimportant. Figure 4.32 shows the distribution of WBO surface events and nuclear recoils from neutron calibration data used to define the cut. The value of d_{χ^2} chosen for a particular detector was a compromise between retaining efficiency for accepting nuclear recoils and retaining efficiency for rejecting surface events. Surface events from the "wide beta" region of open gamma calibration data were used, and nuclear recoils from neutron calibration data were used. Table 4.11 lists the value of d_{χ^2} chosen for each silicon detector. Note that because a value of B was not calculated for each of the four silicon detectors, no scale factors were needed for each of the four silicon detectors.

ZIP	Open leakage	Closed leakage	2σ NR Open leakage	2σ NR Closed leakage
T1Z1 (Ge)*	7(7)	5(5)	5(5)	2(2)
T1Z2 (Ge)	7(4)	8(8)	2(2)	0(0)
T1Z3 (Ge)	3(1)	3(2)	0(0)	0(0)
T1Z4 (Si)	6	10	1	6
T1Z5 (Ge)	6(3)	18(12)	0(0)	0(0)
T1Z6 (Si)*	8	18	5	6
T2Z1 (Si)*	7	4	4	3
T2Z2 (Si)	8	10	5	3
T2Z3 (Ge)	2(1)	5(4)	0(0)	0(0)
T2Z4 (Si)	5	10	4	4
T2Z5	4(2)	1(0)	1(1)	1(0)
T2Z6 (Si)	8	3	1	1
All Ge	22(11)	35(26)	3(3)	1(0)
All Si	42	52	20	23

Table 4.12: List of detector "wide beta" events passing the surface event rejection cuts in gamma calibration data. Values in parentheses are for the stricter cut. An asterisk indicates that the detector's WIMP-search data was not used in the analysis. The 'All Ge' and 'All Si' rows include only the detectors used in the analysis.

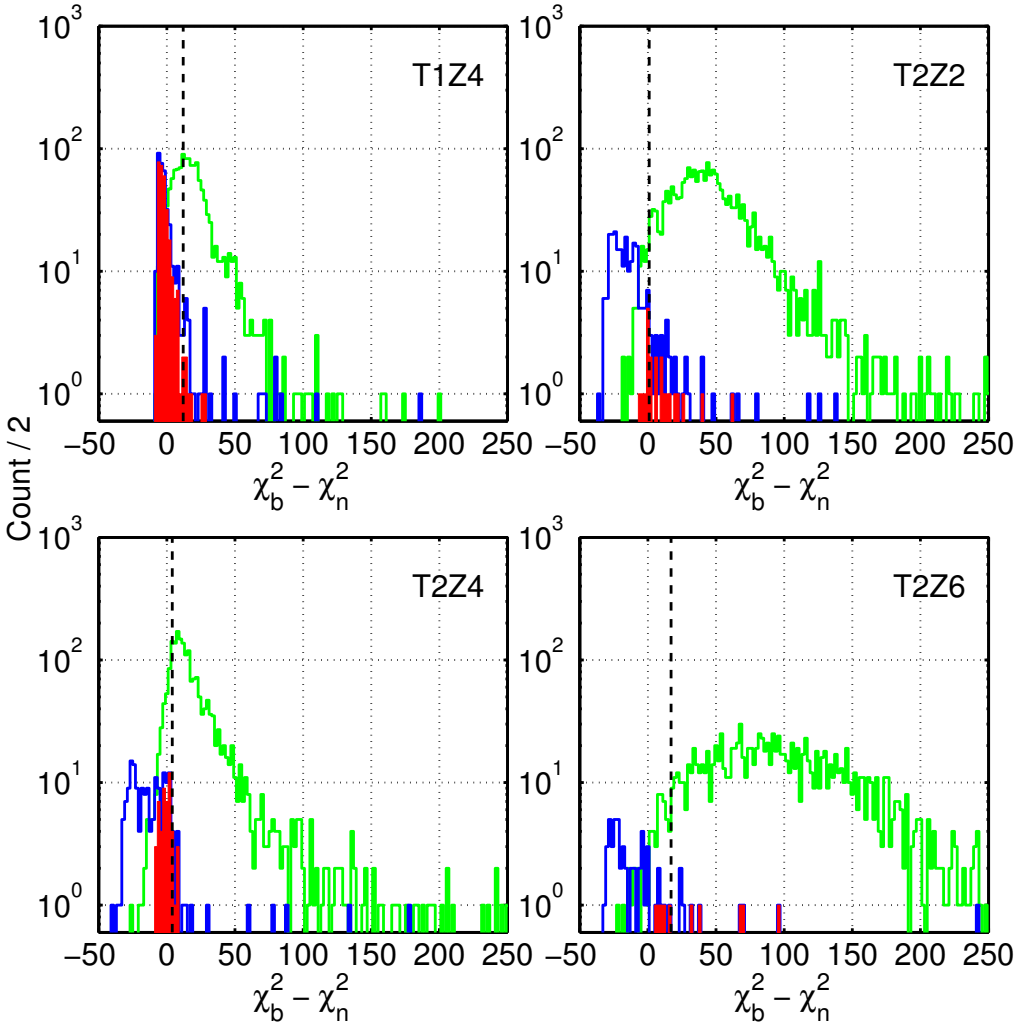

Figure 4.32: Figure showing how the value of d_{χ^2} was chosen for each silicon detector. The blue (green) histogram shows WBO surface events (nuclear recoils from neutron calibration data). The filled red histogram shows WBO surface events that passed the "outlier cut", and thus have sufficiently low χ^2_n. The vertical dashed line denotes the chosen value of d_{χ^2}.

4.12.10 Testing the Background Prediction

Before applying the surface event rejection cut to WIMP-search data and obtaining a result, the CDMS collaboration had agreed to first test the surface event rejection cut on another sample of surface events, the surface events in the "wide beta" region of *closed* gamma calibration data described in Section 4.1. We agreed to use the surface events in closed gamma calibration data in two ways. First, we used them to test if the cut had the expected efficiency for rejecting surface events from a second sample of surface events. This section describes this test. Second, we used them to estimate the WIMP-search surface event background. This is accomplished in Section 4.14.

The expected number of events in the wide beta region of the closed gamma calibration data passing the looser (stricter) cut is:

$$L_c = (L_o + 1) \times \frac{N_c}{N_o} = 25 \pm 5\,(13 \pm 4) \tag{4.12.17}$$

where N_o is the number of WBO surface events, and N_c is the number from closed data. $N_c = 2649$, and $N_o = 2435$.

When the closed gamma calibration data was examined, 35 (26) events in the wide beta region passed the looser (stricter) cut which is 1.9σ (3.6σ) more than expected. Table 4.12 shows the number of events passing the cuts from open and closed gamma calibration data for each detector. Clearly more events passed the surface event rejection cut than were expected particularly for detector T1Z5. The reason for the difference between open and closed gamma

calibration data is uncertain, but is likely due, at least in part, to drift in the experiment.

The primary systematic difference between the open and closed gamma calibration data originates in how the events are split into open and closed samples. As was described in Section 4.1, gamma calibration datasets were typically split such that the first part of most datasets was labeled open and latter part was labeled closed. At the time, it was not fully understood that the time scale on which the gamma calibration data was split ($\sim$2 hours) between open and closed is longer than the time scale ($\sim$1 hour) for the electronic noise environment to change. While a dataset's electronic noise templates [144], which are made from the first 500 events in that dataset, may describe the noise at the beginning of a calibration dataset, they may not be descriptive of the noise environment for later events.

The origin of the primary systematic difference between open and closed gamma calibration data can be verified by resplitting the gamma calibration events in a manner that greatly addresses the impact of drift. Rupak Mahapatra resplit the gamma calibration data into 2 subsets using every other event. He found that the number of surface events that passed the cut in each subset were statistically consistent between the two subsets [143].

We decided to use the closed gamma calibration data to make an unbiased estimate of the expected surface event background. Section 4.14 contains the background estimates including statistical and systematic error. Separate

background estimates are made from closed gamma calibration and side bands of WIMP-search data. After estimating the backgrounds, the blinded region of WIMP-search data will be opened and a result obtained in Section 4.15.

4.12.11 Efficiency

The efficiency of the surface event rejection cut for passing nuclear recoils is estimated from neutron calibration data using events that pass all prior cuts except the single detector hit requirement. The efficiency is fit by the following functional form: efficiency $= C_1 + C_2/(E_r - C_3)^{C_4}$ [122]. Typical values for the fit coefficients are $C_1 = 0.8$, $C_2 =$ -60, $C_3 =$ -6, and $C_4 = 2$. Figure 4.33 shows the fit efficiency of the looser surface event rejection cut for each detector and, Fig. 4.34 shows the fit efficiency of the stricter surface event rejection cut for each detector.

4.13 Combined Efficiencies

Combining efficiencies accurately can be tricky for cuts that are correlated. The majority of cuts in this analysis are uncorrelated. The methods by which the efficiencies of the remaining cuts were calculated were designed to properly account for correlations between cuts. The combined germanium and silicon efficiencies of the analysis are given by multiplying the fit efficiencies of all the cuts discussed in this chapter [122]. Figure 4.35 shows the combined efficiencies

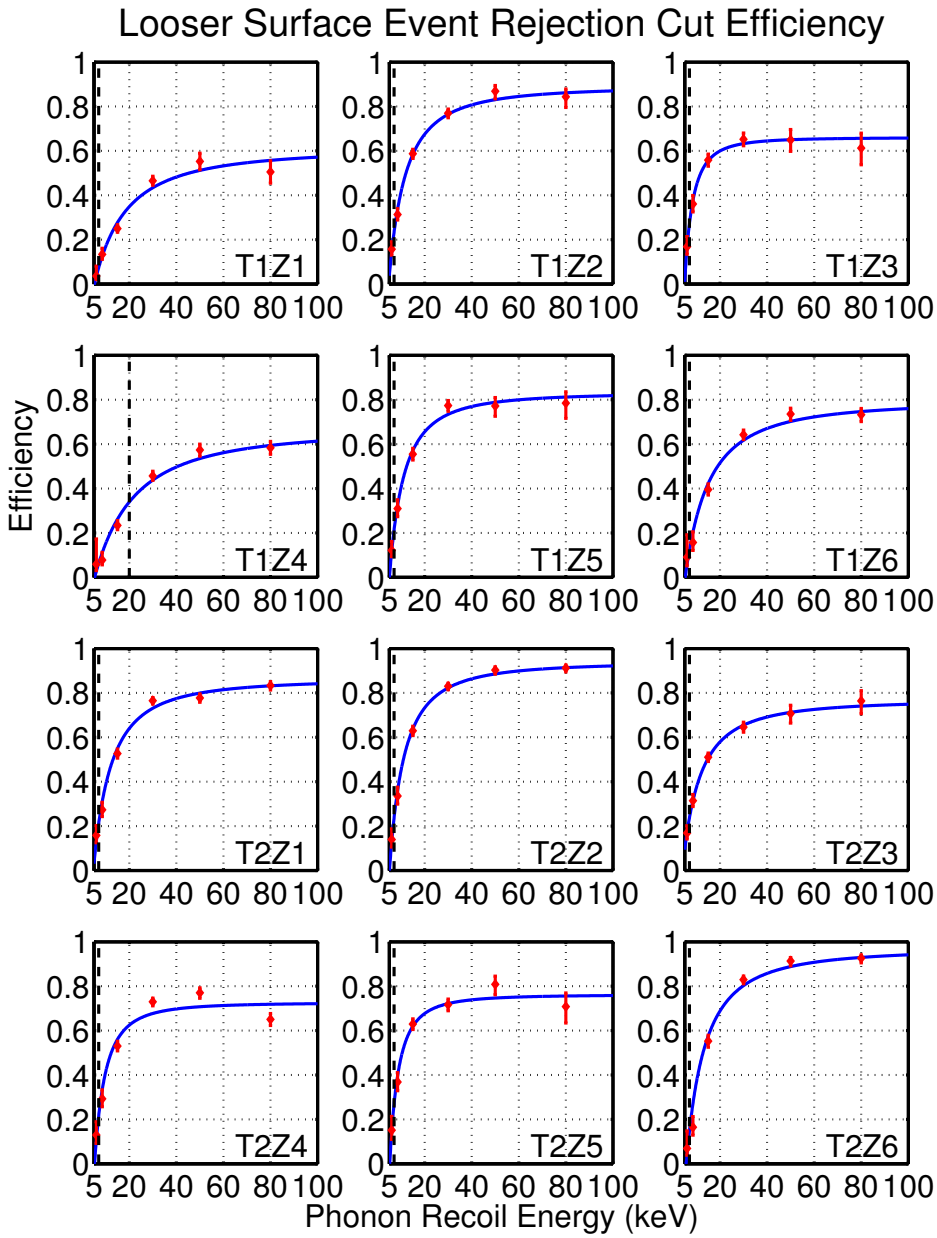

Figure 4.33: The efficiency of the looser χ^2 surface event rejection cut for all twelve detectors. The red data points show the measured cut efficiencies with 1σ error bars. The black, dashed, vertical line shows the energy threshold.

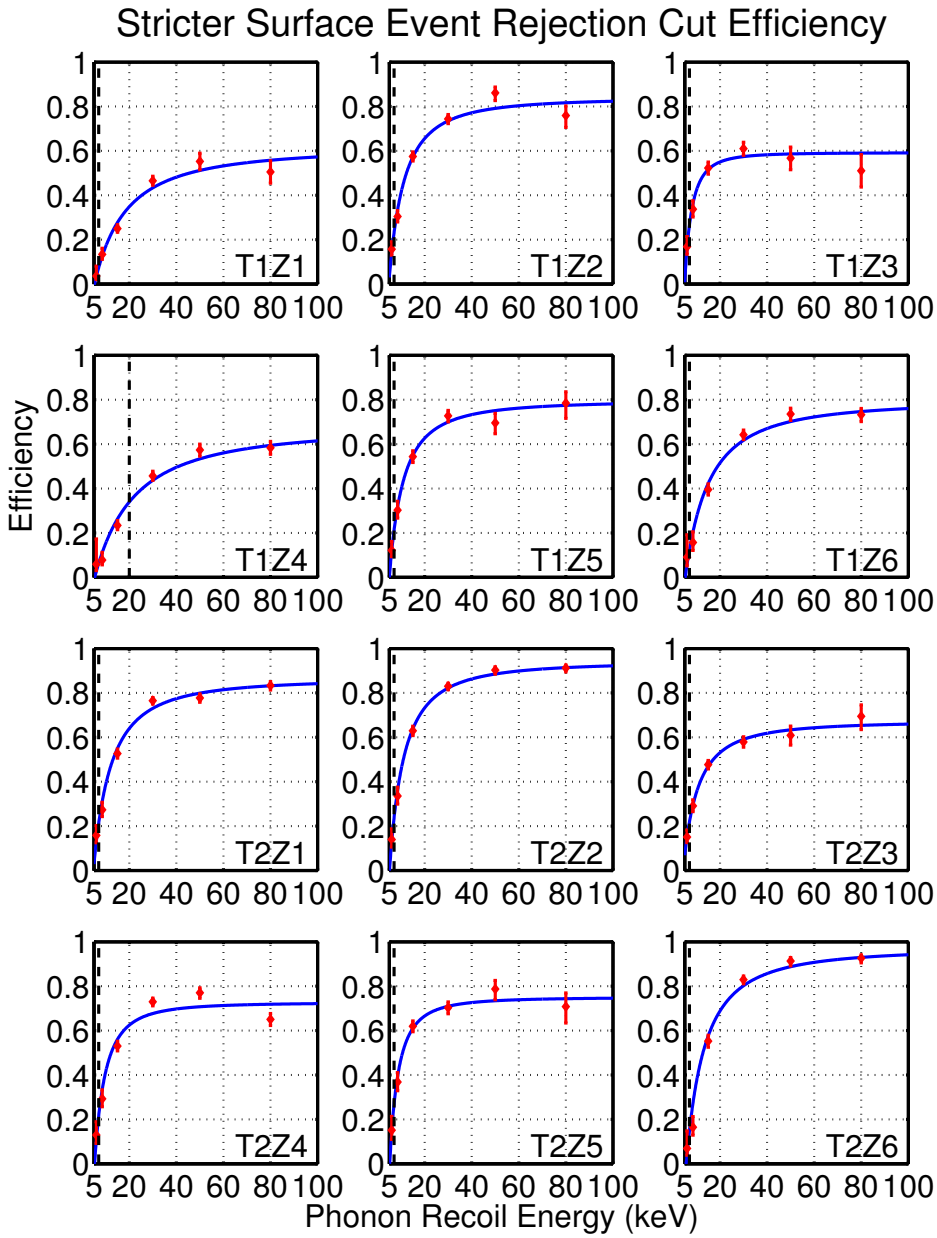

Figure 4.34: The efficiency of the stricter χ^2 surface event rejection cut for all twelve detectors. The red data points show the measured cut efficiencies with 1σ error bars. The black, dashed, vertical line shows the energy threshold.

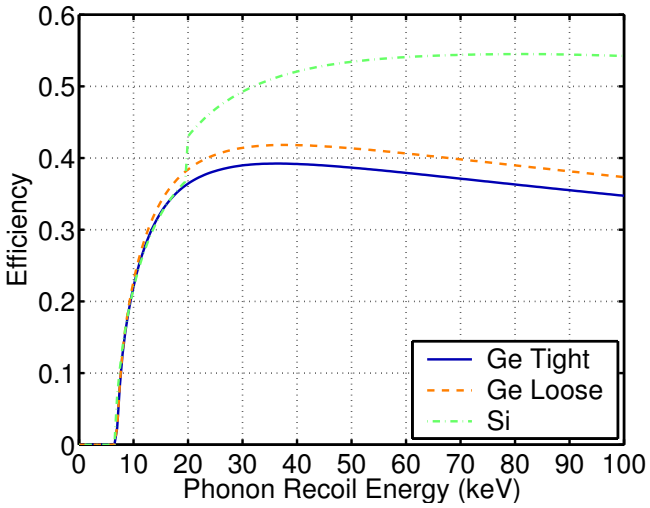

Figure 4.35: The combined efficiencies of the silicon and germanium analyses. The dark blue, solid curve shows the efficiency of the tight germanium analysis, the dashed orange curve for the loose germanium analysis, and the green dot-dashed curve for the silicon analysis. The silicon detector T1Z4 has a 20 keV analysis threshold which leads to the sharp increase in efficiency at a phonon recoil energy of 20 keV.

for the germanium and silicon analyses.

As was discussed in section 2.2.1, nuclear recoils from WIMP interactions deposit energy with an exponential spectrum that is dependent on the WIMP and target masses. The expected number of observed WIMP interactions is

$$N_{\bar{E},i} = M_{s,i}\,\bar{\tau}_{L,i} \int \frac{dR_i}{dE_r}\,\epsilon_i(E_r)\,dE_r \tag{4.13.1}$$

where $E_{\bar{s},1}$ ($E_{\bar{s},2}$) is the germanium (silicon) spectrum-averaged exposure, $M_{s,1}$ ($M_{s,2}$) is the summed mass of the good germanium (silicon) detectors, $\bar{\tau}_{L,1}$ ($\bar{\tau}_{L,2}$) is the mean livetime of the good germanium (silicon) detectors, ϵ_1 (ϵ_2) is the

WIMP Mass (GeV)	10	20	50	100	500
	Expected Number of Interactions				
Silicon	0.09	0.53	0.91	0.74	0.21
Germanium (tight)	0.11	3.64	10.95	9.67	2.74
Germanium (loose)	0.11	3.80	11.55	10.23	2.90

Table 4.13: List of expected number of observed interactions for WIMPs of various masses when assuming standard halo parameters given in Chapter 2 and a cross section of $10^{-42}\,\mathrm{cm}^2$ (normalized to a single nucleon).

recoil energy dependent, combined germanium (silicon) efficiency, and dR_1/dE_r (dR_2/dE_r) is the differential scatter rate in germanium (silicon). Table 4.13 lists the expected number of observed interactions in the CDMS silicon and germanium detectors for WIMPs of various masses with a cross section of $10^{-42}\,\mathrm{cm}^2$ (normalized to a single nucleon); standard halo parameters given in Chapter 2 and the differential scattering rate given by by Eqn. 2.2.25 are assumed to calculate table entries.

The spectrum of recoil energies can also be obtained from Eqn. 4.13.1. The spectrum of recoil energies expected in the CDMS silicon and germanium detectors under the same assumptions is plotted in Fig. 4.36 for a variety of WIMP masses. Changing the WIMP cross section simply scales number events per keV, but does not effect the shape of the spectrum. The number of expected WIMPs observed falls as the WIMP mass is altered away from the target mass. As the WIMP mass falls below about 10 GeV, the silicon detectors begin to have greater sensitivity than the germanium detectors.

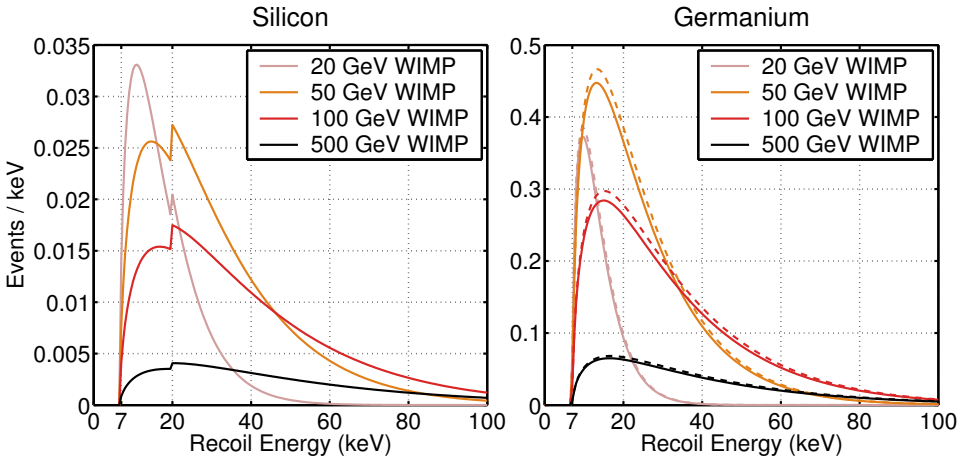

Figure 4.36: The expected spectrum of recoil energies from WIMP interactions. The standard halo parameters given in Chapter 2 and a cross section of $10^{-42}\,\mathrm{cm}^2$ (normalized to a single nucleon) are assumed. The solid (dashed) lines show the spectrum for the stricter (looser) analysis.

The analysis of the data from the silicon detectors has greater sensitivity to very low mass WIMPS while the analysis of the data from the germanium detectors has greater sensitivity to WIMPs with masses above 9 GeV.

4.14 Expected Backgrounds and Systematics

Now that all the cuts in the analysis have been defined and the efficiency of each cut has been determined, it is possible to estimate the background. The background due to surface events and neutrons from cosmogenic activity is estimated.

4.14.1 Surface Event Background

The surface event background is estimated from closed gamma calibration data and also from a side band of WIMP-search data. The estimate from gamma calibration has the advantage of several times higher statistics. Surface events in gamma calibration are almost entirely due to electromagnetic particles incident on the surfaces of a detector and could be systematically different than surface events in WIMP-search data which are due to surface contamination. The estimate from the WIMP-search data has the advantages that the WIMP-search side band surface events are from the same source of surface events as the WIMP-search background surface events, however, the statistics are lower. Both estimates will be described. Estimates of the statistical and systematic errors on the background measurement will also be given.

The sample of surface events used to estimate the background from surface events from closed gamma calibration data is different than the sample of WIMP-search surface events used. For gamma calibration data, all surface events in the "wide beta" region are used. For WIMP-search data, multiple scatters within the "wide beta" region and single scatters outside the 3σ nuclear recoil band but within the "wide beta" region are used.

The technique for estimating the background from surface events is similar for both estimates. The fractions of surface events in the sample that pass the cuts are calculated. We assume that all surface events within the sample have the same probability of passing the cut as single scatter surface events within

the nuclear recoil band. The fractions are then normalized to the WIMP-search signal region: single detector scatters within the nuclear recoil band. We will begin by estimating both normalization factors.

Normalizing Closed Gamma Calibration Surface Events

We deduce an estimate for the number of background events in the signal region of WIMP-search data by extrapolating from the sample of surface events in the "wide beta" region of closed gamma calibration data using a normalization factor S_3. S_3 is estimated using the same method and under the same assumptions as S_1 was calculated in Section 4.12.5.

$$\begin{aligned} S_3 &= R_E \left(\frac{N^p_{2\sigma s}}{N^c_{wba}} \right) \\ &= 2.0 \pm 0.1 \, \frac{8^{+3.9}_{-2.8}}{2615 \pm 51} = 0.0061^{+0.0030}_{-0.0022}. \end{aligned} \tag{4.14.1}$$

where N^c_{wba} is the number of surface events in the "wide beta" region of closed gamma calibration data.

Normalizing WIMP-Search Sideband Events

The number of surface events that were not blinded in the "wide beta" region of WIMP-search data are normalized to the number of single scatter events in the nuclear recoil band under the same assumptions and using a similar method to how S_1 was calculated in Section 4.12.5. In order to have a fair sample, we must remove events that would have been blinded from the

"wide beta" region of the prior WIMP-search run. To do this, we remove the single scatters that fall within the 3σ nuclear recoil. Of the 308 "wide beta" events in the prior WIMP-search run, 291 would not have been blinded. The normalization factor, S_4 is then

$$S_4 = \frac{8^{+3.9}_{-2.8}}{291 \pm 17} = 0.027^{+0.013}_{-0.010}. \tag{4.14.2}$$

Estimating from Closed Gamma Calibration Surface Events

Now that the normalization factors have been obtained, an estimate of the surface event background using closed gamma calibration data can be made. The estimated surface event background is calculated using the normalization factor given in Eqn. 4.14.1

$$B_{\gamma,i} = S_3 \, N_{\gamma,i} \tag{4.14.3}$$

where i is 1 (2) for the looser (stricter) cut. The number of events passing the looser (stricter) cut is $N_{\gamma,1} = 35$ events ($N_{\gamma,2} = 26$ events). The expected background is: $B_{\gamma,1} = 0.21^{+0.11}_{-0.09}$(stat) events for the looser cut, and $B_{\gamma,2} = 0.16^{+0.08}_{-0.07}$(stat) events for the stricter cut.

Estimating from WIMP-search Sidebands

Next, we estimate the surface event background using WIMP-search data. In WIMP-search data, there were 704 events in the "wide beta" region of the germanium detectors that were not blinded as potential signal events. Of these events, the 29 (16) events passed the looser (stricter) cut. Equation 4.14.4

Cut	7 - 10 keV	10 - 100 keV
Ge Stricter	0.7±0.1	0.4±0.2(stat)±0.4(sys)
Ge Looser	1.1±0.2	0.8±0.4(stat)±0.6(sys)
Si	0.9±0.3	0.7±0.2(stat)±0.2(sys)

Table 4.14: List of expected surface event backgrounds.

shows how the surface event background is estimated from WIMP-search data.

$$B_{w,i} = S_4 \, N_{w,i} \tag{4.14.4}$$

where i is 1 (2) for the looser (stricter) germanium analysis. The resulting expected background is given by Table 4.14. The systematic error between 7 and 10 keV is dominated by the systematic error above 10 keV and can be neglected.

For the silicon detectors, there is only one cut. There are 74 events in the side band between 7 and 10 keV and 517 more events between 10 and 100 keV that were not blinded as potential signal events. Using the same technique, the expected background is calculated and given in Table 4.14.

Systematic Uncertainties

The systematic uncertainties were estimated using the simple approach of varying the cuts within reasonable bounds and studying the resulting change in expected background. The cuts that were varied were the cuts to which the background was most sensitive. The inner electrode cut, the nuclear recoil

band cut, and the "wide beta" region were varied. The means of the surface event discrimination parameters were found to be different between surface events in gamma calibration data and WIMP-search data. The means of the surface event discrimination parameters were varied to explore the effect of this difference between gamma calibration surface events and WIMP-search surface events. The resulting systematic errors were conservatively combined by linear addition. Note that the systematic error between 7 and 10 keV is dominated by the systematic error above 10 keV and is be neglected.

Comparing the Background Estimates

The expected backgrounds calculated from closed gamma calibration data and from WIMP-search sidebands differ significantly. The background calculation using WIMP-search data is taken as the more accurate estimate because the surface events are from the same source and would therefore be expected to have the similar fractions of events resulting from surface contamination as surface events within the signal region.

After unblinding it was learned that surface events from gamma calibration data have a different distribution within the detectors than the surface events from WIMP-search data. Because the source is above the detectors, many surface events from gamma calibration data are electrons from gamma interactions in the bottom of one detector freeing an electron which scatters in the top surface, or phonon collection side, of the adjacent lower detector. This

results in about two thirds of surface events in gamma calibration being on the phonon collection side of the detector [145]. Whereas most surface events in WIMP-search data are expected to be uniformly distributed. These surface events result from electrons that are freed by ambient gammas and from ^{210}Pb implanted into detector surfaces by decays of airborne radon. Both sources of surface events are expected to be isotropic and distribute surface events evenly over the surface of the detectors.

Obtaining samples of surface events occurring on the phonon and ionization sides of the detectors allows us to compare our ability to reject surface events on the two sides. By examining surface events in which energy is deposited in two adjacent detectors, events can be segregated into those occurring on the top or bottom of a detector. Using these samples we found that surface events on the phonon side were more easily rejected by a factor of 5 [145].

4.14.2 Neutron Background

Before unblinding, an estimate of the cosmogenically induced neutron background was made. To estimate the neutron background, one needs to know the rate of neutrons that are not rejected by the active veto and scatter in a single detector. The neutron background can be estimated from *in situ* data and from simulation. Estimating the neutron background from data is not particularly effective because the expected rate of veto-anticoincident nuclear recoils is negligibly small. Further the larger sideband of veto-coincident neutron scatters

is also small.

We do look at the veto-coincident data for nuclear recoils. No single scatters were found. A single veto-coincident multiple scatter event was observed where one detector had a nuclear recoil. Another detector had activity that was not consistent with a nuclear recoil. This type of event can result when a muon produces a neutron as well as a shower of particles.

This event was not used to estimate the neutron background for two reasons. First, the scaling between veto-coincident nuclear recoil-other multiple scatters and veto-coincident nuclear recoil single scatters is uncertain. Second, the scaling between veto-coincident scatters and veto-anticoincident scatters was also uncertain.

To estimate neutron background, I use the results of an earlier simulation. The method, by which the estimated neutron backgrounds for germanium, $B_{n,1}$, and silicon, $B_{s,2}$, were calculated, is given by Eqn. 4.14.5.

$$B_{n,i} = R_{ns,i}\tau_{L,i}M_i(1-\epsilon_v)\bar{\epsilon}_{a,i} \tag{4.14.5}$$

where $R_{ns,1}$ ($R_{ns,2}$) is the rate of single scatter nuclear recoils in germanium (silicon), ϵ_v is the efficiency of the muon veto for rejecting neutrons, and $\bar{\epsilon}_{a,1}$ ($\bar{\epsilon}_{a,2}$) is the mean efficiency of the analysis for germanium (silicon). For the earlier 1-tower run at Soudan, Sharmila Kamat had estimated that the rate of unvetoed neutron single scatters with recoil energies between 5-100 keV was 0.0053±0.0001 neutrons/kg/day in germanium and 0.0112±0.0001 neutrons/kg/day in silicon [146]. To be conservative, one can assume that the rate

of neutron induced nuclear recoils from 7-100 keV to be the same that from 5-100 keV. The germanium exposure is 93.2 kg-days and the silicon exposure is 29.8 kg-days.

Recent simulations [101] [147] using GEANT4 and FLUKA have indicated that significantly more than 70% of external neutrons are accompanied a muon or proton that deposits sufficent energy in the veto to allow the neutron to be rejected. To be conservative, $\epsilon_v \equiv 0.70$.

Once the mean efficiencies of the analyses are provided, the neutron background can be calculated. $\epsilon_{a,1} = 0.360$ (.384) for the stricter (looser) analysis, and $\epsilon_{a,2} = 0.491$. Using Eqn. 4.14.5, the estimated, external, unvetoed neutron background is 0.053±0.002 (0.057±0.002) events for stricter (looser) germanium analysis and 0.049±0.001 for the silicon analysis.

4.15 The Result

After all data had been obtained, all cuts had been defined, and all steps in the background estimate had been completed save the normalization, which was taken from the WIMP-search data as was predetermined before unblinding, the WIMP-search data was finally unblinded and the result obtained. No events consistent with WIMP-interactions were observed in any of the germanium detectors. One event in T2Z5 passed the surface event rejection cut, but fell just below the 2σ nuclear recoil band. One potential WIMP-interaction was observed in a silicon detector, T2Z6 which is consistent with the expected

background of 1.6 events from Table 4.14.

The results are customarily shown in yield plots. These are given for T1Z2 to T2Z6 in the upper panes of Fig. 4.37 to Fig. 4.45. Events passing all cuts except the surface event rejection cut and nuclear recoil band cut are denoted by blue dots, events passing the looser (stricter) surface event rejection that are outside the nuclear recoil band are denoted by small (large) green circles. The event that was observed in the silicon detector, T2Z6, is denoted by a red x. The black, dashed, vertical line shows the analysis threshold. A 10 keV ^{68}Ga line from gallium in the germanium detectors is visible.

The distribution of the surface event cut discrimination parameter values, $\chi_b^2 - \chi_n^2$, is shown for T1Z2 to T1Z6 in the center and lower panes of Fig. 4.37 to Fig. 4.45. The figures showing the distribution of $\chi_b^2 - \chi_n^2$ values are described as follows. Center pane: The blue histogram shows events from the wide beta region of WIMP-search data. The light purple histogram shows events that passed the looser cut and are circled in the yield plot. Light purple events to the right of the dashed line also pass the stricter cut. The red histogram in the center pane of Fig. 4.45 shows the single scatter event in the 2σ nuclear recoil band of WIMP-search data that passed the surface event rejection cut. Lower pane: the green (blue) histogram shows the $\chi_b^2 - \chi_n^2$ distribution of nuclear recoils ("wide beta" surface events) from neutron calibration data (closed gamma calibration data). The dot-dashed (dashed) line shows the value of d_{χ^2} for the looser (stricter) cut. Lower pane: The blue (light purple) histogram

shows events from the wide beta region of WIMP-search data (that also pass the looser surface event rejection cut). Light purple events to the right of the dashed line also pass the stricter cut.

The only germanium detector to have a near miss in the yield-phonon recoil energy plane is T2Z5. It is interesting to ask what the surface event discrimination parameter, $\chi_b^2 - \chi_n^2$, was for the near miss in this germanium detector: was it close the the surface event discrimination cut value and similar to the distribution of surface events in gamma calibration data? The distribution of surface events and nuclear recoils from calibration data for T2Z5 can be see in lower pane of Fig. 4.44. In the center pane of Fig. 4.44 the near miss event is denoted by a solid black histogram. The near miss event had a surface event discrimination parameter value, $\chi_b^2 - \chi_n^2$, (105) that lies on the tail of surface events from closed gamma calibration data, but is more consistent with that of a nuclear recoil than a surface event. The event was typical of that of nuclear recoils ($\chi_n^2 = 2$). While interesting, the event does fall outside the 2σ nuclear recoil band and would not be inconsistent with the expected backgrounds had it fallen within the nuclear recoil band.

It is also interesting to examine the potential signal event in the silicon detector T2Z6 as well as the semi-near miss in the silicon detector T2Z2. As can be seen from Fig. 4.45, the surface event rejection parameter for the potential signal event in T2Z6 falls near the distribution of surface events from closed gamma calibration data. Four other simultaneous surface event rejection

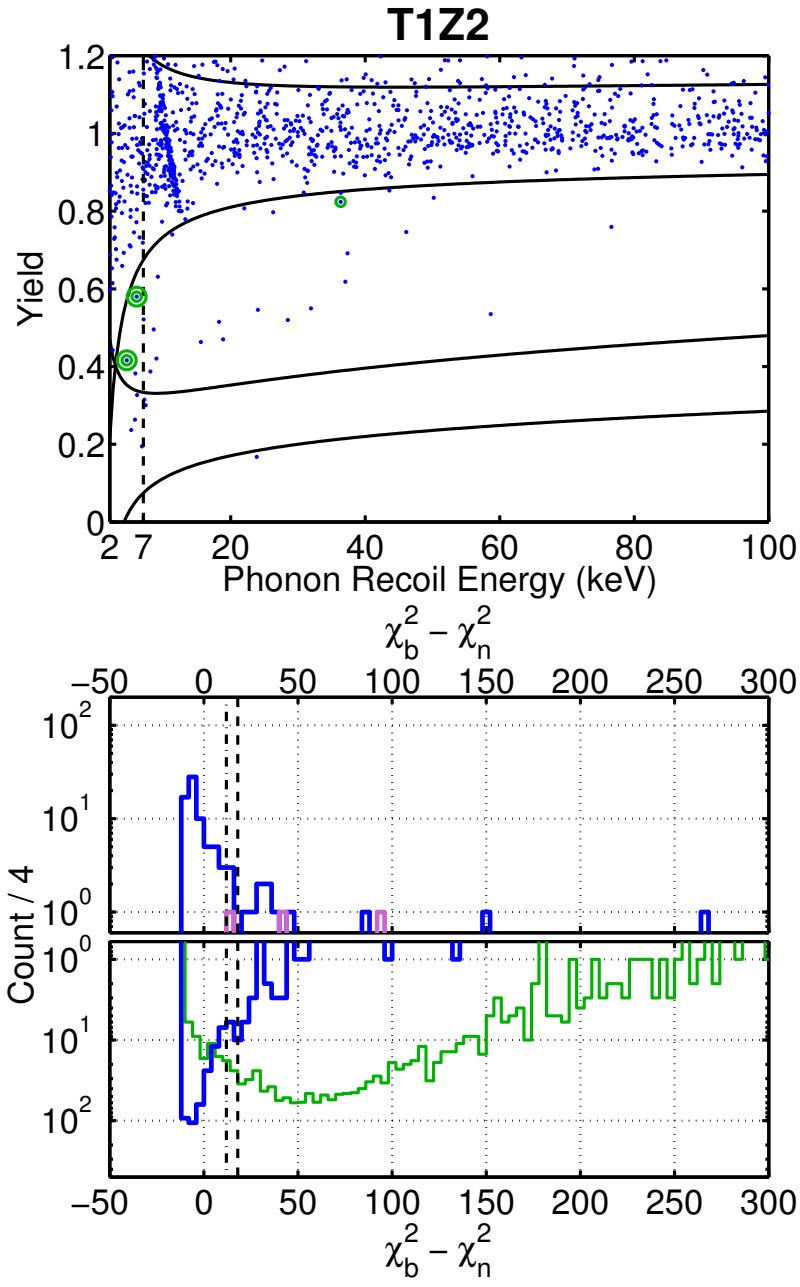

Figure 4.37: Figure showing the result of the analysis after unblinding T1Z2. The figure is described in the text of this section.

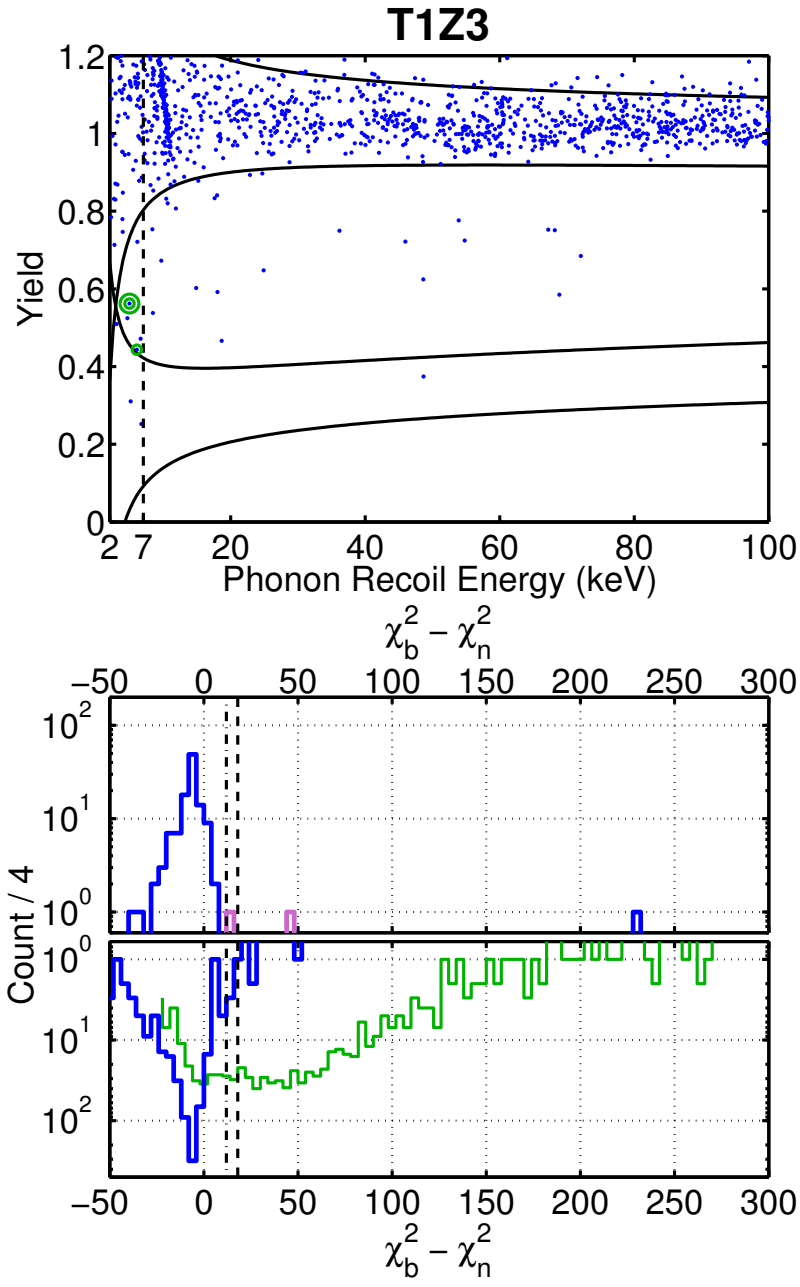

Figure 4.38: Figure showing the result of the analysis after unblinding T1Z3. The figure is described in the text of this section.

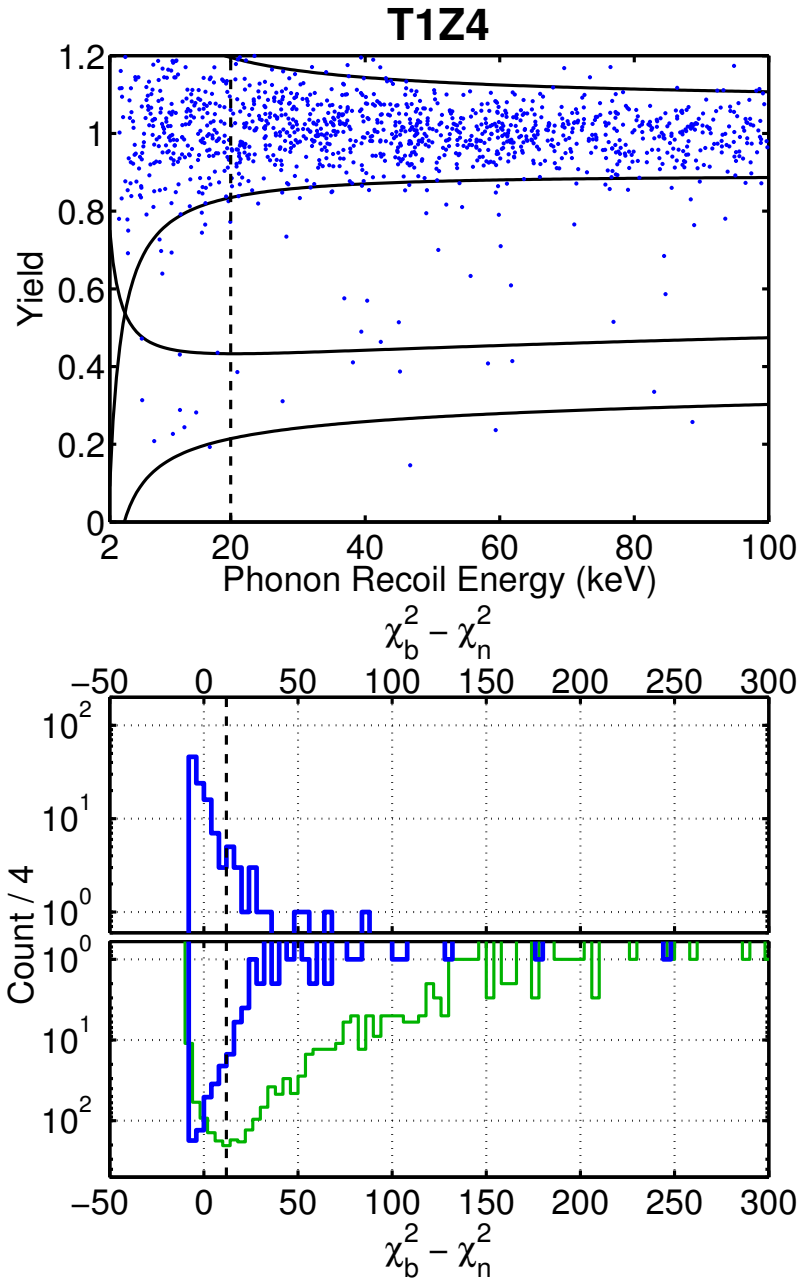

Figure 4.39: Figure showing the result of the analysis after unblinding T1Z4. The figure is described in the text of this section.

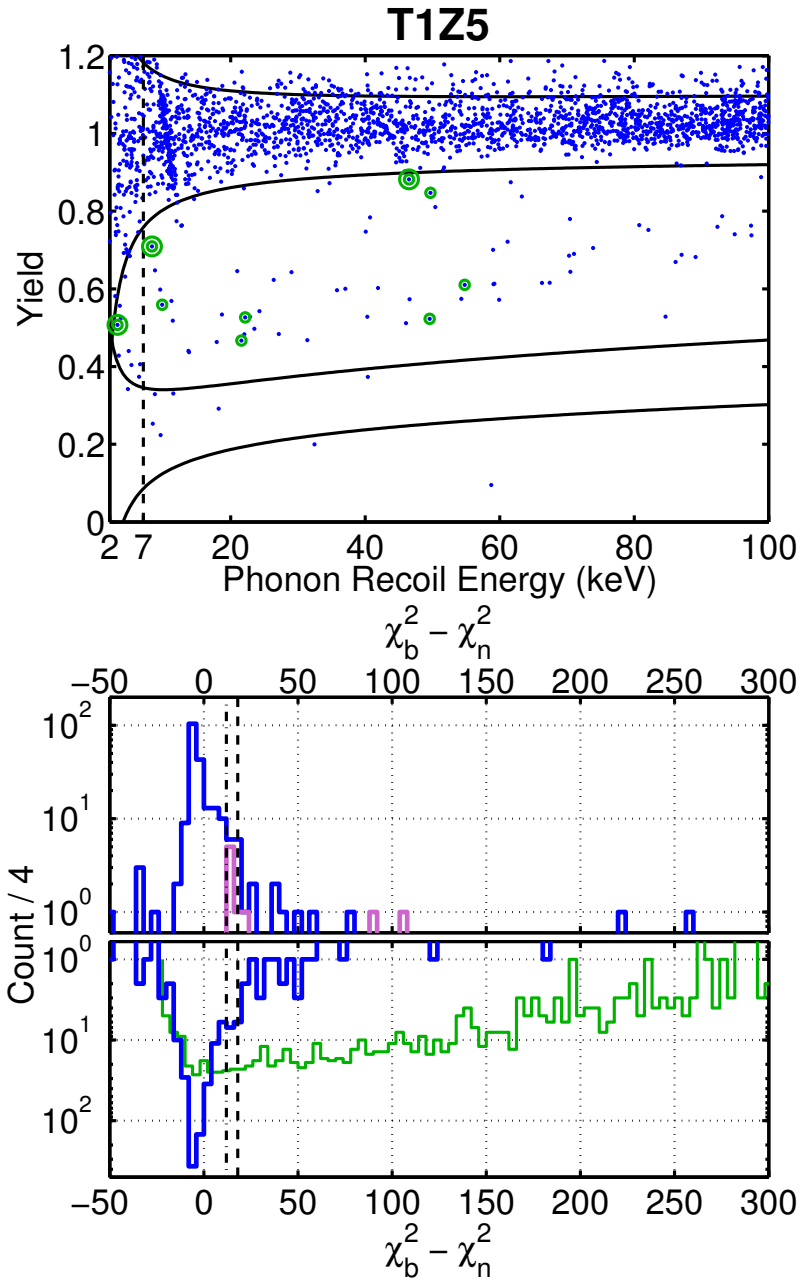

Figure 4.40: Figure showing the result of the analysis after unblinding T1Z5. The figure is described in the text of this section.

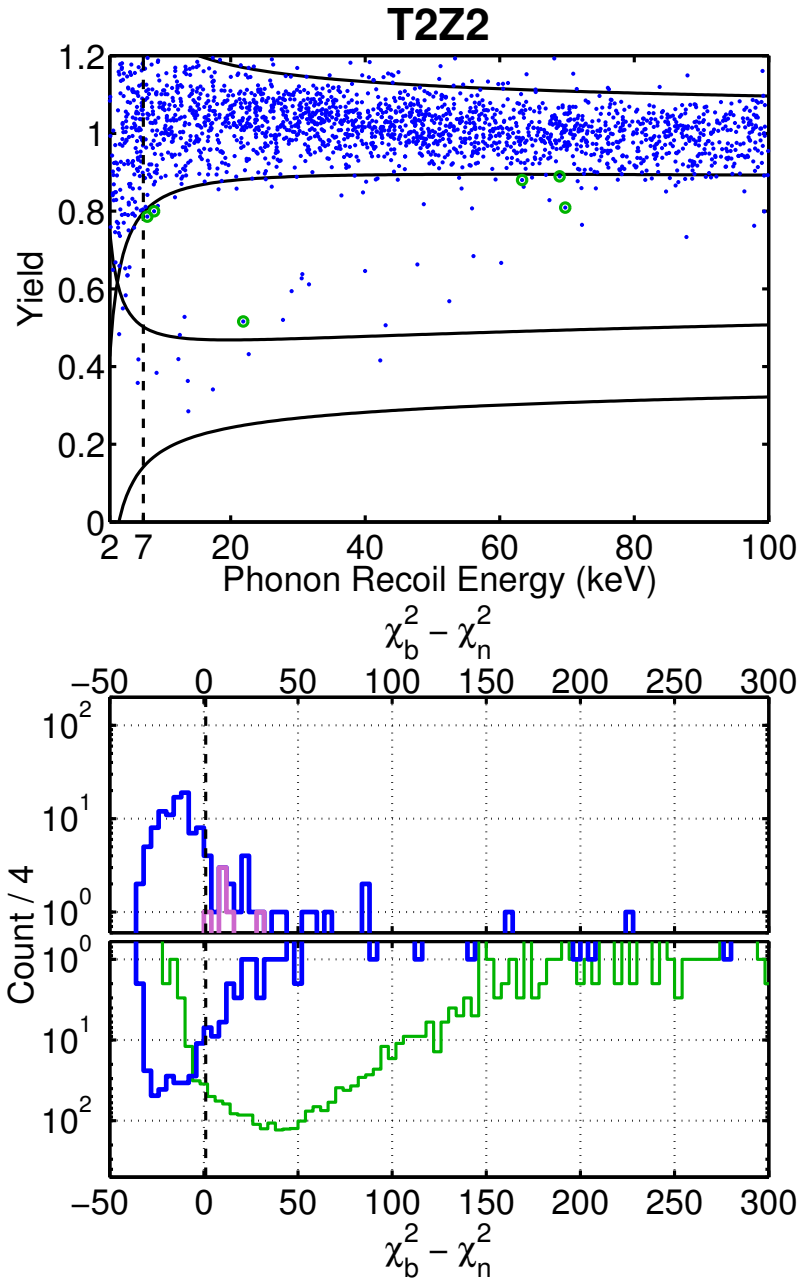

Figure 4.41: Figure showing the result of the analysis after unblinding T2Z2. The figure is described in the text of this section.

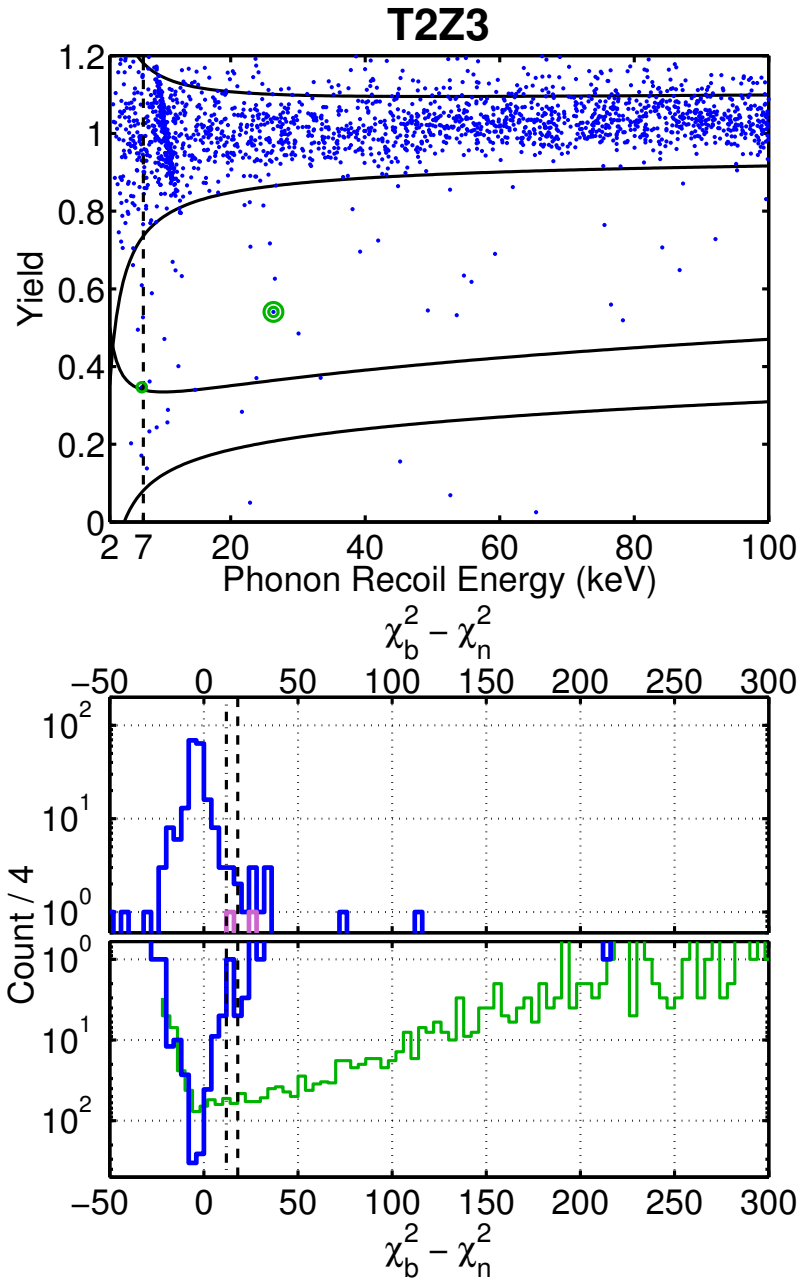

Figure 4.42: Figure showing the result of the analysis after unblinding T2Z3. The figure is described in the text of this section.

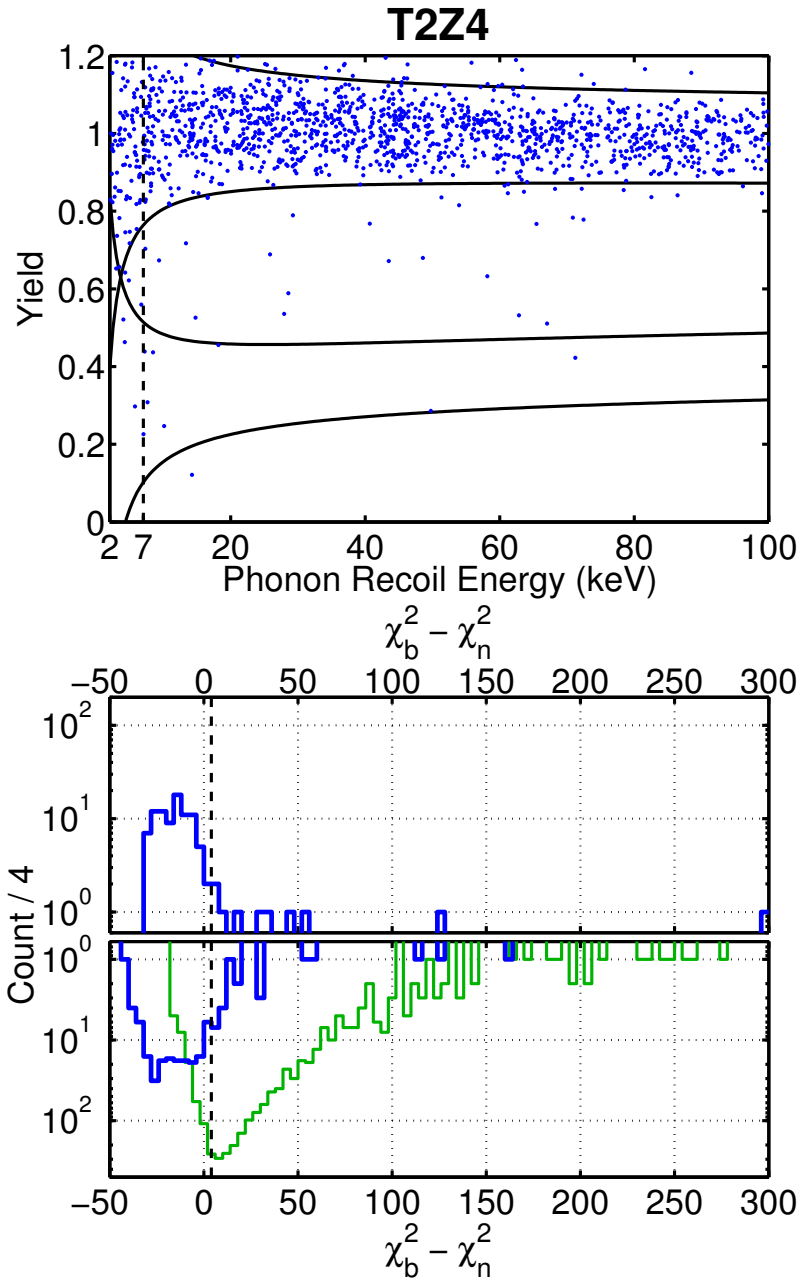

Figure 4.43: Figure showing the result of the analysis after unblinding T2Z4. The figure is described in the text of this section.

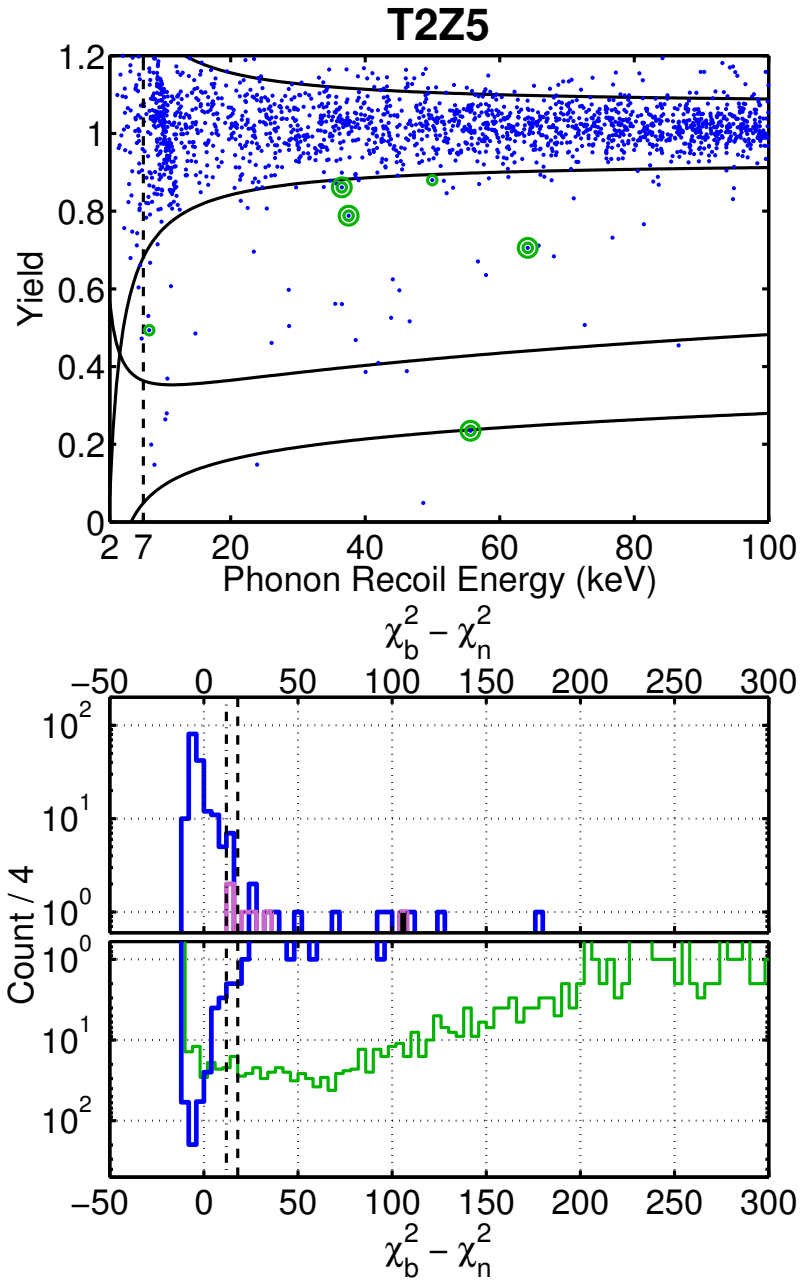

Figure 4.44: Figure showing the result of the analysis after unblinding T2Z5. The figure is described in the text of this section.

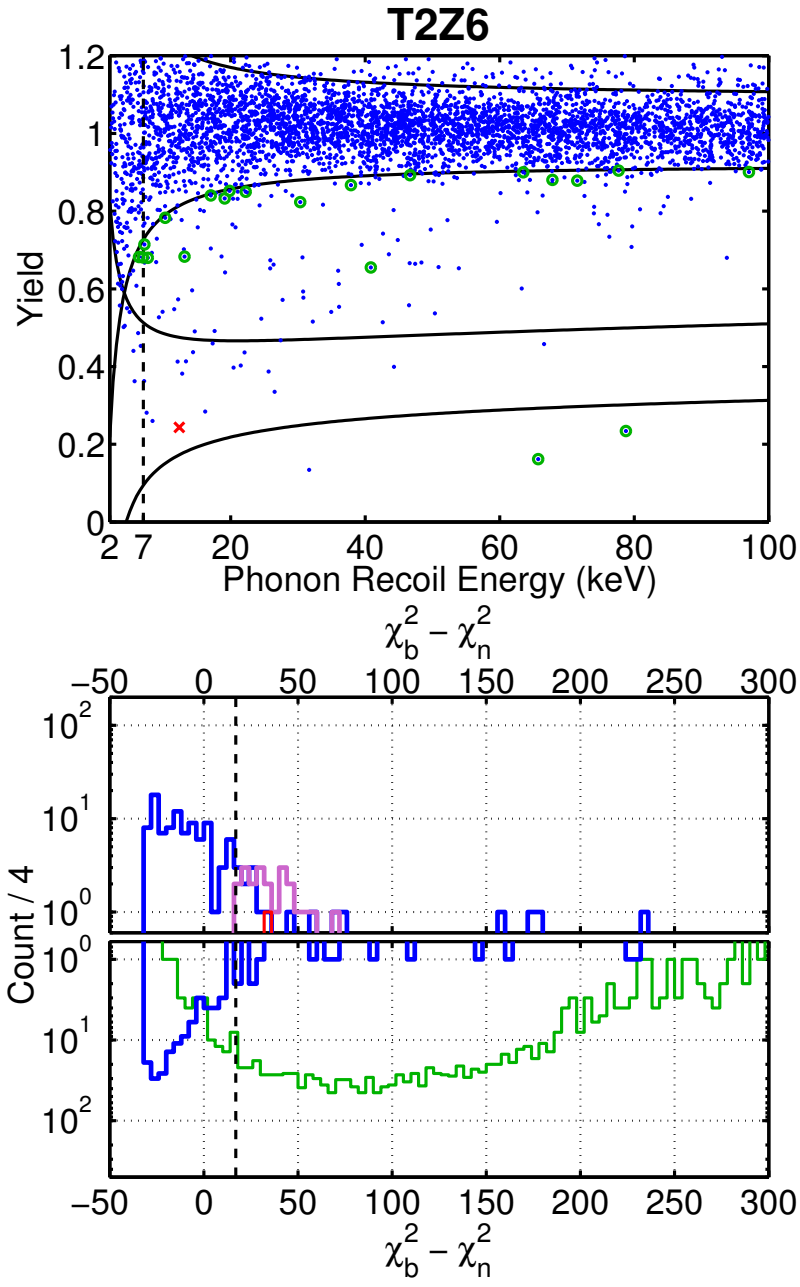

Figure 4.45: Figure showing the result of the analysis after unblinding T2Z6. The figure is described in the text of this section.

cuts were defined using the same calibration datasets by other members of the CDMS collaboration. None one of the other surface event rejection cuts passed this event. The event in T2Z6 is likely a surface event that leaked past the surface event rejection cut. An event that could almost be be described as a near miss in yield-phonon recoil energy plane, see Fig. 4.41, in the detector T2Z2 has a surface event discrimination parameter value of 11.2 which is consistent with the surface event discrimination parameter values of surface events from closed gamma calibration data; this event is likely a surface event. Therefore it is concluded that no WIMP-signal was observed.

In the analysis of the germanium detectors described in the published result [141], one event in T2Z5 passed all cuts. This event did not pass the surface event rejection cut in this analysis. While it had a surface event discrimination parameter value (100) that was not typical of surface events, the event was not typical of nuclear recoils ($\chi^2_n = 16$) and was rejected by the surface event background "outlier cut". This event also occurred during a period in which T2Z5 suffered from poor ionization collection.

4.16 Setting A Limit

When few or no events are observed, one can establish a lower limit on the WIMP cross section for a given WIMP mass under the halo assumptions given in Chapter 2. To establish the lower limit, one asks what the maximum number of events from Poisson statistics that could have been observed at 90%

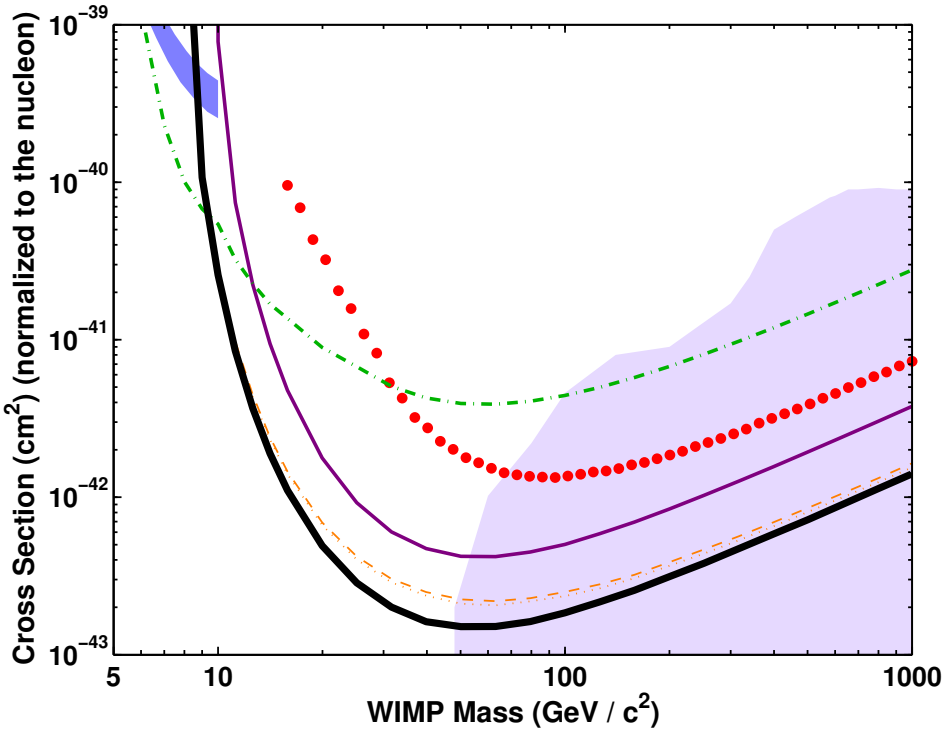

Figure 4.46: Figure showing the resulting limit curves. The thin, dashed, orange line (thin, dotted, orange line) shows the resulting limit from the stricter (looser) germanium analysis as is the resulting limit from the silicon analysis (dot-dashed, green line). The combined result (thick, solid, black line) of the stricter germanium analysis and the last run at Soudan (solid, purple line) is also shown. For comparison the most recent limit from EDELWEISS [114] is shown (large, red circles). Two theoretical regions are also shown. 1) The 90% confidence DAMA allowed region from Gondolo and Gelmini [115] is shown (blue, upper left region) and 2) a prediction from a MSSM model [116].

confidence, N, given the number that were actually observed, n. Then for a given WIMP mass, one calculates the minimum WIMP cross that would be expected to produce N observed events. The process is repeated for a range of WIMP masses to obtain a limit on WIMP cross section as a function of WIMP mass. When there is the presence of significant well-understood background, one can attempt to subtract the expected background from the observed events. This analysis does not incorporate any background subtraction.

Since no significant WIMP-signal was observed, the combined efficiency from Section 4.13, the mean livetime from Section 4.3, the expected WIMP recoil spectrum from Section 2.2.1 and the single event observed in a silicon detector are used to set limits using germanium and silicon data on the WIMP cross section as a function of WIMP mass. Figure 4.46 shows the resulting limit from the germanium detectors (dashed, orange line) as well as the limit from the silicon detectors (dot-dashed green line). The results from the prior run at Soudan can be combined with the results of the stricter germanium analysis to find a "combined result". This is done in four steps:

- Potential signal events from either run are considered to be potential signal events of the "combined run".

- The efficiency and exposure of each run before cuts are applied are obtained.

- The combined efficiency is obtained by taking a weighted average of the

the two run's efficiencies. The exposure of each run before cuts are applied is used for the weighting.

- The exposure of the combined run is obtained by summing the exposure of the two runs.

Combining the germanium limit with the limit from the prior run at Soudan (solid purple line) results in the combined, best germanium limit (thick, solid, black line). Under the assumption of WIMPs interacting via a spin-independent interaction, the current CDMS limit is clearly the best published limit in the world over a large range of WIMP masses and rules out considerable new theoretical parameter space.

WIMP candidates arise from a number of theoretical extensions to the standard model such as the minimal supersymmetric standard model (MSSM) and minimal supergravity models (mSUGRA). A given model can allow a region of WIMP masses and cross sections. Comparison between the results described here and that expected from a given model can be interesting. The silicon result excludes a portion of a region based on an mSUGRA model [115] (dark blue region in the upper left-hand corner of Fig. 4.46) that was designed to reconcile the observed annual modulation observed by the DAMA experiment with the null results of other experiments. The germanium result excludes a portion of a regions of one such MSSM extension by Kim *et al.* [116] (large purple area in the lower right-hand portion of Fig. 4.46).

4.17 A New Neutron Background Estimate

In Section 4.14.2, the pre-unblinding estimate of the neutron background was discussed. After unblinding the WIMP-search data, a better estimate of the neutron background was be made. The new estimate is similar to the prior estimate with three improvements. 1) A neutron background estimate between from 7-100 keV is used instead of from 5-100 keV. Since the neutron background is exponential, removal of the 5-7 keV region is significant. 2) Detectors inside the tower "sandwich" have lower neutron-induced single scatter rates. 3) The dependence of the efficiency on recoil energy is fully accounted for. This is import because the efficiency is lower at low recoil energy while the neutron background falls exponentially with recoil energy. The pre-unblinding estimate assumed a flat efficiency which leads to an overestimate of the neutron background.

The new estimate is made in several steps:

- I obtain an estimate of the neutron-induced recoil spectrum using the plots from Sharmila Kamat's thesis [148] of the recoil spectra from 10-100 keV for Tower 1 detectors. Tower 2 detectors were not included in the simulation.

- Because detectors on the top or bottom of a tower (outer detectors) have a higher rate of neutron-induced recoils than the other detectors (inner detectors), I divide the detectors into three classes: inner germanium, in-

ner silicon, and outer silicon. (The only outer germanium detector, T1Z1, was excluded from this analysis.) A Tower 2 detector of a particular class is assumed to have the same recoil spectrum as a Tower 1 detector of the same class.

- The estimates of the neutron-induced recoil spectrum are used to defined a template recoil spectrum for each class.

- Each detector's recoil spectrum is combined with the detector's analysis efficiency:

$$B_{n,i} = M_{s,i}\,\bar{\tau}_{L,i}(1-\epsilon_v)\sum_j \int S_{n,i,j}\,\epsilon_{i,j}(E_r)\,dE_r \qquad (4.17.1)$$

where $B_{n,1}$ ($B_{n,2}$) is the expected neutron background in the germanium (silicon) detectors, $S_{n,1,j}$ ($S_{n,2,j}$) is the recoil spectrum of the jth good germanium (silicon) detector, $\epsilon_{1,j}$ ($\epsilon_{2,j}$) is the efficiency of the jth good germanium (silicon) detector, and ϵ_v is the veto efficiency. $M_{s,1}$ ($M_{s,2}$) and $\bar{\tau}_{L,1}$ ($\bar{\tau}_{L,2}$) are defined for the germanium (silicon) detectors as in Eqn. 4.13.1.

Sharmila Kamat's thesis [148] gives plots of the recoil spectra from 10-100 keV for Tower 1 detectors that are a more recent than those used in section 4.14.2. A least squares fit to the exponential recoil spectrum be made and extrapolated to lower recoil energy thresholds. The silicon detectors are well fit by a single exponential, while the germanium detectors prefer a double exponential.

Low Threshold (keV)	T1Z1 (Ge)	T1Z3 (Ge)	T1Z4 (Si)	T1Z5 (Ge)	T1Z6 (Si)	$\tilde{Ge}_{in}$	$\tilde{Si}_{in}$	$\tilde{Si}_{out}$
5	2.9	2.5	5.0	2.6	6.6	2.5	5.0	6.6
7	2.6	2.2	4.8	2.3	6.2	2.3	4.8	6.2
10	2.3	1.9	4.4	2.0	5.8	2.0	4.4	5.8

Table 4.15: List of expected neutron backgrounds in 10^{-3} events/kg/day from the extrapolated fits to the simulated single scatter neutron recoil spectra in Sharmila Kamat's thesis[148]. The upper threshold is 100 keV for all detectors and low energy thresholds. T1Z2's y-axis values were unclear, and it was not included in the fit. T1Z2 is not important as T1Z3 and T1Z5 are sufficient to estimate the neutron background for inner germanium detectors. The last three columns give the expected rate for the three classes: inner germanium, inner silicon, and outer silicon respectively.

A double exponential is not unreasonable since Sharmila assumed a neutron flux external to the polyethylene shielding given by one exponential for neutrons below 0.2 GeV and a second exponential above 0.2 GeV. Due to the higher mass of a germanium nuclei, a 2 to 3 times higher energy neutron is typically required to produce a nuclear recoil of the same energy. It is not unreasonable for the a significant number of the resulting below 100 keV nuclear recoils in germanium to be due to the second exponential, while the similar nuclear recoils in silicon were dominated by the first exponential.

Table 4.15 shows the estimated single scatter neutron background in 10^{-3} events/kg/day from the extrapolated fits. Notice that the inner detectors have a lower simulated rate. In this estimate, 3 of the 4 silicon detectors were inner detectors. The mean rates, resulting from my fit to the single scatter,

neutron-induced recoil spectrum from 10-100 keV, are lower than Sharmila's for both germanium and silicon. To be conservative, I rescale my neutron rates to match Sharmila's rates. The expected rate per kg day is $8.8\times10^{-4} \pm 4.2\times10^{-4}$ for the strict germanium analysis, $9.4\times10^{-4}\pm4.5\times10^{-4}$ for the looser germanium analysis, and $3.5\times10^{-3} \pm 1.7\times10^{-3}$ for the silicon analysis. From Eqn. 4.17.1 the expected neutron background is 0.025±0.012 events for the strict germanium analysis, 0.026±0.13 for the loose germanium analysis, and 0.032±0.015 events for silicon analysis. The quoted errors are derived from the simulation's statistical error in the single scatter germanium and silicon event rates given in Sharmila's thesis [148].

4.18 Projecting to the Future

WIMP-search runs in the near future are expected to have substantially improved sensitivity coming from increased detector mass and livetime, reduced exposure to radon, and an improved surface event rejection cut. In the longer term, additional improvement is expected from improved detector design.

- The next CDMS WIMP-search run is already underway. It began in October of 2006 and utilizes an additional 18 detectors. The next WIMP-search run has already accumulated over twice the livetime as this one.

- Care has been taken in handling the additional 18 detectors to minimize their radon exposure. The reduced radon exposure is expected to decrease

the rate of surface events. The additional detectors have been shown to exhibit lower rates of surface events [151].

- The analysis of surface events from this WIMP-search run has lead to increased understanding of surface events. Future runs are expected to be able to use the increased understanding to improve the surface event rejection cut.

In the longer term, additional improvements in sensitivity are expected from improved detector design.

- New detectors that are 2.5 times thicker than the current detectors are being designed. These detectors are expected to reduce the surface event background by a factor of 2.5 in runs planned after the next run.

- A new phonon sensor design that improves phonon sensor coverage and drastically increases reliability.

Bibliography

[1] Zwicky, F. 1933, Helv. Phys. Acta, 6, 110.

[2] Bond, J. R., et al., *In Formation and Evolution of Galaxies and Large Structures* (Reidel, Dordrecht, 1983).

[3] Ap. J. 261:439 Rubin V. C., Ford Jr W. K., 1970.

[4] Milgrom, M., A;. J. **270**, 365 (1983).

[5] E.W. Kolb and M.S. Turner, *The Early Universe* (Addison-Wesley, Redwood City, CA, 1990).

[6] Peebles P. J. E., *Principles of Physical Cosmology* (Princeton University Press, NJ, 1993).

[7] http://pdg.lbl.gov/2005/reviews/bigbangrpp.pdf.

[8] A. Liddle *An Introdcution to Modern Cosmology* (John Wiley & Sons Ltd, West Sussex, England, 2003).

[9] D. N. Spergel et al., Astrophys. J. Supp. **148**, 175 (2003).

[10] Spergel, D. N., Ap. J. S. **170** 377 (2007).

[11] Knop et. al 2003, ApJ, 598:102-137, 2003.

[12] Tully, R. B., Fisher, J. R. 1977, Astron. Ap. **54**, 661.

[13] Jungman, G., Kamionkowski, M., Griest, K., Physics Reports, **267** (1996).

[14] Epstein, R. I., Nature **263**, 198 (1976).

[15] Kirkman, D., et al., astro-ph/0302006.

[16] O'Meara, J. M., et al., Ap. J. **552**, 718 (2001).

[17] Izotev, Y. I., et al., Ap. J. **527**, 757, (1999).

[18] Fields, B. D., Olive, K. A., Ap. J. **506**, 177 (1998).

[19] Spergel, D. N., et al., Ap. J. Supp. **148**, 175 (2003).

[20] Blanchard, A., et al., A.&A. **412**, 35 (2003).

[21] Brindle, S., et al., MNRAS **342**, L72 (2003).

[22] Combes, F., Boissé, P., Mazure, A., Blanchard, A., *Galaxies and Cosmoligy* (Springer-Verlag, Germany, 1995).

[23] Padmanabhan, T., *Cosmology and Astrophysics Through Problems*

[24] Mather, J. C. et al., Astrophys. J. **512**, 511 (1999).

[25] Courteau, S. et al., Astrophys. J. **544**, 636 (2000).

[26] Hu, W., Ph.D. Thesis, U. C. Berkeley, 1995.

[27] http://pdg.lbl.gov/2006/reviews/microwavepp.pdf.

[28] de Oliveira-Costa, A., Tegmark, M., *Microwave Foregrounds* (Astron. Soc. of the Pacific, San Francisco 1999).

[29] Hinshaw, G., et al., astro-ph/0603451.

[30] Ruhl, J. E., et al., Astrophs. J. **599**, 786 (2003).

[31] Scott, P. F., et al., Monthly Not. Royal Astron. Soc. **341**, 1076 (2003).

[32] Pearson, T. J., et al., Astrophys. J. **591**, 556 (2003).

[33] Runyan, M. C., et al., Astrophys. J. Supp. **149**, 265 (2003).

[34] Taillet, R., Cosmic Microwave Background and Cosmological Parameters, *2005 Les Houches Predoctoral School*, http://lappweb.in2p3.fr/conferences/LesHouchesPredoc05/lectures/taillet_houches_b05.pdf.

[35] Astier, P., et al., astro-ph/0510447.

[36] Knop, R. A., et al., Ap. J. **598**, 102-137 (2003).

[37] Abell, G. O., *In Problems of Extragalactic Research*, ed McVittie, G. C., (Macmillan, NY, 1962).

[38] Schechter, P., Ap. J. **203**, 297 (1976).

[39] Propris, R. D., et al., Mon. Not. Roy. Astron. Soc. **342**, 725 (2003).

[40] Faber, S. M., Gallagher, J. S., Ann. Rev. Astron. Ap. **17**, 135 (1979). (Cambridge University Press, Cambridge 1996).

[41] Binney and Tremaine, *Galactic Dynamics* (Princeton University Press, NJ, 1991).

[42] Merritt, D., 1988 *The Minnesota Lectures on Clusters of Galaxies and Large-Scale Structures*, ed. J. M. Dickey (ASP Conf. Ser., 5), p. 175.

[43] Yang et al., MNRAS, **358**, 217 (2005).

[44] Jeltema, T. E., Canizares, C. R., Bautz, M. W., Buote, D. A., Ap. J. **624**, 606 (2005).

[45] Neumann, D. M., Lumb, D. H., Pratt, G. W., Briel, U. G., A&A, **400**, 811 (2003).

[46] Arnaud, Pointecouteau, E., Pratt, G. W., A&A, in press (2005).

[47] Tyson, J. A., AJ, **96**, 1 (1988).

[48] van Waerbeke, L., Mellier, Y., Pello, R., et al, A&A **393**, 369 (2002).

[49] Hoekstra, H., Yee, H. K. C., Gladders, M., Ap. J. **577**, 595 (2002).

[50] Schneider, P., astro-ph/0509252.

[51] Markevitch, M. astro-ph/0511345.

[52] Clow, D., Gonzalez, A. H., Markevitch, M. Ap. J. **604**, 596 (2004).

[53] Clow, D.,et al., Astrophys. J. Lett. **648**, L109 (2006).

[54] Markevitch, M., et al., Ap. J. **606**, 819 (2004).

[55] Rohlfs, K., Kreitschmann, J., A & A, 201, 51 (1988).

[56] Persic, M., Salucci, P., Stel, F., Mon. Not. R. Astron. Soc. **281**, 27 (1996).

[57] Navarro, J. F., Frenk, C. S., White, S. D., Ap. J. **462**, 563 (1996).

[58] Dehnen W., Binney J. 1998, MNRAS, 294, 429.

[59] Evans N. W., Binney J. 2001, MNRAS, 327, L27.

[60] Bissantz N., Debattista V. P., Gerhard O. 2004, ApJ, 601, L155.

[61] Olling, R. P., Merrifield, M. R., astro-ph/9907353.

[62] Harris, G., yCat, 7202 (1997).

[63] Wilhelm, R., et al., AJ, **117**, 2329, (1999).

[64] Morrison H. et al., AJ, **119**, 2254, (2000).

[65] Mateo, M., ARA&A, **36**, 435, (1998).

[66] Battaglia, G., et al., astro-ph/0506102.

[67] Bahcall, J. N., Schmidt, M., Soneira, R. M., Ap. J. **265**, 730 (1983).

[68] Caldwell, R. R., Ostriker, J. P., Ap. J. **251**, 61 (1981).

[69] Turner, M. S., Phys. Rev. D **33**, 889 (1986).

[70] Carignan, C. & Beaulieu, S. Ap.J. **347**, 760 (1989).

[71] Drukier, A. K., Freese, K., & Spergel, D. N. Phys. Rev. D **33**, 3495 (1986).

[72] Cudworth, A. J. Astron. J. **99**, 590 (1990).

[73] Anderson, J. D., et al., ApJ, **342**, 539 (1989).

[74] Anderson, J. D., et al., ApJ, **448**, 885 (1995).

[75] Grøn, Ø., Soleng, H. H., ApJ, **456**, 445 (1996).

[76] Khriplovich, I. B., Pitjeva, E. V., astro-ph/0601422.

[77] Feast, M., and Whitelock, P., Mon. Not. Roy. Astron. Soc., **291**, 683 (1997).

[78] Mihalas, D., Binney, J., *Galactic Astronomy* (Princeton University Press, NJ, 1981).

[79] Baker, D. C., *Unsolved problems of the Milky Way*, IAU Symposium 169, eds Blitz, D. & Teuben, P. (Dordrecht Kluwer, 1996), p 198.

[80] Reid, M. J., Ann. Rev. Astron. & Astrophys., **31**, 345 (1998).

[81] Lewin, J. D., Smith, P. F., Astropart. Phys. **6**, 87 (1996).

[82] Dehnen, W., and Binney, J. J., Mon. Not. Roy. Astron. Soc. **298**, 387 (1998).

[83] Lang, K., *Astrophysical Formulae* (Spring, Berlin, 1980).

[84] Green, A. Phys. Rev. D **68**, 023004 (2003); Erratum-ibid. **69**, 109902 (2004).

[85] Eder, G., *Nuclear Forces* (MIT Press, 1968), Chapter 7.

[86] Engel, J., Phys. Lett. B **264**, 114 (1991).

[87] Helm, R. H., Phys. Rev. **104**, 2357 (1956).

[88] Hahn, Ravenhall, and Hofstadter, Phys. Rev. **101**, 1131 (1956).

[89] Fricke, G, et al., Atomic Data and Nuclear Data Tables 60 (1995) 195.

[90] Brodsky, A., *Handbook of Radiation Measurement and Protection*, (CRC Press, FL 1978) v2, 620.

[91] Enqvist, T., et al., Nuclear Instruments and Methods in Physics Research A, **554**, 286 (1995).

[92] Arnaboldi, C., et al., Astropart. Phys., **20**, 91 (2003).

[93] Esch, E.-I., et al., Nucl. Instr. and Meth. A, **538**, 516 (2005).

[94] http://www.dnr.state.mn.us/state_parks/soudan_underground_mine/index.html

[95] http://nrhp.mnhs.org/property_overview.cfm?propertyID=1

[96] http://www.hep.umn.edu/soudan/brochure.html

[97] Akerib, D. S., et al. [CDMS], Phys. Rev. D **72**, 052009 (2005).

[98] Ruddick, K., MINOS Report No. NuMI-L-210, 1996.

[99] Bauer, D. CDMS Internal Backgrounds Note: Soudan Radon (2004).

[100] J. Sander, CDMS R119 Analysis enote 81.

[101] Reisetter, A., Fraction of Neutrons Vetoed at Soudan, CDMS Monte Carlo Note, Sept. 3, 2004.

[102] Reisetter, A., Study of the Efficacy of an Umbrella Veto at Soudan, CDMS Background Note, July 28, 2006.

[103] Fagaly, R. L., Rev. Sci. Instrum. **77**, 101101 (2006).

[104] http://cdms.berkeley.edu/cdms_restricted/cdmsnotes/clc_thesis.pdf.

[105] http://cdms.berkeley.edu/cdms_restricted/cdmsnotes/mandic_thesis.pdf.

[106] http://cdms.berkeley.edu/cdms_restricted/cdmsnotes/thesis_wang.pdf.

[107] Luke, P. N., Journal of Applied Physics **64**(12), 6858 (1988).

[108] Neganov, B., Trofimov, V., Otkrytia, Izobreteniya **146**, 215 (1985).

[109] Clarke, R. M., et al., Applied Physics Letters **76**(20), 2958 (2000).

[110] http://cdms.berkeley.edu/cdms_restricted/warmelectronics/ZIP_FEB9U/ZIP_FEB9U.html.

[111] http://ppd.fnal.gov/experiments/cdms/

[112] http://www-ppd.fnal.gov/EEDOffice-W/Projects/CDMS/POWER_SUPPLY/cdms8.html.

[113] http://cdms.physics.ucsb.edu/sburke/triglogindex.html.

[114] V. Sanglard *et al.*, (EDELWEISS Collab.), Phys. Rev. D **71**, 122002 (2005).

[115] Gondolo, P., Gelmini, G., Phys. Rev. D **71**, 123520 (2005).

[116] Kim, Y.G., Nihei, T., Roszkowski, L., de Austri, R.R., hep-ph/0208069.

[117] http://www-cdfonline.fnal.gov/r2dm/

[118] R. Mahapatra, CDMS R119 Analysis enotes 22, 55, 61, and 135.

[119] M. Attisha and J. Sander, CDMS R119 Analysis enotes 28, 55, and 59.

[120] J. Sander, CDMS R119 Analysis enote 117.

[121] J. Sander, CDMS R119 Analysis enote 123.

[122] J. Sander, CDMS R119 Analysis enotes 241 and 258.

[123] R. W. Ogburn and J. Sander, CDMS R119 Analysis enotes 89 and 160.

[124] L. Duong, CDMS R119 Analysis enote 105.

[125] J. Sander, CDMS R119 Analysis enote 159.

[126] L.Duong, CDMS R119 Analysis enotes 55, 105, and 132.

[127] L.Duong, CDMS R119 Analysis enotes 21, 112, and 114.

[128] R. W. Ogburn, CDMS R119 Analysis enotes 76 and 128.

[129] A. Reisetter, CDMS R119 Analysis enote 62.

[130] J. Sander and R. Mahapatra, CDMS R119 Analysis enotes 69 and 95.

[131] R. W. Ogburn, CDMS R119 Analysis enotes 55 and 171.

[132] Akerib, D. S., et al., Phys. Rev. D **D68**, 082002 (2003).

[133] Linhard, J., et al., Kgl. Danske Videnskab Selskab Mat. Fys. Medd. **33**, 10 (1963).

[134] Chasman, C., et al., Phys. Rev. Lett. **15**, 245 (1965).

[135] Sattler, R., Phys. Rev. **138**, A1815 (1965).

[136] Shutt, T., et al., Phys. Rev. Lett. **69**, 3425 (1992).

[137] http://titus.stanford.edu/cdms_restricted/backgrounds/060629/BetaPaper_v9.pdf.

[138] A. Reisetter, CDMS R119 Analysis enote 82.

[139] http://cdms.berkeley.edu/cdms_restricted/analysis/BetaCalib/index.html.

[140] Pyle, M., private communication.

[141] Akerib, D. S., et al. [CDMS], Phys. Rev. D **73**, 011102 (2006).

[142] http://titus.stanford.edu/cdms_restricted/Soudan/R18/ebook/040426/CutLeakageR118BS.pdf

[143] R. Mahapatra, CDMS R119 Analysis enote 155.

[144] Attisha, M. and Filippini, J, CDMS R119 Analysis enote 17.

[145] R. Mahapatra, CDMS R119 Analysis enote 297.

[146] S. Kamat, Looking at the External Neutron Simuations for the Soudan Shield Configuration, CDMS internal note: http://cdms.cwru.edu/cdms_restricted/mclinks/sudneuts.htm.

[147] Yellin, S., Examining the Most Recent MC Tables from Raul, CDMS Monte Carlo Note, Oct. 15, 2004.

[148] Kamat, S., Ph.D. Thesis, Case Western Reserve University, 2005, p. 239.

[149] http://pdg.lbl/gov/2007/reviews/statrpp.pdf.

[150] Feldman, G.J. and Cousins, R.D., Phys. Rev. D **57**, 3873 (1998).

[151] R. Mahapatra, CDMS R123 Analysis enote 14.