Exclusion Limits on the WIMP-Nucleon Elastic-Scattering Cross Section from the Cryogenic Dark Matter Search

by

Sunil Ramanlal Golwala

B.A. (University of Chicago) 1993
M.A. (University of California, Berkeley) 1995

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

in

Physics

in the

GRADUATE DIVISION
of the
UNIVERSITY of CALIFORNIA at BERKELEY

Committee in charge:

Professor Bernard Sadoulet, Chair
Professor Marc Davis
Professor Paul L. Richards
Professor William J. Welch

Fall 2000

The dissertation of Sunil Ramanlal Golwala is approved:

Chair Date

Date

Date

Date

University of California at Berkeley

Fall 2000

Exclusion Limits on the WIMP-Nucleon Elastic-Scattering Cross Section from the Cryogenic Dark Matter Search

Sunil Ramanlal Golwala

Abstract

Exclusion Limits on the WIMP-Nucleon Elastic-Scattering Cross Section from the Cryogenic Dark Matter Search

by

Sunil Ramanlal Golwala
Doctor of Philosophy in Physics

University of California at Berkeley

Professor Bernard Sadoulet, Chair

Extensive evidence indicates that a large fraction of the matter in the universe is non-luminous, nonbaryonic, and "cold" — nonrelativistic at the time matter began to dominate the energy density of the universe. Weakly Interacting Massive Particles (WIMPs) are an excellent candidate for nonbaryonic, cold dark matter. Minimal supersymmetry provides a natural WIMP candidate in the form of the lightest superpartner, with a typical mass $M_\delta \sim 100$ GeV c^{-2}. WIMPs are expected to have collapsed into a roughly isothermal, spherical halo within which the visible portion of our galaxy resides. They would scatter off nuclei via the weak interaction, potentially allowing their direct detection.

The Cryogenic Dark Matter Search (CDMS) employs Ge and Si detectors to search for WIMPs via their elastic-scattering interactions with nuclei while discriminating against interactions of background particles. The former yield nuclear recoils while the latter produce electron recoils. The ionization yield (the ratio of ionization production to recoil energy in a semiconductor) of a particle interaction differs greatly for nuclear and electron recoils. CDMS detectors measure phonon and electron-hole-pair production to determine recoil energy and ionization yield for each event and thereby discriminate nuclear recoils from electron recoils.

This dissertation reports new limits on the spin-independent WIMP-nucleon elastic-scattering cross section that exclude unexplored parameter space above 10 GeV c^{-2} WIMP mass and, at $> 75\%$ CL, the entire 3σ allowed region for the WIMP signal reported by the DAMA experiment. The experimental apparatus, detector performance, and data analysis are fully described.

Professor Bernard Sadoulet
Dissertation Committee Chair

To my parents

Acknowledgements

It is truly unfortunate that this dissertation cannot bear the names of all the members of the CDMS collaboration. I am quite fortunate to have been able to participate in bringing so many years of work by so many people to fruition. It would be impossible to list each person's constributions, and I certainly would make omissions. I thank everyone, both those who are currently involved in CDMS and those who have moved on, for their efforts.

My advisor, Bernard Sadoulet, has provided me with a wide variety of opportunities and a great deal of independence. I hope to follow the example he sets in terms of the breadth and depth of his knowledge, from the technical details to the big picture of cosmology and particle physics.

A number of people have played roles as partners-in-crime, teachers, and mentors over the years. My undergraduate advisor, Rene Ong, introduced me to a wide variety of topics in particle physics and astrophysics. His willingness to "leave the fold" of accelerator-based particle physics encouraged me to do so, and it is a road I am happy to have taken. Josef Jochum taught me a great deal about tunnel junctions, quasiparticles, and phonons and made the NIS adventure an interesting one. Rick Gaitskell pulled me into operations and analysis at SUF and guided my work during Runs 15 and 16. I can honestly say things would have turned out much differently without his encouragement and interest. His intellectual appetite and tenacity have been a great example, and always knowing Rick would ask the hard question has kept me on my toes. Tom Shutt has always been an encouraging voice and a fount of knowledge about detectors, cryogenics, electronics, plumbing — you name it. I especially admire his physics judgment and his fortitude in pursuing the dead-layer studies when many of us, myself included, were skeptical that a solution could be found; without the new electrodes, Run 19 would not have been nearly as interesting as it has turned out. I have learned a great deal from Dennis Seitz about electronics, especially how to build things the right way. I can think of no one with whom I would rather fight oscillations, ground loops, and pickup. I thank Richard Schnee for teaching me how to do data analysis properly, maintaining the data flow through DarkPipe, keeping me honest, and putting up with my bad moods and frequent exclamations of dismay and horror at the "peculiarities" of the Run 19 data.

I personally owe a debt of gratitude to a number of people, without whose efforts Run 19 (and this dissertation) would not have happened. John Emes is due special thanks for his heroic efforts and competence in building the Run 19 BLIPs. Garth Smith provided a number of most-excellent mechanical things, in particular a way to wire a lot of tower faces. Storn White designed a bitchin' detector package that was a pleasure to use. Steve Butler, Alex Vaynberg, and many other members of the Physics Machine Shop carefully turned Garth's and Storn's designs into reality; I especially thank George Weber for speeding the Run 19 hardware through the shop as deadlines loomed. Danny Lee and Eric Jones did a beautiful job building the Run 19 FET cards and tower. Dan Bauer kept the DAQ and veto in solid condition, and Jim Perales did innumerable cryogen transfers. Thushara Perera answered my incessant neutron background questions and provided simulation after simulation.

The Berkeley group has provided an enjoyable work atmosphere. I have made innumerable coffee trips and lunch excursions with and learned a great deal from the many members of the group during my stay — Dan Akerib, Eric Aubourg, Peter Barnes, Cherie Brumfield, Alan Cummings, Angela Da Silva, Tim Edberg, Rick Gaitskell, Jochen Hellmig, Chris Hernikl, Josef Jochum, Maria Isaac, Vuk Mandic, Patrizia Meunier, Francois Queinnec, Ron Ross, Dennis Seitz, Paul Shestople, Tom Shutt, Garth Smith, Walter Stockwell, John Taylor, Roy Therrien, Storn

White, Jeff Wilkerson, and Betty Young. Anthony Spadafora has kindly translated schedules and budgets into language I can understand and has provided insight into our management's thought processes. Rose von Thater and Rachel Winheld have been sources of encouragement, and it is unfortunate that I had less time to work with them than I would have liked. I thank Caryl Esteves, Kate Kim, Letty Palomares, Donna Sakima, and Anne Takizawa for cushioning the impact of the university's bureaucracy. Blas Cabrera kindly welcomed me into his group and lab during my time at Stanford. I have enjoyed many conversations and much commiseration with my fellow CDMS graduate students, Roland Clarke, Sae Woo Nam, Thushara Perera, Andrew Sonnenschein, Andrea Snadden, and Roy Therrien. I wish the next generation — Ray Bunker, Clarence Chang, Don Driscoll, Vuk Mandic, Tarek Saab, and a number of others — the best of luck. It's your experiment now!

The guys at the Deakin Street house provided a lively, if somewhat physics-saturated, atmosphere during my early years in grad school. I have enjoyed getting to know the variety of people who have passed through since. John Monnier has been a party to many late-night conversations about life, science, and jobs. Ed Burns has been a cool roommate for the last four years, providing the impetus for many musical excursions. Now if only he'll give me my bass back...

Finally, I thank my parents for letting me go my own way. I know it has been difficult being not quite clear on what your son has been doing for the last seven years, but I think it's going to work out just fine.

Contents

List of Figures

List of Tables

Chapter 1

Dark Matter and WIMPs

1.1 Introduction

The most important goal of cosmology today is the determination of the value and composition of the energy density of the universe. It has been apparent for decades that the dominant component is "dark" — *i.e.*, is not in forms, such as stars or gas, that emit electromagnetic radiation. In the last two decades, with the advent of a precise theory of Big Bang nucleosynthesis of the light elements, it has become increasingly clear that the majority of this dark energy density is not baryonic. In just the last two years, it has been argued that a large component of this dark energy density may not even be in the form of matter, but rather in a vacuum energy density or omnipresent scalar field. Though this "dark energy" may be the majority shareholder in the universe's stock, the arguments for a strong minority stake by nonbaryonic dark matter have also grown stronger. The identity and quantity of this latter component remain significant unanswered questions.

In this chapter, I discuss the evidence for the existence of dark matter and dark energy. I describe various dark-matter candidates, concentrating on Weakly Interacting Massive Particles (WIMPs), the one being searched for by the experiment described in this dissertation. In particular, I focus on the motivation for the neutralino, an excellent WIMP candidate provided by particle physics.

1.2 Cosmological Preliminaries

In this section, I discuss some necessary background material. This discussion follows that of standard cosmology textbooks, such as [1] and [2].

1.2.1 The Homogeneous, Isotropic, Expanding Universe

The energy densities present in the universe determine its expansion rate and ultimate fate of collapse or infinite expansion. This information is summarized by the cosmological equations:

$$H^2 \equiv \left(\frac{1}{a}\frac{da}{dt}\right)^2 = \frac{8}{3}\pi G\rho - \frac{1}{a^2R^2} + \frac{\Lambda}{3} \tag{1.1}$$

and

$$\frac{1}{a}\frac{d^2a}{dt^2} = -\frac{4}{3}\pi G\left(\rho + 3p\right) + \frac{\Lambda}{3} \tag{1.2}$$

where a is the scale factor in the metric for the universe (with $c = 1$):

$$ds^2 = dt^2 - [a(t)]^2\left(\frac{dr^2}{1-(r/R)^2} + r^2 d\theta^2 + r^2\sin^2\theta d\phi^2\right) \tag{1.3}$$

H is the Hubble parameter, indicating the expansion rate of the universe. The metric is the homogeneous, isotropic Friedmann-Robertson-Walker metric, describing an expanding universe that has a uniform energy density and curvature and has no special orientation. The universe contains a homogeneous fluid with time-dependent physical density ρ and pressure p. The Friedmann and acceleration equations, Equations 1.1 and 1.2, are simplifications of Einstein's field equation for general relativity under the assumption of the FRW metric. Λ is the cosmological constant, an extra constant term allowed in Einstein's field equation. It acts like an energy density $\rho_V = \Lambda/8\pi G$ that is not diluted by expansion: it can be identified with a nonzero vacuum energy density having pressure $p_V = -\rho_V$. The expansion of the universe is carried by $a(t)$, the scale factor. (r, θ, ϕ) are unitless, comoving coordinates. A point that is at rest with respect to the expansion has fixed comoving coordinates. The global, time-independent, comoving curvature of the universe is specified by the curvature radius R. The physical radius of curvature is $a(t)R$. For R real, the space is positively curved or spheroidal, the 3-dimensional analogue of the 2-dimensional surface of a sphere. For R imaginary, it is negatively curved (hyperboloidal). For $R \to \infty$, the space is flat or Euclidean. These three categories are also known as closed, open, and flat in reference to their volume: a closed universe has a finite volume because it curves back on itself (but has no edges), while open and flat universes have infinite volumes and no edges. The Freidmann equation describes the energetics of the universe: the kinetic energy of the expansion, expressed by H^2, is equal to the negative of its gravitational potential energy, with contributions from the actual energy density, from curvature, and from vacuum energy. The acceleration equation describes the dynamics: normal matter, with its positive energy and pressure, slows down the expansion, while the cosmological constant accelerates the expansion.

The above equations may be rewritten in terms of the Ω parameters. They simplify under the assumption that the pressure of non-vacuum energy is zero: that is, the normal-matter energy density is dominated by nonrelativistic matter. Vacuum energy is defined by $p = -\rho$, so the vacuum energy's pressure term is retained, yielding

$$\left(\frac{H}{H_0}\right)^2 = \Omega_m\left(\frac{a_0}{a}\right)^3 + \Omega_k\left(\frac{a_0}{a}\right)^2 + \Omega_\Lambda \tag{1.4}$$

$$\frac{1}{aH_0^2}\frac{d^2a}{dt^2} = \Omega_\Lambda - \frac{1}{2}\Omega_m\left(\frac{a_0}{a}\right)^2 \tag{1.5}$$

with

$$\Omega_m \equiv \frac{8\pi G}{3H_0^2}\rho_0 \equiv \frac{\rho_0}{\rho_c} \tag{1.6}$$

$$\Omega_k \equiv -\frac{1}{H_0^2 a_0^2 R^2} \tag{1.7}$$

$$\Omega_\Lambda \equiv \frac{\Lambda}{3H_0^2} \tag{1.8}$$

with 0 subscripts indicating present-day values. The scaling for the matter term follows from the volume dilution of the particle density by the expansion. The Ω's are all constants with their present-day values. Equation 1.4 at $a = a_0$ gives

$$1 - \Omega_k = \Omega_m + \Omega_\Lambda \tag{1.9}$$

That is, the total density parameter of the universe today, $\Omega = \Omega_m + \Omega_\Lambda$, is directly related to its present curvature. A flat universe has $\Omega_m + \Omega_\Lambda = 1$, an open universe has $\Omega_m + \Omega_\Lambda < 1$, and a closed universe has $\Omega_m + \Omega_\Lambda > 1$. The critical density, ρ_c, is the value of the density that delineates the three cases.

The ultimate fate of the universe is determined by the two cosmological equations and the values of Ω_m and Ω_Λ. Consider first the case $\Omega_\Lambda = 0$. The acceleration is always negative, indicating the expansion is always slowing down. If $H^2 > 0$ for all time, the expansion continues forever. This case occurs if $\Omega_k \geq 0$ because the right side of the Friedmann equation is always positive and nonzero. Flat and open universes expand forever. There is a distinction between flat and open universes in their asymptotic behavior. In an open universe, the matter term eventually becomes negligible compared to the curvature term because of their different scalings with a. Solving Equation 1.4 in this case,

$$\frac{da}{dt} = \sqrt{\Omega_k H_0^2 a_0^2} \tag{1.10}$$

That is, the velocity of the expansion velocity goes to a constant. The Hubble parameter goes to zero as $H \propto 1/a \propto 1/t$, but the expansion goes on forever. Intuitively, this is easy to understand: the kinetic energy of the expansion comes to dominate over the gravitational potential of the energy density, so the universe coasts forever, like a projectile launched with velocity greater than the escape velocity. On the other hand, $\Omega_k = 0$ exactly for a flat universe. The expansion rate satisfies

$$\frac{da}{dt} = \sqrt{\Omega_m H_0^2 \frac{a_0^3}{a}} \tag{1.11}$$

which results in

$$a(t) = \left(\frac{2}{3}t\right)^{2/3} \Omega_m^{1/3} H_0^{2/3} a_0 \tag{1.12}$$

In a flat universe, the expansion velocity asymptotes to 0: the kinetic energy of the expansion is exactly balanced by the gravitational potential. This is like a projectile launched with exactly the escape velocity. Finally, in a closed universe, $\Omega_k < 0$, so there is a time when $H = 0$. The acceleration is always negative, so the time when $H = 0$ corresponds to a turnaround point, after which $H < 0$ and the universe recollapses. This is like a projectile launched with velocity less than the escape velocity.

Consider now the case $\Omega_\Lambda \neq 0$. For $\Lambda > 0$, Ω_Λ causes positive acceleration, as is seen in Equation 1.5. In the case of a flat or open universe, $H^2 > 0$ holds for all time and Equation 1.1 asymptotes to

$$\frac{da}{dt} = \sqrt{\frac{\Lambda}{3}}\, a \tag{1.13}$$

yielding an exponential expansion with time constant $\tau = \sqrt{3/\Lambda}$. In the case of a closed universe, the cosmological constant term can dominate if the universe is not too closed: if $|\Omega_k|$ is small enough,

the cosmological constant dominates before the universe can reach the turnaround point. There are closed universes that expand forever. Conversely, if the universe does reach the turnaround point, it continues to collapse because Ω_Λ becomes less important as a becomes small again. There are no cases where the expansion turns around twice. Negative Λ, though usually not considered because there appears to be no observational evidence for it, has a similar nonintuitive effect: it can cause open universes to recollapse by creating a turnaround point. In an open or flat universe, as Ω_m and Ω_k are diluted by the expansion, there is a value of the scale factor at which $H = 0$ because $\Omega_\Lambda < 0$. Once the expansion has been halted and turned around, even an open universe recollapses because of its own gravity: all the acceleration terms are negative. For a closed universe, Ω_Λ may modify the dynamics of the turnaround. In all cases, once recollapse has begun, Ω_Λ becomes less important and the recollapse is dominated by the matter (or radiation) density.

1.2.2 Redshift and the Hubble Parameter

Because it is only the quantity a/a_0 that is relevant in most places, it is useful to define the *redshift* z by

$$1 + z = \frac{a_0}{a} \tag{1.14}$$

$z = 0$ corresponds to today, z increases as t decreases. The redshift is defined this way because it gives the factor by which light emitted at scale factor a from an object at rest in comoving coordinates is redshifted as the universe expands. (Redshifting occurs because the wavelength of the light dilates with the expansion.) For objects with $z \ll 1$, an object's redshift and its present physical distance d are related by the Hubble Law (for $c = 1$),

$$z = H_0 d \tag{1.15}$$

This relation is easy to understand. For objects with $z \ll 1$, the light travel time from the object is $\delta t = d$ (with $c = 1$). In this time δt, the universe expands by $\delta a = a_0 H_0 \delta t$. The light emitted is thus redshifted by $z = -1 + a_0/a(t_0 - \delta t) = H_0 \delta t$, giving the above relation. Hubble's observational discovery of this relation was the first piece of evidence for an expanding universe. The Hubble constant, H_0, is known to a precision of about 15% today; the commonly quoted value is [3]

$$H_0 = 67 \pm 10 \text{ km s}^{-1} \text{ Mpc}^{-1} = 100h \text{ km s}^{-1} \text{ Mpc}^{-1} \tag{1.16}$$

where h parameterizes the present uncertainty on H_0. At larger distances, the expansion of the universe modifies the Hubble law; the Hubble law is the first term of a Taylor expansion in d or z. I return to the more general form later.

1.2.3 A Short History

The universe is not an inert place; it contains a thermal system with reactions continually occuring that maintain the system (hereafter called the "universe," though this is not exactly correct) in thermal equilibrium. The primary characteristic that makes the system tractable is that the entropy per unit comoving volume is conserved during the expansion — the expansion is adiabatic. This can be demonstrated formally; see, for example, [2]. Intuitively, the situation is equivalent to the elementary quantum mechanics problem of a particle in an adiabatically expanding

box. The particle remains in the same quantum level as the box expands if the expansion is slow enough, though the wavefunction and energy of the quantum level change. That the expansion is adiabatic can be seen by putting the Hubble parameter in units closer to physical experience: $H_0 \sim 67$ km s^{-1} Mpc^{-1} $= 2.2 \times 10^{-18}$ m s^{-1} m^{-1}. The Hubble expansion today is incredibly adiabatic, just in case it was not obvious from the world around us! If the expansion is adiabatic today, then it was adiabatic at arbitrary earlier times because the Hubble parameter increases less quickly than the energies of the particles. Given that the expansion is adiabatic, it holds that the temperature scales as a^{-1}, where the temperature is given by T_γ, the photon temperature. Discrete events modify this scaling in discrete steps. For example, when T_γ drops below the electron rest mass, the thermally generated electrons and positrons annihilate, dumping energy in the photons. However, the annihilation process is reversible, so the change in T_γ is set by entropy conservation. The $T_\gamma \propto a^{-1}$ scaling is easily corrected for these discrete events, as is shown below.

There are two clearly defined phases in the expansion history of the universe, depending on what type of energy density dominated:

- *Radiation Domination*: *Radiation* refers generically to all particles that are relativistic, with $E \gg m$. Their energy density scales as a^{-4} because, in addition to dilution of their number density, their energy is also redshifted by the expansion (*e.g.*, photons). The velocities of nonrelativistic particles (*matter*) are also redshifted. This additional energy loss is ignored because the kinetic energy of nonrelativistic particles is small compared to their rest mass, and hence the energy density of nonrelativistic particles scales as a^{-3}. Because of these scalings with a, there is some redshift above which the energy density of radiation dominates over that of matter.

  Relativistic particles have pressure $p = \rho$. The acceleration equation, Equation 1.2, includes a pressure term, yielding stronger deceleration of the expansion during radiation domination. Pressure affects the expansion not because of the mechanical force it exerts, but because it is an energy density also and thus its gravity works against the expansion. During radiation domination, the universe obeys a different growth law.

- *Radiation-Matter Equality*: This is defined to be the time (or redshift) at which the matter and radiation energy densities are equal. For times before this point, radiation dominates and determines the expansion. After this point, matter ($p = 0$) dominates. The redshift of equality is $z_{eq} = 23700\,\Omega_m h^2$, with $T_{eq} = 5.6\,\Omega_m h^2$ eV.

- *Matter Domination*: From z_{eq} to the present, matter has dominated the energy density of the universe and determined the expansion rate. We may be at the point of matter-vacuum energy equality; this is discussed later.

As the universe expands and cools, various phenomena occur: massive particles become nonrelativistic; reaction rates decrease due to the decrease in particle number density and average energy; and reactions with energy thresholds (*e.g.*, creation of massive particles) stop occurring. These phenomena give rise to three specific, significant "events:"

- *Neutrino decoupling*: Neutrinos, being almost massless, act as radiation during most of the history of the universe. However, unlike photons, neutrinos interact only via the weak force, and thus the reactions that maintain them in thermal equilibrium with each other and the rest

of the universe fail to proceed quickly enough once $T_\nu \lesssim 3$ MeV. This point is termed neutrino "decoupling." However, because of the adiabaticity of the expansion, the neutrinos maintain their phase-space distribution, and thus a temperature T_ν can be assigned to them. This temperature falls as $1/a$ due to expansion of the universe. Immediately after decoupling, T_ν tracks the temperature of the photons and the rest of the universe. However, discrete events raise the temperature of the photons (*e.g.*, electron-positron annihilation, mentioned above). The neutrinos are unaffected, so the neutrino temperature differs from the photon temperature at later times. The correction factor due to electron-positron annihilation is

$$T_\nu = \left(\frac{4}{11}\right)^{1/3} T_\gamma \tag{1.17}$$

The factor arises as follows. The annihilation process preserves entropy. The entropy of a relativistic fluid is $\alpha(4\sigma_B/c)T^3$, where σ_B is the Stefan-Boltzmann constant and $\alpha = 4/3$ for bosons and $\alpha = 7/6$ for fermions. The photon entropy thus is increased by a factor $(4/3 + 2(7/6))/(4/3) = 11/4$ by electron-positron annihilation. Since entropy scales as T^3, the photon temperature is increased by a factor $(11/4)^{1/3}$. There are no other strongly or electromagnetically interacting particles with mass less than 1 MeV, so no further modification to this correction factor is needed. Using the above correction, the neutrino temperature today is $T_{\nu,0} = 1.95$ K.

Note that such corrections can be applied generally by calculating an "effective" number of relativistic degrees for freedom, g, with

$$g = N_B + N_F \frac{7/6}{4/3} = N_B + \frac{7}{8} N_F \tag{1.18}$$

The entropy can be calculated at any time using the boson expression and g: $s = (4/3)(4\sigma_B/c)T^3 g$. Additional corrections must be applied if the neutrino and photon temperatures differ, giving

$$g = N_B + \frac{7}{8} N_F + \frac{7}{8}\frac{4}{11} N_\nu \tag{1.19}$$

where N_F are the non-neutrino fermionic degrees of freedom. This expression is used during the discussion of light-neutrino dark matter later. The formula can be generalized to any species that decouples while relativistic.

- *Nucleosynthesis*: At very high temperatures, above a few MeV, nuclei do not exist because the average particle energy is above the nuclear binding energy. As the temperature drops below 1 MeV, protons and neutrons form into light nuclei. I return to this later in conjunction with the baryon density of the universe.

- *Recombination/Photon Decoupling*: As the universe cools below a few eV, photons lose the ability to ionize hydrogen. The universe quickly goes from being a plasma of protons and electrons, with small amounts of heavier nuclei and atoms, to becoming almost completely neutral hydrogen. This process is recombination. Some small residual ionized fraction remains with which the photons may interact. Soon after, the expansion of the universe dilutes the density of free electrons enough that photons can no longer scatter. This point is radiation-matter decoupling and is the time of last scattering of the photons. These photons are the

cosmic background radiation (CBR) observed at mm and cm wavelengths today. They give an image of the baryon distribution at the time of last scattering. The redshift of recombination and decoupling is $z_{dec} \approx 1100$.

I refer to these events and epochs frequently during this chapter.

1.2.4 Mass-to-Light Ratios

When discussing the amount of dark matter inferred to be present in some system, or the universe as a whole, it is useful to consider the *mass-to-light ratio*, Υ, which is the ratio of the mass to the luminosity of the system, in solar units. The luminosity density of the universe in the B band (centered at 445 nm) is observed to be [1]

$$j = 1.0 \times 10^8 \, e^{\pm 0.26} \, h \, L_\odot \, \mathrm{Mpc}^{-3} \tag{1.20}$$

The critical mass density is

$$\rho_c = \frac{3H_0^2}{8\pi G} = 1.88 \times 10^{-29} \, h^2 \, \mathrm{g \, cm}^{-2} \tag{1.21}$$

which, given the above luminosity density, corresponds to a mass-to-light ratio of

$$\Upsilon_c = 2800 \, e^{\pm 0.3} \, h \, \frac{M_\odot}{L_\odot} \tag{1.22}$$

In the cores of galaxies, the observed mass-to-light ratio is $\Upsilon = 12 \, e^{\pm 0.2} \, h$ [1], clearly much less than the critical density. Taking this to be typical for luminous matter and extrapolating to the universe as a whole gives

$$\Omega_{\mathrm{lum}} = 0.0043 \, e^{\pm 0.3} \tag{1.23}$$

Clearly, the energy density in stars is small compared to ρ_c. In this chapter, similar extrapolations are made based on Υ measured in a variety of systems.

1.3 Evidence for the Existence of Dark Matter and Energy

The evidence for the existence of dark matter and energy is indirect, given by observation of their gravitational effects. This evidence is available on scales ranging from galaxies to the universe as a whole. For the smaller, gravitationally collapsed systems, these measurements are only lower bounds, as the measurements are usually insensitive to a homogeneous underlying energy density.

1.3.1 Dark Matter in Galaxies

In this section, I consider the evidence for the dominance of dark matter in the two most important types of galaxies: spirals and ellipticals. Salucci and Persic [4] provide an excellent review of the dominance of dark matter in spirals as well as in other galaxy types. Work by Davis and White [5] and by Loewenstein and White [6] provides convincing evidence for dark-matter dominance in ellipticals.

Two general trends are seen. First, in all systems, dark matter dominates at large radii, beyond the luminous regions, but possesses flat core profiles so that dark and luminous matter give similar contributions inside the luminous regions. Second, dark matter is less dominant in higher-luminosity systems, presumably because the efficiency of star formation is enhanced as the depth of the gravitational potential well is increased.

The optical radius, R_{opt}, of a galaxy is defined to be the radius that contains 83% of the light. The half-light radius, r_e, contains half the light. Surface brightness is the flux received per unit solid angle on the sky and is usually reported in units of magnitudes/square arcsec. The typical luminosity of galaxies is given by the quantity $L_* \approx 2.5 \times 10^{10}\, h^{-2}\, L_\odot$.

Spiral Galaxies

The most striking argument for the existence of dark matter on galactic scales is provided by the rotation curves of spiral galaxies. The visible structure of spiral galaxies is dominated by a luminous disk of stars that corotates about its axis. The disk is usually exponential in radius, having surface brightness obeying

$$I(r) = I_0 \exp(-r/R_D) \tag{1.24}$$

where R_D, the disk radius, defines the scale of the disk. For a disk following the above relation exactly, $R_{opt} = 3.2R_D$. For a L_* galaxy, $R_{opt} \sim 10$ kpc. Spirals also usually have a bulge-like component at the center, with size of order 1 kpc or less, which contributes negligibly to the gravitational potential at larger radii. Clouds of neutral atomic hydrogen are usually visible out to $(2-3)R_{opt} \sim 20-30$ kpc. The velocities of the stars and gas clouds in other galaxies can be measured by observation of the redshifts of various spectral lines emitted by these components. For stars, atomic electronic structure lines in the visible and near-infrared are used. For cold clouds of neutral atomic hydrogen, the fine-structure line at 21 cm $\approx$ 1430 MHz is used.

"Rotation curves" obtained in this way for approximately 1000 spiral galaxies have been collected by Persic and Salucci [7]. They fit a "universal rotation curve" to spiral galaxies over a factor of 75 range in luminosity with a functional form that scales with the disk velocity at the optical radius, $v(R_{opt})$, and depends only on L/L_*. In Figure 1.1, the average rotation curves determined for 11 different luminosity bins are shown. Each bin contains about 50 to 100 galaxies. The averaging is performed after scaling by R_{opt} and $v(R_{opt})$. The small error bars on the curves indicate the averaging is valid; *i.e.*, there is little scatter about the universal curve for each luminosity bin.

The remarkable feature of these curves is that they are flat or nearly flat between 1 and 2 optical radii and show little sign of falling off (except in the highest L bins). It is straightforward to calculate the mass profile needed to produce a flat rotation curve. Equating the gravitational and centripetal accelerations for a particle in a circular orbit of radius r,

$$\frac{[v(r)]^2}{r} = G\frac{M(r)}{r^2} \tag{1.25}$$

where $v(r)$ is the observed rotation curve and $M(r)$ is the mass enclosed in a sphere of radius r. Newton's theorem implies that, if the mass distribution is an axisymmetric disk, it is only the mass enclosed by the radius r that determines the gravitational potential at a point in the plane

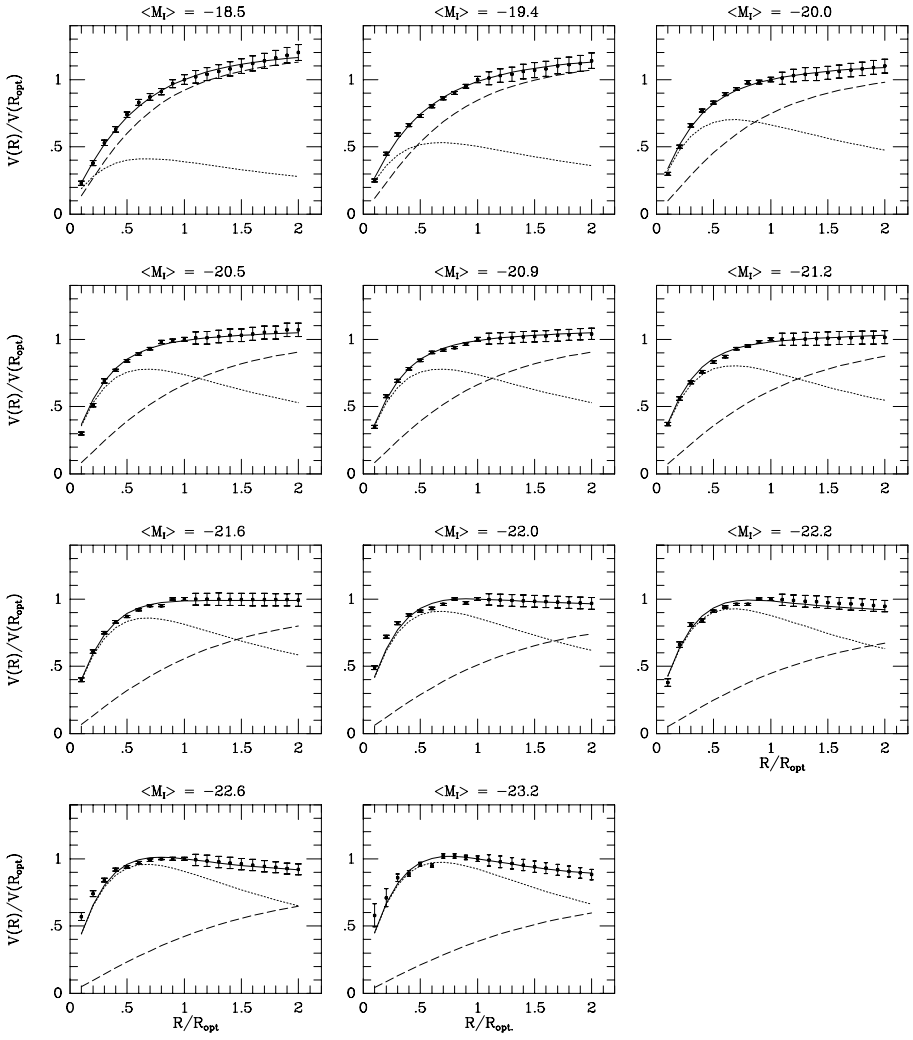

Figure 1.1: Average rotation curves for an ensemble of spiral galaxies. Each bin contains galaxies in the luminosity range indicated by the corresponding plot's title (in absolute magnitudes). The mean luminosities of the lowest and highest bins differ by a factor of ~ 75. There are $50 - 100$ galaxies in each bin (except for the highest luminosity bin). The lines indicate the contributions of various components. Solid: combined disk and halo contribution. Dotted: disk contribution. Dashed: halo contribution. The individual rotation curve contributions add in quadrature. Note that these are averages, not 11 separate galaxies, as was implied in Figure 1.1 of [8]. Figure taken from [4].

of the disk. Taking $v(r)$ constant yields $M(r) \propto r$. However, the disk surface brightness drops exponentially with r, following the exponential disk profile of Equation 1.24. If luminous matter were the only component of such galaxies, $M(r)$ would follow

$$M(r) = 2\pi\sigma_0 R_D^2 \left(1 - e^{-r/R_D} - \frac{r}{R_D} e^{-r/R_D}\right) \tag{1.26}$$

where σ_0 is the central areal mass density of the exponential disk. Clearly, $M(r)$ asymptotes to a constant. If an additional, spherically symmetric halo component is assumed, then $M(r) \propto r$ can be produced if the volume mass density of the spherical component obeys $\rho(r) \propto 1/r^2$ at large r. (Sphericity is assumed for simplicity. In general, a disk or lone, severely flattened ellipsoid is not gravitationally stable.) Clearly, the flat rotation curves observed in spiral galaxies imply the existence of a dark-matter component.

A density profile of the form needed is provided by an "isothermal" halo. An isothermal halo is given by assuming a system of particles of mass m interacting gravitationally that is in thermal equilibrium at temperature T. The phase-space distribution function is

$$f(\vec{r}, \vec{v}) = \frac{\rho_0}{(2\pi kT/m)^{3/2}} \exp\left(-\frac{\frac{1}{2}mv^2 + m\phi(\vec{r})}{kT}\right) \tag{1.27}$$

where ρ_0 is the central density and $\phi(\vec{r})$ is the gravitational potential at $\vec{r}$ determined by the system itself. This form is just a thermal Boltzmann distribution with the gravitational potential energy included. For an isothermal halo containing particles of different masses, it is not the temperature, but rather the velocity dispersion $\sigma^2 = kT/m$, that is common to all the particles. The simplest such distribution is isotropic, so a function of r and v alone. It is nontrivial to solve for the density distribution because of the nonlinear dependence of the phase-space distribution on itself through ϕ; this is done in [9]. Under the assumption that the density distribution is softened so it is flat at small radii (ρ_0 is finite), the density distribution is well approximated by the form

$$\rho(r) \propto \frac{1}{r^2 + a^2} \tag{1.28}$$

where a is the core radius. This provides the desired asymptotic density distribution while maintaining finite density at small radii so galactic bulges and disks dominate the inner regions.

Salucci and Persic find it possible to fit the observed rotation curves with such a form, combined with a simple exponential disk [4]. The forms they use, in terms of the disk and halo contributions to the rotation curve and written in terms of the dimensionless radius $x = r/R_{opt}$, are:

$$\text{disk}: \quad \left[\frac{v_d(x)}{v(R_{opt})}\right]^2 = \beta f(x) \tag{1.29}$$

$$\text{halo}: \quad \left[\frac{v_h(x)}{v(R_{opt})}\right]^2 = (1-\beta)\frac{x^2(1+a^2)}{x^2+a^2}$$

where $f(x)$ is a form that approximates an exponential disk and

$$\begin{aligned} \beta &= 0.72 + 0.44 \log\left(\frac{L}{L_*}\right) \\ a &= 1.5\left(\frac{L}{L_*}\right)^{1/5} \end{aligned} \tag{1.30}$$

These forms fit the data to 1% in spite of their simplicity. With these models, it is possible to calculate global parameters of the halo, such as the mass enclosed. The nominal "edge" of the halo is taken to be R_{200}, the radius at which the halo density is 200 times the background density of the universe. The motivation for this number is that a number in the range 100 to 200 is usually taken to be the value at which an overdensity irreversibly gravitationally collapses. This is likely an underestimate of the true edge because an overdensity accretes additional material from the surrounding, lower-density region. However, with the naive definition, Salucci and Persic find

$$R_{200} \approx 250 \left(\frac{L}{L_*}\right)^{0.2} \text{ kpc} \tag{1.31}$$

$$M_{200} \approx 2 \times 10^{12} \sqrt{\frac{L}{L_*}}\, M_\odot \tag{1.32}$$

Finally, the interesting number, the mass-to-light ratio for such a system:

$$\Upsilon_{spiral} = \frac{M_{200}}{L} \approx 80 \sqrt{\frac{L_*}{L}} \tag{1.33}$$

Extrapolating $\Upsilon(L_*)$ to the universe as a whole gives $\Omega_{m,spiral} = 0.029h^{-1}$. There is a clear need for dark matter in spirals, but this dark matter does not suffice to close the universe.

Note that the above formula implies less luminous galaxies have larger mass-to-light ratios. This is also implied by Figure 1.1, where the rotation curves for less luminous galaxies continue to rise at large radii. In contrast, more luminous galaxies exhibit rotation curves that begin to fall at large radii, indicating the exponential disk component is important in determining the gravitational potential.

Elliptical Galaxies

The existence of dark matter is also inferred in elliptical galaxies. An elliptical galaxy is well approximated as a gas of stars in thermal equilibrium. The velocities of the stars exhibit no coherent directionality of the kind seen in spirals. Ellipticals are anisotropic, exhibiting triaxial ellipsoidal shapes. These galaxies appear to be gravitationally relaxed — *i.e.*, neither collapsing nor expanding. Ellipticals possess halos of ionized, X-ray-emitting gas. The X-ray emission arises from free-free scattering and thermal bremsstrahlung,

$$\begin{aligned} e^- + e^- &\rightarrow e^- + e^- + \gamma \\ e^- + p &\rightarrow e^- + p + \gamma \end{aligned}$$

A relation between the gas density distribution $\rho_g(r)$, the total mass enclosed at r $M(r)$, and the observed X-ray temperature T may be derived under the assumption of hydrostatic equilibrium. This consists of assuming that gas pressure supports the gas against gravitational collapse. Explicitly,

$$\frac{dp_g(r)}{dr} = -G\frac{M(r)\rho_g(r)}{r^2} \tag{1.34}$$

where $p_g(r)$ is the gas pressure. The above relation can be demonstrated by considering a slab of gas of area A and thickness dr. The net force on the slab from adjacent slabs is $F = A(p_g(r +$

$dr) - p_g(r)) = Adrdp_g/dr$. The gravitational force on the slab is $-GM(r)\rho_g(r)Adr/r^2$. Cancelling the slab volume Adr gives the above relation. Under the assumption of thermal equilibrium, the gas pressure and temperature are related by the ideal gas law, $p_g = n_g k_B T = \rho_g k_B T/\mu m_p$ where n_g is the number density of ions, m_p is the mass of the proton, and μ is the average atomic mass of the ions (μm_p gives the average ion mass). Thermal equilibrium implies there is a single T for the entire gaseous halo. Thus,

$$\frac{d}{dr}\left(\frac{\rho_g(r)kT}{\mu m_p}\right) = -G\frac{M(r)}{r^2}\rho_g(r) \tag{1.35}$$

The gas contributes negligibly to the overall mass, so $M(r)$ is independently specified.

Under the assumption of no dark matter, $M(r)$ is determined by the stars alone. These systems appear to be virialized, so the stars should follow a density distribution consistent with thermal equilibrium. If the stars provide the gravitational well containing the X-ray-emitting gas, then the two should have the same velocity dispersion:

$$\sigma_*^2 = \frac{k_B T}{\mu m_p} \tag{1.36}$$

where σ_* is the velocity dispersion of the stars. In general, the density profiles of the stars and gas should follow thermal distributions parameterized by σ_*^2 and T and thus the assumption of thermalization may be tested. However, it is difficult to make position-resolved X-ray luminosity measurements for objects as small as galaxies unless they are very nearby. The X-ray luminosity obeys $j \propto n_e^2\sqrt{T}$ because X-ray emission arises from a two-body scattering process (the dependence on $\sqrt{T}$ is nontrivial — I do not derive it here). The X-ray luminosity therefore falls off quickly with radius. Testing the density distributions directly is impractical.

Fortunately, the values of σ_* and T at the centers of ellipticals are well measured for tens of galaxies [5]. Thus, one can check if Equation 1.36 is satisfied. It is observed that it is not. First, the scaling relation observed is $T \propto \sigma_*^{1.45}$ rather $T \propto \sigma_*^2$. Second, it is observed that $\beta_{spec} \equiv \mu m_p \sigma_*^2/k_B T \approx 0.5$ (note that this β is unrelated to the one defined in Equation 1.30). That is, the stars have about half the velocity dispersion that they should have. The data illustrating these two points are shown in Figures 1.2 and 1.3.

Loewenstein and White show that these observations are best accomodated when dark-matter halos dominate these galaxies [6]. The dark-matter halo provides the gravitational well needed to give the observed X-ray temperature. This is implied by the hydrostatic equilibrium argument above. Their explanation for the disparity between σ_* and its expected value, and for the scaling of β_{spec} with galaxy luminosity, is as follows. At low galaxy masses, the stellar component is both nondominant and is contained in the core radius of the dark halo. The model dark-matter halo is flat enough in the core that the stellar velocity dispersion is characteristic of a small, stellar-dominated component. The halo varies as expected on larger scales. The temperature of the more extended gas is determined by the full structure of the dark halo — infall into the halo potential well heats up the gas. Two effects occur as the galaxy mass is increased. First, the temperatures of all the components increase because the potential well becomes deeper. Second, the stellar component becomes extended and massive enough that it plays an important role in defining the potential. Therefore, the stellar and dark-matter halo velocity dispersions converge at

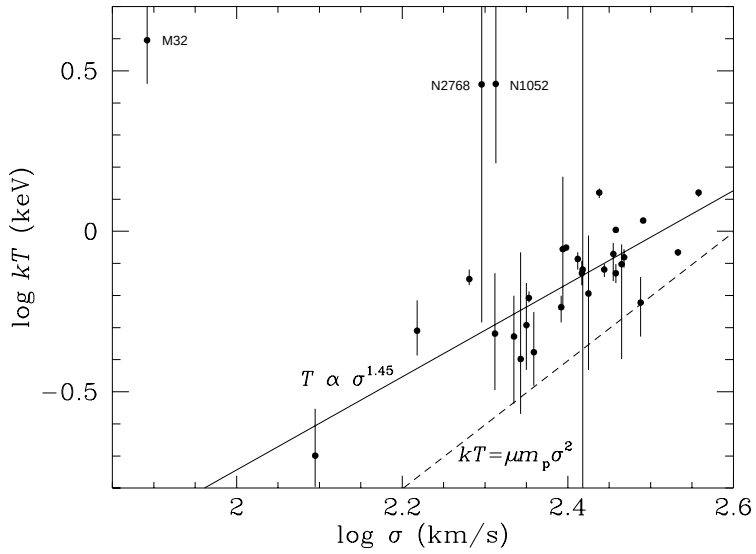

Figure 1.2: X-ray temperature k_BT vs. stellar velocity dispersion σ_* for a set of X-ray-emitting elliptical galaxies. The solid line indicates the best fit exponent, 1.45, and the dashed line indicates the expectation for $T \propto \sigma_*^2$. Figure taken from [5].

high masses. The gas temperature senses the dark-matter halo velocity dispersion, so σ_* and T converge. The increase of β_{spec} with galaxy mass is consistent with this convergence.

Using the mass models that give these results, Loewenstein and White estimate Υ within a radius of $6r_e$ for the systems they analyze, yielding $\Upsilon_{6r_e} \approx 31h$. (Recall that r_e is the half-light radius.) This is about a factor of 2.5 lower than Υ estimated for spirals, but Υ is underestimated because $6r_e \ll R_{200}$ for these systems. The authors hesitate to calculate Υ at R_{200} because their models are far less constrained than the spiral models studied by Salucci and Persic.

1.3.2 Dark Matter in Galaxy Clusters

The dark matter in the space between galaxies can also be studied by its gravitational effects. Galaxies form gravitationally bound ensembles, called groups (few to tens of galaxies) and clusters (tens to hundreds of galaxies). The nearest cluster is the Virgo cluster, at a distance of about 16 Mpc. The Local Group, which includes the Milky Way and the Andromeda galaxy, is falling in toward the Virgo cluster.

There are four well-known methods for measuring the dark-matter content in galaxy clusters. Virial, X-ray-emission, and gravitational-lensing methods give $\Upsilon \approx 250h$ and $\Omega_m \approx 0.09$ while baryon-fraction measurements yield $\Omega_m \approx 0.35$. As stated earlier, these non-global methods only give lower limits on Ω_m, so the discrepancy between these values is not an indication of inconsistency.

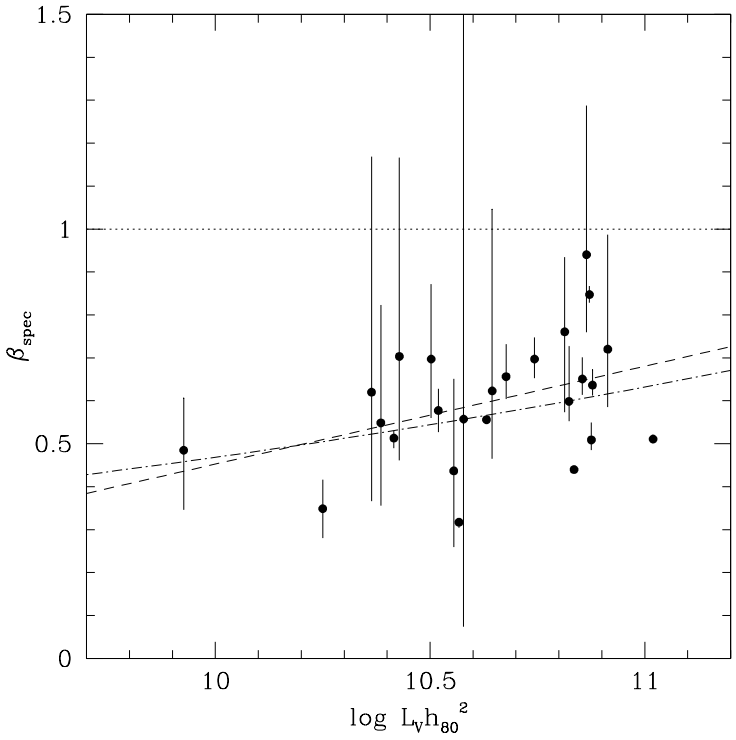

Figure 1.3: $\beta_{spec} = \mu m_p \sigma_*^2 / k_B T$ vs. galaxy luminosity (in the V band, 550 nm). Figure taken from [6].

Virial Theorem

The galaxies in a cluster can be used as test particles to probe the gravitational potential. Under the assumption that the cluster is gravitationally relaxed, the virial theorem holds:

$$\sum_i \frac{1}{2} m_i v_i^2 = G \sum_{i \neq j} \frac{m_i m_j}{r_{ij}} \tag{1.37}$$

where i and j run over all the particles. A virial radius and mass are defined to be a composite of the right side by

$$\frac{1}{2} \langle m \rangle \langle v^2 \rangle = G \frac{M_{vir}}{R_{vir}} \langle m \rangle \tag{1.38}$$

where $\langle v^2 \rangle = \sigma^2$ is the velocity dispersion of the galaxies, $\langle m \rangle$ is the velocity-squared-weighted mean galaxy mass, and R_{vir} is the virial radius, defined circularly as the radius inside of which the assumption of virialization holds. M_{vir} is the cluster virial mass, the mass inside the virial radius. In a simple spherical collapse model for a $\Omega_m = 1$ universe, it can be shown that a system is virialized when the mass density satisfies

$$\rho_{vir} = 18\pi^2 \rho_0 \tag{1.39}$$

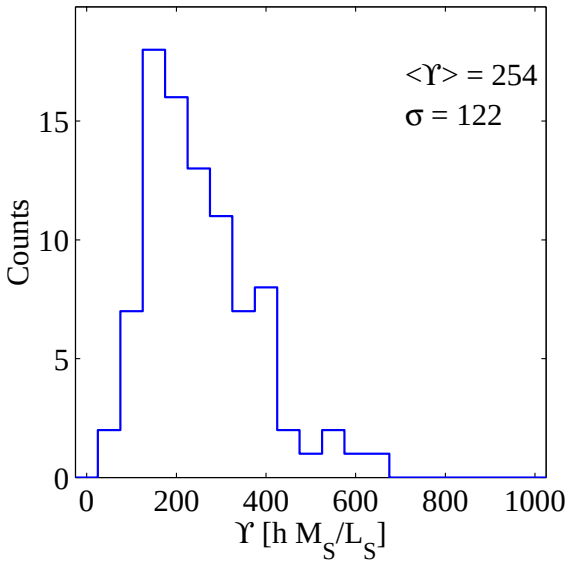

Figure 1.4: Histogram of mass-to-light ratios determined by Girardi *et al.* [11].

where ρ_0 is the background energy density of the cosmology [10]. Cosmologies with Ω_m within an order of magnitude of 1 have similar relationships, with the prefactor modified by factors of a few. Since the virial mass is not known *a priori*, the virial radius must be calculated iteratively: the velocity dispersion is calculated inside an assumed virial radius, the virial mass is determined, Equation 1.38 is checked, and a new virial mass assumed. If there is a radius to which the procedure converges, then presumably the procedure is valid.

Girardi *et al.* apply this procedure to a sample of approximately 100 clusters observed by various surveys [10]. The virial mass calculated in the above way is corrected for the fact that the observations do not include galaxies all the way out the edge of the cluster, the "edge" being defined as the radius at which a test particle falls into the cluster rather than expanding away with the universe, $\sim (5-10)R_{vir}$ [11]. These galaxies have lower velocities (they are in less tightly bound orbits), so neglecting them artificially increases σ and the mass estimate. The masses calculated in this way compare well with masses calculated from the X-ray luminosity (see below).

In another paper, Girardi *et al.* combine these mass estimates with measurements of the optical luminosities of the galaxies to calculate mass-to-light ratios [11]. A histogram of these mass-to-light ratios is shown in Figure 1.4. The distribution appears well behaved, though the large dispersion suggests there is physics in the variation. Regardless, the mean mass-to-light ratio is $\Upsilon \approx 250h$, which extrapolates to $\Omega_m \approx 0.09$. This is a lower limit because it contains only the mass within the virial radius; the turnaround radius is expected to be $5 - 10\,R_{vir}$.

X-ray Emission

As discussed in the context of elliptical galaxies, the gas temperature and density profile are determined by the total mass distribution under the assumption of hydrostatic equilibrium.

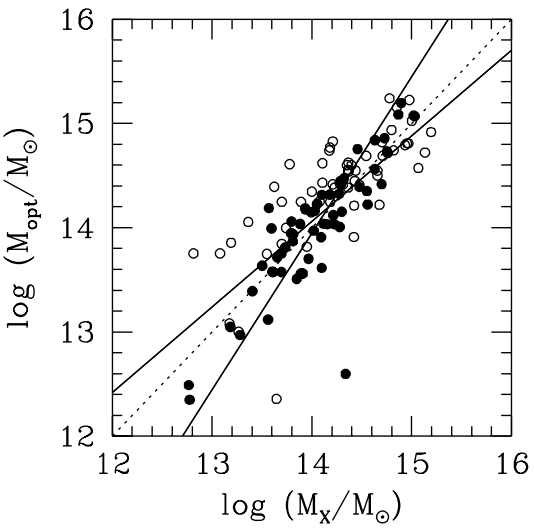

Figure 1.5: Virial masses vs. X-ray-inferred masses. Virial masses calculated from optical galaxy velocity dispersions in [10]. X-ray masses from a variety of sources; X-ray masses for filled dots from [12], discussed in the text. Girardi *et al.* attribute the lower dispersion for clusters with masses derived in [12] to the homogeneity of the analysis method and source of the X-ray data. Figure taken from [10].

Galaxy clusters are extended enough that the X-ray luminosity can be measured out to the virial radius; this is in contrast to elliptical galaxies, where only the central value is measured well in many galaxies. These cluster X-ray profiles can be fit to isothermal models and the cluster mass extracted. Then, using optical data to calculate the luminosity of the cluster, the mass-to-light ratio can be established.

A large fraction of the clusters studied by Girardi *et al.* have X-ray emission data available. White, Jones, and Forman analyzed archival data from approximately 200 clusters and calculated cluster masses [12]. It should be noted that the analysis they used requires the galaxy velocity dispersion as an input in order to define the gravitational potential, so it is not completely independent of the above virial theorem analysis. However, it does provide a consistency check — if the virial masses estimated above are correct, then it should also be possible to construct X-ray-emission temperatures and profiles consistent with the data that give similar masses. The comparison of this method and the velocity dispersion method for the Girardi *et al.* clusters is summarized in Figure 1.5. There is good agreement, supporting the mass-to-light ratios calculated by Girardi *et al.* The ROSAT All-Sky Survey will enlarge the sample of clusters with well-determined masses by a factor of a few. These data are being made available now.

Gravitational Lensing

All the methods discussed up to this point rely on gravitational dynamics to determine galaxy or cluster masses. These arguments require assumptions of virialization and hydrostatic

equilibrium, which are never exactly true for clusters and are sometimes badly violated. Only in the last few years have a sufficiently large number of cluster masses been measured by velocity-dispersion and X-ray-emission methods to really have confidence that one has a picture of the behavior of an ensemble. Gravitational lensing provides a method for estimating cluster masses that is independent of these assumptions.

In the standard application, a foreground cluster at redshift of about 0.5 gravitationally lenses background galaxies at $z \gtrsim 1$. In the case of strong lensing, the images of the background galaxies are obviously affected, with multiple images appearing that may be distorted into arcs or rings. In weak lensing, the background galaxies' ellipticities are slightly modified and their apparent magnitudes slightly increased in a way that is linearly related to the gravitational potential in the cluster along the lines of sight to the background galaxies.

The formalism for strong lensing is too complex to discuss here because the phenomena studied are highly nonlinear. Each such system must be studied in detail in order to reconstruct the mass distribution of the lens. Such systems are also few and far between; while lensing observations determine the mass distributions of these clusters exquisitely well, they require case-by-case analysis and do not promise to be applicable to large numbers of clusters quickly.

Weak lensing is much simpler. Since the distortions are so small, they cause small perturbations of the ellipticity vectors and magnitudes of background galaxies. This distortion is coherent across the lens face, pushing the ellipticity vectors of background galaxies to line up on the tangents to circles around mass concentrations. In the weak limit, the ellipticity and magnification fields depend linearly on the projected mass density. An upcoming review by Bartelmann and Schneider lists the results found for the 12 clusters that have been mapped in this way so far [13]. These analyses confirm, for the most part, the results of virial and X-ray-emission analyses. Two or three clusters exhibit severe discrepancies that appear to remain after redoing the lensing measurement and analysis. In such cases, weak lensing usually estimates a higher mass by a factor of 2 to 3. Such clusters may indeed be out of equilibrium. The main conclusion, however, is that weak lensing verifies the virialization assumption for the majority of galaxy clusters.

Cluster Baryon Fraction

If the ratio of the mass in baryons to the total mass of a galaxy cluster can be established, this fraction can be extrapolated to the universe as a whole using the constraint on the baryon density parameter, Ω_b, from Big Bang nucleosynthesis (discussed later in this chapter). Because they are so large, clusters are expected to be a "fair sample" of the universe — the baryon fraction of clusters is not expected to be enhanced above the background value. The extrapolation to the universe as a whole is through the baryons, rather than the light, but otherwise the idea is the same. Even if the fair-sample hypothesis is violated, the baryon fraction is likely to be enhanced in clusters, so values of Ω_m estimated in this way are lower limits.

Two types of measurements give direct access to the baryon content of intracluster plasma. The mass of this gas is large compared to the baryonic mass in galaxies, containing $\sim 90\%$ of the baryons. As discussed above, the X-ray-luminosity profile directly measures the electron density and thus the proton density. The Sunyaev-Zeldovich effect is also used. The SZ effect is the inverse Compton scattering of the CBR by hot electrons in the intracluster plasma. Depending on the wavelength at which the measurement is done, the CBR temperature appears to either be decremented or incremented. The fractional change in the CBR temperature is simply related to

the line integral of the electron density-temperature product through the cluster

$$\frac{\Delta T_{RJ}}{T} = -2\frac{k_B \sigma_T}{m_e c^2} \int dl\, n_e(l)\, T(l) \tag{1.40}$$

where σ_T is the Thomson cross section. The Rayleigh-Jeans (long-wavelength) limit has been taken to simplify the prefactor. Essentially, this is the line integral of the electron pressure through the cluster. In practice, there are two ways to determine a baryon fraction: compare the X-ray luminosity or SZ decrement to the cluster mass determined by a virial method or by lensing; or construct a self-consistent cluster model using the X-ray and SZ data. Both of these are usually applied to any given cluster for redundancy.

The group led by Carlstrom has used the BIMA and OVRO arrays to make interferometric observations at $\lambda = 1$ cm of the SZ effect [14]. Combined with X-ray data, measurements of 18 clusters give baryon fractions in the range $0.06\,h^{-1}$ to $0.08\,h^{-1}$. An ensemble of 22 clusters was observed by Myers and others using the OVRO 5.5-m telescope, also at 1 cm, and gives a mean baryon fraction of $0.07\,h^{-3/2}$ [15]. Taking a nominal value of $f_B = 0.07h^{-3/2}$, $h = 0.67$, and using the BBN constraint $0.018\,h^{-2} < \Omega_b < 0.022\,h^{-2}$ [3], the allowed range for Ω_m is 0.31 to 0.38.

1.3.3 The Global Value of the Energy Density

There are a powerful set of tests that can be used to determine the global characteristics of the visible universe. In recent years, these tests, in the form of the angular scale of the first Doppler peak of the anisotropy power spectrum of the CBR and the Type Ia supernova magnitude-redshift relation, have come to maturity and are giving significant results on the global parameters of the observable universe. These tests presently give $\Omega \approx 1$ with $\Omega_m \approx 0.3$.

Standard Candles, Standard Rulers, and Distances

The supernova magnitude-redshift test belongs to the class of "standard-candle" tests. In these tests, the flux of light received from a distant object is compared to its known, intrinsic luminosity. In a flat, non-expanding universe, these two are related by the $F = L/4\pi r^2$, where r is the physical distance to the object. In curved universes, this relationship is modified by the non-Euclidean geometry. Consider a light detector on the earth, with physical width d. In a flat universe, two light rays emitted by a point source that hit the edges of the detector are separated by angle $\theta = d/r$. In an open universe, they must be separated by a smaller angle, $\theta < d/r$, in order to hit the edges of the detector because light rays diverge in an open universe. (Formally, the proper distance between two light rays emitted with separation θ increases more quickly with proper distance along the path of the light rays than in a flat universe.) Conversely, in a closed universe, the initial separation angle must be larger, $\theta > d/r$, because light rays converge in a closed universe. Therefore, in open and closed universes, the detector subtends less and more, respectively, solid angle as seen from the source as compared to a flat universe. The emitted luminosity is independent of the type of universe, so the detector sees a smaller and larger flux, respectively, in open and closed universes as compared to a flat universe. Put another way, in the standard-candle test, one places a ruler at one's present location and asks what angle it subtends at the emitting object, whose redshift is known. The ruler is one's telescope mirror. The angle is the angle subtended by

the telescope; it is measured by comparing the flux of light received by the telescope to the known luminosity emitted into all angles.

The measurement of the angular scale of the first Doppler peak of the CBR anisotropy power spectrum belongs to the related, "standard-ruler," set of tests. These are essentially the time reverse of the above. The ruler is an object of known physical size at redshift z. The angle is the angle it subtends at the observer's position.

Let $w(z, \theta_0)$ denote the physical length at the observer today that subtends an angle θ_0 at the emitter at redshift z. The time-reversal argument above ensures that the same function serves to give the angle subtended today at the observer by a physical length at z. Let $r_*(z)$ denote the comoving distance that corresponds to redshift z. That is, r_* is defined so that the light reaching the origin today from objects at comoving distance r_* was emitted at the time corresponding to redshift z. The comoving length corresponding to the physical length w today is w/a_0. The angle subtended by this comoving length at the emitter is independent of the expansion because it is a comoving length. The comoving distance to this emitter is r_* by definition, so for $\theta_0 \ll 1$,

$$\begin{aligned} w(z, \theta_0) &= a_0 r_*(z) \theta_0 \\ &\equiv d_P(z) \theta_0 \end{aligned} \tag{1.41}$$

where $d_P(z)$ is the *proper distance* to an object at redshift z. That is, the physical length today that subtends an angle θ_0 at an object at redshift z is simply related to the present physical distance to that object by the above equation.

The standard-candle test can be written in terms of $d_P(z)$. The power emitted by the source at redshift z into a cone of opening angle θ_0 is

$$P_e = L \frac{\pi \left[\frac{1}{2}\theta_0\right]^2}{4\pi} \tag{1.42}$$

for $\theta_0 \ll 1$. The detected flux is

$$F = \frac{P_e}{(1+z)^2} \frac{1}{\pi \left[\frac{1}{2} w(z, \theta_0)\right]^2} \tag{1.43}$$

The prefactor of $(1+z)^{-2}$ arises because of the expansion: the energies of individual photons are redshifted by $(1+z)^{-1}$, and the rate at which photons arrive is scaled by $(1+z)^{-1}$. This latter effect arises because the source is expanded away as it emits photons: two consecutive photons are received with a larger time difference than that with which they are emitted because the second photon has to make up the distance that the source "moves." The second factor in Equation 1.43 is the geometrical one: the function $w(z, \theta_0)$ gives the physical length today subtended by the angle θ_0 at an emitter whose comoving distance corresponds to redshift z. Again, the physical length w is measured *today*; z is only meant to parameterize the distance of the emitter and the epoch of emission. The flux-luminosity relation resulting from Equations 1.42 and 1.43, in terms of $d_p(z)$, is

$$F = \frac{L}{4\pi (1+z)^2 d_P^2(z)} \tag{1.44}$$

For convenience, the *luminosity distance* is defined by $d_L(z) = (1+z) d_P(z)$, giving

$$F = \frac{L}{4\pi d_L^2(z)} \tag{1.45}$$

The situation is slightly different for the standard-ruler test. In this case, the interesting observable is the angle subtended today by an object at redshift z that has physical size D at that redshift. Therefore, D is analogous to $w(z,\theta_0)/(1+z)$ in the above discussion. This yields

$$\begin{aligned} D &= \frac{a_0 r_*(z)\theta_0}{1+z} \\ &= a(z) r_*(z)\theta_0 \\ &\equiv d_A(z)\theta_0 \end{aligned} \tag{1.46}$$

The second line is what one would have guessed from first principles. The last line defines the *angular diameter distance*. The relationships of the three distances are summarized as follows:

$$a_0 r_*(z) = d_P(z) = (1+z) d_A(z) = (1+z)^{-1} d_L(z) \tag{1.47}$$

$d_P(z)$ gives the present distance to an object at redshift z. Alternatively, it relates the presently observed angular size of an object at redshift z to the physical size it would have now if it expanded with the universe. $d_A(z)$ relates the observed present angular size of an object at redshift z to its physical size when it emitted the light that is observed today, correcting $d_P(z)$ downward for expansion. $d_L(z)$ relates the intrinsic luminosity of an object to the flux observed, correcting $d_P(z)$ for redshifts due to expansion.

All of the above distances depend on $r_*(z)$, the coordinate distance of an object that is observed as having redshift z. $r_*(z)$ can be calculated for an arbitrary cosmology. Light satisfies $ds^2 = 0$, so Equation 1.3 implies

$$\frac{a_0 dr_*}{\sqrt{1-(r_*/R)^2}} = \frac{a_0}{a(t)} dt \tag{1.48}$$

Now, $dt/a = da/(ada/dt) = -dz/(a_0 H)$. So

$$\frac{a_0 dr_*}{\sqrt{1-(r_*/R)^2}} = -\frac{dz}{H(z)} \tag{1.49}$$

Integrating and calculating $d_P(z)$ yields

$$d_P(z) = \begin{cases} \left(H_0\sqrt{|1-\Omega_m-\Omega_\Lambda|}\right)^{-1} \sinh\left(\sqrt{|1-\Omega_m-\Omega_\Lambda|}\int_0^z \frac{dz'}{E(z')}\right) & \Omega_m+\Omega_\Lambda < 1 \\ H_0^{-1}\int_0^z \frac{dz'}{E(z')} & \Omega_m+\Omega_\Lambda = 1 \\ \left(H_0\sqrt{|1-\Omega_m-\Omega_\Lambda|}\right)^{-1} \sin\left(\sqrt{|1-\Omega_m-\Omega_\Lambda|}\int_0^z \frac{dz'}{E(z')}\right) & \Omega_m+\Omega_\Lambda > 1 \end{cases} \tag{1.50}$$

where

$$E(z) = \sqrt{H(z)/H_0} = \sqrt{(1+z)^2(1+\Omega_m z) - z(2+z)\Omega_\Lambda} \tag{1.51}$$

is defined by the Friedmann equation, Equation 1.4, and is dependent only on z and the present-day Ω_m and Ω_Λ. The definition of Ω_k, Equation 1.7, and the present-day version of the Friedmann equation, Equation 1.9 have been used.

The Standard-Ruler Test: the CBR Anisotropy Power Spectrum First Doppler Peak

The standard-ruler test is applied by measurements of the position of the first Doppler peak of the angular power spectrum of the fluctuations of the CBR temperature. The standard ruler is the horizon size at the redshift of decoupling, which defines the size of the largest wavelength oscillation of the photon-baryon plasma before the photons decoupled and the oscillations stopped at $z_{dec} \approx 1100$. The power spectrum of the temperature anisotropy displays a peak at the spherical-harmonic multipole number l_p corresponding to the angular scale θ_p that is subtended by one-half of this length. (A factor of 1/2 arises because, although the wavelength of the mode with the largest power is the above horizon size, the largest amplitude component of the anisotropy power spectrum is at the length scale corresponding to the distance between a peak and trough because this separation gives the largest temperature difference.) I discuss the CBR in some detail later; for now, this standard ruler is assumed. The horizon size at the time of decoupling is dependent on Ω_m and Ω_r, the present-day matter and radiation density parameters. In principle, it is necessary to include Ω_r because, until a redshift of $z = 23700\,\Omega_m h^2$, radiation was dominant and modified the expansion rate. For simplicity, this complication is ignored here; it simply introduces an additional dependence on Ω_r that somewhat modifies the results. Ω_Λ does not enter because it played a negligible role at such high redshifts.

The horizon size at a time t and redshift z is defined by

$$D(t) = a(t) \int_0^t \frac{dt'}{a(t')} \tag{1.52}$$

given by integrating the distance traveled by a photon, with $c = 1$. Using the same tricks used for calculating $d_P(z)$, the expression takes the form:

$$D(z) = \frac{1}{H_0(1+z)} \int_z^\infty \frac{dz'}{E(z')} \tag{1.53}$$

As noted above, the Ω_r and Ω_Λ terms in $E(z)$ are neglected. This leaves $E(z) = \sqrt{\Omega_m(1+z)^3}$, giving for the integral

$$D_{dec} \approx \frac{2}{H_0\sqrt{\Omega_m}}(1+z_{dec})^{-3/2} \tag{1.54}$$

Thus, the horizon size at decoupling, the standard ruler, has simple dependences on H_0 and Ω_m.

The angular diameter distance has a more complicated dependence, shown in Equation 1.50, because Ω_Λ cannot be neglected. Note, however, that H_0 cancels between D_{dec} and $d_A(z_{dec})$, so the angle subtended at an observer today by the length D_{dec} (or, actually, recalling the factor of 1/2, $D_{dec}/2$) at z_{dec} is independent of H_0. Explicitly, the angular scale of the first peak is

$$\theta_p = \frac{\frac{1}{2}D_{dec}}{d_A(z_{dec})} = \frac{1}{\sqrt{\Omega_m(1+z_{dec})^3}} \frac{1}{H_0 d_A(z_{dec})} \tag{1.55}$$

where the quantity $H_0 d_A(z_{dec})$ is independent of H_0. In general, the dependence of the angular scale of the first peak on Ω_m and Ω_Λ is not obvious because of $d_A(z)$. In the limit $\Omega_\Lambda = 0$, $d_A(z)$ can be calculated exactly (the integral is nontrivial):

$$d_A(z) = \frac{2}{H_0\Omega_m^2} \frac{1}{(1+z)^2} \left[\Omega_m z + (\Omega_m - 2)(\sqrt{1+\Omega_m z} - 1)\right] \tag{1.56}$$

(see, for example, [16]). Using this result and making use of $z_{dec} \gg 1$, one finds

$$\begin{aligned} \theta_p &= 0.015\sqrt{\Omega_m}\left(\frac{z_{dec}}{1100}\right)^{-1/2} \text{ rad} \\ &= 0.86^\circ\sqrt{\Omega_m}\left(\frac{z_{dec}}{1100}\right)^{-1/2} \end{aligned} \tag{1.57}$$

This result is the source of the frequently made statement that the first Doppler peak probes the global geometry of the universe. Note that the above derivation only holds in the $\Omega_\Lambda = 0$ case. It turns out that the result $\theta_p \propto \sqrt{\Omega}$ holds in general. Analytically, it is not trivial to show this because of the complication of introducing Ω_Λ into the integrand.

However, it is not difficult to understand why all flat universes give the same θ_p, independent of Ω_m. The integral that gives $d_A(z)$ is the integral of $1/E(z)$ from 0 to z. The integral that gives D_{dec} is the same, except with limits from z to ∞. The first integral receives only a small contribution from $z \ll 1$, where Ω_Λ has an effect. Thus, both integrals are dominated by their Ω_m behavior, which produces a prefactor $1/\sqrt{\Omega_m}$. Since both D_{dec} and d_A have this prefactor, it cancels, leaving no dependence of θ_p on Ω_m. Physically, this can be seen in the following way. D_{dec} is determined by the behavior of $H(z)$ at high z. In a $\Omega_m = 1$ universe, $H(z)$ is larger at high z than in a flat, low-Ω_m universe because the larger Ω_m has decelerated the expansion more in the former case. Thus, D_{dec} is *smaller* in a $\Omega_m = 1$ universe than in a flat, low-Ω_m universe: the expansion rate at high z was slower in the latter case, so the horizon was larger. However, the angular diameter distance is similarly affected: since expansion in a low-Ω_m universe is slower, light can travel further in the same redshift interval, and thus d_P, d_A, and d_L are all bigger in a low-Ω_m universe. In brief, the horizon size at decoupling is physically bigger in a flat, low-Ω_m universe, but decoupling is further away from us in such a universe, so, to first order, the angular size is unchanged.

The strong dependence on the geometry arises from the different functional forms that apply depending on the curvature. The integral $\int_0^z dz'/E(z')$ is not very different for open and flat universes with the same Ω_m, because, as argued above, Ω_Λ has little effect on it. For example, for $\Omega_m = 0.3$, the integral only differs by about 15% for $z = z_{dec}$. However, this integral is then multipied by $\sqrt{|1-\Omega|}$, the sine or hyperbolic sine taken, and then divided by $\sqrt{|1-\Omega|}$. In general,

$$d_A(z) \propto \frac{1}{z\sqrt{\Omega_m \Omega}} \tag{1.58}$$

This dependence yields $\theta_p \propto \sqrt{\Omega/z}$. This particular dependence is not trivial to see. The $E(z)$ function may be written in terms of Ω_m and Ω by using $\Omega_\Lambda = \Omega - \Omega_m$. The integration that yields Equation 1.56 can be generalized to this $E(z)$ and thus yield the above form.

The angular scale of the first peak has been measured with good precision only recently. In spherical-harmonic multipole number l, the position of the first peak obeys $l_p \sim \pi/\theta_p \simeq 200\Omega^{-1/2}$ for θ_p in radians. (Note that, in general, there is no simple exact relation between l_p and θ_p [17].) Figure 1.6 shows data from four recent experiments that determine the peak position: the Mobile Anisotropy Telescope (MAT), Boomerang, Python V, and Viper. MAT observes from the ground on the Atacama Desert in Chile, a very dry site at approximately 15000 feet. The TOCO97 data were taken during the 1997 austral winter using 30 and 40 GHz HEMT receivers [18]. The TOCO98 data were taken during the 1998 austral winter using 144 GHz SIS-mixer receivers [19].

Boomerang is a balloon-borne experiment that uses bolometric receivers at 90 and 150 GHz; the bolometers are operated at 300 mK. The data set shown was taken in 1997 during a single-night flight over Texas [20]. Python and Viper observe from the South Pole using HEMT receivers in the frequency range 37 to 45 GHz. The Python V data were taken during the austral summer 1996–1997 [21]; the Viper data were taken during the 1998 austral winter [22]. All the experiments scan the sky smoothly, rather than chopping, in order to achieve sensitivity to a large range in l. Individually, the four experiments suggest the first peak; together, they provide conclusive evidence and constrain its position to be at $l \approx 200$. The TOCO and Boomerang data, which best constrain the peak, agree well on the position. However, the two experiments exhibit $\sim 2\sigma$ disagreement on the amplitude. The fact that there is disagreement on the amplitude while agreement on the position is not surprising; the normalization of the peak is dependent on absolute calibration of the temperature sensitivity of the two experiments, which is difficult to do precisely, while the angular position depends only on the pointing reconstruction. To summarize, the data clearly indicate a peak at $l \approx 200$ giving $\Omega \approx 1$ and are strongly inconsistent with $\Omega < 0.4$ ($l > 320$) [23].

Very recently, the Boomerang and MAXIMA teams have released new results that give overwhelming evidence for the first peak, also shown in Figure 1.6. The Boomerang data are from a long-duration flight (11 days) over the South Pole [24]. The receiver was reconfigured to have two 90 GHz, six 150 GHz, four 240 GHz, and four 400 GHz channels; the first three frequencies are useful for CBR observations, the 400 GHz channel being used to check for contamination by thermal emission from dust. The results use only one 150 GHz channel, about half the sky coverage, and half of the time-series data. The MAXIMA data are from a summer, 1998, flight over Texas [25]. The MAXIMA experiment also uses bolometers, with eight receivers at 150 GHz and four each at 240 and 410 GHz. The MAXIMA bolometers are operated at 100 mK, rather than 300 mK, and so MAXIMA is able to achieve comparable sensitivity with a single night's data. The data presented are derived from three 150 GHz and one 240 GHz bolometer channels. Clearly, both data sets confirm the previous observations of a first peak at $l \approx 200$. In addition to improving the precision of the first-peak measurement, these two data sets are the first precise probe of higher-l behavior.

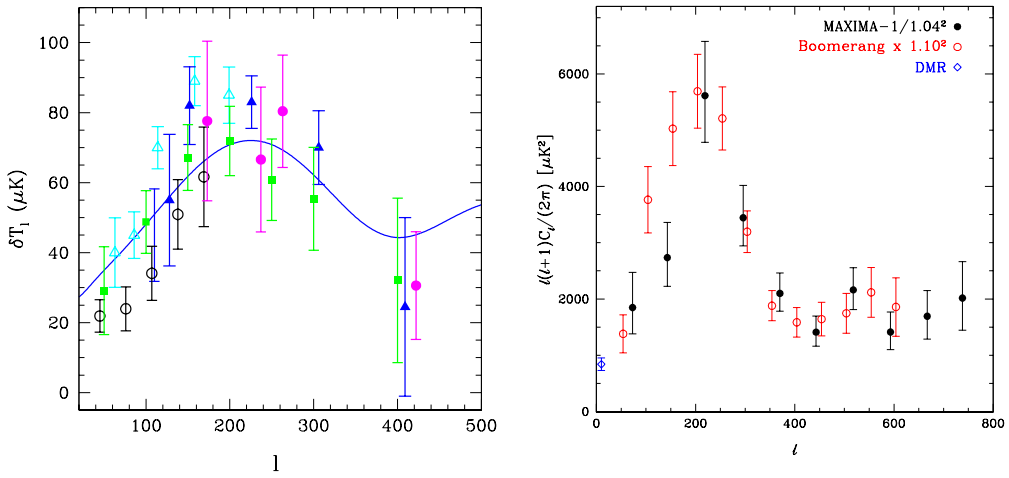

Figure 1.6: Left: Pre-May, 2000, CBR anisotropy data constraining the position of the first peak. The horizontal axis is spherical-harmonic multipole number l. The vertical axis is rms temperature fluctuation. Open triangles: MAT TOCO97 data [18]. Filled triangles: MAT TOCO98 data [19]. Filled squares: Boomerang 1997 data [20]. Open circles: Python V data [21]. Filled circles: Viper data [22]. Line: best fit flat model, $(\Omega_m, \Omega_\Lambda) = (0.31, 0.69)$ [23]. Figure taken from [26]. Right: Recent (May, 2000) Boomerang long-duration flight and MAXIMA-1 results. Note the modification of the vertical axis: the first plot is δT, the second is essentially δT^2. Note also the change in horizontal axis limits. The open circles indicate the Boomerang data, the filled circles the MAXIMA data. Correction factors have been applied to each experiment's temperature calibration to achieve consistency between the two power spectra; these corrections of 10% and 4%, respectively, are equal to each experiment's 1σ temperature-calibration errors. The relatively low amplitude of both data sets at the position of the second peak is now the subject of great interest, but is not important for this discussion. Figure taken from [25].

The Standard-Candle Test: High-Redshift Type Ia Supernovae

The standard-candle test is applied by the Supernova Cosmology Project and the High-z Supernova Search using Type Ia supernovae. Supernovae are massive explosions that occur when a star runs out of fuel and pressure is unable to support the core against gravity. The core collapse yields an explosion that liberates huge amounts of energy and violently ejects the star's outer layers, leaving a compact remnant. Type Ia supernovae appear to have fairly uniform luminosities, leading to their use as standard candles.

It is interesting to discuss some of the astrophysics that explains the uniformity of these objects. The original defining characteristics of Type Ia supernovae are empirical: the complete absence of hydrogen and helium lines and the presence of an absorption line near 610 nm due to a doublet (634.7 nm and 637.1 nm) of singly ionized silicon. The absence of hydrogen or helium implies that the object must be highly evolved — it must have burnt all its hydrogen and helium into silicon, carbon, oxygen or perhaps even heavier elements. I momentarily digress for a discussion of the late stages of stellar evolution, following standard texts such as [27]. Stars may be massive enough that fusion can continue until iron is produced, or may not and fusion stops earlier. In either case, once fusion stops, the star is left with a core composed of elements between C/O and Fe that no longer provides pressure support. Depending on the mass of the star, fusion may continue in shells surrounding the core. Regardless, the core contracts. Eventually, the density in the core becomes so high that electrons are no longer bound to single atoms. The result is a neutral plasma consisting of the ions and free electrons. Due to the high density, the electrons are a degenerate Fermi gas. Pauli exclusion prevents the electrons from losing momentum, and thus their velocity and hence pressure is maintained. However, there is a limit to the core mass that can be sustained by this "degeneracy pressure," the Chandrasekhar limit, 1.4 $M_\odot$. Larger cores require infinite pressure. Thus, if the original star is massive enough, the core exceeds the Chandrasekhar limit and a "core-collapse" supernova occurs, resulting in a neutron star or black hole. It should be noted that the formation of the degenerate core is a gradual process: a very light central core is continually fed as the fusion boundary moves outward. When the core exceeds the Chandrasekhar limit, the core collapses. A great deal of variation is possible among core-collapse supernovae: the core composition can vary anywhere from C/O to Fe, the outer layers of the star may still be burning, with different depths burning different elements, the total mass is free to vary, etc. Core-collapse supernovae cannot be Type Ia SNe because of this variability: one would expect a large fraction to have progenitors with hydrogen and helium envelopes, which should be visible in the supernova spectrum.

For low-mass stars, the Chandrasekhar limit is not exceeded and the core is stable. Such stars are called white dwarfs. They may continue to burn the small remaining amount of hydrogen or helium in their outer layers. The leading progenitor candidate for Type Ia SNe is a white dwarf with no hydrogen or helium left that accretes material from a normal companion star. A subclass of white dwarfs may have lost their hydrogen/helium envelopes during an earlier phase, resulting in a lone degenerate core below the Chandrasekhar mass. If the object is in a binary system with a nondegenerate star, accretion from the companion may allow the white dwarf to reach the Chandrasekhar mass, at which point the white dwarf becomes unstable and a supernova occurs. The important characteristics of this model are that the progenitor has no hydrogen or helium left and that the approach to the Chandrasekhar limit is very specific so the resulting supernovae are presumably quite uniform. The specificity of this pathway is consistent with the observation that

only a few percent of all supernovae are Type Ia.

Independent of the above justification, Type Ia SNe appear, empirically, to be standard candles. The Hamuy *et al.* Calan/Tololo survey discovered and measured redshifts for approximately 30 nearby Type Ia supernovae in the range $z \approx 0.01 - 0.1$ [28]. After applying a cut, defined in the literature prior to the observations, to remove "peculiar" supernovae based on their spectra, the resulting magnitude dispersion is $\sigma_B = 0.25$ mag and $\sigma_V = 0.23$ mag in the B and V optical bands, respectively. These correspond to a fractional luminosity variation between $10^{-0.4\sigma}$ to $10^{0.4\sigma}$, or about 80% to 125%. This level of uniformity is quite impressive for astronomical objects. Furthermore, by applying corrections based on the width or decay time of the supernova light curve (the magnitude vs. time relationship), the dispersion can be reduced and the "standardness" of these candles improved. This "stretch" correction was developed by Perlmutter *et al.* [29]. A more complex correction method, combining light curves obtained in various wavelength bands, was developed by Riess *et al.* and further standardizes the supernovae [30].

Further support for the use of Type Ia SNe comes from the fact that the above dispersions are observed in spite of the fact that the supernovae come from a wide variety of host-galaxy morphologies. A significant worry in using supernovae as standard candles is that supernovae at $z = 0.3$–1 or more are somehow fundamentally different from local supernovae because of the evolution of galactic environments. Galaxy morphology — *i.e.*, whether the galaxy is elliptical or spiral or somewhere in between — is strongly believed to be an indicator of the galaxy's age, with ellipticals being much older. The smallness of the observed dispersions indicates that the galactic environment is not important. This fits in with the white-dwarf model discussed above — there simply are not many free parameters in hydrogen- and helium-free low-mass white dwarfs that can be affected by the galactic environment. However, this argument has not yet been made quantitatively.

Finally, the spectra of Type Ia SNe provide a great deal of information about the explosion, thus providing handles with which to identify abnormal ones. The spectra reflect the abundances of elements produced in the explosion as a function of time. The spectra are strongly time-dependent, but in a pattern that is followed by all normal supernovae.

In practice, the standard-candle test is done as follows. Batches of supernovae are discovered while on the early part of their light curves and followed over the peak and decay until they have disappeared. Observationally, the discovery of large numbers of supernovae is a challenge that has only been tenable in recent years. Large numbers of 4-m class telescopes permit frequent observations of many search fields to high redshift. Large-format CCDs allow the fields themselves to be large and the data to be reduced quickly. CCDs have linear photometric response, allowing relatively easy subtraction of time-separated images of a search field and good photometry in the presence of host-galaxy light. The availability of the Keck 10-m telescope enables the acquisition of spectroscopic data from high-redshift supernovae in a reasonable amount of time.

The supernova light curve thus obtained is corrected for cosmological time dilation to produce a rest-frame light curve. The temporal extent of the supernova (known as the "stretch" factor) is determined by fitting to template light curves from the low-z calibrator data set. The stretch factor is then used to perform a simple linear correction of the observed light curve to standardize to the nominal intrinsic supernova luminosity. Alternatively, the Riess *et al.* multi-wavelength correction method is performed. In either case, each measurement provides a data point $(z, F(z))$ for what is nominally a standard candle at redshift z. In principle, one would then

calculate $d_L(z)$ for each supernova and find the $E(z)$ that best fits the ensemble. This procedure, however, would introduce an uncertainty due to H_0 because $d_L(z)$ depends on H_0. Effectively, even after corrections for cosmological time dilation and stretch, the luminosities of even the calibrator supernovae are only known modulo H_0^{-2}. It is possible to remove H_0 completely by considering the ratio

$$\frac{d_L(z)}{d_L(z_{cal})} = \sqrt{\frac{F(z_{cal})}{F(z)}} \tag{1.59}$$

which depends only on the directly observed quantities, after correction based on the stretch factor. H_0 cancels out and the ratio depends only on Ω_m and Ω_Λ. In practice, this procedure is performed using fluxes in apparent magnitude units, which are logarithmic units such that a difference of 5 magnitudes corresponds to a factor of 100 in received flux. More positive apparent magnitude m corresponds to less received flux. Because magnitudes are logarithmic units, the scaling uncertainty in flux due to H_0 appears as a floating offset in a plot of m vs. z (also known as a Hubble diagram).

The Hubble diagram obtained in this way by the Supernova Cosmology Project for 42 high-z supernovae is shown in Figure 1.7, giving the corrected effective magnitude of each supernova vs. redshift [31]. Expectations for various combinations of $(\Omega_m, \Omega_\Lambda)$ are shown. A nominal value for H_0 has been taken that sets the offset of the magnitude scale. The residuals for the best fit $(\Omega_m, \Omega_\Lambda)$ combination are also shown. Clearly, the data are inconsistent with a flat, $\Omega_m = 1$ universe. Figure 1.8 shows the resulting confidence region in the $(\Omega_m, \Omega_\Lambda)$ plane. The high ellipticity of the contour reflects an interesting degeneracy of the measurement. If the discovered supernovae were all at a single redshift z_{obs}, it is only the particular linear combination of Ω_m and Ω_Λ given in $E(z_{obs})$ that is constrained. The long axis of the confidence contours reflects this degeneracy, mitigated by the fact that the supernovae are observed over a range of redshifts.

The High-z Supernova Search has taken the approach of applying the more complex Riess *et al.* correction method to achieve a lower intrinsic luminosity dispersion of the supernovae. Their 16 high-z supernovae give similar results at comparable precision [32]. Their likelihood contours in the $(\Omega_m, \Omega_\Lambda)$ plane are also shown in Figure 1.8. The final results of the two teams are, under the assumption of a flat universe, $\Omega_m = 0.28^{+0.10}_{-0.09}$ (stat.) $^{+0.05}_{-0.04}$ (syst.) for the Supernova Cosmology Project [31] and $\Omega_m = 0.28 \pm 0.10$ (stat.) for the High-z Supernova Search [32].

There are two primary outstanding issues regarding the use of Type Ia supernovae as standard candles. The first is dust. The effect seen in the Hubble diagram is essentially "cosmological dimming:" in a flat, $\Omega_m = 1$ universe, the supernovae at the observed redshifts would have appeared brighter, as is indicated by Figure 1.7. Host-galaxy dust could dim the high-z supernovae. Upon study, standard dust extinction does not appear reasonable. Such dust does not uniformly decrease the flux at all wavelengths; rather, it scatters blue light more strongly, resulting in a "reddening" of objects seen through dust. Limits can thus be placed on dust extinction by comparing the observed spectra to those of nearby unreddened supernova. The Supernova Cosmology Project estimates that the systematic uncertainty in Ω_m (under the assumption of a flat universe) caused by dust reddening is < 0.03.

Many authors have proposed extinction by "grey" dust, dust whose extinction is wavelength-independent. Such dust dims without reddening. This point requires some explanation. In the limit that the dust-grain size a satisfies $a \gg \lambda$, the dust simply blocks the light and clearly there is no wavelength dependence. When $a \sim \lambda$, a power-law dependence on a is expected. In practice, detailed calculations, modelling dust grains as dielectric spheres, indicate that the boundary

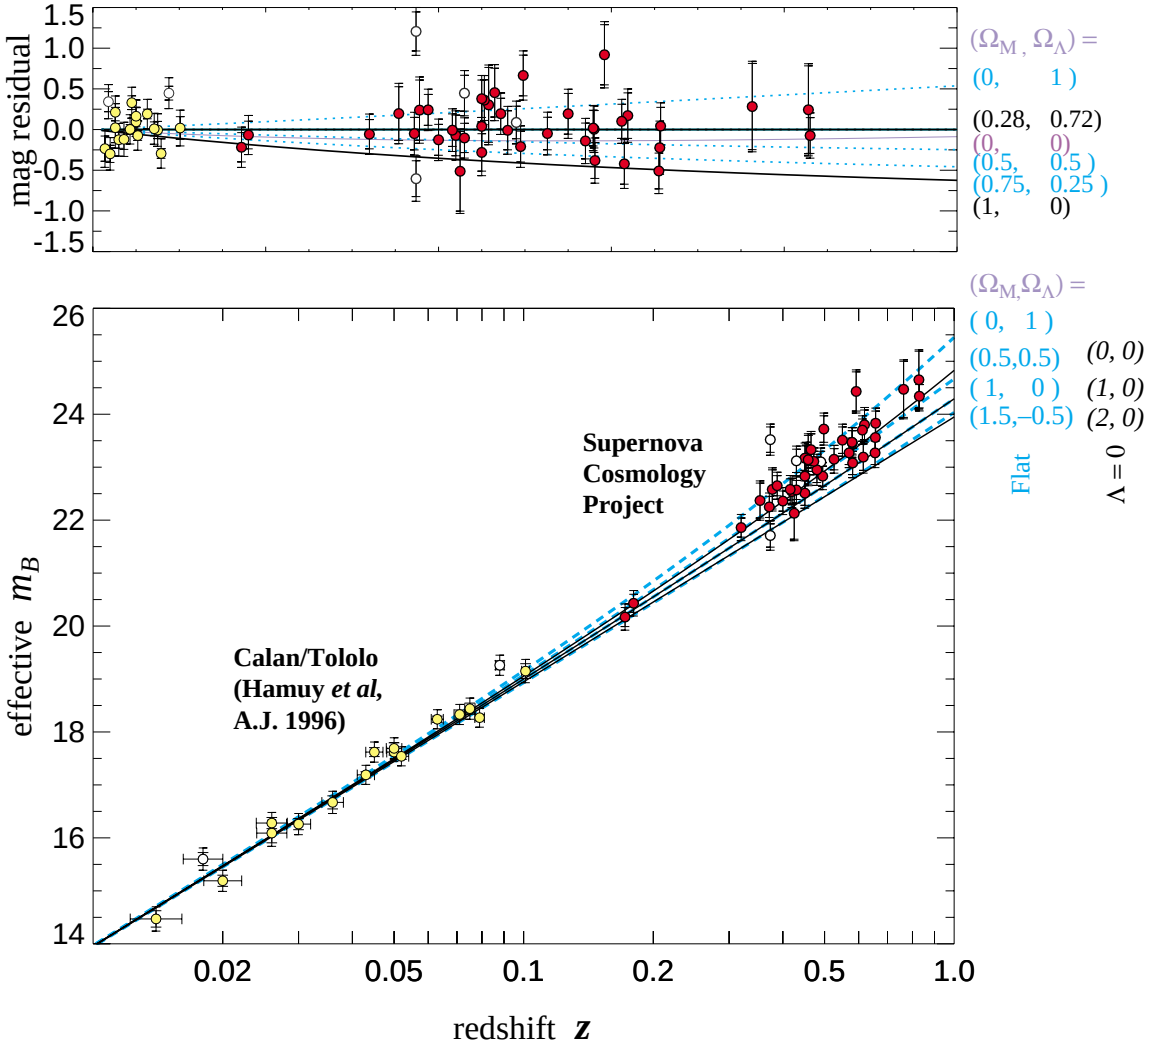

Figure 1.7: Hubble diagram from the Supernova Cosmology Project. Bottom: standard Hubble diagram of apparent magnitude vs. redshift. The slope at low z is 1, set by the low z behavior of $E(z)$. The y-intercept is set by the value of H_0 assumed. The lines indicate the models listed at the upper right corner of the plot; dashed lines are flat universes with $\Omega_\Lambda \neq 0$ allowed; solid lines are universes with $\Omega_\Lambda = 0$. Top: Residuals from the best-fit flat model, $(\Omega_m, \Omega_\Lambda) = (0.28, 0.72)$. Other model universes are also shown, with the same legend as in the bottom plot. Figure taken from [31].

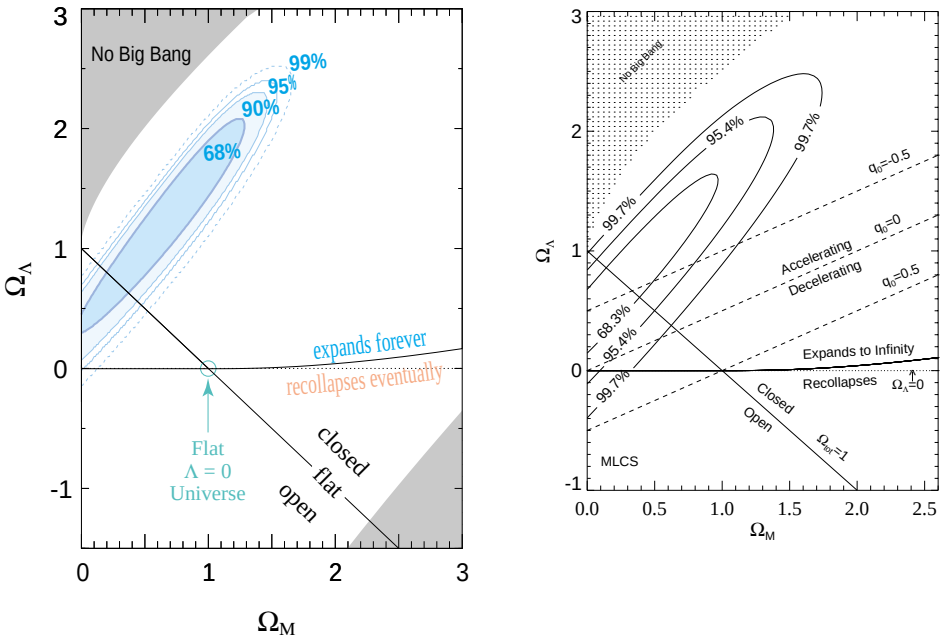

Figure 1.8: Likelihood contours in the $(\Omega_m, \Omega_\Lambda)$ plane from the Supernova Cosmology Project (left) and the High-z Supernova Search (right). Figures taken from [31] and [32].

between the two regimes is at smaller a than one would expect; the opacity appears to be constant out to $\lambda \sim 10a$ [33]. One particular grey-dust model proposes that dust grains smaller than 0.1 μm are destroyed by UV light from stars, leaving only large grains [33]. This extends the λ-independent regime to higher wavelength than for smaller dust grains, allowing the opacity to remain constant for λ up to $\sim 1\,\mu$m. This type of dust would thus dim but not redden. In order for this type of dust to produce the results seen, its abundance would have to increase with z in order to produce a differential effect (the presence of grey dust in all galaxies would dim both low- and high-redshift supernovae, moving the y-intercept of the Hubble diagram but introducing no bias). However, the lines-of-sight to observed supernovae are quite varied, probing many different path lengths through the host galaxies. Thus, the resulting grey-dust extinction should be quite varied, resulting in excess dispersion in the high-z supernova apparent magnitudes. Such an excess dispersion is not seen; the Supernova Cosmology Project reports a dispersion of $\sigma_B = 0.20$ mag for the low-z sample and $\sigma_B = 0.22$ mag for the high-z sample. Furthermore, it will be possible to test the grey-dust hypothesis conclusively because, at $z > 1$, the effect of the cosmological constant on the apparent magnitude of distant supernovae deviates from expectations from even grey-dust extinction. The cosmological constant becomes less significant (as $(1+z)^3$), so higher-z supernovae in a $(\Omega_m, \Omega_\Lambda) = (0.28, 0.72)$ model are measurably brighter than expectations from grey dust.

The second issue is evolution of supernovae. There are a number of qualitative arguments against this. First, the observed supernova light curves match those of low-z supernovae. Second, the spectra on each day match those of the low-redshift counterparts with the same stretch factor.

Third, plausible scenarios for supernova evolution should produce observable effects. For example, younger stellar populations have lower metallicity — abundance of elements heavier than helium — than older ones because the material has seen less stellar processing. Environmental metallicity should affect the C/O ratio in the resulting white dwarf, which should affect the details of the rise of the explosion. The upper limit on the difference of rise times between low- and high-z supernovae is 1 day, small compared to their rise times. Metallicity variation should also cause variations in the light curves and the spectra; no such effects are seen at levels that would affect the final analysis. However, because these arguments are not yet quantitative, they require further work to be convincing.

The possibility of more standard systematic errors is presumably ruled out by the fact that two independent teams are finding the same result. Furthermore, the cultural difference — the Supernova Cosmology Project is based at Lawrence Berkeley National Laboratory, historically a particle-physics lab, while the High-z Supernova Search consists of mainstream astronomers, most at or with pedigrees from the Harvard-Smithsonian Center for Astrophysics — provides additional hope that this is not an example of the mass delusion of a homogeneous community. Moreover, the result is so opposed to most theoretical prejudice that it probably is not an occurrence of getting the answer one wants.

However, there is one sociological aspect that causes pause. The existence of a nonzero Ω_Λ satisfies one long-standing desire of the cosmology community: a flat, $\Omega = 1$ universe. In the mid- to late-1990s, before the supernova results, the tide had been tipping away from $\Omega = 1$: many trustworthy methods were pointing to $\Omega_m \approx 0.3$, in particular measurements of cluster abundances and the observed deficit of the power spectrum of density fluctuations relative to the predictions of the standard $\Omega = 1$ inflation+cold dark matter scenario. After many years and much hoping, there was no evidence for $\Omega_m = 1$. Finally, when people were starting to get used to the idea of an open, $\Omega_m = 0.3$ universe, the supernova results came along, reconciling a low matter density with a flat universe.

There remains a largely philosophical issue, termed the "Nancy Kerrigan Problem" by many in reference to her "Why me, why now?" statement after being hit in the knee with a baseball bat right before the 1996 figure skating World Championships by the thug boyfriend of one of her competitors, Tonya Harding. Why do we happen to be living in a universe with $\Omega_m \sim \Omega_\Lambda$ today? Recall Equation 1.4: the contribution of matter to the energy density of the universe scales as a^{-3} because of volume dilution, while the contribution of vacuum energy is unaffected. Today, matter contributes around 1/4 of the energy density of the universe. Therefore, at $z = 1$, matter dominated, contributing 3/4 of the energy density. When the universe is twice its present size, matter will only contribute 1/20 of the energy density! Why do we happen to be alive and making observations at this special time?

Cluster Abundance

The abundance of galaxy clusters as a function of redshift can be used to gauge Ω_m. Galaxy clusters are believed to be have collapsed from $\sim 3\sigma$ peaks in the density-fluctuation spectrum on the appropriate scale. They are large enough that the evolution of their number (not their internal evolution!) may be treated by linear or mildly nonlinear theory of gravitational evolution. Heuristically, it is easy to see how this depends on Ω_m. At $z > 0$, the expansion rate of a low-Ω_m universe is slower than a $\Omega_m = 1$ universe. This may seem nonintuitive; it arises because both universes are constrained to have the same expansion rate today while the low-Ω_m universe experiences less deceleration and so can have a smaller expansion rate than a $\Omega_m = 1$ universe at $z > 0$. The presence of a cosmological constant has little effect on this because it becomes unimportant by $z \sim 1$. In a more quickly expanding universe, it takes longer for density fluctuations to break away from the expansion and begin to collapse — the expansion works harder to prevent a fluctuation from collapsing. Thus, cluster formation begins later in a $\Omega_m = 1$ universe. However, once collapse begins, it occurs faster: an overdensity of a fixed density contrast $\delta = \delta\rho/\rho_b = (\rho - \rho_b)/\rho_b$ corresponds to a larger physical density fluctuation $\delta\rho$ because the average density ρ_b is larger in a $\Omega_m = 1$ universe. These two effects are seen in Figure 1.9, giving the number density of clusters larger than a certain mass as a function of redshift. Alternatively, one can calculate the number of clusters above a cutoff mass per unit area on the sky and per unit redshift, as is shown in Figure 1.10. The large difference between $\Omega_m = 1$ and $\Omega_m = 0.3$ models is clear, as well as the relative lack of dependence on Ω_Λ.

Bahcall *et al.* show that, in a $\Omega_m = 1$ universe, the physical number density of clusters with $M > 3 \times 10^{14}\, h^{-1}\, M_\odot$ increases by a factor of 40 from $z = 0.5$ to $z = 0$, while in $\Omega_m = 0.3$ cosmologies, the growth factor is only 4. The resulting constraints are $\Omega_m = 0.3 \pm 0.1\,(\Omega_\Lambda = 0)$ and $\Omega_m = 0.34 \pm 0.13\,(\Omega_m + \Omega_\Lambda = 1)$. These results predate the supernova and CBR data by 2 to 3 years!

Velocity-Field Analysis

Another method that uses the evolution of density fluctuations to constrain Ω_m is velocity-field analysis. Again, by smoothing over large enough volumes ($\sim$ few Mpc), density fluctuations in the present universe become small enough to allow linear or mildly nonlinear treatment. In addition to the density-fluctuation field, one can also observe the velocity field — the peculiar (Hubble-flow-subtracted) velocities of clumps of matter. The velocity field is the response to inhomogeneities in the density distribution: matter flows toward overdensities and away from underdensities. The relationship between the velocity field and the density-fluctuation field is dependent on Ω_m for similar reasons as the cluster abundance: if the expansion rate at high redshift is increased by increasing Ω_m, it becomes harder for peculiar velocities to develop.

The relationship can be developed formally by linearly perturbing the equations describing gravitational evolution of a pressureless fluid in an expanding universe. I quote the result (see, for example, [1])

$$\vec{v}(\vec{x}) = \frac{H_0}{4\pi}\frac{a}{D}\frac{dD}{da}\int d^3y \frac{\vec{y} - \vec{x}}{|\vec{y} - \vec{x}|^3}\delta(\vec{y}) \tag{1.60}$$

where $\delta(\vec{y}) = (\rho(\vec{y}) - \rho_b)/\rho_b$ is the density-fluctuation field (ρ_b is the average background density), $\vec{v}(\vec{x})$ is the *peculiar*-velocity field, a is the scale factor, and D is the growing solution for the growth

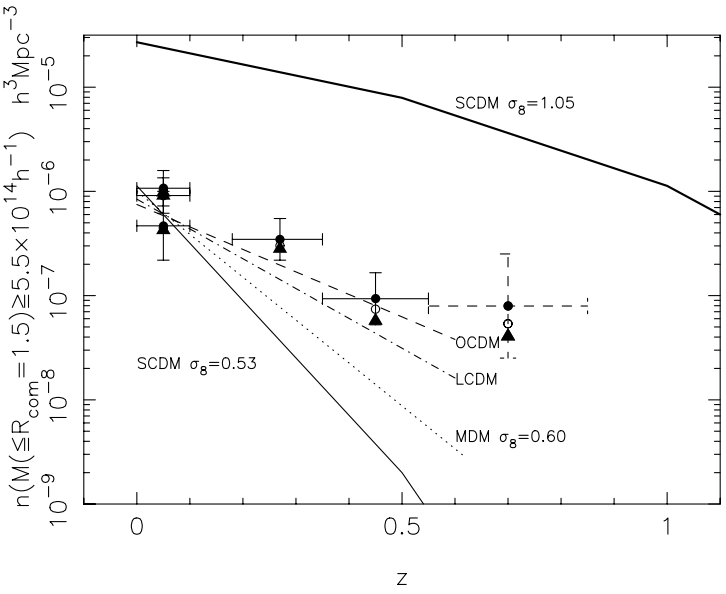

Figure 1.9: Evolution of the abundance of clusters with $M > 5.5 \times 10^{14}\, h^{-1}\ M_\odot$ within the Abell radius of $1.5\, h^{-1}$ Mpc. Four different models are shown: SCDM $(\Omega_m, \Omega_\Lambda) = (1, 0)$, OCDM $(\Omega_m, \Omega_\Lambda) = (0.35, 0)$, LCDM $(\Omega_m, \Omega_\Lambda) = (0.4, 0.6)$, and MDM $(\Omega_m, \Omega_\Lambda) = (1, 0)$ but with a significant fraction of the matter density being in hot dark matter. Figure taken from [34].

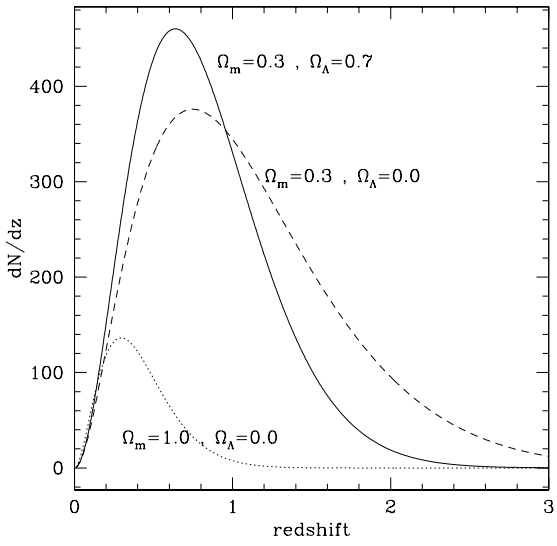

Figure 1.10: Expected number of clusters as a function of redshift with $M > 2 \times 10^{14}\ M_\odot$ in a 12 square degree patch of sky for different cosmological models. Figure taken from [14].

of density fluctuations (given later). This can be written in a slightly more intuitive form

$$\vec{v}(\vec{x}) = \frac{2}{3}\frac{f}{\Omega_m H_0}\vec{g}(\vec{x}) \tag{1.61}$$

with

$$f = \frac{a}{D}\frac{dD}{da} \sim \Omega_m^{0.6} \tag{1.62}$$

That is, the velocity field is given by the gravitational-acceleration field multiplied by a quantity that is approximately the age of the universe. The dependence on the actual value of Ω_m enters in the value that this effective time takes. Ω_Λ is irrelevant to first order, as explained with regard to cluster abundance. Thus, if $\vec{v}(\vec{x})$ and $\delta(\vec{x})$ can be measured, the time factor can be determined.

The velocity field is nontrivial to determine. Redshift gives the sum of peculiar velocity and Hubble flow. To subtract the Hubble flow, an independent distance indicator is needed. There exist relationships, tested in the nearby universe, relating the luminosity of galaxies to their internal parameters, usually velocity dispersion (Faber-Jackson relation for ellipticals) or disk circular velocity (Tully-Fisher relation for spirals). Thus, the peculiar velocity along the line of sight can be determined. To construct the full velocity field, one makes use of the fact that the velocity field is irrotational (it can be shown that curl in the velocity field is diluted by the expansion). The velocity field thus is the gradient of a scalar potential. The scalar potential (not the gravitational potential!) can be calculated from the line integral of the velocity field along the line-of-sight from us to the desired point. The full velocity field can then be derived from the scalar potential.

The density field can be determined from redshift surveys of galaxies. The data set used to determine the velocity field is not good for this purpose because, in the velocity-field set, there is a bias toward high-luminosity galaxies for which one of the distance-indicator methods can be applied. Other surveys of galaxy redshifts, with many more galaxies and geared more toward determining the density field, are used. However, it is possible that even these galaxies may be biased in that they may overpopulate the largest fluctuations. A fudge factor b is introduced: $\delta_m = \delta_g/b$, $b \geq 1$. Independent distances for the galaxy that determine the density field are not available, so the final comparison is done in redshift space rather than physical space.

In practice, there are many gory details of how the above procedure is executed. Most important is obtaining a large catalog of galaxies with both redshifts and distances. New catalogs are being constructed or have recently come out, improving the precision of the determination of the velocity field. The present constraints on $\beta = \Omega_m^{0.6}/b$ range from 0.5 to 0.9, giving Ω_m from 0.3 to 0.8 if no biasing is assumed [35]. The assumption of a single bias parameter is well known to be a drastic oversimplification, and it is this lack of knowledge of how galaxy fluctuations are linked to the underlying dark-matter fluctuations that leaves this type of analysis unable to provide the precise results that other methods give.

Alternatively, the velocity field can be analyzed in a statistical way, similarly to the way galaxy-redshift surveys are used to measure the power spectrum of density fluctuations (see Section 1.4.2). The power spectrum of velocity fluctuations is, in principle, more indicative of the underlying matter power spectrum because velocities respond to the total mass density while galaxy-redshift surveys measure only the luminous matter distribution. Such analyses presently constrain Ω_m to be between 0.3 and 1 at 95% CL [35]. Because velocity-field data are more difficult to acquire than galaxy-redshift data, the precision of even these analyses does not yet match those

of power spectra estimates derived from redshift surveys alone. With increases in the velocity-field data set, this will change soon.

1.3.4 Theoretical Prejudice

There are two good theoretical reasons why Ω should be 1 today, in the past, and forever. Being theoretical, I hesitate to classify them under "evidence," but they do not deserve their own section heading, either!

The Dicke Coincidences Argument

As one can infer from the gravitational-escape analogy, $\Omega = 1$ is the only stable value of Ω. All larger or smaller values deviate from 1 as power laws in time. Thus, the observation that Ω is close to 1 today presents a fine-tuning problem, first pointed out by Robert Dicke: if $\Omega \sim 1$ today, it had to be ridiculously close to 1 in the early universe, yet not exactly. This is also known as the "flatness" problem.

Inflation

In addition to the flatness problem, there are other serious questions posed by observations that have no answer in "classical" cosmology. One issue is the "smoothness" or horizon problem. The cosmic microwave background radiation is incredibly homogeneous — even the largest fluctuations are only 3 parts in 10^5. As shown earlier, the angular size of the horizon at the surface of last scattering is approximately 1°. Thus, at the surface of last scattering, the presently visible universe was composed of $\sim 10^7$ causally disconnected regions that all have the same temperature to 3 parts in 10^5. Similar arguments can be made regarding nucleosynthesis, with an even more extreme result. The other problem is that the source of the density inhomogeneities that gave rise to the structure seen today is not known. As is discussed below, in order to generate the universe's present structure, it is necessary that there were density fluctuations of 1 part in approximately 3000 in the dark matter at the surface of last scattering. Density fluctuations are observed in the CBR. What produced these fluctuations in the early universe?

Inflation, first suggested by Alan Guth in the early 1980s and since modified many times, is a generic method for solving these problems. The central proposition is that, very early in the history of the universe, well before nucleosynthesis, the universe experienced an epoch of exponential expansion driven by a vacuum energy, or something like it. This expansion solves the smoothness problem by changing the horizon size: the entire visible universe was inflated out of a small patch, smaller than the horizon size at the beginning of inflation. Inflation solves the flatness problem by making use of the fact that exponential expansion flattens out the physical curvature. The curvature parameter R of the universe is unchanged, but the scale factor increases exponentially, so the physical radius of curvature today, $a_0 R$, is huge and the observable universe is indistinguishable from flat. To yield an exponential expansion, inflation is usually modeled as produced by a scalar field ϕ with a classical potential $V(\phi)$. The field configuration in which the field ϕ assumes the value ϕ_{min} that minimizes the potential corresponds to the universe's vacuum (lowest energy) state. In the early universe, when T is much larger than the energy depth of the potential, ϕ is driven away from the minimum of V by thermal fluctuations. As T drops, ϕ begins to be affected by the

potential, but it is not likely to be at the minimum. Any field configuration with ϕ away from ϕ_{min} is not the lowest energy state. Therefore, regions with $\phi = \phi_{min}$ expand into regions with $\phi \neq \phi_{min}$. The regions with $\phi \neq \phi_{min}$ thus have negative pressure: the true vacuum state, which by definition has zero pressure, expands into them. There are regions of the universe that have nonzero energy density, but the energy density does not decrease with the expansion of the universe: the classical potential is not diluted by the universe's expansion. This looks like a cosmological constant. Thus, there is a mechanism to generate the exponential expansion phase. To make the universe exit this inflationary phase, ϕ must roll down to the minimum of the potential with $V(\phi_{min}) = 0$. Finally, inflation can provide a density-fluctuation field. The field ϕ possesses quantum fluctuations than can have an effect once inflation has ended, yielding the initial density inhomogeneities required. Inflationary models can even generate the desired scale-free spectrum of fluctuations discussed below.

Inflation has other nice features. In addition to generating the initial density flucuations, it may generate all the matter in the universe. When ϕ rolls down into the minimum of its potential, it oscillates because it acquires kinetic energy during the roll. These oscillations must decay to give $\phi = \phi_{min}$. The decays of the ϕ field thus convert the ϕ kinetic energy into matter and radiation that fill the universe.

So far, consensus has not been reached on a model for inflation. The general idea is a good one and many think it is correct, but a fully workable model does not yet exist. To date, there is only indirect evidence for it. Inflation produces a scale-invariant power spectrum of density fluctuations on super-horizon scales, as is observed (see later discussion the growth of density fluctuations). It produces the necessary temporal correlation of the growth of density fluctuations across the visible universe that is needed to explain the existence of the CBR first Doppler peak (why were all the waves at this particular scale across the whole sky at their maximum amplitude at the same time, decoupling?). It is hoped, though, that one direct signal will be detected: the influence of gravitational waves on the CBR. It is expected that the inflationary field generates relic gravitational waves (gravitons) in addition to simple scalar density fluctuations on super-horizon length scales. It is not possible to distinguish the two in the CBR temperature anisotropy. However, by measuring the polarization anisotropy of the CBR, a signal may be seen. Purely scalar fluctuations would produce an irrotational polarization field; the vector field of polarizations would have to be derived from a scalar potential, and thus would have no curl. Gravitational waves, being a tensor field, could produce curl. The observation of curl in the CBR polarization anisotropy would be the smoking gun of inflation or some similar theory for the generation of density fluctuations. Measurements of the polarization anistropy are approximately 5 to 10 years behind those of the temperature anisotropy because the signal is expected to be 10 to 100 times smaller. There is little doubt, though, that these measurements will be made in the coming decade, finally testing inflation directly.

A final point about inflation — it presents questions also, though perhaps not as bad as those it answers. In addition to setting up a workable model, with a field and potential that do the right thing, one has to ask — who ordered that? That is, is there a reason, other than inflation, to have the field ϕ? Second, there is the question of the naturalness of the potential. Early models required a very extended roll down an almost perfectly flat potential to achieve the needed amount of inflation. One fears that one is trading one fine-tuning for another. Finally, and this is more a philosophical point, inflation is unsatisfying in one way. It implies that we will never be able to

determine the global geometry of the universe. Inflation inflates away the physical curvature of our local patch, and presumably of all the other patches beyond the horizon, so we will never be able to determine what the R in the FRW metric is.

1.3.5 Summary

The current data yield a self-consistent picture with $\Omega \approx 1$ and $\Omega_m \approx 0.3$. These are much larger than the component observed in luminous matter, $\Omega_{\rm lum} \approx 0.004$. Clearly, there is a great deal of evidence for the existence of dark matter and energy.

1.4 The Need for Nonbaryonic Dark Matter

I have discussed above the evidence that a large fraction of the matter in the universe is dark. There is an additional set of arguments that indicate that most of the dark matter must by nonbaryonic.

1.4.1 Big Bang Nucleosynthesis

"Big Bang Nucleosynthesis" (BBN) is a generic name for calculations of the light-element abundances (H, D, He, ^{3}He, and ^{7}Li) within the framework of a hot Big Bang that yields an expanding, cooling universe in local thermal equilibrium. The binding of initially free neutrons and protons into light nuclei occurs because the universe's temperature drops below nuclear binding energies, suppressing the number of photons energetic enough to photodissociate these nuclei. BBN provides a strong constraint on the total baryon density of the universe, indicating the need for nonbaryonic dark matter. This discussion is given in standard textbooks; I follow [2].

There are two basic aspects of nucleosynthesis. First is the establishment of the neutron-to-proton ratio. At high temperature, $T \gtrsim 3$ MeV, neutrinos interact frequently enough to maintain thermal equilibrium between neutrons and protons via the reaction

$$p + e^- \leftrightarrow n + \nu_e \tag{1.63}$$

The neutron-to-proton ratio is thus maintained at the equilibrium value set by the two particles' mass difference:

$$\left(\frac{n}{p}\right) = \exp(-Q/T) \tag{1.64}$$

where $Q = 1.293$ MeV is the mass difference. However, at $T \sim 3$ MeV, neutrinos decouple, halting the above reaction. The n/p ratio becomes fixed, taking on the value $(n/p)_f \approx 1/6$; this ratio is "frozen in." The relative abundance of neutrons and protons is critical for the following.

Second, light elements begin to form because the temperature becomes low enough that photons cannot dissociate nuclei. By $T \approx 0.03$ MeV, the light elements ^{4}He, ^{3}He, ^{3}H, ^{7}Li, and ^{7}Be have established their final abundances. The formation reactions stop when all the neutrons have been used. The (n/p) ratio at freeze-out thus is critical in setting the final abundances. Higher-mass elements do not form because there is a large Coulomb barrier to the necessary reactions. The time development of the abundances is shown in Figure 1.11.

Most of the cross sections for the nucleosynthesis reactions are known precisely. The half-life of the free neutron against β decay is especially important because it gives the rate of the above reaction that maintains neutron-proton equilibrium, determining the freeze-out temperature T_f, the resulting ratio n/p, and thus the final number of neutrons available to form $A > 1$ elements. It also results in the disappearance of neutrons as time increases, limiting the abundances of $A > 1$ elements. Another important parameter is the number of effective relativistic degrees of freedom, g; the universe is radiation dominated at this time, so g sets the density and pressure and thus the expansion rate, which also determines when specific reactions stop. In particular, g is sensitive to the number of light-neutrino species; an impressive success of BBN was the prediction of three (non-sterile) light-neutrino species before the number could be measured directly from the width of the Z boson. Finally, the only free parameter is the density of baryons ρ_b because this, obviously,

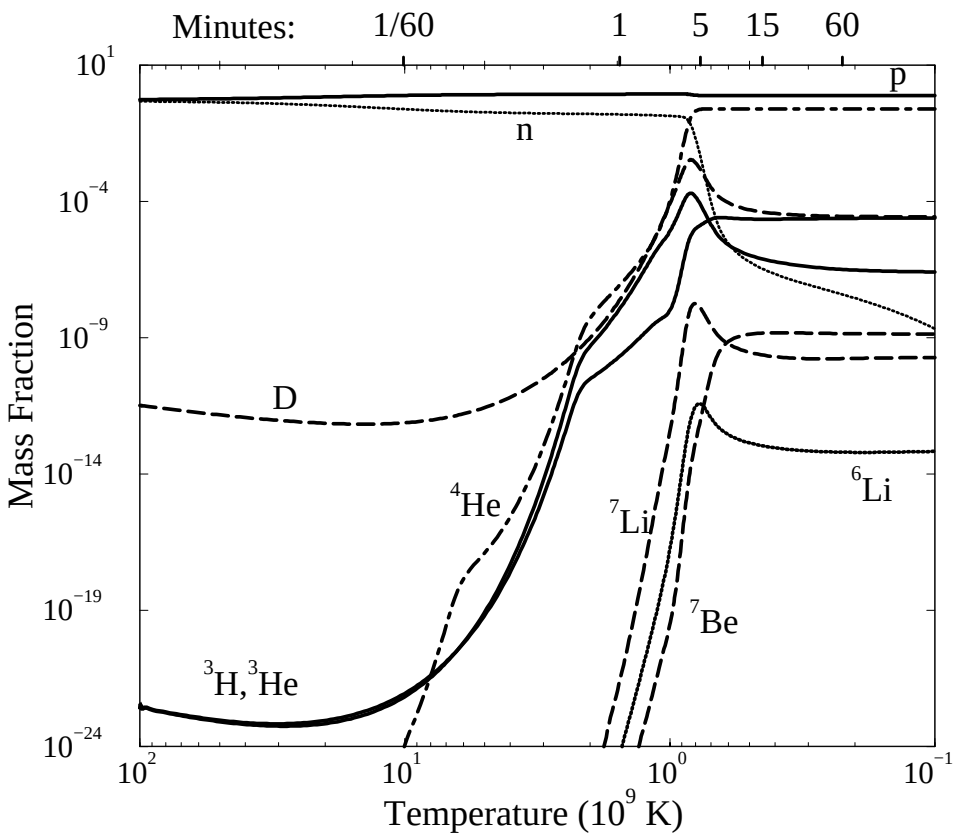

Figure 1.11: Evolution of light-element abundances with temperature and time. Figure taken from [36].

determines possible reaction rates. As the density goes up, the abundance of the very stable element ^{4}He is increased and the abundances of the other, less stable elements decreased. In the literature, $\eta = n_b/n_\gamma$, the baryon-to-photon ratio, is usually used in place of ρ_b; the present photon density is known from the CBR.

Observations of the abundances of the light elements in systems where little or no stellar processing has occured constrain ρ_b. These measurements are summarized in Figure 1.12, which shows the dependence of the abundances relative to hydrogen and the present constraints. The present allowed range for Ω_b is $0.018\,h^{-2} < \Omega_b < 0.022\,h^{-2}$ [3]. h^2 enters because the Ω's are defined in terms of it (see Equation 1.6).

Measurement of light-element abundances is nontrivial because one must look to systems where stars have not affected the abundances by burning some elements or producing others. ^{4}He is relatively easy to measure because of its large abundance. Emission from recombination of ionized ^{4}He in HII regions (essentially, large gas clouds of primarily ionized atomic hydrogen that have not cooled enough to permit star formation) is used. The abundance of elements heavier than helium (*metals*) is used to determine the amount of stellar processing that has occurred in such regions. The ^{4}He abundance vs. metallicity data are fitted to a power law and extrapolated to zero

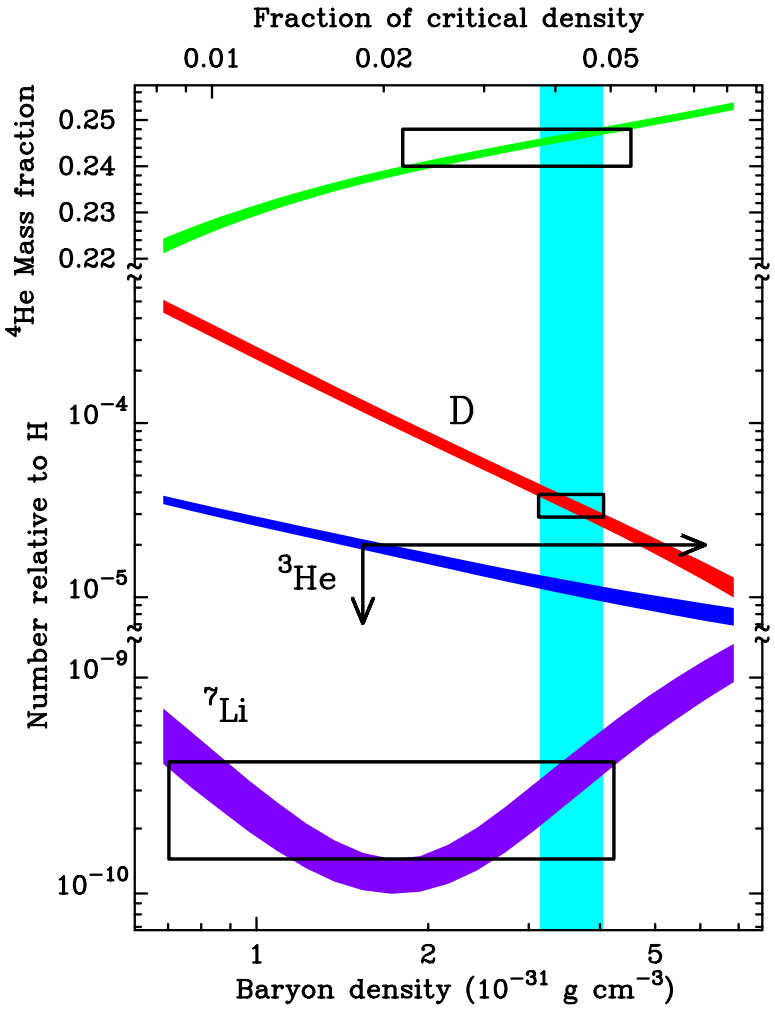

Figure 1.12: Abundances of light elements relative to H as a function of the present baryon density. The top horizontal scale converts ρ_b to Ω_b, with $h = 0.65$ assumed. The widths of the abundance curves indicate 95% CL confidence intervals for the abundance predictions. Measurements of the abundances in primordial environments are shown as boxes. The vertical shaded band indicates the region allowed by deuterium observations. Figure taken from [36].

metallicity to give the primordial ^{4}He abundance.

Deuterium is the next most abundant light element and is much more sensitive to η because D would much prefer to be either H or ^{4}He. It has been especially useful for providing an upper limit on η and thus on Ω_b because it is not created by stars — any measurement of D is a lower limit. It is difficult to observe, however, because, in addition to its low abundance, its emission lines are very close to those of H, separated only by a small isotopic shift due to the different reduced mass of the atom. For example, the Lyman-series transitions (to $n = 1$) of H and D are separated by only 82 km s^{-1}, or $\delta\lambda/\lambda = 2.73 \times 10^{-4}$, 0.033 nm for Ly-$\alpha$ at 121.6 nm. Observations are made by looking at absorption spectra of distant quasars through intervening high-redshift gas clouds. The quasar provides a continuum of ultraviolet radiation that is redshifted as it travels to us. As this light passes through such clouds, Lyman-series absorption lines appear, indicating the presence of H and possibly D in the clouds. The observed absorption of Lyman-series lines is called the "Ly-α forest" because of the forest-like appearance of the spectra. Lines-of-sight with this arrangement are termed "QSO-absorption-line systems."

The difficulties with the observations are numerous. Such a cloud must have a large enough column density (atoms/cm^2) that absorption by D is visible, but not so large that the wings of H absorption wash out the D signal. One must be confident that the system is not made of two clouds that happen to be separated in redshift by the necessary $82(1+z)$ km s^{-1}, or of two neighbor clouds that have peculiar velocities differing by the above amount. The thermal broadening of the lines must also be small enough that the D line can be seen. For these many reasons, only 4 or 5 usable systems have been found despite intensive search over the last decade. The constraint on η from these measurements is shown in Figure 1.12.

^{3}He is not very useful because it is both produced and destroyed in stars. The present measurements are not consistent and plagued by chemical-evolution issues. ^{7}Li is interesting as a cross-check. Empirically, it is observed that the ^{7}Li abundance in the atmosphere of old stars plateaus at high stellar mass, suggesting this value is indeed the primordial abundance. Qualitatively, it is believed that this plateau occurs because the convection currents in low-mass stars penetrate more deeply, pushing the lithium down to depths where it can be burned. The present constraint on ^{7}Li abundance from these observations is also shown in Figure 1.12.

The BBN constraint on Ω_b can be circumvented in two ways. One is to posit segregation of baryons in the early universe, such as in primordial black holes or other impenetrable objects. Another is to require significant inhomogeneities. However, inhomogeneities have difficulty in strongly affecting the abundances because neutrons are free to diffuse since they do not interact with photons: fluctuations in the neutron density are strongly damped. Since the neutron abundance is primary in determining the elemental abundances, the elemental abundances are not much changed.

To summarize, observations of light-element abundances constrain ρ_b and thus Ω_b. The reliability of BBN is based on the fact that the physics is well understood and the results have been tested a number of ways. The constraint it places on the baryon density is important because it is very model-independent. Finally, it is the large discrepancy between the BBN value of Ω_b and large-scale measurements of the matter density Ω_m that indicate a severe need for nonbaryonic dark matter.

1.4.2 Structure Formation and Nonbaryonic Cold Dark Matter

A second line of argument for the existence of nonbaryonic dark matter relies on our understanding of the growth of structure in the universe. This argument is somewhat complex, but is, in the end, fairly simple physics and thus is considered reliable. Furthermore, it leads to the requirement that a significant component of the dark matter be "cold" — non-relativistic at the time of matter-radiation equality. This discussion follows [1] and [2].

The argument is as follows. The early universe was extremely homogeneous. Out of this homogeneity, it is necessary to form the observed clumpy universe. This is possible if small fluctuations in the density were produced in the early universe and grew to the structures seen today. However, the fluctuations in the baryon density imprint themselves in the CBR at the surface of last scattering as anisotropy. With baryons only, the observed fluctuation size is too small to have given rise to the structure seen today. A solution is to posit the existence of matter that does not couple to photons; fluctuations in this matter's density can grow without observably affecting the amplitude of CBR anisotropy. Then, at radiation-matter decoupling, baryons fall into the potential wells prepared by the proposed matter and give rise to the objects seen today.

To flesh out the above argument, some knowledge of how fluctuations grow in an expanding universe is required. The fluctuation in the matter density may be written as $\delta(\vec{x}) = \delta\rho(\vec{x})/\rho_b$, where ρ_b is the average density and $\delta\rho(\vec{x}) = \rho(\vec{x}) - \rho_b$ is the deviation from uniformity at $\vec{x}$. The Fourier transform of $\delta(\vec{x})$ is $\delta_{\vec{k}}$. The autocorrelation function $\xi(\vec{x}, \delta\vec{x}) = \langle\delta(\vec{x})\delta(\vec{x} + \delta\vec{x})\rangle$ can be calculated; homogeneity ensures it is translation-invariant and thus a function only of $\delta\vec{x}$; isotropy ensures that it is function only of δx. The average is taken over all space. The autocorrelation function can be Fourier transformed to yield a power spectrum $P(\vec{k})$. Isotropy ensures that the angular integration is trivial, giving

$$|\delta_k|^2 = P(k) = \frac{4\pi}{k}\int d^3x\,\xi(x)\,\sin kx \tag{1.65}$$

where x is taken to be the argument of ξ rather than δx. Under the assumption that these fluctuations are random and Gaussian distributed, the phases of the transform $\delta_{\vec{k}}$ are random and uniformly distributed over $(0, 2\pi)$. The power spectrum provides all useful information in this case. It is assumed that a primordial power spectrum is provided somehow, usually from quantum fluctuations of the scalar field that drove inflation, with a power spectrum $P(k) \propto k^n$, with $n \approx 1$. Scale-invariance corresponds to $n = 1$ and is attractive because it is not expected that there is any characteristic scale in the early universe. The study of possible non-Gaussianity of the power spectrum is interesting and may point to non-inflationary scenarios for generation of the initial density fluctuations, but is not discussed here.

The growth of fluctuations can be calculated by solving the equations for the evolution of a fluid in an expanding universe, including the effects of gravity and of pressure of any electromagnetically interacting, relativistic component. The relativistic component must interact for it to provide pressure. The case of relativistic, weakly interacting matter is discussed below. At high redshift, perturbations are small and the equations can be linearized. It is necessary to work in comoving coordinates so that it is possible to track a given mode k as the universe expands. Linearity allows independent treatment of different Fourier modes. One finds a second-order differential equation, modified by the expansion of the universe. In the absence of expansion and of fluid pressure, the solutions are exponential growth and decay: overdensities grow exponentially.

The introduction of pressure provides a restoring force against collapse. Energy can stream out of overdensities if the wavelength of the mode is smaller than τc_s where τ is the timescale of collapse and c_s is the sound speed, of order the speed of light for relativistic matter. Thus, with pressure, fluctuations smaller than the horizon size oscillate (exponential with imaginary argument) while larger-scale fluctuations grow. In an expanding universe, the exponential forms are modified to power laws. In addition, calculation of the growth for super-horizon-scale perturbations requires a full general-relativisitic formulation. Full discussion is given in [1] and [2]. I quote the results for the case of a flat universe. This is a reasonable approximation because, for the high-redshift universe considered here, the universe becomes very close to flat.

There are two epochs — radiation domination and matter domination. When the energy density is dominated by radiation, pressure is important. The dividing wavelength between oscillating and growing modes is the Jeans length,

$$\lambda_J = \sqrt{\frac{\pi c_s^2}{G\rho}} \tag{1.66}$$

where $c_s = \sqrt{dp/d\rho}$ is the adiabatic sound speed, $c_s = c/\sqrt{3}$ for a relativistic fluid. The above expression is easy to understand: $(G\rho)^{-1/2}$ is the only expression with time dimension that can be formed from G and ρ; it is the gravitational collapse time of a sphere of uniform density ρ. This would be the time constant of the exponential fluctuation growth in the absence of pressure. The quantity λ_J/c_s is the sound travel time across the wavelength λ_J. Thus, the Jeans length is the largest wavelength for which pressure can act to prevent the collapse of a density fluctuation. For $\lambda > \lambda_J$, sound does not travel quickly enough to drain the energy out of overdensities. The growth laws for fluctuations smaller and larger than λ_J are

$$\delta_k \propto (1+z)^0 \quad \lambda < \lambda_J \qquad\qquad \delta_k \propto (1+z)^{-2} \quad \lambda > \lambda_J \tag{1.67}$$

That is, fluctuations at wavelengths smaller than the Jeans length oscillate and do not grow (and in fact may decay slowly) while larger-scale fluctuations grow.

During matter-domination, in the absence of pressure, the growth laws are modified because the universe expands more quickly — recall that pressure aids in decelerating the expansion of the universe, Equation 1.2. The growth law is

$$\delta_k \propto (1+z)^{-1} \tag{1.68}$$

which holds for all wavelengths because there is no pressure support.

With the above results, the picture is as follows. A primordial power spectrum $P(k) \propto k^n$ is provided in the early universe. During radiation domination, modes with $k > k_h$ do not grow, where k_h is the wavevector corresponding to the horizon size, which is of order the Jeans length. (Recall that k and k_h are in comoving coordinates.) Modes with $k < k_h$ grow. However, the horizon size is continually growing (even in comoving units), so k_h moves to lower k . Modes that are initially super-horizon sized stop growing as they enter the horizon. The high-k end of the power spectrum thus assumes the form of a falling power law, with the exponent fixed by the growth law and the dependence of the horizon scale on redshift.

When the universe passes into matter domination, all modes begin growing with the same growth law. Thus, the shape of the power spectrum at matter-radiation equality is frozen and grows

uniformly. A characteristic scale appears, $k_h(z_{eq})$: there is a peak in the power spectrum at this wavevector because modes at lower k have been growing all through radiation domination while modes at higher k were stalled at an amplitude that decreases as k increases. The length today corresponding to the position of the peak is about $13\,\Omega_m^{-1}h^{-2}$ Mpc, a few times the size of galaxy clusters. The characteristic shape is shown by the power-spectrum model in Figure 1.13. Note that, because of nonlinear evolution of the power spectrum, the peak of the power spectrum appears a factor of a few to lower k than one would expect just based on the horizon size at $23700\,\Omega_m h^2$. The shape is indicative of what the linear power spectrum would look like.

Baryonic matter does not follow the above growth law because there is still significant pressure even during matter domination due to the tight coupling of baryons to photons. In the case of such a mixed fluid, the sound speed is

$$c_s^2 = \frac{c^2}{3}\frac{4\rho_r}{4\rho_r + 3\rho_b} \tag{1.69}$$

which drops as the radiation density drops. The present photon energy density is about 2 orders of magnitude smaller than the baryon energy density; the ratio grows as $1+z$, so they are equal at $z \sim 100$; hence the photon energy density dominates the above expression for $z > z_{dec}$, and so $c_s \sim c/\sqrt{3}$ as held during radiation domination. The result is that sub-horizon-sized perturbations in the baryon-photon fluid oscillate at fixed amplitude from the time they cross inside the horizon until decoupling. Super-horizon-sized perturbations grow according to matter-domination growth law, Equation 1.68.

Consider this from another point of view. Large density fluctuations are observed today. The present-day power spectrum of density fluctuations is shown in Figure 1.13. Of course, the fluctuations have grown well beyond the linear regime. However, on large enough scales, it is possible to calculate the growth of linear or mildly nonlinear fluctuations and thus evolve this power spectrum back to the time of decoupling. This results in an expectation for the fluctuation spectrum at decoupling. Since the baryons are tightly coupled to the photons up to decoupling, the CBR should reflect these fluctuations. The amplitude of the predicted temperature anisotropy on the 0.5 to 1 degree angular scale is $\Delta T/T \approx 4 \times 10^{-4}$ [2]. The observed fluctuations are $\Delta T/T \approx 2.5 \times 10^{-5}$ (see Figure 1.6), a factor of 16 smaller. Clearly, there is a problem with a baryons-only universe.

The problem is solved by disconnecting the CBR from the matter-density fluctuations on angular scales smaller than 1 degree. Assume the existence of a dominant, weakly interacting, nonbaryonic dark-matter component. "Weakly interacting" specifically denotes matter that does not interact electromagnetically or by the strong force; thus, it is not coupled to the photons (the CBR). As discussed above, the fluctuations in this component grow as $1+z$ on all scales during matter domination. On super-horizon scales, the baryon-photon plasma tracks the dark matter because pressure cannot prevent collapse on such large scales. On sub-horizon scales, the baryon-photon plasma oscillates. The stalling of the baryon-photon plasma density fluctuations keeps their amplitude small enough to be consistent with the observed amplitude of the CBR temperature anisotropy on sub-degree scales. When the CBR decouples, the baryons are left pressureless, allowing collapse. The baryons fall into the potential wells prepared by the nonbaryonic dark matter.

Nonbaryonic dark matter solves the problem of excessively large baryon-density fluctuations at decoupling. However, it may overdo the job if it is capable of "free streaming" — essentially,

escaping from density fluctuations because it is weakly interacting. In the standard model of the growth of density fluctuations, discussed above, it is assumed that all relativistic matter is strongly interacting. If this does not hold, then relativistic matter can stream away from overdensities, carrying away energy, without being stopped by interactions with other particles. This phenomenon is called "free streaming." A free-streaming length, λ_{FS}, can be calculated as a function of redshift; fluctuations smaller than λ_{FS} are damped by free streaming. It is most useful to calculate this for z_{eq} because this is the point at which the growth of small-scale density perturbations begins. There are two extreme cases. First is the case of a particle species that becomes nonrelativistic at z_{eq}: this gives the largest possible free-streaming length. This type of particle is termed *hot dark matter*. Obviously, the free-streaming length is just the horizon size at matter-radiation equality, corresponding to the position of the peak in the power spectrum. The power spectrum to the right of the peak is damped by free streaming. The other extreme is *cold dark matter*, a particle species that becomes nonrelativistic early enough that its free streaming can be ignored. Such a species produces no damping of the power spectrum. A pure hot dark-matter particle proves problematic because objects smaller than the horizon size at z_{eq}, which corresponds to about $13\Omega_m^{-1}h^{-2}$ Mpc today, are not produced from primordial fluctuations but must form by the fragmentation of larger objects. There is not space to discuss this here; it suffices to note that this is inconsistent with the power spectrum observed today. The alternative, pure cold dark matter, is more compatible with observations, and a fit of such a model is shown in Figure 1.13. The lack of any free streaming produces more power on small scales than is observed. In terms of particle masses, maximally hot dark matter has a mass of a few eV or less, which is not unreasonable for a light neutrino species. Cold dark matter is less specific; anything with a mass of order 1 MeV or larger suffices. Finally, there is the possibility of *warm dark matter*. A particle with mass of order 1 keV would have a free-streaming length of about 1 Mpc, sufficient to damp power on galactic and sub-galactic scales but having no effect on the scales of clusters and larger. Warm dark matter would solve CDM's problem of too much power on small scales.

To summarize, the existence of nonbaryonic cold or warm dark matter would reconcile the observed size of density fluctuations today with the amplitude of temperature anisotropies in the CBR on sub-degree scales. A universe containing only baryons and photons is inconsistent with the observed CBR anistropy.

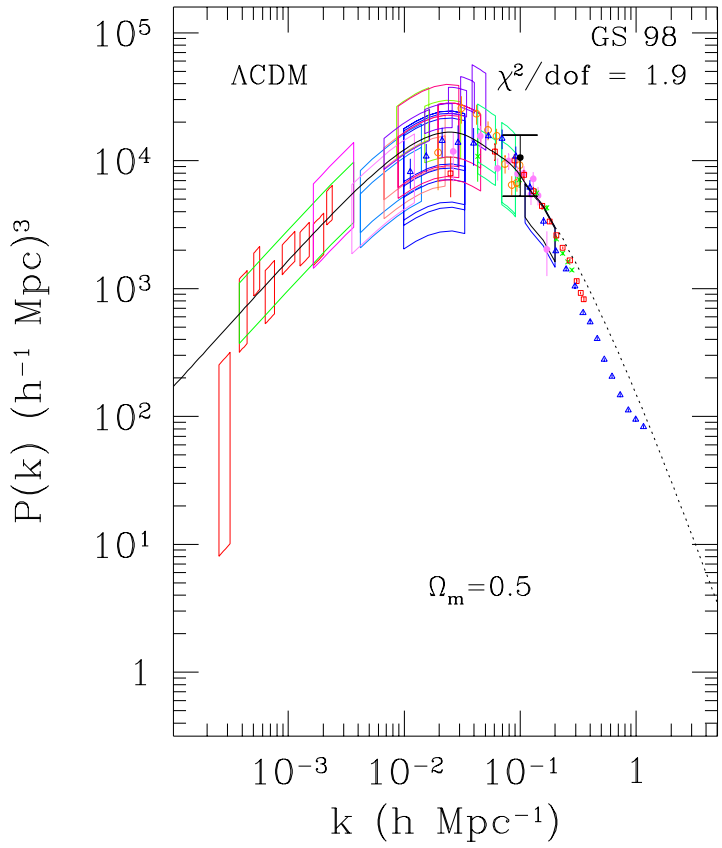

Figure 1.13: The power spectrum of density fluctuations as a function of present-day wavevector. The line is a cosmological model with $(\Omega_m, \Omega_\Lambda) = (0.5, 0.5)$. A variety of measurements are shown: CBR anisotropy (boxes) at low k and over the peak, cluster abundance (large box near the peak, rightmost box), peculiar velocities (point with large error bar), and galaxy-redshift surveys (remaining points with error bars). Figure taken from [37].

1.5 Baryonic Dark Matter

There is a large baryonic dark-matter problem. As described above, BBN constrains the range of Ω_b to be $0.018\,h^{-2} < \Omega_b < 0.022\,h^{-2}$. This is much larger than the observed density of luminous matter, as indicated by Equation 1.23. Where are the dark baryons? Much work has been done on this subject in the last few years, both at low redshift with microlensing searches for MACHOs in our galaxy as well as at high redshift with baryon inventories via QSO-absorption-line systems. Searches for baryonic dark matter in our galaxy are of particular interest here because they impact the nonbaryonic dark-matter content of the halo, which obviously has direct consequences for the detectabilty of nonbaryonic dark matter.

1.5.1 Dark Baryons at High Redshift

The most obvious place to look for dark baryons is in the form of neutral hydrogen in the intergalactic medium. If uniformly distributed and of sufficiently low density, such gas will not have collapsed and cooled and hence would not have produced stars. Since it is neutral, it does not emit. The presence of uniform, neutral hydrogen can be detected by the absorption it produces in the spectra of high-z quasars, the Gunn-Peterson test. Quasars display quite flat continuum emission in the UV, except for a prominent peak due to Ly-α emission at 121.6 nm. As the continuum emission at shorter wavelengths than Ly-α travels to us, it is redshifted so that it may be absorbed by the Ly-α transition of intervening neutral hydrogen at lower redshift. A uniform density of neutral hydrogen thus produces a uniform decrement in the observed spectrum for wavelengths shorter than the Ly-α peak. The emitted continuum flux can be determined from the observed spectrum at λ longer than the Ly-α peak: such longer wavelengths cannot be absorbed by neutral hydrogen, under the assumption that the hydrogen is cold enough that the population of $n > 1$ states is negligible. The resulting limit on the density of uniformly distributed, neutral hydrogen at $z \gtrsim 1$ is quite impressive [1]:

$$\Omega_{\rm HI} \lesssim 2 \times 10^{-7}\,\Omega_m^{-1/2} h^{-1} (1+z)^{3/2} \tag{1.70}$$

(HI is the astronomical term for neutral hydrogen.) The dependence on $\Omega_m h^2 (1+z)^3$ is due to the expansion of the universe. There is no significant amount of uniformly distributed neutral hydrogen.

The next most obvious location for neutral hydrogen is in clouds in the intergalactic medium. A uniform distribution of gas is unstable, so the baryons are expected to clump into clouds. Discrete Ly-α absorption lines indicate the column density of neutral hydrogen in clumped form between us and the quasar. Observations give results similar to the upper limit from the Gunn-Peterson test: a very small fraction of the baryons in the universe is in the form of clumped, neutral hydrogen, of order 10^{-7} to 10^{-6} of the critical density.

Observations of such clouds indicate that the gas is actually quite warm. The thermal broadening of the absorption lines indicates $T \sim 5 \times 10^4$ K. The temperature corresponding to the ionization energy for neutral hydrogen is 13.6 eV$/k_B = 1.5 \times 10^5$ K. Clearly, a large fraction of the baryons may be in the form HII, ionized hydrogen. UV emission from quasars and young stars provides the energy to heat these clouds to such high temperatures (the CBR temperature is $2.73(1+z)$ K, totally negligible in comparison). Though HII is not directly visible in the absorption

line spectrum, and it is too cold and thinly distributed to emit significant amounts of recombination radiation, it is possible to estimate $\Omega_{\rm HII}$ by assuming the clouds have reached thermal equilibrium with the ionizing flux. The ionizing flux is determined from the continuum flux observed longward of the Ly-α peak; the number density of quasars is also observable. The ionization and recombination rates are straightforward atomic physics. From these numbers, the neutral fraction in these clouds can be calculated and thus the observed Ly-α absorption extrapolated to give the total amount of hydrogen in the clouds. A similar analysis can be performed with helium in such clouds. The resulting lower limit on Ω_b from such systems is

$$\Omega_b h^2 \gtrsim 0.018 \tag{1.71}$$

(see, for example, [38]). A lower limit is given because of various effects that cannot be ruled out. The clouds could be hotter than is believed, increasing the ionized fraction and thus the total column density implied by the HI column density. The ionizing flux may be underestimated, also leading to an underestimate of the ionized fraction. It is comforting, however, that the lower limit is fully consistent with the BBN expectation for Ω_b. Thus, at high redshift, the baryons are primarily in the form of ionized hydrogen and helium in the Ly-α forest.

1.5.2 Dark Baryons at Low Redshift

The Ly-α forest disappears by $z \sim 1$. The dark baryons today are not in the form of ionized intergalactic gas with temperature $10^4 - 10^5$ K. Obviously, then, the best place to look for dark baryons is in galaxies.

Some obvious candidates are easily ruled out. Clouds of atomic hydrogen can be seen by hyperfine emission, and, in fact, are used to probe galactic rotation curves at large radii. Clouds of ionized hydrogen and helium are visible via UV thermal-bremsstrahlung radiation and recombination emission. Massive black holes are visible via X-ray and gamma-ray emission from accreting matter. Searches for very dim, red stars have been performed with the Hubble Space Telescope and no significant component has been found [39].

The most popular candidate for baryonic dark matter in recent years has been MAssive Compact Halo Objects (MACHOs). This is a generic term for any dark compact object (*i.e.*, not gas clouds) that inhabits the halo of our galaxy. Brown dwarfs are an excellent MACHO candidate. Brown dwarfs are collapsed, sub-stellar objects that are too light to develop the core pressure and temperature to burn hydrogen (the hydrogen-burning limit is $0.08\,M_\odot$) but that do burn deuterium. Also possible are Jupiter-type objects with masses $\sim 10^{-3}\,M_\odot$. More generally, one may ask whether there is an allowed range for the masses of objects that form from the collapse of gas clouds. An initial uniform distribution of gas of density ρ breaks up into smaller clouds that turn into stars and smaller objects. The dynamics of such a system are set by the ratio of the cooling time to the gravitational collapse time: if a cloud cannot cool by emission of radiation faster than it collapses, then it heats up as it contracts and is unstable against fragmentation. This ratio is [1]

$$\frac{t_{cool}}{t_{coll}} \sim \left(\frac{0.1 M_\odot \rho^{1/10}}{M}\right)^{5/3} \tag{1.72}$$

where the collapse time is set by dimensional analysis, $t_{coll} \sim (G\rho)^{-1/2}$ and the cooling time is

calculated by assuming perfect blackbody emission. The result is that clouds satisfying

$$M \lesssim 0.1\rho^{1/10} M_\odot \tag{1.73}$$

can collapse without fragmentation. This sets a lower limit on MACHO masses: no objects smaller than this size can form via fragmentation. The local halo density, as determined from the rotation curve, is $\rho \approx 10^{-24}$ g cm^{-3}, which yields a lower limit of about 10^{-3} $M_\odot$. The existence of planets in the solar system does not contradict this bound; for these purposes, the solar system should be thought of as one object with $M \approx M_\odot$. Thus, there is a physically plausible range for MACHO candidates of 10^{-3} $M_\odot$ to 0.08 $M_\odot$.

Two major experiments, MACHO and EROS, were mounted during the 1990s to look for MACHOs in the galactic halo. Both employed the technique of microlensing. A massive object near the line of sight to a background star gravitationally deflects light from the star. The deflection deforms the solid angle "emitted" from the star that reaches the observer. With large deflectors, such as galaxy clusters, this results in rings and arcs, as discussed earlier. In the case of nearby stars whose emitted light is deflected by objects of sub-solar mass, the distortions are too small to see; however, the net effect, an increase in the solid angle "emitted" by the star that the observer sees, increases the apparent magnitude of the star. The amplification factor is

$$A = \frac{1 + 2\theta_E^2/\theta_d^2}{\sqrt{1 + 4\theta_E^2/\theta_d^2}} \tag{1.74}$$

where θ_d is the angle between the line from the observer to the lens and the line-of-sight to the star and $\theta_E = \sqrt{4Gm/R}$ is the Einstein radius, a characteristic scale that determines how close to the line-of-sight to the star the deflector of mass m and at distance R from the observer must pass. At $\theta_d = \theta_E$, the amplification is 1.34; these two points on the magnification curve are conventionally used to define the characteristic width of the event. Assuming the MACHOs are distributed in an isothermal spherical halo with rms velocity 270 km s^{-1} (velocity dispersion $\sigma = 220$ km s^{-1}), the effect is detectable as a transient, achromatic increase in flux from the background star. The time profile is symmetric with characteristic width

$$\hat{t} = \frac{2R\theta_E}{v_\perp} \tag{1.75}$$

where $v_\perp$ is the velocity of the MACHO perpendicular to the line-of-sight and assuming the distance R does not vary appreciably during the event. Note that the MACHO velocity, distance, and mass (via θ_E) all enter this expression, but only the combination of the three is determined in any given event. However, with a model for the spatial and velocity distribution of the MACHOs and an ensemble of events, it is possible to constrain or exclude MACHO candidates of a given mass.

The MACHO and EROS collaborations mounted searches for MACHOs by making nightly observations of the Large Magellanic Cloud, Small Magellanic Cloud, and the galactic bulge (the spheroidal component at the center of the galaxy) with dedicated telescopes. Both experiments took images in red and blue filters to take advantage of the achromatic signal. By monitoring millions of stars continuously, it was possible to detect a handful of microlensing events over a few years. After six years of observation, the MACHO team observed 13 to 17 events toward the LMC, depending on one's classification of a few of the events. EROS, with smaller exposure time, found

4 events toward the LMC. In addition, one event toward the SMC was detected by both groups. Two significant conclusions can be made from these data, under the assumption of a standard, isothermal, spherical halo. First, the absence of any short duration events, $\hat{t} < 20$ days, implies that no more than 20% of the halo can be made of MACHOs with masses between 10^{-4} $M_\odot$ and 0.03 $M_\odot$ [40]. Lower mass MACHOs would produce events too short to be detected, but, as discussed above, the existence of MACHOs below 10^{-3} $M_\odot$ is not possible. Second, the detected events indicate the most likely MACHO mass and halo fraction are approximately 0.5 $M_\odot$ and 0.2, with the 68% CL region extending from approximately 0.2 $M_\odot$ to 0.9 $M_\odot$ and from 0.1 to 0.3, respectively. An all-MACHO halo is ruled out at 95% CL. The EROS team chooses to place an upper limit on the halo fraction because they interpret the SMC event as an indication of lack of understanding of possible backgrounds, such as microlensing by stars in the LMC itself; *i.e.*, EROS makes no claim to have detected a MACHO halo. In any case, it is quite clear that MACHOs do not make up the dominant part of the dark halo, though they may indeed exist.

Other arguments exist against the proposition that fairly high-mass stellar remnants contain most of the baryon density today. With the mass range given above, the only reasonable dark stellar-remnant candidate for MACHOs is white dwarfs. However, white dwarfs are end products of stellar evolution, and thus a large mass in white dwarfs requires the other fossils of stellar evolution to be visible. These arguments have been known for some time but were recently resurrected and made precise by Fields, Freese, and Graff [41, 42, 43]. First, white dwarfs are the remnants of much larger stars, so the ejecta from the stars that became the white dwarfs must be present. Second, production of white dwarfs coincides with production of elements heavier than helium (metals); these metals should be visible. Finally, the progenitor stars would contribute to the cosmic infrared background light. Multi-TeV gamma-rays from cosmologically distant sources would interact with this background light, resulting in spectral cutoffs that are not observed. The resulting bound on the contribution of white dwarf MACHOs to the baryon density of the universe is

$$\Omega_{WD} \lesssim 0.006 h^{-1} \tag{1.76}$$

Where does this leave us with regard to the MACHO contribution to the overall baryon density? Fields *et al.* calculated the relationship between halo mass in MACHOs and overall MACHO contribution to the universe [43]. Rescaling their numbers to the final MACHO results gives

$$\Omega_{MACHO} = (0.0016 - 0.0069) h^{-1} f_{gal} \tag{1.77}$$

with $f_{gal} \geq 0.17$ being the fraction of galaxies that contain MACHOs. An all-MACHO halo ($\Omega_{MACHO} \sim 5\times$ higher) would have been consistent with BBN bounds on Ω_b but strongly inconsistent with the arguments regarding by-products of white-dwarf production. The consistency of the MACHO results with the upper bound on the white-dwarf contribution to the mass density suggest the picture holds together. However, with this low halo mass fraction in MACHOs, it appears that roughly 80% of the dark baryons at $z = 0$ remain missing. Furthermore, 80% of the halo dark matter is also missing. There thus remains good motivation to search for nonbaryonic dark matter in our galaxy.

Cen and Ostriker have proposed that the dark baryons are not actually dark, but in fact are so hot that observations to date have been insensitive to them [44]. They demonstrate by simulations that any remaining intergalactic gas is heated to 10^5 to 10^7 K, significantly hotter than Ly-α forest clouds. The neutral hydrogen fraction is so low and thermal broadening so large that

Ly-α absorption by this gas would not yet have been detected. Since this material cannot produce stars, there are no worries about gas ejection, metal production, or infrared background light as with a stellar-remnant solution. The ionized component of the gas may in fact be detectable by extreme UV and soft X-ray (< 1 keV) emission from thermal bremsstrahlung; it would provide a significant component of the soft X-ray background. Cen and Ostriker propose a number of signatures in the X-ray background that would arise from expected structure in the gas. While this proposal is quite plausible, it is fair to say that the dark baryons at low redshift are still hiding. Furthermore, this gas, being intergalactic, has no bearing on the galactic dark-matter problem.

1.6 Nonbaryonic Dark Matter

As argued above, there is significant evidence that the matter density of the universe is dominated by nonbaryonic dark matter. The absence of a significant MACHO halo indicates that this dominance may hold for the halo of our galaxy, too. I discuss some of the candidates here, WIMPs in particular.

1.6.1 Primordial Black Holes

Primordial black holes are a generic class of black holes produced in the very early universe, before nucleosynthesis. By locking up their baryons before nucleosynthesis, the BBN limit on Ω_b is circumvented. Moreover, PBH's do not act like baryonic matter — the baryons are not free to interact with photons, so they circumvent the structure-formation argument also. There are a variety of mechanisms for production of PBH's. Most involve modifications to or special characteristics of the quantum fluctuations at the end of inflation. Others rely on strings or other topological defects. Discussion of these is beyond the scope of this dissertation (and the abilities of the author!).

One simple fact does always hold: when PBHs are created, they accrete essentially all of the mass within their horizons. This relates the mass of the PBHs to their time of production [45]:

$$M_{PBH}(t) \approx 10^{15} \left(\frac{t}{10^{-23}\,\mathrm{s}} \right) \mathrm{g} \tag{1.78}$$

In particular, in order to segregate baryons before nucleosynthesis, which occurs at approximately 1 minute, PBHs may have masses no larger than about $10^6\ M_\odot$. On the high-mass end, more restrictive limits can be placed simply because the radiation emitted by accretion onto a significant population of black holes of mass $\gtrsim 10^4\ M_\odot$ would have already been seen. The MACHO exclusion results also apply, disallowing a significant component of PBHs in the mass range 10^{-4} to $0.03\ M_\odot$ in the galactic halo.

Detection of very low mass primordial black holes is possible by their Hawking radiation. The mechanism for generation of Hawking radiation can, very hand-wavingly, be described as follows. Virtual particle-antiparticle pairs pop in and out of existence all the time. If such an event occurs near the horizon of the black hole and one of the particles crosses through the black hole horizon, the pair cannot annihilate. Thus, the remaining particle becomes real. To conserve energy, the black hole must lose the energy required to create the particle. Effectively, the black hole emits

elementary particles. The emission spectrum is a blackbody spectrum. The temperature of the blackbody and the mass loss rate are related to the mass of the black hole by [45]

$$T \approx \left(\frac{10^{13}\,\mathrm{g}}{M}\right)\,\mathrm{GeV} \tag{1.79}$$

$$\frac{dM}{dt} = -5 \times 10^{-25}\,M^{-2}\,f(M)\,\mathrm{g\,s^{-1}} \tag{1.80}$$

Thermal runaway is thus possible: a hot black hole emits particles, its mass decreases, its temperature increases, the energies of the particles emitted increases, so the mass loss occurs faster. The black hole eventually evaporates in a flash of energetic particles. The lifetime of the black hole is [45]

$$\tau(M) = 6 \times 10^{-27}\frac{M^3}{f(M)}\,\mathrm{s} \tag{1.81}$$

$f(M)$ is related to the number of particle species light enough to be emitted by the hole; it is taken to be 1 for holes larger than 10^{17} g that may only emit massless particles. The mass corresponding to PBHs that are just evaporating now is 10^{15} g. The lifetime is a very strong function of M, so 10^{15} g is effectively a lower limit on the mass of PBHs that can act as nonbaryonic dark matter yet not have flooded the universe with relic radiation.

The constraints on the number density of black holes that are evaporating today, $M \sim 10^{15}$ g, are very stringent. The resulting gamma rays and cosmic-ray antimatter between tens of MeV and a few GeV would be visible. The experiments described in Chapter 2 capable of indirectly detecting WIMP annihilation through gamma rays or cosmic-ray antimatter place stringent limits on Ω_{PBH} of $10^{-9} - 10^{-8}$ [45]. However, increasing the PBH mass by three orders of magnitude decreases the mass-loss rate by a factor of 10^6 and the average energy of the emitted particles to 10 keV, rendering the Hawking radiation undetectable.

When the initial MACHO results came out, suggesting a large ($\sim 50\%$) halo fraction, some authors proposed PBHs of order $0.5\,M_\odot$ would provide the necessary microlensing candidates without messy stellar remnants [46]. Such PBHs would have formed at 10^{-5} s, much earlier than nucleosynthesis and thus still circumventing the BBN bound. However, with the downward revision of the MACHO halo fraction to $\sim 20\%$, PBHs are no longer needed.

PBHs are thus perfectly plausible as "nonbaryonic dark matter." A significant objection to PBHs is that there is not yet any stand-out theory with predictive power: PBHs can be produced at any time during the very early universe, so there is no particular mass that is preferred. This lack of a concrete prediction is disappointing. Finally, there is a large range of PBH masses that are reasonable and yet essentially undetectable, ranging from 10^{18} g $= 10^{-15}\ M_\odot$ to 10^{28} g $= 10^{-5}\ M_\odot$. PBHs ranging from 10^{32} g $= 0.1\ M_\odot$ to 10^{36} g $= 10^3\ M_\odot$ are presently allowed, and may indeed be detectable in large astronomical survey programs by their accretion-induced radiation.

1.6.2 Axions

The axion, a particle proposed to solve the strong CP problem of particle physics, is also a good candidate for nonbaryonic cold dark matter. This discussion of the motivation, properties, and detectability of the axion follows that of [2]. The strong CP problem arises because the $SU(3)_C$ gauge group of the strong interaction has a nontrivial vacuum structure — there are multiple,

degenerate vacua. This is possible in $SU(3)_C$ because there are more than three gauge fields. The true vacuum is a superposition of the degenerate vacua, $|n\rangle$:

$$|\Theta\rangle = \sum_n \exp(-in\Theta)|n\rangle \tag{1.82}$$

where Θ is *a priori* arbitrary. For reasons too technical to be described here, the effects of the multiple vacua can be rewritten as a single non-perturbative term in the QCD Lagrangian

$$\mathcal{L}_v = \Theta \frac{g^2}{32\pi^2} G\widetilde{G} \tag{1.83}$$

where G is the gluon field strength and $\widetilde{G}$ is its dual. Such terms do not occur in the electroweak theory. This term is CP-violating, leading to an electric dipole moment of the neutron well above experimental limits. Θ must therefore be very close to zero, but there is no apparent reason.

This problem is solved by the Peccei-Quinn axion. The Θ parameter is taken to be a dynamical variable. A new global symmetry, $U(1)_{PQ}$, is assumed to exist but to be spontaneously broken at an energy scale f_{PQ}. As with electroweak-symmetry breaking, a massless Goldstone boson, the axion (a), is generated. Again, for reasons too technical to go into here, perturbative effects give the axion a mass and also yield a term in the Lagrangian that forms a classical potential for the axion with the original $G\widetilde{G}$ term. The value of a, $\langle a\rangle$, that minimizes this potential results in cancellation of the two $G\widetilde{G}$ terms. The quantized oscillations about the minimum of the potential are the axion; the axion has absorbed the Θ degree of freedom, but in a way that all remaining $G\widetilde{G}$ terms include an axion field (*i.e.*, appear as axion-gluon-gluon vertices) and no $G\widetilde{G}$ term with a constant coefficient remains. The axion satisfies an interesting relation with the pion due to the way its mass is generated:

$$f_{PQ} m_a = C f_\pi m_\pi \tag{1.84}$$

where C is a factor of order 1, $f_\pi = 93$ MeV is the pion decay constant, and $m_\pi = 135$ MeV is the pion mass. The resulting axion mass is proportional to $1/f_{PQ}$ and is *a priori* arbitary because f_{PQ} is arbitrary; this arbitrariness introduces no problems, as any axion mass solves the strong CP problem. The axion is a scalar (there was no reason to introduce additional degrees of freedom) and may have interactions with a variety of particles.

Cosmologically, axions are produced in a very interesting way. In the early universe, the axion temperature is large compared to the energy scale of it classical potential and the axion field may take any value. As the universe cools through the QCD scale, $\Lambda_{QCD} \sim 300$ MeV, the axion field falls into its potential and oscillates about the minimum. For reasons not to be explained here, the oscillations are coherent over the horizon size at the time this occurs. Coherence is equivalent to small momentum: the phase of the axion field does not vary over a length equal to the horizon. In fact, the axions form a Bose condensate. Axions thus produced act as cold dark matter because they are completely nonrelativistic despite the high temperature of production relative to their mass. Moreover, because they form a Bose condensate, it is very difficult for them to acquire kinetic energy — any excited axion falls back into the condensate. Calculating the relic density of axions produced this way is nontrivial and leads to the expectation (see [2])

$$\Omega h^2 = 0.85 \times 10^{\pm 0.4} \left(\frac{\Lambda_{QCD}}{200\,\mathrm{MeV}}\right)^{-0.7} \left(\frac{m_a}{10^{-5}\,\mathrm{eV}}\right)^{-1.18} \tag{1.85}$$

which holds only if axions are produced after inflation. The inverse dependence on m_a leads immediately to the overclosure bound $m_a \gtrsim 10^{-6}$ eV. This expression does not include the contribution of axionic strings.

Other limits can be derived due to the axion's interaction with other particles. The axion looks very much like a pion, except it possesses additional interactions with electromagnetic particles, again due to the technical details of the way its mass is generated. Thus, the axion may decay to two photons, $a \to \gamma\gamma$, as the pion does. It may also participate in

$$\gamma + e^- \to a + e^- \qquad \text{Compton} \tag{1.86}$$

$$e^- + Z \to e^- + Z + a \qquad \text{bremsstrahlung} \tag{1.87}$$

$$e^- + e^- \to e^- + e^- + a \qquad \text{free} - \text{free} \tag{1.88}$$

$$\gamma + Z \to a + Z \qquad \text{nuclear Primakoff} \tag{1.89}$$

$$\gamma + e^- \to a + e^- \qquad \text{electron Primakoff} \tag{1.90}$$

where Z is a nucleus of charge Z, not the Z boson. It is these reactions that distinguish the axion from the pion. These reactions have matrix elements proportional to $1/f_{PQ}$. There are two primary axion models. The DFSZ axion couples to hadrons, photons, and electrons at tree level. The "hadronic" axion has no tree-level coupling to the electron. The pion-decay-like process $a \to \gamma\gamma$ still occurs, so the axion has a smaller coupling to the electron through loop diagrams.

Because of the axion's similarity to the pion, rare decays of kaons and quarkonium may yield axion end products. Since the axion coupling scales as $1/f_{PQ} \propto m_a$, such rare-decay rates increase as the axion mass increases. Present limits on the existence of these rare decays limit the axion mass from above at $\sim$10 keV.

The other reactions noted above also provide bounds on the axion mass. The axion can be produced in stars and carry energy away. As the axion mass increases, these reaction rates increase as $1/f_{PQ}^2 \propto m_a^2$. In order to produce the same photon output as is observed, the Sun's core must therefore increase in temperature as the axion mass increases. No direct effect would be seen because the photon output would be unchanged. However, the neutrino output, which is observable, would also be increased. The present knowledge of the solar-neutrino flux limits the axion mass from above at $\sim$1 eV. This argument fails if the axion is heavier than $\sim$10 keV because the solar core temperature is only a few keV.

Axion-mass limits are also available from red giants. Hadronic axions would speed the cooling of red giants. In addition, DFSZ axions, with their stronger coupling to electrons, might prevent the thermal-runaway process that gives rise to "helium flash," the ignition of fusion of helium in the core. The lower limits are of order 1 eV for hadronic axions and of order 0.01 eV for DFSZ axions. Finally, the observed neutrino burst from supernova 1987A constrains the speed with which the core could have cooled after collapse. Axions with masses greater than about 10^{-3} eV would have shortened the neutrino burst.

The resulting open range for the axion mass is about 10^{-6} eV to 10^{-3} eV. The favored search mechanism is $a \to \gamma\gamma$, but with one photon provided by a magnetic field via the coupling $\vec{E} \cdot \vec{B}$, where $\vec{B}$ is provided by a static magnetic field and $\vec{E}$ is the photon produced. To boost the signal relative to the amplifier noise, the searches are done in resonant radiofrequency cavities. The

power emitted into the cavity due to axion decay has a fairly simple form [2]:

$$P_{nl} \approx 2 \times 10^{-22}\,\mathrm{W}\frac{V}{10\,\mathrm{L}}\left(\frac{B}{6\,\mathrm{T}}\right)^2 C_{nl}\frac{m_a}{10^{-5}\,\mathrm{eV}}\min\left(\frac{Q_{nl}}{Q_a},1\right) \tag{1.91}$$

where the nl subscript indicates the cavity mode (*e.g.*, TM_{10}), V is the cavity volume, B is the magnetic-field strength, C_{nl} is a factor of order unity dependent on the mode, Q_{nl} is the cavity quality factor for the nl mode, and Q_a is the intrinsic energy dispersion of the axions, $Q_a = E_a/\Delta E_a \sim 10^6$. Since resonant cavities are used, the cavity tuning frequency must be swept to probe different axion masses. The signal-to-noise ratio is given by

$$\frac{S}{N} = \sqrt{\frac{t}{\Delta\nu}}\frac{P_{nl}}{k_B T_n} \tag{1.92}$$

where t is the time spent integrating at a given tuning point, $\Delta\nu$ is the bandwidth over which the signal and noise are integrated at that point, and T_n is the noise temperature of the receiver mounted to the cavity. T_n is defined so that $k_B T_n\sqrt{\Delta\nu/t}$ is the root-mean-square power received in bandwidth $\Delta\nu$ averaged over a time t, so the relation above is the signal to noise. The dependence on bandwidth is nontrivial. $Q_{nl} \approx \nu_{nl}/\Delta\nu$: the cavity is resonant only for the bandwidth $\Delta\nu$, so there is no point in accepting a larger bandwidth from the receiver. So, nominally, S/N scales as $\Delta\nu^{-3/2}$. Decreasing the bandwidth appears to help. But, since the axion mass is unknown, and the search is performed over some fixed mass range Δm_a, the length of the experiment τ scales as $\Delta m_a t/h\Delta\nu$: the smaller $\Delta\nu$ is, the more times the resonant frequency has to be changed. The more interesting figure of merit is the signal to noise over a fixed mass range that can be achieved per year the experiment runs:

$$\frac{S}{N\tau} = \sqrt{\frac{1}{t\Delta\nu}}\frac{(P_{nl}\Delta\nu)}{T_n\Delta m_a h} \tag{1.93}$$

Note that P_{nl} scales as $1/\Delta\nu$, so the quantity $(P_{nl}\Delta\nu)$ is independent of bandwidth. The figure of merit improves as the bandwidth is decreased, but not as fast as suggested by the previous equation. The figure of merit is actually hurt by increasing t; one should make t as small as possible while optimizing against fake detections that have to be rescanned.

The present best results are from the LLNL-MIT team, using HEMT receivers in a standard cavity search. The present exclusion-limit plot is shown in Figure 1.14. Hadronic (KSVZ) axion models are being probed, but DFSZ models remain untouched. The quickest ways to gain in the figure of merit are to increase the magnetic field, the volume of the cavity, cavity Q, or to decrease the noise temperature of the receiver. Since price is a stronger than linear function of cavity size and magnetic field, most groups have concentrated on receiver improvements. The LLNL-MIT team are collaborating with John Clarke to switch from HEMT receivers, with $T_n \sim$ few K, to RF SQUID receivers, with noise temperatures in the hundreds of mK. A Japanese group are attempting to use Rydberg atoms as receivers. Rydberg atoms are pumped to a very high excited state, passed through the axion-detection cavity, and hit with a microwave photon after emerging. If the atom absorbed an axion-generated photon while in the detection cavity, the electron would be in a higher energy state and the atom becomes ionized by the last microwave photon; otherwise, the electron is driven to a higher bound state. A photoelectron detector determines whether ionization occurred. Such a technique has the advantage of having no "quantum limit" — in conventional-receiver experiments, because both the phase and amplitude of the photons are measured, the

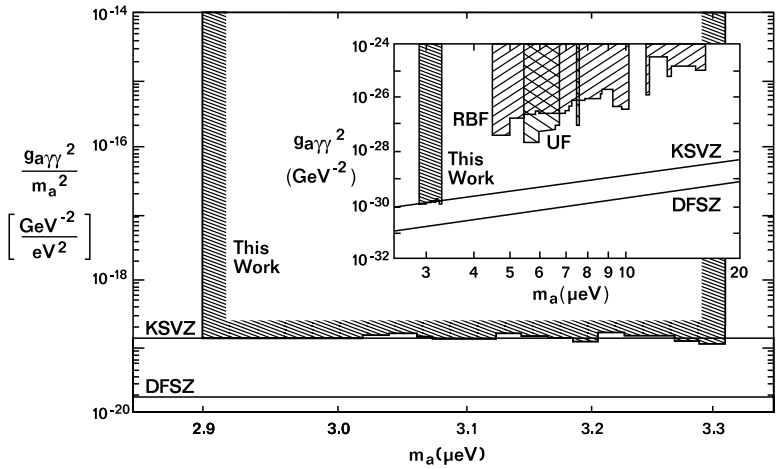

Figure 1.14: Exclusion limits on the axion mass and coupling from the LLNL-MIT axion search. The vertical axis is the square of the ratio of the axion-photon-photon coupling to the axion mass, the horizontal axis is the axion mass. Predictions for the KSVZ model (a particular hadronic-axion model) and the DFSZ model are shown. Note that the KSVZ hadronic axion has a *larger* photon coupling than the DFSZ axion. Figure taken from [47].

uncertainty principles imposes a minimum possible noise. The Rydberg-atom technique makes no phase determination, but simply detects particles. Both techniques should probe the DFSZ axion models in the next few years.

1.6.3 Light Neutrinos

Light neutrinos, with masses of a few eV to tens of eV, are disfavored as the dominant weakly interacting dark matter because they damp density fluctuations on scales smaller than the horizon size at radiation-matter equality, $\lambda = 13\,\Omega_m^{-1}h^{-2}$. However, a small admixture of neutrinos may reduce the power on small scales to the level required by observation, so I briefly discuss their relic abundance. The following argument is known as the Lee-Weinberg bound; I follow the derivation given in [1].

Neutrinos belong to the class of *hot relics* — particles that are relativistic when they fall out of thermal equilibrium with the rest of the universe. This process is called "freeze-out." The general criterion for freeze-out to occur is that the annihilation rate (per particle) of the particle species, which maintains equilibrium with other particle species, falls below the expansion rate:

$$\Gamma_A < H \tag{1.94}$$

That is, the average time between annihilations becomes comparable to the age of the universe. By detailed balance, this implies that the rate of the reactions producing such particles also becomes too low to maintain thermal equilibrium. The annihilation rate is given by

$$\Gamma_A = n\langle\sigma_A v\rangle \tag{1.95}$$

where n is the particle density, σ_A is the annihilation cross section, and v is the relative velocity of two particles. The angle brackets indicate averaging over the thermal distribution. The typical interaction cross section is given by a simple weak interaction (the relative velocity v is $c = 1$),

$$\sigma_A \propto G_F^2 E_\nu^2 \tag{1.96}$$

where G_F is the Fermi coupling constant that characterizes weak interactions at energies much below M_W and E_ν is the neutrino energy and arises from phase-space considerations, assuming annihilation to massless species. The number density of neutrinos while still in equilibrium is easy to calculate because they are still relativistic at freeze-out (as is seen below); it is

$$n_\nu = \frac{3\zeta(3)}{2\pi^2}\left(\frac{k_B T_\nu}{\hbar c}\right)^3 \tag{1.97}$$

which is derived in the usual manner from the Fermi-Dirac distribution. As mentioned very early in this chapter, the neutrino and photon temperatures are not the same for $T \lesssim 500$ keV due to electron-positron annihilation. The neutrino temperature scales as $1 + z$, so the present-day value of 1.95 K may be used to extrapolate to earlier times. Doing the algebra yields

$$n_f\langle\sigma_A v\rangle_f H_f^{-1} \sim 0.026\left(\frac{T_f}{1\text{ MeV}}\right)^3 \tag{1.98}$$

where f subscripts indicate values at freeze-out. T_f is the same for neutrinos and photons because it is assumed that neutrino decoupling occurs before electron-positron annihilation. Requiring the above to be unity gives $T_{\nu,f} = 3.4$ MeV, validating the assumptions $T_{\nu,f} \gtrsim 500$ keV and $T_{\nu,f} \gg m_\nu$.

For hot relics, the relic abundance is easy to calculate because their number density up to and through freeze-out is given simply by a blackbody distribution for massless particles — the

actual time of freeze-out plays no role except to establish that the particles are indeed relativistic at freeze-out. Applying Equation 1.97 with $T_\nu = T_{\nu,0} = 1.95$ K gives $n_{\nu,0} = 113$ cm^{-3}. The resulting density parameter in neutrinos is

$$\Omega_\nu = \frac{8\pi G}{3H_0^2} n_{\nu,0} m_\nu = 0.011 \left(\frac{m_\nu}{1 \text{ eV}} \right) h^{-2} \tag{1.99}$$

The overclosure bound on the neutrino mass is thus $91h^2$ eV. A neutrino mass much larger than a 1 eV is ruled out by structure-formation arguments.

Experimentally, there is good motivation for non-zero neutrino mass, though not of the size needed to have significant cosmological effects. The solar-neutrino problem is the deficit of ν_e from the Sun as compared to expectations from solar models. A variety of the fusion reactions in the Sun produce neutrinos in numbers closely linked to the overall luminosity of the Sun. The deficit has been measured numerous times using different neutrino-detection techniques over the last 30 years and is a solid result. It could be solved by invoking the oscillation $\nu_e \leftrightarrow \nu_\mu$ with $\Delta m^2 \sim 10^{-5}$ eV^2 [8]. The atmospheric neutrino problem is usually stated in terms of the deficit of the ratio $(\nu_\mu + \overline{\nu_\mu})/(\nu_e + \overline{\nu_e})$ relative to expectations. Atmospheric neutrinos are generated in cosmic-ray air showers by the process $\mu^- \rightarrow e^- + \overline{\nu_e} + \nu_\mu$. While the flux of any single neutrino species is subject to uncertainties in the cosmic-ray flux and shower development, the ratio of the muon-neutrino flux to the electron-neutrino flux should be robust because of the link through the production reaction. This deficit would be solved by $\nu_\mu \leftrightarrow \nu_\tau$ with $\Delta m^2 = 2 \times 10^{-3}$ eV^2 [8]. Thus, a consistent picture would arise with $m_{\nu_e} = 0$, $m_{\nu_\mu} = 0.003$ eV, and $m_{\nu_\tau} = 0.045$ eV. This scenario is not of particular cosmological interest, however.

The LSND experiment claims to have observed the oscillation $\nu_\mu \leftrightarrow \nu_e$ via appearance of ν_e from a beam of ν_μ [48]. Their result is much disputed, primarily because of doubt in their estimate of the background ν_e flux into their detector. However, if taken at face value, the resulting Δm^2 is in the range 0.2 eV^2 to 15 eV^2, which would yield nonnegligible neutrino energy densities. The nominal mass hierarchy would be a massless or negligibly massive electron neutrino followed by degenerate muon and tau neutrinos with mass between 0.45 eV and 3.9 eV. A sterile light neutrino must be invoked to solve the solar-neutrino problem. The high-mass end of this scenario is disfavored by structure-formation arguments. Furthermore, the portion of the LSND 90% CL allowed region above 2 eV^2 is ruled out by the KARMEN2 experiment at 90% CL [49]. However, a mass between 0.45 eV and about 1.5 eV is still permitted. The MiniBooNE experiment, currently in preparation at Fermilab, will test the entire LSND allowed region [50].

1.6.4 WIMPs

The final widely discussed category of nonbaryonic dark-matter candidates is WIMPs, or Weakly Interacting Massive Particles. "Weakly interacting" specifically means that these particles interact only via the weak interaction (and gravity). "Massive" means sufficiently massive to solve the structure-formation problems brought on by a baryons-only, or even a neutrinos+baryons, universe. This would require a WIMP mass of at least 1 keV. The argument given below, also a Lee-Weinberg-type bound, demonstrates that WIMPs are well motivated only if their mass is above a few GeV. I follow the derivation of [1].

The essence of the argument is that the relic abundance of a *cold relic*, a particle δ that falls out of thermal equilibrium after it has become nonrelativistic, is directly related to its annihilation

rate at freeze-out and that this annihilation rate is characteristic of the weak interaction if the particle is cosmologically significant ($\Omega_\delta \sim 1$). As discussed above in the case of BBN and light neutrinos, freeze-out occurs when the rate of the annihilation reactions that create and destroy the particle δ drops below the expansion rate of the universe:

$$\langle \sigma_A v \rangle_f n_f \sim H_f \tag{1.100}$$

where the f subscript indicates values at freeze-out, σ_A is the typical annihilation cross section for $\delta\overline{\delta} \rightarrow$ anything, v is the WIMP relative velocity, the angle brackets indicate thermal averaging, and n_f is the WIMP number density. To calculate $\langle \sigma_A v \rangle_f$, it is necessary to determine n_f and H_f.

The abundance at freeze-out, n_f, can be related to the relic energy density contributed by δ:

$$\Omega_\delta = \frac{8\pi G}{3H_0^2}\rho_\delta = \frac{8\pi G}{3H_0^2}\frac{n_f M_\delta}{(1+z_f)^3} \tag{1.101}$$

where z_f is the redshift of freeze-out. The redshift of freeze-out is related to the freeze-out temperature through the present temperature of the CBR

$$T_f = \left(\frac{1+3\frac{7}{8}\frac{4}{11}}{g_f}\right)^{1/3}(1+z_f)T_{CBR} \tag{1.102}$$

The prefactor results from changes in the number of degrees of freedom contributing to the entropy of the universe; today, the photon and 3 neutrino families are the only degrees of freedom carrying significant entropy, giving the numerator, while at freeze-out there are more; electrons and positrons are possible if $T_f \gtrsim 1$ MeV, others if T_f is higher. The effective number of degrees of freedom, g, was defined by Equation 1.19. Combining Equations 1.101 and 1.102 gives

$$n_f = \frac{3}{8\pi G}\frac{\Omega_\delta H_0^2}{M_\delta}\left(\frac{g_f}{1+3\frac{7}{8}\frac{4}{11}}\right)\left(\frac{T_f}{T_{CBR}}\right)^3 \tag{1.103}$$

The abundance at freeze-out is related to the relic abundance by volume dilution, which is given explicitly by the ratio of the temperature at freeze-out to the temperature today. The Hubble parameter at freeze-out, H_f, has a straightforward form because the universe is radiation dominated at the high temperatures of interest:

$$H_f^2 = \frac{8\pi G}{3}\rho_{r,f} = \frac{8\pi G}{3}\frac{4\sigma_B T_f^4}{c^3}g_f \tag{1.104}$$

The energy density is just due to radiation, with $u = 4\sigma_B T_f^4/c$, with g_f giving the effective number of relativistic degrees of freedom. Solving Equation 1.100 for $\langle \sigma_A v \rangle_f$ thus gives

$$\langle \sigma_A v \rangle_f = \frac{H(z_f)}{n_f} \sim \frac{10^{-27}}{\Omega_\delta h^2} g_f^{-1/2}\frac{M_\delta}{T_f}\,\mathrm{cm}^3\,\mathrm{s}^{-1} \tag{1.105}$$

where M_δ and T_f are in units of GeV. The annihilation cross section is related to the relic density and the freeze-out temperature in units of M_δ.

The ratio M_δ/T_f is well constrained under the assumption that the WIMP has a significant relic abundance. This can be seen by relating n_f, the abundance at freeze-out, to Ω_δ, the present

abundance. n_f is approximately given by the thermal Boltzmann distribution for a massive particle at the freeze-out temperature, T_f:

$$n_f = 4\left(\frac{M_\delta k_B T_f}{2\pi\hbar^2}\right)^{3/2} \exp\left(-\frac{M_\delta c^2}{k_B T_f}\right) \tag{1.106}$$

(M_δ and T_f in kg and K, respectively). Because the particle is massive, a classical Boltzmann distribution suffices. The particles are assumed to be energetic or sparse enough that Pauli exclusion or Bose condensation are not important. Equations 1.103 and 1.106 relate the mass, freeze-out temperature, and relic density of the particle δ: the first relates its abundance today to its abundance at freeze-out via the ratio of the present temperature and the freeze-out temperature, the second relates the freeze-out abundance and temperature. Combining these two equations gives

$$\frac{M_\delta}{T_f} - \frac{3}{2}\log\left(\frac{M_\delta}{T_f}\right) = 18.9 + \log\left[\frac{M_\delta}{\Omega_\delta h^2 g_f}\right] \tag{1.107}$$

where M_δ and T_f are again in GeV. The constant term comes from the various constants. The logarithmic term on the right side may be neglected if its argument is not too different from 1, which holds if the WIMP is cosmologically interesting and $M_\delta \sim 1$ GeV. Under this assumption, $M_\delta/T_f \sim 24$. This fixes the relation between the relic density and annihilation cross section at freeze-out, Equation 1.105. For $\Omega_\delta h^2 \sim 1$, the annihilation cross section must be $\langle\sigma_A v\rangle_f \sim 10^{-26}$ cm^3 s^{-1}.

It can be demonstrated that the above annihilation cross section is typical of a particle that interacts by the weak force (and gravity). Let the annihilation process have matrix element $\mathcal{M}$. Then, by Fermi's Golden Rule, the annihilation rate (per unit particle number density) is

$$\langle\sigma_A v\rangle = \Gamma = \frac{2\pi}{\hbar}|\mathcal{M}|^2\rho(E) \tag{1.108}$$

where $\rho(E)$ is the density of states for the final state. Assuming δ annihilates primarily to relativistic particles, the density of states can be calculated trivially:

$$\rho(E) = 4\pi(2\pi\hbar c)^{-3}E^2 N \tag{1.109}$$

where N is the number of particle species δ may annihilate to. The energy available in the final state is $M_\delta c^2$. Solving for $|\mathcal{M}|^2$ gives

$$\begin{aligned} |\mathcal{M}|^2 &= \frac{\langle\sigma_A v\rangle_f \hbar^4}{M_\delta^2 c N} \\ &= 24\frac{10^{-27}\text{cm}^3\,\text{s}^{-1}}{\Omega_\delta h^2 g_f^{-1/2}}\frac{\hbar^4}{M_\delta^2 c N} \\ &\approx \frac{10^{-14}\,\text{GeV}^4\,\text{fm}^6}{\Omega_\delta h^2 g_f^{-1/2} N}(M_\delta \times 1\,\text{GeV})^{-2} \end{aligned} \tag{1.110}$$

where M_δ is now in GeV. For weak interactions, the characteristic matrix element is $G_F = (\hbar c)^3 \times 1.17 \times 10^{-5}$ GeV^{-2} $= 8.95 \times 10^{-8}$ GeV fm^3 $\sim |\mathcal{M}|$ if the assorted undefined factors are taken to be

unity. That is, a 1-GeV cold relic with a cosmologically interesting relic density has an annihilation cross section characteristic of the weak interaction.

An important part of the result, that there is only weak dependence of the argument on the particle's mass alone, derives from the dependences of the two important quantities on M_δ and T_f:

$$\begin{aligned} \Omega_\delta h^2 \propto \frac{n_f M_\delta}{T_f^3} &\propto M_\delta \left(\frac{M_\delta}{T_f}\right)^{3/2} \exp\left(-\frac{M_\delta}{T_f}\right) \\ \frac{1}{\langle \sigma_A v \rangle_f} \sim \frac{n_f}{H(z_f)} &\propto M_\delta \left(\frac{M_\delta}{T_f}\right)^{1/2} \exp\left(-\frac{M_\delta}{T_f}\right) \\ \Rightarrow \frac{\Omega_\delta h^2}{\langle \sigma_A v \rangle_f} &\propto \frac{M_\delta}{T_f} \end{aligned} \tag{1.111}$$

That is, everything depends mainly on M_δ/T_f, and if M_δ/T_f changes slightly, the Boltzmann factor exponentially suppresses or amplifies the relic density. M_δ/T_f is fixed, but otherwise, M_δ may vary. This robustness relies on the assumption of the WIMP being nonrelativistic at freeze-out, introducing the all-important exponential factor. Note that the relic abundance is *inversely* proportional to the annihilation cross section, modulo a mass factor. That is, the higher the annihilation cross section, the lower the relic abundance. This is perfectly reasonable: the larger the annihilation cross section, the fewer WIMPs left today. This point is important when discussing the relation between WIMP-nucleon interaction cross sections and relic abundance in Chapter 2.

At face value, almost any freeze-out annihilation rate picked out by the desire for an interesting relic density would have given a WIMP mass — if the annihilation rate were 12 orders of magnitude higher or 6 lower, WIMP masses of 1 keV or 1 TeV, respectively, would still yield a matrix element characteristic of weak interactions. However, recall Equation 1.107; if the WIMP mass moves away from 1 GeV by more than 2 or 3 orders of magnitude, the logarithmic correction terms pushes M_δ/T_f away from 24; the Boltzmann factor quickly results in either overclosure or irrelevance. Thus, masses in the GeV range are especially picked out by the argument.

An important aspect of the motivation for WIMPs lies in the fact that WIMP masses of 1 to 1000 GeV are also special from the point of view of particle physics for completely independent reasons. The Standard Model of particle physics is unstable in a theoretical sense. All the particle masses are below 100 to 200 GeV, yet radiative corrections should pull them up to the next interesting mass scale, the energy at which the electroweak and strong forces are joined by a Grand Unified Theory (GUT), $M_{\rm GUT} \sim 10^{16}$ GeV. The disparity between known particle masses and these scales suggests there is some intermediate scale, M_{new}, at a few hundred GeV to a few TeV, above which a symmetry is restored that stabilizes particle masses. Generically, symmetry-breaking energy scales in particle physics are associated with new particles. The electroweak scale, ~ 100 GeV, where electroweak symmetry is broken and the electromagnetic and weak interactions take different forms, is set by the masses of the W and Z bosons. The QCD scale, ~ 300 MeV, where chiral symmetry is broken, corresponds to the masses of the π and η mesons. Thus, generically, new particles are expected in the same mass range expected for WIMP dark matter.

It is this last incredible fact — that cosmology and particle physics independently point to the existence of a weakly interacting particle species with mass between a few GeV and a few TeV

— that motivates WIMP dark matter and, in particular, has convinced many to search for WIMPs via their weak interactions. I discuss methods for detection of WIMPs in detail in Chapter 2.

1.7 Supersymmetry and WIMPs

As stated above, particle physics provides a good generic rationale for the existence of WIMPs. In this section, I discuss in some detail supersymmetry, a well-motivated and presently most-favored extension to the Standard Model of particle physics, and its implications for WIMPs. The lightest new particle predicted by supersymmetry would make an ideal WIMP.

1.7.1 Why Supersymmetry?

There are two main reasons to suspect supersymmetry is the theory in which the Standard Model is contained: the solution it provides for quadratic divergences and the coupling unification it yields. Assorted other heuristic reasons exist, including some from string theory, but are more speculative and out of the scope of this discussion.

Quadratic Divergences

As mentioned above, it is surprising that the Standard Model particle masses are so low compared to $M_{\rm GUT}$ and $M_{\rm Pl}$. The crux of this problem is the Higgs boson, the only fundamental scalar.

To see why the Higgs is a problem, first consider particle masses in a simple theory, quantum electrodynamics (QED) with only electrons, positrons, and photons. The electron mass acquires corrections through diagrams of the form shown in Figure 1.15. A propagating free electron should include these extra contributions to its mass. The piece of the diagram corresponding to the loop and the associated vertices is (ignoring all niceties of γ matrices)

$$\begin{aligned}\pi_e(p=0) &\sim \int^{\Lambda} d^4k(-ie)\frac{i(k+m)}{k^2-m^2}(-ie)\frac{-i}{k^2}\\ &\sim -e^2 m\int^{\Lambda}\frac{d^4k}{k^2(k^2-m^2)}\\ &\sim -e^2 m\log\frac{\Lambda}{m}\end{aligned}$$

where the electron is taken to be at rest for simplicity. The term containing k in the numerator vanishes because it is odd. The issue that arises is that the cutoff Λ is arbitrary. That is, there is a correction to the electron propagator that is logarithmic in the cutoff energy, and is formally infinite because we do not know if there is a cutoff energy. In renormalization, this type of diagram is absorbed into the electron mass by assuming the "bare" electron mass is also infinite; the two infinite terms cancel to give a finite physical mass. One can imagine diagrams with more loops; however, one can see intuitively that all such diagrams yield corrections of the form $\log(\Lambda/m)^N = N\log\Lambda/m$. That is, all higher diagrams can be absorbed by a single renormalization.

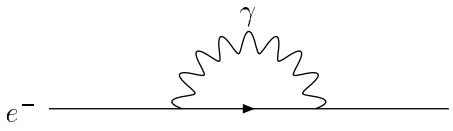

Figure 1.15: Electron self-energy correction (from [51]).

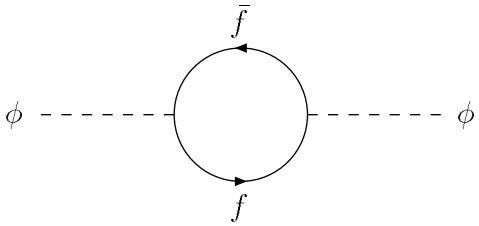

Figure 1.16: Higgs self-energy correction (from [51]).

Now, consider renormalization of the Higgs mass. The Higgs boson is subject to corrections of the form shown in Figure 1.16. The contribution is

$$\begin{aligned} \pi_\phi^2(p=0) &\sim \int^\Lambda d^4k i\lambda_e \frac{i(k+m)}{k^2-m^2} i\lambda_e \frac{i(k+m)}{k^2-m^2} \\ &\sim -\lambda_e^2 \int^\Lambda d^4k \frac{k^2+m^2}{(k^2-m^2)^2} \\ &\sim -\lambda_e^2 \left[\frac{\Lambda^2}{m^2} + 2m^2 \log\frac{\Lambda}{m}\right] \end{aligned} \tag{1.112}$$

where unimportant constant prefactors have not been carefully followed. The term with k in the numerator vanishes again. This time, there is a term that is quadratically dependent on Λ. This is problematic; even if this term is renormalized away, the next order diagram contributes $(\Lambda/m)^4$, which requires a second renormalization, and so on. A single renormalization does not suffice. This is unsatisfactory. If the existence of a true, physical cutoff is assumed, such that Λ is finite, then the Higgs mass stays finite. However, in the Standard Model, the only plausible cutoff scales are $M_{\rm GUT}$ or $M_{\rm Pl}$; corrections of these sizes would drive the Higgs mass up to the corresponding scale! This problem goes under the rubric "quadratic divergences" for obvious reasons. Note that it is the scalar nature of the Higgs that is the source of the problem; as noted above, fermions do not suffer from quadratic divergences. It can be shown that gauge bosons are also safe.

Supersymmetry (SUSY) aims to solve the above problem by introducing, for each Standard Model particle, a partner of the opposite fermionic or bosonic nature. Scalars and spin-1 particles acquire spin-1/2 partners, spin-1/2 particles acquire spin-1 partners. Including gravity, spin-2 particles acquire spin-3/2 partners. In unbroken supersymmetry, the couplings of a given particle and its "superpartner" to other particles are the same. That is, the Lagrangian is forced to be invariant under transformations that exchange particles and superparticles. The result is that

wherever a fermionic loop exists, a bosonic loop appears, and vice versa. Fermionic and bosonic loops have opposite signs, however, so the loops cancel in the case that supersymmetry is satisfied exactly. In the real world, no superpartners have been seen yet, so supersymmetry must be broken at some scale $M_{\rm SUSY}$ such that superpartner masses are fixed to be of order $M_{\rm SUSY}$; thus, the cancellation is imperfect and leaves quadratic corrections of the form discussed above, but with $\Lambda \sim M_{\rm SUSY}$ rather than $\Lambda \sim M_{\rm GUT}$ or $\Lambda \sim M_{\rm Pl}$. In order for SUSY to keep these corrections small, $M_{\rm SUSY}$ must be reasonably low, of order a few TeV.

Coupling Unification

It is strongly believed that there is a scale $M_{\rm GUT}$, where the strong and electroweak forces unify. To indicate what is meant by "unify," I digress momentarily on the subject of running couplings.

Any theory of particle physics has coupling constants that are applied at vertices in diagrams. These couplings indicate the strength of the interaction — they are coefficients of various terms in the Lagrangian. For reasons I do not discuss here, these couplings must be dimensionless (when appropriate units are used) to give a renormalizable theory. The argument is similar to that made regarding logarithmic vs. quadratic divergences. When the couplings are made dimensionless, it is possible to compare the strengths of different interactions in a sensible way.

These couplings are not, however, energy independent. Just as the electron propagator contains corrections due to higher-order diagrams that appear everywhere an electron propagator appears, so do interaction vertices. As with the electron mass, these additional diagrams can be summed to give one lump contribution that is applied everywhere. However, the result is dependent, effectively, on the energy of the interaction — how much energy is transferred at the vertex. As a result, the effective coupling strength is energy dependent. This dependence is termed "running," and one often speaks of the "running of the couplings." The strong coupling strength decreases with energy; this gives rise to "asymptotic freedom," whereby at high energies the strong force binding quarks into nuclei is weakened so free quarks may exist. Unification of the strong and electroweak interactions in some GUT can only occur if the couplings meet at some energy $M_{\rm GUT}$.

It is this last point that provides strong support for supersymmetry. In GUTs without supersymmetry, the couplings almost unify at $\sim 10^{19}$ GeV, but not quite; see Figure 1.17. The improvement in the precision of the measurements of the couplings in recent years (from experimental data from LEP, LEP2, and the Tevatron) has made it clear that non-supersymmetric GUTs do not unify the couplings. The running of the couplings depends on the gauge-group structure, so a more complicated gauge structure than $SU(5)$ would likely cause even greater disagreement. Inclusion of supersymmetry in a minimal way appears to fix the problem; see Figure 1.17. Generically, in order to provide this unification, $M_{\rm SUSY}$ must be in the range 100 GeV to about 10 TeV. If it is lower, it should have been seen already; if it is higher, it does not affect the running of the couplings enough to yield unification.

1.7.2 The Supersymmetric Particle Spectrum

Introducing supersymmetry in a formal manner is well outside the scope of this discussion. However, the particle spectrum and the rationale for its structure can be explained. The expected lightest superpartner, the neutralino, and its possibility as a WIMP candidate are of particular

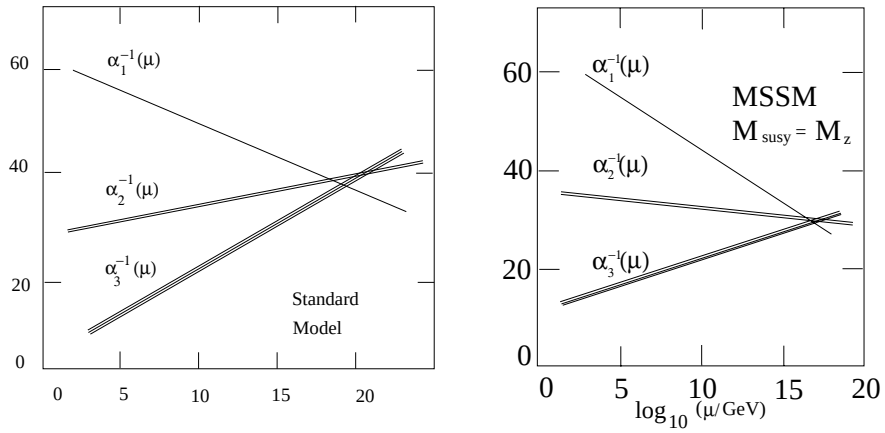

Figure 1.17: Running of couplings. α_1, α_2, and α_3 are the $U(1)$ hypercharge, $SU(2)$ electroweak, and $SU(3)$ strong couplings, respectively. Left: Standard Model with $SU(5)$ GUT. Right: Minimial supersymmetric $SU(5)$ GUT with $M_{SUSY} = M_Z$. Bands indicate uncertainties, set by precision of coupling measurements at low energy. Figures taken from [52].

interest. Only the minimal supersymmetric standard model (MSSM), the minimal extension to the Standard Model consistent with supersymmetry, is considered.

In the MSSM, partners for each of the normal matter particles are introduced with different Fermi/Bose statistics. The standard naming scheme is that scalar partners of Standard Model fermions acquire a "s" prefix (*e.g.*, selectron) while fermionic partners of Standard Model scalar and vector bosons acquire the "ino" suffix (*e.g.*, higgsino, gluino). The symbols for supersymmetric partners are given by placing a tilde over the Standard Model particle symbol (*e.g.*, $\widetilde{l_L}$ is the superpartner of the standard model lepton doublet l_L). The MSSM particle spectrum is as follows:

- Fermions and sfermions. In the Standard Model, the fermions are separated into helicity states (*e.g.*, e_L and e_R). The left-handed fermions are grouped into weak-isospin doublets l_L. For example ν_L and e_L are grouped together (the e subscript on ν is dropped for clarity), with weak isospin $1/2$ and $-1/2$, respectively. The right-handed states are observed to not interact by the weak interaction, so e_R and ν_R are weak isosinglets (under the assumption neutrinos have mass, which forces the existence of a right-handed neutrino — a recent modification of the Standard Model). The other lepton generations are grouped similarly. The quarks are also grouped this way, with the weak isodoublet q_L containing u_L and d_L with weak isospin $+1/2$ and $-1/2$, respectively. The right-handed quarks are also weak isosinglets. In the MSSM, each Standard Model particle gains a scalar partner. This introduces some confusion because, for example, there exist superpartners $\widetilde{l_L}$, $\widetilde{e_R^c}$, and $\widetilde{\nu_R^c}$. That is, in spite of being scalars, there are "left-handed" and "right-handed" sleptons (scalar leptons). The c subscript on the right-handed particles is a technical detail, indicating the need to conjugate them into left-handed particles before defining the superpartners. Similar superpartners exist for all the other fermions, collectively known as "sfermions."

- Gauge boson and gauginos. The gauge bosons acquire spin-1/2 partners, called "gauginos." The B and W_i (i = 1,2,3) bosons that become the photon A and the weak bosons $W^{\pm}$ and Z are given partners called binos $\widetilde{B}$ and winos $\widetilde{W}_i$. The gluons yield gluinos $\widetilde{g}$, one for each of the 8 different gluons. Though not discussed within the context of the MSSM, the graviton would give a gravitino, but with 3/2 spin because the graviton is spin-2. The gauginos are Majorana particles — they are their own antiparticles. This definition places a relationship between the two different helicity states in order to reduce the number of degrees of freedom to 2, as for each gauge boson. The implications of the Majorana nature of gauginos are described in detail in Chapter 2.

- Higgses and higgsinos. With apologies to Higgs, I write all higgs particles with a lowercase "h." For technical reasons, it is not possible to construct a supersymmetric model with a single higgs weak isodoublet. Therefore, two normal-particle higgs weak isodoublets are introduced. Recall that each Standard Model higgs doublet is complex and thus has 4 degrees of freedom. In the Standard model, three of the higgs degrees of freedom are "eaten" by the weak gauge bosons to give them longitudinal polarizations (and, therefore, mass). In the MSSM, there are 8 higgs degrees of freedom, of which only 3 are eaten by the weak gauge bosons. The remaining higgs particles are the $H^{\pm}$ charged scalar higgses, two neutral scalar higges h and H, and a neutral pseudoscalar higgs A. In the MSSM, two neutral and two charged spin-1/2 higgsinos are added, giving 8 superpartner degrees of freedom to match the 8 normal-particle degrees of freedom. The higgsinos are also Majorana fermions. This arises from the way the superpartners are "constructed." Intuitively, one can see this by analogy to the fermions and sfermions. The fundamental Standard Model fermions are two-component Weyl fermions having a single helicity state. This is expressed by the division of the Standard Model Dirac fermions into separate left- and right-handed components with different interactions. Each two-component fermion acquires a single complex sfermion partner. Conversely, each complex scalar Higgs acquires a single Weyl-fermion partner. Being neutral and Weyl, these particles must be their own superparticles; they are Majorana by definition. The two charged higgsinos are not Majorana because they are charged; they instead combine to form a four-component Dirac fermion.

MSSM Lagrangian

The Standard Model Lagrangian is extended simply by duplicating the Standard Model interactions with two particles replaced by their superpartners. In particular, the interactions of the superpartners are governed completely by the original Standard Model interactions. Supersymmetry is thus not as complicated as it might seem at first glance. Initially, no new parameters are introduced except for the additional higgs particles.

Such a Lagrangian would be supersymmetric, but superparticles are not observed. Therefore, SUSY-breaking terms must be added. There are a small set of new terms that can be added without disturbing the quadratic-divergence cancellation engendered by unbroken SUSY [53]. Most of the additional terms are just additional mass terms for the superpartners. All the sfermions gain mass terms of the form $\phi^{\dagger}\mathbf{M}\phi$. The masses $\mathbf{M}$ are matrices to allow intergenerational mixing. The gauginos acquire mass terms of the form $M_1\widetilde{B}\widetilde{B}/2$; note that these are not like the Standard Model fermion mass terms because these are Majorana particles. The higgs scalars are given additional

mass terms, rather than giving them to the higgsinos. It is arbitrary whether the higgsinos or the higgs scalars are given the additional terms; this choice is more convenient. The sfermion Yukawa couplings are modified also. Recall that the Standard Model contains Yukawa terms of the form, for example, $H\overline{q_L}u_R$ and $H\overline{q_L}d_R$, that couple the fermions to the higgs. Supersymmetric versions of these are automatically produced in unbroken SUSY; for example, $H\widetilde{q_L}^\dagger\widetilde{u_R^c}$ and $H\widetilde{q_L}^\dagger\widetilde{d_R^c}$. To break supersymmetry in these terms, copies are added with new, arbitrary coefficents. The coefficents are 3×3 complex matrices, as in the unbroken case, with off-diagonal terms allowed to give intergenerational mixing. Additional higgs-higgs interaction terms are also added. While this may sound complex, it should be remembered that these new terms are analogous to what already exists in the unbroken SUSY Lagrangian; SUSY-breaking is instituted by modifying the coefficients.

Finally, in minimal supersymmetry, there are no interactions that permit decay of a superparticle to two normal particles. Recall that, in constructing the unbroken SUSY Lagrangian, the supersymmetric analogues were found by replacing *two* Standard Model particles by their superpartners. The soft SUSY-breaking terms do not modify this statement. Effectively, then, there is a conserved multiplicative quantum number that is different for particles and superparticles, termed "R-parity." Two superparticles may annihilate to two normal particles and vice versa, but a single superparticle cannot decay to only normal-particle products. This conservation law renders the lightest superparticle stable and thus of cosmological interest.

The Neutralino

The supersymmetric particle spectrum is, of course, quite complex. It generically holds that the superpartners must have masses of tens of GeV or higher; otherwise, they would have already been seen. The lowest motivated value for $M_{\rm SUSY}$ is about 100 GeV, so no superparticles are expected significantly below this mass.

General considerations allow one to establish that the lightest superpartner is most likely the neutralino (or, rather, the lightest neutralino), a combination of the bino, neutral wino, and two neutral higgsinos. As stated above, assuming R-parity is conserved, as it is in the MSSM, the lightest superpartner is stable. This automatically rules out all charged superparticles — a charged, stable, heavy superparticle would have been detected as an ultraheavy charged particle in a variety of searches. The gluino is also ruled out because it interacts via the strong force. The sneutrino is, in principle, a possibility. However, because its interactions are related to those of the neutrino, its cross section on normal matter is known and is, in fact, already essentially ruled out [54]. This leaves only the assorted neutral gauginos and higgsinos that interact only via the weak interaction.

As noted above, there are four such "neutral-inos." They may mix because they have the same charges under all the gauge groups. (More accurately, they are chargeless under all the gauge groups except $SU(2)_{EW}$, under which they are all the $T_3 = 0$ component of weak isotriplets formed with their charged analogues (by analogy to their normal partners).) The mass eigenstates are written as the neutralinos, $\widetilde{\chi}_i^0$, $i = 1, 2, 3, 4$. The lightest one is $\widetilde{\chi}_1^0$. The four neutralinos are found by diagonalizing the mass matrix of the four neutral-ino particles, which is given by (in $\widetilde{B}$,

$\widetilde{W}_3$, $\widetilde{H}_1^0$, $\widetilde{H}_2^0$ space)

$$\mathbf{M} = \begin{pmatrix} M_1 & 0 & -m_Z \sin\theta_W \cos\beta & +m_Z \sin\theta_W \sin\beta \\ 0 & M_2 & +m_Z \cos\theta_W \cos\beta & -m_Z \cos\theta_W \sin\beta \\ -m_Z \sin\theta_W \cos\beta & +m_Z \cos\theta_W \cos\beta & \delta_{33} & -\mu \\ +m_Z \sin\theta_W \sin\beta & -m_Z \cos\theta_W \sin\beta & -\mu & \delta_{44} \end{pmatrix} \quad (1.113)$$

where $\tan\beta = v_2/v_1$ is the ratio of vacuum expectation values for the two Higgs bosons, $M_{1,2}$ are the gaugino masses introduced by the soft SUSY-breaking terms, and m_Z is the mass of the Z boson. θ_W is the Standard Model weak-mixing angle. The δ's are radiative corrections important when the two higgsinos are close in mass. In general, three basic configurations exist:

- Gaugino: If M_1 and M_2 are much smaller than μ, then the lightest neutralino is gaugino dominated. The ratio of M_1 to M_2 is fixed by grand unification and the renormalization-group equations that establish the running of M_1 and M_2 with energy, giving $M_1 < M_2$ at low energy. Thus, a gaugino-type neutralino is mostly bino.
- Higgsino: If $\mu \ll M_1, M_2$, then the lightest neutralino is primarily higgsino.
- Mixed: If $M_1 \sim \mu$, the neutralino is a roughly equal mixture of bino and higgsino.

In Chapter 2, I discuss in detail how the composition of the neutralino affects its interactions. For now, it suffices to note that the bino, wino, and higgsinos all have different interactions, so the composition affects both the annihilation cross section and thus the relic density as well as the cross section for neutralino-quark scattering.

In conclusion, the lightest MSSM neutralino is a good WIMP candidate: it interacts only by the weak force, is massive (tens of GeV or more), and stable. Because of the excellent motivation for SUSY with $M_{\rm SUSY} \sim 1$ TeV, and because it is a fairly well developed and predictive framework, the MSSM neutralino is taken to be the straw-man WIMP from the point of view of detection.

Chapter 2

Detection of and Experimental Constraints on WIMPs

2.1 Introduction

The indirect evidence for the existence of nonbaryonic cold dark matter is significant and varied. However, it is critical to demonstrate directly the existence of a specific species of nonbaryonic cold dark matter by observing it in some direct way and measuring its properties. In this chapter, I discuss various techniques used to search for WIMP dark matter, which fall under the categories of indirect and direct detection. In particular, I discuss the expected signal characteristics and concentrate on how these aspects determine the sensitivity of various detection methods. I focus on WIMPs in the context of the MSSM, though parts of the discussion are more generally applicable.

2.2 Indirect Detection of WIMPs via WIMP Annihilation

Detection of WIMPs by observation of the products of WIMP-pair annihilation is referred to as "indirect" detection. The annihilation rate scales as n_δ^2 because the process requires two WIMPs. Thus, it most useful to search for emission of WIMP-annihilation products from regions in which the WIMP density is enhanced by gravitational capture. Typical source candidates are the center of our galaxy, the Sun, and the Earth.

WIMP annihilation can produce a variety of particles. Given that WIMPs are very cold today, the available energy is the WIMP rest mass, thus restricting present-day annihilation products to particles lighter than M_δ. This constraint rules out other superparticles as annihilation products because the WIMP is expected to be the lightest superpartner. For a WIMP with a mass of tens of GeV c^{-2} or higher, many normal-particle annihilation products remain. Electroweak gauge bosons and the neutral and charged Higgses are allowed; however, because they are unstable, they are only detectable via their decay products. Even if the product particles are stable, they must be able to reach us and must be distinguishable from background sources of such particles. These constraints leave four candidates for indirect detection. Photons and neutrinos are especially useful because they are neutral and thus point back to their astrophysical sources. Cosmic-ray

antiprotons and positrons are also interesting because the background fluxes are small enough that it may be possible to observe excesses due to WIMP annihilation.

A note should be made on model predictions. Indirect detection offers the apparent advantage over direct detection that the annihilation cross sections used to calculate relic abundance are more directly related to the cross sections yielding detectable annihilation products. There is less intrinsic model dependence in calculating annihilation rates to detectable particles than in calculating WIMP-nucleon scattering rates, so the range of event rates predicted for indirect-detection experiments tends to be more constrained than for direct-detection experiments. However, two points must be remembered. First, indirect-detection methods use specific annihilation channels and it is not necessary that these annihilation channels be "typical;" the annihilation rate to neutrinos, for example, might be particularly suppressed because of some property of the WIMP. Such an effect would have little effect on the total annihilation cross section in the early universe and thus on the WIMP relic abundance, but would greatly affect detectability. Second, especially for WIMP annihilation in the Earth, WIMPs must lose energy by elastic scattering with nuclei in order to be gravitationally captured. Thus, while the annihilation cross sections are less model-dependent than elastic-scattering cross sections, the annihilation *rate* in the Earth does depend on the elastic-scattering cross section with normal matter. Thus, there can be a strong correlation between expected indirect- and direct-detection rates.

2.2.1 Gamma Rays

Photons are produced by WIMP annihilation via $\delta\delta \rightarrow \gamma\gamma$ and $\delta\delta \rightarrow \gamma Z$. Because WIMPs are nonrelativistic today, the gamma rays are monoenergetic, with $E_\gamma = M_\delta$ or $E_\gamma = M_\delta(1 - M_Z^2/4M_\delta^2)$, respectively. Thus, one expects a line source of gamma rays at energies of tens to hundreds of GeV. The only background at such high energies is the continuum flux of high-energy (HE) and very-high-energy (VHE) gamma rays emitted by point sources and the galactic plane. The galactic halo is an interesting candidate source because the continuum is suppressed relative to the plane, but the galactic center is also plausible because the WIMP-density enhancement may outweigh the disadvantage of the higher continuum background. A difficulty with this technique is that, because there is no $\delta\gamma$ coupling at tree level, the above processes occur through loop diagrams and are accordingly suppressed.

High WIMP masses, tens of GeV c^{-2} and higher, are probed today by ground-based gamma-ray observatories. These experiments detect a VHE gamma ray by the Cerenkov light emitted by the air shower the gamma ray produces when it interacts at the top of the atmosphere. Examples in the 50–500-GeV range are the CELESTE and STACEE projects, which use the large mirror areas available at solar-power plants. At higher energies, above 200 GeV, a dedicated single mirror or small array of mirrors is used by experiments such as Whipple, CAT, CANGAROO, and HEGRA. Both types of experiments have the important ability to distinguish gamma-ray-induced from hadron-induced showers by the Cerenkov light profile of the showers, allowing rejection of the enormous hadronic cosmic-ray background. However, such observatories are not ideal for a WIMP search. First, their energy resolution is not very good, of order 30%. Second, they have small fields-of-view, typically fractions of a square degree; to observe gamma rays from the galactic center, they must look directly at it. This is not usually done because the most interesting sources for gamma-ray astronomy at these energies are elsewhere. Finally, the galactic center is visible a much larger fraction of the time if the observatory is in the southern hemisphere; only CANGAROO

satisfies this requirement. Other observatories can look for annihilation in the galactic halo, but the rate is greatly suppressed by the n_δ^2 dependence. Another technique, employed by the MILAGRO experiment, is to detect the charged particles in the air shower at ground level using a large water-Cerenkov detector. The advantage of such a detector is that it is not pointed and so can look at many sources simultaneously. However, it has little or no hadronic cosmic-ray rejection near threshold, so the background is much higher. MILAGRO is also in the northern hemisphere. Unfortunately, none of the present detectors is yet sensitive enough to probe MSSM parameter space, and the coming generation of telescope arrays will only test a small fraction of such models [55].

At present, the Energetic Gamma-Ray Experiment Telescope (EGRET) on the Compton Gamma-Ray Observatory has the highest energy reach of any space-based experiment, observing in the energy range 20 MeV to 30 GeV. In addition to a line source, WIMP annihilation is expected to yield a continuous spectrum of gamma rays arising from the decay of pions produced in hadronization of quark-antiquark annihilation products. The EGRET detection of the galaxy's gamma-ray halo may be consistent with a very light WIMP (2-4 GeV c^{-2}) [56, 57]. However, a WIMP of this mass is ruled out in the context of the MSSM based on accelerator constraints [58]; moreover, there are far less exotic explanations of such a halo. In general, EGRET is 1 to 2 orders of magnitude too insensitive to detect even the most optimistic MSSM WIMPs considered under normal scans of MSSM parameter space [59]. The Gamma-ray Large Area Space Telescope (GLAST) will search for higher-mass WIMPs by detecting gamma rays in the energy range 20 MeV to 300 GeV. It is expected to launch in 2005. Its intended energy resolution is 10% and field-of-view is 1 to a few steradians. GLAST will begin to probe MSSM WIMP models, but will not cover much of the parameter space either [55].

2.2.2 Neutrinos

Another interesting avenue is detection of neutrinos produced by WIMP annihilation in the Sun and Earth. Though WIMPs scatter very infrequently (for example, the interaction length of a WIMP with 10^{-42} cm^2 WIMP-nucleon cross section is of order 10^{10} meters of germanium), they do scatter off nuclei in the Earth and Sun, losing energy and thus becoming gravitationally bound. The density enhancement increases the annihilation rate to possibly detectable levels. Neutrinos resulting from WIMP annihilation would be observed with a continuous, rather than line, spectrum for two reasons. First, such neutrinos are produced not just through $\delta\delta \to \nu\overline{\nu}$ but also by $\delta\delta \to f\overline{f}$, where f may decay, or hadronize and decay if f is a quark, emitting a neutrino. Second, the mean free path of neutrinos in the Sun is small, so a neutrino loses energy as it exits the Sun. This latter effect is not important for WIMP annihilation in the Earth. Neutrino rates from the Earth and Sun are typically expected to be much larger than from the galactic center because of proximity and solid angle.

Muon neutrinos are the neutrino flavor of interest for indirect detection. Muon neutrinos are detected by their byproducts, upward-going muons: a muon neutrino passing through the Earth may interact via W-boson exchange and produce a muon. For high-energy neutrinos, the secondary muon is very energetic (typically carrying about $0.5E_\nu$) and is well collimated with the neutrino direction (typical angular separation of $\sqrt{1\,\mathrm{GeV}/E_\nu}$ radians). The energetic muon may survive to the Earth's surface, in which case it can be detected as an upward-going muon pointing back to its source. A key element of indirect detection is the E_ν^2 enhancement factor: the neutrino-interaction cross section is proportional to E_ν and the range of the daughter muon is roughly proportional to

E_μ and thus to E_ν. Thus, it is not problematic, and can be advantageous, to use a high-threshold detector. The only known background at these high energies is atmospheric neutrinos produced in cosmic-ray air showers on the opposite side of the Earth as the detector; solar neutrinos have energies less than tens of MeV.

Other flavors of neutrinos are unsuitable for indirect detection. Electron neutrinos, while produced in copious quantities, would prove undetectable because the secondary electron is quickly stopped in the Earth. Tau neutrinos are produced mainly via the direct annihilation reaction $\delta\delta \rightarrow \nu_\tau\overline{\nu_\tau}$; the high (2-GeV c^{-2}) mass of the τ greatly suppresses production by τ decay or by decay of other particles.

A number of experiments built for other purposes (primarily proton decay and solar or atmospheric neutrinos: IMB, Baksan, Kamiokande, and MACRO) have demonstrated that interesting limits on the upward-going-muon flux can be set in spite of the much larger (a factor 10^6) flux of downward-going muons produced in cosmic-ray showers. Their current best limits on the flux of upward-going muons (due to annihilations in the the Earth or the Sun) are $\sim 10^4$ km^{-2} yr^{-1} above 1 GeV [60, 61, 62, 63, 64, 65]. These experiments have low thresholds (1 GeV) and small active areas ($\sim 10^3$ m^2) and so are not ideal for WIMP indirect detection. A number of new experiments (AMANDA, ANTARES, SuperKamiokande) with larger active areas ($\sim 10^4$ m^2) and higher thresholds ($\sim$ 25–50 GeV, due to sparser instrumentation) are being built or operating. Significant gain in sensitivity will only be possible with the construction of $\sim$ 1-km^2-active-area (km^3-active-volume) neutrino telescopes with $\sim$ 25-GeV threshold, in the planning stages now. Limits from the present 10^3-m^2-area, 1-GeV-threshold detectors and expected sensitivities for 1-km^2-area, 25-GeV-threshold detectors are shown in Figure 2.1.

Though the Earth and Sun have traditionally been considered the best candidate sources for indirect detection, Gondolo and Silk have recently shown that there may be an amplification of the WIMP-annihilation rate at the galactic center [66]. There may be a density "spike" at the galactic center because of the existence of a $2 \times 10^6\ M_\odot$ black hole there [66]. The observational evidence for such a black hole has become overwhelming in recent years, primarily through observations of proper motions of stars and gas in the vicinity of the galactic center [67, 68]. The resulting density spike follows a power law in radius, $\rho \propto r^{-\gamma}$ with $\gamma \geq 3/2$. The $\gamma = 3/2$ case corresponds to a halo with a flat core density; *e.g.*, a softened isothermal sphere. If the core density is allowed to become singular, as in a halo with a central cusp, the amplitude and exponent of the spike become larger. In any case, because the WIMP-annihilation rate scales as n_δ^2, there is an increase in annihilation at the edge of the black hole. For a density distribution with a softened core, favored by observations (see Chapter 1), the enhancement from the spike is small. However, for a halo with a central cusp, as is expected from simulations, the enhancement by the spike is typically five orders of magnitude or more. For this final case, the present limit on the upward-going-muon flux of $\sim 10^4$ km^{-2} yr^{-1} is on the verge of excluding some neutralino models. The expected fluxes are well above the atmospheric neutrino background for higher-threshold detectors like ANTARES. Note that, because the source is the galactic center, the upward-going-muon observatory must now be in the *northern* hemisphere. Expectations for the upward-going-muon fluxes for this technique are shown in Figure 2.2. If new limits exclude a significant fraction of models in the future, one should probably take this as more evidence that dark-matter halos have softened cores rather than interpreting this as ruling out neutralino WIMPs. I return to this point below.

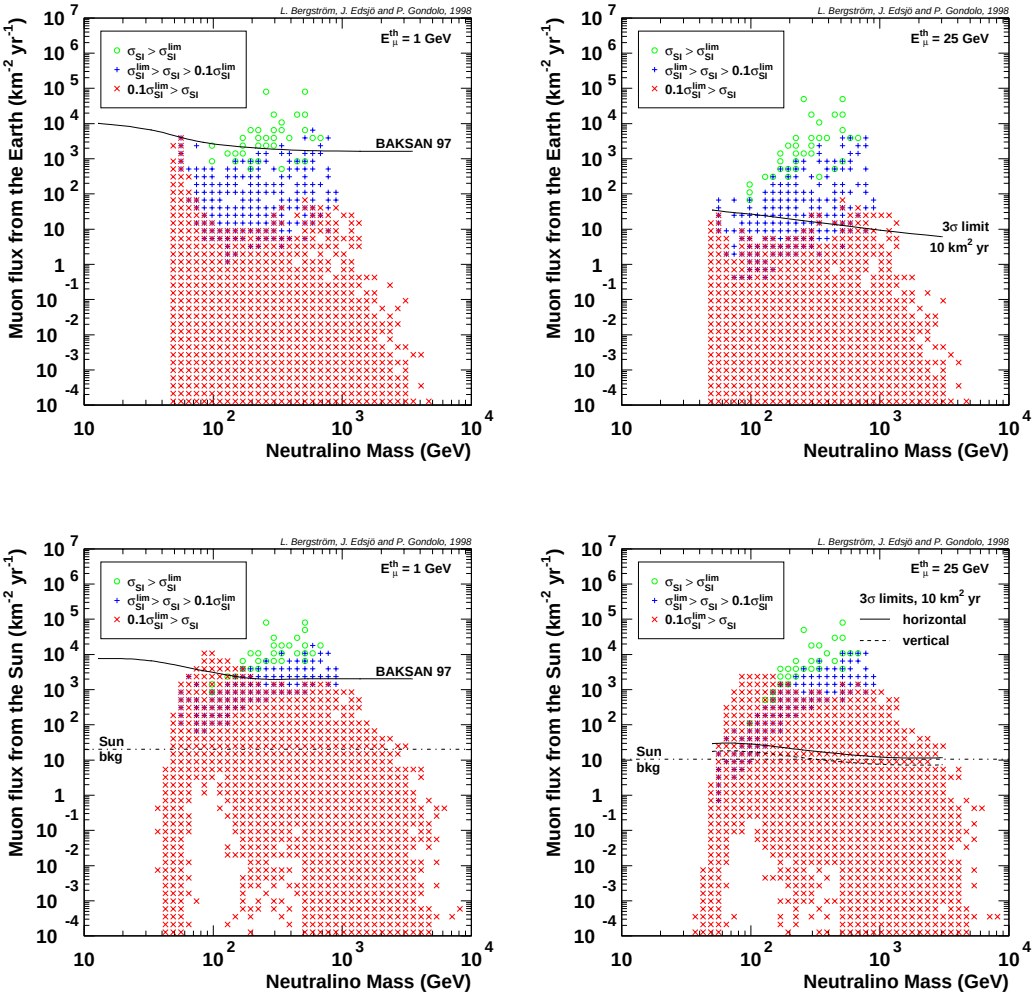

Figure 2.1: MSSM predictions and present and expected limits on the flux of upward-going-muons due to WIMP annihilation in the Earth and Sun. Upper left: flux for annihilation in the Earth for $E_\mu > 1$ GeV with limit from Baksan [61]. MACRO has a similar upper limit [65]. Upper right: flux for annihilation in the Earth for $E_\mu > 25$ GeV with expectation for 10 km^2 yr exposure for a km^3 detector. Lower left and lower right: similar plots for annihilation in the Sun. The "Sun bkg" lines indicate the background of neutrinos from cosmic-ray interactions in the Sun. Figure taken from [59].

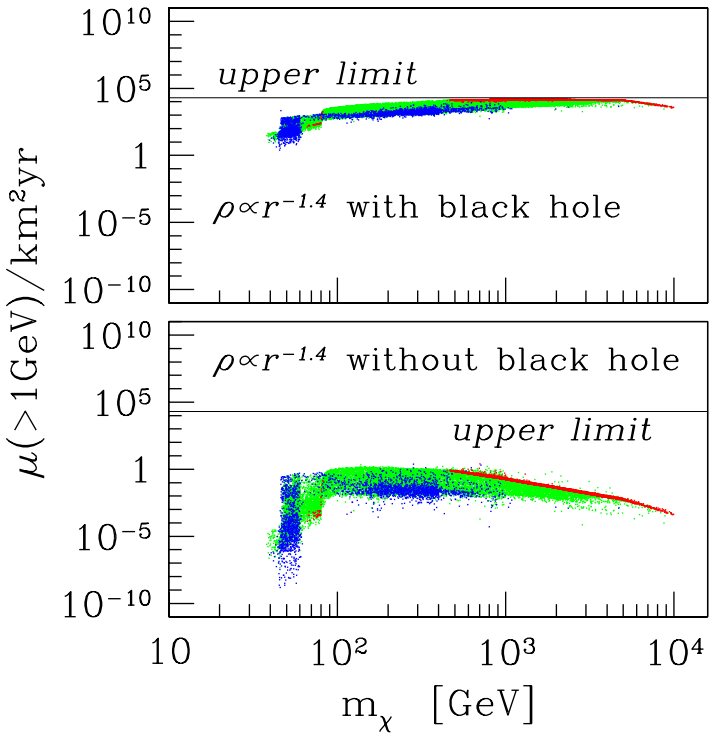

Figure 2.2: Expectations and present sensitivities to WIMP annihilation at the galactic center with and without a central black hole and assuming a halo density distribution with a cusp. Figure taken from [66].

2.2.3 Synchrotron Emission

A novel technique has recently been studied by Gondolo: detection of synchrotron emission by electrons and positrons produced by WIMP annihilation at the galactic center [69]. A magnetic field of strength $B \sim 1$ mG exists in the vicinity of the galactic center (much enhanced over the typical galactic magnetic field of $B \sim 1$ μG). This field would cause charged annihilation products to emit synchrotron radiation with a frequency spectrum rising until a few hundred GHz and remaining flat into the soft-X-ray regime. Based on the observed radio flux from the galactic center at 408 MHz, the exponent γ of the inner halo density, $\rho \propto r^{-\gamma}$, is constrained to be very small, less than about 0.01. The constraint is weakened to $\gamma < 0.1$ if the magnetic field is reduced to the strength expected from equipartition. These constraints are displayed in Figure 2.3. Thus, either neutralino WIMPs do not exist, or the halo has a flat core, in contradiction with simulations. As was discussed in Chapter 1, dynamical evidence gives no evidence for cuspy halos either, so all these observations of flat halo cores are a problem for simple WIMP-dark-matter halo models.

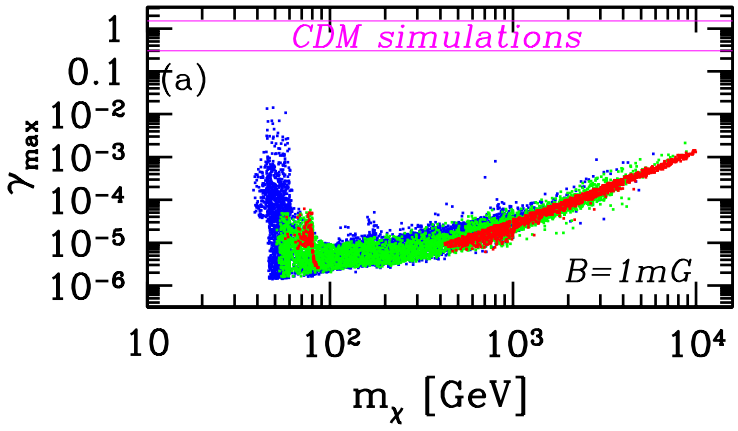

Figure 2.3: Maximum allowed cusp power-law exponent consistent with observed 408 MHz flux from galactic center for MSSM models. A 1-mG magnetic field is assumed. The power-law exponents expected from CDM simulations are indicated by the band at the top. Figure taken from [69].

2.2.4 Cosmic-Ray Antimatter

As noted above, WIMP pairs may also annihilate to quarks and charged leptons. The normal-matter component would be undetectable given the huge background of normal-matter cosmic-ray protons and electrons. However, the flux of antimatter particles may be visible because the background flux is much smaller.

Cosmic-ray antiprotons are normally produced in spallation reactions of high-energy cosmic-ray protons and nuclei with protons in the interstellar medium. The flux of *low-energy* ($\lesssim$ 1-GeV) antiprotons is greatly suppressed by kinematics and the cosmic-ray spectrum. Consider a collision between a primary cosmic-ray proton and a proton at rest in the ISM. In the center-of-mass frame, the collision is isotropic. However, the products are boosted forward when viewed from the lab frame. In order to produce a low-energy antiproton, the antiproton must have momentum almost exactly opposite that of the incoming proton in the center-of-mass frame. The incident proton must be very energetic ($E_p \gg m_p$) to have a significant probability of producing a backward antiproton. Since the primary spectrum falls steeply with energy, this happens infrequently and production of $\lesssim$ 1-GeV antiprotons is suppressed. Therefore, it is, in principle, possible to detect WIMP annihilation by the observation of a significant excess of antiprotons above expectations. However, model expectations rely on simulation of cosmic-ray production, propagation, spallation, and confinement time in the galaxy and so are not precise.

Detecting cosmic-ray antiprotons is conceptually straightforward. The type (electron/positron, proton/antiproton, or nucleus), sign of charge, and momentum or energy of the incident cosmic rays must be determined. Balloon or satellite experiments are necessary to determine these properties. The charge sign and momentum of the incident particles are usually measured using a drift chamber in a magnetic field; the curvature radius and direction of the particle's path give the desired quantities. A method of measuring particle energies and distinguishing protons from electrons and nuclei is also required. Kaons, produced by interactions with the small amount of atmosphere above the balloon, must also be rejected. BESS, a balloon-borne experiment [70], uses scintillation counters to measure the incident particles' dE/dx and times-of-flight across the apparatus. A Cerenkov radiator is used to further reject leptons, which produce larger Cerenkov signals due to their higher velocity at the same energy.

BESS's 1995 and 1997 flights determined the $\overline{p}$ flux to a precision of 20% to 50% over the region 0.2–3.5-GeV *kinetic* energy [70]. These measurements were the first to clearly demonstrate the presence of a peak at 2 GeV that arises from the details of the cross sections and kinematics. Figure 2.4 shows these data, overlaid with data from previous experiments as well as various theoretical predictions for the $\overline{p}$ flux arising from cosmic-ray spallation only. Figure 2.5 shows the BESS measurements in the context of MSSM expectations calculated by Bergstrom, Edsjo, and Ullio [59]. The measurements exclude a small number of low-Ω_δ models and are consistent with others. However, there is no reason to believe that the antiproton flux seen is due to anything other than secondary $\overline{p}$'s produced by cosmic-ray spallation: as seen in Figure 2.4, the spectrum is perfectly consistent with models without WIMP annihilation. By increasing exposure times using long-duration balloon flights or placing such an experiment on a satellite, the antiproton flux will be more precisely measured in the future. However, given the large uncertainties on the production and propagation of the cosmic-ray secondary antiprotons (Bergstrom *et al.* quote a factor of few [71]), it is unlikely that antiproton-flux measurements will further constrain WIMP models.

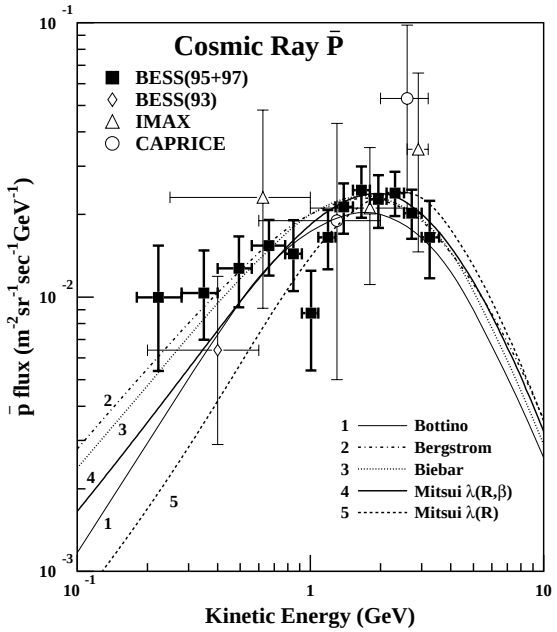

Figure 2.4: Flux of $\overline{p}$ as a function of kinetic energy. Assorted theoretical calculations for the expected flux from cosmic-ray spallation are also shown, using either diffusion or leaky-box models for cosmic-ray propagation. Figure taken from [70].

Cosmic-ray positrons are also a possible sign of WIMPs [72]. It must be noted that direct electron-positron-pair production from WIMP annihilation is suppressed because MSSM WIMPs are Majorana particles: any given annihilation matrix element is proportional to the mass of the product (this fact is explained further in the context of WIMP-quark scattering below). If the WIMP is more massive than the W and Z bosons, WIMPs may annihilate to W and Z pairs. The gauge bosons may decay to e^+e^- or $e^{\pm}\nu$, yielding an enhancement of the positron flux at $M_\delta/2$. Other decay products of the gauge bosons may themselves decay, yielding positrons with energies near $M_\delta/20$. These features may be visible as enhancements in the cosmic-ray e^+/e^- ratio.

Positrons are detected using methods similar to antiprotons, except that care must be taken to ensure high rejection of the much larger flux of hadronic cosmic rays. HEAT, another balloon experiment, uses a similar configuration to BESS [73]. It contains a magnetic spectrometer also. Instead of scintillation counters, HEAT contains a transition radiation detector (TRD) above the spectrometer and an electromagnetic calorimeter below it. The calorimeter measures the profile of the electromagnetic shower produced by the particle in the calorimeter, which is quite different for hadrons and electrons of the same momentum because of their hugely different masses. The TRD signal is dependent on the particle's Lorentz factor γ; electrons produce a signal in the TRD while protons and nuclei do not because electrons and hadrons of the same momentum have very different Lorentz factors.

HEAT data from 1994 and 1995 show a bump in the positron fraction at about 7 GeV that is reminiscent of the lower-energy feature one would expect from WIMP annihilation [72]; see

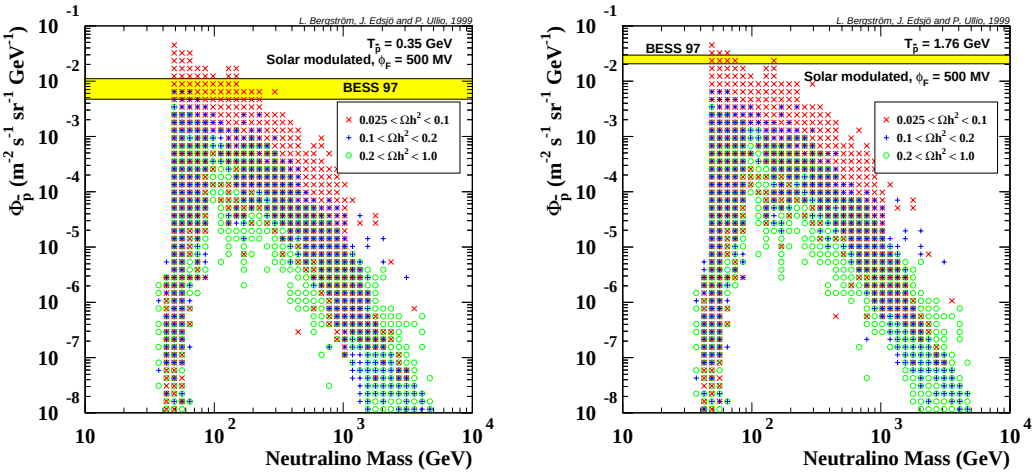

Figure 2.5: Comparison of BESS 1997 $\overline{p}$ flux measurements to expectations for MSSM models in the 0.35- and 1.76-GeV energy bins. A few (low-Ω_δ) models are excluded as giving higher $\overline{p}$ fluxes than observed. Some models are consistent with supersymmetry under the assumption that the entire flux arises from WIMP annihilation. However, more conventional explanations are consistent with the data, as seen in Figure 2.4. Figure taken from [59].

Figure 2.6. The experimenters claim that the effect cannot be produced by systematics. For example, though secondary electrons and positrons can be produced in the small amount of atmosphere above the experiment, the effect is corrected for and an error in the correction would appear as an overall normalization discrepancy rather than a feature of the kind seen. The feature can, in fact, be fit by a few different WIMP models. However, it is possible to reproduce the bump using unconventional, but not implausible, new astrophysical sources of positrons. Both classes of interpretations are shown in Figure 2.6. In addition, in the WIMP interpretation, the experiment only sees the rise of the bump and not its fall. Data in the 50-GeV range are necessary to demonstrate the fall of the bump. Furthermore, if the feature corresponds to the $M_\delta/20$ positrons, then a second peak at hundreds of GeV should also be observable.

2.3 Direct Detection of WIMPs via WIMP-Nucleon Elastic Scattering

"Direct detection" consists of observing WIMP-nucleus scattering events. The plausibility of direct detection was first demonstrated by Goodman and Witten [74] following the proposal of a neutral-current neutrino-detection method by Drukier and Stodolsky [75]. In the 15 years since, both the theoretical framework for calculating interaction rates and sufficiently sensitive experimental techniques have been developed. In this section, I describe the canonical model for expected WIMP-interaction cross sections, event rates, and deposited energy spectra. I describe the various detector technologies used for WIMP direct-detection experiments, with their various

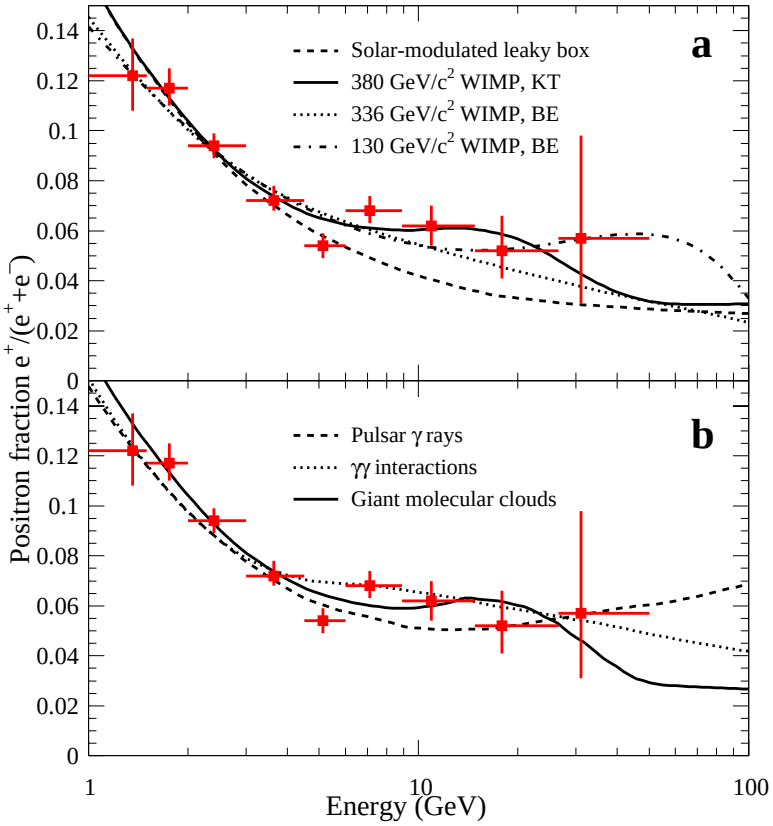

Figure 2.6: Positron spectrum measured by HEAT. Top: with best-fit expectations from WIMP models. Bottom: with possible astrophysical explanations. Figure taken from [72].

advantages and disadvantages.

There are nonnegligible astrophysical uncertainties on the parameters of the halo. I discuss expected event rates in the context of a standard, "spherical-cow," isothermal halo.

2.3.1 WIMP-Nucleus Scattering Cross Section

The fundamental process involved in WIMP-nucleus scattering is WIMP-quark scattering. This interaction, summed over the quarks present in a nucleon, gives an effective WIMP-nucleon cross section. In the low momentum-transfer limit, the contributions of individual nucleons are summed coherently over all nucleons to yield a WIMP-nucleus cross section, allowing properly for the breakdown of coherence as higher momentum transfers probe lengths smaller than the nuclear radius.

Each of the above steps has given rise to a large theoretical literature. In this section, I give an overview of this material, concentrating on the points that are particularly nontrivial or that are of particular interest to the experimentalist.

Cross-Section Estimate by Crossing Assumption

One can quickly derive a rough estimate of the interaction rate of WIMPs with normal matter based on an argument similar to that used to demonstrate that WIMP-annihilation rates in the early universe are typical of weak interactions. I follow the discussion in [1].

The rate of the scattering of WIMPs with nuclei can be calculated by Fermi's Golden Rule

$$\sigma_s v = \Gamma = \frac{2\pi}{\hbar}|\mathcal{M}_s|^2 \rho(E) \tag{2.1}$$

where $\mathcal{M}_s$ is the matrix element for scattering, σ_s is the scattering cross section, v is the WIMP velocity, and $\rho(E)$ is the density of states for the final state. Assuming WIMP annihilation and WIMP-nucleus scattering are governed by the same type of weak interaction, the annihilation and scattering matrix elements can be taken to have similar values, $\mathcal{M}_s \sim \mathcal{M}_A$. This is the "crossing assumption," referring to the fact that, generically, annihilation and elastic-scattering Feynman diagrams are related by exchange of their time and space axes. The validity of this assumption is discussed below. The density of states for s-wave (isotropic) scattering of a nonrelativistic WIMP of energy $E = M_\delta v^2/2$ is

$$\rho(E) = 4\pi(2\pi\hbar)^{-3} m_r^2 v \tag{2.2}$$

where m_r is the reduced mass of the WIMP and the nucleus. For $M_\delta \ll M_N$, $m_r \sim M_\delta$. Recalling Equation 1.109, which gave the density of states for WIMP annihilation in the early universe, and rewriting explicitly in terms of the WIMP mass,

$$\rho_A(E) = 4\pi(2\pi\hbar c)^{-3} E^2 = 4\pi(2\pi\hbar)^{-3} M_\delta^2 c \tag{2.3}$$

The phase space for the two reaction rates differ only by a factor of order v/c. Under the crossing assumption, the remainder of the Golden Rule calculation is the same as in the case of annihilation; therefore (taking $m_r \approx M_\delta$)

$$\frac{\sigma_s v}{\langle \sigma_A v_A \rangle} \sim \frac{m_r^2 v}{M_\delta^2 c} \sim \frac{v}{c} \tag{2.4}$$

or

$$\sigma_s \sim \frac{\langle \sigma_A v_A \rangle}{c} \sim \frac{1}{c}\frac{10^{-27}}{\Omega_\delta h^2} g_f^{-1/2} \frac{M_\delta}{T_f}\,\mathrm{cm^3\,s^{-1}} \sim 10^{-36}\,\mathrm{cm^2} \tag{2.5}$$

As is calculated below, this cross section corresponds to an event rate of order 1 $\mathrm{kg^{-1}\ d^{-1}}$, a sensitivity achievable with the detector technologies to be discussed later in this chapter.

The above crossing assumption links the WIMP relic abundance to the WIMP-nucleus scattering rate, rendering WIMPs plausible as a particle that can have sufficient relic abundance yet still interact weakly enough to be the as-yet-undetected nonbaryonic cold dark matter. Crossing has also been used to suggest that there are natural lower limits to the scattering cross section of cosmologically interesting WIMPs — if the scattering cross section becomes too small, the annihilation cross section also goes down and the WIMP becomes so abundant that it overcloses the universe.

However, the crossing argument makes the significant assumption of "democracy" among all annihilation processes; in particular, it is assumed that annihilation to u, d, and s quarks (those responsible for WIMP-nucleon scattering) is typical of all annihilation channels. This assumption is immediately invalidated by the fact that WIMP annihilation to quarks is proportional to m_q.

Hence, c, b, and possibly t quarks dominate annihilation to quarks, with matrix elements 1 or more orders of magnitude larger than those for u, d, and s final states. Moreover, annihilation to quarks may not be typical. For example, in pure-higgsino WIMP models, when the WIMP mass exceeds M_W, annihilation to W-boson pairs via neutral Higgses greatly increases the annihilation cross section but has no effect on the WIMP-quark scattering cross section. More generally, there are a variety of annihilation paths to Higgs, W, and Z bosons that can increase the annihilation cross section while having no effect on the quark-scattering cross section. Alternatively, it is shown below that WIMP-quark scattering, which proceeds only via Higgs, Z, and squark exchange, can be suppressed if the WIMP is a pure gaugino or higgsino. This suppression has little effect on the annihilation cross section because many other diagrams contribute. Finally, there are couplings that contribute to annihilation that vanish in the nonrelativisic limit used for WIMP-nucleon scattering; *e.g.*, $V_\mu A^\mu$ (see below). The main lesson is that, while the crossing argument provides motivation that direct detection of WIMPs is plausible, one should not rely on crossing arguments to establish hard upper or lower theoretical limits on the scattering cross section; one has to look at the specific particle-physics models.

General Considerations on Allowed Interactions

Some general points can be made. First, the WIMP-quark interaction may be written as a generic four-fermion interaction. Because a weak interaction is assumed, a massive particle (with mass M) must be exchanged. Because WIMPs are cold, the momentum transfer involved is small. Therefore, the propagator factor $(q^2 + M^2)^{-1}$ asymptotes to M^{-2}, leaving only a constant multiplying four fermion fields in the interaction Lagrangian. This low-energy limit for weak interaction was first introduced by Fermi and gives G_F for beta-decay processes involving the W boson.

Four-fermion interactions may be written as current-current interactions of the two particles' currents. The Lagrangian must be a Lorentz scalar, so this limits the possible currents and combinations. In the nonrelativistic limit considered here, there are only a finite set of combinations of fermion fields ψ and γ matrices to form particle currents:

$$\begin{array}{rcl} \overline{\psi}\psi & S & \text{scalar} \\ \overline{\psi}\gamma^5\psi & P & \text{pseudoscalar} \\ \overline{\psi}\gamma^\mu\psi & V & \text{vector} \\ \overline{\psi}\gamma^\mu\gamma^5\psi & A & \text{axial vector} \\ \overline{\psi}\sigma^{\mu\nu}\psi & T & \text{tensor} \end{array} \tag{2.6}$$

with

$$\sigma^{\mu\nu} = \frac{i}{2}[\gamma^\mu, \gamma^\nu] \tag{2.7}$$

These combinations were also introduced by Fermi. The above terms simplify further in the nonrelativistic limit. This is easiest to see through the use of helicity-projection operators, $P_L = \frac{1}{2}(1-\gamma^5)$ and $P_R = \frac{1}{2}(1+\gamma^5)$. A pure right-handed state has its spin vector along its momentum, while a pure left-handed state has it spin vector opposite to its momentum. Writing in terms of helicity projections gives

$$S \quad : \quad \overline{\psi}\psi = \overline{\psi}(P_R + P_L)\psi \tag{2.8}$$

$$\begin{aligned} P &: \overline{\psi}\gamma^5\psi = \overline{\psi}(P_R - P_L)\psi \\ V &: \overline{\psi}\gamma^\mu\psi = \overline{\psi}\gamma^\mu(P_R + P_L)\psi \\ A &: \overline{\psi}\gamma^\mu\gamma^5\psi = \overline{\psi}\gamma^\mu(P_R - P_L)\psi \\ T &: \overline{\psi}\sigma^{\mu\nu}\psi = \overline{\psi}\sigma^{\mu\nu}(P_R + P_L)\psi \end{aligned}$$

When a particle is nonrelativistic, its momentum is very small compared to its mass. Therefore, the mass term coupling left- and right-handed projections dominates, mixing them: the amounts of left- and right-handed components are about the same. Thus, the S term is much bigger than the P term. The V and A terms involve multiplication by γ^μ, which is basically the Lorentz vector $(1, \vec{\sigma})$, where $\vec{\sigma}$ are the three 2×2 Pauli spin matrices. In the nonrelativistic limit, the time component of V is large but the space components vanish because the left- and right-handed components give opposite signs when σ operates. Conversely, the time component of A vanishes while the space components are large and proportional to the particle spin. For the T term, the tensor operator $\sigma^{\mu\nu}$ reduces to

$$\sigma^{\mu\nu} = \begin{cases} 0 & \mu = 0 \text{ or } \nu = 0 \\ -\epsilon^{ijk}\sigma_k & \text{otherwise} \end{cases} \tag{2.9}$$

Thus, the T term gives no time-time or space-time components and gives space-space components proportional to the spin vector. The left and right projections cancel as with the space components of V, making the T current small compared to the S, V, and A currents.

Considering only the S, V, and A currents left in the nonrelativistic limit, the only Lorentz scalar combinations are SS, $V_\mu V^\mu$, $A_\mu A^\mu$, $V_\mu A^\mu$ and $A_\mu V^\mu$. The VA and AV terms vanish because V's space components vanish and A's time component vanishes, leaving only SS, VV, and AA.

A final point is that the vector current vanishes for Majorana fermions, which the LSP is in the context of the MSSM. A Majorana fermion is defined by $\psi^C = \psi$, with $\psi^C = -i\gamma^2\gamma^0\overline{\psi}^T$. This can be made more obvious using the notation

$$\psi = \begin{pmatrix} \phi_L \\ \phi_R \end{pmatrix} \tag{2.10}$$

where ϕ_L and ϕ_R are two-component spinors. This notation yields the expected relations

$$P_L\psi = \begin{pmatrix} \phi_L \\ 0 \end{pmatrix} \quad \text{and} \quad P_R\psi = \begin{pmatrix} 0 \\ \phi_R \end{pmatrix} \tag{2.11}$$

With these definitions, a Majorana fermion satisfies

$$\phi_R = i\sigma_2\phi_L^* \tag{2.12}$$

where σ_2 is a 2×2 Pauli spin matrix. Now, consider the time component of the V term. It is

$$\begin{aligned} \overline{\psi}\gamma^0\psi &= (i\sigma_2\phi_L)^\dagger\phi_L + \phi_L^\dagger(i\sigma_2\phi_L) \\ &= -i\phi_L^\dagger\sigma_2\phi_L + i\phi_L^\dagger\sigma_2\phi_L \\ &= 0 \end{aligned} \tag{2.13}$$

Thus, there are only two types of terms allowed in the effective Lagrangian: SS and AA.

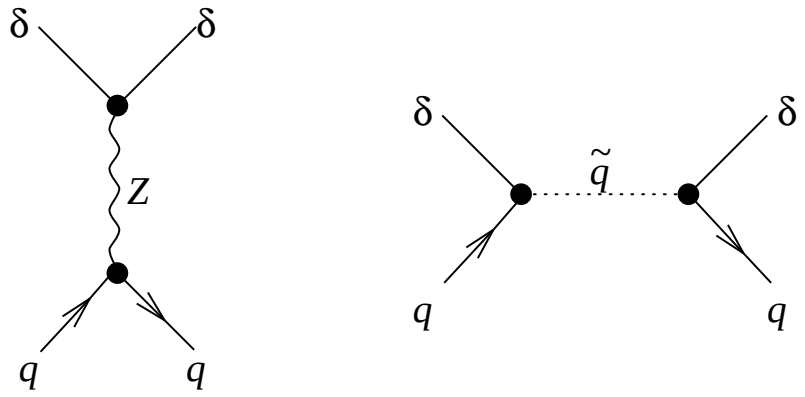

Figure 2.7: Spin-dependent WIMP-quark scattering interactions in the MSSM (from [54]).

The above simplifications leave two classes of WIMP-quark interactions in the non-relativistic limit, spin-independent and spin-dependent. The SS term gives spin-independent interactions because the WIMP and quark spins are absent. The VV term, in the case of non-Majorana WIMPs, gives a similar interaction because, as noted above, only the time component of a fermion's vector current remains in the non-relativistic limit. Both scalar and vector interactions are thus collectively grouped as spin-independent interactions. On the other hand, the AA term gives a spin-dependent interaction proportional to $\vec{s}_\delta \cdot \vec{s}_q$.

At the one-loop level, WIMP-gluon interactions also become important. These interactions are spin-independent because there is no second fermion spin with which to form a $\vec{s}_\delta \cdot \vec{s}_f$ term.

WIMP-Quark and WIMP-Gluon Scattering in MSSM Models

Specific forms for the above interactions are given by specifying a particle-physics model. Most theoretical work on WIMPs is done in the context of the MSSM, discussed in the previous chapter. I discuss the allowed interaction terms in this model, reminding the reader that MSSM neutralinos are Majorana particles, so their vector current vanishes in the nonrelativistic limit. The material presented in this section is derived directly from Jungman, Kamionkowski, and Griest (JKG) [54] and from discussions between Paolo Gondolo and Rick Gaitskell, Vuk Mandic, and myself, unless otherwise specified.

MSSM interactions are in general given by considering analogues of Standard Model interactions with two normal-particle fields replaced by their superpartners. This invariance under partner-superpartner transformations is the essence of supersymmetry. For example, the Standard Model $\overline{q}qH$ terms give rise to $\widetilde{q}^c\widetilde{q}H$ and $\overline{q}\widetilde{q}\widetilde{H}$ terms. Additional supersymmetry-breaking terms are added, but such terms involving normal particles modify only the coefficients of the above terms.

The diagrams of interest for spin-dependent interactions are shown in Figure 2.7. The first diagram, Z-boson exchange, can yield both vector and axial-vector currents for an arbitrary fermion; by the arguments made earlier, it is only the AA term that contributes, and hence Z exchange only contributes a spin-dependent interaction. The second diagram, squark exchange, yields a spin-dependent interaction when the initial- and final-state quarks have the same helicity. Explicitly, this helicity requirement results in factors $\overline{q}P_LP_Lq$ or $\overline{q}P_RP_Rq$ in the matrix element, which reduce

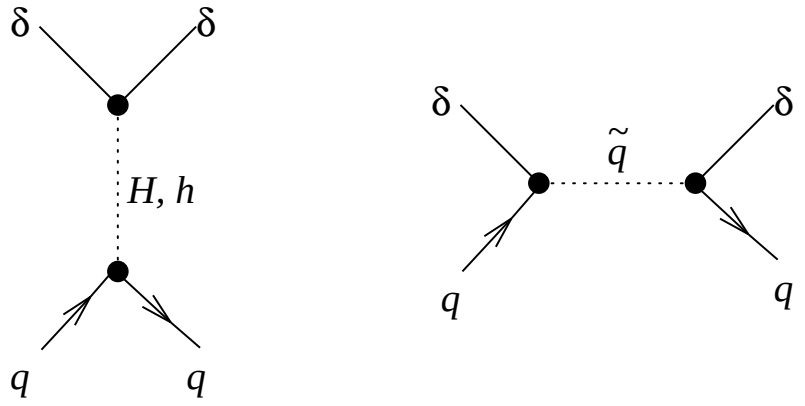

Figure 2.8: Spin-independent WIMP-quark scattering interactions in the MSSM (from [54]).

to $\overline{q}P_Lq$ and $\overline{q}P_Rq$ and give spin operators, as discussed earlier.

The primary diagrams yielding spin-independent interactions are shown in Figure 2.8. The first diagram is exchange of the two neutral Higgs scalars H and h. These particles have no spin, so obviously the interaction must be spin-independent. Exchange of the neutral Higgs pseudoscalar A would yield pseudoscalar currents, which, as seen earlier, vanish in the nonrelativistic limit. The second diagram is squark exchange. In the case that the incoming and outgoing quarks have opposite spin, the contributing factor is $\overline{q_L}m_qq_R + \overline{q_R}m_qq_L = \overline{q}m_qq$. The quark-mass coefficient comes from the $\overline{q}\widetilde{q}\widetilde{H}$ vertex: it is a supersymmetric analogue of the $\overline{q}qH$ vertex that yields the quark mass term after electroweak-symmetry breaking, and thus must have the same coefficient m_q. No spin operators appear because the vertices contain no helicity-projection operators. Note also that the direct analogues of the spin-dependent terms, $\overline{q}P_LP_Rq$ and $\overline{q}P_RP_Lq$, vanish exactly. Furthermore, even the spin-dependent interactions containing factors $\overline{q}P_Lq$ and $\overline{q}P_Rq$ may give significant spin-independent contributions if L and R squarks mix; this is discussed in more detail below.

A second set of spin-independent interactions arises from interactions with gluons through quark and squark loops. It is not particularly instructive to discuss these in detail. However, their contributions can be important because these interactions can be mediated by gluons, Higgs bosons, and quarks alone, which can be much lighter than the squarks that mediate the tree-level diagrams.

It is worthwhile to explore the tree-level diagrams' dependence on the neutralino composition in more detail. As discussed in the previous chapter, the MSSM neutralino is a mixture, or superposition, of gauginos and higgsinos. The various scattering diagrams are therefore actually coherent sums over all possible initial and final states composed of gauginos and higgsinos, each weighted by the appropriate coefficients that define the neutralino in terms of the gauginos and higgsinos:

$$\delta \equiv \widetilde{\chi}_1^0 = N_{11}\widetilde{B} + N_{12}\widetilde{W}_3 + N_{13}\widetilde{H}_1^0 + N_{14}\widetilde{H}_2^0 \tag{2.14}$$

In considering the component diagrams, diagrams with $\widetilde{W}_3$ can be ignored because the gaugino component is dominated by its bino piece, as was discussed in Chapter 1. The following points can be made:

- Z exchange: $\widetilde{B}\widetilde{B}Z$, $\widetilde{H}\widetilde{B}Z$, and $\widetilde{H}_1^0\widetilde{H}_2^0Z$ do not exist because there are no Standard Model

analogues. $\widetilde{H}_i^0 \widetilde{H}_i^0 Z$ is, however, nonzero. Thus, bino and mixed-state neutralinos have smaller Z-exchange contributions than a pure-higgsino neutralino.

- Higgs exchange: $\widetilde{B}\widetilde{B}H$ and $\widetilde{H}\widetilde{H}H$ also do not appear because there are no Standard Model analogues. Only the $\widetilde{B}\widetilde{H}H$ vertex is allowed. Therefore, a pure bino or higgsino cannot interact via Higgs exchange, while a mixed bino/higgsino can.

- Squark exchange: As discussed above, diagrams where the incoming and outgoing quarks have the same helicity are allowed and are spin-dependent, while diagrams where the incoming and outgoing quarks have opposite helicities are spin-independent. The latter diagrams require exactly one $\overline{q}\widetilde{q}\widetilde{H}$ vertex to yield the helicity flip and so only occur for mixed-state neutralinos. Therefore, pure-state neutralinos nominally have strong spin-dependent interactions but weak spin-independent interactions via squark exchange, and vice versa for mixed-state neutralinos. However, as noted above, mixing of L and R squarks (*i.e.*, $\widetilde{q_L} \leftrightarrow \widetilde{q_R}$ is significant) makes even these spin-dependent diagrams give spin-independent contributions. The suppression of spin-independent interactions with pure-state neutralinos is not as absolute as it appears at face value.

It follows that every channel of spin-independent neutralino-quark interaction is suppressed if the neutralino is pure bino or pure higgsino. Also, note that, if the neutralino is pure higgsino, then the spin-dependent cross section is larger than in mixed or pure bino cases because Z exchange is allowed; however, the spin-dependent cross section is still likely to be smaller than the spin-independent cross section. The effects of the above suppressions can be seen in a plot of the expected WIMP-nucleon cross section in different MSSM models, as is discussed later in this section.

WIMP-Nucleon Scattering

Complications arise in converting the above matrix elements for interactions of fundamental particles into a cross section on the nucleon $\sigma_{\delta \mathrm{n}}$. The main problem is uncertainty about the flavor and spin content of the nucleon: in addition to the u and d valence quarks, there are significant contributions from the s quark, heavier quarks, and even gluons. These issues are discussed in detail in [54], and I summarize their discussion here.

Consider the spin-independent case first. The expression describing the piece of the matrix element representing the transition of the quark q in the nucleon from its initial to final state is the quark scalar current:

$$\langle n | m_q \overline{q} q | n \rangle = M_n f_{Tq}^{(n)} \tag{2.15}$$

where m_q and M_n are the quark and nucleon masses and q ranges over the light quarks, u, d, and s. $f_{Tq}^{(n)}$ parameterizes the model dependence of the matrix element. Note that the quark-mass factor is present in all spin-independent interactions discussed above. For Higgs exchange, this factor arises from the $\overline{q}qH$ and $\overline{q}qh$ vertices that also give the quarks mass. For squark exchange, m_q comes from the left-right coupling at one of the vertices, the superpartner analogues of $\overline{q}qH$ and $\overline{q}qh$.

The contributions of the heavy quarks (c, b, and t) to the WIMP-nucleon cross section can be related to the gluon contribution by QCD in the limit that the heavy quarks are very massive. The gluon contribution is related to the nucleon mass and the light-quark contributions by QCD theorems. Of the light quarks, the main contribution is, interestingly enough, from the s sea quarks.

This reflects the mass dependence of the spin-independent interaction. In fact, setting f_{Tu} and f_{Td} to zero in calculations has little effect. The major contributions to the spin-independent cross section can be expressed in terms of the light quark f_{Tq}'s, and, in practice, only f_{Ts} is important. Unfortunately, to quote JKG, "the determination of the pion-nucleon sigma term from the data is fraught with significant uncertainties, which lead to uncertainties in the parameters f_{Tq}." Note that this has nothing to do with uncertainty about the WIMP properties! JKG quote two different values for f_{Ts}, 0.14 and 0.46, and proceed to use 0.14 with no strong justification.

In the spin-dependent case, the situation is even worse. The coupling is no longer proportional to quark mass, so the heavy-quark contributions are negligible and the u and d quark contributions are important. The contribution of the axial-vector current of the quark q to the spin of the nucleon may be written

$$\langle n|\overline{q}\gamma_\mu\gamma^5 q|n\rangle = 2s_\mu^{(n)}\Delta q^{(n)} \tag{2.16}$$

where $s_\mu^{(n)}$ is the spin of the nucleon n and $\Delta q^{(n)}$ parameterizes the contribution of the quark q to the spin of the nucleon n. In general, the overall spin of the nucleon is written schematically as

$$\frac{1}{2} = \frac{1}{2}\Delta\Sigma + L_q + \Delta g + L_g \tag{2.17}$$

with $\Delta\Sigma = \Delta u + \Delta d + \Delta s$ giving the contributions of the light quarks, Δg giving the gluon contribution, and L_q and L_g giving the contributions of the orbital angular momenta of the quarks and gluons. It is only the Δu, Δd, and Δs terms that contribute to WIMP-nucleon spin-dependent scattering. The values inferred from polarized-lepton-nucleus inelastic-scattering data are $\Delta u^{(p)} = 0.81 \pm 0.01$, $\Delta d^{(p)} = -0.45 \pm 0.01$, $\Delta s^{(p)} = -0.11 \pm 0.01$, and $\Sigma = 0.25 \pm 0.04$ [76], though these remain the subject of some debate. The corresponding values for the neutron can be determined by a flavor isospin rotation, giving $\Delta u^{(n)} = \Delta d^{(p)}$, $\Delta d^{(n)} = \Delta u^{(p)}$, and $\Delta s^{(n)} = \Delta s^{(p)}$. At low q^2, the remainder of the nucleon spin is believed to be given by the L_q term. The above values are used to calculate the contributions of the u, d, and s quarks to the overall spin-dependent WIMP-nucleon cross section; it must be noted that each quark axial-vector current has a different coefficient based on its weak isospin (in the case of Z exchange) or the couplings and propagators for the squarks to which it couples (for squark exchange), so the total matrix element is not simply proportional to $\Delta\Sigma$.

Calculating the WIMP-Nucleus Cross Section

The final step is to calculate the WIMP-nucleus elastic-scattering cross section. The standard review reference on the subject is Lewin and Smith [77].

In the case of spin-independent interactions, the procedure is fairly straightforward. The spin-independent interaction is very simple, with an effective WIMP-nucleon interaction

$$\mathcal{L}_{\delta n}(\vec{x}) = f_n\,\overline{\delta}(\vec{x})\,\delta(\vec{x})\,\overline{n}(\vec{x})\,n(\vec{x}) \tag{2.18}$$

where δ is the neutralino and n is the nucleon. In minimal supersymmetry, the couplings of WIMPs to neutrons and protons are the same and take a common value f_n. The nucleus can be approximated as a collection of nucleons with overlapping wavefunctions, so an effective WIMP-nucleus interaction may be written:

$$\mathcal{L}_{\delta N}(\vec{x}) = f_n\,\overline{\delta}(\vec{x})\,\delta(\vec{x})\sum_{k=1}^{A}\overline{n}_k(\vec{x})\,n_k(\vec{x}) \tag{2.19}$$

Integrating this interaction-Lagrangian density over space to get the matrix element for a transition from a nucleus at rest to a nucleus with momentum $\vec{q}$ gives

$$M(\vec{q}) = f_n \sum_{k=1}^{A} \int d^3x \, |n_k(\vec{x})|^2 \, e^{i\vec{q}\cdot\vec{x}} \tag{2.20}$$

where the WIMP factor of the matrix element gives 1 and the nucleon factor imposes the requirement that the initial and final states differ by the acquired momentum $\vec{q}$. The quantity $|n_k(\vec{x})|^2$ is the probability of the kth nucleon being at $\vec{x}$. The probability of any nucleon being at $\vec{x}$ is the sum over all nucleons given above. Obviously, if all nuclei were the same size, this would increase with A. Therefore, the A dependence is factored out, leaving the probability $\rho(\vec{x})$ of a given nucleon being at $\vec{x}$, averaged over all nucleons. This yields

$$\begin{aligned} M(\vec{q}) &= f_n A \int d^3x \rho(\vec{x}) e^{i\vec{q}\cdot\vec{x}} \\ &\equiv f_n A F(\vec{q}) \end{aligned} \tag{2.21}$$

where $F(\vec{q})$ is the Fourier transform of $\rho(\vec{x})$. $F(\vec{q})$ is called the "form factor" because it expresses the dependence of the interaction on the shape of the nucleus, given by $\rho(\vec{x})$. Note that A has been scaled out, leaving $F(\vec{q}=0) = 1$.

More generally (*i.e.*, outside the MSSM), the matrix element takes the form

$$M(\vec{q}) = [Z f_p + (A - Z) f_n] \, F(\vec{q}) \tag{2.22}$$

where the WIMP-proton and WIMP-nucleon couplings, f_p and f_n, may differ.

Under the assumption that electric charge is distributed in the same way as the nucleons (*i.e.*, there is no segregation of neutrons and protons), low-energy lepton-nucleus elastic scattering can be used to determine $F(\vec{q})$ and $\rho(\vec{x})$. It is observed that low-energy scattering is isotropic, so $F(\vec{q}) = F(q)$ and $\rho(\vec{x}) = \rho(x)$. Lewin and Smith recommend the following parameterization, first suggested by Helm, for the form factor:

$$\begin{aligned} F(qr_n) &= 3 \frac{j_1(qr_n)}{qr_n} \, e^{-(qs)^2/2} \\ &= 3 \frac{\sin(qr_n) - qr_n \cos(qr_n)}{(qr_n)^3} \, e^{-(qs)^2/2} \end{aligned} \tag{2.23}$$

with

$$\begin{aligned} r_n^2 &= c^2 + \frac{7}{3}\pi^2 a^2 - 5s^2 \\ c &= 1.23 A^{1/3} - 0.60 \, \mathrm{fm} \\ a &= 0.52 \, \mathrm{fm} \\ s &= 0.9 \, \mathrm{fm} \end{aligned} \tag{2.24}$$

where q is the momentum transfer of the scattering event, r_n is the nuclear radius, a and s are nuclear skin-thickness parameters, and j_1 is a spherical Bessel function. This form may be taken at face value as a good fit to the lepton-nucleus-scattering data down to $A \approx 10$.

Calculation of the differential cross section from the above matrix element and Fermi's Golden Rule gives

$$\frac{d\sigma_{\delta N}}{dq^2} = \frac{1}{\pi v^2} |M(\vec{q})|^2 = \frac{1}{\pi v^2} f_n^2 A^2 F^2(q^2) \tag{2.25}$$

The factor $(\pi v^2)^{-1}$ arises from final-state density of states and the standard Golden Rule $2\pi/\hbar$ factor. This may be written in the form

$$\frac{d\sigma_{\delta N}}{dq^2} = \frac{\sigma_0}{4m_r^2 v^2} F^2(q) \tag{2.26}$$

where $m_r = M_N M_\delta/(M_N + M_\delta)$ is the reduced mass of the nucleus and WIMP and v is the relative WIMP-nucleus velocity. σ_0 gives the total cross section if there were no form-factor suppression:

$$\sigma_0 = \int_0^{4m_r^2 v^2} \left.\frac{d\sigma_{\delta N}}{dq^2}\right|_{q=0} dq^2 \tag{2.27}$$

where the limits of integration are fixed by kinematics. It is also useful to write this in terms of the WIMP-single-nucleon cross section $\sigma_{\delta \mathrm{n}}$ to make the A dependence explicit. It holds that

$$\frac{d\sigma_{\delta N}}{dq^2} = \frac{d\sigma_{\delta n}}{dq^2} \frac{m_{r,\delta N}^2}{m_{r,\delta n}^2} A^2 F^2(q) \tag{2.28}$$

where $m_{r,\delta n}$ and $m_{r,\delta N}$ are the WIMP-nucleon and WIMP-nucleus reduced masses. The single-nucleon differential cross section $d\sigma_{\delta n}/dq^2$ is independent of momentum transfer because the nucleon is, for the momentum transfers of interest, pointlike. The ratio of reduced masses is necessary because these are center-of-mass cross sections. The total WIMP-nucleon cross section is given by

$$\sigma_{\delta n} = \int_0^{4m_{r,\delta n}^2 v^2} \frac{d\sigma_{\delta n}}{dq^2} dq^2 = 4m_{r,\delta n}^2 v^2 \frac{d\sigma_{\delta n}}{dq^2} \tag{2.29}$$

The fact that the single-nucleon cross section is energy independent has been used. The above definitions yield

$$\frac{d\sigma_{\delta N}}{dq^2} = \frac{\sigma_{\delta n}}{4m_{r,\delta n}^2 v^2} A^2 F^2(q) \tag{2.30}$$

and

$$\sigma_0 = \sigma_{\delta n} \frac{m_{r,\delta N}^2}{m_{r,\delta n}^2} A^2 \tag{2.31}$$

This form is useful because it puts all the particle physics in $\sigma_{\delta n}$ and the kinematics and nuclear physics in the remaining factors, allowing comparison of cross sections for different targets independent of the particular particle-physics model.

The spin-dependent case is more complicated. I follow [54]. It is customary to write the differential cross section as

$$\begin{aligned} \frac{d\sigma}{dq^2} &= \frac{8}{\pi v^2} G_F^2 \Lambda^2 J(J+1) F_S^2(q) \\ &= \frac{1}{\pi v^2} \left(\frac{g^2}{4M_W^2}\right)^2 4\Lambda^2 J(J+1) F_S^2(q) \end{aligned} \tag{2.32}$$

where g is the fundamental $SU(2)_{EW}$ coupling, M_W is the W-boson mass, and v is the relative WIMP-nucleus velocity. The prefactor $1/\pi v^2$ arises from phase space and kinematics as in the case of the spin-independent cross section. The 4 arises from spin factors; in the case of Z exchange, the virtual Z may have two polarization states. Λ is a coefficient that depends both on particle physics and on the particular assumption for the distribution of the nuclear spin J among the nucleons. If the interaction is through Z exchange, the particle-physics piece of Λ is 1 while the nuclear-spin piece is nontrivial; if the interaction is via exchange of a different particle, such as a squark, Λ may also contain correction factors due to different couplings and propagators. Finally, $F_S^2(q)$ is a form factor; in the literature, $F_S^2(q)$ is sometimes replaced by $S(q)/S(0)$.

The most naive model to use for calculating Λ is the single-particle shell model. All nucleons are assumed to be paired into spin singlets except for possibly one unpaired proton and one unpaired neutron. It is assumed that the total spin $\vec{J}$ of the nucleus is given by the spin and orbital angular momenta of these unpaired nucleons. The coupling is assumed to be via Z exchange so Λ contains only spin factors. The WIMP-quark coupling is proportional to $\vec{s}_\delta \cdot \vec{s}_q$ and thus, to high accuracy, the WIMP-nucleon coupling is given by $\vec{s}_\delta \cdot \vec{S}_n$. Consider the case of a single unpaired nucleon, so that only one nucleon's $\vec{S}$ and $\vec{L}$ determine the nuclear $\vec{J}$. The nuclear Hamiltonian depends on the square of the total angular momentum, $J^2 = \left(\vec{L} + \vec{S}\right)^2$. In such a system, the commuting angular momentum operators are J^2, L^2, S^2, and J_z, as given by elementary quantum mechanics; the nuclear spin $\vec{J}$ determines the quantization axis. The projection of the nucleon-spin operator along the WIMP spin axis, $\vec{s}_\delta \cdot \vec{S}$, does not necessarily commute with the Hamiltonian. However, the projection of $\vec{S}$ along $\vec{J}$, given by $\vec{S} \cdot \vec{J}$, satisfies

$$\vec{S} \cdot \vec{J} = \frac{1}{2}\left(J^2 - L^2 + S^2\right) \tag{2.33}$$

and so commutes with the Hamiltonian; that is, the expectation value $\langle \vec{S} \cdot \vec{J} \rangle$ is nonzero, whereas the expectation values of the other projections of $\vec{S}$ vanish. It is therefore this projection to which the WIMP spin may couple. Thus, the WIMP-nucleus spin-dependent coupling satisfies

$$\begin{aligned} \Lambda^2 &= \left[\frac{\vec{S} \cdot \vec{J}}{J^2}\right]^2 \\ &= \left[\frac{J^2 - L^2 + S^2}{2J(J+1)}\right]^2 \\ &= \left[\frac{J(J+1) - L(L+1) + \frac{3}{4}}{2J(J+1)}\right]^2 \end{aligned} \tag{2.34}$$

where $S^2 = 3/4$ for a nucleon and the $J(J+1)$ factor is necessary to yield a cross section proportional to spin^2.

The above formula may be generalized to

$$\Lambda = \frac{1}{J}\left[a_p \langle S_p \rangle + a_n \langle S_n \rangle\right] \tag{2.35}$$

where a_p and a_n satisfy

$$a_p = \sum_{q=u,d,s} \frac{d_q}{g^2/4M_W^2} \Delta_q^{(p)} \qquad a_n = \sum_{q=u,d,s} \frac{d_q}{g^2/4M_W^2} \Delta_q^{(n)} \tag{2.36}$$

$\Delta_q^{(p,n)}$ are the contributions to the spin of the corresponding nucleon from the quark q (discussed early); $\Delta_q = 1$ would imply that the quark q fully accounts for the spin of the given nucleon. d_q is the coefficient of the term in the effective Lagrangian giving rise to the interaction. $\langle S_{p,n} \rangle$ are the expectation values of the proton and nucleon spins for the nuclear ground state. The single-particle shell model with Z exchange only has $d_q = 1$ for all q so that $a_p = a_n = 1$ and, for an odd-n nucleus, $\langle S_p \rangle = 0$ and $\langle S_n \rangle = J[\vec{S} \cdot \vec{J}/J^2]$.

In general, $a_{p,n}$ and $\langle S_{p,n} \rangle$ take on more complex forms. A more accurate approximation is to assume that all nucleons of the same type as the unpaired nucleon contribute while the other type of nucleon does not. Now the total $\vec{J}$ is determined by superposition of the $\vec{S}$ and $\vec{L}$ vectors of all the nucleons in the "odd group." This odd-group model determines the expectation value of the spin of the odd group, $\langle S_{odd} \rangle$, from the observed nuclear magnetic moment via the formula

$$S_{odd} = \frac{\mu_{mag} - g^l_{odd} J}{g^s_{odd} - g^l_{odd}} \tag{2.37}$$

where $g^l_p = 1$, $g^l_n = 0$, $g^s_p = 5.586$, and $g^s_n = -3.286$, for which I give no justification [54]. More generally, the expectation values $\langle S_{p,n} \rangle$ can be calculated from the full nuclear wavefunctions if they are known.

There appear to be no general rules regarding spin-dependent WIMP-nucleus cross sections except that one should use odd-p, odd-n, or odd-p/odd-n nuclei in order to maximize the unpaired nucleon spins. I neglect spin-dependent interactions from this point on because present direct-detection techniques are orders of magnitude away from being sensitive to MSSM WIMPs via spin-dependent interactions and because the natural-germanium detectors used here are composed primarily of spin-0 isotopes, with only a 7% contribution from ^{73}Ge.

Expected WIMP-Nucleon Cross Sections

A number of authors have calculated the WIMP-nucleon cross section for MSSM WIMPs. Because SUSY has a number of undetermined free parameters (115 even for minimal supersymmetry), it is necessary for authors to make some set of assumptions regarding relationships among these parameters to reduce the set to a manageable number. In Appendix A, I delineate these different "frameworks" for reducing the parameter freedom.

Free parameters remain, so authors calculate the WIMP-nucleon cross section and the WIMP mass for a range of input values of the free parameters. The scans are typically logarithmic in the free parameters, with hard limits specified by experimental constraints or by author judgment. Once a WIMP model has been calculated, constraints on the WIMP mass or relic abundance are applied. Results from accelerator searches for neutralinos place lower limits on the neutralino mass. Overclosure places an upper limit on $\Omega_\delta h^2$, with the exact choice of the limit determined by the author's belief in various cosmological observations. Cosmological relevance places a lower limit on $\Omega_\delta h^2$.

To illustrate the typical expected parameter space, Figure 2.9 shows the regions of parameter space filled in by the MSSM models computed by Gondolo *et al.* [78] and by Corsetti and Nath [79]. Gondolo *et al.* use the "General MSSM" framework, which is one of the least restrictive ones. The allowed ranges for the various parameters are very large: the upper bounds on the higgs-mixing parameter, the gaugino mass, the higgs-pseudoscalar mass, and the sfermion masses are

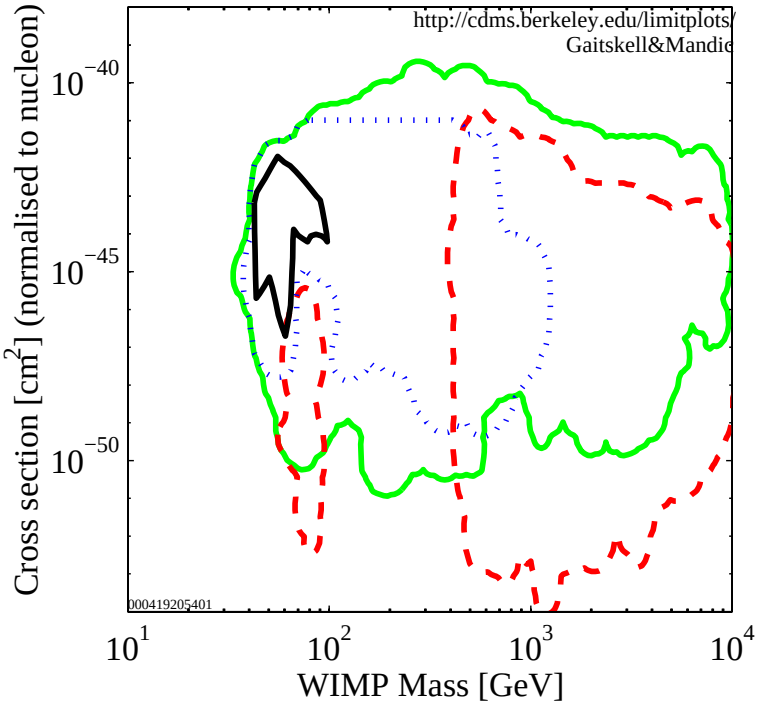

Figure 2.9: WIMP-nucleon spin-independent cross section vs. WIMP mass in MSSM models computed by Gondolo's DarkSUSY code [78] and by Corsetti and Nath [79]. Dark solid: Corsetti/Nath mSUGRA predictions. Dotted: Gondolo gaugino models. Dashed: Gondolo higgsino models. Light solid: Gondolo mixed models. Figure provided by [80].

$\mu < 300$ TeV c^{-2}, $M_2 < 300$ TeV c^{-2}, $m_A < 10$ TeV c^{-2}, and $m_s < 50$ TeV c^{-2}, respectively. The calculated models are divided into three groups for display: pure ($> 99\%$) higgsino, pure gaugino, and mixed higgsino/gaugino. In contrast, Corsetti and Nath use the "mSUGRA" framework, which is very restrictive because it applies unification relations between the different mass parameters. The most reasonable set of models is probably somewhere in between.

The structure of the expected regions is framed by the cosmological constraints imposed and the particle physics of the MSSM. First, the cosmological-relevance constraint tends to place an upper limit on the scattering cross section. As mentioned earlier, the crossing argument indicates that the annihilation and scattering cross sections scale together, so a high $\sigma_{\delta n}$ implies a high annihilation cross section and thus a low relic density. Second, the overclosure constraint eliminates high WIMP-mass models. To see this, it is necessary to rearrange the standard freeze-out equation. Recall from Equation 1.107

$$\frac{M_\delta}{T_f} - \frac{3}{2}\log\left(\frac{M_\delta}{T_f}\right) = 18.9 + \log\left[\frac{M_\delta}{\Omega_\delta h^2 g_f}\right] \tag{2.38}$$

where T_f is the freeze-out temperature, Ω_δ is the present relic density parameter in WIMPs, h is the Hubble parameter in units of 100 km s^{-1} Mpc^{-1}, and g_f is the effective number of relativistic

degrees of freedom at T_f. M_δ and T_f are in GeV. Solving for Ω_δ gives

$$\Omega_\delta h^2 = \frac{M_\delta}{g_f}\left(\frac{M_\delta}{T_f}\right)^{3/2}\exp\left(18.9 - \frac{M_\delta}{T_f}\right) \tag{2.39}$$

As discussed in Chapter 1, the quantity M_δ/T_f must remain very close to 24 in order for the relic abundance to be in the neighborhood of 1. However, the relic abundance has an additional linear dependence on the WIMP mass. Thus, as the WIMP mass increases, the relic density increases so that high WIMP masses violate the overclosure bound when all other parameters are fixed. Third, the overclosure constraint also defines the lower edge of the expected regions because it limits the annihilation cross section from below. In spite of this limit, the scattering cross section may still be suppressed by many orders of magnitude because of failure of the crossing argument. Lastly, the large missing swath of higgsino models between 100 and 400 GeV c^{-2} occurs because the threshold for annihilation to W pairs occurs at $M_\delta = M_W$: the annihilation cross section increases, pushing the relic density below the cosmological-relevance cutoff.

2.3.2 Expected Recoil-Energy Spectra and Event Rates

Once a WIMP-nucleus cross section has been calculated by the above procedure, it is straightforward to calculate the expected spectrum of recoil energies due to WIMP-nucleus scattering events. Lewin and Smith again is the standard reference [77]. I follow their derivation, albeit with some modifications that I hope clarify some points. Particularly critical for the experimentalist is understanding the dependence of the event rate on target nucleus and recoil energy. Finally, I calculate some typical numbers for event rates and energy depositions.

Calculation of Expected Recoil-Energy Spectrum

In general, the rate of interactions per unit mass of target material of a particle with interaction cross section σ is given by

$$R = \frac{N_0}{A}\,\sigma v n \tag{2.40}$$

where n is the number density of incoming particles having speed v relative to the target of atomic mass A. N_0 is Avogadro's number. The total WIMP interaction rate is given by integration over the WIMP velocity distribution. To determine the recoil-energy spectrum, it is necessary to consider the above as a function of momentum transfer.

The WIMP velocity distribution is given by the WIMP phase-space distribution. For simplicity, it is standard to assume the WIMPs occupy an isothermal, isotropic phase-space distribution, appropriate for a fully gravitationally relaxed population of WIMPs. It is known that this is not completely accurate. However, these uncertainties are small compared to the theoretical uncertainties in the WIMP-nucleon cross section; moreover, a simple standard model ensures that results presented by direct-detection experiments are easily interpreted. The full phase-space distribution function for an isothermal halo composed of particles of mass M_δ is

$$f(\vec{x},\vec{v})\,d^3x\,d^3v \propto \exp\left(-\frac{M_\delta v^2/2 + M_\delta\phi(\vec{x})}{k_B T}\right) \tag{2.41}$$

which is just a Boltzmann distribution with the gravitational potential $\phi(\vec{x})$ included. At a given point, the position-dependent factor is fixed and can be included in the normalization, leaving a simple Maxwellian velocity distribution:

$$f(\vec{v}, \vec{v}_E) = \exp\left(-\frac{(\vec{v}+\vec{v}_E)^2}{v_0^2}\right) \tag{2.42}$$

where $\vec{v}$ is the WIMP velocity relative to the Earth and $\vec{v}_E$ is the Earth's velocity relative to the nonrotating halo of the galaxy. Note that the definition of $\vec{v}$ has been changed relative to Equation 2.41. The quantity v_0^2 is characteristic of the WIMP kinetic energy, $k_B T_\delta = M_\delta v_0^2/2$ and has value $v_0 \approx 220$ km s^{-1}. It corresponds to the most probable velocity and is known as the velocity dispersion. It is related to the root-mean-square WIMP velocity v_{rms} by $v_{rms} = v_0\sqrt{3/2}$, as holds for any Maxwellian distribution. To add to the confusion, the one-dimensional root-mean-square velocity is $v_{rms}/\sqrt{3} = v_0/\sqrt{2}$; this quantity is not used here. Including all the normalization factors gives

$$dn = \frac{n_0}{k} f(\vec{v}, \vec{v_E}) d^3v \tag{2.43}$$

with

$$n_0 \equiv \int_0^{v_{esc}} dn \tag{2.44}$$

and

$$k = \int_0^{v_{esc}} v^2 dv \int d\Omega_v f(\vec{v}, \vec{v}_E) \tag{2.45}$$

The Maxwellian distribution is cut off at $|\vec{v}+\vec{v}_E| = v_{esc}$ by the halo escape velocity. Note that the cutoff is isotropic in the galactocentric WIMP velocity $\vec{v}+\vec{v}_E$, not in the Earth-centric $\vec{v}$. In practice, the escape velocity $v_{esc} \approx 650$ km s^{-1} is so large compared to v_0 that it has negligible effect. k_0 is defined to be the value of k when $v_{esc} \to \infty$; it can easily be shown that $k_0 = (\pi v_0^2)^{3/2}$.

With the above velocity distribution, the interaction rate as a function of q^2 and WIMP velocity in the target's reference frame has the form

$$\frac{d^2R}{dq^2 dV_v}(q^2, v) = \frac{N_0}{A}\left[\frac{d\sigma}{dq^2}(q^2, v)\right] v \left[\frac{n_0}{k(v_{esc})} f(\vec{v}, \vec{v}_E)\right] \tag{2.46}$$

where the cross section in the general expression Equation 2.40 has been replaced by the differential cross section and the number density by the distribution function from Equation 2.43. The quantity $k(v_{esc})$ is the normalization k from Equation 2.45 with the dependence on v_{esc} shown explicitly. The differential volume element in velocity space is given by $dV_v = d^3v$. (Note that the differential cross section is independent of reference frame: it depends only on the momentum transfer q^2, the reduced mass $m_r \equiv m_{r,\delta N}$, and the relative WIMP-nucleus velocity v, all of which are independent of reference frame.) The above must be integrated over WIMP velocity to give the full differential rate. The upper limit of integration is given by the escape velocity v_{esc}. The lower limit is set by the kinematics of the collision: in order to produce a momentum transfer q^2, the incoming WIMP must have some minimum velocity v_{min}. Conservation of energy and momentum in the target frame yield $v_{min} = q/2m_r$. Substituting from the expression for the differential cross section, Equation 2.26, yields

$$\frac{dR}{dq^2}(q^2) = \frac{N_0}{A}\frac{\sigma_0}{4m_r^2} F^2(q^2) \frac{n_0}{k(v_{esc})} \int_{q/2m_r}^{v_{esc}} \frac{f(\vec{v}, \vec{v}_E)}{v} d^3v \tag{2.47}$$

The integral is rendered nontrivial by the fact that v_{esc} is a limit on $|\vec{v}+\vec{v}_E|$, not on v alone. Consider the case $\vec{v}_E = 0$ for simplicity. The integrand becomes isotropic. Substituting for f and performing the angular integration gives

$$\left.\frac{dR}{dq^2}(q^2)\right|_{\vec{v}_E=0} = \frac{N_0}{A}\frac{\sigma_0}{4m_r^2}F^2(q^2)\frac{n_0}{k(v_{esc})}4\pi\int_{q/2m_r}^{v_{esc}} e^{-v^2/v_0^2}v dv \tag{2.48}$$

where a factor of v^2 has come from the volume element d^3v. Rewriting in terms of the recoil energy $E_R = q^2/2M_N$ and integrating, with $v_{esc}\to\infty$ for simplicity, gives

$$\left.\frac{dR}{dE_R}(E_R)\right|_{\vec{v}_E=0,v_{esc}\to\infty} = \frac{2}{\sqrt{\pi}}\frac{N_0}{A}\,n_0\,\sigma_0\,v_0\frac{e^{-E_R/E_0 r}}{E_0 r}F^2(q^2=2M_N E_R) \tag{2.49}$$

where $E_0 \equiv M_\delta v_0^2/2$, $r \equiv 4m_r^2/M_\delta M_N = 4M_\delta M_N/(M_\delta+M_N)^2$ and $k(v_{esc}\to\infty) = k_0 = (\pi v_0^2)^{3/2}$ has been used. In the case of $F^2(q^2)=1$, integrating the above over recoil energy gives the total rate

$$R_0 = \frac{2}{\sqrt{\pi}}\frac{N_0}{A}\,n_0\,\sigma_0\,v_0 \tag{2.50}$$

which allows the spectrum to be rewritten in the form

$$\left.\frac{dR}{dE_R}(E_R)\right|_{\vec{v}_E=0,v_{esc}\to\infty} = \frac{R_0}{E_0 r}e^{-E_R/E_0 r}F^2(q^2=2M_N E_R) \tag{2.51}$$

That is, for $\vec{v}_E = 0$, the recoil-energy spectrum is a simple exponential modified by the form factor. Retracing the steps, it is straightforward to see how the above form came about. The WIMP "beam" flux per unit WIMP energy has the form $E\exp(-E/E_0)$ from phase space and the Boltzmann distribution. The differential cross section goes like $v^{-2}\propto 1/E$ and is independent of $q^2\propto E_R$ when the form factor is ignored. Therefore, the product of the beam flux and differential cross section is an exponential in E. This product is integrated from the minimum possible E that can yield E_R to infinity. This lower limit is proportional to E_R, leaving an exponential in E_R with normalization factors set by the WIMP velocity (energy) distribution.

In the general case, with nonzero $\vec{v}_E$ and finite v_{esc}, the energy dependence becomes more complex. I quote the full formula, Equation 3.13 of [77]:

$$\left.\frac{dR}{dE_R}(E_R)\right|_{\vec{v}_E,v_{esc}} = \frac{k_0}{k_1}\frac{R_0}{E_0 r}\left(\frac{\sqrt{\pi}}{4}\frac{v_0}{v_E}\left[\mathrm{erf}\left(\frac{v_{min}+v_E}{v_0}\right)-\mathrm{erf}\left(\frac{v_{min}-v_E}{v_0}\right)\right]-e^{-v_{esc}^2/v_0^2}\right) \tag{2.52}$$

where $v_{min} = q/2m_r = v_0\sqrt{E_R/E_0 r}$ and $k_0 = (\pi v_0^2)^{3/2}$ are as defined before and

$$k_1 = k_0\left[\mathrm{erf}\left(\frac{v_{esc}}{v_0}\right)-\frac{2}{\sqrt{\pi}}\frac{v_{esc}}{v_0}\,e^{-v_{esc}^2/v_0^2}\right] \tag{2.53}$$

For calculation of the expected spectrum, it is necessary to specify v_E. The Earth's velocity can be broken into three pieces: the velocity of the galactic disk ($\vec{u}_r$), the velocity of the Sun with respect to the disk as measured by motion relative to nearby stars ($\vec{u}_S$), the motion of the Earth with respect to the Sun ($\vec{u}_E$):

$$\vec{v}_E = \vec{u}_r + \vec{u}_S + \vec{u}_E \tag{2.54}$$

For a disk with a flat rotation curve, it holds that $u_r = v_0$. In galactic coordinates, the disk rotation and Sun velocities are

$$\vec{u}_r = (0, 220, 0) \tag{2.55}$$
$$\vec{u}_S = (9, 12, 7) \tag{2.56}$$

The Earth orbits the Sun with a speed of approximately 30 km s^{-1}. Its orbit is very close to circular. A radial line from the center of the galaxy is almost in the plane of the orbit (only 5° off) and the plane of the orbit is tilted 60° from the plane of the galaxy. Therefore, the component of the Earth's velocity in the direction of the disk's rotation is annually modulated:

$$u_{E,y} \approx 30 \cos(60^\circ) \cos\left(2\pi \frac{t - 152.5}{365.25}\right) \text{ km s}^{-1} \tag{2.57}$$

where t is in days (00:00 January 1 corresponds to $t = 0$). The relative WIMP-Earth velocity is dominated by the y-component in this coordinate system because the orthogonal velocities add in quadrature, so it is not necessary to consider the x and z components. Thus, for substitution in Equation 2.52,

$$v_E = 232 + 15 \cos\left(2\pi \frac{t - 152.5}{365.25}\right) \text{ km s}^{-1} \tag{2.58}$$

Because of the sinusoidal variation in v_E, there is a small modulation of the recoil-energy spectrum on an annual basis, which is discussed in some detail below. The mean spectrum can be calculated by setting $v_E = \langle v_E \rangle = 232$ km s^{-1}.

Dependence of Event Rate on Target Nucleus

As can be seen from combining Equations 2.31 and 2.50, the expected total event rate per unit mass, in the case of spin-independent WIMP-nucleon scattering, naively scales as $m_{r,\delta N}^2 A^2/A \propto A^3$ for $M_N \ll M_\delta$. For $M_N \gg M_\delta$, this is modified to A because the reduced mass factor has no effect. Thus, it appears highly advantageous to use high-A materials for WIMP detection.

However, two factors counter this gain. First, Equation 2.61 indicates that, for $M_N \gg M_\delta$, the average recoil energy $E_0 r$ scales as $M_\delta^2/(100 M_N)$ where all masses are in GeV c^{-2}. That is, as the target A increases, the average deposited energy drops. This is just a kinematic effect. The momentum transfer depends only on the WIMP because it can range from zero to the full WIMP momentum. The WIMP momentum scales as M_δ because the velocity distribution is fixed. The energy transfer is $q^2/2M_N \propto M_\delta^2/M_N$, giving the above scaling. Second, the form factor suppresses the event rate at high energies. Recall that the form factor is a function of $x = qr_n$, where $r_n \sim A^{1/3}$ (Equations 2.23 and 2.24). Since $E_R = q^2/2M_N$, the form factor is a function of $r_n\sqrt{2M_N E_R}$ at fixed E_R, and therefore $x \propto A^{5/6}$. But F^2 goes as x^{-4}, so the form factor suppresses the rate at fixed E_R as $\sim A^{-3}$. These two effects leave the differential rate at $E_R = 0$ unchanged but cause the spectrum to fall more quickly as A increases.

These effects are illustrated in Figure 2.10 for a 100-GeV c^{-2} WIMP. The change in the energy dependence is especially critical because detectors have thresholds in the few-to-tens-of-keV region. A germanium target offers a clear advantage over a silicon target, but iodine yields no benefit over germanium unless the recoil-energy threshold is below 10 keV. The peak at $E_R = 0$, which scales as A to A^3, is essentially inaccessible. Typical thresholds for detectors using phonons, ionization, and scintillation are discussed later in this chapter.

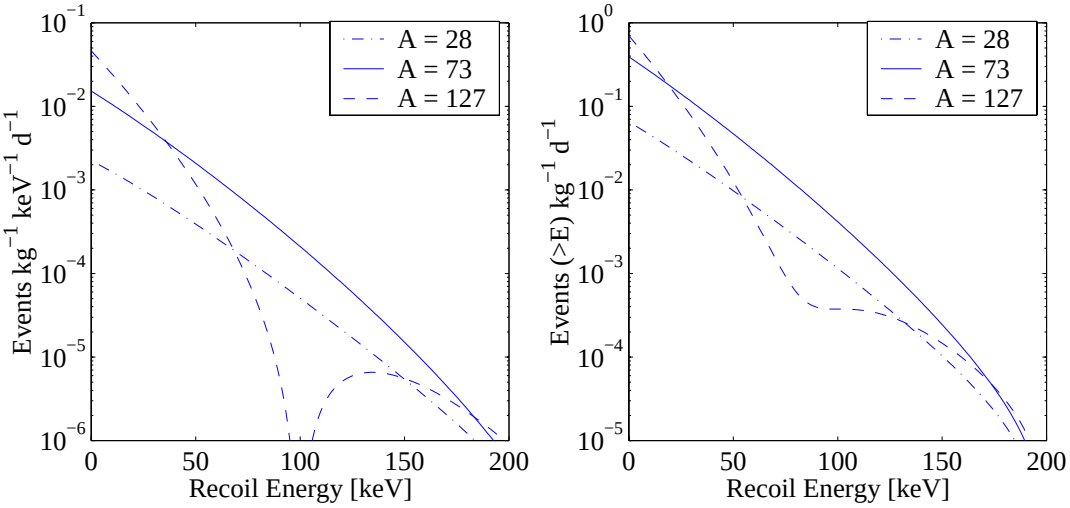

Figure 2.10: Variation of recoil-energy spectrum with A. Left: differential spectrum. Right: integral of spectrum from E_R to 200 keV ($\sim \infty$), indicating the integral rate above a given threshold. Spectra are shown for silicon ($A = 28$), germanium ($A = 73$), and iodine ($A = 127$). Note that aluminum ($A = 27$) and xenon ($A = 131$) have responses very similar to silicon and iodine, respectively. All three effects discussed in the text are included.

Numerical Results

The above formulae can be investigated numerically to get a feel for the energies and event rates involved. Lewin and Smith calculate a useful form for R_0,

$$R_0 = \frac{361}{M_\delta M_N}\left(\frac{\sigma_0}{10^{-36}\,\mathrm{cm}^2}\right)\left(\frac{\rho_\delta}{0.3\,\mathrm{GeV\,c^{-2}\,cm^{-3}}}\right)\left(\frac{v_0}{220\,\mathrm{km\,s^{-1}}}\right)\mathrm{kg^{-1}d^{-1}} \tag{2.59}$$

where M_δ and M_N are in GeV c^{-2}. Recall that R_0 is the integral rate with no form-factor correction and so is only indicative of the detectable event rate. A typical WIMP mass of interest today is $M_\delta = 100$ GeV c^{-2}. The target we use is Ge, with $A \approx 73$. Halo parameters typically used are $\rho_\delta = 0.3$ GeV c^{-2} cm^{-3}, $v_0 = 220$ km s^{-1}, and $v_{esc} = 650$ km s^{-1}. This v_0 corresponds to a rms WIMP velocity of $v_{rms} = 270$ km s^{-1}. An event rate within the reach of present experiments is 1 kg^{-1} d^{-1}; this corresponds to $\sigma_0 \approx 20 \times 10^{-36}$ cm^2, or $\sigma_{\delta n} \approx 2 \times 10^{-42}$ cm^2 (using Equation 2.31). This event rate accesses some MSSM WIMP models, as seen in Figure 2.9.

The average energy deposition is given by

$$E_0 r = \frac{1}{2}\left(\frac{M_\delta}{\mathrm{GeV\,c^{-2}}}\right)\left(\frac{v_0}{c}\right)^2 \frac{4M_\delta M_N}{(M_\delta + M_N)^2}\,\mathrm{GeV} \tag{2.60}$$

$$= \left(\frac{M_\delta}{100\,\mathrm{GeV\,c^{-2}}}\right)\left(\frac{v_0}{220\,\mathrm{km\,s^{-1}}}\right)^2 \frac{4M_\delta M_N}{(M_\delta + M_N)^2} \times 26.9\,\mathrm{keV} \tag{2.61}$$

For the standard halo parameters given above, $M_\delta = 100$ GeV c^{-2}, and a germanium target, this gives $E_0 r \approx 26$ keV. For reference, R_0 and $E_0 r$ are calculated for a number of common targets and WIMP masses of 10, 100, and 1000 GeV c^{-2} in Table 2.1.

	R_0 [kg^{-1} d^{-1}]			$E_0 r$ [keV]		
	M_δ [GeV c^{-2}] =			M_δ [GeV c^{-2}] =		
target	10	100	1000	10	100	1000
H ($A = 1$)	3.9×10^{-6}	3.9×10^{-5}	3.9×10^{-7}	0.8	1.0	1.0
Si ($A = 28$)	7.8×10^{-2}	5.5×10^{-2}	8.1×10^{-3}	2.2	17.6	26.6
Ge ($A = 73$)	3.0×10^{-1}	5.4×10^{-1}	1.3×10^{-1}	1.2	25.9	64.1
I ($A = 127$)	5.8×10^{-1}	1.7×10^{0}	6.4×10^{-1}	0.8	26.7	101.7

Table 2.1: R_0 and $E_0 r$ for hydrogen, silicon, germanium, and iodine and for WIMP masses of 10, 100, and 1000 GeV c^{-2}.

2.3.3 Non-WIMP Backgrounds and WIMP Sensitivity

Equal in importance to the expected WIMP interaction rate is the issue of how interaction rates of "background" particles factor in. In any detector, there are non-WIMP interactions due to such background particles. These are primarily due to electromagnetically interacting particles — photons, electrons, and alpha particles — from natural radioactivity of surrounding materials or the detectors themselves. At a much lower level, neutrons arising from muon interactions in the surroundings or from natural radioactivity (fission and (α, n) reactions) may also interact in detectors. The sensitivity of a WIMP search is determined by its detectors' ability to see a WIMP spectrum of the form shown in Figure 2.10 above a spectrum of energy depositions by these background particles. There are four ways to deal with this challenge.

Consider first the case of a detector that cannot distinguish WIMPs from background particles. Any such detector observes some spectrum of background-particle interactions. Such a detector can detect a WIMP by observing the expected exponential recoil-energy spectrum on top of its background spectrum while conclusively showing that such a spectrum could not arise from interactions of any other particle. Conversely, such a detector can exclude a WIMP model if the model predicts a spectrum with a higher event rate than the observed background. This is depicted naively in Figure 2.11. In such an experiment, if no counts are observed, such a detector's sensitivity improves as MT where M is detector mass and T is exposure time — the 90% CL upper limit on 0 observed counts is 2.3 events, so the event rate excluded at 90% CL in some energy bin is $2.3/MT$. However, once background events are seen, the sensitivity improvement stops. Once enough counts have been collected that the statistical error on the background spectrum is small compared to the spectrum itself, the detector's WIMP sensitivity is fixed by the background spectrum, as shown in Figure 2.11. Increasing the mass or the integration time yields no more improvement. The only way to improve WIMP sensitivity is to reduce the background spectrum.

The natural next step is background subtraction. If one has independent knowledge that predicts the background spectrum, then it is possible to subtract the background spectrum and look for a WIMP spectrum in the residuals. For example, in germanium ionization-mediated detectors, this can be done by finding photon lines in the high-energy ($>$ few-hundred-keV) spectrum, using these to set contamination levels for the associated radioisotopes, and simulating the $<$ few-hundred-keV continuum background expected from these radioisotopes and the associated decay chains. If the background model fits the observed spectrum perfectly, then the residual number of counts in an energy bin increases as $\sqrt{MT}$. The exposure increases as MT, so the residual event rate decreases

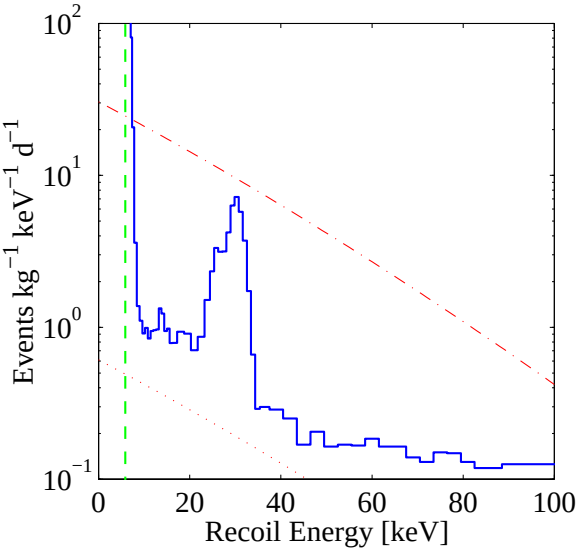

Figure 2.11: Energy spectrum observed by the UCSB/LBNL/UCB Oroville experiment, which used Ge ionization-mediated detectors [81]. The two lines indicate recoil spectra for a 100-GeV c^{-2} WIMP with WIMP-nucleon cross sections of $\sigma_{\delta \mathrm{n}} = 2 \times 10^{-39}$ cm^2 and $\sigma_{\delta \mathrm{n}} = 4 \times 10^{-41}$ cm^2; the top spectrum is excluded, the bottom spectrum allowed by these data. Note that QF ≈ 0.3 for this data, so the line at about 30 keV is the 10.4-keV gallium X-ray due to internal activation.

as $1/\sqrt{MT}$: the WIMP-nucleon cross section being probed decreases as $1/\sqrt{MT}$. A similar type of subtraction will be discussed later in this dissertation, where the CDMS neutron background is measured by the nuclear-recoil rate in Si detectors (which are far more sensitive to neutrons than WIMPs) and by the multiple-scatter nuclear-recoil rate in Ge detectors and is subtracted from the single-scatter nuclear-recoil rate in Ge detectors. In any such analysis, the accuracy required of one's background model also grows with time: the fractional systematic error on the background subtraction must decrease as $1/\sqrt{MT}$ in order for statistical errors to continue to dominate. Most experiments of this type eventually reach the point where the systematic errors cannot be decreased and the sensitivity stops improving.

Sensitivity can also be improved by outright background rejection. This takes many forms, but always is based on finding some signature that uniquely distinguishes WIMP events from background events. One form that has been pursued by many groups is nuclear-recoil discrimination. Most electromagnetic-background particles interact with electrons, producing electron recoils. Only WIMPs and neutrons interact with nuclei, producing nuclear recoils. Therefore, rejection of the electron-recoil background puts one back into the first situation above, but with a much lower background spectrum. However, no electron-recoil-rejection technique yet discovered is perfect, so there can be misidentification of electron recoils as nuclear recoils. When such leakage results in nuclear-recoil candidates, the experiment can enter the background-subtraction phase: the misidentification rate is measured using a calibration data set (usually a photon source producing electron recoils in the detector) and this misidentification rate and the electron-recoil background are used to predict the number of misidentified electron recoils. This prediction is subtracted from the

observed nuclear-recoil rate. As discussed above, this works until systematics dominate.

A second class of background-rejection techniques utilizes the unique kinematics of WIMP interactions. One such signature is diurnal variation. Because the Earth and Sun move at approximately 232 km s^{-1} through an isotropic WIMP halo (the 220 km s^{-1} disk velocity plus the additional velocity of 12 km s^{-1} relative to the disk), WIMPs primarily come from the direction into which the solar system is moving. Because of the Earth's rotation about its axis, this direction completes a circle in the lab frame once per day. A detector that is sensitive to the direction of the recoiling nucleus will see a diurnal variation in the angular distribution of nuclear recoils with a specific phase.

A similar technique is to observe an annual modulation in the recoil-energy spectrum. As discussed earlier, the Earth's velocity with respect to the halo is sinusoidally modulated due to the Earth's orbit around the Sun. The lab-frame velocity of the WIMPs thus changes with time of year, yielding a small variation in the recoil-energy spectrum. It is not necessary to measure the recoil direction, only the spectrum. The modulation amplitude is typically a few percent of the total WIMP interaction rate. Diurnal- and annual-modulation searches are particularly intriguing because the known expected phase of the signal provides an additional handle for rejecting background events and instrumental effects. On the other hand, such experiments require incredible stability because the signal is a small modulation on a large DC background signal.

The above arguments can be stated formally in terms of a given experiment's parameters. Let $\sigma_{\delta n}$ be the WIMP-nucleon cross section expected for a given WIMP model. The expected integral event rate over all recoil energies in the absence of halo parameter corrections is given by Equations 2.31 and 2.59:

$$S_{tot} = \left(\frac{2}{\sqrt{\pi}} N_0 n_0 v_0\right)\left(\sigma_{\delta n}\frac{m_{r,\delta N}}{m_{r,\delta n}} A \int_0^\infty dE_R\, F^2(q^2 = 2M_N E_R)\right) \quad (2.62)$$

$$\equiv S_n F_N \quad (2.63)$$

with units of kg^{-1} d^{-1}. S_n is the event rate on a target with $A = 1$; F_N contains the factors dependent on target material. Halo-parameter corrections have a small effect and are common to all experiments, so it is reasonable to take this expression as a starting point. Let α denote the fraction of S_{tot} that the experiment can observe; this factor contains reductions due to the choice of energy range, efficiencies of cuts that define the accepted signal, etc. Obviously, $\alpha < 1$. The number of events the experiment expects to see is

$$\langle N_{sig}\rangle = \alpha S_{tot} MT = \alpha S_0 F_n MT \quad (2.64)$$

If a portion of the background events the experiment observes can leak through the signal-acceptance cut, the expected number of leakage events is

$$\langle N_{leak}\rangle = \beta BMT \quad (2.65)$$

where B is the background rate — *i.e.*, the rate of events failing the signal-acceptance cuts — and β is the ratio of the number of background events that leak through the signal-acceptance cuts to the number of correctly identified background events. As mentioned above, β is determined from a calibration data set. It is assumed that that the uncertainty on β is negligible. If the experiment has no background-discrimination capability, then $\beta = 1$. In the following, the assumption is made that

no signal is seen and the experiment sets an upper limit on S_0. This simplifies the analysis, as there is no signal in the events both passing and failing the signal-acceptance cuts. The exclusion limits thus calculated are indicative of the experiment's sensitivity to WIMP interactions. The analysis can be generalized to the case where signal is seen; this typically contributes factors of order 1. The following analysis is incorrect for "pathological" experiments, where background subtraction is attempted with $\alpha \ll 1$ and β close to 1. The more general case is analyzed in [82].

As discussed above, there are three regimes in which an experiment can operate. If no events are seen in the signal-acceptance region, then the 90% CL upper limit on the number of events seen is 2.3. The resulting 90% CL upper limit on the experiment-independent parameter S_0 is

$$S_{0,90} = \frac{2.3}{\alpha F_n MT} \tag{2.66}$$

The limit on the event rate, and thus on the WIMP-nucleon cross section, decreases as $1/MT$. Better sensitivity is also achievable by increasing F_n and α. Annual- and diurnal- modulation experiments fall into this regime. However, for annual-modulation experiments, F_n is reduced by the ratio of the modulated to the DC signal. This ratio is usually in the range 0.01 to 0.05. Hence, large MT is required when performing an annual-modulation search. For diurnal modulation, the modulated signal may be of the same size as the DC signal, though this depends on the detector technology. This is discussed later in the context of the DRIFT TPC.

If the experiment observes events in the signal-acceptance region but believes them to be dominated by leakage of background events, then the experiment's 90% CL upper limit on S_0 is determined by the 90% CL upper limit on the number of events seen. In the absence of a WIMP signal, the expected number of events accepted is the expected number of leakage events, $\langle N_{leak} \rangle$. In the limit that N_{leak} is Gaussian distributed, the 90% CL upper limit on the number of events is given by $\langle N_{leak} \rangle + 1.28\sigma_{leak} = \langle N_{leak} \rangle + 1.28\sqrt{\langle N_{leak} \rangle}$. Thus, the expected 90% CL upper limit on S_0 is

$$\begin{aligned} S_{0,90} &= \frac{\langle N_{leak} \rangle + 1.28\sqrt{\langle N_{leak} \rangle}}{\alpha F_n MT} \\ &= \frac{\beta BMT + 1.28\sqrt{\beta BMT}}{\alpha F_n MT} \\ &= \frac{\beta B}{\alpha F_n} + \frac{1.28}{\alpha F_n}\sqrt{\frac{\beta B}{MT}} \end{aligned} \tag{2.67}$$

Once the fractional uncertainty on the observed number of leakage events, $1/\sqrt{\beta BMT}$, becomes small compared to 1, the first term dominates and the experiment's sensitivity is not improved by further exposure.

In the case of an experiment that performs subtraction of the expected background leakage into the signal region, the expected number of leakage events $\langle N_{leak} \rangle$ is subtracted from the number of events N_{acc} passing the signal-acceptance cuts. In the absence of a WIMP signal, the result is consistent with zero. Assuming the uncertainty on both on β and B vanish, the leakage prediction is perfect and the 90% CL upper limit is set by the statistical uncertainty on the subtraction because the number of leakage events fluctuates about the expected value. It is nontrivial to determine the exact value of the upper limit because one has to account for the physical boundary at $N_{sig} = 0$: one cannot have a negative number of WIMP events. The Feldman-Cousins approach is the correct

way to deal with this [83], but does not lend itself to analytical formulae. For these purposes, only an approximate result is needed, so $\langle N_{leak} \rangle = N_{acc}$ is assumed. The Bayesian probability distribution for N_{sig} is a Gaussian of width $\sqrt{\langle N_{leak} \rangle}$ truncated at $N_{sig} = 0$. The distribution is renormalized to yield an integral of 1 over the physical region $N_{sig} \geq 0$. Thus, the 90% CL upper limit is given by the 95% CL upper limit of an untruncated Gaussian, yielding

$$\begin{aligned} S_{0,90} &= \frac{1.64\sqrt{\langle N_{leak} \rangle}}{\alpha F_n MT} \\ &= \frac{1.64}{\alpha F_n}\sqrt{\frac{\beta B}{MT}} \end{aligned} \tag{2.68}$$

With background subtraction, the experiment's exclusion limit decreases as $1/\sqrt{MT}$.

There are cases of background subtraction where the uncertainty on B is significant. A typical case is an experiment that uses a multiple-scatter data set to determine B, as is the case for the data presented in this dissertation. A number of of multiple scatters N_m is observed, and a fraction f_m are expected, yielding an expected total background rate $\langle B \rangle = (N_m/f_m)/(MT)$. B is not known exactly, so the uncertainty on B enters the upper-limit calculation. In the high statistics limit, this is $\sigma_B = (\sqrt{N_m}/f_m)/(MT)$. Making the same approximation regarding the physical boundary as was made above gives

$$\begin{aligned} S_{0,90} &= \frac{1.64\sqrt{\langle N_{leak} \rangle + \beta N_m/f_m^2}}{\alpha F_n MT} \\ &= \frac{1.64\sqrt{\beta BMT + \beta BMT/f_m}}{\alpha F_n MT} \\ &= \frac{1.64}{\alpha F_n}\sqrt{\frac{\beta(1+1/f_m)B}{MT}} \end{aligned} \tag{2.69}$$

Because $f_m < 1$, it is the second term that dominates. The upper limit still decreases as $1/\sqrt{MT}$, albeit with a bigger coefficient than the previous case.

As noted before, experiments performing background subtraction eventually reach a point where systematics dominate and the above scaling fails. For example, systematic differences between calibration and background data may limit the accuracy with which a subtraction can be performed: there is some irreducible uncertainty on β. Another example is that the simulation used to calculated f_m has limited precision, either due to insufficient computation time or due to uncertainties in the scattering processes involved.

The above arguments are only approximate, primarily because of the single-bin simplification made. In most experiments, maximum-likelihood statistical analysis is used to make use of differences between the energy spectrum of the background and of the expected signal. The results should not differ significantly, and the scaling laws certainly still hold.

2.3.4 WIMP Detector Technologies

Because of the numerous ways to achieve sensitivity to WIMPs — high A, high target mass, low threshold, low background, background-rejection and -subtraction techniques — a variety

of different technologies are presently in use for WIMP searches. In this section, I discuss a few of these methods.

It is necessary to discuss the phenomenon (improperly) called "quenching." Most detectors are calibrated and provide a linear response to electron recoils. However, WIMPs produce nuclear recoils, which tend to produce less ionization, scintillation, etc. than electron recoils. The quenching factor, or QF, is the ratio of observed energy for nuclear and electron recoils of the same recoil energy. For example, in germanium ionization-mediated detectors described below, a nuclear recoil produces approximately 1/3 the ionization as an electron recoil of the same recoil energy. This effect is quite important, as it raises the threshold for WIMP recoils by 1/QF. Because detectors are usually calibrated using photon sources that yield electron recoils, energies and thresholds for such detectors are usually quoted in "keV electron-equivalent" or "keVee," which must be multiplied by 1/QF to obtain the total recoil energy (just a Jacobian transformation). An additional point is that the QF may be energy-dependent, though this effect is usually small. For example, Ge ionization-mediated detectors follow $\mathrm{QF} = 0.14 E_R^{0.19}$ [84].

Detection of ionization produced by the recoiling particle was the first technique used for direct WIMP searches [85]. A recoiling nucleus travels through the target material before stopping, ionizing atoms as it passes. In liquid or gaseous targets, this results in free electrons and ions. In a semiconducting crystal, electron-hole pairs are created, as is discussed in Chapter 4. In either case, these charges may be drifted across the target by application of an electric field and detected using a integrator circuit. The first WIMP dark-matter searches employed detectors using high-purity germanium crystals with masses of order 1 kg. These detectors were initially used to look for the spectral-line signature of neutrinoless double beta decay of an isotope in the detector itself. Because of this unique signature, such detectors were optimized for very good energy resolution. As a result, they also had very low thresholds, 5 to 10 keVee. Since QF $\sim 1/3$ for ionization in Ge, the recoil-energy threshold for WIMP interactions is approximately 15 to 30 keV. This is interesting for WIMP searches, as the WIMP spectra shown in Figure 2.10 demonstrate. The first group to use this technology was the PNL-USC collaboration [85]; the UCSB/LBNL/UCB collaboration, involving a number of the present members of CDMS, also used such detectors [81]. The final background spectrum from the latter experiment is shown in Figure 2.11. In recent years, germanium ionization-mediated detectors optimized for better resolution and thus lower threshold have been employed in dedicated WIMP searches [86]. The detector noise scales with the detector capacitance and thus its size, so such detectors are typically less massive than higher-threshold detectors. More recently, the best limits have been achieved by the Heidelberg-Moscow and IGEX experiments [84, 86]. These detectors continue to be useful as their intrinsic and external backgrounds are gradually improved. A proposal for a new experiment, GENIUS, takes this concept to its logical extreme by suspending such detectors in a large liquid nitrogen tank with only kevlar supports and electronic readout wires; nitrogen can be made extremely pure and thus provides an excellent shield against external electromagnetic backgrounds.

Scintillation-mediated detectors are also used. In this case, the recoiling particle directly or indirectly excites electrons to energy levels above their ground states, from which they may decay by emission of ultraviolet or visible photons. Photomultiplier tubes view the scintillator directly or through light guides. It is possible to purify scintillator materials to achieve low radioactive-contamination levels. Large masses (100 kg and higher) are also easily achieved. The photomultiplier tubes pose an issue in that they are made of glass, which typically has high uranium and

thorium content. However, by separating the tubes from the scintillator by light guides and using specially manufactured phototubes with "cleaner" glass, these backgrounds can be reduced. The most popular scintillators are NaI(Tl) and Xe because there is already a large literature on their properties, they have high A, and they can be easily purified. CaF_2(Eu) is also being studied because F has large spin, providing a useful target for WIMPs with spin-dependent interactions. The QF for scintillators varies; for NaI, QF ≈ 0.3 for interactions on Na and QF ≈ 0.09 for interactions on I. The latter is the important one for WIMPs with masses of tens of GeV c^{-2} and higher because of the higher cross section and better kinematic matching. For Xe, the QF for scintillation is not well known; a value of 0.2 is usually quoted (primary + secondary scintillation, see below), though this is debated by some [87]. For CaF_2, QF ≈ 0.08 for Ca and QF ≈ 0.12 for F. Thresholds are in the range of 2 to 5 keVee, corresponding to $\sim$ 20-to-50-keV recoil energy. Because of the reasonably low backgrounds and large masses achievable, scintillators are naturally used for annual-modulation searches. The DAMA collaboration has operated a NaI scintillator experiment and claim to see such a signal [88]; this signal is discussed in detail in Appendix E.

Phonon-mediated detectors have also been proposed and used because they provide very low (1-to-few-keV) energy thresholds with no quenching. However, given the difficulties of operation and readout, a phonon-mediated detector with no nuclear-recoil discrimination capability is of interest only for very low mass ($<$ 20-GeV c^{-2}) WIMPs, which have been essentially ruled out in the context of the MSSM by accelerator searches.

The detection technologies discussed above offer little or no direct background-rejection capability, with the exception of annual modulation. An important background-rejection technique possible with many technologies is nuclear-recoil discrimination. In NaI scintillators, nuclear-recoil discrimination is possible by pulse shape — nuclear recoils yield more quickly decaying scintillation pulses than electron recoils, and this effect is accentuated as the crystals are cooled to 100 K. The UKDM collaboration has been concentrating on developing pulse-shape discrimination in NaI using cooled crystals [89]. The DAMA collaboration has also demonstrated pulse-shape discrimination; however, their 1996 "pulse-shape-analysis" exclusion limit turns out to be dominated by their 2–3-keVee (22–33-keV-recoil) energy bin, in which they have essentially no discrimination capability due to low signal-to-noise.

Liquid xenon with scintillation readout also provides nuclear-recoil discrimination by pulse-shape analysis. A recoiling particle deposits energy in two ways. First is direct electronic excitation of Xe_2 molecules, resulting in scintillation. Second is ionization of Xe_2. These two excitations may be detected in different ways. One method, scintillation-only, utilizes the fact that the electrons and ionized molecules recombine to an excited state of Xe_2, which may then decay by scintillation. The timing of the two pulses differs for nuclear and electron recoils because of the different dE/dx along the recoil track. In the second method, scintillation+ionization, the ionization is drifted to a a gaseous acceleration region where a proportional scintillation pulse can be produced. The direct scintillation is also observed; the relative sizes of the two scintillation pulses can be used to perform nuclear-recoil discrimination. A collaboration consisting of the UKDM, UCLA, Torino, and ITEP groups is undertaking experiments of both types [89].

Combination of phonon readout with a second excitation mechanism, such as ionization or scintillation, offers nuclear-recoil discrimination also. The phonon signal is "unquenched," providing a measurement of the total recoil energy. The second channel is quenched, providing a signal to discriminate electron recoils from nuclear recoils. As mentioned above, the phonon-signal threshold

is very low. With CDMS detectors, which use ionization as the second channel, the resolution in the second channel can be comparable because the detectors are small (100–300 g) and thus have small capacitance. The low mass is compensated by the discrimination capability. An ionization dead layer (a thin surface layer with suppressed ionization yield) reduces the discrimination capability for low-energy electrons from external sources, but this has recently been corrected using a new ionization-electrode technology and athermal-phonon pulse-shape discrimination against surface events. CDMS BLIP detectors have demonstrated useful thresholds of as low as 4-keV *recoil energy*, below which the ionization signal is indistinguishable from noise and thus no discrimination can be performed. It may be possible to reduce this by a factor of 2 for CDMS ZIP detectors.

CRESST use phonon-mediated detectors and are pursuing the use of scintillation as the second channel in $CaWO_4$. Rather than using phototubes to observe the scintillation, a second phonon-mediated detector with a light absorber is placed adjacent to the primary detector to observe the scintillation light. The resolution on the scintillation signal is so far much worse than in the phonon channel because of poor light collection, resulting in a useful recoil-energy threshold of $\sim$ 15–20 keV. It has also not yet been demonstrated that the scintillation yields for nuclear recoils of the three constituent elements are the same, or even that nuclear recoils of all three elements yield scintillation. Also, only a 6-g prototype has been demonstrated, though 300-g detectors are under construction. However, the technique has the impressive advantage that it appears to have no scintillation dead layer.

Another technique capable of nuclear-recoil discrimination is superheated droplets. Droplets of superheated freon are suspended in a viscous gel. The energy deposition by an electron recoil is too distributed to cause a bubble to boil, but nuclear recoils can. The boiling of a bubble is detected acoustically using piezoelectric sensors. The electron-recoil rejection of such detectors has been shown to be excellent. Though no energy information is available from a single detector, the energy threshold can be tuned by the operating conditions, allowing one to use many detectors with different thresholds to reconstruct an energy spectrum. The fluorine present in freon provides a good cross section for WIMPs with spin-dependent interactions. There are a number of disadvantages, however: the lack of a high-A target, the low target density, and the lack of energy resolution in a single detector.

Along similar lines, the UKDM collaboration is developing scintillators with sensitivity to the recoiling particle's track length. The detector consists of submicron-size CaF_2 granules suspended in liquid scintillator. The scintillation pulse from the liquid scintillator is much shorter than that from CaF_2. Nuclear recoils have short track lengths and thus excite only CaF_2, while electron recoils traverse both CaF_2 and the liquid scintillator, yielding pulses with a fast component.

Finally, direction sensitivity is perhaps the holy grail of WIMP searches because it provides a way to harness the huge, unique signal provided by diurnal modulation. The distinction from annual modulation lies in the size of the modulation: recall that the annual modulation is only a few % of the DC signal; with diurnal modulation, the signal can be much larger because of the extreme directionality of the WIMP "wind." To make use of diurnal modulation, the experiment must measure the direction of the recoiling particle. The DRIFT collaboration is pursuing direction sensitivity using a Xe-gas time-projection chamber (TPC) doped with CS_2. A TPC operates by applying a drift field to a low-pressure gas. Any ionization track produced is drifted to a detection grid at one end of the chamber. The projection of the track parallel to the grid is determined by the position of the track hitting the grid. The projection perpendicular to the grid is determined

by the relative time of arrival of the track at the grid points. However, the ionization track diffuses significantly as it drifts because electrons are very light. Recently, the introduction of CS_2 doping has made it possible to maintain the track as it drifts across the chamber: CS_2 is very electronegative, so it immediately picks up any free electrons and turns into a negative ion. Diffusion of the heavy ions during drift is suppressed by their large mass. The electrons are stripped from the CS_2 ions when they reach the anode. Energy can also be determined by integrating dE/dx along the track, and obviously event position is also determined. However, there are a number of deficiencies [87]. First, as with any TPC, the target mass per unit volume is very small. Second, the threshold is very high: $\sim$ 50–100-keV recoil energy is needed in order to measure a track direction. Third, at present, it is not possible to distinguish the start and end of track, though this is in principle possible by the variation in dE/dx over the track length. Therefore, only a forward/back vs. side asymmetry can be measured, reducing the modulation effect to only about 20% [90]. Finally, reduction of threshold is not helpful because the asymmetry decreases as the threshold is decreased: low-velocity WIMPs begin to dominate. While interesting, this technology has not yet demonstrated its competitiveness.

2.3.5 Current Exclusion Limits

Figure 2.12 shows exclusion limits on the spin-independent WIMP-nucleon cross section from the leading experiments. The curves are: UKDM cooled NaI with pulse-shape discrimination; IGEX and Heidelberg-Moscow Ge ionization-mediated detectors; and DAMA 1996 NaI "pulse-shape-analysis" limit. The curves assume the halo model and parameters discussed above, with $v_0 = 220$ km s^{-1}, $\rho_\delta = 0.3$ GeV c^{-2} cm^{-3}, and $v_{esc} = 650$ km s^{-1}. Decreasing ρ_δ scales all the curves upward. Changes in v_0 modify the curves' shape, shifting toward lower mass as v_0 is increased because the average recoil energy for a WIMP of a given mass increases. The curves have a characteristic shape that reflects the experimental sensitivity, which is best when the target nucleus and WIMP mass are matched to give the largest $E_0 r$.

Also shown is the claimed detection by the DAMA experiment, which observes an annual-modulation signal in its NaI scintillation-mediated detectors. This result is discussed in detail in Chapter E. At this point, it is of interest to note that the lower edge of their allowed region is determined by the sensitivity curve of NaI; if no signal were seen and a limit were placed, its lowest section would have the same shape as the contour shown.

It is important to note that comparison of the various experiments in this way is essentially model independent, assuming only the general phenomenon of coherent scattering scaling as A^2 and the Helm form factor from Equation 2.23. As noted, the actual shape of the curves is halo-parameter dependent, but all the curves move in similar ways as halo parameters are varied.

Clearly, the various methods are all quite competitive. Furthermore, if a WIMP signal is detected, observation using different techniques will be important in substantiating the signal.

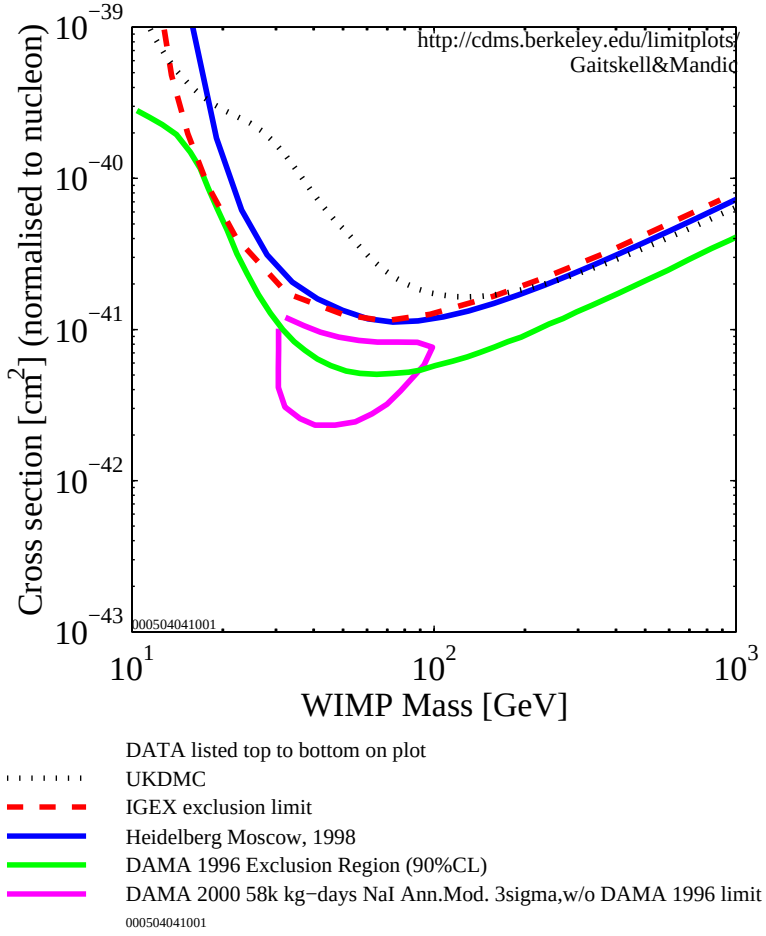

Figure 2.12: Exclusion limits on the spin-independent WIMP-nucleon elastic-scattering cross section from the currently most sensitive experiments (May, 2000). All limits are at 90% CL. Also shown is the DAMA annual-modulation 3σ CL allowed region. Figure from [80].

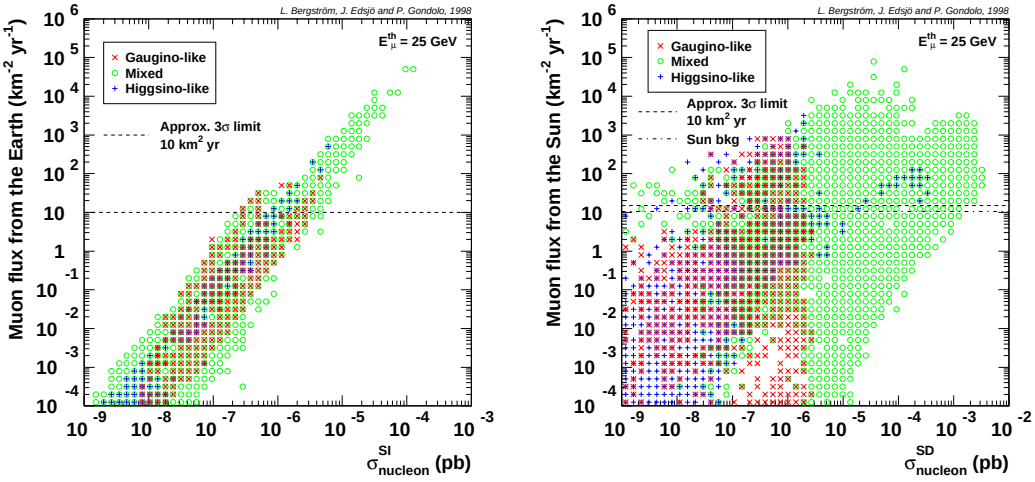

Figure 2.13: Comparison of direct- and indirect-detection event rates for MSSM WIMPs. Left: upward-going-muon flux for annihilation in the Earth vs. spin-independent WIMP-nucleon cross section. Right: upward-going-muon flux for annihilation in the Sun vs. spin-dependent WIMP-nucleon cross section. Figure taken from [59].

2.4 Indirect vs. Direct Detection

For a number of years, there has been debate as to the relative sensitivities of indirect- and direct-detection methods. The main problem is that comparison of the two methods is model dependent. The indirect-detection rate is determined by the model for the gravitational infall of the WIMPs into the Earth, Sun, or galactic center, by the WIMP-annihilation cross section, and, in some cases, the model for propagation of the products to the detector. On the other hand, the direct-detection rate is determined by the global halo model and the WIMP-quark scattering rate, which is not trivially related to the WIMP-annihilation rate.

Restriction to comparison of indirect detection via neutrinos emitted by annihilation in the Earth and Sun to direct detection reduces the model dependence. One such comparison has been made recently by Bergstrom, Edsjo, and Gondolo, shown in Figure 2.13 [59]. For spin-independent interactions, there is a strong correlation between direct- and indirect-detection event rates, with only 1 to 2 orders of magnitude of dispersion. This correlation arises because, to be indirectly detected, the WIMPs must become gravitationally bound to the Earth, which is only possible if the WIMPs can lose energy by elastic scattering with nuclei in the Earth. Current direct-detection experiments are reaching sensitivities of 10^{-42} cm^2 (10^{-6} pb); Figure 2.13 indicates that models with such a cross section will be probed by neutrino telescopes looking for annihilations in the Earth only after many years or significant size increase. For spin-dependent interactions, the situation is reversed. Current direct-detection experiments have limits in the range of 10^{-37} to 10^{-36} cm^2 (0.1 to 1 pb), two orders of magnitude or more above the highest cross-section WIMP models considered. On the other hand, the nominal 10 km^2 yr indirect-detection exposure on the Sun used in [59] promises to probe a nonnegligible fraction of models. Moreover, there appears to be almost

no correlation between the spin-dependent WIMP-nucleon cross sections and the indirect-detection rates.

It appears fair to say that direct-detection efforts will in general be more sensitive to WIMPs with spin-independent interactions, while neutrino telescopes will win for WIMPs with only spin-dependent interactions. Though the latter group has become more disfavored by minimal supersymmetry, the two methods are clearly quite complementary and both are being vigorously pursued.

Chapter 3

The Cryogenic Dark Matter Search Low-Background Facility and Cryostat

3.1 Introduction

The remainder of this dissertation describes the experimental apparatus of and results obtained during Run 19 of the shallow-site stage of the Cryogenic Dark Matter Search, CDMS I. In this chapter, I discuss the low-background aspects of CDMS I: the expected background-particle rates at the Stanford Underground Facility (SUF), the shield and muon veto designed to attenuate these fluxes, and the radiopure cryostat that holds the detectors. The detailed work to assess background-particle fluxes at SUF during the design of the shield and veto was done by a group consisting of Angela Da Silva from UCB/UBC, Brian Dougherty from Stanford, and Steve Yellin from UCSB. These studies defined the veto and shield design and helped to determine the expected sensitivity of CDMS I. Angela Da Silva's dissertation thoroughly describes this work [91]. The veto and shield were built by the UCSB group, including Dan Bauer, Dave Hale, Al Lu, Andrew Sonnenschein, and Steve Yellin and have been maintained by Dan Bauer, Andrew Sonnenschein, and Ray Bunker. Andrew Sonnenschein's dissertation covers in detail the performance of the veto as built [8]. During the operation of CDMS detectors at SUF between 1996 and 2000, simulation work for CDMS I has been taken over by Thushara Perera and Dan Akerib with significant contributions from Steve Eichblatt and Steve Yellin. Pre-screening of most of the materials placed inside the shield was performed by Donna Hurley, Dick McDonald, and Al Smith of LBNL's Low-Background Counting Facility.

CDMS's unique cryostat, the Icebox, was designed by a team consisting of many members of the Berkeley, LBNL, and Stanford groups. The Icebox was built at LBNL by a team led by Ron Ross and John Taylor in the Building 77 shops. Assembly and testing was done in the UCB campus labs, also by a large team. The design and construction of the Icebox comprised a significant part of Peter Barnes's graduate work and is fully documented in his dissertation [92].

It is questionable how much of this work to discuss here. Clearly, I cannot do as comprehensive or as high-quality a job as those who were intimately involved in the work. More appropriate is a brief discussion of the "why's" of the design of the veto, shield, and Icebox. However, I present

the topic of neutron backgrounds in great detail. As is discussed later, statistical subtraction of this background is performed via simulations. It is thus important to discuss the physics that produces these neutrons and to demonstrate that the simulations are reasonable. The consistency of the simulations and the data are discussed in Chapters 7 and 8. Thushara Perera will give a comprehensive and definitive presentation of this material in his dissertation [93]; I simply hope to demonstrate that the neutron background is sufficiently well understood to justify use of the simulations for statistical subtraction.

I note that the CDMS I sensitivity goal is 0.3 kg^{-1} d^{-1}, which is the event rate for a 100 GeV c^{-2} WIMP with WIMP-nucleon cross section $\sigma_{\delta \mathrm{n}} \approx 10^{-42}$ cm^2. Given the corresponding recoil-energy spectrum is exponential with $\langle E \rangle = 27$ keV, this integral rate corresponds to a differential event rate of ~ 0.01 keV^{-1} kg^{-1} d^{-1} at 10 keV.

3.2 The Stanford Underground Facility

The Stanford Underground Facility (SUF) consists of three tunnels and an associated loading platform at the south end of the Hansen Experimental Physics Laboratory's End Station III. Diagrams of SUF are displayed in Figure 3.1. The tunnels are unused beam dumps for the HEPL linear accelerator. Tunnel A has been extended and widened and houses the experimental apparatus. Tunnel B contains pumps and gas canisters necessary for operating the dilution refrigerator. Tunnel C is used primarily for storage of low-radioactivity materials: on the surface, these materials would experience cosmogenic activation by the hadronic component of cosmic-ray air showers. It also houses a low-background alpha-particle counter set up by Andrew Sonnenschein.

The top of tunnel A is approximately 16 meters-water-equivalent (10.7 meters of earth) below the surface. It is located underneath Panama Street near where it meets Morris Way. A concrete floor provides a stable platform for the experiment. The walls and roof of the tunnel are covered with shotcrete, a type of concrete that can be applied by "shooting" it against the surface to be covered, to retain the dirt walls of the tunnel. Masonry waterproofing sealant and paint provide a final sealing layer.

An air conditioning unit situated near the mouth of the tunnel feeds cooled air from the surface to the closed end of the tunnel; this ensures the experimental area is positively pressurized, both to reduce the buildup of radon gas and to minimize the flow of dirt and dust from the (extremely dirty) main hall of End Station III. The air is filtered with standard and electrostatic filters; HEPA (High Efficiency Particulate Air) filters would be preferable but are quite expensive for such a large air-flow rate.

A diagram of the cryostat and shield is given in Figure 3.2. As is discussed later, the cryostat containing the detectors is located inside the veto and shield; a complex coldfinger connects the cryostat, through holes in the shield and veto, to the dilution refrigerator. A plastic curtain divides the tunnel into two regions. The volume containing the cryostat, veto, and shield is maintained as a clean area; personnel working in this area are required to wear full cleanroom gowning. The remainder of the tunnel contains the dilution refrigerator and associated plumbing, a work bench, storage cabinets, and computers and electronics for data acquisition. This region is kept fairly clean by weekly mopping and use of clean room footwear, but no other precautions are taken.

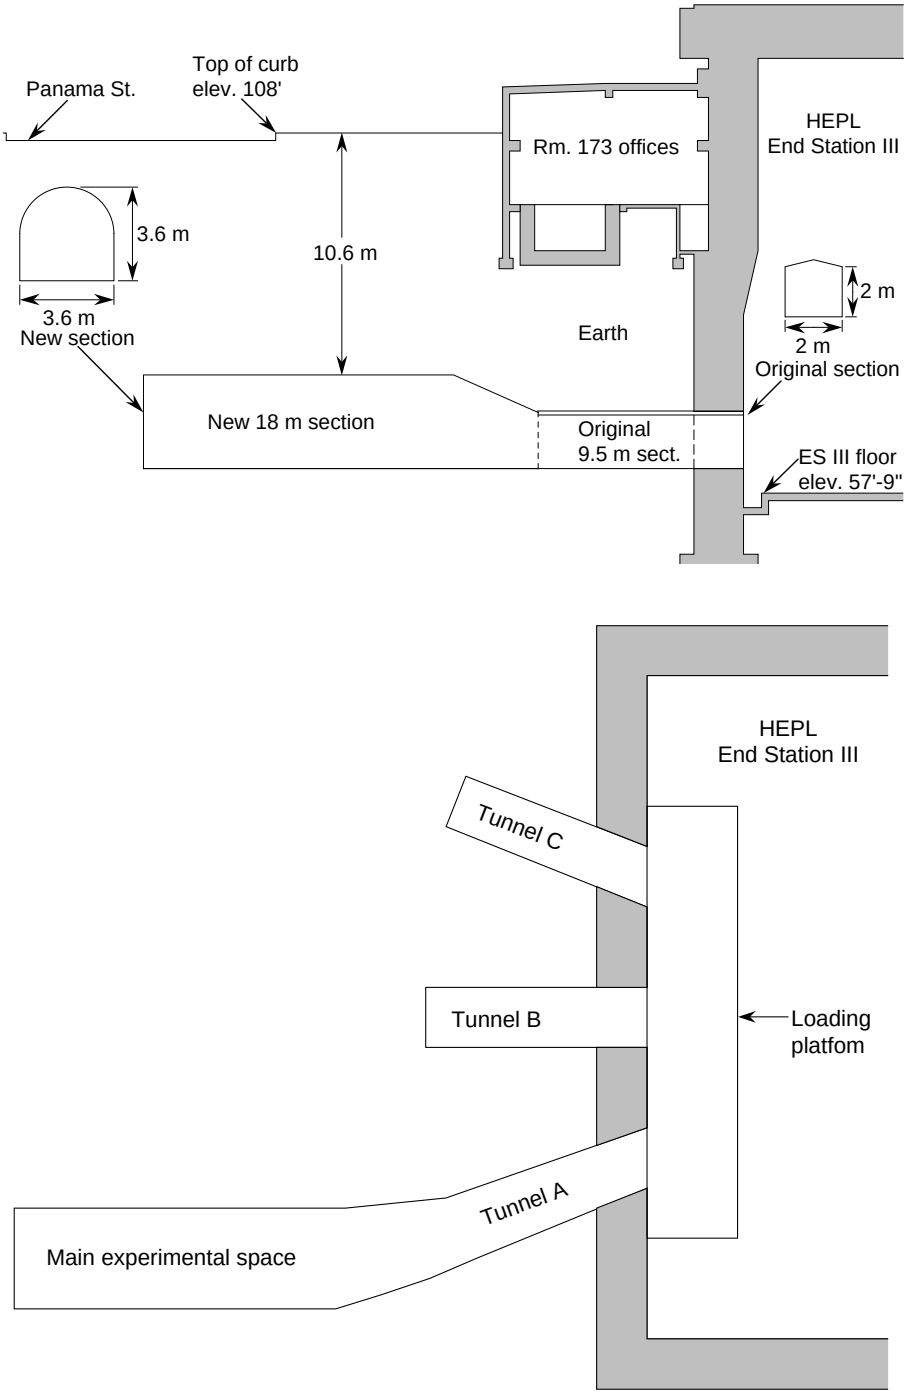

Figure 3.1: Elevation and plan views of the Stanford Underground Facility. Figures taken from [91].

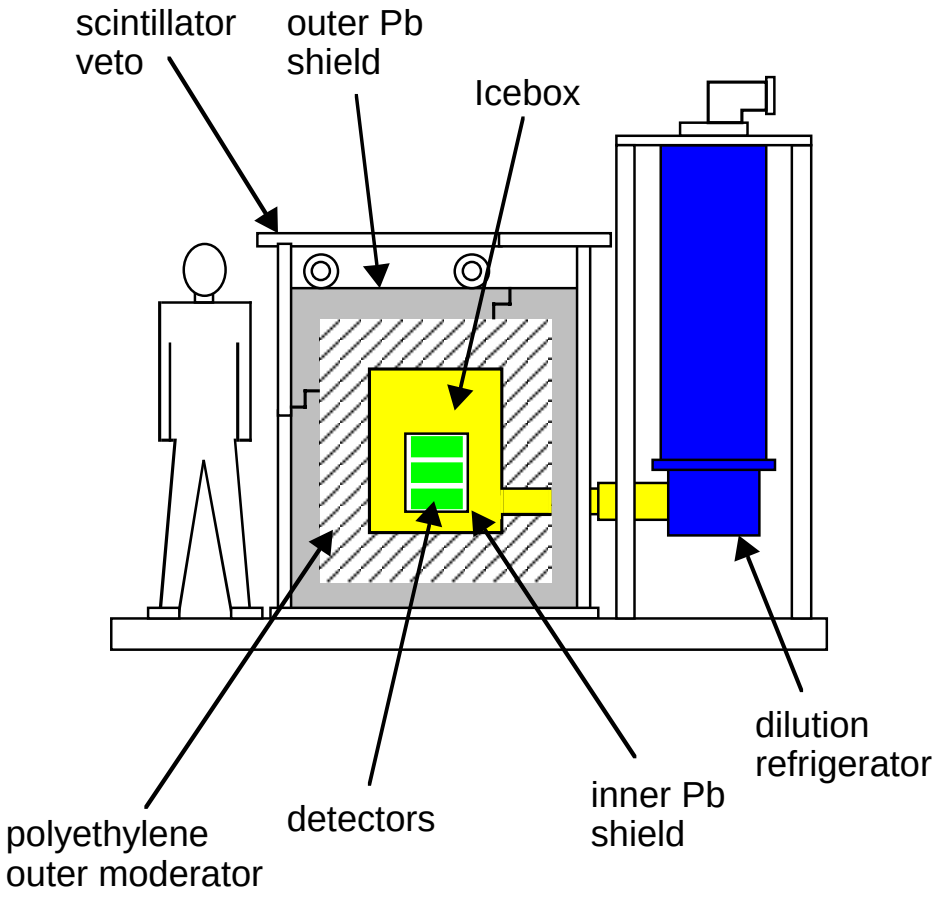

Figure 3.2: Schematic diagram of shield and Icebox.

3.3 Sources of Background Particles at SUF

In general, background-particle sources in low-background experiments can be divided into two categories: muon-induced and non-muon-induced. "Muon-induced" denotes particles that are produced promptly by a muon interaction in the experimental apparatus or surrounding material. These muons are produced in cosmic-ray-induced air showers. This division is useful to make because it highlights the differences between CDMS I and low-background experiments operating at deep sites and defines those backgrounds that will be greatly reduced in CDMS II, the next stage of CDMS, being built in the Soudan Mine at 2050 mwe. Muon-induced backgrounds are reduced by orders of magnitude at deep sites.

At this point, it is useful to summarize the structure of the CDMS I shield (Figure 3.3) because it defines the importance of various shielding components. Outermost is a plastic-scintillator veto to identify muons passing through the apparatus. Inside the veto is a 15-cm-thick lead shield to attenuate the photon flux from external sources. A 25-cm thickness of polyethylene between the outer shield and the cryostat moderates neutrons with energies $\lesssim$ 50 MeV. The cryostat, constructed almost entirely of high-purity copper, lies at the center of the polyethylene moderator. Inside the innermost layer of the cryostat, surrounding the detectors, is a 1-cm-thick lead shield cast from low-activity lead.

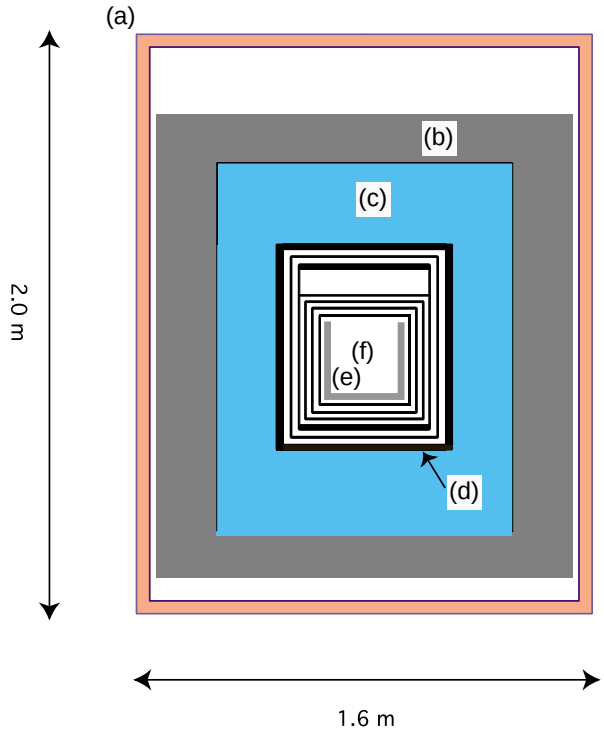

Figure 3.3: Diagram of CDMS I shield. a) muon veto, b) outer lead shield, c) polyethylene moderator, d) cryostat, e) inner lead shield, f) detector location. Figure taken from [8].

3.3.1 Muon-Induced Particles

Cosmic-ray muons arise from air showers initiated by high-energy cosmic-ray primaries interacting at the top of the atmosphere. These muons produce secondary particles by two different classes of processes.

Fast-Muon Interactions

Above several-hundred-GeV muon energy, muon interactions are dominated by radiative processes. These radiative processes are:

- *bremsstrahlung*: Like bremsstrahlung emitted by electrons, this consists of photons emitted due to the muon's acceleration by the Coulomb fields of nuclei it passes.
- *pair production*: Similar to bremsstrahlung, except that the photon is virtual and becomes an electron-positron pair in the final state. This is detailed in [94, 95].
- *photonuclear reactions*: Real photons are available from bremsstrahlung or as secondaries in showers initiated by δ-rays (see below) or pairs. The real photons, and virtual photons of the muon's electromagnetic field, may undergo photonuclear reactions via direct interaction with the electric dipole moment of the nucleus. When the process occurs through a virtual photon, it is usually termed "muon-nucleus inelastic scattering." In either case, the photon is usually absorbed in an E1 transition of the nucleus to an excited state. The excited state can decay to the ground state by photon emission. However, if the incident photon is of sufficient energy, there may be an allowed strong-interaction transition to a different nucleus, with the emission of one or more neutrons or other hadrons. At very high energies, the photons interact with individual nucleons and produce hadrons. The resulting high-energy hadrons may themselves be energetic enough to initiate hadronic showers via interactions with other nuclei. The products of photonuclear reactions are described in more detail below.

Fast-muon interactions are also understood to include

- *δ-ray production*: In high-q^2 Coulomb-scattering interactions of muons with electrons, the electron acquires significant energy; these electrons are called δ-rays or knock-on electrons. δ-ray production is not a significant energy-loss mechanism for fast muons, but the energy spectrum of δ-rays becomes harder as the muon energy increases and electromagnetic showers are initiated by δ-rays. In fact, at the SUF depth, δ-ray production is more important than bremsstrahlung and pair production in generating showers [96].

In brief, fast-muon interactions give rise to energetic electromagnetic and hadronic particles and showers.

Capture of Slow Muons

Muons lose energy by both the above fast-muon interactions and by ionization of atoms. Below several hundred GeV, ionization dominates. Once a muon has slowed down sufficiently, it may be captured by an atom, taking up residence in a Bohr orbital and quickly decaying down to the $1s$ level. The muon may be captured by the nucleus via $\mu^- + p \rightarrow n + \nu_\mu$. The resulting excited nucleus deexcites by direct emission or evaporation of neutrons.

Products of Muon Interactions

A variety of particles are produced by muon interactions, consisting of single electromagnetic and hadronic particles and of showers caused by these single particles. I discuss the particles produced by these processes in more detail.

Electromagnetic showers arising from fast-muon interactions or decay of stopped muons provide primarily a flux of electrons, positrons, and photons.

Photonuclear reactions result in emission of hadrons. There are three energy regimes [97]:

- At low photon energies ($<$ 10–30 MeV), the photon is absorbed in a transition of the nucleus between discrete energy levels and the nucleus may decay by emission of 1 or 2 nucleons (transitioning to a discrete level of a nucleus at $A-1$ or $A-2$). The nucleons have energies of only a few MeV, the typical nuclear binding energy per nucleon.

- In the 10–30-MeV region, there is typically a "giant resonance" due to the nuclear electric dipole moment. The resonance corresponds to a large increase in the photonuclear cross section. The nuclear states at these energies form a continuum; the excited state may still decay to lighter nuclei by emission of 1 or 2 nucleons. Again, the nucleons emitted only have energies in the few MeV range.

- Above the giant resonance, the photonuclear cross section is fairly flat. However, new processes become possible when the photon energy becomes high enough, increasing neutron production (multiplicity) in spite of the dropping cross section. For example, if the photon energy exceeds the pion mass (140 MeV), the excited nucleus may emit pions. This occurs through the following inelastic charge-exchange reactions:

$$\gamma + (Z, A) \rightarrow (Z-1, A-1) + n + \pi^+ \tag{3.1}$$
$$\gamma + (Z, A) \rightarrow (Z, A) + \pi^+\pi^- \tag{3.2}$$

  Above about 200 MeV, the photon wavelength becomes comparable to the size of the nucleus and photons scatter off individual nucleons. The following further inelastic charge-exchange reactions on individual nucleons become possible:

$$\gamma + p \rightarrow n + \pi^+ \tag{3.3}$$
$$\gamma + n \rightarrow p + \pi^- \tag{3.4}$$

  In addition to the protons and neutrons produced by the above processes, the pions themselves may also interact with the initial nucleus or with a different nucleus by the reaction

$$\pi^- + (Z, A) \rightarrow (Z-1, A-1) + n \tag{3.5}$$

  also known as a "π-nuclear reaction." This last reaction has a very low neutron production rate, but produces the highest-energy neutrons, the ones that "punch through" the shield. There are thus direct pathways for production of high-energy neutrons. Typically, low-energy neutrons are also emitted as the nucleus relaxes to its ground state by evaporative and direct-emission processes discussed below in the context of slow-muon capture.

Photons available in electromagnetic showers have a spectrum falling as E_γ^{-2} for $E_\gamma \ll E_\mu$ [98]. Virtual photons available from muons have a spectrum falling as E_γ^{-1} [99]. Thus, real and virtual photons primarily lie in the first two categories. The neutrons produced by the first two processes are soft, having energies of only a few MeV or less. The energy spectrum of such neutrons is similar to those produced by muon capture since similar processes, evaporation and direct emission, are responsible. As is discussed later, these neutrons are easily moderated, and thus it is the rare high-energy photonuclear reactions that are responsible for the production of high-energy neutrons in the tunnel walls that can punch through the shield.

The high-energy hadrons produced by the above high-energy photonuclear reactions may themselves initiate hadronic showers by interacting with other nuclei. Since the binding energy available in a nucleus is only a few MeV per nucleon, the energy of a hadronic shower is determined entirely by the energy of its primary. A high-energy photon is needed to start the process, and high-energy hadrons must be produced for the shower to propagate. Note that low-energy charged hadrons quickly stop due to ionization loss; low-energy neutrons can continue to propagate, but do not have enough energy to produce any other particles by nuclear interactions.

Finally, as noted above, stopped muons that are captured by nuclei give rise to direct emission and evaporation of neutrons.

The neutrons produced by the above processes can be modeled with three energy spectra. Neutrons produced by slow-muon capture in lead follow [91]

$$dN(E) = \begin{cases} 0.812\ (E/1\ \text{MeV})^{5/11}\ \exp(-E/1.22\ \text{MeV})\ dE & E < 4.5\ \text{MeV} \\ 0.018\ \exp(-E/9.0\ \text{MeV})\ dE & E > 4.5\ \text{MeV} \end{cases} \tag{3.6}$$

This spectrum is also used for copper, though it is not expected to be exactly correct; unfortunately, measurements of neutron production in copper are not forthcoming. The first term describes "evaporation" of neutrons — thermalized deposited energy heats the nucleus so that transitions to states differing by a single neutron become thermally activated. The second term describes direct emission: the nucleus is excited to a higher nuclear state and a direct transition to a state with fewer neutrons is possible.

For slow-muon capture in the tunnel walls, a muon-capture spectrum appropriate for the elemental composition (mainly silicon) of the rock is used [91]:

$$dN(E) = \begin{cases} 0.38\ (E/1\ \text{MeV})^{5/11}\ \exp(-E/1.7\ \text{MeV})\ dE & E < 4.5\ \text{MeV} \\ 0.10\ \exp(-E/7.0\ \text{MeV})\ dE & E > 4.5\ \text{MeV} \end{cases} \tag{3.7}$$

which has similar components, modified for the smaller Si nuclei.

It is generally accepted to use the muon-capture spectra for low-energy photonuclear reactions arising from fast muons because it is believed that the neutron production processes involved, evaporation and direct emission, are the same. In fact, because of the similarity of the spectra, some authors refer to muon-capture and low-energy photonuclear reactions collectively as muon-capture reactions [96].

The collective neutron production rates for slow-muon capture and for low-energy photonuclear reactions caused by fast muons at the SUF depth are as follows:

- lead: $243 \pm 9\ \text{kg}^{-1}\ \text{d}^{-1}$, measured precisely in [91].

- copper: ≈ 100 kg^{-1} d^{-1}. This number is only an estimate. An uncertainty of $\sim 20\%$ can be assigned to this rate based on Figures 4.11 and 4.13 and Tables 4.4 and 4.5 of [91].

- tunnel walls (rock): 41 kg^{-1} d^{-1}, from [100] as cited in [91] but not measured precisely *in situ*. As is discussed below, slow-muon capture in the rock does not contribute significantly to the expected neutron flux at the detectors, so precise knowledge of the production rate is unnecessary.

Finally, for high-energy neutrons produced by high-energy photonuclear reactions and associated hadronic showers, the production spectrum used is

$$dN(E) \propto \begin{cases} 6.05\ \exp(-E/77\ \text{MeV})\ dE & E < 200\ \text{MeV} \\ \exp(-E/250\ \text{MeV})\ dE & E > 200\ \text{MeV} \end{cases} \tag{3.8}$$

which is taken from Figure 4 of [101]. The authors of [101] note that "the spectra do not depend on the projectile (π, p , n, γ) and its energy provided the latter is greater than 2 GeV;" hence, this single two-component spectrum is used for the many different high-energy photonuclear and hadronic shower processes. The production rate is taken to be ≈ 4 kg^{-1} d^{-1}, yielding an integral flux of these neutrons into the tunnel of 6.6×10^{-5} cm^{-2} s^{-1}, but this value is quite uncertain; the true flux could as much as two times larger or smaller [96].

3.3.2 Non-Muon-Induced Particles

This category consists entirely of particles resulting from radioactive decay of unstable isotopes. The products are primarily photons, electrons, and positrons. Neutrons are produced in small amounts by fission and (α, n) reactions, with the α's coming from radioactive decays.

Relevant Radioisotopes

There are three elements possessing radioactive isotopes with large naturally occurring abundances and long-lived decay products: uranium, thorium, and potassium. Uranium and thorium possess complex decay chains that include α, β, and γ emission. These elements also have fissionable isotopes. ^{40}K, potassium's only naturally occurring radioisotope, decays via β decay or electron capture, the latter accompanied by a high-energy photon. Radioisotopes of these three elements are present in metallic construction materials because of their derivation from mined ores. Copper is the lucky exception because it is routinely purified through repeated electrolytic dissolution and redeposition; redeposition of uranium, thorium, and potassium is chemically disfavored. Plastics tend to be quite low in such radioisotopes because they are produced from natural gas. Possible contaminants would be ^{3}H and ^{14}C produced by nuclear spallation by cosmic rays. However, natural gas is mined from very old petroleum deposits, so such isotopes are expected to have decayed away. Germanium and silicon also tend to be quite clean because of the extensive purification performed during production of detector-grade crystals.

Potassium is, of course, present in human beings in large amounts. Potassium is transferred through fingerprints, sweat, and other human excretions. The use of gloves and other cleanroom wear during final cleaning steps and handling of cleaned materials is necessary to prevent deposition of ^{40}K.

Other radioactive isotopes are created by cosmic-ray spallation. Of particular interest are Mn, Fe, and Co isotopes created in copper and Ge isotopes created in the detectors. Stock material for the cryostat and detectors is stored underground to minimize creation of such isotopes, but complete prevention is impossible because of the need for some surface exposure during processing or machining. There are no fission or α-decay modes in these lower-Z radioisotopes, so they only give rise to electromagnetic background.

Neutrons from Natural Radioactivity

Neutrons produced by fission and (α, n) reactions follow a steeply falling spectrum [91]:

$$dN(E) \propto (E/1\text{ MeV})^{1/2} \exp(-E/1.29\text{ MeV})\, dE \tag{3.9}$$

These neutrons are easily moderated by the polyethylene shielding between the outer shield and the cryostat. Cryostat construction materials and the inner lead shield, which lie inside the polyethylene shielding, are not significant neutron sources because of their intrinsically low uranium and thorium (and daughter) levels. In her dissertation [91], Angela Da Silva demonstrated the above claims are indeed true.

Electromagnetic Backgrounds from Natural Radioactivity

Though CDMS detectors are able to discriminate electron recoils from nuclear recoils, such discrimination is not perfect, so it is necessary that the electromagnetic backgrounds do not exceed certain levels. At the initiation of CDMS, the detector photon-rejection goal was 99%. Therefore, to reach the CDMS I goal of 0.01 keV^{-1} kg^{-1} d^{-1} of nuclear-recoil candidates, the photon-background goal was set at 1 keV^{-1} kg^{-1} d^{-1}.

As Angela Da Silva discusses comprehensively in her dissertation [91], the uranium and thorium chains and ^{40}K yield photons with energies from tens of keV to 2.614 MeV. At low energies (< 100 keV), these photons are photoelectrically absorbed in the detectors. At higher energies, Compton scattering dominates and these chains yield a smooth spectrum of electron recoils up to the highest energy photon at 2.614 MeV. For any low-background experiment, these photons are the initial worry.

In addition, it is necessary to avoid introducing, or else shield, any materials containing significant levels of other radioisotopes. For example, steels contain ^{60}Co due to its use in the liners of production furnaces as a thickness monitor. Lead used for shielding against uranium, thorium, and potassium in the environment typically contains ^{210}Pb, which gives rise to a significant photon continuum. In general, any new material introduced into the apparatus must be prescreened for radioactivity. The measures taken against such backgrounds are discussed later in this chapter.

3.4 Expected Rates of Background-Particle Interactions

In this section, I discuss briefly the simulations done to determine expected background-particle rates for detectors run at SUF. Thushara Perera's disseration [93] will provide detailed description and checks of the simulations. I concentrate on neutron-background simulations because, in practice, these are the limiting background at SUF. Angela Da Silva's dissertation [91]

describes the electromagnetic-background studies and simulations performed during the design of the shield that created the present, pleasant situation, in which electromagnetic backgrounds do not limit the experiment's sensitivity.

3.4.1 Strategy for Simulation of Expected Backgrounds

There are clearly a large number of possible background sources at the SUF. To simulate all of them is not practical, or even necessary. A number of simplifications can be made. The discussion here is informed by the data acquired at SUF and so should be viewed as what we should worry about for our running experiment rather than what we would worry about if we were designing a new experiment.

Electromagnetic Backgrounds

It is not necessary to simulate muon-induced electromagnetic backgrounds because it is observed that the dominant electromagnetic background is due to natural radioactivity. The observed electromagnetic backgrounds coincident and anticoincident with veto are 60 keV^{-1} kg^{-1} d^{-1} and 2 keV^{-1} kg^{-1} d^{-1}. During the majority of the data set described here, the veto efficiency for muons that pass through the detectors was $> 99.9\%$. These muons probe the center of the muon veto. The efficiency of the veto near its edges is likely even higher because these regions are closer to the phototubes. Under the reasonable assumption that such a muon sample is representative of muons that produce electromagnetic showers that hit the detectors, one expects the muon-induced electromagnetic background to drop by a factor of 10^4 when veto anticoincidence is required. Clearly, this does not happen. The large thickness of shielding provided by the lead shields and the cryostat ensure that electromagnetic particles produced by muons outside the shield do not reach the detectors. Thus, the residual, veto-anticoincident electromagnetic background is not muon-induced either inside or outside the shield.

The non-muon-induced photon background does not limit the sensitivity of the experiment at present and is not likely to do so in the near future. The single-scatter photon-background rate in the 10-to-100-keV region is ≈ 0.5 keV^{-1} kg^{-1} d^{-1}. As discussed in Chapter 4, the photon rejection of BLIP detectors is $\gtrsim 99.9\%$. Thus, the expected photon leakage into the nuclear-recoil sample is $\approx$ 0.0005 keV^{-1} kg^{-1} d^{-1}, negligible compared to the CDMS SUF goal of 0.01 keV^{-1} kg^{-1} d^{-1}. Some simulation work has been done to understand this photon-background level assuming reasonable levels of radioisotopes in the construction materials but so far has been unable to yield a rate as high as that observed. Because of the secondary impact of the photon background on the CDMS SUF sensitivity, I do not discuss these simulations further.

The non-muon-induced low-energy-electron background is more troubling because of the CDMS detectors' ionization dead layer (see Chapter 4). The low-energy-electron rejection of BLIP detectors using α-Si/Al-Schottky electrodes is $> 95\%$ above 10 keV (shown in Chapter 4), so low-energy-electron rates of only 0.2 keV^{-1} kg^{-1} d^{-1} can limit the CDMS I sensitivity. Simulation of low-energy-electron backgrounds is extremely difficult given the complicated atomic physics involved. However, in future runs, CDMS will use ZIP detectors, which employ athermal-phonon pulse-shape rejection of surface events, yielding low-energy-electron rejection of $> 99.7\%$ above 20 keV [102, 103]. Work remains to demonstrate this high rejection down to desired thresholds of 5 to 10 keV.

Neutron Backgrounds

It is not necessary to simulate all sources of neutrons. A few simplifications can be made. As mentioned earlier, the rate of neutrons from natural radioactivity of materials inside the shield is negligible because of the careful choice of construction materials. Neutrons from natural radioactivity in the tunnel walls and outer lead can be ignored: their spectrum is softer than that of neutrons produced by muon-capture and low-energy photonuclear reactions, so they become negligible in comparison above about 3 MeV. Also, their soft spectrum ensures that they are well moderated by the polyethylene [96].

Muon-induced neutrons may be classified into a number of subgroups:

- Neutrons produced by muon capture and low-energy photonuclear reactions in the cryostat copper and inner lead shield — "internal neutrons." Based on the production rates per unit mass given earlier and the masses of the cryostat copper and the inner lead shields, the total production rates of neutrons in these components by slow-muon capture and low-energy muon-induced photonuclear reactions are estimated to be 3.2×10^4 d^{-1} and 1.0×10^4 d^{-1}, respectively. Since there is no moderator inside the cryostat or inner lead shield, neutrons produced in these two structures are not significantly moderated. As many as half of the neutrons might make it to the outer polyethylene, but this is a relatively small factor in this context.

- Neutrons produced by high-energy photonuclear processes and hadronic cascades in the cryostat and inner lead. Their production rate is not expected to be too different from the production rate of such neutrons in the rock, which is ≈ 4 kg^{-1} d^{-1}; this is negligible compared to the 100 kg^{-1} d^{-1} and 243 kg^{-1} d^{-1} production rates of neutrons by muon capture and low-energy photonuclear reactions. In addition, because these neutrons have such high energies, their cross section for interacting in the detectors is suppressed.

- Neutrons produced by muon capture and low-energy photonuclear reactions in the outer lead. The production rate of such neutrons is very high, 4.8×10^6 d^{-1}. The 25-cm-thickness of polyethylene significantly attenuates the flux of these neutrons. With this attenuation, these neutrons make up only about 3% of all muon-capture neutrons that reach the detectors — they are negligible.

- Neutrons produced by high-energy photonuclear reactions and hadronic cascades in the outer lead. *A priori*, one might think it necessary to simulate these neutrons because the high-energy tail (> 50 MeV) is not efficiently moderated: such neutrons may punch through the moderator, leading to a veto-coincident high-energy tail to the neutron-scattering spectrum in the detectors. Low-statistics simulations indicate the rate of detector interactions due to such neutrons is about 19 kg^{-1} d^{-1}, 2.3 times the rate due to punch-through neutrons from the rock (see next item and Table 3.1). This is approximately 5% of the internal-neutron rate from the cryostat and the inner lead (see Table 3.1) and so can be ignored. Also, empirically, the observed rate and spectrum of muon-coincident neutrons are reproduced without inclusion of such processes.

- Neutrons produced by muon interactions outside the veto (in the tunnel walls surrounding the experiment) — "external" neutrons. These form a separate class because they are anticoinci-

dent with the veto — the muon does not pass through the veto. Neutrons from muon capture and low-energy photonuclear reactions in the rock are moderated by the outer moderator and thus are expected to yield a negligible rate in the detectors. However, neutrons from high-energy photonuclear reactions and high-energy hadronic showers, and their secondaries produced in the outer lead shield, can punch through the moderator. This punch-through flux is the main source of veto-anticoincident neutrons.

Given the above considerations, I only discuss in detail simulations of neutrons created by: a) muon capture and low-energy photonuclear reactions in the cryostat copper and inner lead shield ("internal" or "veto-coincident"); and b) high-energy photonuclear reactions and hadronic cascades outside the veto ("external" or "veto-anticoincident").

3.4.2 Simulation of Neutron Backgrounds at SUF

Neutron backgrounds are simulated using the MICAP and FLUKA extensions to the well-known GEANT particle-physics simulation package. The results presented here are for the Run 19 geometry, consisting of the veto, outer lead shield, outer moderator, cryostat, inner lead shield, and detector package, shown in Figure 3.3.

The MICAP and FLUKA packages track neutrons above and below 20 MeV, respectively. At present, no attempt is made to simulate the production of the neutrons. GEANT/GHEISHA and standalone FLUKA can in principle do this, though at present GEANT/GHEISHA gives a neutron production rate off by a factor of 10. The use of FLUKA for this purpose will be investigated in the future. Presently, the neutron-production rates and spectra given earlier are assumed and only the propagation of the neutrons and their interactions in the detectors are simulated.

As noted above, the neutron-production rate by muon capture and low-energy photonuclear reactions in the cryostat and the inner lead is known precisely. Absolute predictions of the simulation can be compared to the data to check the accuracy of the simulation and of the data analysis. On the other hand, the normalization of the external-neutron production spectrum, Equation 3.8, is not well determined at the SUF depth of 16 mwe; there is a factor of ~ 4 uncertainty in the production rate alone, as noted earlier. Additional uncertainty arises because some fraction ($\sim 40\%$) of external neutrons interact in the veto scintillator and are thus tagged as muon-coincident. An unknown fraction of the hadronic showers associated with neutron production may also trigger the veto. For these reasons, the overall normalization of the external-neutron simulation is completely disregarded; only the spectral shape and other normalization-independent quantities are used in later analysis. Figure 3.4 shows the production spectra for internal and external neutrons. The simulated external neutrons are generated in the tunnel walls over the first 5-meter depth, so the simulated spectrum of neutrons that emerge from the rock into the tunnel is also shown.

Figure 3.5 shows the simulated energy spectra of the neutrons that are incident on the detectors for both external and internal neutrons. Note that, due to the high energy of the external neutrons, the neutrons thrown by the simulation outside the shield (the primaries) generate a significant flux of secondary neutrons through inelastic scattering with and spallation of nuclei in the lead shields and cryostat. The neutrons designated as "external" in the plot are dominated by these secondaries and tertiaries. These primaries, secondaries, and tertiaries are collectively referred to as "external" when it causes no confusion. The external-neutron spectrum shows a

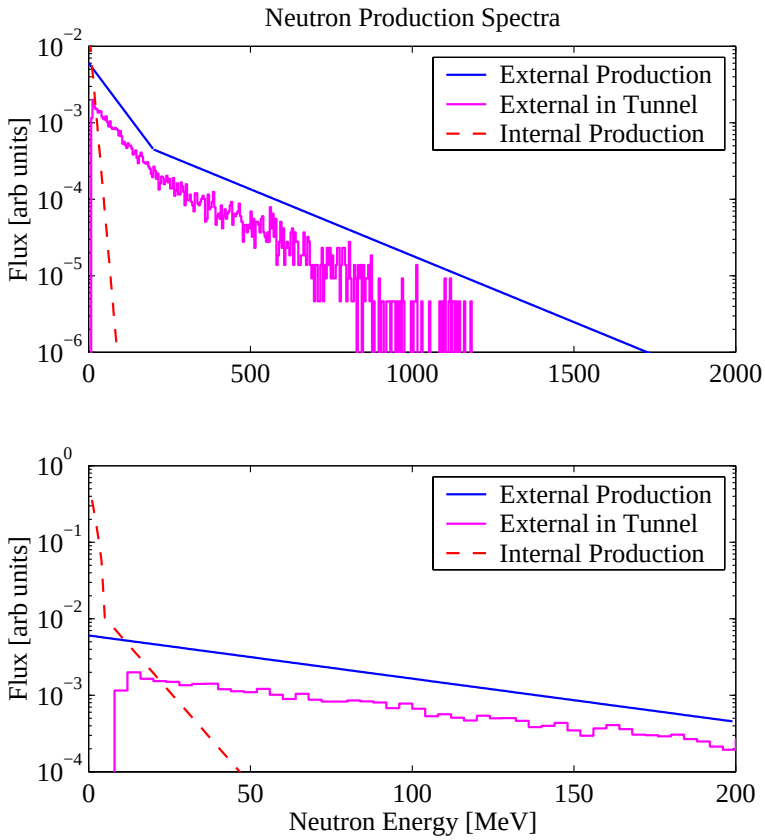

Figure 3.4: Production spectra of internal and external neutrons with two different energy scales. Also shown is the spectrum of external neutrons after propagation through the tunnel walls. The three spectra are each arbitarily normalized; the relative amplitudes should not be used.

significant tail to high energies that reflects the harder spectrum used. However, because of the large number of secondaries produced, the shapes of the incident spectra of internal and external neutrons are very similar below a few MeV: large numbers of secondary and tertiary neutrons are produced by the same evaporation and direct-emission processes that give rise to internal neutrons after slow-muon capture.

To illustrate the effect of the variation of the neutron-interaction cross section with energy, Figure 3.6 shows the energy spectrum of the neutrons that *interact* in the detectors. Only neutron energies for "first hits" — the neutron that produces the first detector interaction out of all neutrons produced by a given primary — are shown. These neutrons are selected to ensure that the histogram is not filled repeatedly for an incident neutron that multiply scatters. Comparing Figure 3.6 to Figure 3.5, the effect of the rise in the scattering cross section at low neutron energies is apparent: the first-hit spectrum is more steeply rising at low energies than the incident spectrum. The feature in both spectra at 20 MeV is, presumably, an artifact due to the cross-over between using MICAP and FLUKA. As is seen below, this feature has little effect on the sub-200-keV detector recoil-energy spectrum.

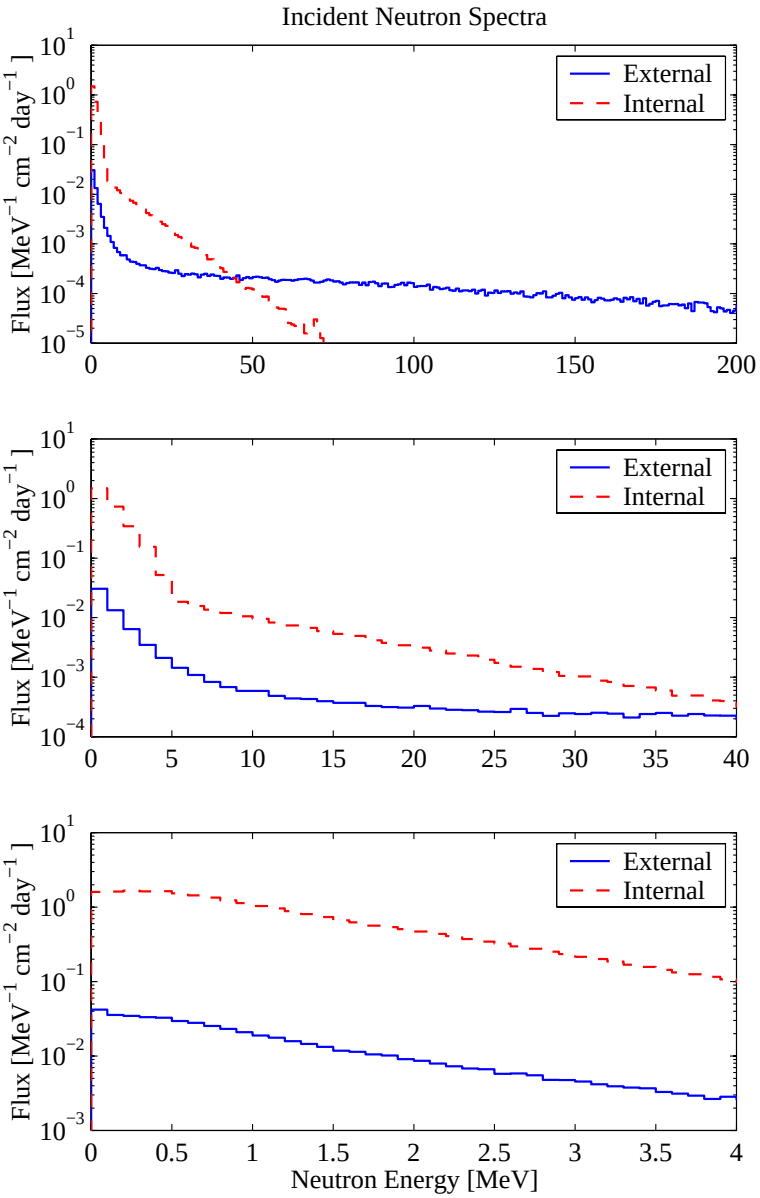

Figure 3.5: Energy spectra of neutrons incident on the detectors. The number of decades shown on the vertical axes is the same as in the Figure 3.6.

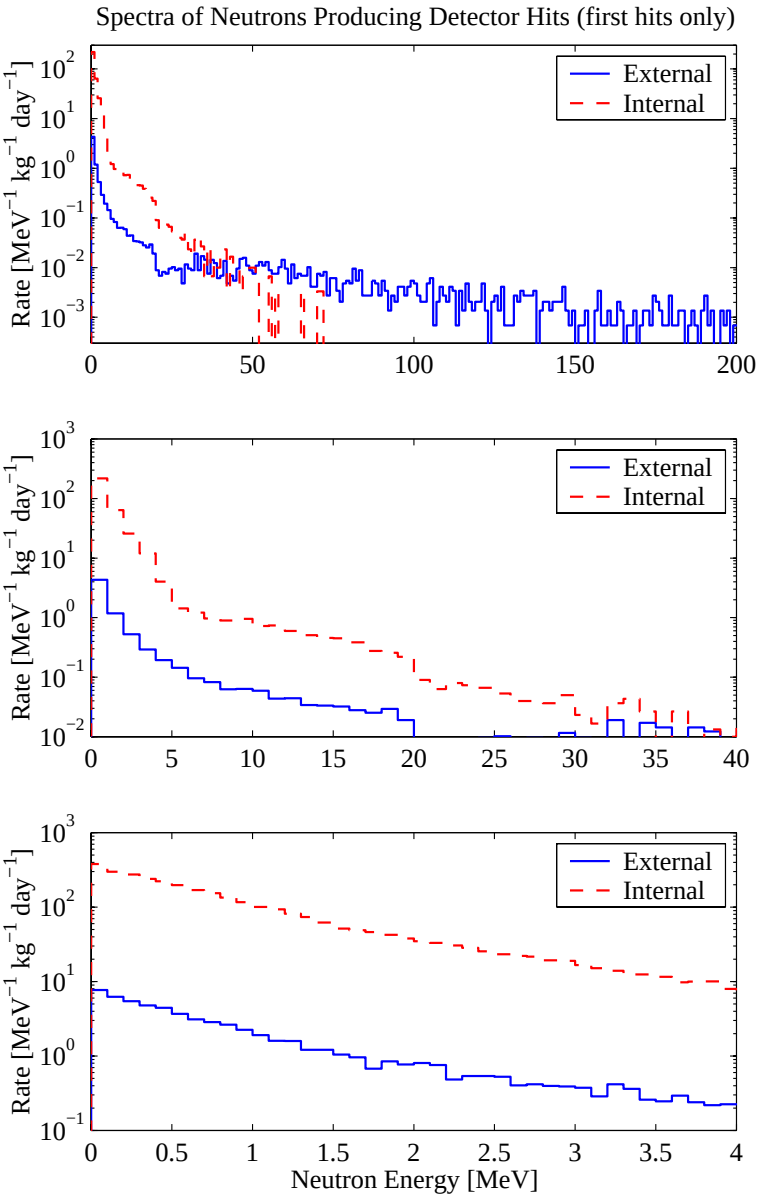

Figure 3.6: Energy spectrum of "first-hit" neutrons that interact in the detectors. The number of decades shown on the vertical axes is the same as in the Figure 3.5.

To see the rise in cross section more clearly, Figure 3.7 shows the ratio of the "first-hit" spectrum to the incident spectrum. The interaction rate of neutrons in the detectors for an incident neutron flux F is

$$R = F\sigma nV \tag{3.10}$$

where σ is the neutron-scattering cross section, n is the target material density, and V is the volume of target material. Figure 3.7 plots R/F and so gives the cross section modulo an energy-independent normalization. The feature at the MICAP-FLUKA cross-over point is again visible. Ignoring this dip, a gentle rise in the cross section as the energy decreases from 100 MeV to a few MeV is apparent. There is a significant additional smooth rise below 1 MeV. The precipitous increase in the first bin is an artifact; the low-energy detector recoil-energy spectra show no sign of such a rise except in their very lowest (0–2-keV) bin.

Interactions of neutrons with detectors are dominated by neutron-nucleus elastic scattering. At the low neutron energies ($\lesssim$ few MeV) of interest, the scattering is s-wave, or isotropic, so the spectrum of energies of the recoiling nucleus is flat as for WIMP-nucleus scattering. The maximum energy that can be given to the nucleus is

$$E_{R,\max} = 4\frac{M_n}{M_N}E_n \tag{3.11}$$

For Ge, the prefactor is approximately 0.05. This simple scattering yields a direct correspondence between the "first-hit" neutron spectrum and the detector recoil-energy spectrum. The recoil-energy spectra for internal and external neutrons are shown in Figure 3.8. As one would expect from the similarity of the first-hit neutron energy spectra below 20 MeV and the rise in the detector-interaction rate at low energy, the recoil-energy spectra produced by internal and external neutrons are extremely similar up to 1-MeV recoil energy. The high-energy tail of the incident external neutrons is suppressed by the high-energy drop in the interaction cross section, so this tail is barely visible in the recoil-energy spectrum. The region of interest for WIMP searches is below 200-keV recoil energy, where the two spectra are almost identical.

The primary conclusion of the above is that the sub-200-keV recoil-energy spectrum produced by external neutrons is very similar to that produced by internal neutrons because of the large numbers of secondaries and tertiaries produced by the high-energy external primaries. The shape of the $<$ 200-keV portion of the external-neutron recoil-energy spectrum is insensitive to the particulars of the external neutrons' high-energy tail — changing the tail presumably would change the number of secondaries and tertiaries produced, renormalizing the spectrum, but would have little effect on the shape. Since only *normalization-independent* predictions from the external-neutron simulation are used in later analysis, these predictions are as trustworthy as the internal-neutron simulation.

Based on these simulations, a table of expected neutron-background rates is given in Table 3.1. Note that the integral rates are integrated from 0 to ∞ and so are only indicative, not exact predictions, of observed detector rates. When comparing the observed data to the simulations in Chapters 7 and 8, the simulated spectra are corrected for thresholds, efficiencies, etc. to make precise predictions.

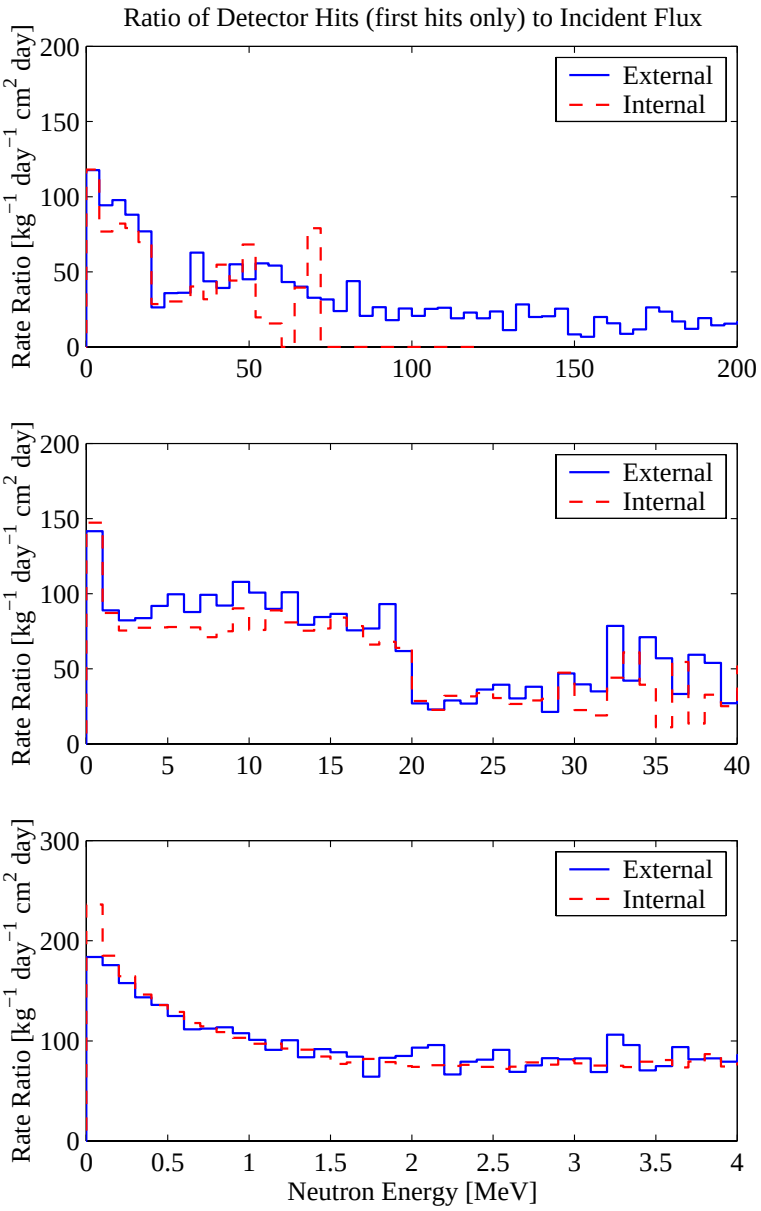

Figure 3.7: Ratio of spectra of "first-hit" and incident neutrons. Note the rise in the ratio ($\propto$ cross section) below 1-MeV neutron energy.

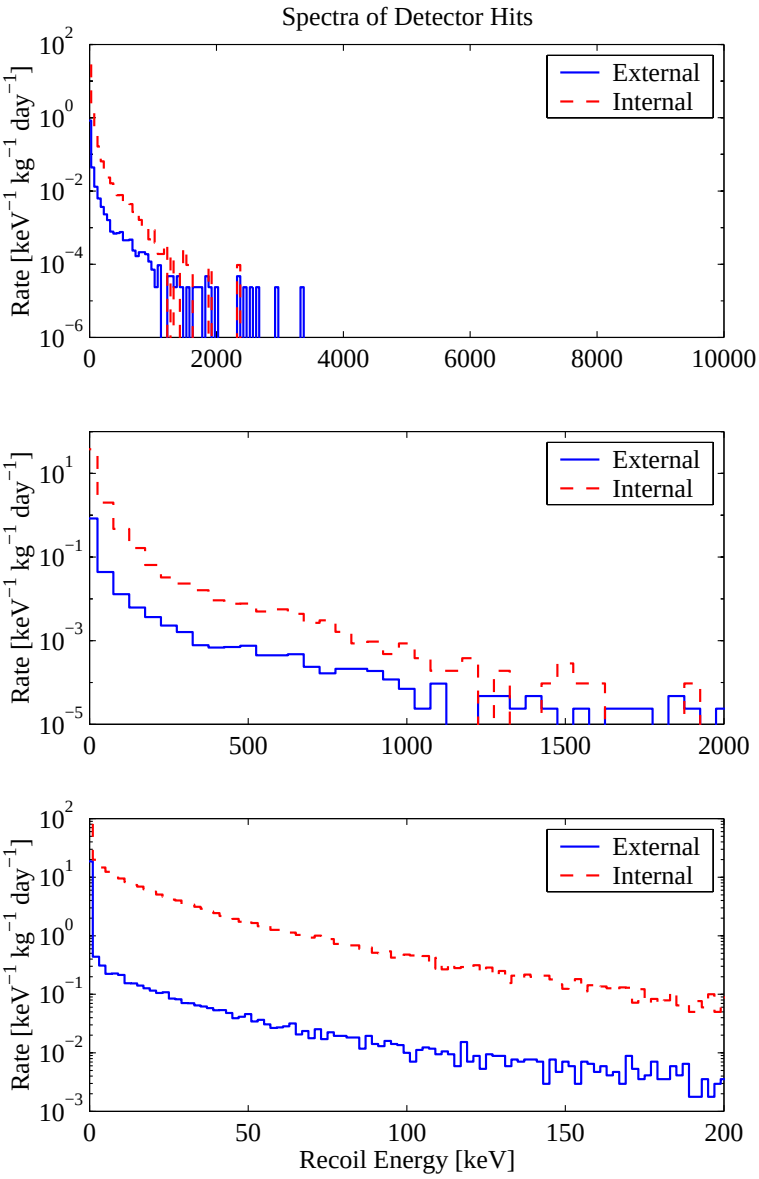

Figure 3.8: Recoil-energy spectra produced by internal and external neutrons. The energy scales are 5% of the scales shown for the incident and "first-hit" spectra to make clear the neutron-energy-recoil-energy correspondence.

Source	Rate[1] [kg^{-1} d^{-1}]	Rate[2] [kg^{-1} d^{-1}]	Rate[2]at 10 keV [keV^{-1} kg^{-1} d^{-1}]
Internal			
cryostat copper	189±1	219±1	5.4±0.1
inner lead	177±1	204±1	4.7±0.1
total	366±1.4	423±1.4	10.1±0.14
outer lead	~ 11	~ 13	~ 0.3
External			
rock	8.4±0.1	9.8±0.1	0.22±0.01

Table 3.1: Expected neutron-interaction rates at SUF. [1]: neutron rate: neutrons that multiply scatter are counted only once. [2]: detector-interaction rate: neutrons that multiply scatter are counted once for each detector interaction produced. The total internal-neutron rate combines only the cryostat-copper and inner-lead neutrons. Rough values for the neutron rate from the outer lead are shown, based on low-statistics simulations; these neutrons are not included in the spectra shown. Numbers provided by T. Perera.

3.5 Veto and Shield Design and Performance

I discuss briefly the design of the veto and shield and their performance during Run 19. The basic design of the system has been discussed earlier and is shown in Figure 3.3. It consists of five layers.

3.5.1 Muon Veto

The experiment is surrounded by a 5-cm-thick plastic-scintillator veto. Each exposed side of the experiment is covered by 1 or 2 large panels of scintillator. Waveshifter bars are glued to at least three edges of each panel. The waveshifter bars are viewed by phototubes at their ends. Since essentially all muons are downward going, complete coverage of the bottom face is unnecessary. Four slabs of scintillator, each with a single phototube at one end, cover approximately 70% of the bottom face. Finally, a crack counter, consisting of a single slab of scintillator viewed by a phototube at each end, covers the joint between the two top face counters. An illustration of the 13 counters is given in Figure 3.9.

Two kinds of particles produce events in the veto counters. Muons pass through the scintillator, depositing approximately 10 MeV. The deposition varies with the angle of incidence. Photons from natural radioactivity outside the experiment yield energy depositions of up to 2.614 MeV through photoabsorption, Compton scattering, and pair production. The pulses from the phototubes for each counter are summed and thresholds on the summed signal set to, in principle, trigger on muons and not on photons. However, due the large size of the counters, light collection varies across the counters and perfect discrimination is not possible. In practice, the thresholds are set to accept some photons in order to maintain high efficiency for muons. However, a total veto-trigger rate of 10 kHz yields 25% dead time when a veto-coincidence window of 25 μs is used, so the threshold can not be lowered arbitrarily. The total veto-trigger rate during normal operation is approximately 6 kHz.

As noted above, even with the polyethylene moderator, approximately 400 kg^{-1} d^{-1} of neutron interactions are expected in the detectors. To keep the anticoincident (untagged) internal-neutron rate below the 0.3 kg^{-1} d^{-1} sensitivity goal, the veto has to be $> 99.93\%$ efficient. The efficiency of the veto for muons that produce neutrons is nontrivial to measure because it is impossible to acquire an internal-neutron sample with enough events and with no misidentified anticoincident photon background leaking into the sample. It is, however, easy to measure the veto efficiency for muons that hit the detectors. The highest energy photon from natural radioactivity lies at 2.614 MeV, while normally incident muons deposit about 13 MeV in the detectors. α-particles or fission daughters can be quite energetic, but are easily distinguished by their low ionization yield. Therefore, any event with recoil energy above about 10 MeV is a muon. The veto efficiency is given by the fraction of these that are tagged as veto-coincident.

These muons are a reasonable sample to use to measure the veto efficiency. First, the efficiency of a veto counter is expected to be worst at its center because the phototubes are at the corners. Detector-tagged muons sample these regions reasonably well because the detectors lie in the center of the shield. Second, as mentioned above, 97% of internal neutrons are generated by muons stopped in the cryostat or inner lead shield. Such muons pass close to the detector and thus have similar trajectories to detector-tagged muons. Finally, the veto-counter pulse-integral spectra of detector-tagged muons and of muon-coincident neutrons are similar, as is shown in Figure 3.10;

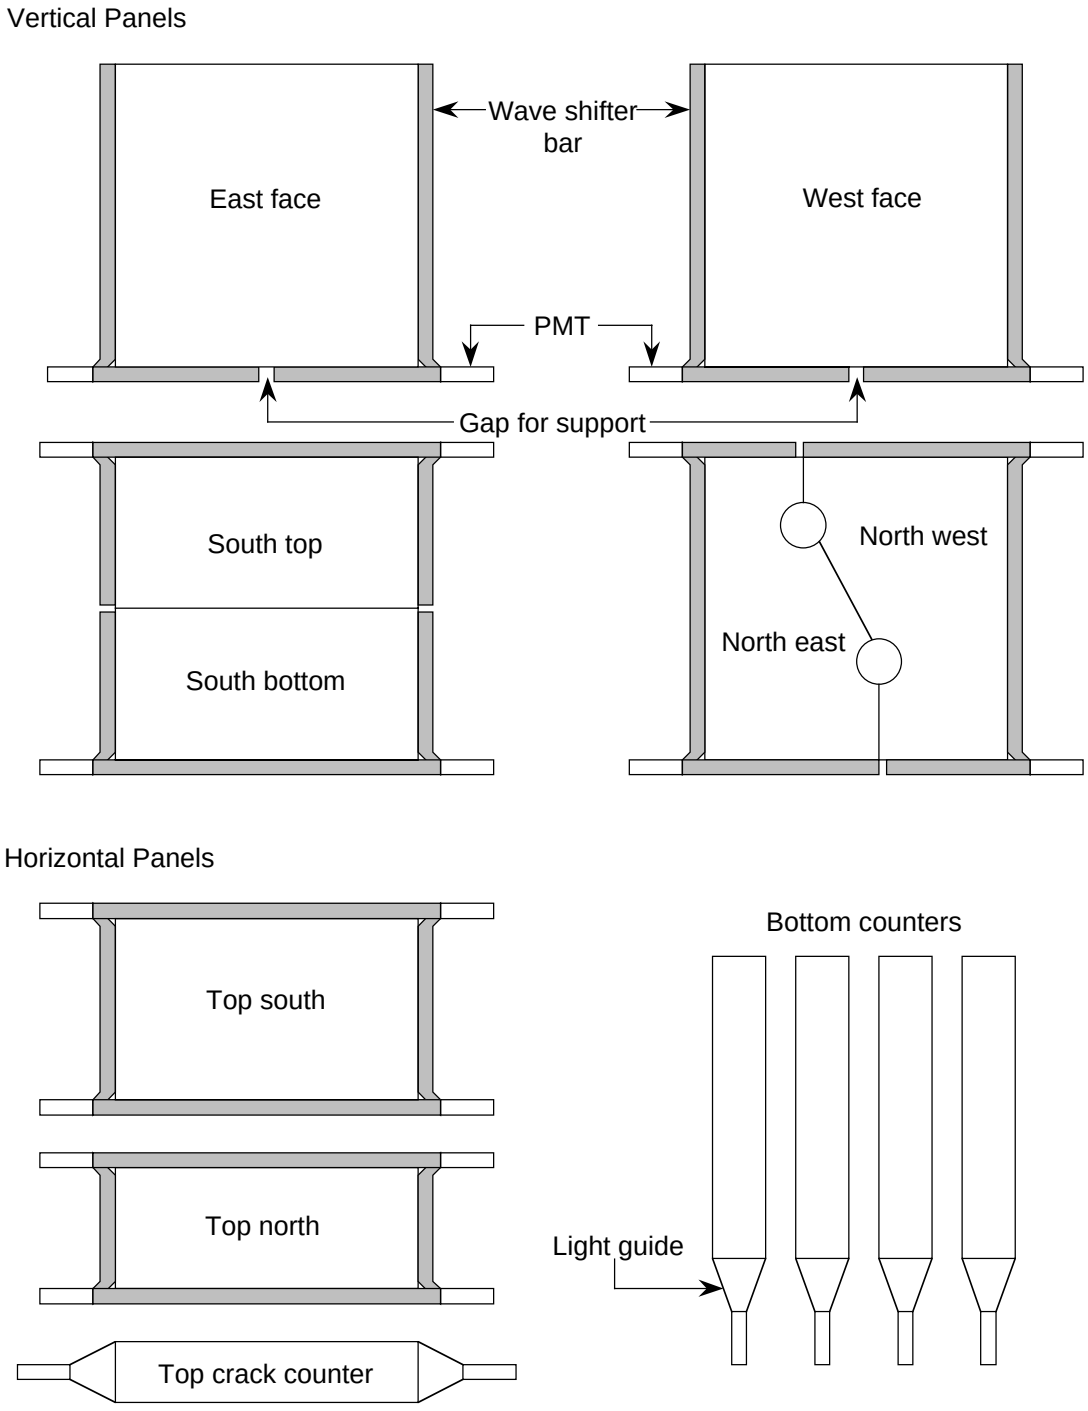

Figure 3.9: Schematic of the 13 veto counters. Figure taken from [91].

that is, internally generated neutrons do not appear to be produced by atypical muons. None of these arguments is ironclad: seeing that the bulk of muons that give rise to internal neutrons behave similarly to detector-tagged muons says nothing about possible tails of the pulse-integral distribution of the former. These arguments do give some confidence that the detector-tagged-muon efficiency is not a gross misestimate of the efficiency for muons that generate internal neutrons.

Using the above method, the veto efficiency has been measured to be at least 99.99% when all the counters are in good repair. However, it is difficult to maintain such a state. Helium gas from the daily liquid-helium transfers diffuses into phototubes and decreases their gains over time. Glue joints may fail due to age or abuse. At various points during Run 19, the veto efficiency dropped to 99.7%. The mean value of the veto efficiency was 99.9%. The variation of veto efficiency with time and the possibility that the neutron-scattering events observed in the Run 19 data set are due to veto inefficiency are discussed in Chapter 8.

3.5.2 Lead Shields

The purpose of the outer lead shield is to shield the detectors from the radioactivity of the tunnel and various pieces of the experimental apparatus. The dilution refrigerator contains high levels of ^{60}Co in stainless steel and other metallic components, contains long-lived cosmogenic isotopes of Ag in the silver solder used for many joints, and uses indium (whose most abundant isotope decays by β emission) for cryogenic seals. The veto-counter phototubes are made of glass. In general, the outer lead shield allows us to be fairly cavalier about these sources because they are outside the shield.

An important element of the shield design is the dual inner/outer lead shields. The outer shield is made of off-the-shelf lead, which contains the isotope ^{210}Pb (the other uranium and thorium chain daughters are removed during smelting). ^{210}Pb gives rise to the chain

$$^{210}\mathrm{Pb} \xrightarrow{\beta^-} {}^{210}\mathrm{Bi} \xrightarrow{\beta^-} {}^{210}\mathrm{Po} \xrightarrow{\alpha} {}^{206}\mathrm{Pb} \tag{3.12}$$

The first β decay has $Q = 63.1$ keV, but the second decay has $Q = 1.1615$ MeV. The resulting high-energy electron emits bremsstrahlung as it traverses the lead. Thus, increasing the thickness of lead shielding yields a saturating photon-background level related to the intrinsic ^{210}Pb level of the lead. The 15-cm thickness of the outer lead shield is close to the saturation thickness [91]. To further reduce the photon background, an additional 1-cm-thick lead shield cast from "ancient" lead is installed inside the innermost layer of the cryostat. This lead, referred to as Nantes lead, was recovered from the ballast of a French ship that sunk in the 1800s; the age of the lead ensures that the ^{210}Pb level has decayed significantly ($T_{1/2} = 22.3$ years).

A thin layer of copper lines the inside of the outer lead shield. It provides structural support and prevents lead dust and shavings from migrating into the cryostat.

3.5.3 Polyethylene Moderator

As has been discussed, polyethylene moderator is necessary to attenuate the flux of neutrons produced by muon capture and low-energy photonuclear reactions in the massive outer lead shield. It also moderates neutrons produced by similar reactions outside the shield. Fission neutrons and (α, n)-reaction products created outside the shield are also moderated. Polyethylene fills

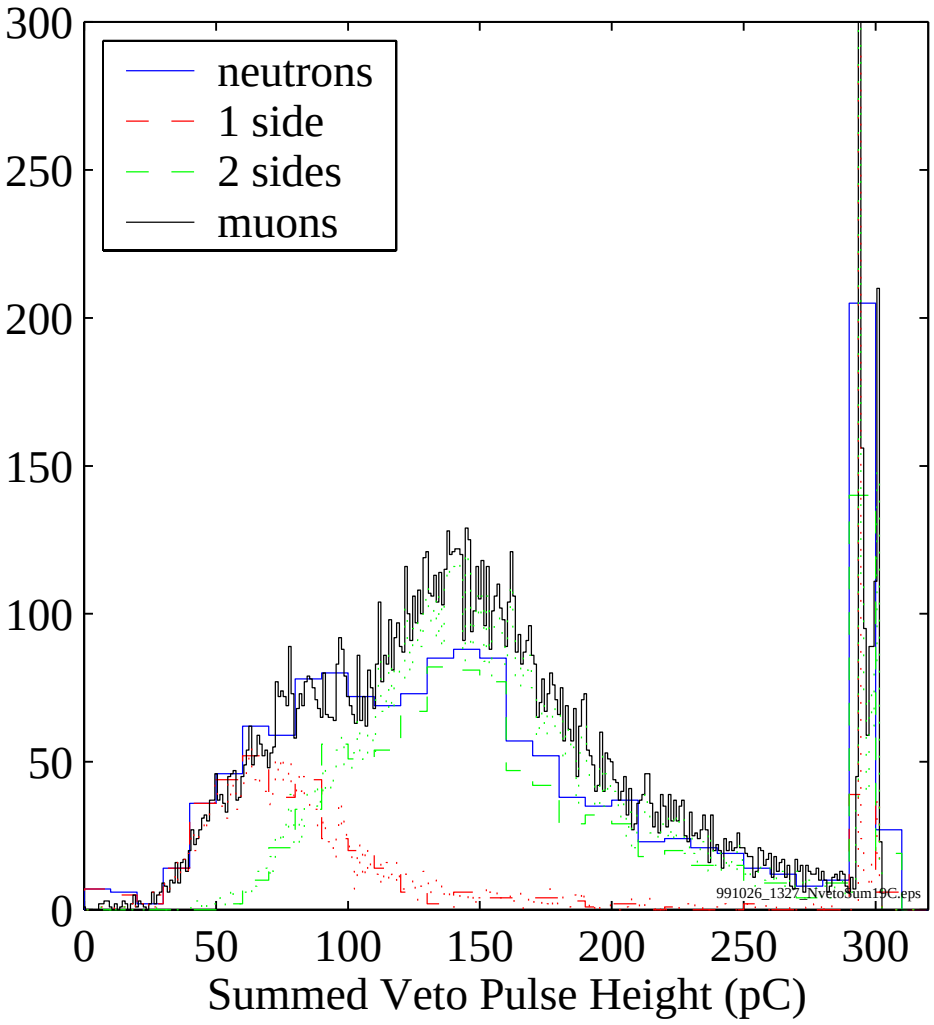

Figure 3.10: Veto-sum pulse-integral spectra for detector-tagged muons and for muon-coincident detector-tagged neutrons. Finely binned histograms are for muons, coarsely binned for neutrons. Dark, finely binned, solid line: all detector-tagged muons. Dark, coarsely binned, solid line: all detector-tagged neutrons. Left, finely binned, dotted line: detector-tagged muons that trigger only one veto side. Left, coarsely binned, dashed line: detector-tagged neutrons where the muon triggers only one veto side. Center, finely binned, dotted: detector-tagged muons that trigger two veto sides. Center, coarsely binned, dashed: detector-tagged neutrons where the muon triggers two veto sides. The similarity of the veto pulse-integral spectra for the muon-coincident neutrons and detector-tagged muons indicates detector-tagged muons can be reasonably used to measure the veto efficiency for muons that generate neutrons. Note, however, at very low pulse integral, there are events associated with detector-tagged neutrons but many fewer events associated with detector-tagged muons. Figure provided by R. Schnee.

the entire volume between the cryostat and the outer lead shield. The thickness varies from 20 cm to 28 cm; 25 cm is taken as the nominal thickness. At this thickness, neutron production in the wax itself starts to become significant; further increase of the thickness would result in saturation of the internal-neutron rate [91].

Polyethylene can be a source of electromagnetic background. Samples of the polyethylene used were prescreened for contamination by standard radioisotopes and no significant levels of uranium, thorium, or potassium were found. However, as shown in [91], neutron capture on hydrogen results in a 2.223-MeV photon line that increases the continuum background below 300 keV by a factor of a few. This effect has not been tested in the final configuration, but there is no reason to expect it to be absent.

An additional layer of polyethylene shielding may be placed inside the cryostat, inside the inner lead shield. This provides significant moderation of the internal neutrons produced in the inner lead shield and the cryostat, thus relaxing the 99.93% veto-efficiency requirement. More importantly, it moderates the secondaries produced by external neutrons; the simulation indicates that such additional shielding reduces the event rate due to external neutrons by a factor of 2.3. As is discussed later, this neutron background presently limits the sensitivity of CDMS. Internal polyethylene was installed in February, 2000, in preparation for Run 20. Confirmation of the reduction awaits data from Run 20.

3.5.4 The Icebox

The cryostat itself is part of the shield. Its main role is to provide a radiopure volume that is separated from the radioactive dilution refrigerator. To this end, a dogleg design is used (see Figure 3.6) so that there is no direct line of sight from the bulk of the refrigerator to the detector volume. However, to serve its purpose, the cryostat itself must have low radioactivity.

The Icebox is fabricated from oxygen-free electronic (OFE) copper. Higher purity is available by electroforming of parts, but at significantly higher cost. By prescreening samples using sensitive photon spectrometers at LBNL's Low-Background Counting Facility (LBF), it was checked that lower-cost OFE copper is sufficiently clean. Fasteners present a special issue: copper screws would no doubt break easily, so brass is used for screws. Brass is an alloy of copper and zinc with a few percent lead and so contains ^{210}Pb. However, a source of fairly low-activity brass was found, and the mass of brass is quite small. Copper gaskets are used for all cryogenic vacuum seals. Fiberglass and G-10, standard cryogenic supports having low thermal conductivity, are avoided because of their known high contamination levels. Instead, kevlar loops are used to suspend the various temperature layers of the Icebox concentrically from the outermost layer; kevlar has low radioactivity and incredible strength. Nylon standoffs are used at many points to separate different layers while warm; they are not in contact when cold.

3.5.5 Other Low-Background Precautions

As mentioned earlier, a number of operational precautions are taken to prevent contamination of the detector volume by human activity or by dirt and dust that can carry radioisotopes. All materials used inside the cryostat are also prescreened by the LBNL LBF; primary among these are the materials used in the detector mounts and wiring.

^{222}Rn must be suppressed because one of its daughters is ^{210}Pb. ^{222}Rn occurs in the uranium chain and has a 4-day half-life. It is a gas and so diffuses out of the tunnel walls. Because of its short half-life, it quickly decays and its daughters plate out on surfaces. A significant increase in ambient background due to the presence of radon was seen in [91]. To prevent radon from entering the shield, N_2-boiloff gas from the the dilution refrigerator's liquid-nitrogen bath is vented into the shield. This gas is free of radon or ^{210}Pb because they freeze out in the bath.

Detector handling is an aspect of these precautions. Persons handling detectors wear gloves and full cleanroom gowning. Mechanical parts near the detectors are etched to remove surface contamination whenever possible. Recently, detectors have been stored in LN-boiloff-purged, radon-free environments when not being handled or tested.

3.6 The Icebox

The Icebox consists of six nested copper cans and stems corresponding to six temperatures stages: room temperature (RT), liquid nitrogen (LN), liquid helium (LHe or 4K), still (ST), cold plate (CP), and mixing chamber (MC). A mechanical drawing is shown in Figure 3.6. The first three stages nominally operate at 300 K, 77 K, and 4 K, respectively. The last three connect to the corresponding stages of the Oxford 400S dilution refrigerator's dilution unit; they nominally operate at 600 mK, 50 mK, and 5 mK, respectively. The six cans are nested concentrically. Each can has a mating flange on its side at the can bottom, as does each layer of the dilution refrigerator. These flanges are joined together by nested tubes, collectively know as the "cold-stem" or "C-stem." Signals are carried out to room temperature through an auxiliary set of tubes known as the "electronics-stem" or "E-stem" on custom stripline cables. I discuss here the design innovations; radiopurity considerations were mentioned above. The definitive reference on the Icebox is Peter Barnes's dissertation [92].

3.6.1 Cryogenic-Engineering Considerations

The primary challenge of the Icebox is the cryogenic engineering involved.

First, heat must be transferred from the Icebox to the refrigerator across a distance of approximately 1 meter. This is accomplished by extending every layer of the refrigerator with a copper tube. Right-angle joints are made at the bottom of the refrigerator and near the bottoms of the Icebox cans. The innermost layer of the refrigerator is extended by a solid copper rod attached to the refrigerator's vertical cold finger; it is connected to the innermost Icebox can by a flexible copper braid. The most challenging layers from the heat-load perspective are the LN and LHe layers [92]. The LN layer blocks blackbody radiation from the RT outer layer. The emissivity of the LN layer is required to be very low, $\epsilon = 0.01$ at the cans. The LN can is gold-plated and 20-layer superinsulation installed between the LN and RT cans. The striplines are also heat sunk to this layer on their way to the LHe layer, giving an additional 1.8 W heat load. The total LN layer power load is 36 W. The LHe layer's intrinsic heat load, from radiation and conduction from warmer layers, is 420 mW, also assuming gold-plating. Each stripline adds 12 mW and each BLIP FET card (4 FETs) adds 7 mW [104]. During Run 19, operating with 20 striplines and 24 FETs, the total LHe layer heat load was ≈ 800 mW.

Second, thermal contractions significantly perturb the Icebox. Because the dilution re-

frigerator's layers are fixed at the top of the refrigerator, the points of contact of the dilution refrigerator to the C-stem move upward when the refrigerator cools. The clearance between stems must be sufficient to accomodate this movement. Because thermal contractions are power laws in absolute temperature, the primary relative thermal contraction is between the room temperature and LN layers. The differential contraction between the LN and LHe layers is also significant. Because the LHe layer is already at 4.2 K, the contractions of the inner layers relative to the LHe layer are insignificant. In addition to the vertical contraction, the stems themselves contract as they cool, pulling the cans toward the refrigerator.

A porch swing design is used to accommodate these effects. Each can is hung at two points using Kevlar loops from the next warmer layer. These two points are separated by 90 degrees from the C-stem so that, as the stems contract, the cans tilt like swings. On the refrigerator side, the bottom vertical sections of the refrigerator (the "tails") are split with vertical bellows sections. The LN layer has a simple hinge. (Hinges are not used everywhere because they have very poor thermal conduction; a cooling loop carries LN down below the hinge to cool the LN tail). The rigid connections to the stems are made below the bellows sections: the refrigerator above the bellows contracts vertically upward while the sections below the bellows tilt due to the weight of the cans pulling through the stems. The stems themselves are rigid. In the E-stem, which connects the LN and LHe cans to the fixed room-temperature electrical feedthrough, bellows sections are also used. These flex as the cans are pulled toward the refrigerator during cooling. They also have low thermal conduction to the room-temperature anchor points.

A number of other cryogenic-engineering innovations were developed to ensure good conductivity across flexible joints, to isolate the cans microphonically from the refrigerator, to make cryogenic vacuum seals using only copper mating surfaces and gaskets, etc. A particularly significant effort went into the design and assembly of a temperature-monitoring system for the Icebox. This was primarily necessary for diagnosis of problems during initial testing of the Icebox. It has continued to prove useful in understanding variations in cooldown times. This system is also described in great detail in Peter Barnes's dissertation [92].

3.6.2 Leak Repair

Dilution refrigerators leak — this is an inevitable occurrence. The Icebox was designed to permit disconnection of the refrigerator from the Icebox for repair without disassembly of the Icebox, but this procedure had never been tried during assembly and testing. When the dilution unit began to leak during a cooldown in November, 1997 (the leak was found on my 25th birthday), a procedure for an alternate disassembly was developed, a process in which I (unfortunately) was personally involved.

The following procedure is followed for normal assembly:

1. The refrigerator is assembled, hanging from its frame. The flanges to be connected to the C-stems are accessible at the side port.

2. The C-stems are nested and connected to the refrigerator tails, MC layer first.

3. The cans are situated and attached to the C-stems, warmest layer first. That is, the RT can is set on its stand and joined to its C-stem, then the LN can is placed inside the RT can and

attached to its C-stem, which penetrates through the RT can, and so on. The screws for the flanges that mate the cans to the C-stems are inserted from the inside of the cans.

4. The E-stems are nested and connected to the cans; the E-stem only has RT, LN, and LHe layers.

5. The shield and veto are assembled around the completed Icebox.

Obviously, in order to open the refrigerator proper to repair a dilution unit leak, one could reverse the whole procedure at the cost of months of work. The procedure developed for separation of the refrigerator from the Icebox takes much less time. It is as follows:

1. The tail covers are removed to allow access to the flanges attaching the tails to the refrigerator. (Note: these are not the flanges attaching the stems to the tails. These flanges are *above* the bellows sections.)

2. A custom tripod is set up to rigidly support the RT tail. A staircase support fixture is attached to the RT tail. The inner tails are disconnected from the refrigerator and allowed to rest on and are attached to the staircase support. The weight of the tails and the C-stem is carried by the supported RT tail through the staircase support.

3. The refrigerator is lifted vertically using its suspension plate with a custom forklift. The refrigerator must be lifted about 2 feet to allow the MC vertical cold finger to clear the tails.

4. The refrigerator is translated using the forklift. Full access to the refrigerator is possible now, so the leak can be repaired in this position. The tails remain in place, supported by the tripod.

5. The above procedure is reversed to return the refrigerator to its normal position. While vertically dropping the refrigerator back in place, guide rods are inserted through the screw holes in the tail flanges that mate to the refrigerator to ensure alignment (the tail flanges are not well constrained because of the flexible bellows). This step involves remaking the LHe-bath seal, consisting of two indium O-ring seals that are made simultaneously.

The critical component of the procedure is that the refrigerator and tails must be returned very precisely to their initial positions ($\lesssim$ 1 mm offset) to avoid introducing new stresses in the system. The procedure was successful and the refrigerator was recooled in mid-January, 1998. In the future, such a repair can probably be done in one month now that the procedure and necessary hardware are available.

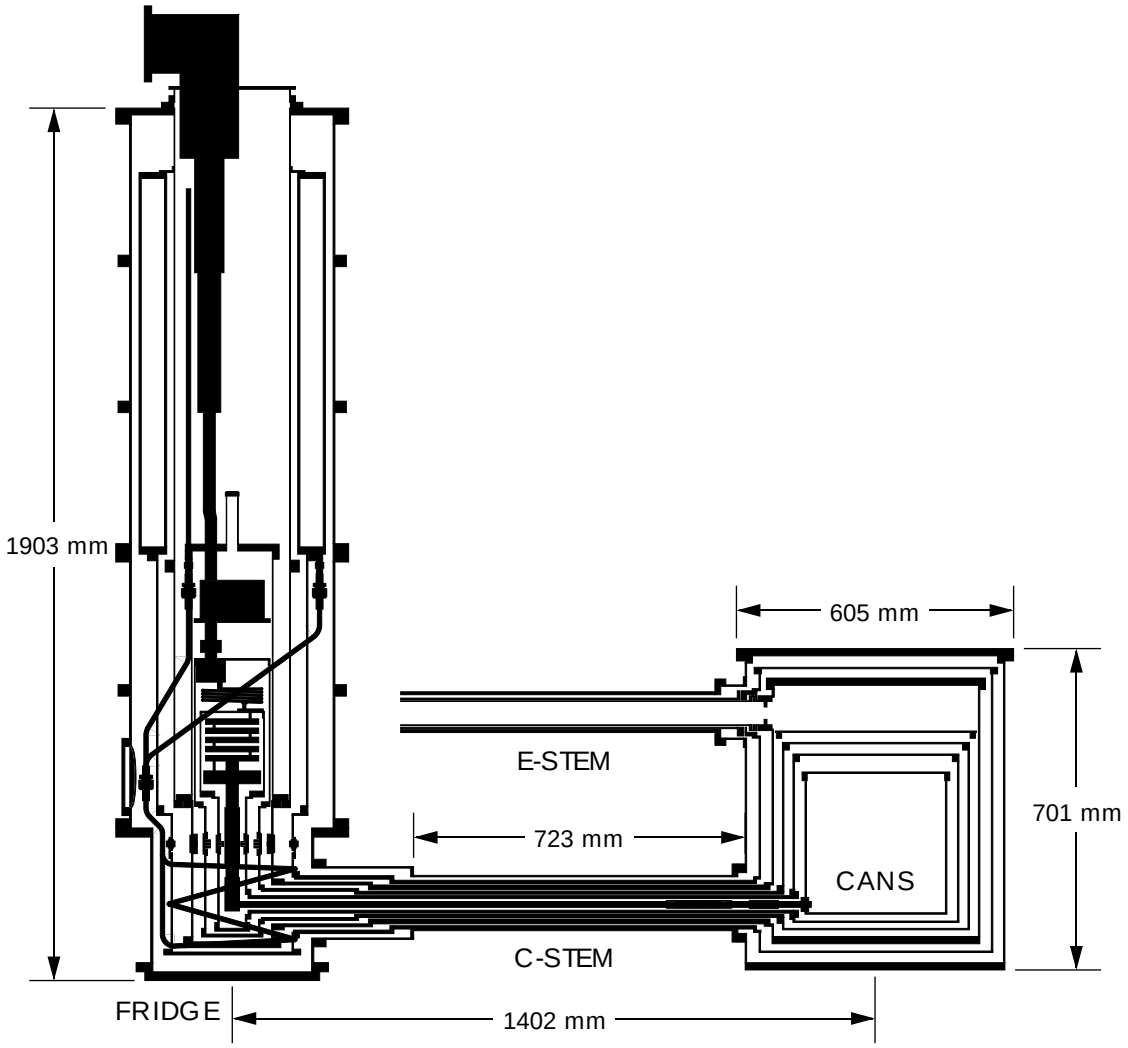

Figure 3.11: Side view of the Icebox and Oxford 400S dilution refrigerator. The "tails" are the narrowed sections at the bottom of the refrigerator that attach to the C-stem.

Chapter 4

NTD Ge-Based, Thermal-Phonon-Mediated BLIP Detectors

The data presented in this dissertation were taken primarily with Berkeley Large Ionization- and Phonon-mediated (BLIP) detectors. As mentioned in Chapter 2, the WIMP sensitivity of this detector technology arises from its ability to discriminate nuclear recoils from electron recoils based on measurement of both the phonons and electron-hole pairs produced by a recoiling particle. These data were acquired during CDMS Run 19 using a 4-detector array of BLIPs (BLIPs 3, 4, 5, and 6) in a close-packed geometry. The mechanical details will be discussed in Chapter 5. For now, it is sufficient to know that the four detectors are stacked face-to-face, with no intervening material, with BLIP3 on top and BLIP6 on the bottom.

In this chapter, I discuss the operating principles and observed capabilities of BLIP detectors. I review background material that has been presented in more detail elsewhere and concentrate instead on topics that can be uniquely addressed by these data. I present the principles of the phonon and ionization measurements and their observed energy resolutions. I discuss the dead-layer problem, the phenomenon of reduced ionization collection for particle interactions near the detector surfaces, and the recent advances that have mitigated this problem. I explain how the detectors used in this run were fabricated, and finally I present analysis of the *in situ* nuclear-recoil acceptance and photon and electron misidentification.

The BLIP detector technology was developed in the Sadoulet group at Berkeley and is the subject of dissertations by Ning Wang, Tom Shutt, Alan Cummings, and Walter Stockwell [105, 106, 107, 108]. The BLIP prototype, E5, and the first two BLIPs were operated at SUF in Runs 13 through 16 and 18, with detector operations and analysis led by Rick Gaitskell. The performance of the detectors in these runs is described in various internal notes [109, 110, 111, 112] and in Andrew Sonnenschein's dissertation [8]. More recently, Jochen Hellmig, Bernard Sadoulet, and Tom Shutt developed a new ionization-electrode technology that solves, to a large extent, the dead-layer problem.

In this and later chapters, I will allude to CDMS Z-sensitive Ionization- and Phonon-mediated detectors (ZIPs), which have been developed more recently by Blas Cabrera's group at Stanford University. In ZIP detectors, the phonon signal is measured by sensing of high-energy

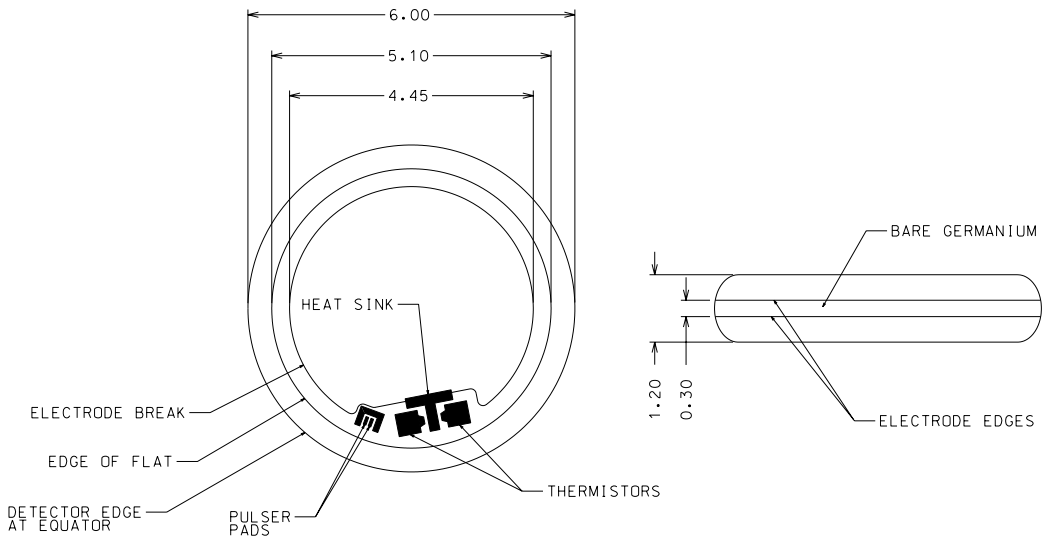

Figure 4.1: A BLIP detector. Dimensions are in centimeters. The ionization-electrode breaks are indicated. The thermistors are not shown on the side view; they are ≈ 0.26-cm high.

(few-K) athermal phonons. The athermal-phonon pulse shape's rise time is sensitive to event transverse position and depth. The latter allows rejection of dead-layer events. The performance of the first full-size ZIP detector is described in Roland Clarke's dissertation [102] (where the detectors are referred to by their original name, Fast Large Ionization- and Phonon-mediated detectors, or FLIPs). In addition, the ionization-electrode technology developed to solve the dead-layer problem is applicable to ZIPs and has been used on the detectors being currently operated in CDMS Run 20. ZIPs are the baseline detector technology for CDMS II, the next phase of CDMS at a deep site and with an order of magnitude increase in detector mass.

A note on plots: BLIPs 3, 4, 5, and 6 have been assigned "detector numbers" 103, 104, 105, and 106. Plots of individual detector data are labeled by detector number.

4.1 Detector Description and Operation

In this section, I discuss the physics, readout schemes, pulse shapes, and noises of the phonon and ionization measurements. A BLIP detector consists of a cylindrical crystal of germanium, 165 g in mass, 6 cm in diameter, and 1.2-cm thick, with two neutron-transmutation-doped (NTD) germanium thermistors (approximately $3.1 \times 3.1 \times 2.6$ mm^3) attached to perform the phonon measurement and electrodes deposited on the top and bottom faces to provide the ionization drift field and contact to the ionization bias and amplifier. See Figure 4.1.

4.1.1 NTD-Ge-Based Phonon Measurement

There are two primary motivations for making use of a phonon-mediated technology such as the BLIP detector. First, large masses can be achieved while maintaining energy resolution ~ 1 keV. Second, the phonon measurement gives the full energy of the interaction, regardless of the recoil type. These two facts are important for direct detection of WIMPs — a recoil-energy threshold well below the average energy of a typical WIMP-induced nuclear-recoil spectrum can be achieved.

A thermal detector like the BLIP technology relies on the fact that the heat capacity of an insulating crystal drops as T^3 at cryogenic temperatures. Thus, very small depositions can, in principle, be sensed because they cause large temperature rises. In general, the heat capacity of germanium at $T \ll \Theta_D = 374$ K is given by the Debye result

$$\begin{aligned} C &= \frac{12\pi^4}{5} N k_B \left(\frac{T}{\Theta_D}\right)^3 \\ &= 3.2 \times 10^{12} \left(\frac{M}{\mathrm{g}}\right) \left(\frac{T}{\mathrm{K}}\right)^3 \ \mathrm{eV/K} \end{aligned} \tag{4.1}$$

For a 165-g BLIP operated at 20 mK, a 10-keV deposition gives a temperature rise of 2.4 μK, which is measurable. While the above is useful as motivation, it should be noted that the BLIP detector is more complicated: the thermistor and crystal heat capacities are comparable, and it is only the temperature change of the thermistor that can be directly sensed. The precise behavior is discussed in more detail below.

The use of NTD-Ge thermistors for particle detection is discussed in detail in Ning Wang's dissertation [105]. As the name indicates, NTD-Ge thermistors respond to a change in temperature by a change in resistance. A single thermistor may be used for direct particle detection. In BLIP detectors, the thermistor is instead used to measure the temperature rise of the attached large Ge target crystal.

Properties and Physics of NTD Germanium

Neutron-transmutation-doped (NTD) germanium is produced by exposing germanium to neutrons from a nuclear reactor. Nuclear reactions convert Ge to Ga, making the material strongly p-type. Neutrons have a long enough mean free path that the doping is uniform, even over samples of few-mm dimension. The impurity density is high enough that electrons may hop between impurity sites, allowing the material to conduct. The hopping is thermally-assisted, so the conduction is highly dependent on the electron temperature, following the behavior [105]

$$R = R_0 \exp\left(\sqrt{\frac{\Delta}{T}}\right) \tag{4.2}$$

For the thermistors used with BLIPs, $\Delta \approx 7$ K. A figure of merit is the logarithmic derivative, $\alpha = (T/R)dR/dT = -(1/2)\sqrt{\Delta/T} = -9.4$ at 20 mK, which relates a fractional temperature change to a fractional resistance change by $(\delta R/R) = \alpha(\delta T/T)$.

Electron-Phonon Decoupling and Power Flow in a NTD-Ge System

Central to the physics of the NTD-Ge thermal system at low temperature is electron-phonon decoupling. Electron-phonon interaction rates become low enough that the time needed for the electron and phonon systems of the thermistor to share energy deposited in one system is significant compared to the two systems' internal thermalization times. If a significant DC power is deposited in the electron system and the thermistor heat-sunk via its phonons, a significant steady-state temperature difference arises. Ning Wang describes this phenomenon in great detail in her dissertation [105]; the final result for NTD germanium is

$$P_{ep} = g_{ep}(T_e^6 - T_p^6) \tag{4.3}$$

The exponent arises from one power of the created or destroyed phonon's energy, one power from the fundamental matrix element for phonon creation and destruction, and the remainder from phase space/density of states. The last factor differs for metals, yielding T^5 instead of T^6 (electron-phonon decoupling in metals is described in Fred Wellstood's dissertation [113]). Electron-phonon decoupling and the heat capacities of the electron and phonon systems determine the timescales for power flow in BLIP detectors.

An important element of electron-phonon decoupling is the dependence on thermistor volume. The coefficient g_{ep} has the form ΣV, where Σ is a purely material-dependent parameter and V is the thermistor volume. The electron-phonon power flow can be calculated by assuming a deformation-potential approximation, which is valid for long phonon wavelengths: in this model, an electron's state changes because phonons distort the unit cell, modifying the Bloch states enough to yield a quantum transition. In this approximation, the fundamental electron-phonon matrix element has the form

$$|M|^2 = \frac{\hbar}{2\rho c_s}\left(\frac{2E_f}{3}\right)^2 \frac{q}{V} \tag{4.4}$$

where ρ is the crystal mass density, $\vec{q}$ is the phonon wavevector, c_s is the phonon sound speed (the slope of the dispersion relation at $\vec{q} = 0$), E_f is the Fermi energy, and V is the crystal volume. Calculation of the overall rate of energy flow between the electrons and phonons requires a phase-space integral over the phonon wavevector, bringing in one power of V. Integration over all electrons brings in a density-of-states factor $D(E_{\vec{k}})$ where $\vec{k}$ is the electron wavevector. Since it is $D(E_{\vec{k}})/V$ that is volume independent, an additional power of V is introduced. The overall result is thus proportional to V, with the coefficient Σ depending on the material parameters and the electron and phonon dispersion relations but not on the volume. The scaling of the coupling with volume is important because it suggests the use of large thermistors to overcome electron-phonon decoupling.

Thermal Model of BLIP Detectors

A simplified thermal model for BLIP detectors, including only one thermistor, is shown in Figure 4.2. One block of the system is the detector phonon system, combining the phonon systems of the crystal and of the thermistors. The thermistors are eutectically bonded to the crystal. Eutectic bonding provides a reproducible, reliable interface that is transparent to low-energy phonons [105, 114]. The other block is the thermistor's electron system. The electron system is coupled to the detector phonon system via the electron-phonon coupling. The detector is coupled to the refrigerator via a gold heat-sink pad deposited on the detector that is connected to

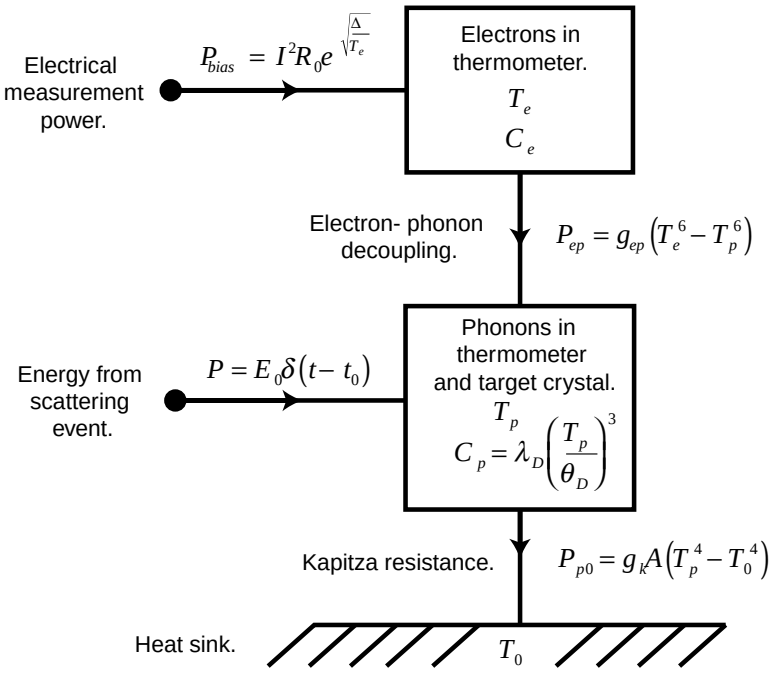

Figure 4.2: BLIP thermal model. The top box is the electrons and the bottom the crystal/thermistor phonons. The heat sink is shown at the bottom. The power inputs, heat capacities, and thermal couplings are described in the text. Figure taken from [8].

the refrigerator via a gold wirebond. The dominant impedance in the latter coupling is the acoustic-mismatch resistance between the crystal and the heat-sink pad, which is proportional to the pad area. This is usually (incorrectly) referred to as Kapitza resistance. The thermal impedances inside the heat-sink pad, along the wirebond, and to the refrigerator are negligible in comparison because these systems are metallic. The refrigerator operates at about 20 mK. Bias-power dissipation causes the thermistor and, to a lesser extent, the crystal, to heat above the refrigerator temperature. The thermistor is usually a few mK warmer.

Schematically, the power flows are as follows. A thermistor-bias current I_b produces a measurable voltage $I_b R$. This dissipates power $I_b^2 R$ in the thermistor. (A current bias is needed to prevent thermal runaway because $dR/dT < 0$.) This power flows to the heat sink via the phonon system. An event in the crystal gives a δ-function energy deposition in the *phonon* system. The phonons heat up, warming the electrons via the electron-phonon coupling and yielding a measurable change in resistance. The energy flows out of the system via the connection to the heat sink. The couplings are chosen so the electron system senses the phonon-system temperature rise before the energy can leave the detector.

In general, the power-flow equations for the electron and phonon systems are

$$\frac{dU_e}{dt} = I_b^2 R(T_e) - g_{ep}(T_e^6 - T_p^6) \quad (4.5)$$

$$\frac{dU_p}{dt} = g_{ep}(T_e^6 - T_p^6) - g_K(T_p^4 - T_s^4) \quad (4.6)$$

where U_e and U_p are the energy contents of the electron and phonon systems, I_b is the thermistor-bias current, $R(T_e)$ is the temperature-dependent thermistor resistance, T_s is the heat-sink/refrigerator temperature, g_{ep} is the electron-phonon coupling, and g_K is the Kapitza coupling. The parameters I_b, T_s, g_{ep}, and g_K are constant in this analysis. The quiescent state is given by setting the left side of both equations to 0, yielding T_{e0} and T_{p0}. These are nontrivial functions of the other parameters because of the power-law dependences; it is not particularly instructive to calculate them explicitly. Clearly, though, $T_{e0} > T_{p0}$ because of the Joule power dissipated in the thermistor.

To determine the response to a particle interaction, the above equations are Taylor-expand around T_{e0} and T_{p0}, giving

$$C_e \frac{d\delta T_e}{dt} = -\left(G_{etf} + G_{ep}\right)\delta T_e + G_{pe}\delta T_p + P_{in,e}(t) \tag{4.7}$$

$$C_p \frac{d\delta T_p}{dt} = G_{ep}\delta T_e - \left(G_{pe} + G_K\right)\delta T_p + P_{in,p}(t) \tag{4.8}$$

where $P_{in,e}(t)$ and $P_{in,p}(t)$ are the power inputs to the electron and phonon systems by the particle interaction. The heat capacities of the phonon and electron systems are given by C_p and C_e. They are taken to be constant here; their temperature dependences contribute terms of second order in $d\delta T_e/dt$ and $d\delta T_p/dt$ and so can be neglected. In general, they have significant temperature dependences that must be taken into account when optimizing the detector parameters. For crystal interactions, $P_{in,e}(t) = 0$ and $P_{in,p}(t) = E\delta(t)$ — all the energy is input to the phonon system. The bias power and electron-phonon and Kapitza couplings give rise to coefficients called conductances:

$$\begin{aligned} G_{etf} &= -I_{bias}^2 \left.\frac{dR}{dT_e}\right|_{T_{e0}} \\ G_{ep} &= 6 g_{ep} T_{e0}^5 \\ G_{pe} &= 6 g_{ep} T_{p0}^5 \\ G_K &= 4 g_k T_{p0}^3 \end{aligned} \tag{4.9}$$

G_{etf} is defined with a negative sign to make it positive. The point made earlier about the dependences on volume is apparent now: a crystal interaction causes an increase in phonon temperature δT_p via Equation 4.8, which then is coupled into the electron system by Equation 4.7. The rate of change of the temperature of the electrons is given by $\delta T_p/(C_e/G_{pe})$, which is independent of thermistor volume (in the limit that $C_e \ll C_p$; this does not hold exactly).

Because the system is somewhat complicated, solving the equations explicitly is not particularly illuminating; the full solutions are given in [107]. Intuitively, there are a number of obvious time constants given by dimensional analysis and the generic structure of the equations:

$$\begin{aligned} \tau_{etf} &= C_e/G_{etf} \\ \tau_{eep} &= C_e/G_{ep} \\ \tau_{epe} &= C_e/G_{pe} \\ \tau_{pep} &= C_p/G_{ep} \\ \tau_{ppe} &= C_p/G_{pe} \\ \tau_K &= C_p/G_K \end{aligned} \tag{4.10}$$

These correspond to the time constants, respectively, for: the electrons to cool by reduction in Joule heating ("electrothermal feedback"); the electrons to cool by power flow into the phonons; the electrons to heat by power flow from the phonons; the phonons to heat by power flow from the electrons; the phonons to cool by power flow into the electrons; and, the phonons to cool by power flow into the heat sink. Note that τ_{eep} and τ_{epe} are independent of thermistor volume. C_p is dominated by the crystal and so τ_{pep} and τ_{ppe} are crystal-volume dependent. Though the few mK temperature difference between the electrons and phonons may not appear important at face value, the large power-law exponents in the G's amplify this difference: if $T_{e0} = 25$ mK and $T_{p0} = 20$ mK, the error incurred by taking them equal is $1.25^5 \approx 3$.

A comment should be made regarding the above simplistic model and the role of the eutectic bond. The use of a thermal model is an approximation, though apparently a good one. Because the eutectic bond is not transparent to the high-energy phonons that emerge from a particle interaction on sub-μs timescales, these phonons both experience enough reflections at the crystal surfaces and have enough time to down-convert (by anharmonic decay) that they lose memory of their original distribution. By the time the phonons can penetrate into the thermistor, they have become thermal enough that the above model predicts the detector behavior reasonably well [106, 107, 108]. Though it is not transparent to high-energy phonons, a eutectic bond is used because it yields more reproducible detector behavior than gluing, probably for two reasons. First, the eutectic interface, while not transparent to high-energy phonons, is reproducibly transparent to low-energy phonons, yielding a reproducible coupling. Second, gluing appears to stress thermistors more than eutectic bonding, which reduces the reproducibility of thermistor properties.

Two thermistors are used to to provide rejection of interactions in the thermistors. Use of two thermistors also decreases the phonon readout noise by $1/\sqrt{2}$. For crystal interactions and assuming the two thermistors are identical, the temperature-evolution solutions have the same form as a one-thermistor system: the two thermistors can be treated thermally and electrically as a single thermistor. For interactions in a thermistor, the symmetry is broken and the results become more complicated. Also, the current bias is imperfect and the thermistors possess capacitances, both of which modify the electrothermal feedback slightly.

Phonon Readout

In principle, the thermistor signal is a negative-going voltage pulse given by the product of the fixed bias current and the resistance decrease arising from an energy deposition. A low-noise voltage amplifier can be used to measure this signal. In practice, the time constants are slow enough that a significant component of the signal lies at low frequencies. The rise and fall times of the BLIP phonon signals are ~ 5 ms and ~ 50 ms (see Figure 4.6), corresponding to poles in the pulse frequency spectrum at ~ 3 Hz and ~ 30 Hz. Below 500 Hz, $1/f$ noise (in the JFET, thermistor, or electrical connections) and 60 Hz from power supplies become significant; see Figure 4.3. It is therefore advantageous to use a lockin technique. The DC current bias is replaced by a 1-kHz sine-wave bias; the noise environment around 1 kHz is very clean, as is seen in Figure 4.4. To keep R constant, the bias frequency must be high enough that its period is small compared to the system's thermal time constants. Though the power dissipation oscillates at 1 kHz, the electron and phonon temperatures cannot respond this quickly, so the system effectively sees a constant power input $\langle I_b^2 \rangle R$ and maintains a constant resistance R. The bias frequency is limited from above by $(2\pi RC)^{-1}$, where R and C are the thermistor resistance and capacitance; the bias is shunted

through by the capacitance at higher frequencies. The typical thermistor operating resistance is 1–2 MΩ; the capacitance is dominated by the FET-gate capacitance, $C_{FET} = 40$–60 pF [115], so the cutoff frequency can actually be uncomfortably close to 1 kHz. "Mixing" occurs in the thermistor by multiplication of the slow variation in the resistance by the 1-kHz current bias:

$$v(t) = i(t) \times r(t) = \mathcal{FT}^{-1}\left(\mathcal{FT}\left[i(t)\right] * \mathcal{FT}\left[r(t)\right]\right) \tag{4.11}$$

where $*$ indicates convolution and $\mathcal{FT}$ and $\mathcal{FT}^{-1}$ indicate forward and inverse Fourier transforms. The signal seen by the amplifier thus appears around 1 kHz, away from the low-frequency noise. The signal is square-wave demodulated using the zero crossings of the signal itself to determine the switching times. Prior to demodulation, the signal is low-pass filtered with a 2-pole 1.85-kHz-cutoff filter to prevent square-wave demodulation from mixing higher-frequency noise into the band of interest. The demodulated signal is low-pass filtered with a 8-pole 450-Hz-cutoff Butterworth filter to remove the frequency components near 2 kHz, 4 kHz, etc. Further details of the implementation of the lockin technique are discussed in Chapter 5. The circuit is depicted schematically in Figure 4.5.

The demodulation scheme requires discussion. The zero crossings of the carrier, rather than of the phonon bias, are used for demodulation because the thermistor signal is phase-shifted by capacitances in the electrical system. For example, the current bias is provided by voltage biasing the series combination of a large bias resistor ($R_b = 40\,\mathrm{M\Omega}$) and the thermistor, so capacitances of both elements enter. Use of the signal carrier for demodulation is possible because the carrier is large compared to the signal — many mV versus the signal pulse height of $\sim$ 160 nV/keV. Square-wave demodulation is used for simplicity: it can be done by simply switching an inverter in and out of the signal path. In the limit that the noise of the modulated signal, in the band of interest, dominates over noise introduced by the demodulator or the following filters, the signal-to-noise is independent of whether square-wave or sine-wave demodulation is used, since the noise and signal are shifted in frequency space in the same way. This limit holds in practice because the modulated signal is amplified by 500 before demodulation. In this implementation, both sine-wave and square-wave schemes yield no degradation relative to combined sine/cosine demodulation because the demodulation is automatically synchronized with the incoming signal — the component of the signal 90 degrees out of phase with the demodulation vanishes.

The phonon pulse shapes are shown in Figure 4.6 and the noise spectra with the lockin in use are shown in Figure 4.7. The significant improvement in low-frequency noise over Figure 4.3 is apparent. In these spectra, the noise is dominated by thermistor Johnson noise. The thermistor-bias parameters and operating resistances are given in Table 4.1. Expected and observed energy resolutions are discussed in Section 4.4.

One minor disadvantage of the lockin technique is that the large thermistor bias ($\sim$ 10 mV across the thermistor) cross-talks to the ionization amplifiers. This can be mitigated by a factor of 10 by biasing the two thermistors out of phase by 180°. It is not clear which of the many couplings between the phonon and ionization circuits provides the path for this cross-talk. Two possibilities can be ruled out. Refer to Figure 4.10 for a schematic of the combined electronics circuit. One possibility is that the large voltage appearing at the top of the thermistor on the detector is capacitively coupled to the ionization electrode via the detector. This is not plausible because the detector ground plane lies between the two, making the capacitance between them negligible. In addition, the thermistor is not a stiff voltage source because of its high impedance. The second

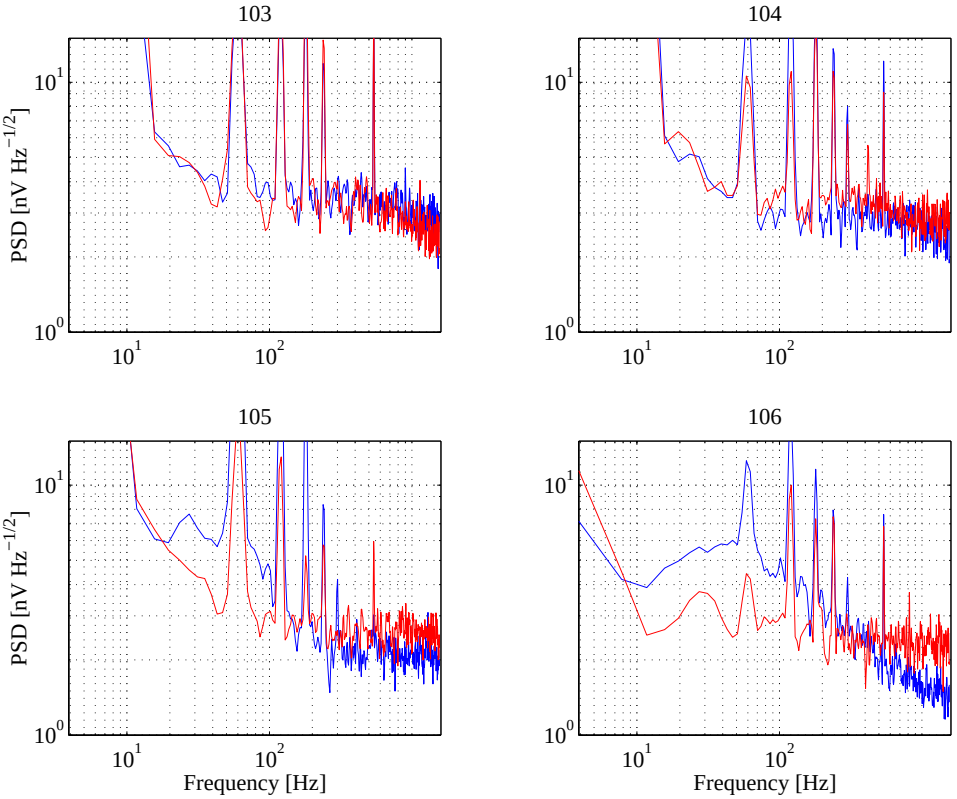

Figure 4.3: Phonon-channel noise spectra without lockin, logarithmic scales. Dark: phonon sensor 1. Light: phonon sensor 2. The continuum noise, above about 100 Hz, is dominated by thermistor Johnson noise — the FET contributes ~ 1 nV/$\sqrt{\text{Hz}}$. The spectral lines are 60 Hz and harmonics. The significant increase in "smooth" noise and in 60 Hz and harmonics at low frequencies motivates the use of a lockin technique: recall that the fall and rise times of the phonon pulses correspond to ~ 3 Hz and ~ 30 Hz, so the phonon signal is all below 30 Hz. See Figure 4.4 for plots with linear scales. Figure taken from [116].

Parameter	BLIP3		BLIP4		BLIP5		BLIP6	
	P1	P2	P1	P2	P1	P2	P1	P2
bias current (nA)	6.9	6.9	6.7	6.1	7.7	7.6	7.6	6.5
$V_{carrier}$ (mV)	15	15	15	16	13	14	15	15
R_{op} (MΩ)	2.2	2.2	2.2	2.6	1.7	1.8	2.0	2.3

Table 4.1: Thermistor operating parameters. The carrier voltages are only shown to 1-mV precision because they vary by ~ 1 mV over the run. The bias currents are calculated assuming 40 MΩ bias resistors. The thermistor operating resistances are calculated assuming no capacitive shunting. The FET-gate capacitance is likely between 40 and 60 pF, corresponding to impedances of 4 MΩ and 2.7 MΩ at 1 kHz, so there is certainly some capacitive shunting.

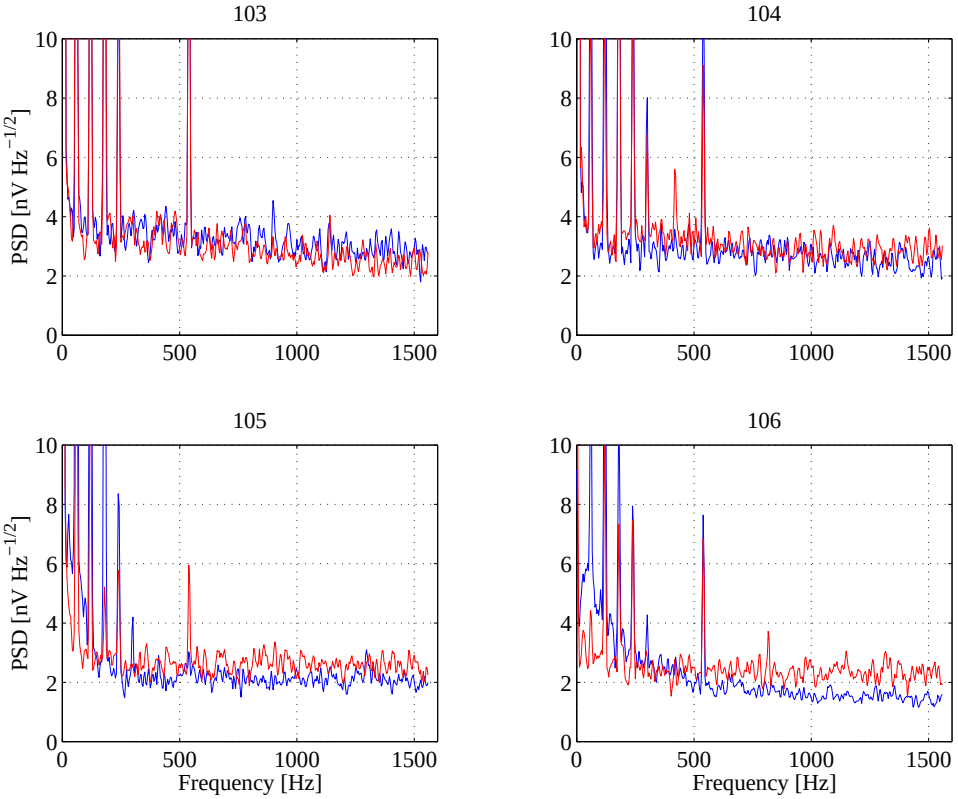

Figure 4.4: Phonon-channel noise spectra without lockin, linear scales. Dark: phonon sensor 1. Light: phonon sensor 2. These plots indicate the significant difference in the noise environment below 100 Hz and at 1 kHz, the phonon-bias frequency. See Figure 4.7 for noise spectra with the lockin in use.

possibility is that, because of wiring resistance, the detector ground plane itself oscillates at 1 kHz relative to the amplifier grounds on the electronics boards. This resistance could be as high as a few Ω; the phonon-bias current is ∼ 7 nA per thermistor, so a voltage between 10 and 50 nV could appear on the detector ground plane. This drives a current through the detector capacitance directly into the input of the ionization amplifier and thus produces a signal. The expected voltage at the output of the ionization amplifier is therefore

$$\begin{aligned} V_{out} &= I_b R_{lead} \left(\frac{1}{\omega C_d}\right)^{-1} R_f \\ &\sim (7\text{ nA})(5\ \Omega)(8\text{ M}\Omega)^{-1}(40\text{ M}\Omega) \\ &\sim 175\text{ nV} \end{aligned} \tag{4.12}$$

The cross-talk observed is about 1000 times larger, so clearly this explanation fails also. Some evidence suggests that the cross-talk occurs primarily via capacitive coupling between the thermistor-bias lines and the ionization-amplifier gate lines on the Side Coax Assembly Board (defined in

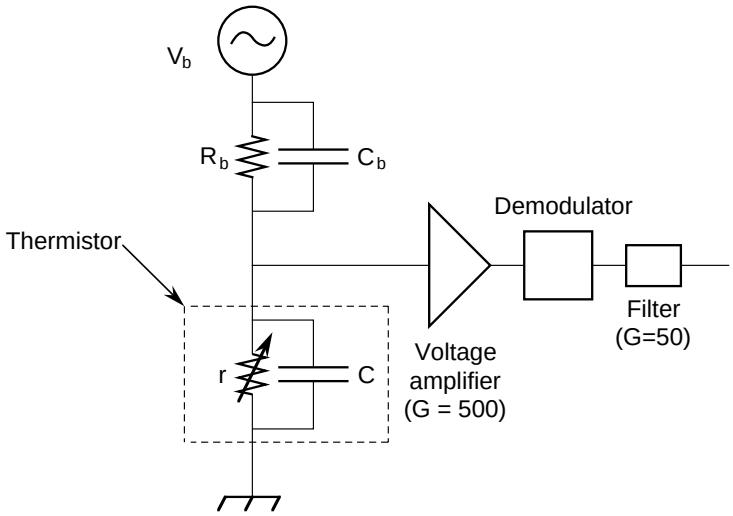

Figure 4.5: Thermistor-voltage-measurement circuit. The 1-kHz bias is indicated at top. The bias resistor, $R_b = 40\,\mathrm{M}\Omega$, converts the voltage bias to a current bias. The amplifier, demodulator, and filter stages are indicated. The 2-pole 1.85-kHz-cutoff filter between the amplifier and the demodulator is not shown. Figure taken from [91].

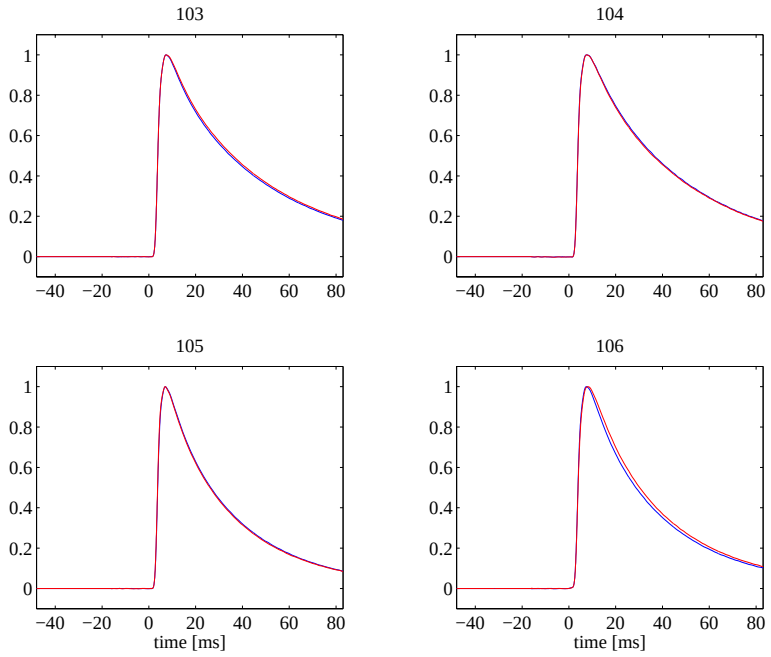

Figure 4.6: Run 19 phonon-channel pulse shapes. Dark: phonon sensor 1. Light: phonon sensor 2. The typical energy conversion (peak height) is 160 nV/keV.

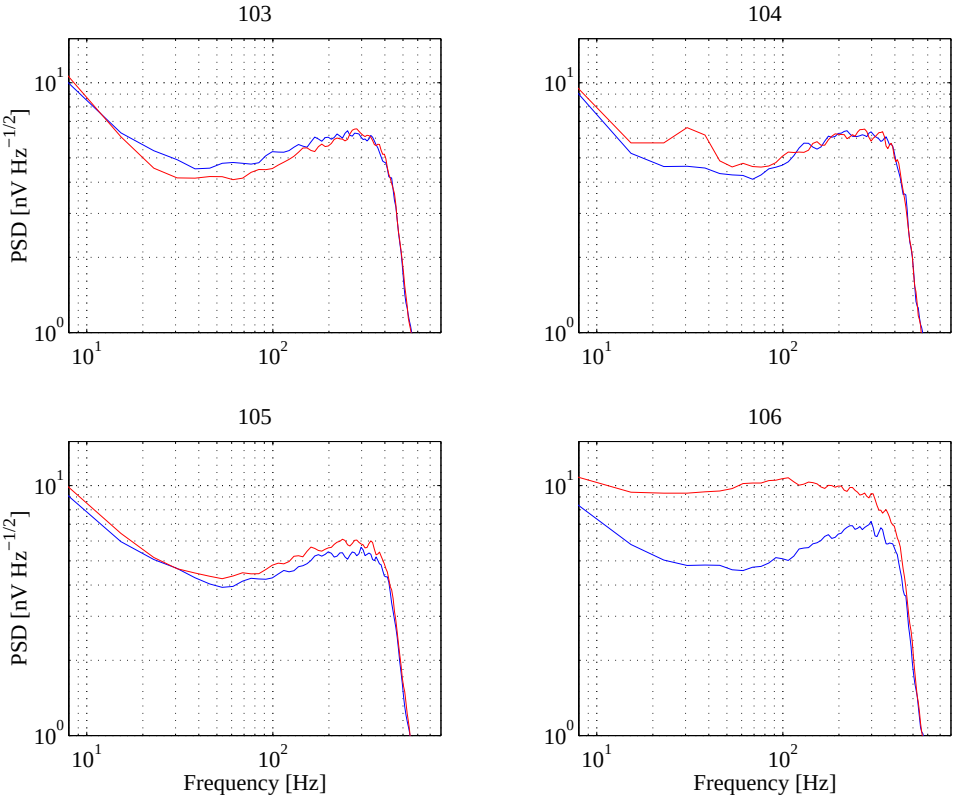

Figure 4.7: Phonon-channel noise spectra (with lockin). Dark: phonon sensor 1. Light: phonon sensor 2. The sharp rolloff at 450 Hz is due to the 8-pole Butterworth filter. Clearly, the phonon traces are heavily oversampled. The anomalously high noise of BLIP6 P2 is not understood.

Chapter 5). With optimal filtering, the impact of this cross-talk can be almost completely nullified, so it has not been investigated further.

Thermistor-Contained Events, Tritium Contamination, and Baking

Particle interactions may occur in the thermistors themselves. The resulting pulses in the two thermistors are very different from crystal-interaction pulses. An interaction yields a recoiling electron or nucleus. Being charged, the recoiling particle interacts primarily with the thermistor electron system. Therefore, to first order, all the energy is deposited instantaneously in the electron system. The thermal evolution obeys the energy-transport equations given above. The thermistor's electron temperature immediately increases, giving a sharp drop in resistance. This pulse has a first decay time of $\tau_{eep} = C_e/G_{ep}$, corresponding to the flow of power out of the electrons into the phonons. The phonon system warms with a time constant $\tau_{pep} = C_p/G_{ep}$. Energy then flows into the second thermistor. The temperature of the second thermistor rises with a shape approximately given by the convolution of the phonon-temperature rise with the thermistor-temperature profile for crystal interactions; the latter gives the system's response to an impulsive energy deposition in

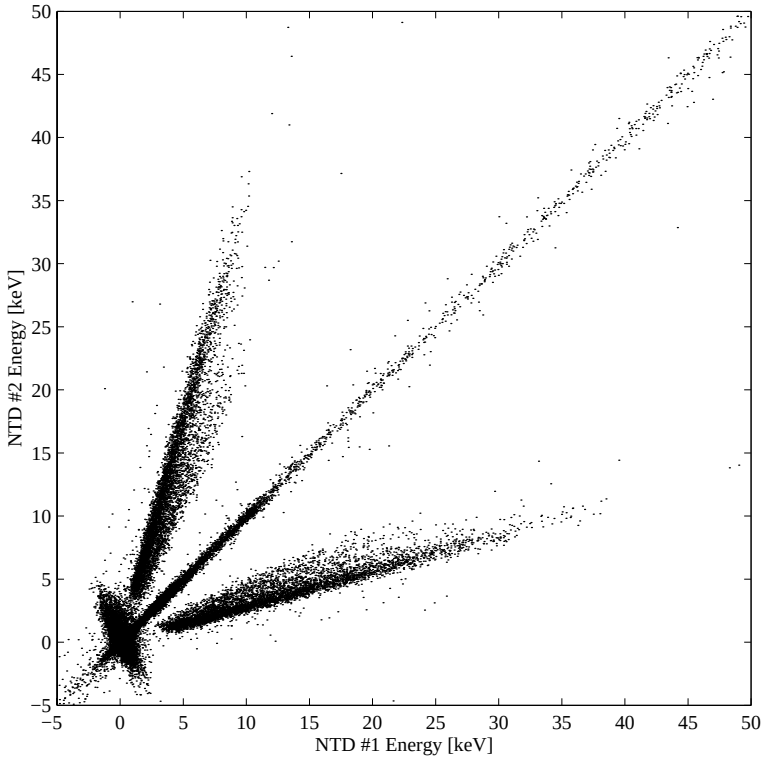

Figure 4.8: Scatter plot of individual thermistor energies for BLIP2, from Run 18. The two distributions off the diagonal are thermistor interactions in thermistor 1 (lower right) and thermistor 2 (upper left). Figure taken from [8].

the crystal. Since the first thermistor has already been heated, this is not exactly correct, but the conclusion is accurate: the rise time of the second thermistor's temperature is slower than that of a crystal-interaction pulse. Because the electrons and phonons are more coupled to each other than to the heat sink, the three systems equilibrate before energy flows to the heat sink, so the decay time of pulses is similar to the decay time of crystal-interaction events. Thus, the first thermistor shows a very sharp pulse with one fast and one normal decay time and the second thermistor shows an anomalously slowly rising pulse with a single, normal decay time.

Thermistor interactions are easily identifiable because, when fitted with a standard pulse template, they result in extremely different pulse heights in the two thermistors. Figure 4.8, shows a scatter plot comparing the pulse heights in the two sensors, also called a "dagger plot." The data displayed in this figure are taken from Run 18 BLIP2 data. The large number of thermistor events is incompatible with the photon background observed by the full detectors after scaling by volume. This class of events was first discovered during Run 14 in the E5 detector and confirmed during Runs 15 and 16 in BLIP1. The events are fit well by a tritium β-decay spectrum, with an endpoint of 18.6 keV, allowing for a scaling factor to correct the pulse height. The tritium is believed to have been created by nuclear reactions that occurred during neutron-transmutation doping of the thermistor material. The rate of such events is 570 d^{-1}, or 2.1 $\times$ 10^6 kg^{-1} d^{-1} [111].

It was determined that "baking" the thermistors decreases the tritium content by a factor of 100 [117]. The thermistors used for the Run 19 BLIPs were baked at 650 C for two hours with a continous-flow argon atmosphere. The baking process was observed to change somewhat the NTD properties, but not significantly. There may be evidence that the R_0 parameter increases, but Δ, the parameter that determines the responsivity, appears unchanged [118]. The pulse heights observed in the Run 19 BLIPs, $\sim$ 160 nV/keV, are no different from those observed in previous runs. At this point, it is no doubt unclear to the reader why the thermistors were baked, given that thermistor-contained events can be easily separated from detector interactions. I return to this issue, and the motivation for baking, in Chapter 8.

4.1.2 Ionization Measurement

The physics of ionization collection at 20 mK in BLIP detectors was investigated by Tom Shutt, Betty Young, and Bernard Sadoulet and detailed by Tom Shutt in his dissertation [106]. I discuss the microphysics of the appearance of an ionization signal, how ionization-loss mechanisms reduce this signal, the ionization readout, and the geometry of the electrodes.

Fundamentals of Ionization Collection

To discuss the ionization measurement, it is necessary to understand how ionization is generated and collected in low-temperature ionization-mediated detectors. Particle interactions result in a recoiling electron or nucleus that creates an ionization track along its path through the crystal — the recoiling charged particle pulls valence electrons off lattice sites, or, from the band-structure perspective, excites electrons from the valence band into the conduction band, creating electron-hole pairs. Though the bandgap energy of germanium is 0.743 eV, 3.0 eV of deposited energy is needed to create an electron-hole pair on average; most of the deposited energy is dissipated in other ways. For example, a 59.54-keV ^{241}Am photon produces $\sim$ 20000 charge pairs, not $59.54 \times 1000/0.743 \approx 80000$ charge pairs. Electrodes deposited on the two disk-shaped faces of the germanium crystal are held at different voltages to supply a drift field, accelerating the electrons toward one face and holes to the other.

A drifting charge quickly reaches a "terminal velocity." The usual mechanisms for this are elastic scattering with impurities, which operates in metals at low temperatures, and inelastic absorption and emission of thermal phonons, which mimic elastic scattering and are important at normal temperatures. Given the high purity of the detector crystals, elastic scattering with impurities is rare. The low operating temperature ensures that absorption of thermal phonons, which have only μeV energies, has little effect. Instead, the drift speed is set by phonon drag: a drifting charge is quickly accelerated to a speed at which emission of relatively hot ($\sim$ 1 K) phonons balances acceleration by the electric field, resulting in a fixed drift speed that is close to crystal's phonon (sound) speed. The fixed drift speed is determined by the kinematics of phonon emission. This phenomenon has been explored theoretically in an internal note by Blas Cabrera [119].

In discussing the collection of the drifting charges at the ionization electrodes and the appearance of a signal on the amplifier, I simplify the ionization circuit by voltage biasing one side of the crystal and using the amplifier to provide a virtual ground for the other side; see Figure 4.9. The ionization amplifier operates as a current integrator; the signal observed is the voltage drop across the feedback capacitor, which collects a charge equal in number to the number electron-hole

pairs drifted across the detector: $V = Q/C_f$. A feedback resistor in parallel with the feedback capacitor discharges the capacitor with time constant $\tau_f = R_f C_f$. Another way of understanding the signal size is to equate the work done by the amplifier in supplying charge to the detector to the work done on the drifting charges, yielding

$$e\frac{V_b}{d}(N_Q d_e + N_Q d_h) = Q_f V_b \tag{4.13}$$

where N_Q is the number of charge pairs generated and drifted, V_b is the detector-bias voltage, d is the crystal thickness, d_e and d_h are the distance that electrons and holes are drifted in the field, and Q_f is the charge appearing on C_f, equal to the charge supplied by the amplifier to the crystal. The factor V_b/d is the magnitude of the electric field in the crystal. The sum of the hole and electron drift distances is $d_e + d_h = d$, thereby giving $Q_f = eN_Q$. A microphysical picture is also interesting. Just after creation of an electron-hole pair, the hole's electric field lines terminate on its electron counterpart, yielding a tightly contained dipole field. As the two charges are drifted apart, their field lines switch over and become attached to the electrodes to which they are drifting. Gauss' Law requires that surface charge densities appear on the electrodes to terminate these field lines. For the drifting hole, these charges are drawn from the amplifier; for the electrons, they are drawn from the voltage bias. When the hole finally reaches the electrode, all its field lines have terminated on the electrode and one unit of electrical charge has appeared on the feedback capacitor. Note that it is necessary for the electron to also drift completely to the positive electrode; otherwise, some of its field lines reach across to the negative electrode and draw holes into it, cancelling part of the signal arising from the drifted hole. During drift, a fixed current is drawn from the amplifier, given by taking the time derivative of Equation 4.13:

$$I = \frac{dQ_f}{dt} = \frac{e}{d}(N_Q v_e + N_Q v_h) \tag{4.14}$$

where v_e and v_h are the electron and hole drift speeds, fixed by phonon drag.

From these basics, one can see that all created charges must drift completely to their respective electrodes for the full charge to appear on the amplifier's feedback capacitor. This is immediately apparent from above relation: if $d_e + d_h < d$, then Q_f is less than its nominal value. This can also be seen microphysically. Consider the case of a drifting hole. If a hole is trapped by a charged impurity site along the way, the hole's field lines terminate on the site, not on the negative electrode, so the hole becomes invisible to the amplifier. The electrons initially drawn to the electrode by the hole return to the amplifier. However, the charge accumulated on the feedback capacitor up this point is not lost. The impurity site's field lines also disappear. These field lines held an equal number of holes on the same electrode — equal because the impurity has the same position in the crystal as the trapped hole. Thus, an equal number of electrons and holes return to the amplifier, so there is no net current. Whatever charge appeared on the feedback capacitor during the hole drift remains, giving a partial signal. An analogous argument can be made for the case of a hole that diffuses into the "incorrect" electrode (the electrode that the field was trying to drift it away from). Both arguments work for electrons and for a drift field of opposite polarity. Therefore, in all cases, loss of a drifting charge results in decreased signal.

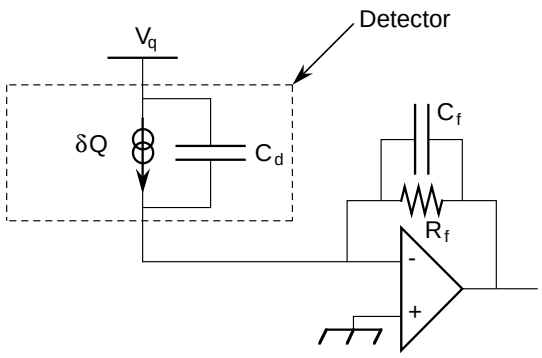

Figure 4.9: Simplified version of ionization-readout circuit. One side of the detector is held at 0 V by the ionization amplifier circuit, which maintains its input terminal as a virtual ground. The detector is voltage biased to provide a drift field. Electrically, the detector can be viewed as a current source with a parallel capacitance $C_d \approx 20$ pF. The feedback component values are $C_f \approx 1$ pF and $R_f = 40\,\mathrm{M}\Omega$. Figure taken from [91].

Ionization Readout

I describe the scheme used in practice to read out the ionization signal from BLIP detectors. As noted above, the standard readout technique for ionization-mediated detectors (Ge-diode detectors, photoconductors, etc.) is shown in Figure 4.9. When phonon and ionization measurements are combined on the same detector, small changes are necessary. One terminal of the phonon sensor is biased and the other is grounded. Therefore, it is necessary to establish a true ground on the detector, rather than the virtual ground provided by an integrator. The simplest approach is to ground the side of the detector that previously was connected to the ionization amplifier and connect the ionization amplifier to the biased side through a coupling capacitor C_c. This scheme is shown in Figure 4.10. The bias resistor R_b on the ionization side is necessary to prevent the voltage bias from competing with the amplifier in sourcing holes to the electrode. The resulting pulse shapes, which are not visibly different from what would be seen from the simplified readout circuit given earlier, are shown in Figure 4.11.

The ionization measurement is complicated by this setup. First, the voltage bias does not actually bias the detector during an event. At zero frequency, the voltage bias sets the voltage drop across the detector — no current flows through the detector ($C_d \approx 20$ pF) nor through the coupling capacitor $C_c \approx 330$ pF. However, on the timescale of the pulse, the voltage bias is irrelevant: the time constant for its response is no smaller than $R_b C_d \approx 800$ μs, far slower than the μs response time of the amplifier or the pulse decay time $\tau_f = R_f C_f \approx 40$ μs. It is the coupling and detector capacitances that maintain the drift field across the detector during the pulse. They are capable of doing this: of order 10^{10} charges are stored on $C_c + C_d$ at 6-V bias, so an energy deposition of order 1 GeV is required to significantly depress the voltage bias. Second, the signal size is modified by this setup. In the limit $C_c \gg C_d$ and an infinitely fast ionization amplifier, the amplifier provides all necessary charge to the detector electrode via C_c. In practice, a fraction $C_d/(C_c + C_d)$ of the charge is provided by the detector capacitance. One way to think of this is to view the detector as

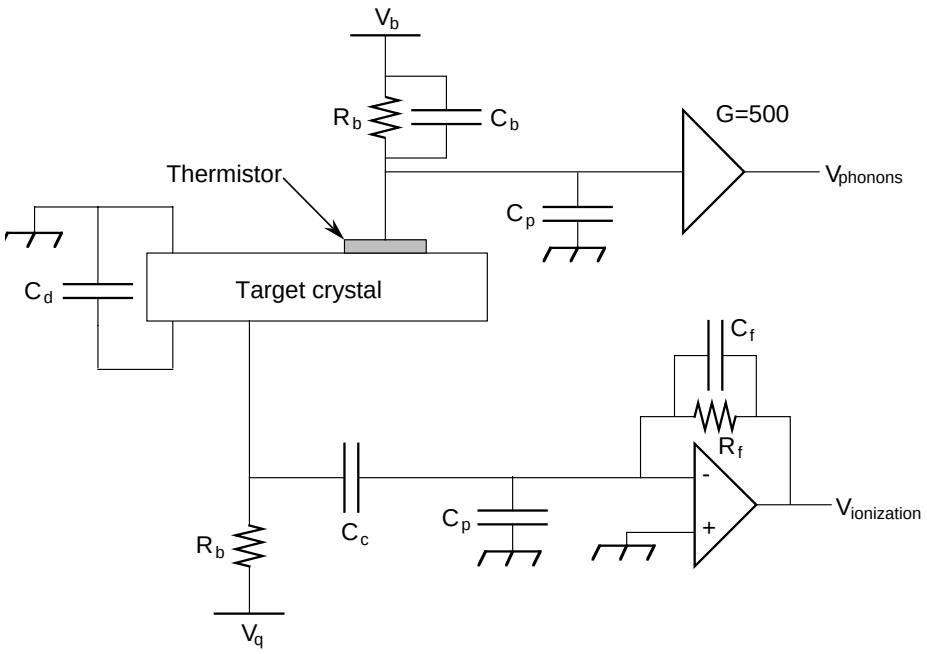

Figure 4.10: Combined ionization- and phonon-readout circuit. The phonon side of the detector is a true ground. The ionization amplifier connects to the biased side via a coupling capacitor $C_c \approx 330$ pF. The detector capacitance is $C_d \approx 20$ pF. The ionization-bias resistor is $R_b = 40\,\mathrm{M}\Omega$. The parasitic capacitance C_p is dominated by FET capacitance and is in the range 40–60 pF. The other component values are given in the text and in Figures 4.5 and 4.9. Figure taken from [91].

a current sink with C_d and C_c in parallel with it and each other. A fraction of the current sunk by the detector flows from each capacitance, but only the fraction flowing from C_c is visible to the amplifier. One can give a microphysical explanation, in terms of terminating field lines, also. This type of loss is not damaging in the same way as loss due to trapping or drift into the wrong electrode: the lost fraction is the same for all events, so the loss only degrades the signal-to-noise ratio of the measurement.

The situation is even more interesting at very early times. The above description assumes the amplifier has zero rise time. In fact, until the amplifier reacts, the terminal of the coupling capacitor attached to the amplifier appears connected to nothing: the amplifier provides no charge to the detector on sub-μs timescales. The detector capacitance provides all the charges to terminate field lines of drifting charges at early times. In the current-sink picture, initially the current sink only sees C_d at early times. When the amplifier reacts, C_c becomes visible and the charged-up coupling capacitor can also source current. The total charge provided by the amplifier is the same as for the infinite-rise-time case, so this effect is merely a curiousity. The only case in which the signal-to-noise ratio might be degraded is if the bias time constant R_bC_d is comparable to the amplifier rise time — the bias could then provide the necessary charges and the amplifier would see no signal.

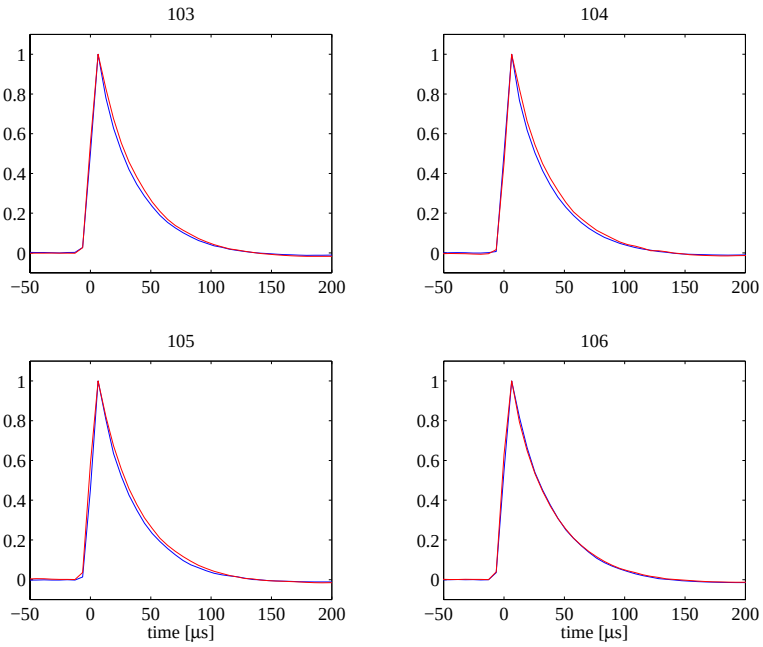

Figure 4.11: Run 19 ionization-channel pulse shapes. Dark: inner electrode. Light: outer electrode. The typical energy conversion (peak height) is 45 μV/keV.

The next step is the discharge of the feedback capacitor through the feedback resistor, with a time constant $\tau_f = R_f C_f$. This has no effect on what happens at the detector because the amplifier input continues to be a virtual ground.

Finally, on a longer timescale, the detector capacitance is recharged by the voltage bias. The coupling capacitor is recharged to its quiescent voltage also — recall that its detector-side voltage dropped as necessary to share the duty of sourcing charges to the detector. Current is now *sunk* by C_c (and C_d); to maintain its input at ground, the amplifier's output must go negative to sink this current through R_f (C_f can be ignored on these long timescales). Thus, a *negative* excursion appears on the timescale needed to charge the two capacitors through the bias, $R_b(C_d + C_c)$.

One might ask if this is the optimal circuit. It appears to be the only practical one. One alternative would be to place the phonon sensor on the bias side. The problem then arises that the FET's source and drain must float to values near the detector-bias voltage, which is difficult to do for arbitrary bias voltage. Another alternative would be to connect the low side of the phonon sensor to the biased side of the detector through a capacitor, but this is physically impossible because the low side of the phonon sensor is the phonon-transmission interface. If one attempts to use the amplifier to bias the detector by engineering the feedback so that it maintains its input terminal at the bias voltage, the FET source and drain must also float to near the bias voltage. A final option is to measure the change in the voltage drop across the detector capacitance; essentially, use the detector capacitance as a charge-to-voltage converter and attach a voltage amplifier. This again requires the front-end FET to float to the ionization bias. This technique was used during Run 18 because of shorted coupling capacitors; information can be found in Andrew Sonnenschein's

dissertation [8]. This introduces minor issues. The ionization pulse shape is determined by the detector capacitance and bias resistance and so varies if the detector geometry is changed. The decay time constant of the pulse would be $R_b C_d \sim 800$ μs ($f_{3dB} = 200$ Hz), pushing the signal to lower frequency where $1/f$ noise can be important. The noise requirements on the voltage bias are more stringent because the voltage bias becomes visible to the amplifier as the detector recharges, which occurs on a faster timecale than for the integrator readout because $C_d \ll C_c$.

Geometry of BLIP Detectors

As discussed in Section 4.1.1, the thermal nature of the detectors implies that any 165-g chunk of germanium suffices for the phonon measurement. The geometry of the BLIP detectors is determined by ionization-measurement constraints. The electrode dimensions used for BLIPs 3 through 6 are shown in Figure 4.1.

The first constraint is detector capacitance. In the simple configuration where C_c is absent, the detector capacitance has no effect on the pulse height — pulse height and shape are determined by C_f and R_f, as described above. The detector capacitance does affect the noise of the ionization measurement. I follow the discussion in Tom Shutt's dissertation [106]. The FET noise can be modeled as a voltage noise at the FET gate, e_g, with amplitude ~ 1 nV/$\sqrt{\text{Hz}}$ [120]. The amplifier tries to keep the FET gate at ground: it responds to a change in the gate voltage by driving current through the feedback loop to cancel it. Therefore, the amplifier tries to enforce the following condition, where e_g is the noise at the FET gate and e_o is the amplifier output:

$$e_g + \frac{e_o}{Z_f + Z_i} Z_i = 0 \tag{4.15}$$

where Z_f is the feedback impedance and Z_i is the impedance between the FET gate and ground, which is the parallel combination of the detector capacitance, the FET's input capacitance, and parasitic capacitance. Z_i and Z_f form a voltage divider with the amplifier output at the top of the divider, the bottom of the divider at ground, and the midpoint (above Z_i) attached to the FET gate, giving the above equation. Solving for e_o gives

$$e_o^2 = e_g^2 \left| 1 + \frac{Z_f}{Z_i} \right|^2 \tag{4.16}$$

Note that the gain factor is not Z_f/Z_i.

At high frequencies, the input and feedback impedances are dominated by capacitances, giving

$$e_o^2 = e_g^2 \left(1 + \frac{C_i}{C_f} \right)^2 \tag{4.17}$$

The feedback resistance rolls off the noise gain at low frequency, $f < (2\pi R_f C_f)^{-1}$. Thus, it is desirable to decrease the detector capacitance until the FET input capacitance dominates. For the BLIP detector geometry, where the inner electrode (see below for discussion of inner-outer segmentation) has a radius of 2.2 cm and the electrode separation is 1.2 cm, the detector capacitance is calculated to be ≈ 20 pF, assuming a dielectric constant of 16 for germanium [106]. This is small compared to the FET gate capacitance, $C_{FET} = 40$–60 pF [115], and so this constraint is met. The resulting noise spectra are shown in Figure 4.12.

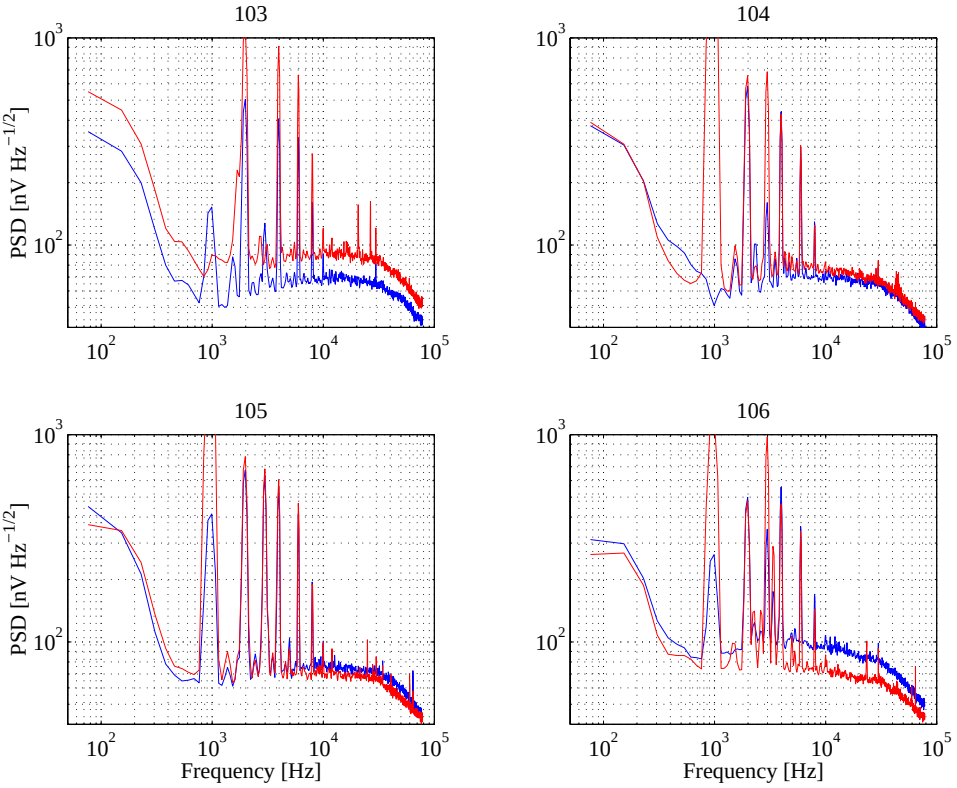

Figure 4.12: Ionization-channel noise spectra. Dark: inner electrode. Light: outer electrode. The peaks are due to cross-talk of the 1-kHz phonon bias and its harmonics. The dip at 1 kHz is due to a notch filter in the electronics chain. The high-frequency rolloff is set by the anti-alias filters (see Chapter 6). Note that, because the phonon-bias cross-talk is coherent (its phase is stable and nonrandom), power spectral density is not the appropriate measure; the Fourier amplitude (in nV, not nV/$\sqrt{\text{Hz}}$) should be used. This can be obtained by using the spectrum's frequency-bin width, $\Delta f = 76.3$ Hz.

Including the coupling capacitor C_c results in a fractional signal loss $C_d/(C_c + C_d)$, so it is necessary to make C_c large compared to C_d. This can, in principle, be done for arbitrary C_d. In practice, thin-film coupling capacitors much larger than a few hundred pF are susceptible to electrostatic damage. This constraint also makes smaller C_d desirable. For Run 19, $C_c = 330$ pF was used, so $C_d \ll C_c$ is also satisfied.

Another constraint is that it is desirable to segment the detectors radially, yielding a disk-shaped inner electrode and an annular outer electrode. This provides some sensitivity to event position via the relative inner- and outer-electrode pulse heights, always useful in a low-background experiment. Furthermore, it provides a guard ring if ionization collection in the outer electrode is poor, as discussed next.

Finally, it is desirable to deal with the bare Ge walls of a standard cylindrical geometry. Studies made during 1992 and 1993 demonstrated that these bare surfaces are susceptible to ioni-

zation loss. One explanation for this is that the bare Ge surface has hanging bonds that result in mid-gap energy levels; it is energetically favorable for drifting charges to become trapped on these sites, in analogy to the case of ionized impurities. As a result, later detectors are curved and the electrodes are wrapped around the curve, leaving only a 1-mm break between top and bottom. E5 and BLIPs 1 and 2 were fabricated in this way and all had excellent ionization collection in their outer electrodes. With BLIPs 3 through 6, the first full-size detectors using α-Si/Al-Schottky electrodes, technical considerations to be discussed in Section 4.3 made a larger break width, ~ 3 mm, more convenient. The photon misidentification in the outer electrodes of these detectors is measurably degraded, as shown in Section 4.5.

Neutralized of Ionized Impurities

In BLIP detectors, charged impurities can trap drifting charges and thereby reduce ionization yield, which can degrade the detectors' nuclear-recoil-discrimination capability. In this section, I discuss why impurities tend to be ionized and the procedure used to neutralize them.

A simplified diagram of the band structure of high-purity germanium is shown in Figure 4.13. The valence and conduction bands are separated by an energy gap of 0.743 eV. Impurities result in donor and acceptor states in the gap. The number of acceptors in the crystal is N_a and the number of donors N_d. For the germanium used for BLIPs, the crystals are p-type, $N_a \gg N_d$, with $n_a - n_d \approx 6 \times 10^{10}$ cm^{-3} [121]. The dominant acceptor levels lie at $\epsilon_a \approx 12$ meV above the valence band. At higher temperatures, the large numbers of available hole states in the valence band and electron states in the conduction band ensure almost complete ionization of the impurities, with $\sim N_a$ holes in the valence band and $\sim N_d$ electrons in the conduction band: the Boltzmann factor is 1 for all states because of the high temperature and thus favors neither impurity or free states; the many more degrees of freedom in the conduction/valence bands makes it far more likely for a charge to occupy one of these states. However, when cooled to 20 mK, free charge is Boltzmann suppressed by a factor $\exp(-\epsilon_a/kT) \sim e^{-5800}$ — there is no free charge. Not all the impurity sites are filled, however. Rather, it is energetically favorable for the N_D electrons to fall onto acceptor sites rather than to bind to the N_D donor sites — they annihilate with holes bound to the acceptor sites. This results in N_D ionized donors and N_D ionized acceptors. If left this way, ionization collection in the crystal is hopeless — the $\mathcal{O}(10^9)$ cm^{-3} ionized impurity sites trap charges generated by events; recall, only about 330 charge pairs are created per keV of energy deposited (for electron recoils). This trapping can be overcome by applying an ionization bias so large that drifting charges acquire too much kinetic energy to be trapped, but the resulting Neganov-Luke-phonon generation (see below) would swamp the intrinsic phonon signal.

Trapping can be prevented by neutralizing the ionized impurities. While neutralized impurities are lattice defects and cause elastic scattering, they are orders of magnitude less likely to trap charges than ionized impurities [106]. The detectors are neutralized by exposing them to a LED while at low temperature and with no ionization bias applied. Photons from the LED produce electron-hole pairs; the absence of a drift field allows these free charges to either recombine or be trapped on ionized impurities. Germanium, with its low (10^{10} cm^{-3}) impurity concentration, can be neutralized with a LED "bake" of an hour or less. Silicon requires a few hours of baking to reach neutralization. It is necessary to keep the detector well below the the activation temperature of the impurity levels during and after baking to maintain neutralization; this is easy, given that 12 meV corresponds to approximately 140 K.

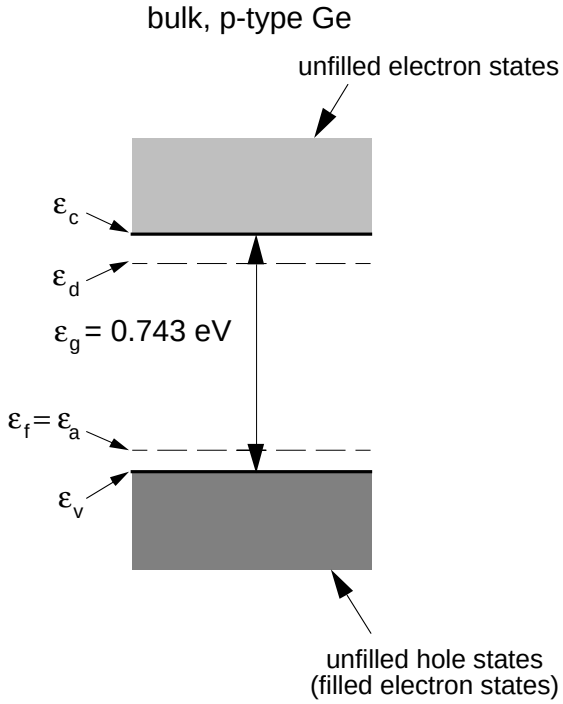

Figure 4.13: Simplified band structure of germanium. The acceptor and donor levels are indicated.

Neutralization degrades with time, presumably due to liberation of trapped charges by scattering with drifting charges. In germanium, restoration of neutralization is accomplished by grounding the detectors for a brief period; particle interactions create the necessary free charge to refill the traps. The Run 19 BLIPs, with α-Si/Al-Schottky electrodes, show no signs of degraded ionization collection with a 50-minute-biased/5-minute-grounded cycle. This has not been optimized, so it may be possible to increase the biased period or decrease the grounded period. It is interesting to note that the required grounded fraction is, in general, independent of the average event rate. For example, at SUF vs. the surface, there are fewer events during the grounded period, creating less free charge, but there are also fewer events during the biased period, so fewer impurities become ionized by scattering with drifting charges.

Neganov-Luke Correction, Ionization "Energy" and Yield

As noted above, charges drifting across a detector crystal under influence of the applied electric field reach a terminal velocity. The work being done on the charges by the field is dissipated in the form of phonons emitted by the drifting charges, which contribute to the phonon signal:

$$E_P = E_R + eV_bN_Q \tag{4.18}$$

where E_R is the recoil energy, the energy deposited by the particle interaction, V_b is the bias voltage across the detector, and N_Q is the number of charge pairs created. The additional phonon energy generated is the number of charge pairs created and drifted across the crystal multiplied by the

work done by the field per charge pair. Note that the work done is independent of the position of the creation of the charge pairs because the sum of the path lengths of the two charges of a pair is always the thickness of the crystal (in the absence of trapping). Neganov and Luke discovered this effect independently [122, 123]; the phonons emitted by the drifting charges are called Neganov-Luke phonons. To obtain the recoil energy from the observed phonon and ionization signal, it is necessary to invert the above expression. This inversion is known colloquially as "Luke correction." Note that, since the noises on E_P and N_Q are independent because they are dominated by electronics noise, the recoil energy thus calculated is noisier than each individual measurement.

The above equation assumes a nontrivial fact: that the entire recoil energy appears as phonons. This is true because even the energy used to create electron-hole pairs appears in the phonon channel. When these electrons and holes recombine with free charges at their respective electrodes, each created pair releases the bandgap energy in phonons. This fails if charges are trapped on ionized impurities — they are prevented from recombining.

The quantity "ionization energy" is defined by

$$E_Q = N_Q \times 3.0 \text{ eV} \tag{4.19}$$

That is, the ionization energy is the recoil energy inferred from the number of charge pairs N_Q under the assumption that the event is an electron recoil. E_Q is *not* the recoil energy if the event is a nuclear recoil. E_Q is defined simply for convenience and should not be further interpreted.

Finally, the ionization yield is defined by

$$Y_R = \frac{E_Q}{E_R} \tag{4.20}$$

An alternative definition for ionization yield is $Y = N_q/E_R$; the energy-ratio definition used here is always subscripted with R to indicate the distinction. Y_R is the parameter used to disriminate nuclear recoils from electron recoils. For bulk electron recoils, $Y_R \approx 1$. For bulk nuclear recoils, $Y_R \sim 1/3$, with a weak energy dependence. Because of their suppressed ionization collection, dead-layer electron-recoil events lie at $Y_R < 1$; their exact position is dependent on the type of ionization electrode used and is discussed in detail in the next section.

4.2 The Dead Layer

The dead-layer phenomenon — reduced ionization collection for events occurring in a thin surface layer — has limited the WIMP sensitivity of CDMS in all runs up to Run 19; see Figure 4.14 and [8]. In this section, I discuss the history of the problem, the strategies taken to understand and overcome the phenomenon, and the performance of detectors using α-Si/Al-Schottky electrodes, the present best solution to the problem.

The new α-Si/Al-Schottky ionization electrodes were developed by Jochen Hellmig, Bernard Sadoulet, Tom Shutt, and Betty Young, with advice from Blas Cabrera and Paul Luke. I present a discussion of this material here because the good electron rejection ($> 95\%$) of the Run 19 BLIP detectors (discussed in Section 4.5.3) is due to this new ionization-contact technology.

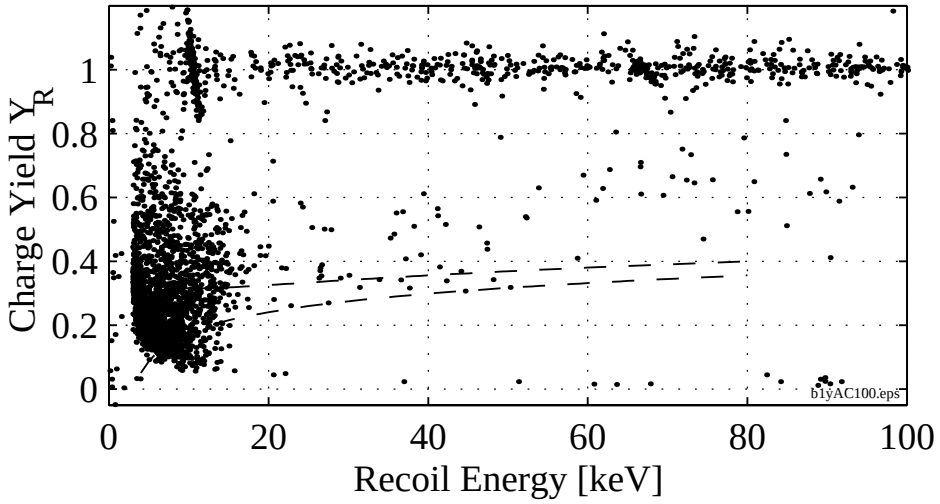

Figure 4.14: Run 18 veto-anticoincident data from BLIP1. A band of dead-layer events intrudes into the nuclear-recoil band (dashed) at low energy, < 50 keV. The dashed lines indicate only the $\pm 1\sigma$ width. The region below 20 keV is dominated by tritium events, which are discussed in Chapter 8. Figure taken from [8].

4.2.1 Discovering the Dead Layer

Even before CDMS began operations at SUF, it had been observed that BLIP detectors with implanted contacts suffer from reduced ionization collection for low-energy photons incident on the detectors externally. This is easily seen using low-energy (14-, 18-, and 21-keV) X-rays available from ^{241}Am, as is shown in Figure 4.15. This fact was not a primary concern at the time because there was little reason to expect there would be a significant low-energy-photon flux in the Icebox. There are simple arguments behind this expectation. One possible source of low-energy photons is Compton scattering of much higher-energy photons within the last tens of microns of surfaces surrounding the detectors. Only this skin layer can produce low-energy photons that can reach the detectors because of the small penetration length of such photons: the low-energy ^{241}Am lines have penetration lengths of ~ 17, 31, and 51 μm in Ge. Furthermore, only a small fraction of Compton scatters yield low-energy photons, providing another large suppression factor, $\sim E_{max}/E_\gamma$, where E_{max} is the maximum energy of the low-energy-photon range of interest and E_γ is the incoming photon energy, typically hundreds of keV to ~ 2 MeV. Thus, Compton scattering is not expected to generate enough low-energy photons to be a problem.

A second source would be a surface or bulk radioactive contaminant of the surrounding materials that emits low-energy X-rays. Low-energy X-rays do not come directly from nuclear transitions — the energy scale (10–20 keV) is small compared to the separation of excited nuclear states (^{241}Am is an anomaly with its 60-keV gamma emission). Rather, such X-rays are emitted when a K-shell electron disappears due to a nuclear-decay process and electrons from higher atomic energy levels cascade down to fill the empty state. For "run-of-the-mill" radioactive contaminants, such as the uranium and thorium chains, these X-rays have energies of $\sim$ 100 keV. For ^{40}K,

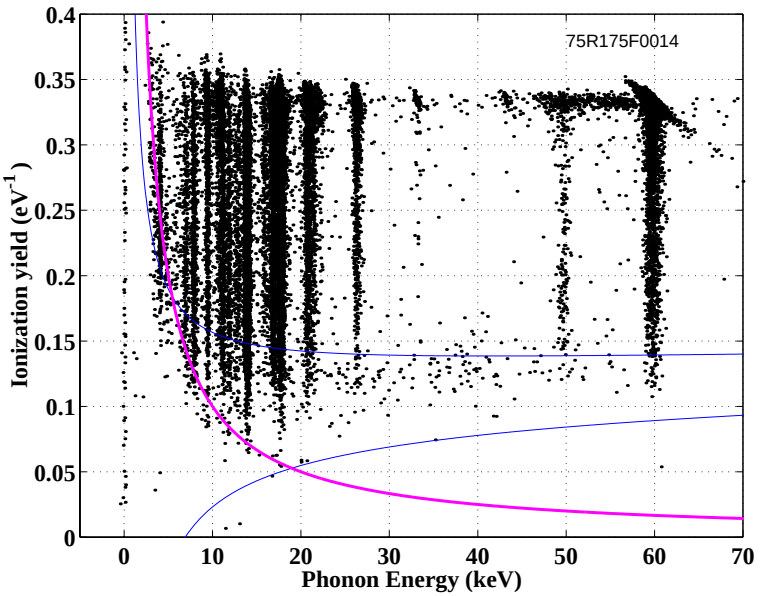

Figure 4.15: Data from exposure of a p-p$^+$-type test device to ^{241}Am. The vertical axis is the ratio of the number of electron-hole pairs, N_Q, to the recoil energy, E_R. Recall $E_Q = N_Q \times 3.0$ eV, so bulk electron recoils lie at 1/3 and nuclear recoils at about 1/9. Figure taken from [121].

its K-shell X-rays are very low energy, all below 4 keV. Photons of these energies would not be problematic — the former have large penetration depths, the latter are near or below threshold. Moreover, previous low-background experiments saw no evidence of low-energy X-rays incident on the detectors externally (though internal low-energy X-rays due to cosmogenic activation of the detector material were apparent; see Figure 2.11).

The observation of a clear population of reduced-ionization-yield events in the run of the first BLIP detector at SUF in summer, 1996, was thus quite unexpected. More recent data taken with BLIP 1 during Run 18 (spring/summer, 1998) are qualitatively the same and are shown in Figure 4.14. The lack of photon lines indicates that these events do not arise from low-energy X-ray emission. Arguments presented above imply that Compton scattering of high-energy photons is not a reasonable explanation either. Experimentally, as is seen in Figure 4.15, a photon explanation would lead one to expect a continuous distribution in ionization yield, "raining" down from the photon band; such a distribution is not seen in Figure 4.14. Low-energy electrons provide a reasonable suspect. There are a number of continuum sources of low-energy electrons — low-energy beta decay (^{210}Pb, ^{14}C), Auger electrons, etc. Moreover, because of their low mass, electrons severely multiply scatter and thus their penetration depth is very small. The 90% stopping lengths (or practical ranges) in germanium are 0.5 μm and 10 μm at 10 and 60 keV, respectively [124]. Why were low-energy electrons never considered in estimates of expected background-particle rates? CDMS has drawn the majority of its expectations of background-particle rates from a number of experiments using Ge ionization-mediated photon spectrometers operated at 77 K because of their superficial similarity to the CDMS detectors. Such detectors employ a coaxial electrode geometry,

with the outer electrode being 1-mm thick and thus dead. With such a thick dead layer, low-energy electrons and photons are fairly invisible — they may be present, but they produce no signal and the detector does not trigger. Furthermore, any such events were masked by the $\sim$1 keV^{-1} kg^{-1} d^{-1} photon background at low energies.

4.2.2 The Physics of the Dead Layer in p-Type Implanted Contacts

Though it was not possible to devote large resources to it at the time, the physics of the dead layer was explored using small, ionization-only devices prior to the first CDMS I run. The primary documentation of this work is in the undergraduate thesis of Rodney Diaz, who analyzed a large fraction of the data collected with these devices [125]. More recently, data taken with miniature BLIP-style devices have clarified the physics of the dead layer in the p-type implanted contacts used in the pre-Run 19 BLIPs [121].

The ionization contact used in BLIP detectors up through BLIPs 1 and 2 is implanted boron. The characteristic thickness of these contacts is $\approx$ 450 nm. The implantation is to the point of degeneracy: the impurity density is so high that charges hop from site to site easily, and the number of free charges is large compared to the unimplanted case, so the material is effectively metallic. A band of new states forms in the region between the nominal acceptor level and the nominal valence band edge. The resulting detectors have a p^+-p-p^+ structure — the two contacts are degenerate p-type while the bulk is semiconducting p-type material.

If thermal equilibrium is assumed, then the structure of the bulk-contact interface is simply given by requiring alignment of the Fermi levels of the contact and the bulk. This is illustrated in Figure 4.16. The Fermi level in the bulk lies at $\epsilon_a \approx 12$ meV because of the low temperature; no holes are excited into the valence band. The Fermi level in the contact lies slightly above the valence band edge, near the edge of the band of new states created by degenerate doping. Alignment of the Fermi levels requires a slight bending of the conduction and valence band edges, as illustrated in Figure 4.16. A charge dipole forms at the interface to provide the electrostatic potential that creates this bending. The shift in the bands is very small, approximately ϵ_a minus the difference between the Fermi level and valence band edge in the contact.

When an event occurs near the surface of the detector, the drift field drives electrons one way and holes the other. Consider the face that is positively biased. Holes must be drifted away from this contact in order to be collected by the other, negative electrode. However, the charges produced are hot and so may diffuse before the drift field has a significant effect on their motion. The charge cloud produced by a recoiling particle may also shield itself because the separating charges have dipole fields that counter the drift field. In any case, holes may diffuse into the positive contact because there is no barrier preventing them from doing so (and, if anything, a small potential well inducing them to fall into the contact). Ionization is lost — a dead-layer is born. For electrons, the situation may be slightly better because the $\approx$ 12-meV band-bending provides a barrier. However, empirically, it is seen that this barrier is either nonexistent or negligible on the energy scale of the hot, diffusing charges. If such a barrier existed, then ionization collection for low-energy photons incident on the negatively-biased side of the detector would be better than for the positively-biased side. In the first case, holes drift toward the "source" electrode (the electrode exposed to low-energy photons) and electrons away. Diffusion of holes does not result in lost charge because the source is near the "correct" (negative) contact for holes. The barrier would prevent diffusion of electrons into the "incorrect" (negative) contact. In the opposite bias polarity, the situation is reversed:

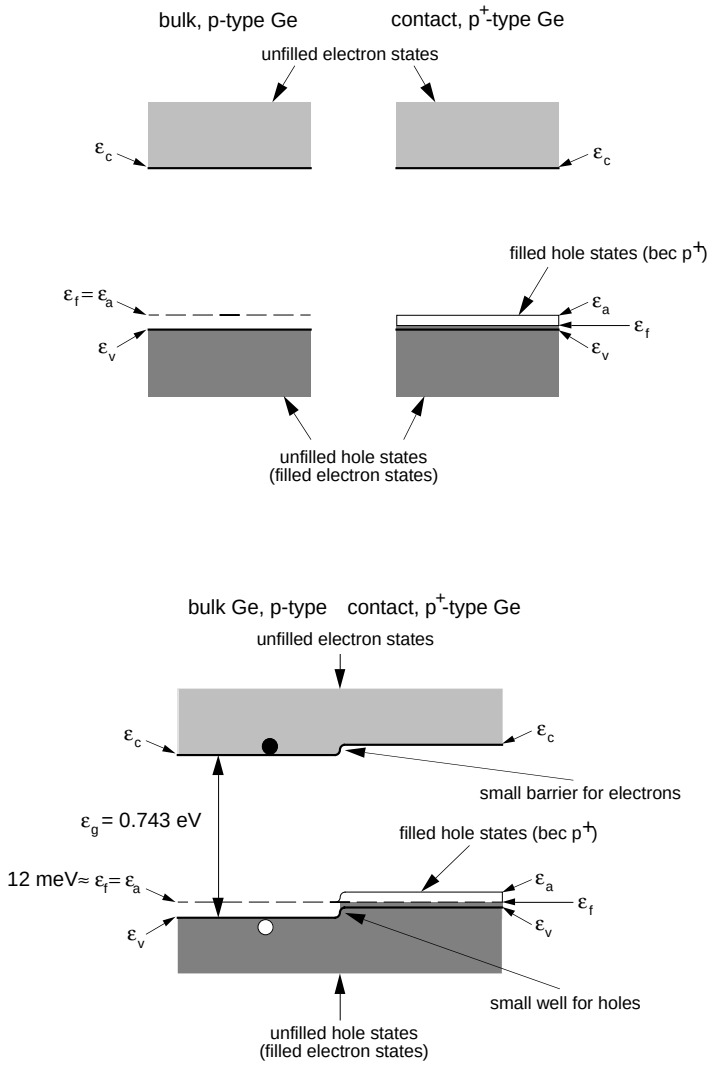

Figure 4.16: Diagram of p-p$^+$ interface under assumption of thermal equilibrium. Top: "before" connection: conduction and valence bands aligned, Fermi levels misaligned. Note that, in the degenerate contact, there are states between the nominal acceptor level ϵ_a and the valence band edge. The Fermi level lies close to the valence band edge. Bottom: "after" connection: Fermi levels aligned, conduction and valence bands offset by the difference between the Fermi level and the valence band in the contact.

holes would see no barrier against diffusion into the incorrect, positive contact, while the barrier would be unnecessary for the electrons. No such difference in dead-layer behavior is observed in p-p^+-type devices [106]

In addition, the above thermal-equilibrium picture is invalid because of the low operating temperature and neutralization process. It is therefore likely that the band-bending is smaller than expected and no significant barrier exists. This explanation of the dead layer — "back-diffusion" of charges into the wrong contact — gives the correct length scale of the problem. The observed length scale is inferred from the ^{241}Am data by a deconvolution method developed by Tom Shutt and Steve Yellin, the results of which are shown in Figure 4.19. Clearly, the length scale is of order 10 μm, large compared to the thin (450-nm) contact thickness and better explained by a diffusion model. This model for the dead layer survives tests of other contact structures and explains the success of the α-Si/Al-Schottky contacts eventually developed.

4.2.3 Hydrogenated-Amorphous-Silicon/Aluminum-Schottky Contacts

To address the dead-layer problem, CDMS began testing a number of new contact technologies in early 1997. The present best reference on the ionization-contact work is [121]. I present the arguments made in that paper, modified somewhat by recent discussions with Tom Shutt.

The dead-layer phenomenon arises because there is no barrier against diffusion of charges into the implanted contact for events near the detector surface. The clear way to counteract back-diffusion is to provide a barrier. For detectors operating at higher temperatures, two solutions present themselves.

The first possibility is a p-n^+ contact structure, much like a diode. (n^+ denotes degenerately doped n-type material.) When in thermal equilibrium, a charge dipole forms at the interface to provide a potential change that aligns the Fermi levels of the bulk p-type material and the n^+ contact. The Fermi level in the bulk is $\approx$ 12 meV above the valence-band edge. The Fermi level in the degenerate contact is a slightly below the conduction-band edge, in analogy with the p^+ contact. In aligning the Fermi levels, the valence and conduction bands must therefore exhibit a sharp change in energy at the interface: they must shift by almost the entire bandgap energy of 0.743 eV. This provides a barrier against back-diffusion of holes from the bulk to the contact, a "built-in field." This structure has the disadvantage that it only presents a barrier to one sign of charge, so it only solves the dead-layer problem on one side of the detector.

Moreover, the high-temperature picture of the p-n^+ contact fails at low temperature. Given the energy gaps involved, thermal equilibrium cannot be attained because charges are not thermally activated at 20 mK. Neutralization of the bulk drives the system further from equilibrium. Test devices using p-n^+ contacts operated at low temperature show evidence for a dead layer not very different from p-p^+ contacts [121].

The second possibility is a Schottky-barrier contact. In such a contact, a metal is deposited on the surface of the semiconducting crystal. In a way similar to the p-n^+ interface, a charge dipole is formed at the interface in order to align the Fermi levels of the metal and the semiconductor, resulting in a bending of the conduction- and valence-band edges. Such a barrier rejects charges of one sign only also. Schottky-barrier devices produced by the Stanford group, using Al on bulk germanium, show good ionization collection when the face that sees the photon source is biased at positive voltage; they show much worse collection in the opposite polarity [126]. With positive voltage on the source side, electrons drift toward the source side and holes away; good collection

in this polarity indicates the Schottky barrier has a built-in field that blocks holes well. Bad collection in the other bias polarity confirms this sign of the built-in field. Thus, in contrast to the p-n^+ structure, the Schottky barrier qualitatively displays thermal-equilibrium behavior, even at low temperature. This may reflect the fact that the metal provides many more free charges than the n^+ contact in the p-n^+ structure, allowing the system to reach thermal equilibrium in spite of the low temperature. If one could make Schottky-barrier contacts on the two faces of a detector with built-in fields that block electrons at one face and holes at the other, this technology would eliminate the dead layer on both detector faces for one bias polarity. This is, in principle, possible by using metals with appropriate work functions (the position of the Fermi level with respect to vacuum). In practice, the vast majority of Schottky-barrier contacts bend only in the way seen above; this is believed to occur because hanging bonds at the semiconductor surface create mid-gap states; the Fermi level in the metal, rather than lining up with the bulk-semiconductor Fermi level, is pinned to these mid-gap states [127, 128].

An alternative to the above is a hydrogenated-amorphous-silicon contact, referred to as α-Si in the following. "Amorphous" refers to the noncrystalline nature of the material; it is deposited by sputtering or evaporation rather than being grown as a crystal. The material is hydrogenated by deposition in a partial-hydrogen atmosphere to seal dangling bonds arising from the amorphous nature of the material. α-Si possesses a bandgap of 1.2 eV, almost twice as large as bulk germanium. If the bands of the bulk germanium and a deposited layer of α-Si are centered on each other, the α-Si blocks diffusion of charges of both polarities with large barriers, of order tenths of an eV; see Figure 4.17 for a schematic illustration. There is uncertainty about the details of the band structure of α-Si; there may be a significant population of mid-gap states. Furthermore, the details of the α-Si-Ge interface are not known; surface states may also yield mid-gap states. It is also not clear whether the bandgaps of the α-Si and bulk germanium are centered relative to one another.

The history of the α-Si contact is an interesting footnote. The use of α-Si was among a number of suggestions made after thinned and lightly doped implanted contacts proved no better than standard contacts. The α-Si contact was first suggested by Paul Luke of LBNL because it is a standard solution to dead-layer problems in 77 K Ge diode detectors. However, the physics of the dead layer at 20 mK is quite different, and it was Tom Shutt who realized that the above blocking phenomenon is possible and developed an α-Si contact for the BLIP detectors.

Empirically, contacts using α-Si effectively solve the dead-layer problem. Data taken with ABL1, a device with a 27.5-nm α-Si layer on both faces, during exposure to an ^{241}Am source are shown in Figure 4.18. The improvement of ionization collection between Figure 4.15 and this figure is clear: the α-Si electrode pulls low-yield events out of the nuclear-recoil band. The ionization yield as a function of depth (calculated by deconvolution of the ^{241}Am data as mentioned earlier) for p-p^+ and α-Si contacts is shown in Figure 4.19. The α-Si electrode yields uniformly higher ionization collection at all depths. Most importantly, the minimum value of the ionization collection of the α-Si contact is significantly higher than that for nuclear recoils, solving the problem of discriminating low-energy electrons from nuclear recoils.

Comparison of data taken with opposite bias polarities may give some indication of the relative alignment of the α-Si and bulk-Ge bandgaps. In Figure 4.20, data taken during exposure of ABL1 to a ^{14}C electron source are shown; the bias sign of the electrode exposed to the source is changed between the two data sets. Two effects are seen. First, when the source side is at negative bias, the median ionization collection at low energies is *better* than when the source side is

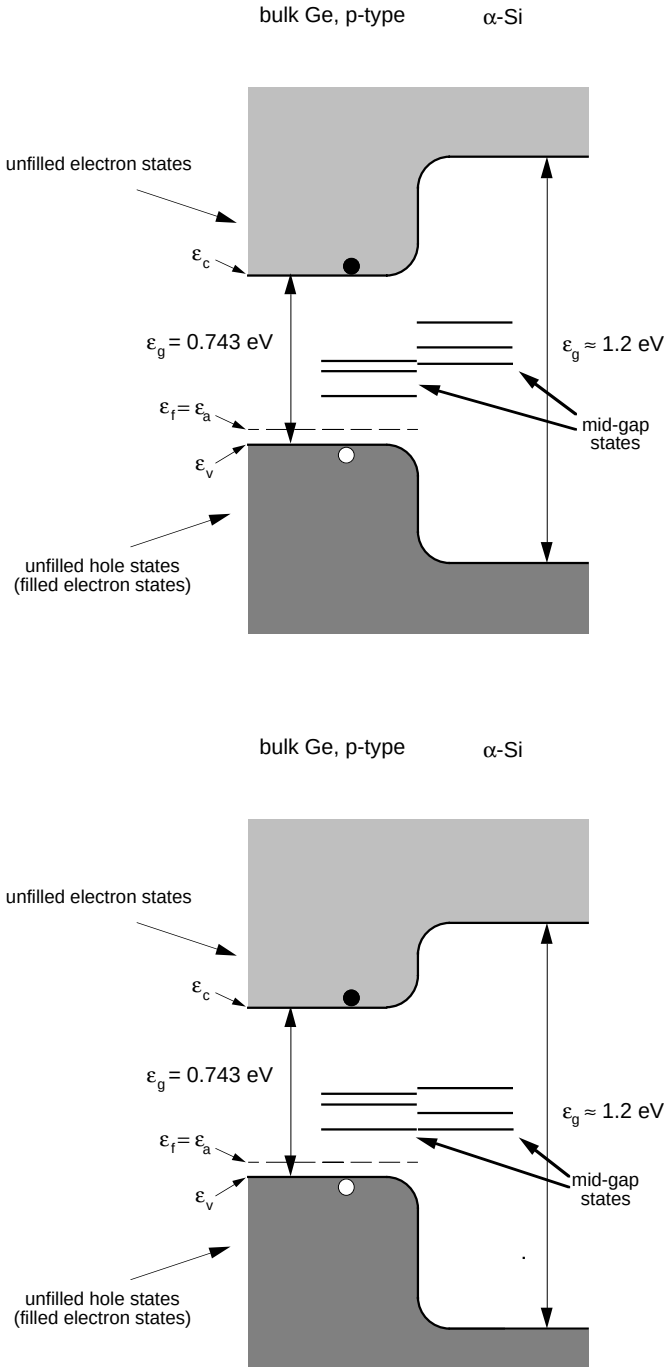

Figure 4.17: Schematic illustration of bulk-Ge/α-Si interface. Top: Assuming alignment of the bandgap centers. Bottom: qualitative misalignment suggest by ABL1 data (see Figure 4.20). Mid-gap states that may serve to define the alignment are schematically indicated.

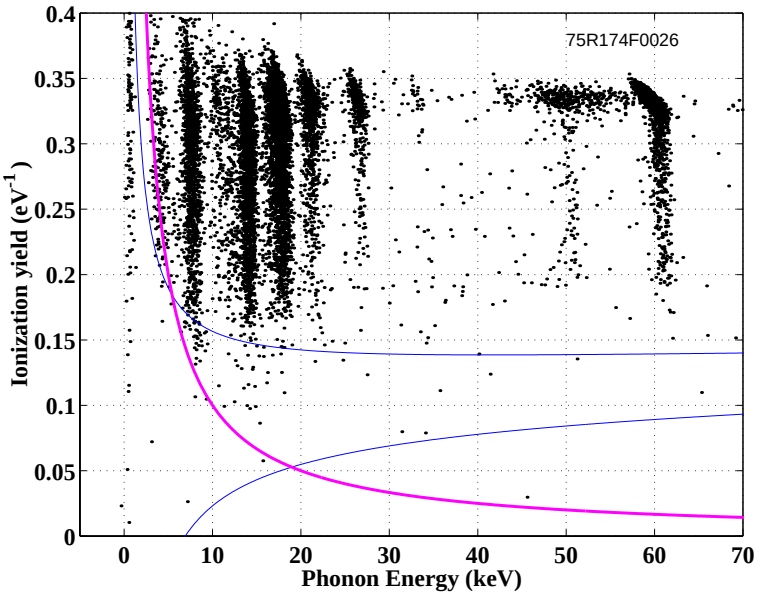

Figure 4.18: ^{241}Am source data taken with ABL1, a test device using an α-Si contact. The vertical scale is the same as in Figure 4.15. Figure taken from [121].

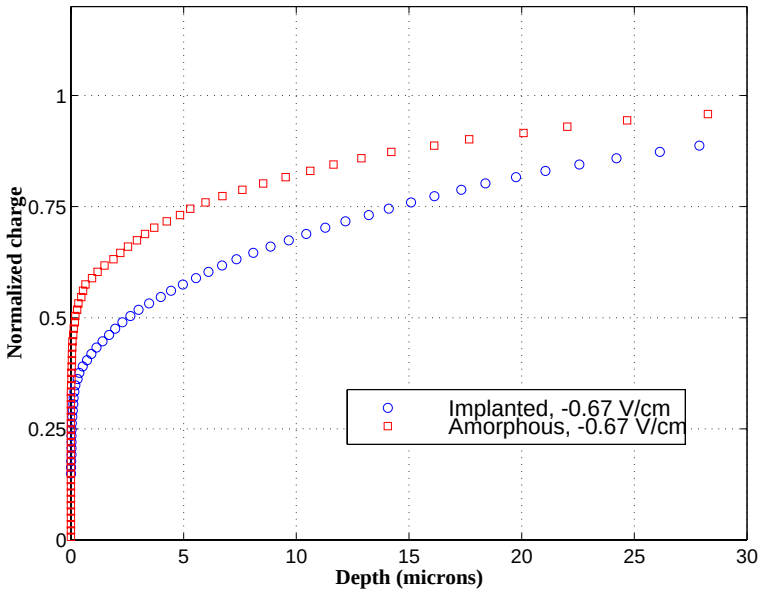

Figure 4.19: Ionization collection vs. interaction depth for p-p$^+$ and α-Si contacts, as deconvolved from Figures 4.15 and 4.18. Nuclear recoils lie at normalized charge of ≈ 0.3 and thus are mimicked by dead-layer events in a p-p$^+$-type contact. Figure taken from [121].

at positive bias. Second, the tail to low collection when the source side is at negative bias is *worse* than when the source side is at positive bias. I concentrate on the second point, since the median position in either case is above the nuclear-recoil band.

In the bias direction where the source side is negative, electrons are drifted away from and holes toward the source side; back-diffusion of *electrons* into the source-side electrode yields depressed ionization collection. Thus, the α-Si layer is less effective at blocking electrons than holes. One explanation of these data is that the center of the α-Si bandgap is lower (in electron energy — the standard polarity of bandgap diagrams) than the center of the bulk-Ge bandgap, and thus the barrier seen by electrons is smaller than the barrier seen by holes; see Figure 4.17. Data were also taken with a device that used amorphous germanium rather than amorphous silicon, ABL2. α-Ge is expected to have a bandgap half as large as α-Si. ABL2 showed no significant degradation compared to ABL1, unexpected if the remaining "deadness" of ABL1 were due to an insufficiently large barrier. However, it is likely that the center of the α-Ge bandgap is not in the same place, relative to the bulk, as the α-Si bandgap because the effect of surface states differs in the two cases; the additional shift may be just enough to leave the barrier to electrons unchanged. The barrier to holes would be decreased, but it is not clear that this would have an observable effect.

Another explanation of this phenomenon may be provided by the additional Al layer deposited. α-Si is conducting, but barely so, with GΩ/square [121]. To make electrical contact to the electrode, a layer of Al is deposited on top of the α-Si. The α-Si layer is very thin, 27.5 nm, so charges can tunnel through the blocking barrier to the Al layer [121]. Note that this tunneling is much more likely for the correct-sign charges because the drift field pushes them against the barrier, providing them many chances to tunnel; incorrect-sign charges are quickly pushed away by the drift field and therefore have few chances to tunnel. The Al layer on ABL1 is 250-nm thick. The Al layer naturally forms a Schottky-barrier contact with the α-Si layer. Because the α-Si layer is so thin, the built-in field created may penetrate slightly into the bulk germanium [121]. As discussed earlier in the context of simple Al-Ge Schottky-barrier devices, the built-in field helps ionization collection for one bias polarity and hinders it for the other. Thus, the ABL1 data can be interpreted as indicating the Schottky barrier formed at the α-Si-Al interface has a built-in field that penetrates into the bulk germanium and blocks holes but attracts electrons. The polarity is in agreement with the previously cited tendency of Schottky-barrier contacts. Two more devices, ABL3 and ABL4, both of which used a p^+ contact on the side away from the source, exhibited similar behavior in that the amorphous contact also blocked holes better than electrons.

Based on the above data, α-Si/Al-Schottky contacts were deposited on the Run 19 BLIP detectors. The Al-layer thickness was reduced to 10–20 nm; it was feared that the tail to low ionization collection in even the positive-bias ABL1 data might be due to electrons stopping in the thick (250 nm) Al layer; being metallic, this layer is probably completely dead.

It is not yet clear whether the behavior of the ABL devices is explained by misaligned bandgaps, Schottky barriers, or something else. In any case, it may be possible to control the exact behavior by doping the amorphous material during deposition. It may be possible to reverse the polarity of either the bandgap misalignment or the Schottky barrier and thereby produce a detector that has "no-tails" behavior at both electrodes for one bias polarity. It should be noted that there are a number of variables — surface states and surface preparation, hydrogenation of the α-Si, doping — so that further development will not be trivial. Schedule constraints during preparation of the Run 19 detectors left too little time to make such tests, but, clearly, there is

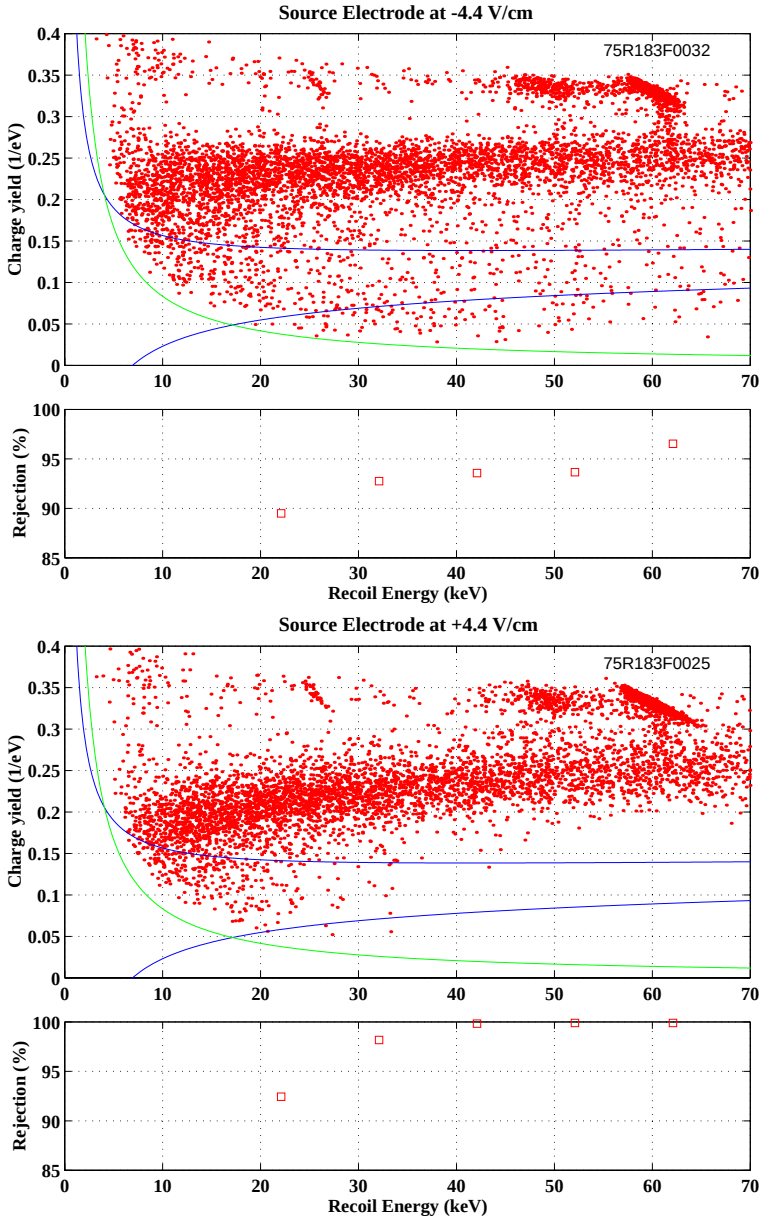

Figure 4.20: ^{14}C source data taken with ABL1 with two different bias polarities. Top: source-side electrode at negative bias, so holes drift toward and electrons away from source-side electrode. The contact being studied with the source blocks electrons. Bottom: source-side electrode at positive bias; electrons drift toward and holes away from the source-side electrode, so the contact being studied blocks holes. The vertical scale of the ionization-yield plots are the same as for Figures 4.15 and 4.18. Figure provided by T. Shutt.

much interesting parameter space that requires investigation.

4.3 Detector Fabrication

The majority of BLIP-style detectors have been fabricated by John Emes, a technician at the Lawrence Berkeley National Laboratory, using the equipment of Eugene Haller in LBNL Building 2. Many of the steps are modifications of standard germanium processing steps for the unusually shaped BLIP crystals.

- *Wafer Slicing*: A boule (a cylindrical crystal grown from a seed) is received from the supplier, in this case EG&G Ortec. The boule is grown along the $\langle 100 \rangle$ axis. The $\langle 110 \rangle$ facets are marked for later alignment of the thermistor and crystal axes. The boule is sliced into 1.2-cm-thick disks perpendicular to the growth axis using an inner-diameter saw with a diamond blade. Witness samples are taken from the ends and middle of the boule.

- *Rounding, Lapping, and Polishing*: The sliced disks are given to the LBNL Optical Shop for grinding of the curved edges. The grinding machine is much like a lathe. The tool bit has a curved edge — this edge determines the 0.6-cm radius of curvature of the crystal edge. The crystal is spun and slowly ground down using the tool and a lubricating slurry. The ionization side of the curved detector is hand-lapped to produce a uniformly mechanically damaged surface and then polish-etched to remove this mechanical damage. The ionization side is protected with wax and the thermistor side mechanically lapped. In both cases, lapping removes about 200 μm of material. The thermistor side is then optically polished. The crystal is placed on a turntable-like wheel, thermistor side down, with a large cylindrical weight glued to the ionization face using wax. A slurry containing water and a fine grit is flowed onto the turning wheel while the crystal is held stationary by attachments to the weight. As the grit is made finer, the finish of the crystal surface improves. The later eutectic bonding requires the crystal surface to be flat to better than 100 nm over the 3.1-mm length of a thermistor. This is checked using a Newton's Rings-type interferometer. No precautions are taken to ensure that the flat smoothly joins the curved edge; there is always a slight bevel at this point. There is no evidence that this causes an ionization-collection problem.

- *Implantation*: Prior to BLIPs 3 through 6, all BLIP detectors were implanted with boron to create a p^+-p-p^+ electrode structure, as discussed above. Even in the case of the α-Si/Al-Schottky electrodes used for BLIPs 3 through 6, it is necessary to perform a boron implant in the region to which the thermistors are bonded to provide a reasonable electrical contact to the p-type NTD Ge. A second small, implanted patch is needed to provide a resistive element for thermal pulsing. During ion implantation, the crystal is attached to a wheel tilted at 45° to the ion beam. The wheel spins to help make the implantation uniform. The crystal is tilted to prevent channeling. Channeling is a phenomenon by which, if the beam is aligned with a crystal axis, the ion penetration depth is anomalously large because each crystal plane shadows successive layers. Were the crystal perpendicular to the beam, the entire flat face would suffer from channeling.

- *Thermistor Preparation*: The (baked) NTD-Ge wafer to be used for thermistors is lapped and polished as the crystals are. The wafer faces are implanted with boron and metallized with

gold to provide contact for wirebonds. A thin (20-nm) palladium sticking layer is deposited before the gold. Both the palladium and gold are deposited by thermal evaporation. The wafer is diced into individual thermistors. The thermistors are approximately 3.1 mm by 3.1 mm by 2.6 mm.

- *Eutectic Bonding*: The use of eutectic bonding was discussed earlier. The technique is described in great detail in Ning Wang's dissertation [105]. A 200-nm-thick pad of Au-Ga (2% Ga) is thermally evaporated over the region where the bond will be made (a 20-nm-thick Pd sticking layer is also evaporated) using a shadow mask. The thermistor is placed on the pad, with its $\langle 110 \rangle$ crystal axis aligned with the detector crystal's. The alignment is done by eye and so is accurate to only a few degrees. With pressure applied to the thermistor using a weighted spring-and-pin assembly, the thermistor-Au-Ge-crystal sandwich is heated to the eutectic point of Au-Ga-Ge, where the mixture at the interface melts but the surrounding germanium does not. The system is then slowly cooled, allowing the interface to recrystallize. There was some worry that this step would be problematic for BLIPs 3 through 6, for which the thermistors are placed close to the edge of the optical flat rather than at the center. The polishing process produces more curvature toward the crystal edge. The crystals and thermistor wafer were polished to yield surface variations of less than 100 nm over 3.1 mm (the length of one side of the crystal-thermistor interface), which is about half the thickness of the Au-Ga pad deposited. There are no obvious indications of stress problems, though the anomalous behavior of BLIP6's P2 sensor may be explained in this way. The temperature achieved during eutectic bonding also serves to anneal the damage to the lattice incurred by implantation and thermally activates migration of the boron ions from interstitial sites to lattice sites.

- *Removal of Implanted Layer*: Prior to deposition of the α-Si/Al-Schottky electrodes, the implanted layer is etched off the detectors except for small islands around the thermistors and the pulser pads. The positions of the thermistors and pulser pads are indicated in Figure 4.1. The etch depth is 1 μm, enough to remove the implanted region. The Run 18 data indicate that removal of a very deep layer ($\sim$ 30 μm) has no effect on the observed tritium rate (*c.f.* Chapter 8), and so a deeper etch is not performed. The regions to be left unetched (the thermistors, especially) are masked with picene wax [129]. The etch is performed by placing the detector in a small, teflon bowl, pouring the etch in, gently agitating by hand for the desired time, and then quenching. The detector is then rinsed and blown dry with nitrogen gas that is derived from liquid-nitrogen boiloff and is thus radon-free. Because of this manual procedure, the etch depth is not uniform – typically, one sees a sunflower-type pattern. Fine scratches appear on the detector, presumably due to small pieces of abrasive material present in the etch solution. Clearly, this is not an ideal procedure, and the techniques used for fabrication of ZIP detectors at the Stanford Nanofabrication Lab are far superior in these respects.

- *Electrode Deposition*: The picene wax mask is removed and a new mask, with slightly less coverage of the thermistors and more coverage of the pulser pads, is painted on prior to deposition of the electrodes. The slight reduction of the masked region around the thermistors allows the deposited material to overlap the implanted regions to make electrical contact. The increase in mask size over the pulser pads ensures they remain electrically isolated from the

Detector/	α-Si		Al	
Face	thickness (nm)	composition	thickness (nm)	composition
BLIP3 P	22.0	Si (4.5% Ar)	15.0	$Al_{0.70}O_{0.30}$
BLIP4 P	20.0	Si (4.5% Ar)	13.0	$Al_{0.70}O_{0.30}$
BLIP5 P	19.0	Si (5% Ar)	11.5	$Al_{0.75}O_{0.25}$
BLIP6 P	19.0	Si (4.5% Ar)	11.0	$Al_{0.65}O_{0.35}$
BLIP3 Q	21.0	Si (4.5% Ar)	11.5	$Al_{0.70}O_{0.30}$
BLIP4 Q	20.0	Si (4.5% Ar)	11.0	$Al_{0.65}O_{0.35}$
BLIP5 Q	19.0	Si (4.5% Ar)	9.0	$Al_{0.70}O_{0.30}$
BLIP6 Q	19.0	Si (5.5% Ar)	11.0	$Al_{0.80}O_{0.20}$

Table 4.2: Thickness and composition of the electrode layers for BLIPs 3–6. These values determined by analysis of witness samples, which were not always exactly at the same height as the detectors. The α-Si layers are too thin to determine their hydrogen content. Table provided by Kin Man Yu.

deposited material. The α-Si and Al depositions for the Run 19 BLIPs were performed using Eugene Haller's LBNL Building 2 sputtering machine using the same Si target and hydrogen canister as used for fabrication of ABL1 and other test devices in Paul Luke's sputtering machine in LBNL Building 70. The effect of changes in the purity of the Si target or in the hydrogen atmosphere had not yet been explored, so extreme efforts were made to ensure the deposition environment was as close to that of the test devices as possible. The sputtering parameters are as follows [130]:

Atmosphere : 10 mT Argon, 18% H_2

Machine : RF sputterer

RF power : 400 watts

Time : 40 seconds

Nominal deposition rate : 26.7 nm/minute

Deposition target : cast silicon

Witness samples were deposited at the same time as the detectors and their thicknesses and compositions were determined by Kin Man Yu at LBNL's Ion Beam Analysis Facility; the results are in Table 4.2. The deposition thickness and compositions are uniform in an absolute sense. However, because it is not known how dependent the band structure and Schottky barrier are on the exact materials parameters, little more can be said.

- *Electrode Definition*: Electrode definition consists of defining breaks between the top and bottom electrodes at the equator and between the inner and outer electrodes on the ionization size of the detector. Since the phonon side is only a ground plane, it requires no inner-outer break. In the case of BLIPs 3 through 6, the inner electrode is deformed to move the thermistors and pulser pads to the outer electrode, as is shown in Figure 4.1. This was done because implanted material was left in both locations (intentionally, as described above) and

the Run 18 data alone did not rule out the implantation process as the source of tritium contamination (see Chapter 8 for a discussion of tritium contamination).

For prior BLIPs, picene wax was used as a mask during electrode-definition etching steps. Elimination of picene wax from the procedure is desirable because it may be a source of surface contamination — picene wax is a hydrocarbon derived from coal tar and petroleum, and thus presumably can leave ^{14}C behind if mask removal is imperfect. The use of picene wax in steps before electrode deposition is allowed because cleaning-etch steps (not listed here) are performed just prior to electrode deposition. Such etching cannot be done after electrode deposition because of the delicacy of the α-Si and Al layers.

In this vein, John Emes and I attempted to develop a shadow-mask technique to define the electrode breaks. The top-bottom break was defined by placing a mask with a circular hole on the detector. The hole is slightly smaller than the detector diameter at the equator and the mask rests on posts on a detector-mounting jig to place it at the correct height. The mask defines a height below which material cannot be deposited. A break width of 3 mm was used. Previous BLIPs had used a 1-mm break, but it was feared that edge definition could be a problem with the shadow mask. It was not believed that a 3-mm break would be problematic for photon misidentification; this belief has proven to be incorrect, as is discussed in Section 4.5.2. An attempt was made to define the inner-outer break using shadow masks, but shadows blurred the edges and left enough conductive material to short-circuit the 0.5-mm gap. The inner-outer break was, in the end, produced by masking with picene wax and etching. A first etch was done to remove only the Al; it was hoped that leaving the α-Si in place would help protect the bare germanium and ensure good charge collection in the break. It was observed that even the α-Si caused a short-circuit and it too was etched away from the inner-outer break. This series of repeated etchings was performed on BLIP 3, the first detector of the series, and is presumably the cause of the surface contamination seen in the Run 19 data (see Chapter 7).

- *Contact Pad Definition*: A final gold deposition is necessary to make contact pads for the ionization electrodes, the pulser pads, the detector ground plane, and the thermistors themselves — wirebonding directly to the very thin α-Si/Al-Schottky electrode is expected to be prone to failure. Gold pads (with a thin Pd sticking layer) are deposited by thermal evaporation through a shadow mask. Wirebonds are made between the detector ground-plane contact pad and the gold in contact with the thermistor-detector interface to ensure connection of the implanted island under the thermistor to ground. The pulser-pad region is intentionally left electrically disconnected from the phonon-side ground plane. It is grounded by the pulser electronics when the pulser is not in use.

It is important to note the extensive use of picene wax as an etch mask. As noted above, picene may contain ^{14}C, and thus any residual wax may be a surface contaminant. Incomplete mask removal may explain the large surface contamination seen on BLIP3, discussed in Chapter 8.

4.4 *In Situ* Noise and Energy Resolution

Though I will discuss the pulse-fitting algorithms in detail in Chapter 6, it is useful to present here results on detector noise and energy resolution. The numbers are summarized in

Table 4.3.

4.4.1 Baseline Resolution

To establish baseline noise performance, a 5% fraction of random triggers are taken during normal data acquisition. These "randoms" are not truly random in the sense that they are generated by the DAQ computer at only certain times during the DAQ cycle, but they are random enough for the purpose of determining noise spectra and baseline energy resolution. These traces are fit as if they were real events; the width of the distribution of fitted energies is the baseline energy resolution. Figures 4.21 and 4.22 show histograms of the noise fits for the phonon and summed-ionization energies for the entire data set. The offset from 0 in the phonon-channel histogram is due to the time-offset search performed in the fit; this search is discussed in detail in Chapter 6. The search has little effect on the width of the distribution because the search range is small compared to the phonon-pulse length. Note also the tails on the phonon-channel histograms; these presumably are due to non-Gaussian thermal disturbances. The ionization-energy histograms are, in contrast, very clean, reflecting the stability of the ionization noise.

The resolutions shown here are significantly worse than the 0.375-keV FWHM obtained for BLIPs 1 and 2 in Run 18 (average of both phonon sensors) [8]. Significant work was done to maximize the pulse height with respect to the thermistor-bias currents and refrigerator temperature; the $\sim$ 160-nV/keV pulse height for BLIPs 3 through 6 is the same as was obtained in previous runs. The refrigerator was operated between 17 and 19 mK, lower than in previous runs. This was done because modest pulse-height gains were seen (5–10%) at lower temperatures [131]. However, little care was taken to check the thermistor noise; since the carrier amplitude was not seen to vary significantly, it was assumed that the resistance and thus the Johnson noise were also not varying. This may have been incorrect: the resistance may have been increasing due to decreased Joule heating, with the bias current being increasingly shunted by the thermistor and parasitic capacitances, producing a relatively slow variation in the carrier amplitude. However, the noises observed are ≈ 5 nV/$\sqrt{\text{Hz}}$ (see Figure 4.7), which would require a 25 MΩ operating resistance if due to Johnson noise! This, unfortunately, is not impossible: a 80-pF gate capacitance corresponds to a resistance of 2 MΩ at 1 kHz, similar to what is observed (see Table 4.1). It is hard to believe, however, that such a large operating resistance would not have shown itself in the pulses somehow. For example, given the strong dependences of the thermal couplings on temperature, the pulse time constants should have been much longer if the thermistors were so cold. In the end, the poor resolution was not investigated because the ionization-energy noise dominates the recoil-energy resolution because of the large Neganov-Luke correction at 6-V ionization bias. Lower phonon noise may have been useful in better resolving photon lines in the phonon-energy spectrum, however.

To investigate the possibility of correlated or anticorrelated noise in the separate ionization channels, the inner- and outer-electrode energy histograms are displayed in Figure 4.23. As remarkable at it may seem, the summed-ionization-energy widths are fully consistent with no correlation between the inner- and outer-electrode energies, as can be seen by comparing the summed-ionization widths with the quadrature sums of the individual channel widths, all given in Table 4.3. This is explained by the use of optimal filtering. Optimal filtering and its application to these detectors are discussed in detail in Appendix B and Chapter 6. In performing the energy fit, optimal filtering weights frequency components by the inverse square of the expected noise in

energy	BLIP3	BLIP4	BLIP5	BLIP6
Phonon Energy				
baseline	0.66	0.66	0.54	0.68
baseline (P1)	0.85	0.80	0.66	0.80
expected	0.60	0.56	0.73	0.63
opt. filt.	0.38	0.37	0.45	0.43
baseline (P2)	0.82	0.85	0.68	0.92
expected	0.62	0.59	0.79	0.75
opt. filt.	0.37	0.41	0.47	0.58
31.2 keV short	1.20	0.89[1]	1.32	1.29
31.2 keV long	1.62	2.36	1.44	1.44
1533 keV[2]	71	68	64	61
Summed Ionization Energy				
baseline	0.85	0.97	0.94	1.11
10.4 keV short	0.68	1.06	1.06	1.22
10.4 keV long	1.60[3]	1.67	1.32	1.60
511 keV	14.8	17.7	11.1	13.0
Inner Ionization Energy				
baseline	0.59	0.64	0.61	0.85
opt. filt.	0.51	0.55	0.59	0.73
10.4 keV short	0.82	0.64	0.64	0.73
10.4 keV long	1.22	1.41	0.97	1.30
Outer Ionization Energy				
baseline	0.66	0.80	0.75	0.80
opt. filt.	0.67	0.57	0.56	0.60
10.4 keV short	0.57	0.52	1.08	1.13
10.4 keV long	1.51	1.67	1.08	1.30

Table 4.3: Energy resolutions for phonon energy and summed, inner-electrode-only, and outer-electrode-only ionization energies. Energy resolutions are FWHM in keV as given by Gaussian fits. Numbers are given separately for the "short" and "long" data sets (defined in the text) for the gallium line to indicate the importance of drifts. The "opt. filt." resolutions are those expected for optimal filtering. For phonons, they give the best resolutions one could expect given the noise. For ionization, they should match the observed resolutions. The observed widths are slightly higher. The "expected" resolutions for phonons are the resolutions expected to be given by the non-optimal fit that is performed. The fitting algorithms are discussed in Chapter 6. The formalism for calculating expected resolutions for optimal and non-optimal filtering is presented in Appendix B. Notes: [1]: artificially decreased by not including satellite peak in fit. [2]: all peaks are too low in phonon energy by about one half-width; BLIP4 is low by 100 keV, almost 1.5 FWHM. [3]: likely is slightly increased by events to left of peak; picking FWHM off histogram gives 1.40.

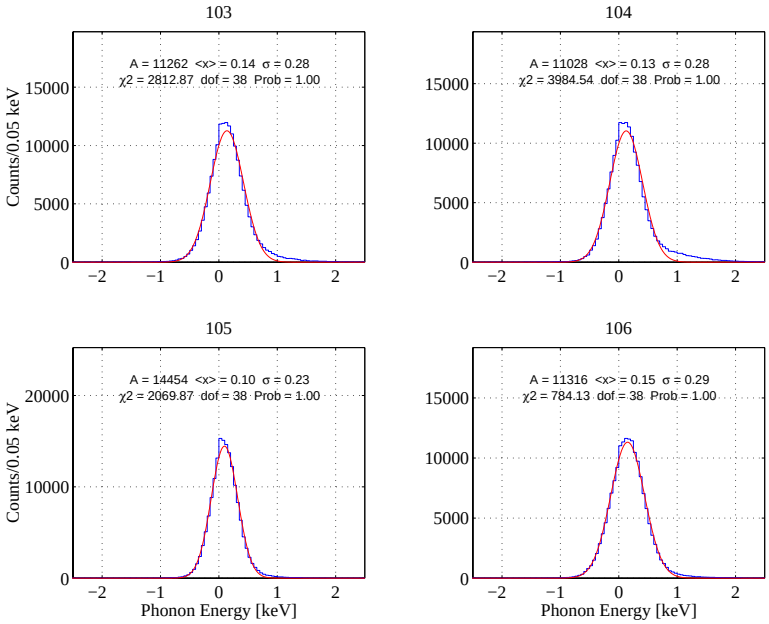

Figure 4.21: Histograms of phonon energies (mean of P1 and P2) for random triggers.

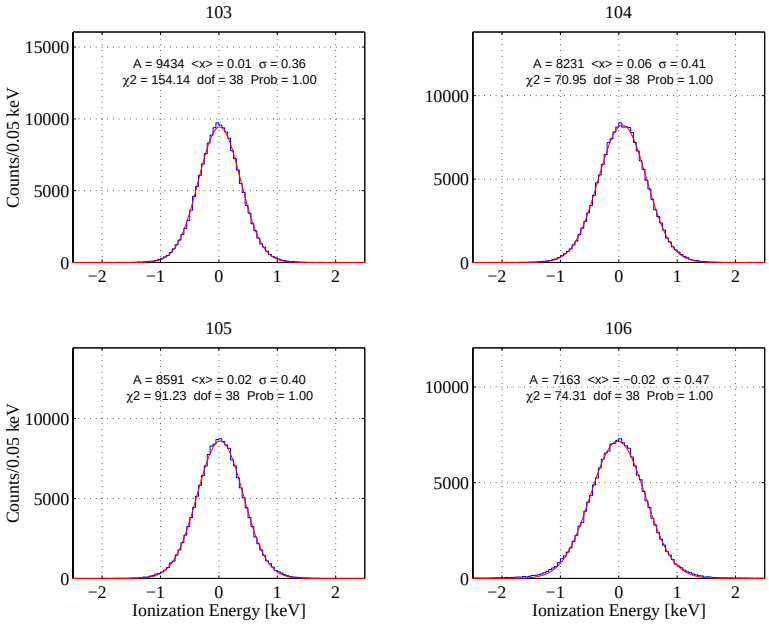

Figure 4.22: Histograms of summed-ionization energies for random triggers.

the frequency component. Correlated and anticorrelated noises tend to appear as spectral lines because they usually arise from pickup or feedthrough of a strong, monochromatic signal. In the BLIP detectors, the primary such spectral lines are cross-talk from the 1-kHz phonon bias, 60 Hz, and harmonics of both. By removing these frequency components, optimal filtering leaves only the FET noise and ionization-bias-resistor Johnson noise, which are truly not correlated between the two channels. In Table 4.3, the expected resolutions, based on optimal filtering, are also shown. The observed resolutions are somewhat high, typically 10% to 20% and as high as 30%. This has not been investigated.

Correlations between the noises in the two phonon sensors may be investigated in a similar manner. In principle, the use of two sensors improves the resolution by a factor of $\sqrt{2}$ because their individual noises are dominated by Johnson noise, which is uncorrelated. In the past, significant correlations and anticorrelations have been observed. The baseline noise histograms for the two sensors of each detector are shown in Figure 4.24. Table 4.3 shows the baseline resolutions for the two sensors separately. The noise on the mean of the two sensors is generally about 10% higher than expected assuming no correlation: there is a small amount of correlated noise in the two sensors. The phonon traces are not optimally filtered for practical reasons, so the presence of some correlation may not be surprising. The noise spectra of the phonon channels exhibit no spectral lines. A more likely cause is low-frequency thermal effects, perhaps of a kind similar to those which produce the high-energy tails on the phonon-channel noise distributions. Since the effect is only about 10%, it has not been investigated.

Expected resolutions for the individual phonon channels for the non-optimal fit that is performed and for optimal filtering are shown in Table 4.3. The observed resolutions are typically higher than expected by about 30%. This has not been investigated. The expected optimally filtered resolutions are about 30% lower than the expected resolutions for the non-optimal fit. Not much would be gained by using optimal filtering for the phonon channels. As is discussed briefly in Chapter 6, there are technical reasons that optimal filtering for phonons is inconvenient and difficult.

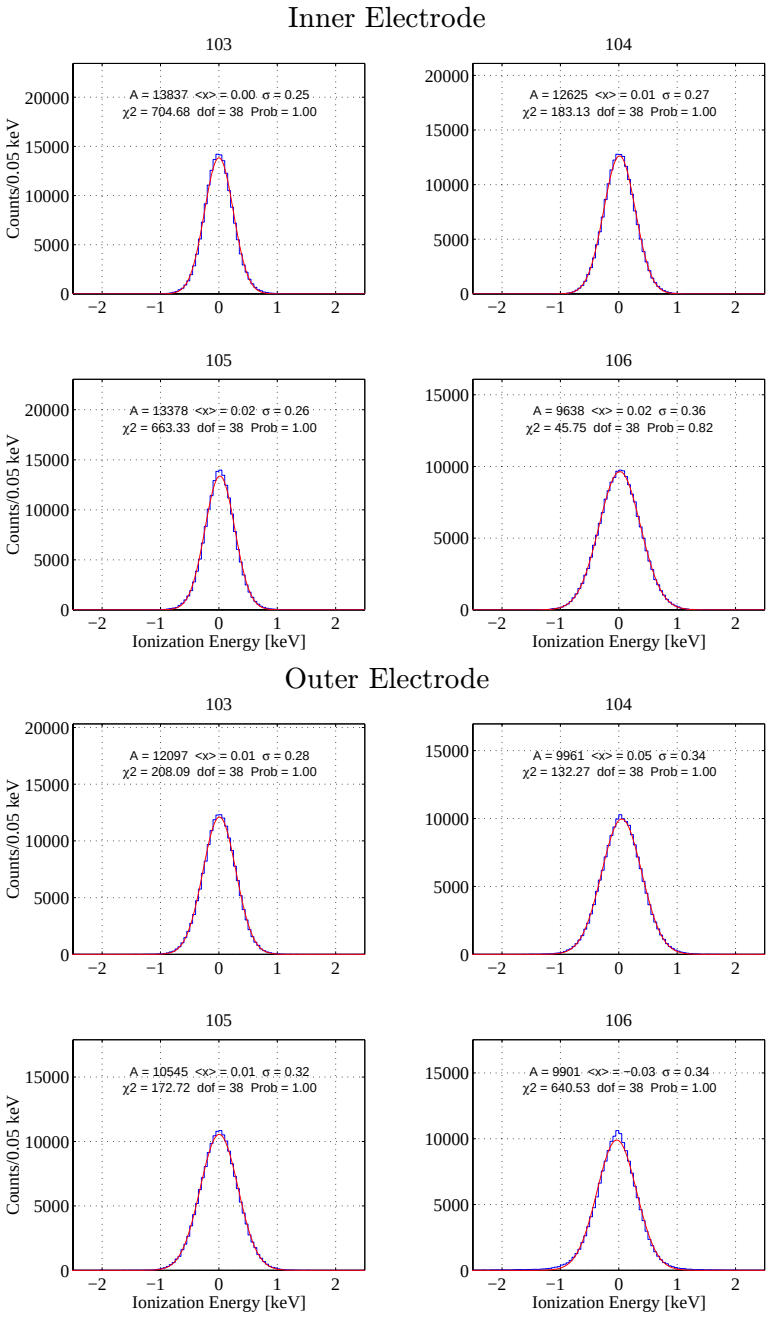

Figure 4.23: Histograms of inner- and outer-electrode ionization energies separately. Top plots: inner electrode. Bottom plots: outer electrode.

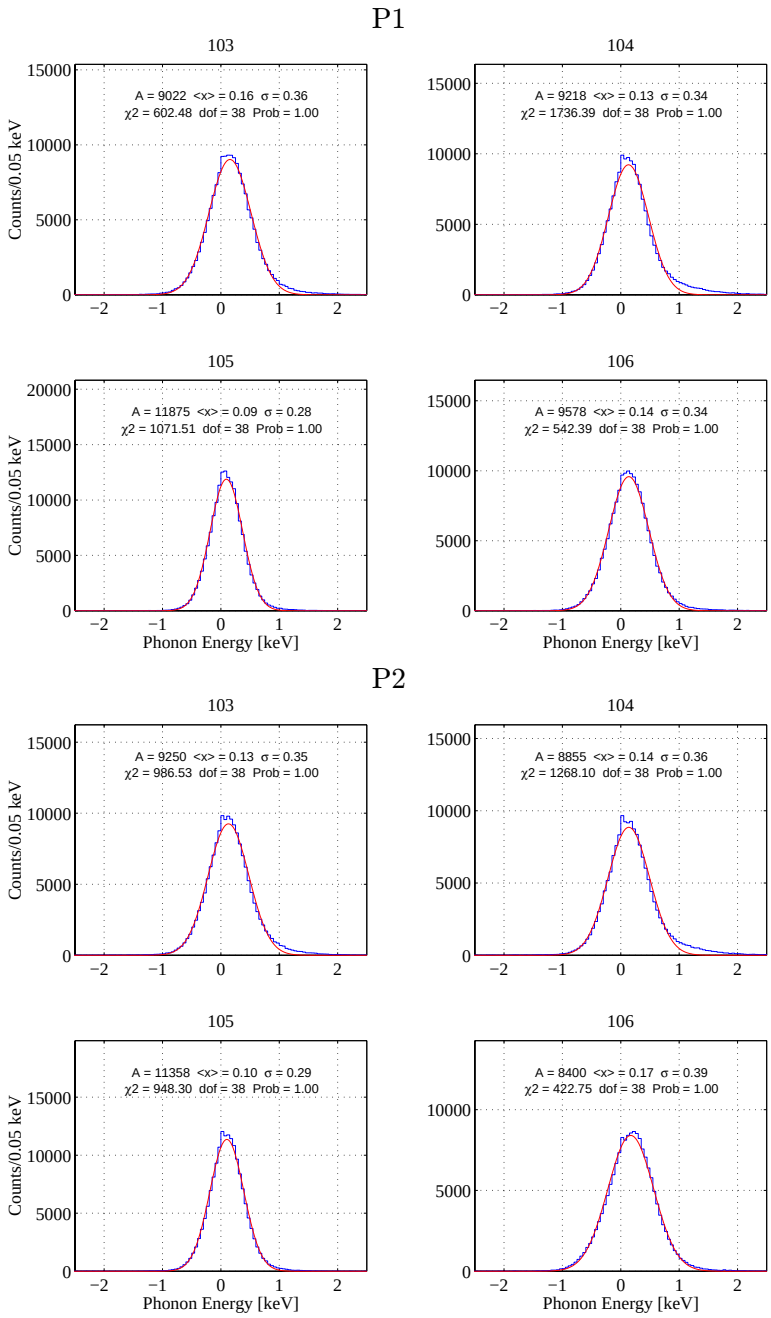

Figure 4.24: Histograms of individual phonon-sensor energies. Top plots: P1. Bottom plots: P2.

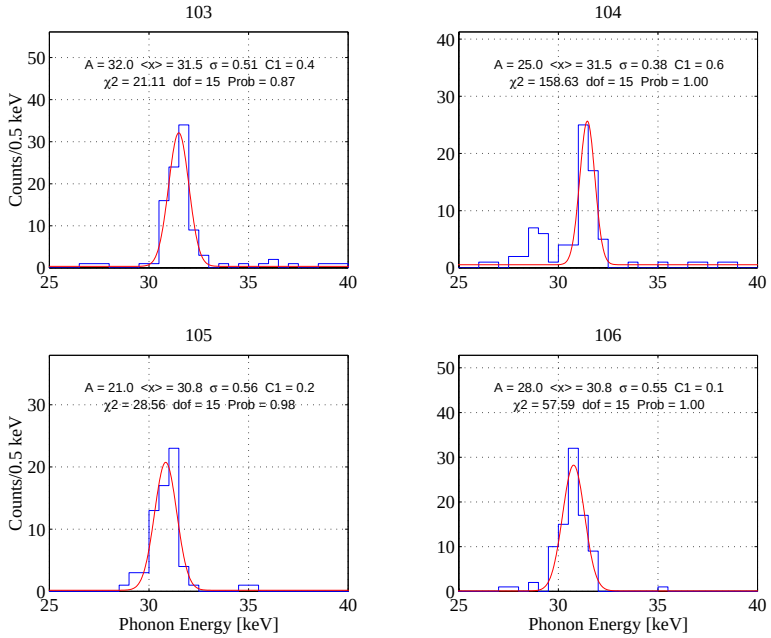

Figure 4.25: Gallium line in phonon energy. The plot is for a 5-day period following the April neutron calibration. The fourth fit parameter is an arbitrary flat background. BLIP4 exhibits short term drifts that displace several of the events from the central distribution, as can be seen by the satellite peak at lower energy. The fit ignores these events. When a larger data set is used, the distribution widens and the satellite peak joins the main distribution.

4.4.2 Low-Energy Resolution

The low-energy resolution (at nonzero energy) can be determined using the 10.4-keV Ga X-ray, emitted by daughters of ^{68}Ge and ^{71}Ge produced cosmogenically and by inelastic scattering during neutron calibrations. These isotopes are created uniformly through the crystal. Being internal, they are not affected by the dead layer.

Figures 4.25 and 4.26 show histograms of the gallium line in phonon energy and ionization energy. Note that the 10.4-keV line appears at 31.2 keV in phonon energy because a 6-V ionization bias is used; 20.8 keV of Neganov-Luke phonons are generated. Figure 4.27 show separate plots for inner- and outer-electrode-contained events, discarding the other electrode's signal. This gives the noise of each electrode separately. These plots are made for short periods, only about 5 days, immediately following a neutron-calibration data set. This exploits the high rate of gallium X-rays provided by ^{71}Ga, a daughter of ^{71}Ge, which is produced by inelastic neutron scattering and has a 11.4-day half-life. By considering only a short period, long-term drifts and calibration errors are removed, giving the energy resolution as limited by the detectors and electronics and short-term drifts. Table 4.3 lists resolutions for all the detectors and channels. The ionization energy resolutions are not significantly degraded relative to the baseline resolutions, while the phonon energy widths are about twice as large as the baseline phonon widths.

The energy resolution is degraded by uncorrected long-term drifts. As discussed in Chap-

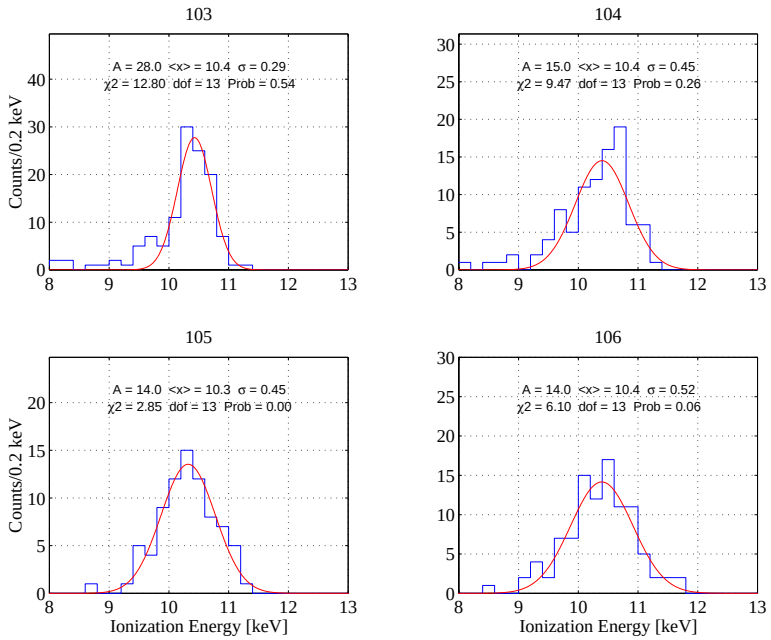

Figure 4.26: Gallium line in ionization energy. The time period is restricted as for Figure 4.25.

ter 6, drift corrections are performed, but they are clearly imperfect. Figures 4.28, 4.29, and 4.30 show the energy spectra for the gallium line for approximately the second half the Run 19 data set. The data set used in the previous plots is removed to prevent it from dominating and artificially improving the resolution. Table 4.3 lists these long-term resolutions also. The first half of the Run 19 data set exhibits poorer resolution because work was being done to improve the electronics, optimize the phonon bias, and minimize cross-talk of the phonon bias to the ionization channels; I do not believe the first half of Run 19 is indicative of what can be achieved over extended data runs.

Drift does not appear to degrade the phonon-energy resolution significantly ($\sim$ 10% at 31.4 keV); this may be more accurately interpreted as the fact that thermal drifts tend to occur over fairly short time periods (few hours to 1 day), much smaller than the time period included in the "short" data set. The ionization-energy resolution is significantly degraded by drift, yielding non-Gaussian distributions whose widths are between 20% and a factor of 3 larger than observed in the short data set! This is extremely disappointing and difficult to understand. The ionization-amplifier feedback components are mounted at 20 mK and thus are presumably very stable. There was only one thermal cycle to room temperature during the "long" data set, so one would expect bimodal distributions if changes in these components were responsible. The FETs also are hard to blame: since their 4 K heat sink is so much colder than the FETs themselves, they are insensitive to variations in the 4 K layer temperature. Their temperature may change if their power dissipation changes, which is set by the room-temperature amplifier circuit. It is possible that the use of marginally sampled ionization traces (see Chapter 6 for a discussion) makes the fitting algorithm especially sensitive to minor variations in pulse shape, but the pulse shapes must actually change

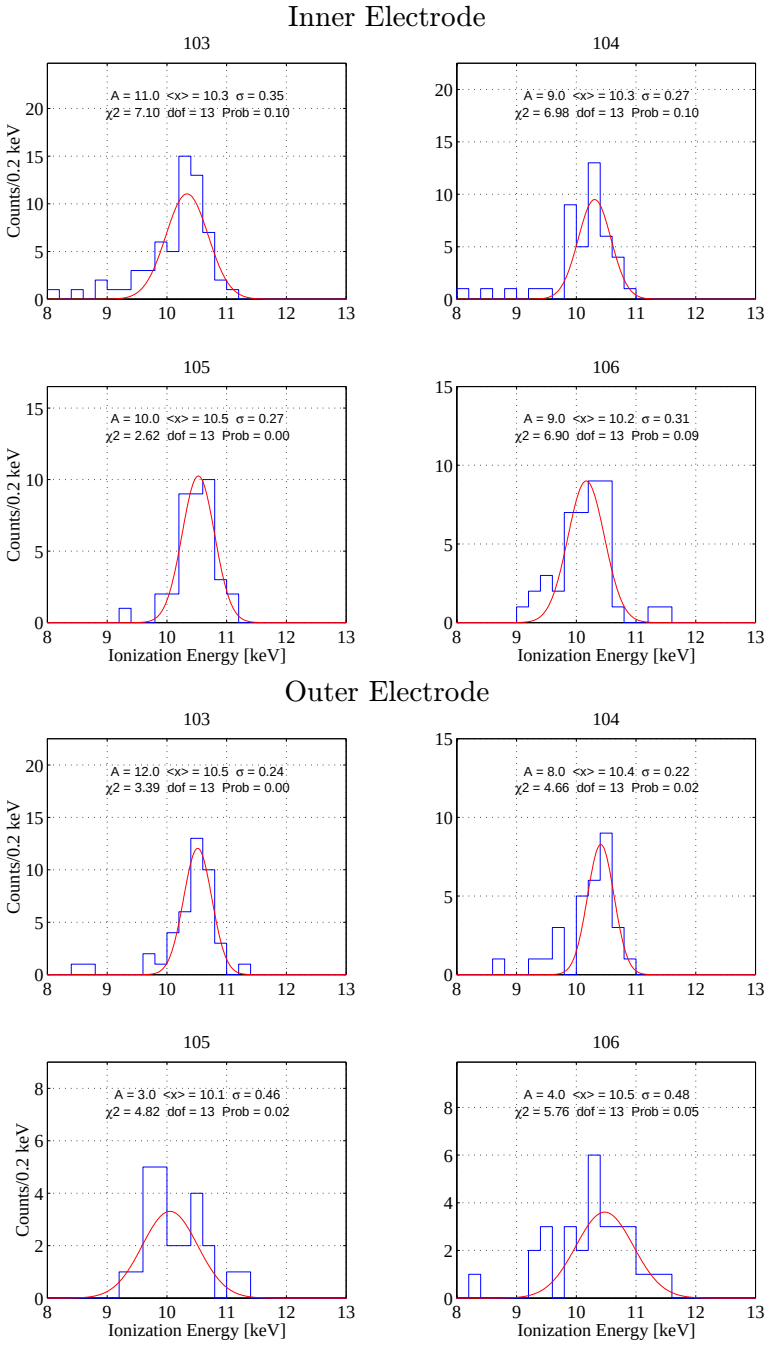

Figure 4.27: Gallium line in ionization energy for inner- and outer-electrode-contained events; for these events, only the energy of the corresponding electrode is used (though the fit is done jointly). The time period is restricted as for Figure 4.25.

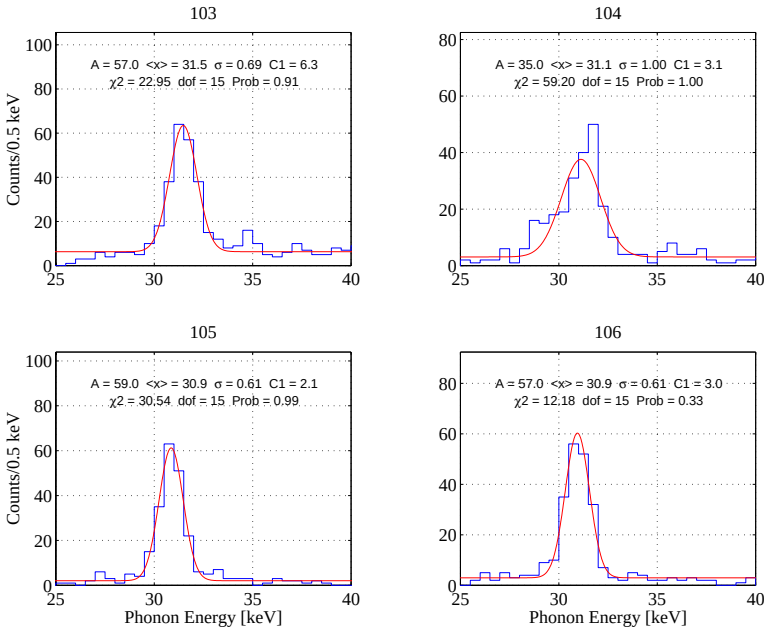

Figure 4.28: Gallium line in phonon energy over the second half of the run. The fourth fit parameter is an arbitrary flat background.

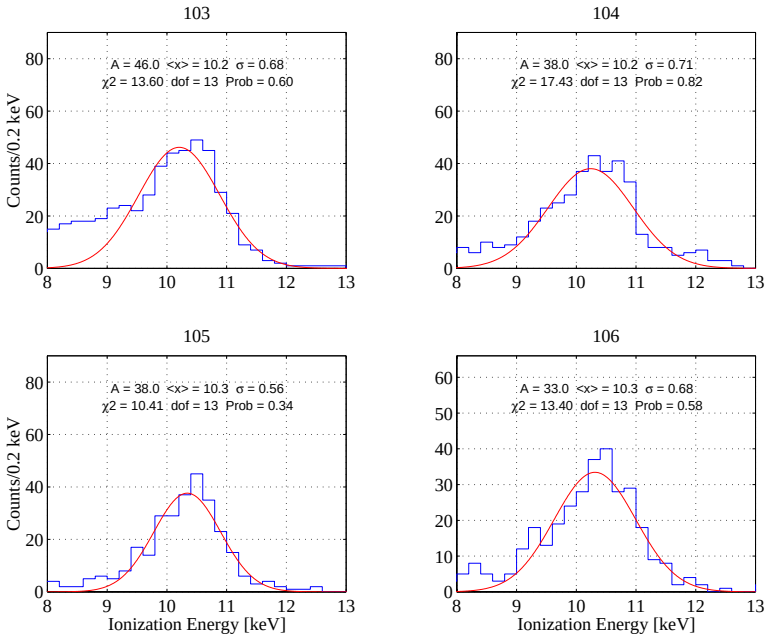

Figure 4.29: Gallium line in ionization energy. The time period is the same as for Figure 4.28.

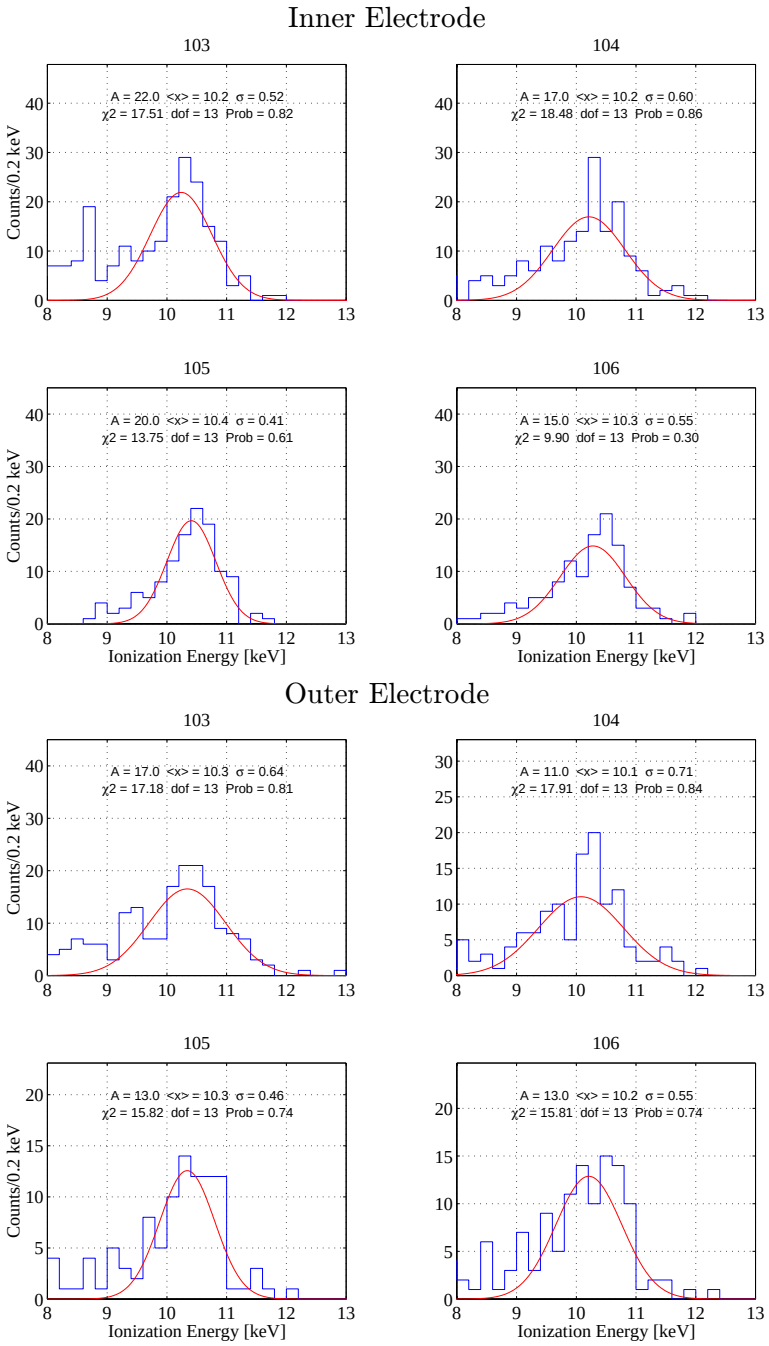

Figure 4.30: Gallium line in ionization energy for inner- and outer-electrode-contained events. The time period is the same as for Figure 4.28.

for this to occur. The room-temperature front-end electronics are thus the most likely cause.

The main lesson to be learned for future data sets is that stability is as important as intrinsic energy resolution: the excellent energy resolutions achieved on short timescales are significantly degraded by long-term drifts. Do these drifts significantly degrade the experiment's WIMP sensitivity? They do not affect the threshold: the baseline noises are quite stable, as discussed in Section 4.4.1. However, by increasing the ionization noise, which dominates the width of the nuclear-recoil band, such effects do force the use of a larger acceptance band and degrade the separation of nuclear recoils from surface events and photons. Drift is a difficult problem to solve because it is entirely due to the imperfect nature of the measurement electronics, but it also seems to be one that should draw significant attention.

4.4.3 High-Energy Resolution

Though not particularly important for dark-matter analysis, the detectors' high-energy resolution is important for diagnosing electromagnetic backgrounds. In general, the photon background below 100 keV in low-background experiments arises from Compton scattering of much higher-energy photons emitted by decays of radioisotopes. The rate of full-absorption events of these high-energy photons indicates which isotopes are important and can be used to predict their Compton-scattering low-energy contribution. Alpha particles from radioisotopes can also be useful; for example, as mentioned in Chapter 3, ^{210}Po, a daughter of ^{210}Pb, emits a 5.3-MeV α. Detection of this α constrains the presence of ^{210}Pb, which beta decays with endpoints of 16.6 keV and 63.1 keV: it is a possible source of low-energy electrons.

The high-energy resolution is determined using the 511-keV positron-annihilation photon line. This line appears at 1533 keV in the phonon channel because of the 6-V ionization bias used. Figure 4.31 shows spectra around the 511-keV line. For simplicity, only the phonon and summed-ionization-energy spectra are shown. Unfortunately, as is seen in the figure, the resolution degrades at such high energies. The phonon signal saturates the digitizers by this energy, making it more difficult to do accurate fits; pulse-shape changes also contribute. As is discussed in Chapter 6, the ionization-trace-fit-χ^2 distribution indicates the ionization pulse shape is also modified at such high energies, though saturation does not occur. With more work, and possibly a second set of digitized traces with lower gain to prevent saturation, it should be possible to improve the resolution. For example, the Heidelberg-Moscow experiment reports a FWHM resolution of 2.37 keV at 727 keV for their 2.76-kg germanium-diode ionization-mediated detectors [84].

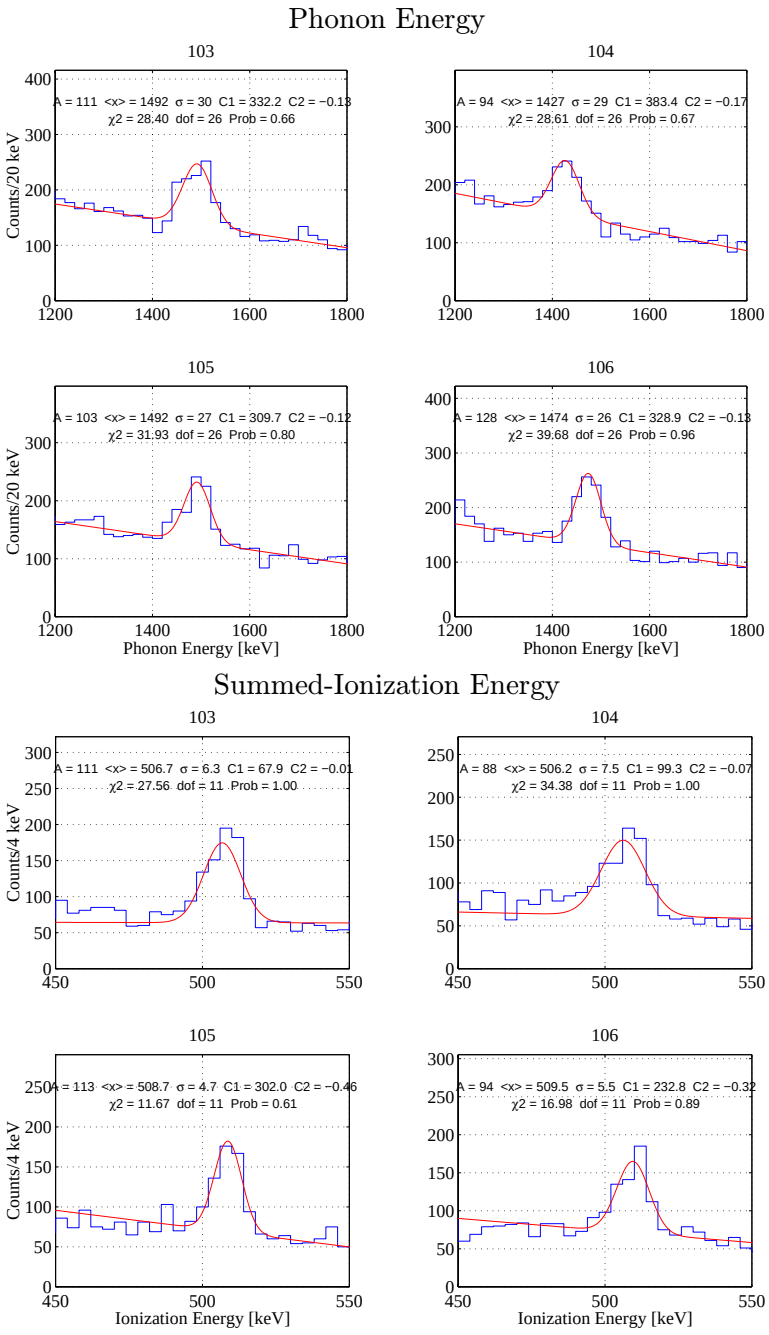

Figure 4.31: 511-keV positron-annihilation line in phonon energy and summed-ionization energy. The fourth and fifth fit parameters are an arbitrary background with a flat and a linear component. The time period is the same as for Figure 4.28.

4.5 *In Situ* Nuclear-Recoil-Discrimination Performance

Three quantities summarize the intrinsic performance of CDMS detectors with regard to dark-matter detection: nuclear-recoil acceptance α, photon misidentification β_γ, and electron misidentification β_β; these parameters were defined generically in Chapter 2. The first two numbers are established *in situ* by exposing the detectors to neutron and photon sources placed external to the cryostat. Electron-response calibration is best performed at a test facility using an electron source placed adjacent to the detector. In practice, the contamination source present on BLIP3 (see Chapter 7) serves as a source of electrons incident on BLIP4 that can be used as a low-statistics electron calibration. In this section, I describe how the three parameters α, β_γ, and β_β are determined from the *in situ* neutron, photon, and electron calibrations.

As noted earlier, the photon-misidentification parameter is higher for the outer electrode than for the inner electrode, presumably due to the large top-bottom electrode gap. To study this, three different ionization-partition cuts are defined. "cQinOnly" refers to events fully contained in the inner electrode; no signal is observed in the outer electrode. "cQoutOnly" conversely requires complete outer-electrode containment. "cQShare" requires the event be shared: signal is observed in both the inner and outer electrodes. (The "c" prefix indicates that the name refers to a cut.) The exact definition of these cuts is discussed in Chapter 7. Finally, the neutron- and photon-calibration data sets were taken at 6-V ionization bias. Because it is uniformly distributed over the run, the electron calibration contains some data taken at 4 V.

4.5.1 Nuclear-Recoil Acceptance

To determine the position of the nuclear-recoil-acceptance region in ionization as a function of recoil energy, two neutron calibrations were performed during Run 19, one in April, approximately midway through the run, and a second in September, at the end of the run. I discuss the cuts used on these data sets in Chapter 7. In this section, I describe the method used to define the nuclear-recoil-acceptance regions and to calculate their acceptance.

A complication is introduced by a discrete shift in the apparent position of the nuclear-recoil band at the time of a power outage on April 3, 1999, just after the first neutron calibration. Efforts to correct this are discussed in Chapter 6; they are not completely successful. The most straightforward method to deal with this is to define different nuclear-recoil bands for the pre- and post-April 3 data sets using the two neutron calibrations. Chapter 7 contains a discussion of checks made to ensure the nuclear-recoil acceptance is stable in spite of this shift.

Another minor complication is that a small part of the Run 19 data was taken at 4-V ionization bias. It is conservative to use 6-V neutron-calibration data: the width of the nuclear-recoil band increases because of the increased Neganov-Luke correction at 6 V, so the acceptance of the band in the 4-V data set is underestimated slightly. This should be a small effect and no attempt is made to correct for it.

Definition of Nuclear-Recoil-Acceptance Region

To define the nuclear-recoil-acceptance region, it is necessary to determine the position and width of the nuclear-recoil band in ionization energy as a function of recoil energy. Figures 4.32, 4.33, 4.34, 4.38, 4.39, and 4.40 show plots of ionization yield (Y_R = ionization energy/recoil energy)

vs. recoil energy for the four detectors and the two neutron calibrations. Plots are shown for all events and for cQinOnly and cQoutOnly events separately. The neutron and photon bands are quite clear.

A "cleaning cut" (named "cClean") that encloses the neutron band and excludes data that are clearly not neutrons is made. Events passing cClean are binned in recoil energy. An initial guess at the nuclear-recoil line (the center of the nuclear-recoil distribution as a function of energy) is made by eye. The deviation (in ionization energy) of the data from the nuclear-recoil line is calculated. In each recoil-energy bin, the mean and standard deviation of the data are calculated and data 2 standard deviations from the mean are discarded. The process is repeated until the standard deviation stabilizes. A Gaussian fit to a histogram of the remaining data in each recoil-energy bin is performed. The mean and standard deviation of the distribution are returned by the fit; uncertainties are calculated from the covariance matrix. The ionization energy is reconstructed from this mean deviation for each recoil-energy bin and a χ^2 fit to a power-law function of recoil energy is done, using the uncertainties provided by the histogram fits. In performing the fit, the bins below 10 keV are discarded; the fit is constrained at low energy by the requirement that it intersect the origin. This fit is fed back into the beginning of the above process and the fit redone to ensure that it is robust.

Bins below 10 keV are discarded in the fit because, in this energy range, the nuclear-recoil band intersects the ionization-search threshold. As is discussed in Chapter 6, for events in which the detector trigger came from the phonon channel, a search is performed in the ionization trace for an ionization signal; the ionization trace is long enough that if there is an ionization signal, it will be found. If no ionization signal is present, a noise energy is found. Because of the search, the distribution of the noises is offset from zero; this is discussed in Chapter 6 and Appendix B. Thus, only an upper limit on the ionization is available. It is not possible to determine whether such an event is a low-energy nuclear recoil or a low-energy electromagnetic event with suppressed ionization (an electron incident on the dead layer or a photon incident on the top-bottom electrode break), and, regardless, the ionization value is not physical. Clearly, such events should not be included in the fit. However, the distribution left after applying this threshold is a truncated Gaussian whose mean and width are not accurately fitted by the simple fitting algorithm, especially at energies where the truncation leaves less than half the Gaussian. Therefore, bins below 10 keV are discarded; the power law must go through the origin, constraining the fit in this region. Points above 100 keV are also not used in the fit because there are very few neutrons at such high energies. These fits are displayed in Figures 4.35, 4.36, 4.37, 4.41, 4.42, and 4.43. Though the fit is performed in ionization energy vs. recoil energy, it is more convenient to display it in ionization yield.

The width of the nuclear-recoil band is, in principle, more complicated. The issue is that the noises in ionization and recoil energy are not independent. Rather than attempt to calculate theoretically the width, it is easier to just lump all the noise into ionization and use the empirically derived standard deviation in each recoil-energy bin calculated by the aforementioned Gaussian fit. These data points are fit to a line and the line is then used to parameterize the width of the nuclear-recoil band in ionization as a function of recoil energy. Table 4.4 displays the fit values for both the center and width of the nuclear-recoil band.

Using the above position and width of the nuclear-recoil band, the acceptance region is defined. Somewhat arbitrarily, it was decided that a 90% acceptance band would be adequate at high energy (> 50 keV), where photon and electron leakage into the nuclear-recoil band are

Channel/	BLIP3		BLIP4		BLIP5		BLIP6	
Data Set	a	b	a	b	a	b	a	b
Nuclear-Recoil Line: aE_R^b								
First Calibration								
Summed	0.200	1.128	0.201	1.132	0.149	1.211	0.162	1.183
Inner	0.176	1.156	0.173	1.173	0.095	1.331	0.143	1.214
Outer	0.225	1.114	0.213	1.124	0.179	1.172	0.150	1.209
Second Calibration								
Summed	0.185	1.143	0.200	1.132	0.216	1.111	0.197	1.131
Inner	0.171	1.167	0.181	1.158	0.194	1.140	0.186	1.145
Outer	0.191	1.141	0.169	1.181	0.210	1.125	0.146	1.214
Nuclear-Recoil Band Width: $aE_Q + b$								
First Calibration								
Summed	0.063	0.52	0.088	0.48	0.081	0.57	0.064	0.54
Inner	0.052	0.43	0.068	0.55	0.089	0.37	0.061	0.49
Outer	0.040	0.40	0.092	0.49	0.046	0.80	0.056	0.62
Second Calibration								
Summed	0.072	0.45	0.076	0.40	0.068	0.50	0.066	0.58
Inner	0.050	0.48	0.077	0.34	0.056	0.45	0.081	0.36
Outer	0.072	0.45	0.076	0.40	0.068	0.50	0.066	0.58

Table 4.4: Fit values for the nuclear-recoil band. The top half of the table gives parameters for a power-law fit to the center of the nuclear-recoil band of the form aE_R^b. The bottom half gives parameters for a linear fit to the standard deviation, in ionization energy, of the band, with the form $aE_Q + b$. The width is parameterized as a function of ionization energy rather than recoil energy because it is dominated by the noise of the ionization-pulse fit. The width can be written as a function of recoil energy using the fit to the band center. Note that the a and b fit parameters have different meaning for the two kinds of fits.

negligible. A nominal 90% acceptance band is given by a region that extends from 1.28σ above to 3σ below the fit to the nuclear-recoil line, with σ given by the corresponding line fit to the nuclear-recoil-band width. The band is truncated from below at the ionization-search threshold because it is not possible to discriminate nuclear recoils from photons or electrons (or even from mechanical impulses) below this value; this threshold is shown by the dashed lines in the plots of the neutron-calibration data. At lower energies, one would like to optimize the high-ionization edge of the band against photon and electron leakage based on calibration data. An electron calibration with enough events is not available. A photon calibration was not available at the time the acceptance region was first defined. Instead, it was defined by optimizing against photon leakage under the assumption the photon leakage distribution was determined only by electronics noise. This is, of course, not true, and yields too large an acceptance band, but it was the best choice available at the time. As is discussed in Section 4.5.2, it turns out the acceptance region defined in this way gives essentially no photon leakage in the calibration data set, so there is no need to further optimize the band. The band so defined is shown in the plots of the photon-calibration data, Figures 4.46, 4.47,

and 4.48. Note that, because the position and center of the nuclear-recoil band shift at the time of the April 3 power outage, the bands shown in these figures are only the post-April 3 bands. In the discussion of the final data set, in Chapter 8, the bands for both the pre- and post-April 3 data are shown.

Efficiency of Nuclear-Recoil-Acceptance Cut

The nuclear-recoil efficiency is calculated as follows. A grid in the plane of ionization energy vs. recoil energy is set up. For each bin in recoil energy, a Gaussian distribution in ionization is calculated using the position and width of the nuclear-recoil band at that energy as determined above. The integral of the Gaussian in each recoil-energy bin is 1. The Gaussian is then zeroed for ionization energies outside the nuclear-recoil-acceptance region. The ionization-search threshold is also applied, zeroing the Gaussian below this threshold. To calculate the nuclear-recoil efficiency at a recoil energy E_R, the modified Gaussian is integrated along ionization energy.

It is also possible to calculate the nuclear-recoil efficiency empirically. The cClean cut removes all other particles, leaving a sample of neutrons only. The data are binned in recoil energy and the fraction of events accepted in each recoil-energy bin is calculated. To calculate the efficiency of the combined inner-electrode-containment and nuclear-recoil-acceptance cuts, the fraction of events that are simultaneously accepted by the cQinOnly and nuclear-recoil cuts is calculated. This approach is valid because neutron-scattering events are distributed approximately uniformly through the detector; the neutron mean free path is large compared to the detector thickness. The assumption of uniformity is not exactly correct; however, since the normalization of the calculated efficiency is a free parameter, this assumption does not affect the comparison of the calculated and empirical efficiencies. A correction for the lack of uniformity is made later. A similar calculation may be done for the cQShare cut. At low energy, this empirical estimate is a lower limit. To calculate the efficiency conservatively, cClean accepts noise-charge events at energies where the nuclear-recoil band merges with the noise-charge band (< 10 keV). Not all noise-charge events are neutrons, so the "raw" number of nuclear recoils is overestimated and the efficiency underestimated.

Figures 4.44 and 4.45 show the results of the efficiency estimates for each detector with no containment cut and with the cQinOnly cut. Since it is not possible to calculate the acceptance of the inner-electrode cut from first principles — this relates to the solid-state physics of the size of the ionization cloud, a matter of some debate — the calculated efficiency is normalized to match the empirical efficiency at high energy. The calculated and empirical curves are grossly similar, though the empirical efficiency is systematically low below 10 keV, as expected. In later analysis, a threshold of 10 keV is imposed because surface-event leakage into the nuclear-recoil band is important below this energy, so the accuracy of the efficiency below 10 keV is not important. Note that the figures show the efficiency of only the nuclear-recoil and cQinOnly cuts; other efficiencies (trace-fit-χ^2 cuts, trigger efficiencies) must also be included to give the overall efficiency. An absolute check of the efficiency calculation is given by comparison against the neutron-calibration recoil-energy spectrum provided by the simulation (discussed in Chapter 3); this comparison is made in Chapter 7, where cut efficiencies are discussed in detail.

It is important to note that the efficiency above 100 keV is not well known simply because there are so few neutrons that the position of the nuclear-recoil band cannot be determined. Later analysis is restricted to events below 100 keV for this reason. This does not significantly degrade the

detectors' WIMP or neutron sensitivity because both particle types produce recoil-energy spectra that are approximately exponential with $\langle E \rangle \sim 20$–$30$ keV.

One can determine the fraction of the detector volume accepted by the cQinOnly and the cQShare cuts using the efficiencies of the various nuclear-recoil-acceptance cuts at high energy, where thresholds have no effect. The high-energy efficiency of the full-detector nuclear-recoil-acceptance band gives the normalization; the ratio of the high-energy efficiency of the cQinOnly or cQShare cuts to this normalization gives the physical volume fraction. The high-energy efficiencies are 35%, 22%, and 88% for cQinOnly, cQShare, and all events, yielding 40% and 25% for the volume fractions.

Two straightforward corrections must be made. First, according to the simulation, 6% of neutrons scatter once in each electrode of a given detector, yielding a shared event. Thus, the physical volume corresponding to the cQShare cut is 19%, rather than 25%, of the detector. The simulation shows that the probability of inner-to-outer and outer-to-inner scattering is the same, so it is reasonable to divide this 6% equally between the inner and outer electrodes. This yields physical volumes of 43%, 19%, and 38% for the cQinOnly, cQShare, and cQoutOnly cuts. Second, the simulation shows that the probability of a neutron interacting in the outer electrode is higher than expected from the volume fraction, simply due to self-shielding. The simulation models the detectors as 180 g cylinders with equal inner- and outer-electrode volumes. Of the simulated events, 50% are outer-electrode-contained and 44% inner-electrode-contained. (Note that the simulation does not simulate sharing of the charge cloud — simulated events are shared only if interactions occur in both electrodes.) Thus, the neutron calibration additionally overestimates the outer-electrode/inner-electrode volume ratio by a factor $0.50/0.44 = 1.14$. The inner and outer electrode volumes must each be corrected by a factor $\sqrt{1.14} = 1.07$, yielding 46%, 19%, and 35%. The inner electrode nominally contains 56% of the detector volume, so these numbers are consistent with the cQShare volume being geometrically equally divided between the inner and outer electrodes, as one would expect.

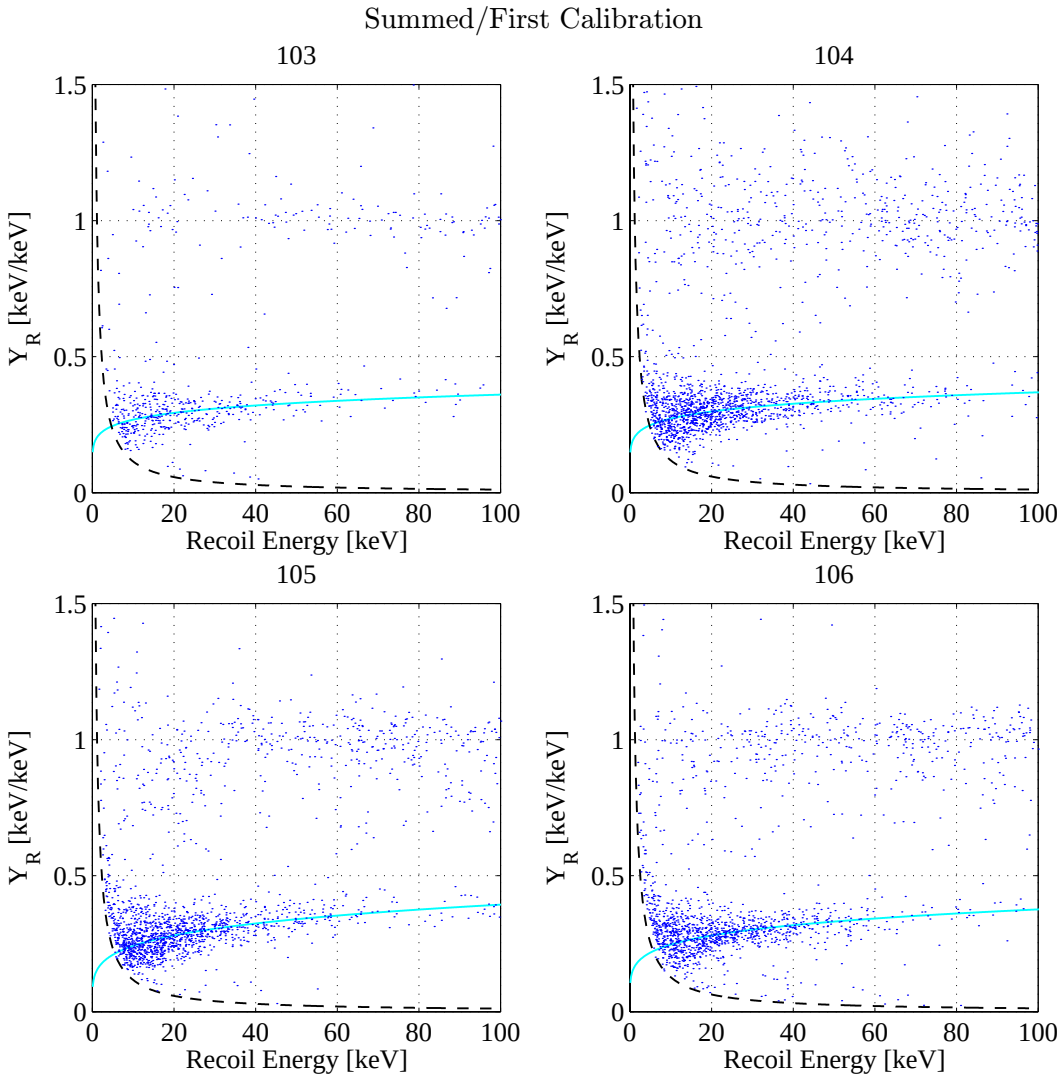

Figure 4.32: Ionization yield vs. recoil energy for the first neutron calibration. Light solid line: fitted nuclear-recoil line. Dashed line: ionization-search threshold; if shown, noise-charge events would form a band having the same shape as the dashed line but just below it.

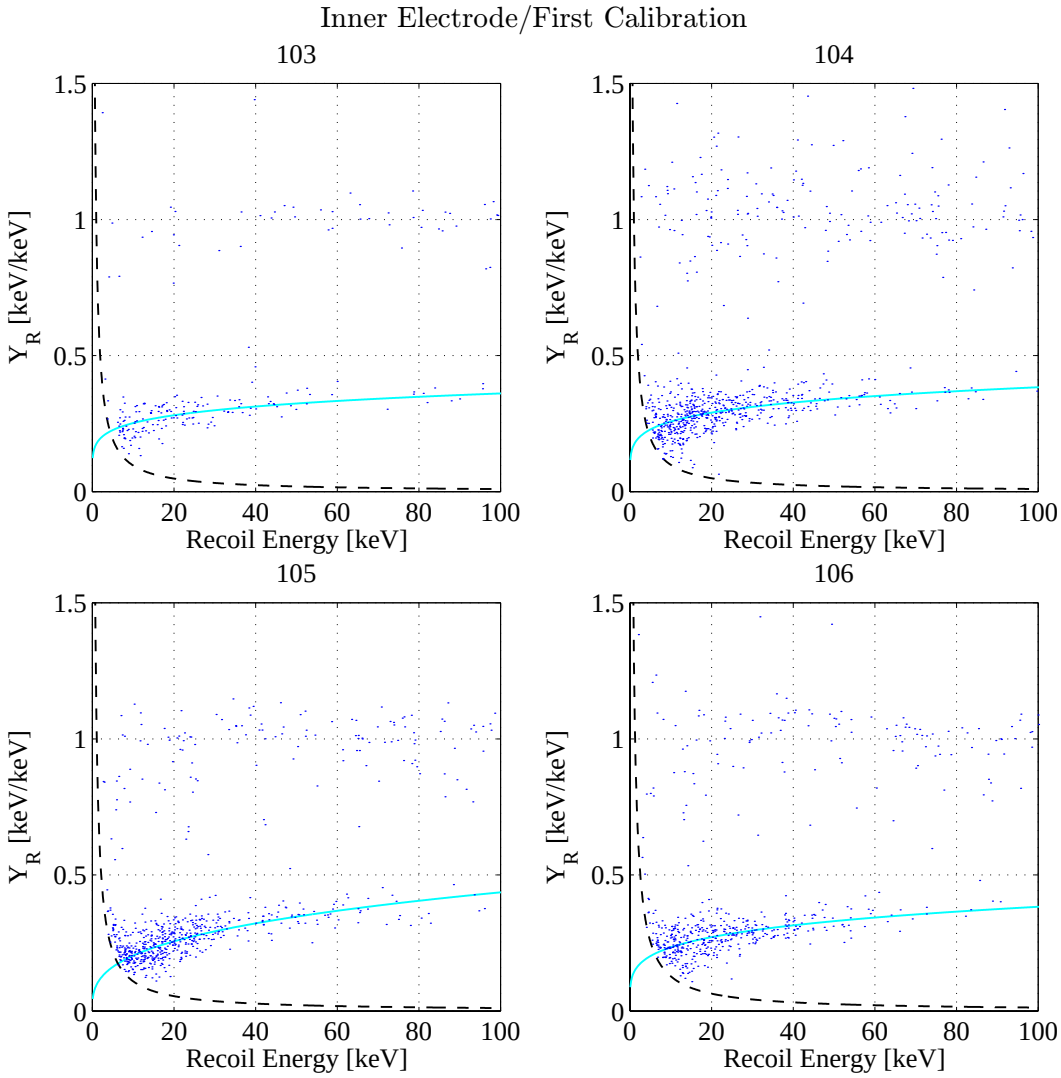

Figure 4.33: Inner-electrode ionization yield vs. recoil energy for the first neutron calibration, cQinOnly events only. Legend as in Figure 4.32.

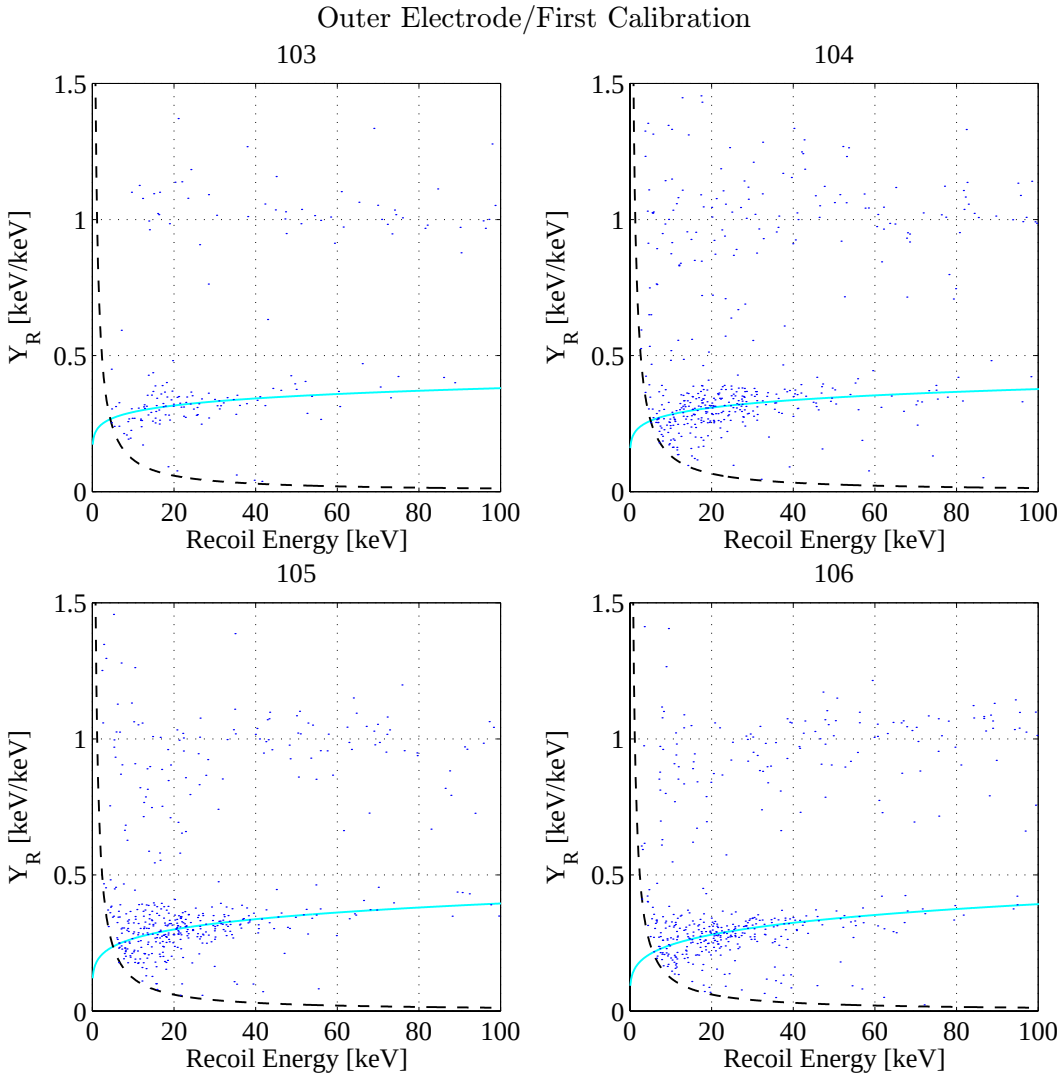

Figure 4.34: Outer-electrode ionization yield vs. recoil energy for the first neutron calibration, cQoutOnly events only. Legend as in Figure 4.32.

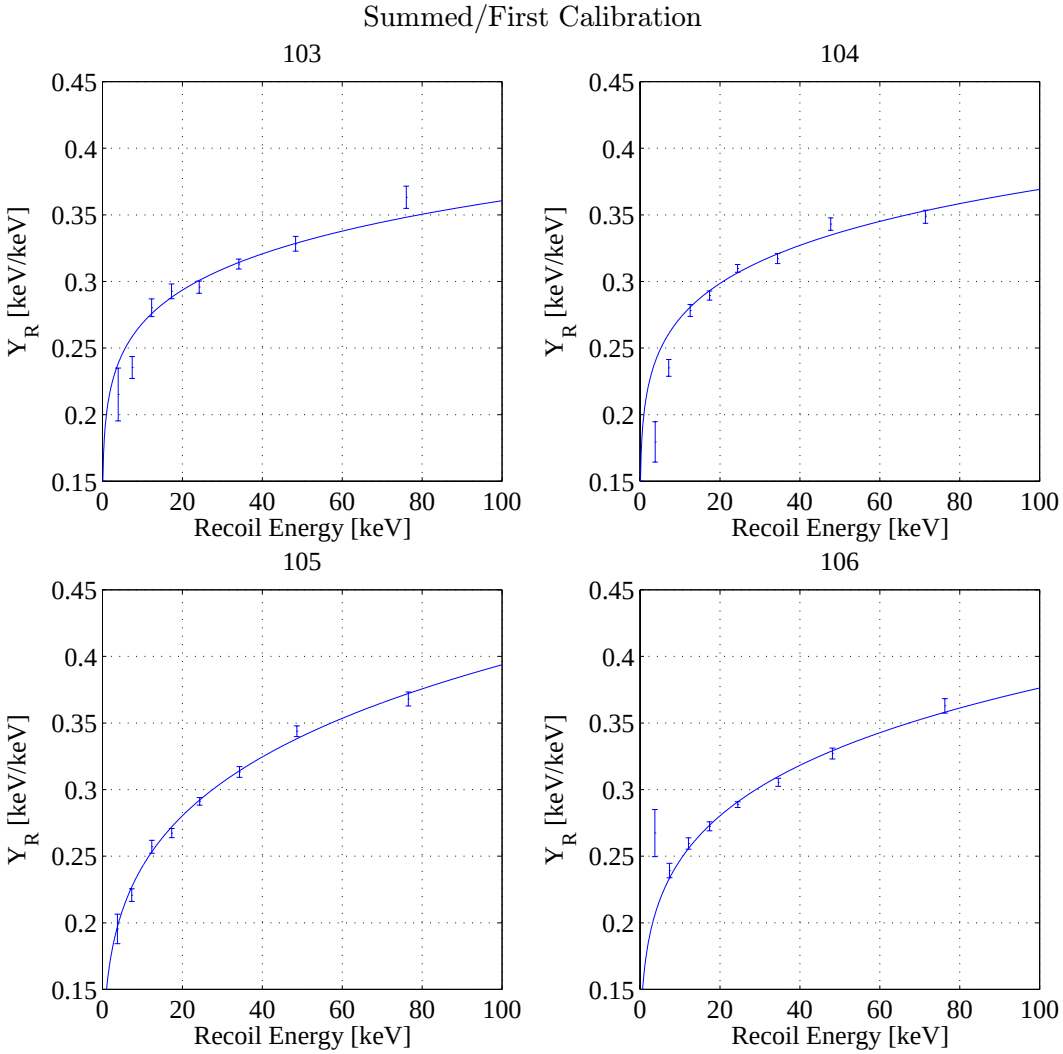

Figure 4.35: Nuclear-recoil-line data points and fits for the first neutron calibration.

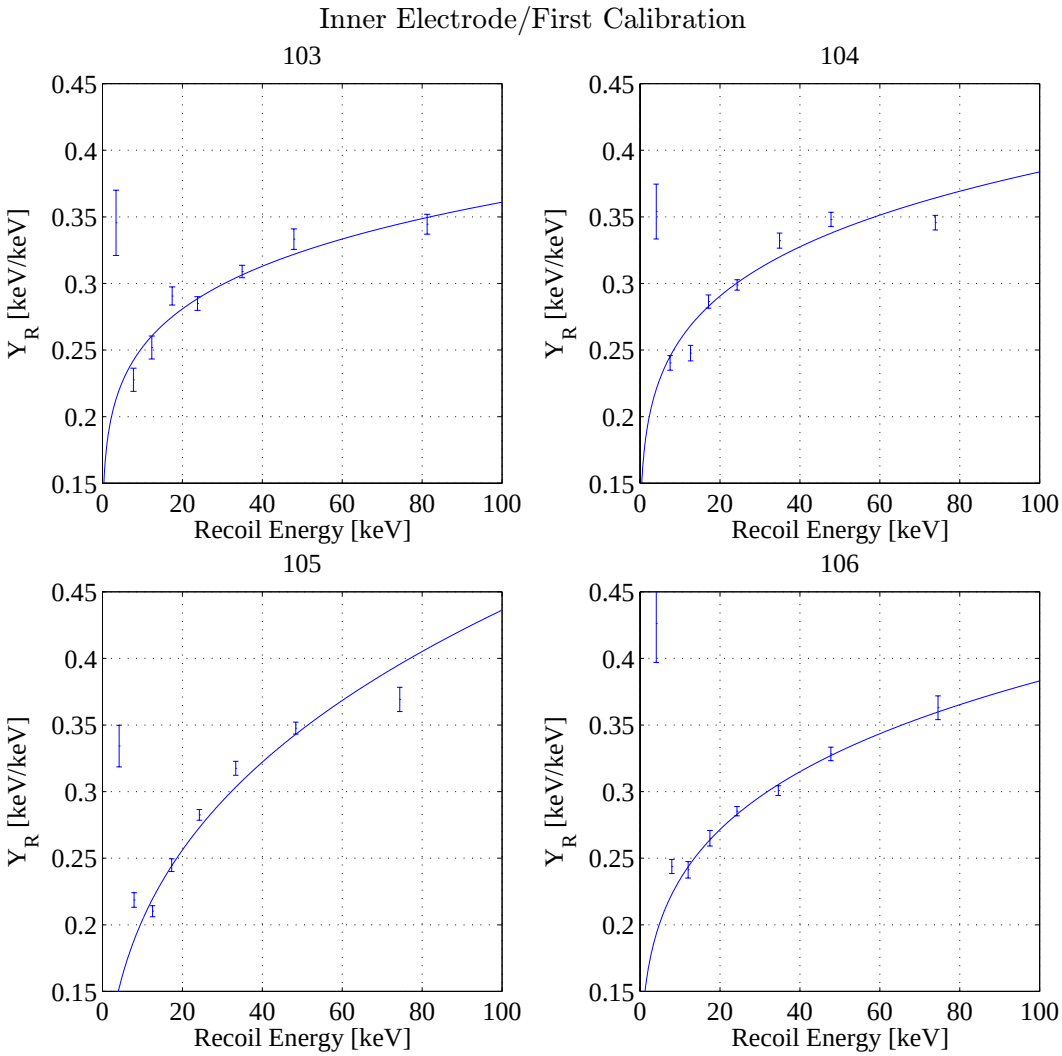

Figure 4.36: Inner-electrode nuclear-recoil-line data points and fits for the first neutron calibration.

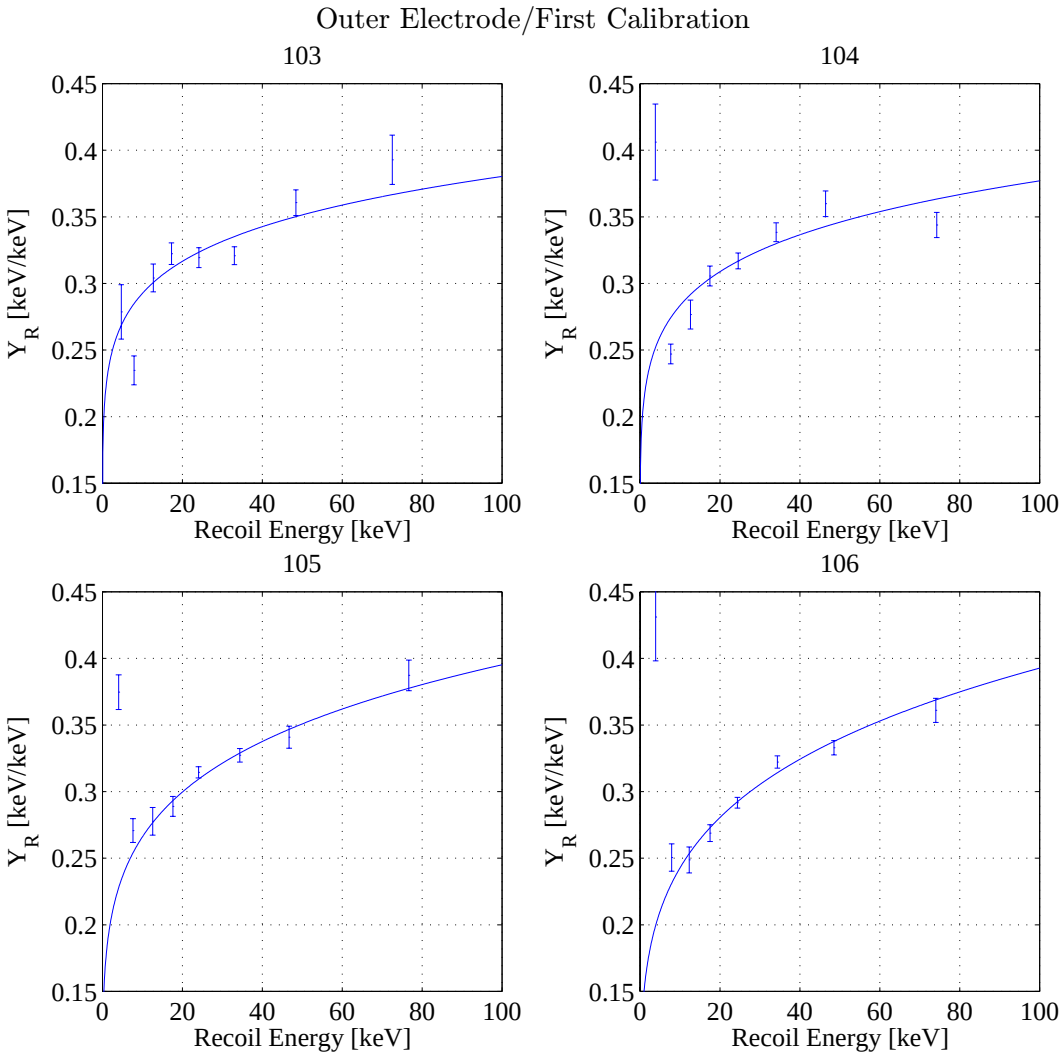

Figure 4.37: Outer-electrode nuclear-recoil-line data points and fits for the first neutron calibration.

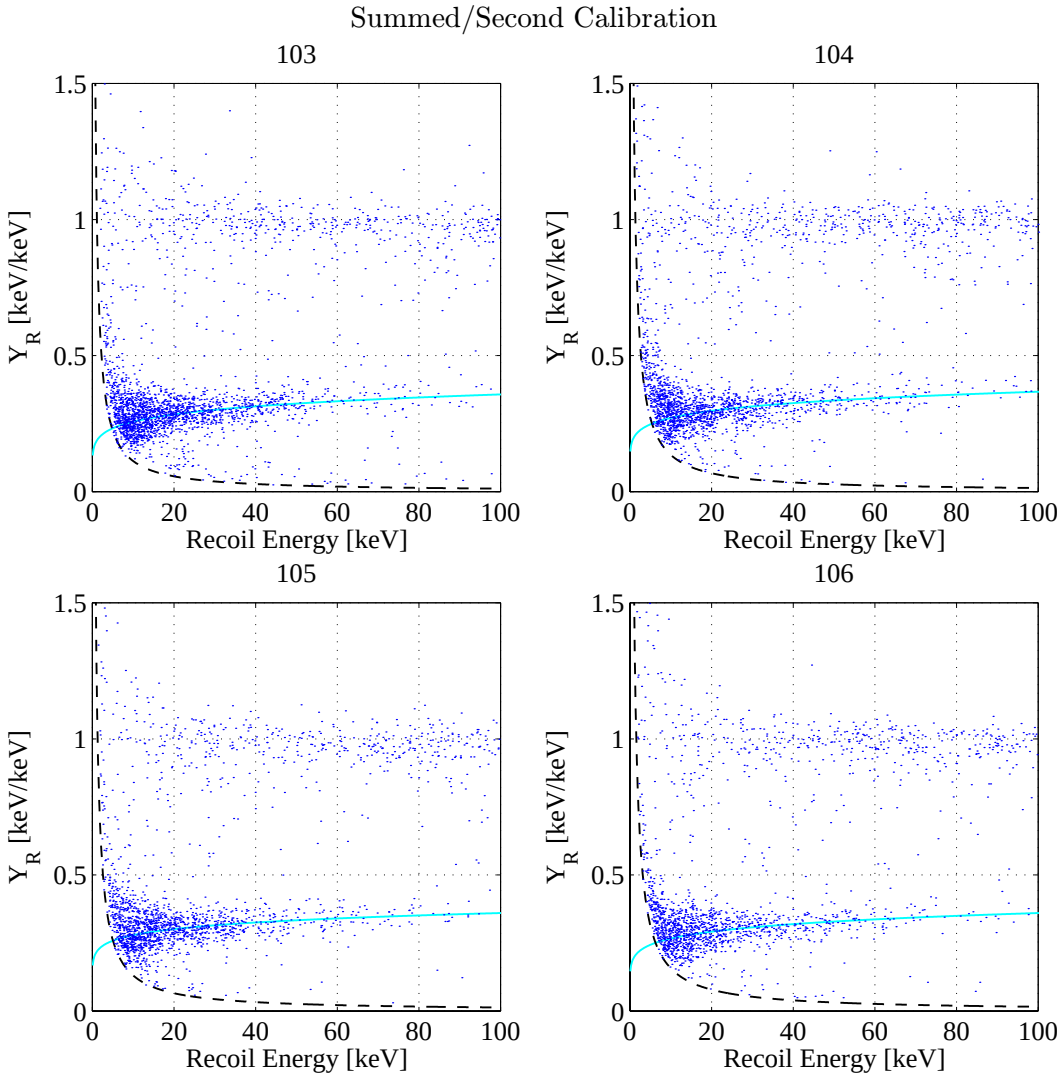

Figure 4.38: Ionization yield vs. recoil energy for the second neutron calibration. Legend as in Figure 4.32.

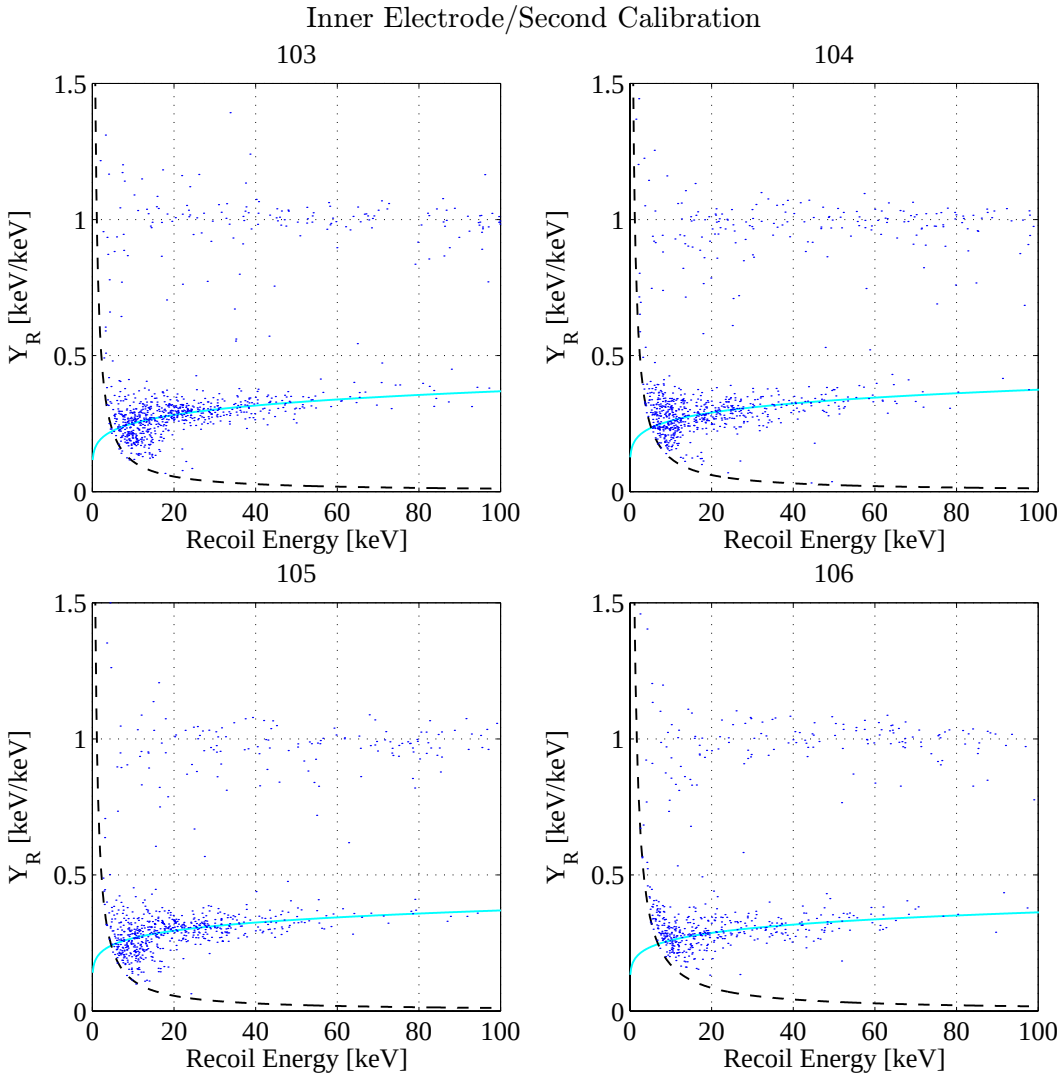

Figure 4.39: Inner-electrode ionization yield vs. recoil energy for the second neutron calibration, cQinOnly events only. Legend as in Figure 4.32.

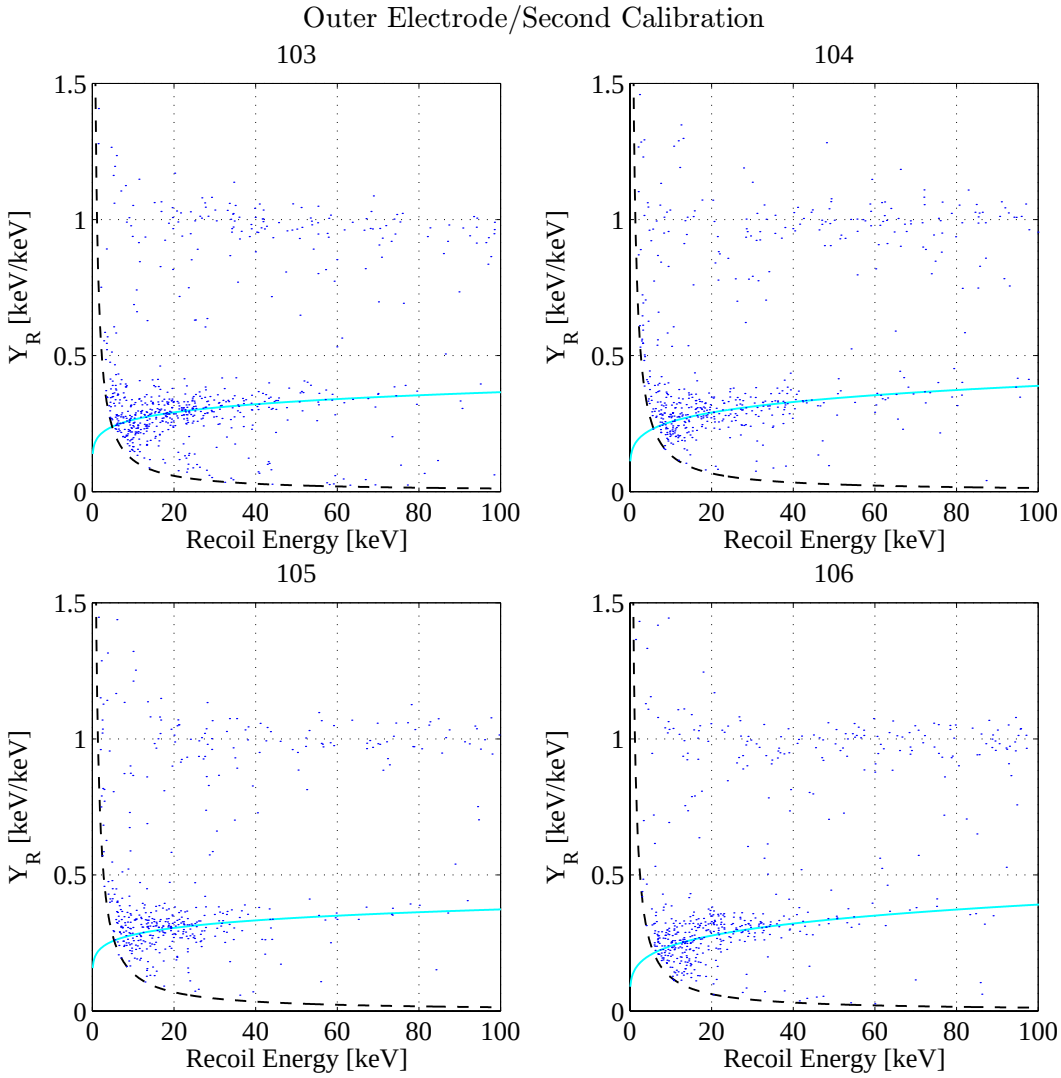

Figure 4.40: Outer-electrode ionization yield vs. recoil energy for the second neutron calibration, cQoutOnly events only. Legend as in Figure 4.32.

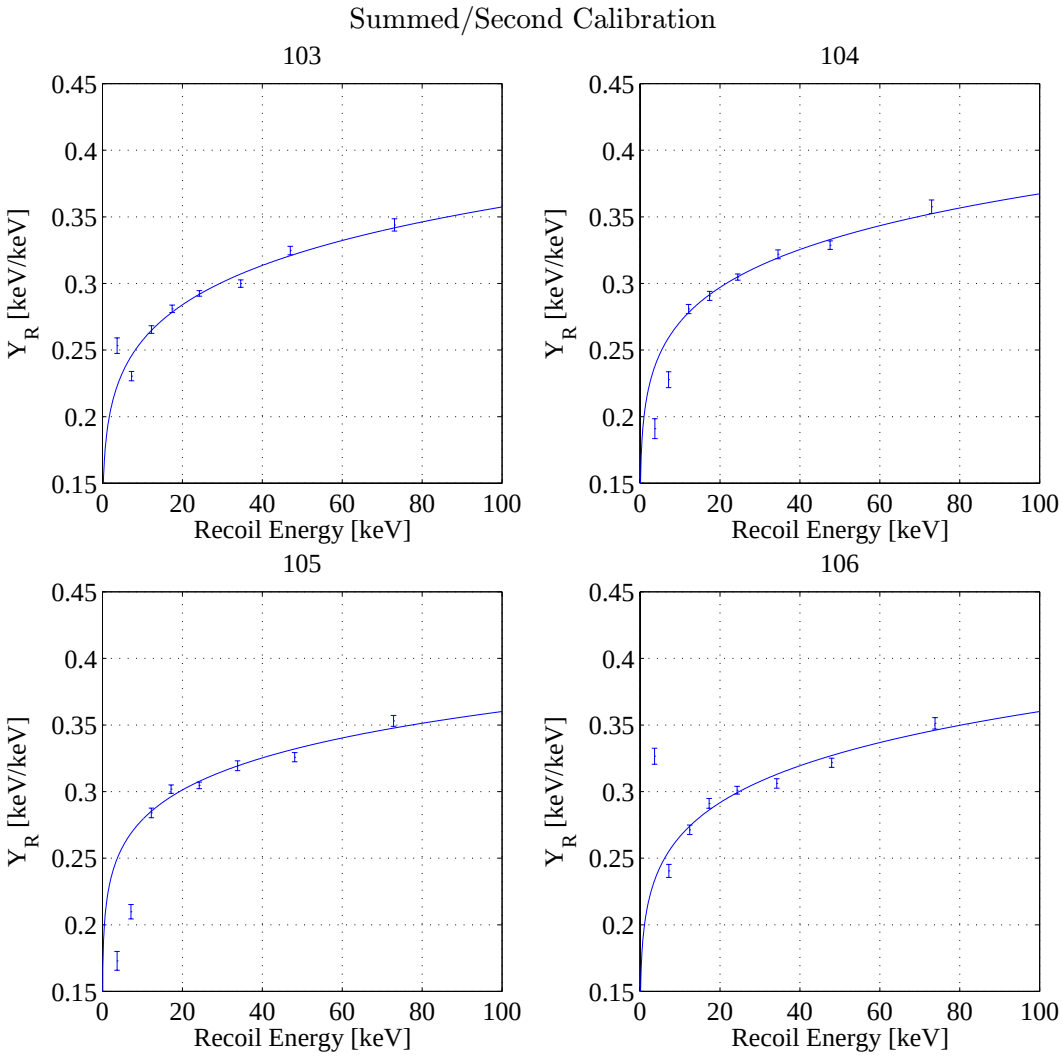

Figure 4.41: Nuclear-recoil-line data points and fits for the second neutron calibration.

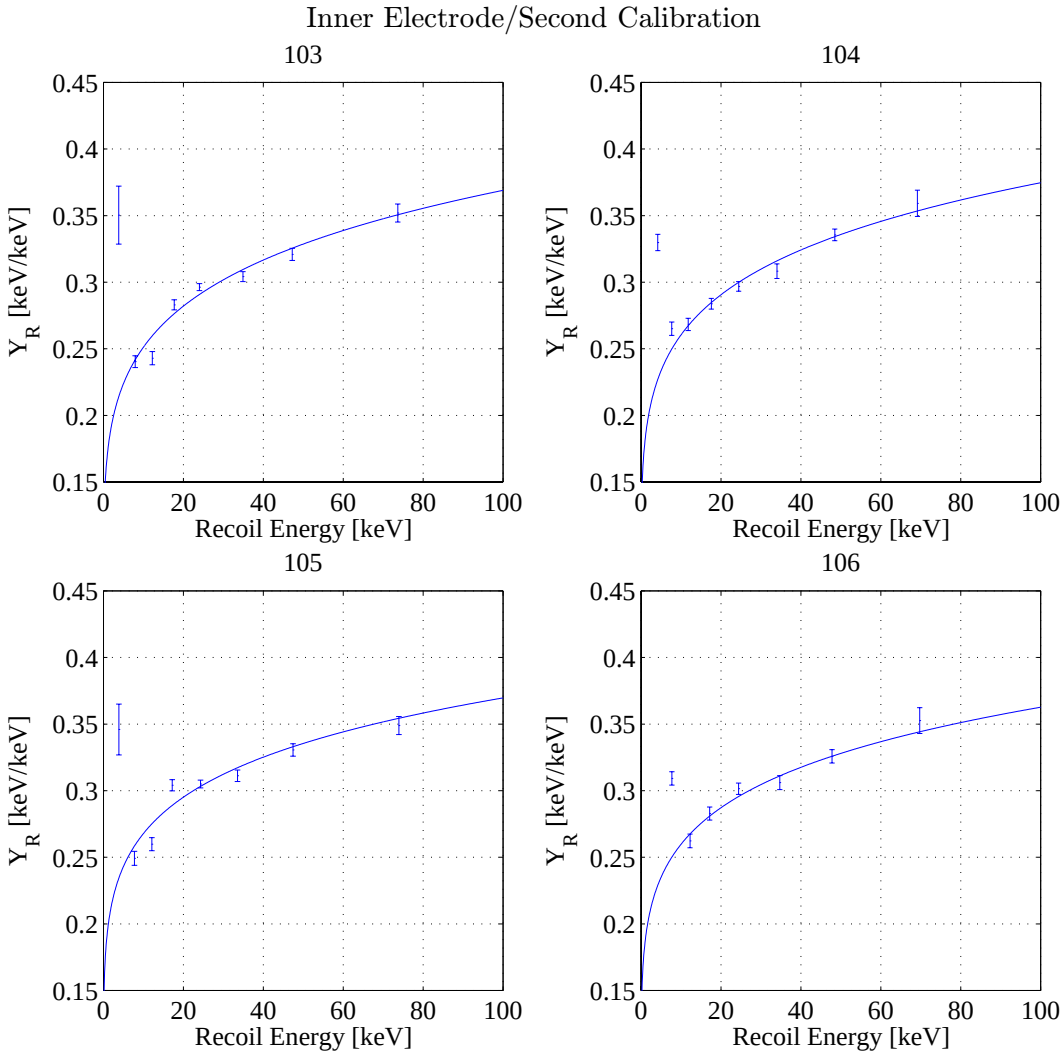

Figure 4.42: Inner-electrode nuclear-recoil-line data points and fits for the second neutron calibration.

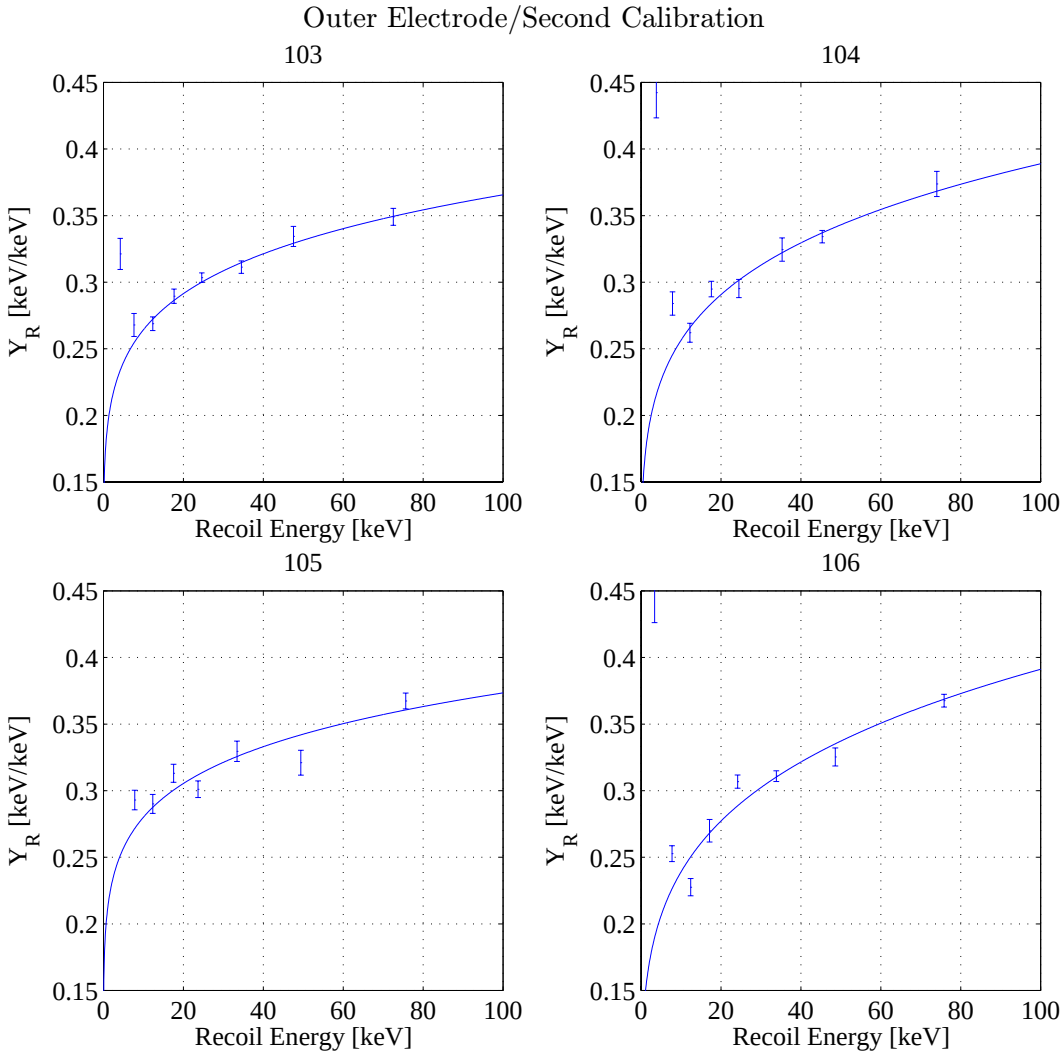

Figure 4.43: Outer-electrode nuclear-recoil-line data points and fits for the second neutron calibration.

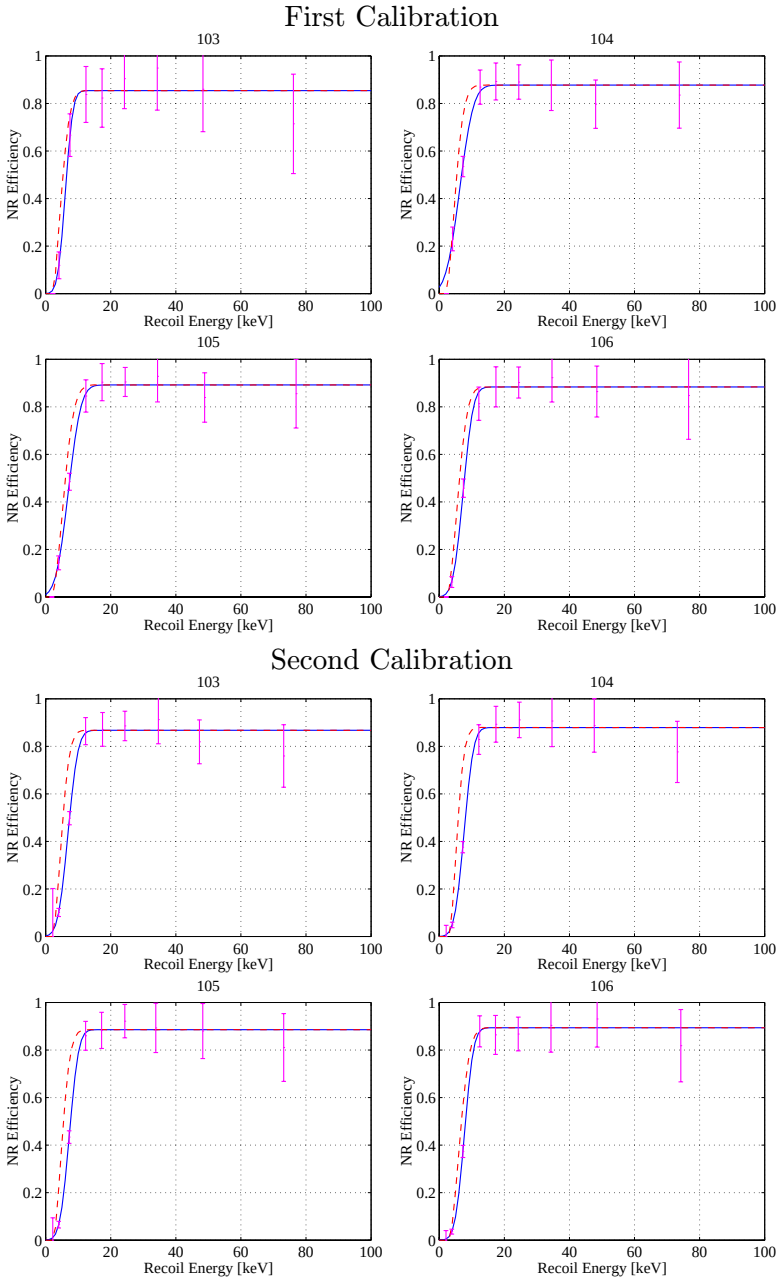

Figure 4.44: Nuclear-recoil efficiency with no ionization-partition cut. Error bars: calculated from data. Solid line: fit to an error-function form. Dashed-line: calculation as described in text. The two lines coincide above 20 keV.

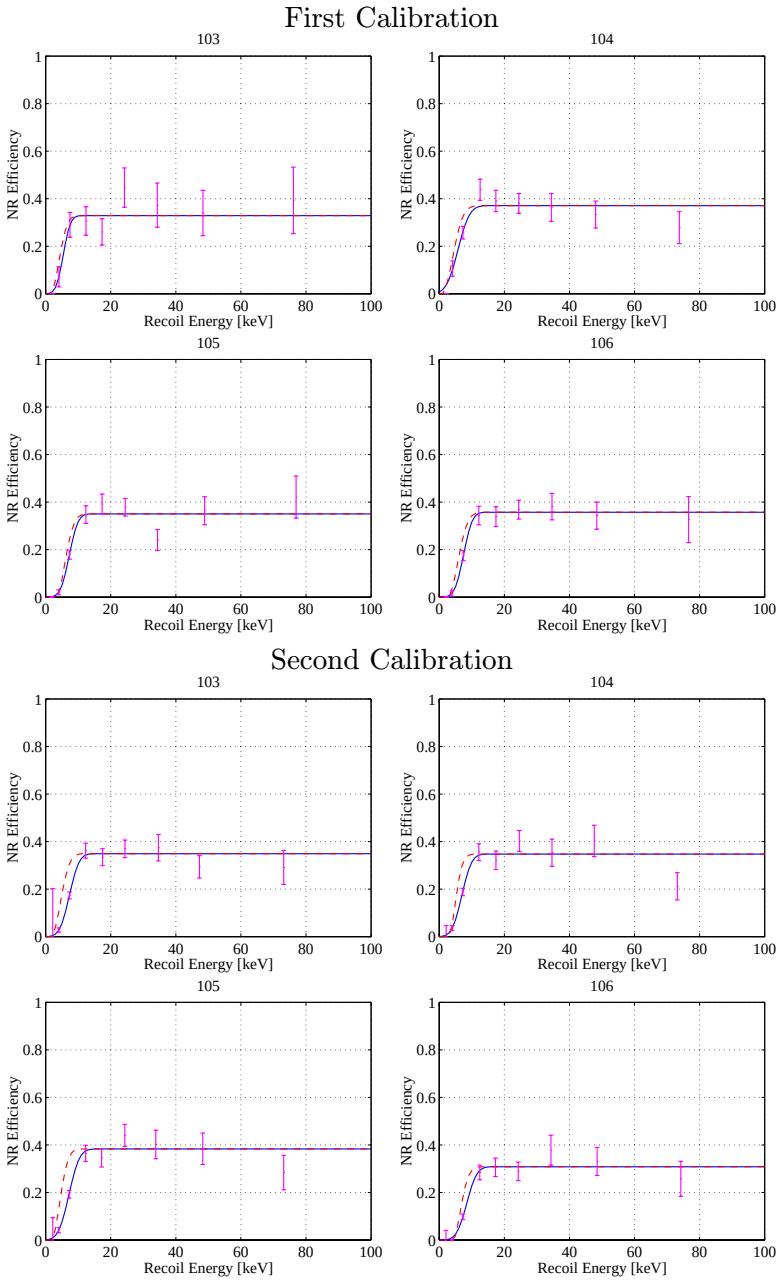

Figure 4.45: Nuclear-recoil efficiency for cQinOnly events. Legend as in Figure 4.44.

4.5.2 Photon Misidentification

The photon-misidentification parameter, β_γ, is given by comparing the total number of events observed in a photon-calibration data set to the number that are accepted by the nuclear-recoil-acceptance cut. The photon calibration is done by inserting a ^{60}Co source through a small hole in the lead shield. ^{60}Co emits two high-energy photons, at 1173 keV and 1332 keV. These photons Compton scatter uniformly through the detectors (the mean free path of such high-energy photons is large compared to a detector), yielding a sample of bulk electron recoils with possibly a small fraction of low-ionization-yield events. Such events arise from electrons both ejected from surrounding materials by high-energy photons as well as from electrons kicked into the dead-layer by Compton-scattering events *inside* the detector [132]. The latter effect was found, via simulation, to dominate the observed fraction of low-yield events in the Run 18 photon-calibration data.

In performing the photon calibration, a trigger veto against high-energy detector events is implemented. Events above about 100 keV are not useful for determining the photon misidentification at low energies. The details of the implementation of this veto are discussed in Chapter 6. The calibration data are analyzed in the same way as the normal data stream. A larger fraction of events are cut due to pileup, but this is not a concern: the photon misidentification is determined by beginning with a set of events that pass all data-quality cuts and calculating the fraction that pass the nuclear-recoil-acceptance cut; the overall normalization of the efficiency has no effect, as long as there are no data-quality cuts that have significant dependence on ionization yield.

Ionization yield vs. recoil energy plots of the photon-calibration data are shown in Figures 4.46, 4.47, and 4.48. Data passing the three different electrode-containment cuts — cQinOnly, cQShare, and cQoutOnly — are shown separately to study the effect of the inner-outer and top-bottom electrode breaks on photon misidentification. Histograms of ionization yield in two recoil-energy bins, 10–30 keV and 30–100 keV, are shown in Figures 4.49, 4.50, and 4.51 to display the shape of the low-yield tails.

The data were taken during July, 1999, so the nuclear-recoil band calculated from the September, 1999, neutron calibration is used. A veto-anticoincidence cut is made to avoid contamination of the data set with muon-coincident neutrons and electrons. Events with any saturated detectors are also discarded — saturated detectors can electrically cross-talk to other detectors, producing pulses just above the ionization threshold that may enter the nuclear-recoil band. The WIMP multiple-scatter probability is vanishingly small, so no bias for the final dark-matter analysis is introduced. The photon-misidentification parameter is $\beta_\gamma \equiv N_l/(N_t - N_l)$, where N_l is the number of misidentified photons (leaking into the nuclear-recoil band) and N_t is the total number of events. Note that the denominator is *not* N_t, though in this case it makes negligible difference. The statistical analysis is straightforward (there are no uncertainties that enter), and Feldman-Cousins 90% CL intervals on the photon-misidentification parameter β_γ are calculated and shown in Table 4.5. Note that β_γ is overestimated in the 30–100-keV bins because the high-energy veto takes effect at a fixed ionization energy, not fixed recoil energy; this bias is visible in the ionization yield vs. recoil energy plots. This bias has little qualitative effect on the conclusions.

The upper limits to the photon-misidentification parameters for cQinOnly and cQShare events are both approximately 0.1–0.5% and equal within statistical uncertainty. cQoutOnly events exhibit significantly higher photon misidentification, at the level of 1–2%. The cQinOnly and cQShare numbers are clearly Poisson limited, while the cQoutOnly numbers are much more precisely determined, so the true value of β_γ for cQinOnly and cQShare likely is near 0.1%. It is also clear

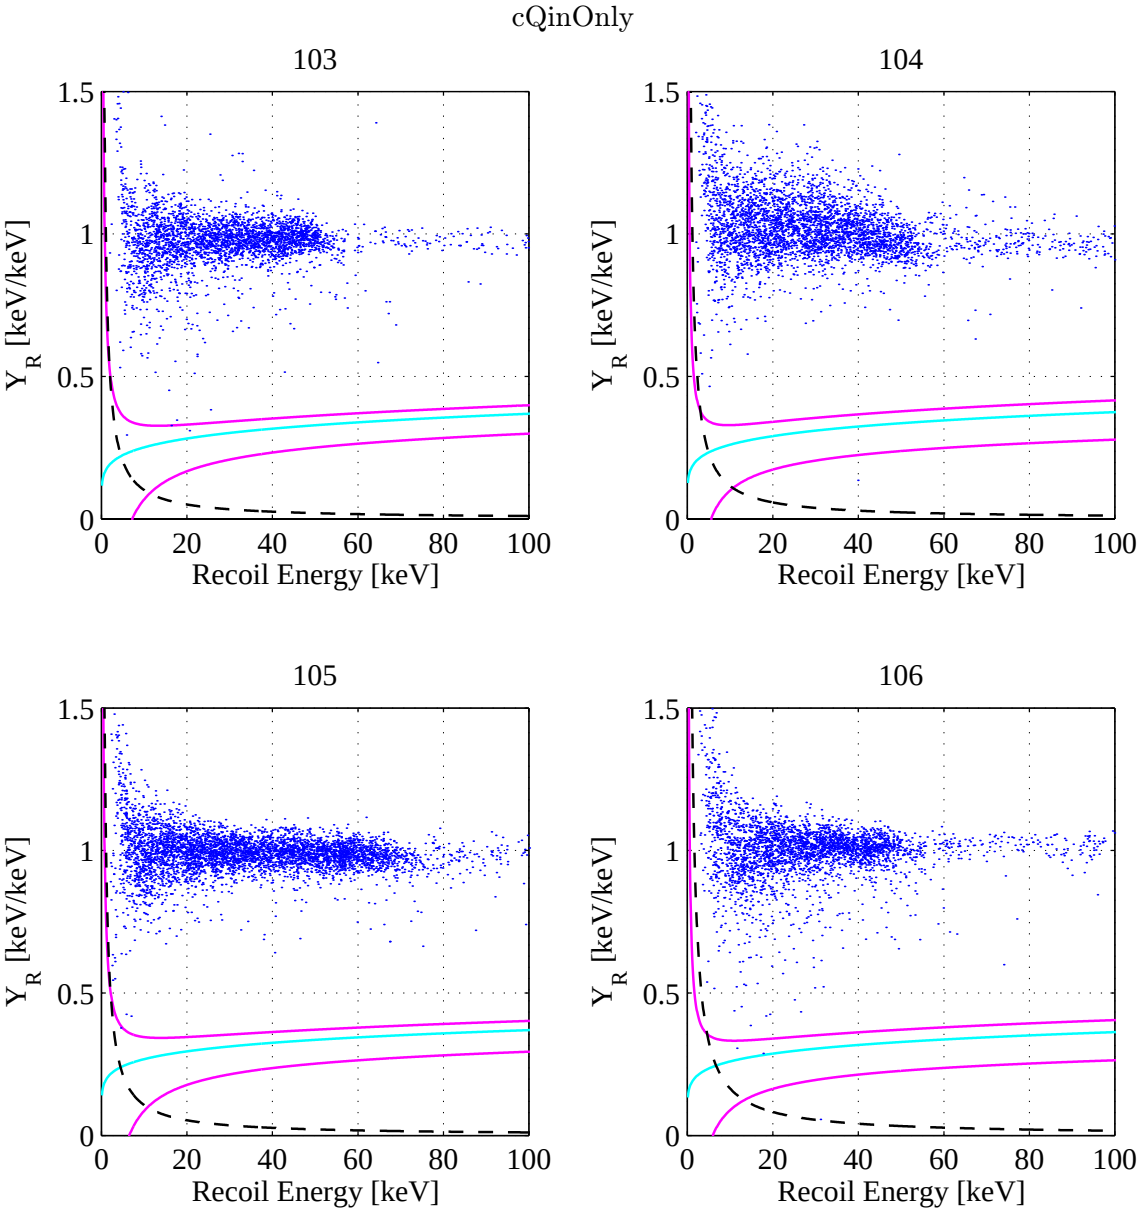

Figure 4.46: Ionization yield vs. recoil energy for cQinOnly photon-calibration data. Light line: center of nuclear-recoil band (power-law fits determined from neutron-calibration data). Dark lines: upper and lower edges of nuclear-recoil-acceptance region; note the asymmetric definition. Dashed: ionization-search threshold. The sharp cutoff at ~ 60 keV is due to the high-energy veto. The large width of the photon band for BLIP4 is due to its poor phonon-energy resolution.

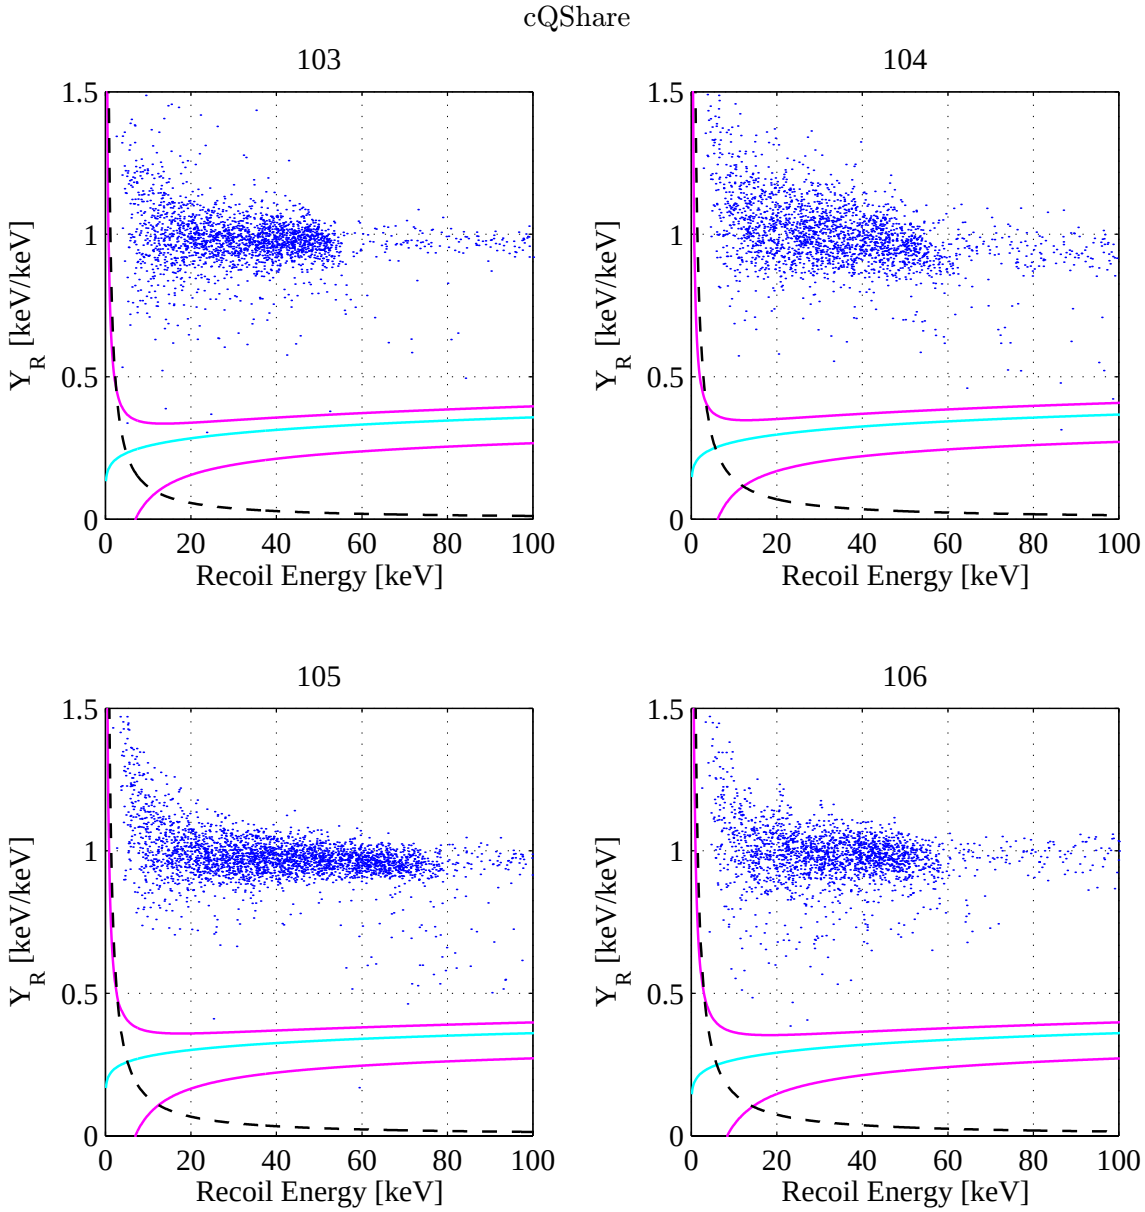

Figure 4.47: Ionization yield vs. recoil energy for cQShare photon-calibration data.

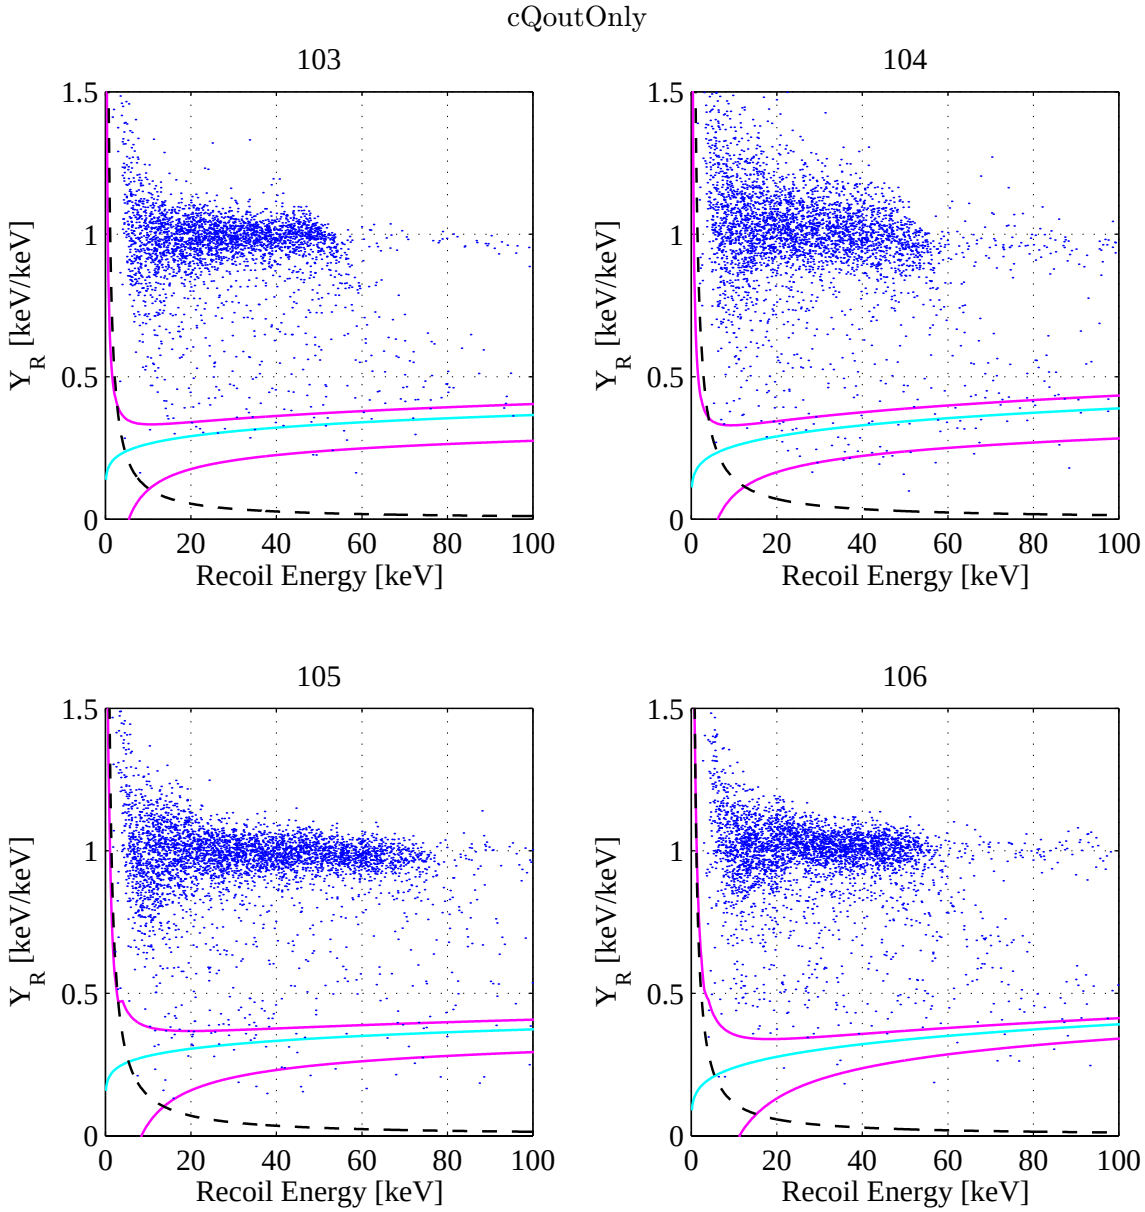

Figure 4.48: Ionization yield vs. recoil energy for cQoutOnly photon-calibration data.

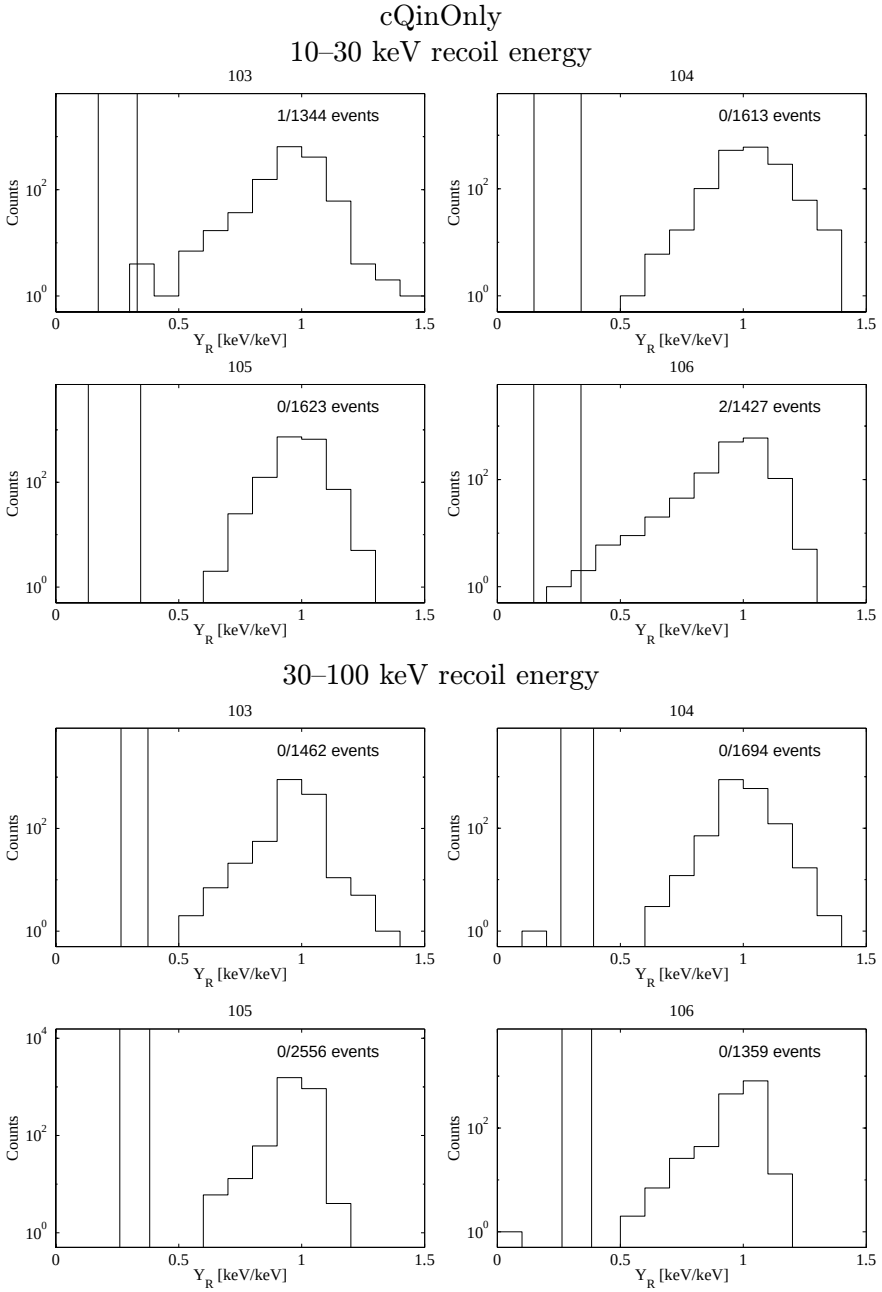

Figure 4.49: Ionization-yield histograms for cQinOnly photon-calibration data. The lines indicate the position of the nuclear-recoil-acceptance region, averaged over the energy bin. The text in the plots gives the total number of events and the number falling in the nuclear-recoil-acceptance region; this last number is determined using the fully energy-dependent acceptance region, not just the lines shown in the plots.

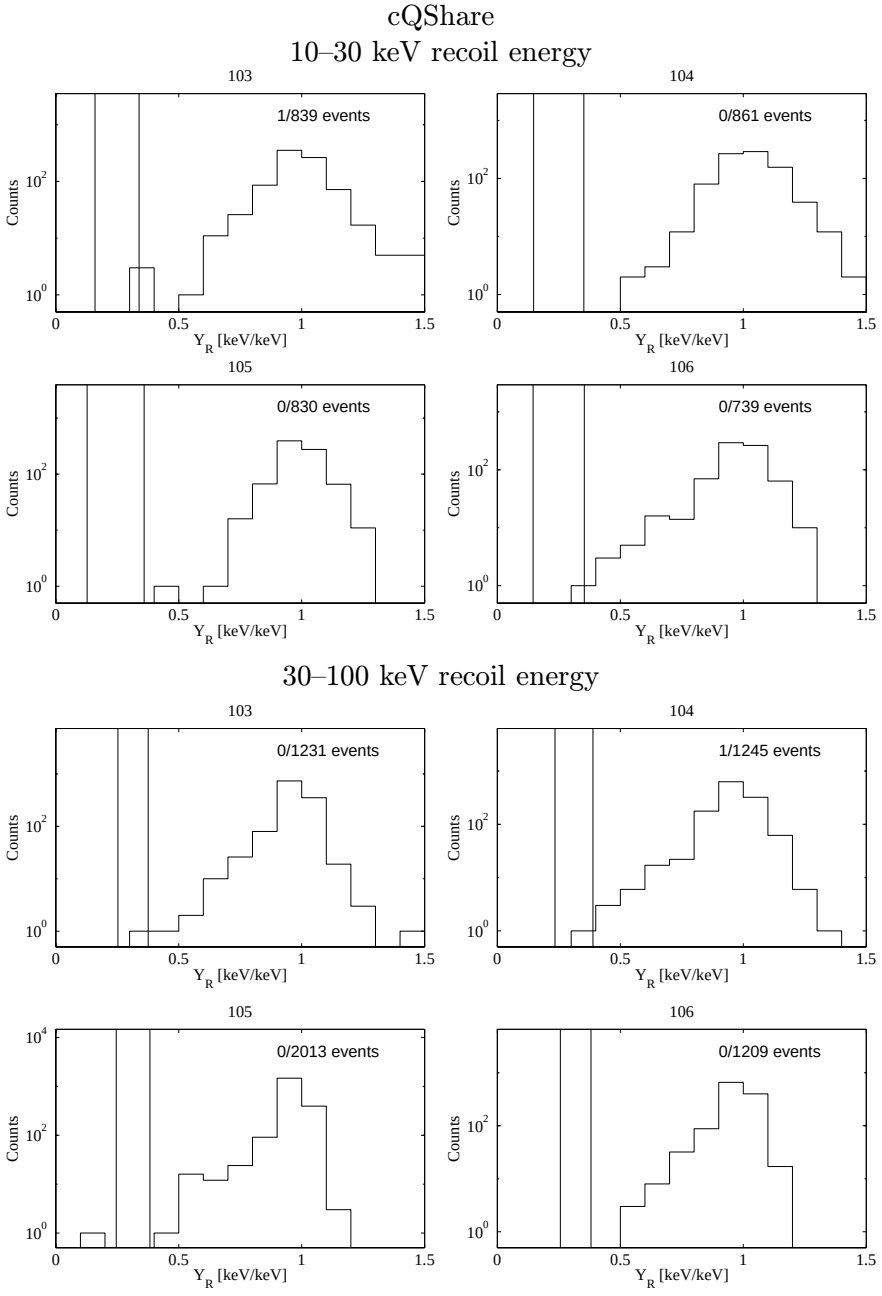

Figure 4.50: Ionization-yield histograms for cQShare photon-calibration data.

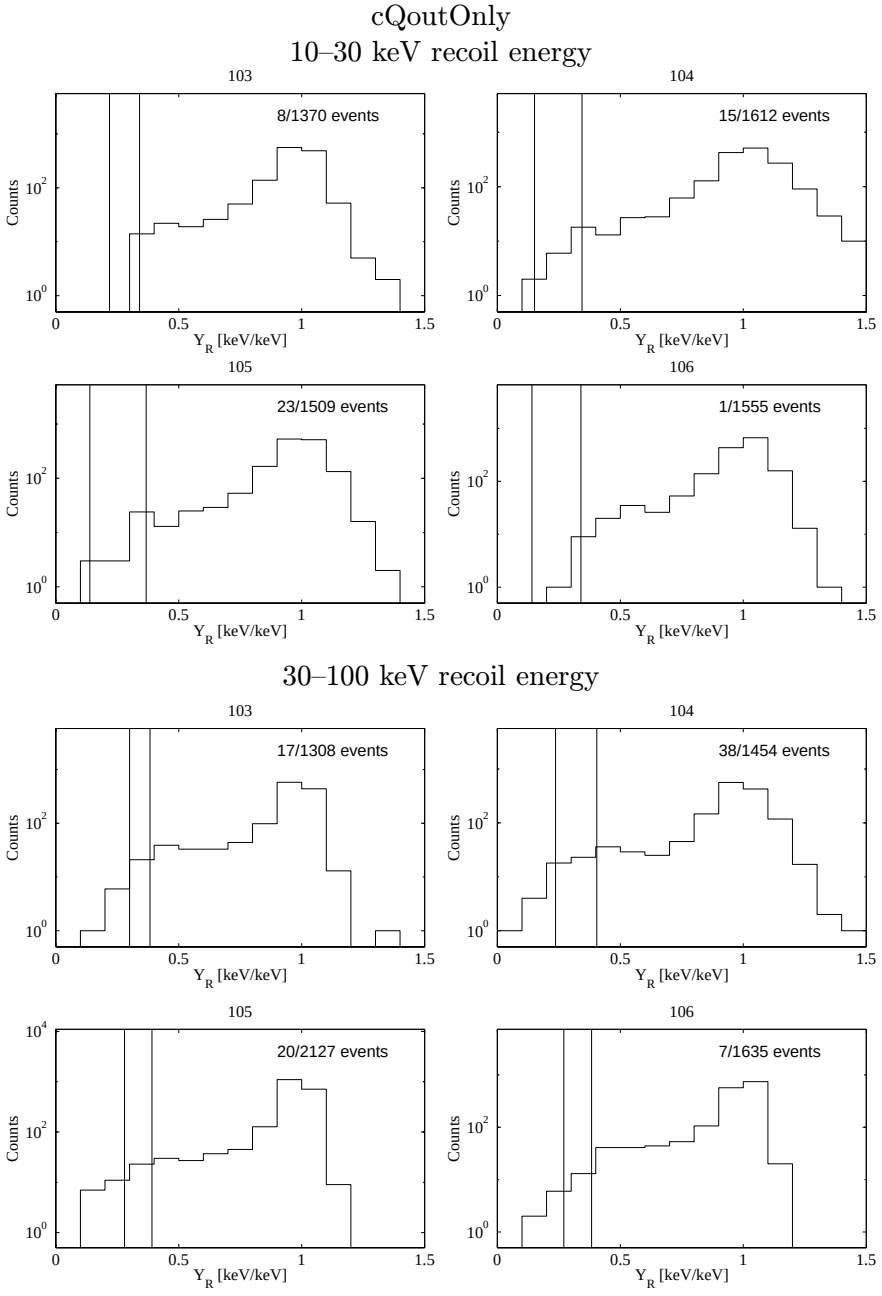

Figure 4.51: Ionization-yield histograms for cQoutOnly photon-calibration data.

Energy Range	BLIP3		BLIP4		BLIP5		BLIP6	
cQinOnly								
10–30	0.0082	0.32	0.0	0.15	0.0	0.15	0.037	0.41
30–100	0.0	0.17	0.0	0.14	0.0	0.095	0.0	0.18
cQShare								
10–30	0.013	0.52	0.0	0.28	0.0	0.29	0.0	0.33
30–100	0.0	0.20	0.0088	0.35	0.0	0.12	0.0	0.20
cQoutOnly								
10–30	0.29	1.0	0.036	1.3	0.044	2.0	0.0071	0.28
30–100	0.046	1.8	0.050	3.3	0.030	1.3	0.22	0.77

Table 4.5: Feldman-Cousins 90% CL intervals on β_γ in *parts per hundred*. Note that many of the entries correspond to 0 observed events and so the corresponding confidence intervals are one-sided.

that the ionization-yield distributions of the cQoutOnly data are qualitatively different from the other two sets; there is a significant shoulder to low yield. Two conclusions can therefore be drawn: the top-bottom break clearly is more susceptible to low-yield events; and including cQShare events incurs no appreciable photon-misidentification penalty.

A number of cuts normally applied to the low-background data have been relaxed in order to improve the statistical precision of this analysis. First, the "cPreQual(1)" cut, described in Chapter 6, is relaxed. This cut requires that no detector have pretrigger pileup; in photon-calibration mode, with an elevated event rate, this cuts a significant fraction of the data. In this analysis, each individual detector is required to have no pretrigger pileup for an event to be included, but the state of other detectors is not considered (aside from the no-saturation requirement). Also, both single-scatter and multiple-scatter events are included in this data set. In the dark-matter analysis, only single scatters are included because WIMPs are very unlikely to double scatter. Multiple scatters are only likely to degrade the photon misidentification by including events where a photon Compton scatters near the surface of one detector and ejects an electron that interacts in the dead-layer of an adjacent detector. Photon-misidentification estimates have been made with these fully restricted data sets and are statistically consistent with Table 4.5 because of the lower total numbers of events.

4.5.3 Electron Misidentification

The electrodes deposited on BLIPs 3 through 6 were never fully tested in the lab. As discussed earlier in Section 4.2, small devices were prepared with variants of the electrode used and tested, but, simply due to a lack of time, the exact electrode used on the detectors was never pre-tested. Therefore, no laboratory electron calibration is available.

The photon calibration contains a very small fraction of electrons, 0.66% in the 10-to-100-keV range according to the photon-calibration simulation [132]. The typical number of events observed in this energy range in the calibration is ~ 2000. Therefore, only about 15 electrons are expected. This is clearly insufficient for placing any useful limit on electron misidentification. The veto-coincident data also contain electrons. However, there are also a large number of neutrons (a few per day), thus providing a background for the measurement of electron misidentification. It would be necessary to subtract the expected neutron contribution.

The veto-anticoincident data provide an electron calibration. BLIP3 appears to be heavily contaminated with an electron source that results in clear electron bands in BLIPs 3 and 4. The contamination is likely due to the multiple processing steps performed on BLIP3 during fabrication (see Section 4.3), which may have left picene wax or some other material. The contamination argument is made in Chapter 8 using scatter plots of ionization yields for double scatters among detector pairs. It is useful here to reproduce this plot for the BLIP3/BLIP4 pair, shown in Figure 4.52. In making this plot, the event selection is designed to pick events due to the BLIP3/BLIP4 contamination source: these are veto-anticoincident events coincident between the two detectors and for which no other detector triggers. Because an analysis threshold of 10 keV is imposed for dark-matter analysis, both events in each pair are required to be above this threshold here. The sample is also restricted to events where both energies are below 100 keV; this reduces contamination by Compton-scatter photons coincident with electrons but keeps double-scatter low-energy electrons. Finally, at least one of the interactions in a double-scatter pair must pass cQinOnly or cQShare — this restricts to events on the faces of the detectors, well away from the top-bottom break. This data set is completely independent of the event set used for dark-matter analysis because all multiples are discarded in the WIMP candidate set and BLIP3 is discarded from the double-scatter-neutron sample because of its extreme contamination (see Chapter 8). Figure 4.53 shows ionization yield vs. recoil energy in the two detectors for the calibration data set. The surface events form a clear band in ionization yield, as is seen in the α-Si/Al-Schottky test device (see Figure 4.20). The bulk of the events are concentrated at low recoil energy, so this data set probes energies where electron misidentification is worst.

The definition of the data set may seem strange. The motivation is is follows. It is already clear from the photon-calibration data that cQoutOnly events suffer from larger photon misidentification than cQinOnly or cQShare events, presumably due to the top-bottom electrode break. Therefore, in the dark-matter analysis, single-scatter cQoutOnly events are discarded. However, requiring that both interactions of double-scatter pairs be cQinOnly or cQShare events greatly reduces the efficiency for observing double-scatter neutrons; it is desirable to avoid such a reduction. One would expect that loosening the cut, by requiring only one detector of double-scatter pairs to have a cQinOnly or cQShare event, would still strongly reject electrons interacting in the top-bottom break because there are no lines of sight between the detector faces and the top-bottom breaks. To evaluate this expectation, this looser cut is applied to the electron-calibration data.

In Figures 4.54, 4.55, and 4.56, histograms of ionization yield for BLIPs 3 and 4 for

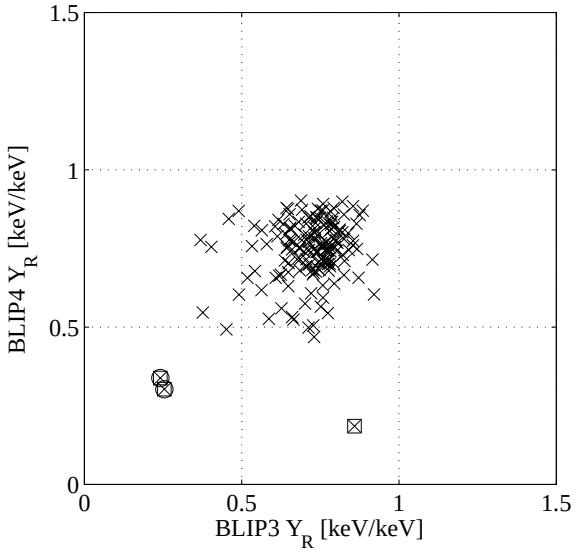

Figure 4.52: BLIP4 ionization yield vs. BLIP3 ionization yield for the electron-calibration data set. Circles indicate events tagged as nuclear recoils in BLIP3; squares indicate events tagged as nuclear recoils in BLIP4. The cuts defining this set are described in the text. Note that events that appear as bulk electron recoils in either detector ($Y_R \sim 1$) are cut. The large separation from the main distribution of the two events tagged as nuclear recoils in both BLIP3 and BLIP4 suggests they are, in fact, neutrons (no cut can be made to prevent neutron leakage); in the analysis in the text, they are conservatively assumed to be misidentified electrons. Similarly, the event tagged as nuclear recoil in BLIP4 only may be a photon-neutron coincidence, with somewhat low ionization yield for a photon in BLIP3. It also is conservatively assumed to be a misidentified electron.

cQinOnly, cQShare, and cQoutOnly events are shown. For a given event, only one of the two detectors is required to pass cQinOnly or cQShare, so the second detector may have a cQoutOnly event. It should be noted that every event that appears in Figure 4.52 appears in one of the histograms — no additional cuts are applied; the events are just divided up among the histograms.

The electron-misidentification parameter, β_β, is defined by

$$\beta_\beta \equiv \frac{N_l}{N_\beta} \tag{4.21}$$

where N_l is the number of events appearing in the nuclear-recoil band (subscript l for "leakage") while N_β is the number of events appearing between the nuclear-recoil band and the bulk electron-recoil (photon) band. This definition was made in Chapter 2; this particular denominator is used, rather than $N_l + N_\beta$, because in practice one wants to use N_β to predict N_l. Table 4.6 displays N_β and N_l for both detectors and for the two ionization-partition cuts. Feldman-Cousins 90% CL confidence intervals on β_β are calculated using these numbers. These intervals are not strictly correct because the statistical uncertainty on N_β has not been included; this has a nonnegligible effect for the entries with $N_\beta \lesssim 10$.

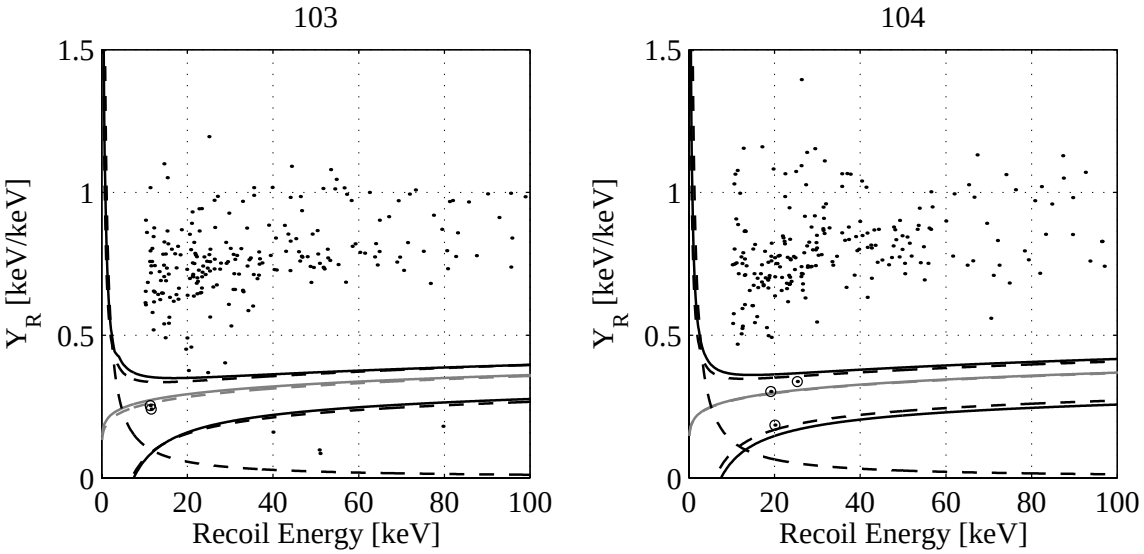

Figure 4.53: Ionization yield vs. recoil energy for the electron-calibration data set. Nuclear-recoil candidates are circled. Hyperbolic dashed line: ionization-search threshold. Light solid line: center of nuclear-recoil band, pre-April 3. Light dashed line: center of nuclear-recoil band, post-April 3. Dark solid line: nuclear-recoil-acceptance region, pre-April 3. Dark dashed line: nuclear-recoil-acceptance region, post-April 3.

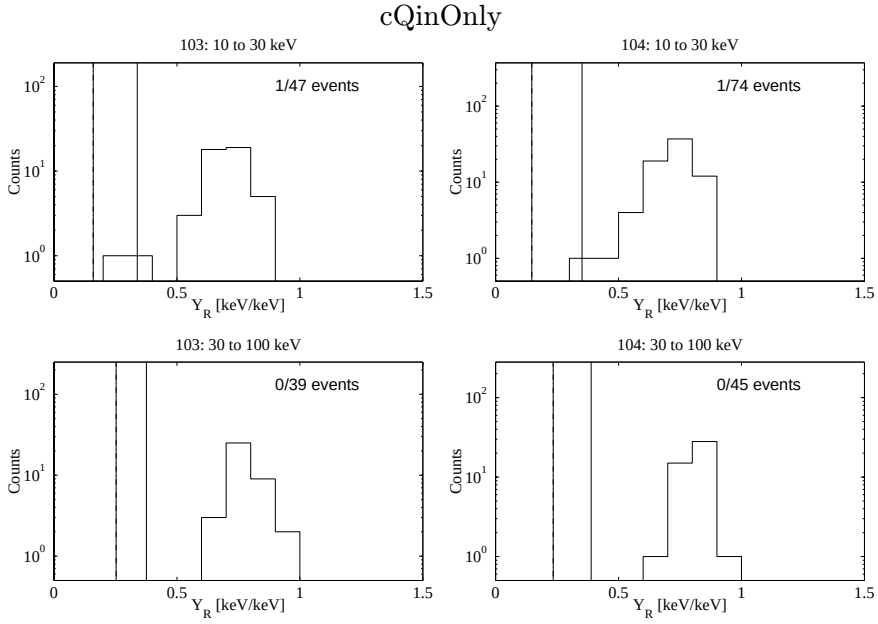

Figure 4.54: Ionization-yield distributions for BLIP3 and BLIP4 for cQinOnly scatters from the electron-calibration data set.

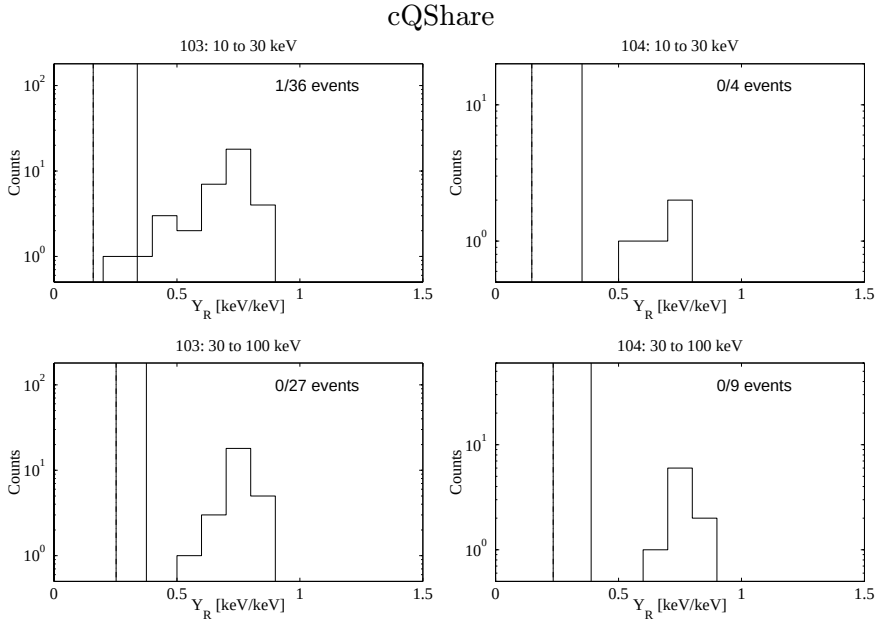

Figure 4.55: Ionization-yield distributions for BLIP3 and BLIP4 for cQShare scatters from the electron-calibration data set.

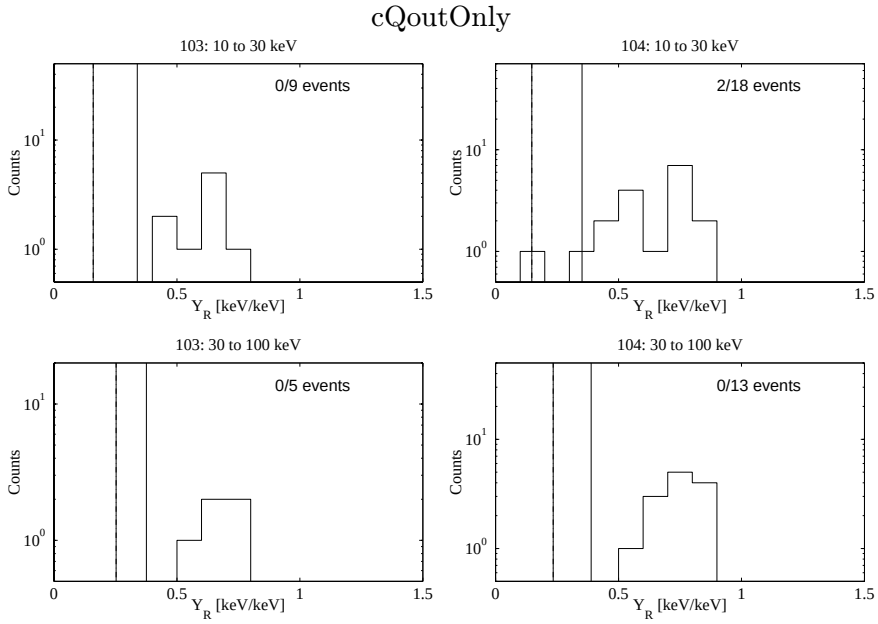

Figure 4.56: Ionization-yield distributions for BLIP3 and BLIP4 for cQoutOnly scatters from the electron-calibration data set.

Energy Range	BLIP3				BLIP4			
	N_l	N_β	β_{90}		N_l	N_β	β_{90}	
cQinOnly								
10–30	1	46	0.0024	0.095	1	73	0.0015	0.060
30–100	0	39	0.0	0.063	0	45	0.0	0.054
10–100	1	85	0.0013	0.051	1	118	0.0009	0.037
Inner-Outer-Shared								
10–30	1	35	0.0031	0.12	0	4	0.0	0.61[1]
30–100	0	27	0.0	0.090	0	9	0.0	0.27[1]
10–100	1	62	0.0018	0.070	0	13	0.0	0.19[1]
cQoutOnly								
10–30	0	9	0.0	0.27[1]	2	16	0.033	0.37[1]
30–100	0	5	0.0	0.49[1]	0	13	0.0	0.19[1]
10–100	0	14	0.0	0.17[1]	2	29	0.018	0.20

Table 4.6: Electron leakage confidence intervals. These are Feldman-Cousins 90% CL confidence intervals. [1]: these values underestimated: fluctuation in N_β has not been included.

Two conclusions can be drawn from these data. First, at the available statistical precision, BLIP3 and BLIP4 exhibit no difference in electron misidentification. This statement is fairly strong for cQinOnly given the large number of events, but much weaker for cQShare and cQoutOnly. Figure 4.20 indicates that BLIP4 should exhibit better electron rejection because its positively biased face is exposed to BLIP3's negatively biased face. If they truly had the same rejection, BLIP4 should probably have seen another leakage event in the 10–30-keV cQinOnly sample, but the fluctuations are large with such small numbers of events. The second conclusion is that there appears to be no worsening of electron misidentification in going from the cQinOnly to the cQShare cut, at least for BLIP3 where there are comparable numbers of events in the two data sets. In going to cQoutOnly, the misidentification rate is statistically consistent with the other cuts (the 90% CL intervals overlap), but again this statement is very weak because of the small numbers of events.

These misidentification-parameter estimates can be used in two ways in the final analysis. First, the cQinOnly numbers can be used to estimate the leakage of single-scatter electrons into the single-scatter nuclear-recoil set used for the WIMP search. This estimate is, in the end, not useful because it is rather high, limited by the statistical precision achievable with the low-statistics electron calibration. Second, and more importantly, the above analysis can be used to estimate the leakage of double-scatter electrons into the double-scatter-neutron data set. In that case, the calibration sets limits on electron leakage that are well below the number of double scatters observed. These two topics, as well as the validity of the application of this electron calibration to the final analysis, are discussed in detail in Chapter 8.

Chapter 5

Detector Mounting and Readout

5.1 Introduction

In this chapter, I discuss in detail the detector-mounting hardware and readout electronics. Figure 5.1 shows a tower with 4 BLIP detectors in their Run 19 positions. The detector mount and tower provide mechanical support and electrical connections to the detectors, support the cryogenic components of the front-end amplifiers, and heat-sink the wires travelling between the two.

The mounts hold the detectors and present the various detector signal and biasing lines on their outer faces. The detector package, consisting of the four detector mounts stacked together with endcaps attached, is suspended rigidly from the tower. Coaxial wire assemblies provide electrical connections from the mounts to the tower. The tower consists of four temperature stages: mixing-chamber (MC), cold plate (CP), still (ST), and LHe (4 K), corresponding to temperature stages of the Icebox. Thermal standoffs separate the stages. Wires run up the sides of the tower from the MC stage to the 4 K stage, with heat sinking at each stage. For each detector, a printed-circuit board holding JFETs is mounted at 4 K. The JFETs (referred to simply as "FETs" in the following) are the input terminals of the front-end amplifiers. These "FET cards" interface to copper-kapton striplines that connect to a room-temperature hermetic electrical feedthrough. An electronics crate is mounted near the hermetic feedthrough and houses front-end electronics boards holding detector-biasing circuits and the remainder of the front-end amplifiers.

The detector-mount, tower, and FET-card designs discussed here are the work of Dan Akerib, Bernard Sadoulet, Dennis Seitz, Tom Shutt, Garth Smith, Walter Stockwell, Roy Therrien, Storn White, and myself. The tower and FET-card thermal design was developed by Dan Akerib, Tom Shutt, and Walter Stockwell and is discussed in detail in Walter Stockwell's dissertation [108]. Dan Akerib, Dennis Seitz, and Storn White were responsible for design and production of the striplines. I review some of this material. The detector-mount and tower designs have been changed since last documented in [108], so I discuss these new features in some detail. The front-end electronics boards were developed and have been used in the Berkeley group since BLIP detector development began in the late 1980s. Contributions have been made by a large group of people: Alan Cummings, Yannick Giraud-Heraud, Bernard Sadoulet, Dennis Seitz, Tom Shutt, Walter Stockwell, and Dominique Yvon. The 9U-format boards used in Run 19 were produced under the supervision of Mike Crisler and Steve Eichblatt at Fermilab.

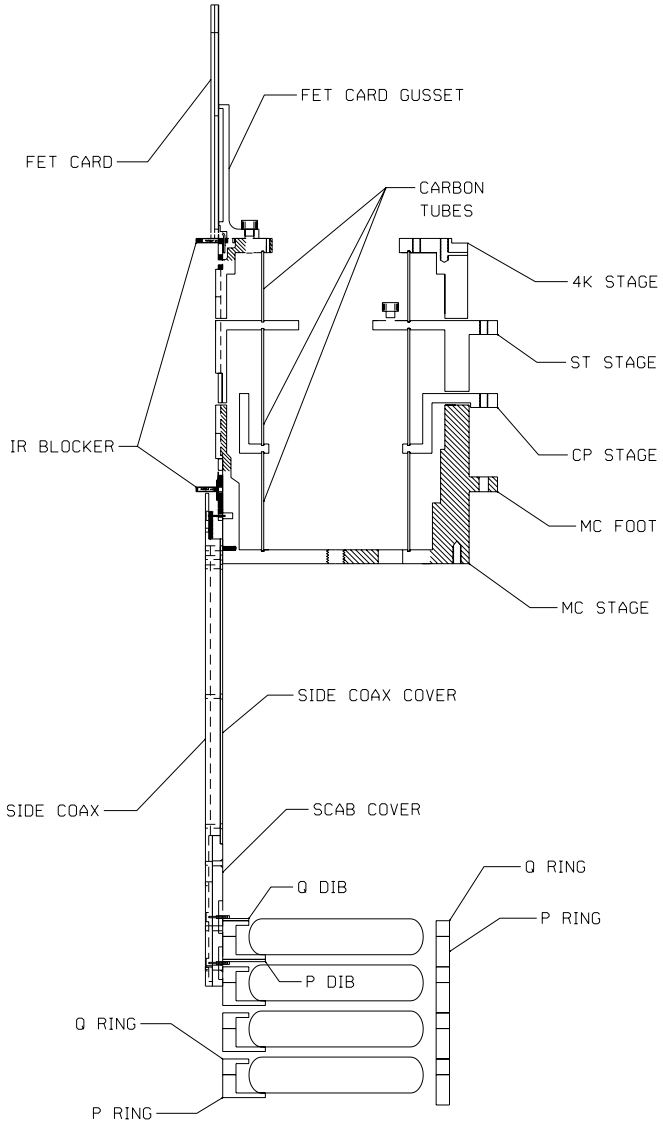

Figure 5.1: Detector mount and tower. The close-packed detector mounts for 4 detectors are shown at the bottom. The endcaps of the detector assembly are not shown. The separate Q (ionization) and P (phonon) rings that make a single close-packed detector mount are indicated. The Q and P Detector Interface Boards (DIBs) are shown for one detector. The thin vertical pieces are coaxial wire assemblies (side coaxes). The four temperature stages of the tower are shown, separated by graphite-tube thermal standoffs. The tab sticking out the side of the MC stage ("MC foot") indicates the mount points to the MC lid of the Icebox. The tabs sticking out the sides of the other stages receive thermal straps from the Icebox. The tower infrared-blocker boards are indicated. A FET card is shown at the top. A stripline (not shown) connects to the top of the FET card.

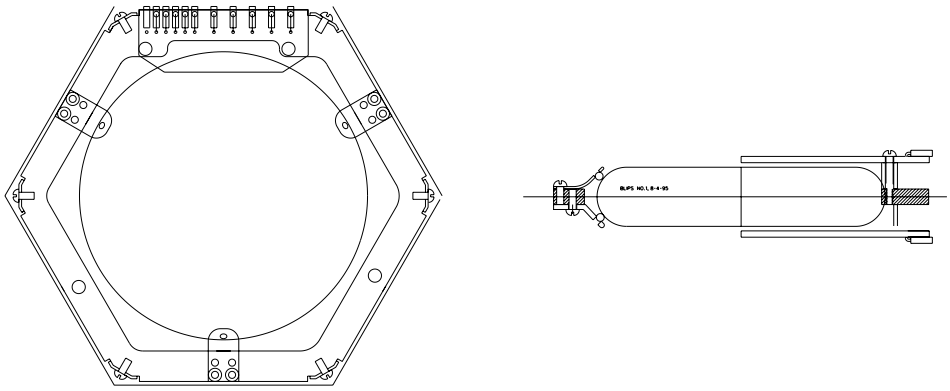

Figure 5.2: Minimal-mass detector mount. The circle in the center is the detector. The positions of the support feet are visible. The Y's are not shown in the top view, but are partially visible in the side view. The sapphire balls are shown in the side view.

5.2 Detector Mounts

The detector mounts used for BLIP detectors have evolved significantly from the original design. I present some of the history here to motivate the design of the present mounts.

5.2.1 Minimal Mass

The initial design goal for the detector mounts was to minimize the mass of material near the detectors. Any material near the detectors is a low-energy-photon source (via Compton scattering of higher-energy photons) as well as a neutron source. The original mount design is elegant, consisting of a low-mass hexagonal copper ring with support feet holding the detector at 6 points, as shown in Figure 5.2. The detector rests on sapphire balls at these points; sapphire has low thermal conductance. The sapphire balls are coated with gold on the side touching the detector to prevent the sapphire from cracking the germanium when cooled. Electrical connections are provided by wirebonds from the detector to thin copper-kapton printed-circuit boards mounted above and below the detector (called "Y"'s because of their shape).

Tests of this mounting scheme at the UCB Test Facility demonstrated that, because of the open design, infrared radiation emitted by the FETs (which self-heat to 120 K, yielding a blackbody spectrum peaked at ≈ 35 μm) reaches the detectors at 20 mK and degrades the ionization measurement by ionizing neutralized impurities (see Chapter 4). A number of modifications were made to absorb and block infrared radiation. Primary among these is the use of a "sock," a 0.25-mm-thick copper foil assembled around the detectors and attached to the MC stage of the tower. Residual gaps are sealed with copper tape.

The prototype Ge BLIP detector, E5, was operated at SUF in Runs 13 and 14 during summer and fall, 1996, in this type of mount. An exposure of 1.2 kg d was collected using this 60-g detector. A significant rate of veto-anticoincident low-ionization-yield events at low energy (< 100 keV) was observed: about 3 keV^{-1} kg^{-1} d^{-1} in the range 20 to 40 keV [109]. The physics of these "dead-layer" events is discussed in Chapter 4; they are caused by low-energy electrons

incident on the detectors. The minimal-mass mount provides no protection from sources of low-energy electrons on surfaces visible to the detector.

5.2.2 β-Coffins

In response to the above high rate of low-yield events, the mount was modified to completely enclose the detector in a small copper box, a design which provides a number of advantages. First, it clearly delineates the surfaces that the detector sees — only the inner faces of the box. These faces can be treated to remove low-energy-electron sources. Second, the detector can be mounted in its package under full cleanroom conditions, reducing the cleanliness required when attaching the detector to the tower; it is this latter step that is performed more frequently and could lead to detector contamination. Finally, by bringing these visible surfaces closer to the detector, the design increases the likelihood that particles associated with the generation of low-energy electrons also hit the detector, moving the events out of the 0-to-100-keV region. An example would be high-energy photons that eject electrons from the copper.

This design, somewhat morbidly named the "β-coffin" by Storn White, is shown in Figure 5.3. Copper sheets of 1-mm thickness cover the top and bottom surfaces of the detector, with small holes to allow wirebonding. The sides of the mount are constructed from thick copper. The Y's are glued to the outside faces of the copper sheets. "Umbrellas," consisting of small copper sheets on standoffs, cover the holes in the mount where wirebonds enter. The sapphire ball supports were abandoned, primarily for convenience. Pyramids of kapton are attached to the copper support feet such that the pointed end makes contact to the detector. The point has a small area and thus low thermal conductance.

During Runs 15 and 16, in winter and spring, 1997, two BLIP detectors were operated in these new mounts at SUF. One of the detectors was E5, remounted but otherwise unmodified. The second detector was BLIP1, the first full-size (165-g) Ge BLIP detector. Exposures of 0.9 kg d and 2.5 kg d, respectively, were collected using these detectors. Low-ionization-yield event rates of 3 keV^{-1} kg^{-1} d^{-1} and 0.5 keV^{-1} kg^{-1} d^{-1} in the range 20 to 40 keV were observed in E5 and BLIP1, respectively [111]. That is, E5's rate was unchanged while BLIP1 showed a rate lower by a factor of approximately 6. The large discrepancy is most easily explained by the fact that E5 was a prototype and suffered numerous cooldowns and remountings at UCB before being run at SUF, while BLIP1 was a relatively virgin detector. E5's low-yield event rate is due to surface contamination accumulated during these numerous remountings. The E5 low-energy-electron rate is thus not indicative of achievable rates. Since no control data set is available, the source of the $>$ 20-keV low-yield events in BLIP1 is not clear. The copper of the β-coffin and detector surface contamination are equally valid suspects.

5.2.3 Close Packed

Given the above results, a new mount was designed, shown in Figure 5.4, with the goal of using the detectors themselves to shield each other from external (*i.e.*, not on the detector surfaces) low-energy-electron sources. The detectors are packed close together with no intervening material; copper endcaps close the ends of the detector stack. This design is very different from the β-coffin mounts, in which each detector sees copper in all directions. In this "close-packed" geometry, a large fraction of the solid angle visible to the inner electrode of a detector is occupied by another

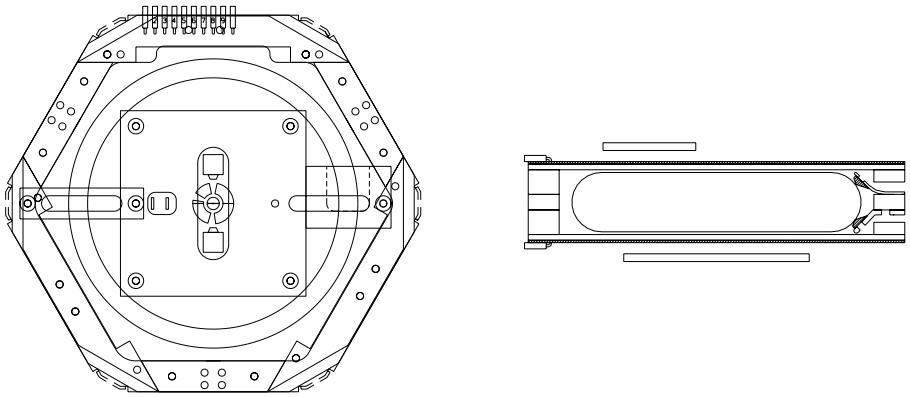

Figure 5.3: β-coffin. The detector is visible at the center. In the top view, the inner circle indicates the edge of the detector flat. The pie-shaped pattern is the thermal heat sink and electrical connection to the detector ground plane. The squares above and below indicate the thermistor positions. The pads to the left of center provide contact to the resistive boron implant for applying heat pulses. The open slot to the right allows light from the LED package (above the slot) to reach the detector. The phonon-side umbrella is the large square at the center. In the side view, the walls and top and bottom faces of the detector mount are shown. The floating copper pieces are the umbrellas (top is ionization).

detector. With the inner-electrode radius being 22 mm, the expected rate reduction for inner-electrode events due to an external low-energy-electron source as a function of detector separation is shown in Figure 5.5. A 3.5-mm detector separation is used, which is the smallest spacing that can accomodate the $\approx$ 2.5-mm thermistor thickness and the associated wirebonds, giving an expected reduction of a factor of 5. On the other hand, if the residual rate is due to surface contamination on the detectors themselves, no rate reduction is expected.

In addition to this "self-shielding," the close-packed configuration allows low-energy electrons to multiply scatter. The back-scattering probability for electrons in the energy range 10 to 100 keV is about 30% [87]. While this is not large enough that vetoing multiple-scatters would yield a significant rate reduction, it is large enough to see the effect. Furthermore, multiple-scatter phenomena can be used to check an alternate hypothesis, that low-energy electrons are ejected from surrounding materials by high-energy photons rather than being emitted by a β-decaying radioisotope. The former phenomenon was first noticed by Thushara Perera and Dan Akerib in their simulations of photon-calibration data sets [132]. The two phenomena are distinguishable by comparing the ionization yields observed for multiple-scatter events in adjacent detectors; the electrons-only mechanism results in low-yield events in both detectors, while the photon/electron mechanism should result in significant fraction of high-yield/low-yield pairs.

A third design change was to cover the inside of the copper detector housing with 2-mm-thick wafers of detector-grade germanium. The wafers are etched using the same procedure as for the detectors and thus are expected to be equally clean (or dirty). In fact, since no picene wax mask is used on the wafers, they may be even cleaner than the detectors. If the electron source is surface contamination of the detectors, there should be no reduction of the low-energy-electron

rate, while electrons emitted by surface contamination of the copper should be stopped by the Ge wafers. The photon/electron mechanism mentioned above would also show no reduction with the addition of Ge shielding. It is possible to cover approximately 90% of the exposed copper of the housing with germanium wafers, whose shapes are shown in Figure 5.6. The wafers are not rigidly attached because the relative thermal contractions of copper and germanium would cause the germanium to shatter. Rather, notches are cut in the germanium wafers, as if the germanium were to be screwed down to the copper, but the screws are left loose so the wafers are free to slide and contract.

Innovations were also introduced in the mechanical aspects of the design. Previous designs were not light-tight. The copper-foil sock used to block IR radiation must be disassembled and reassembled whenever a detector package is removed and reinstalled. The sock requires copper tape to seal cracks and so is time-consuming to build. In contrast, the close-packed assembly uses infrared-blocker boards to provide electrical feedthroughs while sealing the package against light leaks, as shown in Figure 5.7. An IR blocker consists of a copper-kapton-copper circuit board with metallized penetrating holes isolated electrically from the ground plane of the board. Sockets or pins are soldered into the holes to permit electrical signals to travel through the board. Even at the break between the holes and the ground plane, the circlex substrate of the board helps block IR radiation because it is capable of (though not particularly good at) absorbing IR [108]. All wirebonds are made to Detector Interface Boards (DIBs) *inside* the package; these boards have Mill-Max sockets soldered to them that penetrate the blocker boards. These penetrations are made light-tight by soldering the sockets into the blocker boards.

Another problem seen with previous designs is that the copper feet supporting the detector are thin (0.5 mm) and relax over time. This allows the detector to "jiggle" in the mount, yielding mechanical impulses that disturb the detector thermal system, adding noise, producing pulses, or, in some cases, making the detectors inoperable. In the close-packed mount, the detector sits on three 1.5-mm-thick ledges, as seen in Figure 5.4; the ring and ledges are machined from the same block of copper. On the top side, the detector is pushed against two stiff copper blocks with slanted faces that make contact to the curve of the detector, also shown in Figure 5.4. At the third point of the top face, a restraint with a similar profile to the stiff supports but with two tensioned Kevlar straps is pressed against the detector, shown in Figure 5.8. The Kevlar remains tensioned while cold. Since the germanium contracts less than the surrounding copper ring, the detector is held even more tightly while cold. Electrical isolation from the copper ledges and blocks is provided by 0.13-mm-thick (5-mil) kapton at all points of contact. The larger thermal conductance resulting from the increase in contact area appears to have had no serious effects on the phonon-channel pulse height or time constants, which were discussed in Chapter 4.

A final modification is the removal of all electrical components from the detector mount. The detector biasing and signal circuits use thin-film resistors and capacitors; connections are made using wirebonds. In the earlier mounts, these components are placed on the Y's above the detectors. Repair of damaged components requires that the whole detector mount be removed from the tower. Furthermore, the circuit boards are not protected. Given the large number (≈ 40) of components and bonds necessary, electrostatic damage to the components or accidental damage to wirebonds has occurred frequently enough to be inconvenient (two ionization coupling capacitors were shorted this way in Run 18). In the close-packed assembly, the circuit board holding all the components is moved to the basement side coax (see next section). The only electrical component

on the DIB is the LED used for detector neutralization; see Figure 5.9. The wirebonds between the detector and the DIB are completely protected once the detectors are stacked and the endcaps attached. The other electrical components lie under a protective cover on the Side Coax Assembly Board (or SCAB) at the detector end of the basement side coax; see Figure 5.10. The side coax can be assembled, all electrical components attached and wirebonded, and the whole assembly cryogenically tested independently of the detector mount.

One disadvantage of the close-packed design is that the capacitance between detectors becomes very large, leading to two cross-talk effects. First, the thermistor bias applies a $\sim$ 10-mV-rms sine wave across the thermistors, only 1 mm from the adjacent detector's ionization electrodes. There is significant cross-talk of this bias. This can be reduced by a factor of 10 by biasing the two thermistors 180 degrees out of phase, but perfect cancellation is not possible because the two thermistors posses somewhat different capacitances to the adjacent detector and different electrical characteristics. Second, a large ionization pulse in one detector produces significant cross-talk to adjacent detectors because the large inter-detector capacitance lets the detector draw current from the adjacent detector's ground plane. The capacitance between the inner electrode of one detector and the adjacent detector's ground plane can be estimated assuming a simple cylindrical capacitor with radius 22 mm (the inner-electrode radius) and plate separation 3.5 mm (the detector separation), giving $C = 4$ pF. The coupling capacitor to the amplifier has $C_c \approx 330$ pF, so $\sim 1\%$ cross-talk is expected. Thus, large pulses can disturb the adjacent detector. As is discussed in the next chapter, it is necessary to discard events where one detector has a large ionization pulse. Fortunately, WIMPs should interact in only one detector and produce $<$ 100-keV pulses, so such a cut has no effect on the experiment's sensitivity.

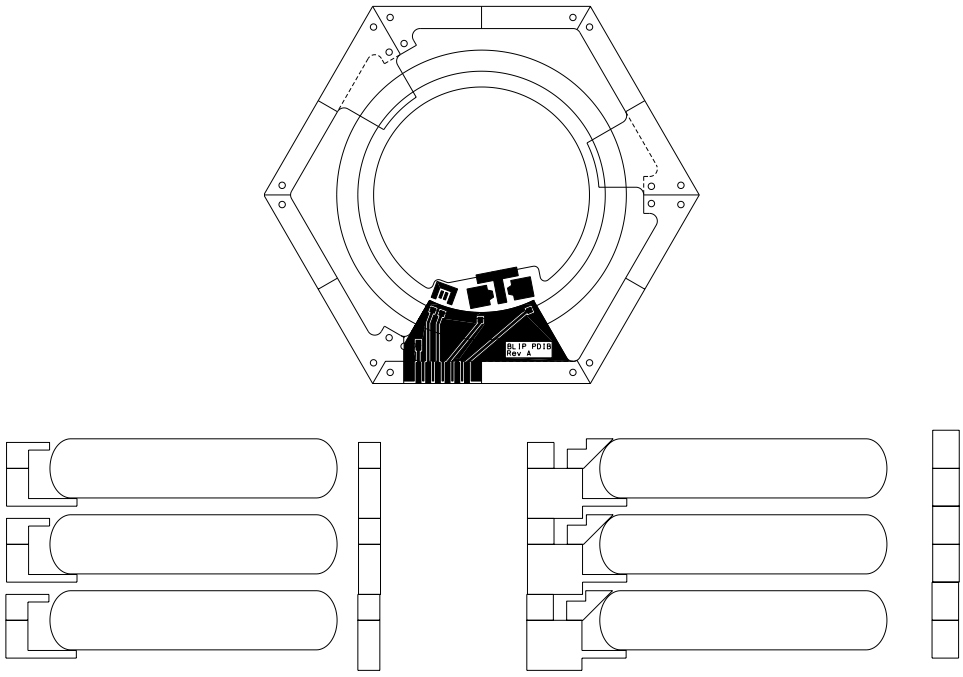

Figure 5.4: Close-packed detector mount. Top: top view, phonon side. The detector is shown in the center, phonon side up. The two circles indicate the edge of the detector and the edge of the detector flat. The circle with the indentation is the inner-outer-electrode break (etched into the contact on the ionization side). The three ledges on which the detector sits are shown (one is hidden underneath the DIB). The large squares on the detector are the thermistor positions. The "T" between them serves both as the connection to the detector ground plane and as the thermal heat sink. The two small rectangles inside the "U" make contact to a resistive patch of boron implant used for applying heat pulses. The phonon DIB is also shown. The chip-LED pair is connected to the left-most bond pad. Bottom left: side view along a midline running bottom to top. The phonon side is at the bottom. The detector separation is ≈ 3.5 mm. Bottom right: side view along a midline running left to right. A stiff support foot is shown.

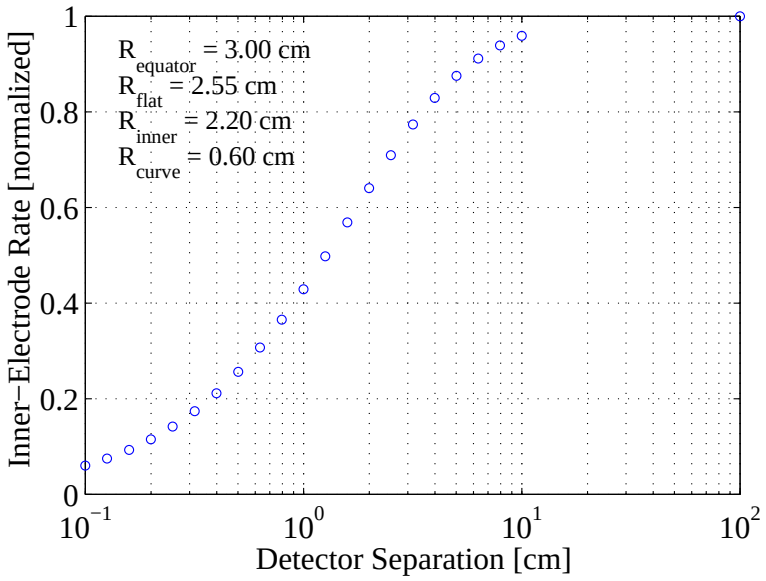

Figure 5.5: Expected inner-electrode event rate for an isotropic, external, non-penetrating radiation source as a function of close-packed detector separation. The detector parameters assumed are listed on the plot. The effect of the curved BLIP edges is taken into account. The rate is normalized to the rate in the absence of other detectors (separation $\rightarrow \infty$). The separation used for the close-packed BLIP package is 0.35 cm. Calculation by T. Shutt.

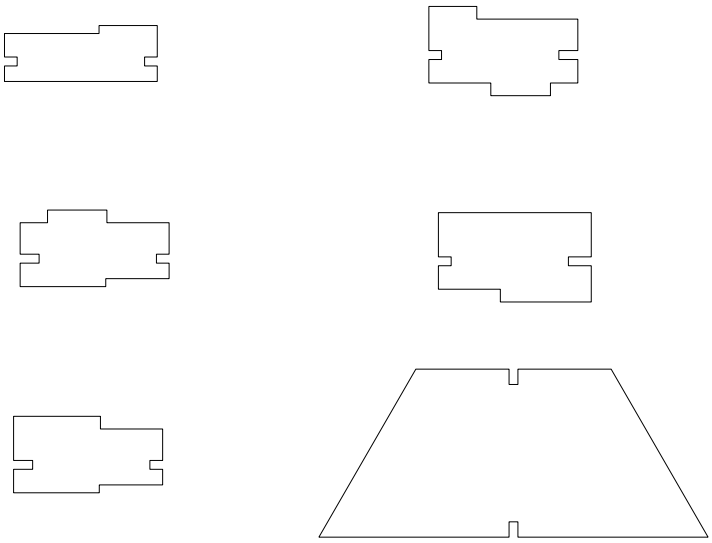

Figure 5.6: Germanium shield pieces. The lower right piece covers half of an endcap face; the other pieces cover inner faces of the detector rings (there are two ring faces that use the same type of piece). The notches accept 0-80 screws; they are slotted to allow for thermal contraction and for ease of cutting.

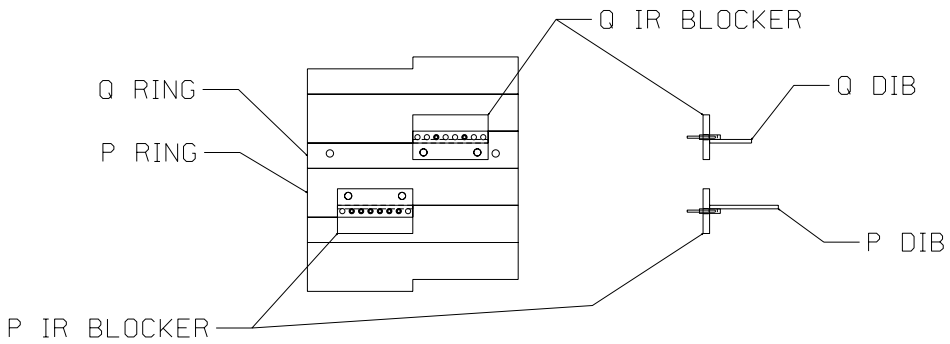

Figure 5.7: Face-on and side views of infrared-blocking electrical feedthroughs. The sockets feeding through the blockers are also soldered to the DIBs, as seen in the side view. The faces of the adjacent detector rings are shown to indicate how the blocker boards cover the holes in the ring walls. Unused blocker board holes are filled with solder.

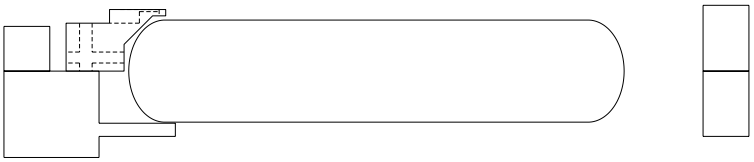

Figure 5.8: A detector and ring with a "kevlar" foot. Kevlar straps (not shown) loop through the holes in the foot and make contact to the detector on its curve.

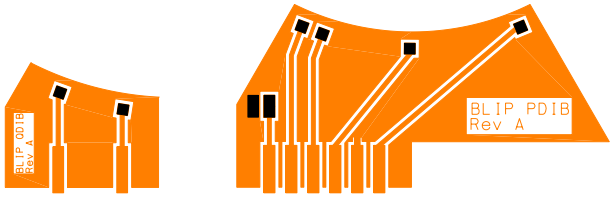

Figure 5.9: Detector Interface Boards (DIBs), Q (left) and P (right). The bond points are the large squares near the curved top edges. A chip LED is placed on each of the two leftmost pads. Mill-Max sockets attach to the large rectangles at the bottom edge.

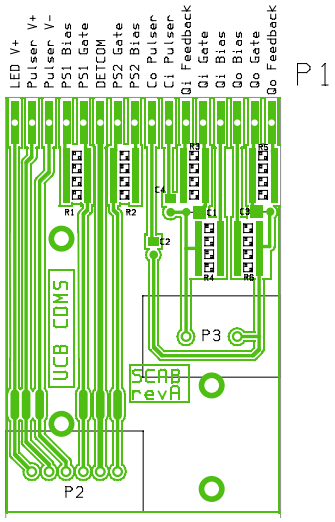

Figure 5.10: Side Coax Assembly Board (SCAB). The positions of the 40 MΩ resistors are indicated by squares with two dots (R1 to R6, 24 total). The ionization coupling and feedback capacitors sit on the squares marked C1–C4. The circular points marked P2 and P3 are the positions of the Mill-Max pins that mate to the sockets on the DIBs. P1 indicates the points where the tensioned NbTi wires of the basement side coax are attached; the signal names are given. Figure taken from [133].

5.3 Tower and Cold Electronics

The detector tower serves a number of purposes. It:

- provides a removable mount for the detectors and cryogenic electronic components;
- provides low-capacitance, microphonically clean, electrical connections between the FET card and the detectors;
- provides a path for the power dissipated by the front-end FETs (approximately 7 mW per FET) to flow to the LHe bath;
- heat-sinks the aforementioned connecting wires to maintain reasonable power loads on the MC, CP, and ST refrigerator stages in spite of the large power dissipation at the 4 K stage and the natural power flow between the various temperature stages;
- and is modular: many components (detectors, basement side coaxes, FET cards) can be removed and replaced without significant disassembly. All connections are made using high-quality connectors: soldering or screw terminals are not used for connections that are frequently remade.

I discuss qualitatively the tower design. Walter Stockwell's dissertation presents the numbers behind the design [108].

5.3.1 Tower Design

A schematic of the tower is shown in Figure 5.1. This side view shows the four temperature stages. Each temperature stage is machined from a single piece of copper and is mostly empty space. The four stages are separated by graphite tubes. The tower is hexagonal in cross section. This design permits close packing of many towers. A view of a single face is shown in Figure 5.11. Each face has 16 vertical channels in which wires a strung. Also visible are the heat-sink boards at each stage.

One purpose of the tower is to provide low-capacitance, microphonically clean, electrical connections. As discussed in Chapter 4, capacitance to ground between the detectors and the FET gates is problematic. The biggest capacitance in the readout circuit is the FET-gate capacitance of 40–60 pF, so any wiring capacitances should be small in comparison. This requirement rules out standard cryogenic coaxial cables, which have capacitances of tens of pF per foot.

Standard coaxial cables are also disallowed because they are susceptible to microphonic noise. The usual mechanism by which microphonics produces noise is discussed in Tom Shutt's dissertation [106]. Microphonics arises from the triboelectric effect, in which a difference in the work functions of two materials in contact results in the buildup of charge at the interface [134]. The interface between the inner conductor and dielectric in a coaxial cable is susceptible to this effect. It is difficult for the surface charge on the insulator to dissipate because it is insulating. If there is microscopic relative motion of the conductor and the dielectric, some of the surface charge may be coupled directly or capacitively into conductor. Microphonics thus acts like a current source. Given the high impedances seen by the lines connecting the detector to the front-end FETs in both the BLIP ionization and phonon measurements, this current source can produce a significant voltage

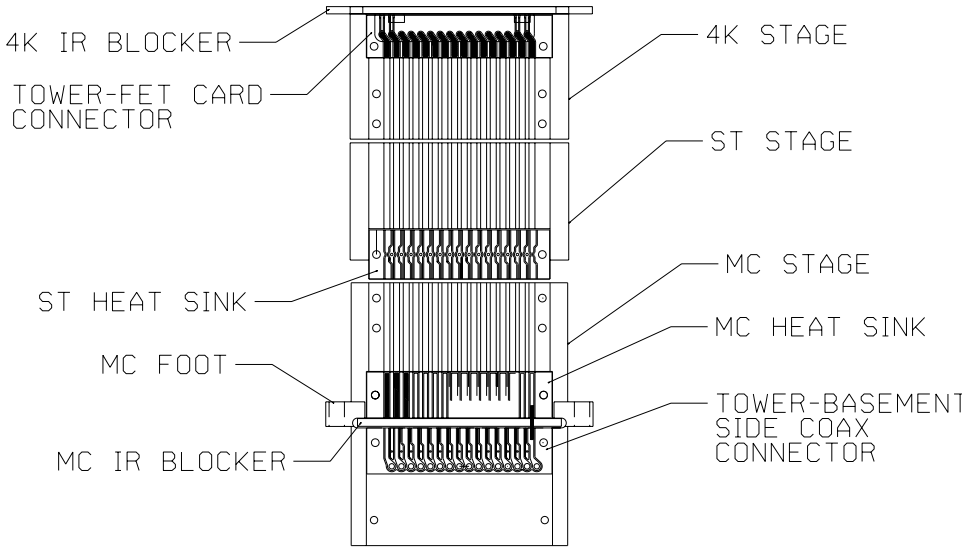

Figure 5.11: A tower face. The wire channels are visible on each stage's face. The gaps between stages provide thermal isolation between stages; the mechanical connection is through graphite tubes (behind the face, not shown), as described in the text. The CP stage has no face (it only has feet) and so is not shown (its feet stick out through the indentations in the ST stage). The ST-stage feet are also not shown. The ST and MC heat sinks and the tower-basement side coax and tower-FET card connectors are shown. The tower IR blockers are seen edge-on.

signal. In contrast, for low-impedance devices and measurement schemes, such as a transition-edge-sensor coupled to a SQUID current amplifier, microphonics is irrelevant because the current coupled is very small.

The main driving force for microphonics in a dilution refrigerator is the flow of helium through the needle valve that feeds the 1 K pot from the LHe bath. The resulting vibrations occur at many kHz to 20 or 30 kHz. Microphonics is primarily a problem for the ionization measurement because of its large bandwidth. It can be a problem for the thermistor measurement if it extends down to 1 kHz. The obvious solution to microphonics is to remove all insulators. The tower design accomplishes this by stringing tensioned wires in metal channels. Metal covers are placed over the channels to complete the outer conductor and to protect the wires. Though these channels can be thought of as transmission lines at higher frequencies, it is not particularly useful to do so in this case; the travel time for an electromagnetic wave down the tower is less than 1 ns, much shorter than the amplifier response times, so the tower can be treated electrically as a lumped element.

The heat-sinking problem brought on by the above solution is challenging. Given the need to tension the wires, long wire runs must be mechanically supported along their entire path; *i.e.*, one cannot simply use a long, flexible cable to provide a large thermal impedance between stages. Given the relatively short distance between the lids of the various cans in the Icebox (1 to 2 inches, typically), the available vertical height for such wire runs is small. Thus, it is necessary to repeatedly heat-sink the wires running from the 4 K to the MC stage to maintain a low power load

on the MC stage. Heat-sink boards, consisting of 52-mil-thick cirlex (thick kapton) with copper on both sides, are used. The top side of each board is patterned with traces aligned with the wire channels. As the wires pass over the board, they are soldered down. Heat carried down the wire can flow through the cirlex to the underlying tower stage rather than flowing down the wire to lower-temperature stages. These boards are shown in Figure 5.11. An important part of the heat-sinking scheme is the use of NbTi wires, which superconduct at 9.5 K. When a superconductor is significantly below its transition temperature, a large fraction of the electrons are bound in Cooper pairs that form a Bose condensate; these pairs cannot transport heat. This greatly reduces the thermal conductivity of the wire; only phonons and broken pairs can transport heat.

Another problem brought on by the short stage separations in the Icebox is that the temperature stages of the tower are themselves separated by only 1 to 2 inches also. Graphite tubes are used to separate the stages of the tower because of their intrinsically low thermal conductance. G-10 would have been preferable, but is disallowed because of its high radioactivity. Because the thermal conductance of a tube goes like A/h where A is the cross-sectional area (of the graphite alone) and h is the tube height, it is desirable to minimize the tube-wall thickness and maximize the height. The tubes have a wall thickness of 1.5 mm (60 mils), set by what can be reasonably machined. The height thus determines the conductance. As one can see, the tower stages have been designed to increase the tube heights beyond what is allowed by the Icebox temperature-stage separations alone. The CP stage of the tower is used only as a thermal intercept for the graphite tubes; wires are not heat-sunk at this stage.

Another issue is the thermal connection of the various temperature stages of the tower to the corresponding stages of the Icebox. As discussed in Chapter 3, the Icebox cans are not rigidly connected to one another. Therefore, the tower may be connected to the Icebox rigidly at only one layer. The MC stage of the tower is rigidly connected to the MC can lid of the Icebox at the six "feet" shown in Figure 5.1. The other temperature stages are connected to the corresponding can lids by flexible heat straps made from welding cable. The conductivity requirements on these straps are nontrivial, especially at the 4 K stage due to the large power dissipation of the FETs.

Because the wires travel down empty channels, these channels act as light leaks for the blackbody radiation emitted by the FETs. To address this problem, infrared-blocking boards are implemented at the MC and 4 K stages. The tower IR blockers are very similar to those used in the detector mounts. Blockers of this type are installed at the bottom of the FET card, where it plugs into the tower, and at the MC stage of the tower, between the heat sink that receives wires from warmer stages and the connector to the basement side coax. The MC-stage blocker abuts the MC-stage tower-mating flange, sealing the gap. At the 4 K stage of both the Icebox and the test facility refrigerators, copper-kapton-copper radiation shields that overlap the 4 K IR blockers are attached to the corresponding flanges; these shields seal the gap at 4 K without a rigid connection to the tower.

Finally, to provide convenient mate-demate connectors for the FET card at the top of the tower and the basement side coax at the bottom, Mill-Max sockets are soldered to the boards. At the 4 K stage, the sockets are laid on their sides so the socket openings point upward toward the FET card; these mate to the pins sticking through the IR blocker board attached to the bottom of the FET card. This design allows FET cards to be easily mated to or demated from a tower that is already installed by moving the cards vertically. In practice, this is important when installing a tower with many detectors. The sockets at the MC stage point horizontally outward from the

tower itself. Pins on the basement side coax mate to these sockets.

5.3.2 Tower Construction

Assembling the tower proper is fairly trivial. Circular grooves at the points of attachment of the graphite tubes to the tower stages locate the tubes. Stycast 1266 epoxy is applied in the grooves using a syringe, the tubes placed in the grooves, and the stages stacked vertically. Alignment fixtures are used to ensure all four stages are centered and azimuthally aligned. The gap between stages is set by inserting spacers at the faces. Assembly is performed in two steps, epoxying the bottom ends of the tubes in one pass and then the top ends after the first set of epoxy joints have cured; this prevents epoxy from leaking out of the grooves and running down the tubes while upside-down.

Wiring the tower is nontrivial. There are many issues to consider. First, the wires must be selectively etched. Solder does not wet to NbTi, so the wires are clad in Cu, which is solderable. This cladding is a pure (not alloyed) normal metal, so it is a good thermal conductor. The cladding must therefore be removed on the sections of wires between heat sinks, which can only be done by etching with nitric acid. Second, the wires must be tensioned: since there is no insulator in the wire channels, the wires must be prevented from touching the channel walls. Copper shrinks more than NbTi when cooled, so the wires must be overtensioned to remain taut when cold. Third, even though low-activity solder is used [135], it is undesirable to perform soldering on the tower itself: flux may leak under the heat-sink boards, and even into the tower itself. Any solvents used to clean solder joints inevitably leak into the tower also.

To address these problems in a way amenable to mass production, Garth Smith and I developed a new wiring procedure in preparation for Run 19. It was realized that the wires themselves are strong enough to support the heat sinks. Thus, rather than solder the wires to the heat sinks on the tower, soldering is done away from the tower and the cleaned wire/heat-sink assembly transferred to the tower. A fixture was designed to place the heat-sink boards at their nominal separations. NbTi wires are tensioned and soldered to the heat-sink boards and to printed-circuit boards on a removable transfer frame that surrounds the fixture. The heat sinks are then unscrewed from the fixture and the transfer frame disconnected, leaving a transfer frame with 16 wires strung across it and with heat sinks suspended from the wires. Apiezon Wax W [129] is thinned with toluene and painted over the heat sinks to protect them during the etch step. The thinned wax has a high surface tension and sticks to the heat sinks but does not wick onto the wires between heat sinks; when dry, the wax forms a solid mask. The entire assembly is dipped in nitric acid and the copper cladding on the wires between the heat sinks is etched away. The mask is removed using xylene. The etched assembly is cleaned in alcohol and laid over a tower face. The heat sinks are screwed down. With the heat sinks in place, the wires to the transfer frame can be clipped, leaving a fully wired tower face in place. In addition to satisfying the aforementioned needs, this procedure leaves the tower unoccupied during the soldering and etching steps (which take 2 to 3 days). Multiple faces can be prepared simultaneously, or replacement faces can be prepared while a tower with a damaged face continues to be used.

5.3.3 Basement Components

As mentioned earlier, a wire assembly, named the "basement side coax," is used to connect the detector mounts to the wires at the MC stage of the tower. A basement side coax is shown in Figure 5.12. The design is essentially the same as that of the tower side coaxes. Vertical channels are cut in a block of copper. Copper-kapton circuit boards are placed at the two ends to receive the wires. Prior to assembly, Mill-Max pins are soldered into both circuit boards; these mate to the sockets presented by the DIB/IR-blocker assemblies in the detector mount and the tower MC-stage connector. The basement side coax wiring procedure is similar to the tower wiring procedure.

As mentioned earlier, the circuit board at the detector end, the SCAB, also holds electronic components for the detector readout. These are attached and bonded after the side coax is assembled. Covers are placed over the SCAB and over the wire channels for mechanical and electrical protection. The covers are intricately machined with recesses for the components to minimize inter-trace capacitances.

A final point is that, even though the temperatures of the two ends of the basement side coax are the same, the wires are etched anyways. At the MC stage, a second stage of heat sinking is necessary to ensure the electrical lines connecting directly to the detector are cold. No heat sink is perfect, so even though the wires are heat-sunk to the MC stage of the tower, the electrons in the wire are not fully cooled. The only thermal isolation in the electrical path between the SCAB and the detector is aluminum wirebonds, which are rather short and thus do not provide a large impedance (in spite of being superconducting). Thus, if the electrons in the traces on the SCAB were slightly elevated in temperature due to power flow down the wires, they would deposit this power directly into the thermistors. The thermal impedance provided by the long (3 inches or more) run of etched NbTi wire in the basement side coax ensures that the power flow into the thermistors is negligible.

In the design used in Run 19, the basement side coaxes are machined from 0.25-inch-thick copper blocks. They are quite strong and so are used to provide structural connection between the tower and the detector package. In hindsight, this proves to be somewhat problematic because stress is placed on the basement side coax pins while attaching the detector package to the tower. This issue has been remedied in the basement design for Run 20 by the use of a separate mechanical support for the detector package.

Finally, for the Run 19 and previous towers, a hexagon of Nantes lead (old lead, whose ^{210}Pb content has decayed away) is screwed to the bottom of the MC stage of the tower to block the hole in the inner lead shield through which the tower is inserted. It is not clear whether this is necessary; for convenience, it has been abandoned for the Run 20 tower.

5.3.4 FET Cards

A central component of the detector readout is the FETs mounted on FET cards at the 4 K stage of the tower. As discussed in Chapter 4, the FET gates are the input terminals of the ionization and voltage amplifiers. It is advantageous to mount the FETs inside the cryostat. This reduces the length of wire between the detector and the FETs and thus the capacitance at the FET gate. Also, FET noise improves as the FET is cooled because it is Johnson noise of the drain-source channel [115]. Below ≈ 100 K, the carriers freeze out and the FET stops operating. For the FETs used, InterFET IF4500 [115], selected for their noise characteristics, the optimal operating

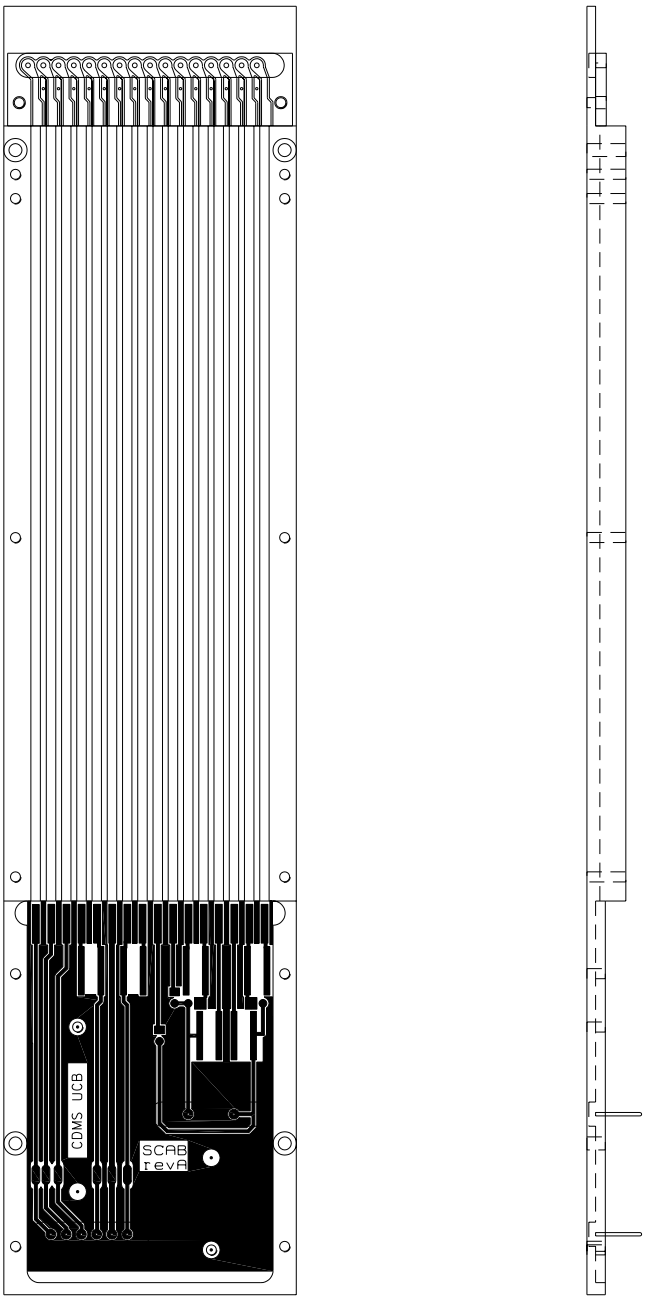

Figure 5.12: A basement side coax. The circuit board at the top hold pins for connection to the tower. The bottom circuit board is the SCAB (see Figure 5.10 for a detail view). The pins mating to the DIBs are shown in the side view.

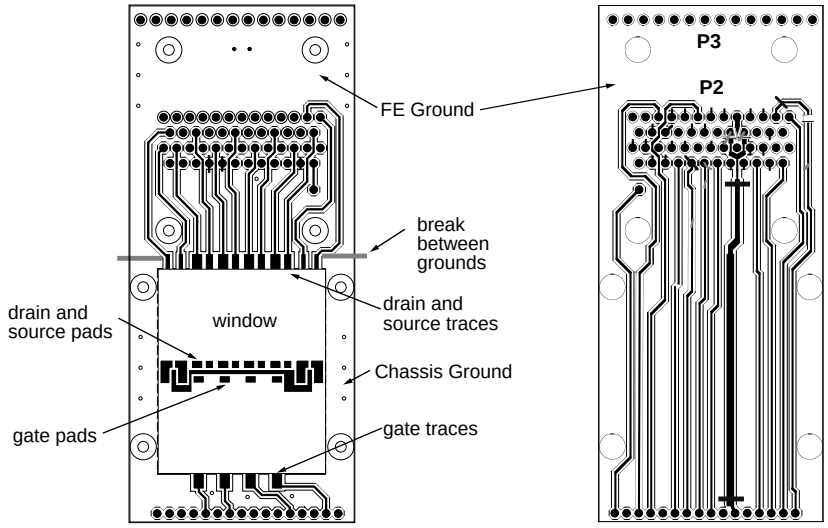

Figure 5.13: FET card, front (left) and back (right). The FETs are soldered to the copper pads at the center of the window. The vertical edges of the window are cut. The window is covered by the FET card gusset. Figure taken from [133].

temperature is approximately 120 K [108].

FET-Card Design

A schematic of the FET card is shown in Figure 5.13. The card is built up from many layers of kapton and adhesive (these cards were fabricated before cirlex was available). There are three layers of copper: one at the center and one on each face. The center layer is a ground plane. A window in the board is made during fabrication and a single layer of kapton laid over the board before the top copper layer is deposited. The edges of the window are cut so the window acts as a thermal standoff. Copper pads at the center of the window provide mount points for FETs, which are soldered in place. Platinum traces are sputtered on the window to link the pads to traces on the rest of the board. The thin platinum traces and insulating window provide a large thermal impedance, allowing the FETs to self-heat to ~ 120 K in the center of the window while the board is mounted to 4 K. A heater and a temperature-sense diode are also placed on the window near the FETs; the heater is needed to preheat the window when first turning on the amplifiers (the FETs are frozen out at 4 K and do not function until preheated). The temperature-sense diode is part of a feedback loop with the heater to fix the window temperature at its optimal value, though in practice this loop has not been used.

The remainder of the front side of the card carries the gate traces from the bottom of the window to the tower-connector pins at the bottom of the card and the drain and source traces to the stripline connector. The back side of the card carries the detector-biasing, detector-pulsing, and LED traces from the stripline connector to the tower connector. Some subtleties arise regarding grounding; there are two different grounds on the FET card, a chassis ground connected directly to

the metal of the tower, and a front-end (FE) ground connected to FE-ground traces on the stripline but isolated from chassis. The rationale for these is discussed in Section 5.4.

To protect the FETs and to mount the card to the tower, a right-angle copper gusset attaches to the front face of the FET card; see Figure 5.1. The inside of the gusset is painted with a Stycast 1266 loaded with carbon black to absorb infrared radiation emitted by the FETs; this mixture has been used by groups using cryogenic bolometers for astronomical measurements for many years [136]. Small gaps between the gusset and the card are blocked with thin strips of kapton to electrically insulate signal traces leaving the FET window from the gusset and to block IR emission.

The soldering of the FETs in place is recent innovation. For years, it had been believed that soldering these sensitive FETs damages them. Exhaustive tests by Tom Shutt and Storn White showed that there is no evidence for this hypothesis, provided the soldering is done quickly and carefully. Furthermore, solder joints are seen to be far more reliable than joints made using silver epoxy, especially with respect to $1/f$ noise usually attributed to poor electrical connections. All the FET cards used for BLIPs 3 through 6 have their FETs soldered in place; none of these 16 FETs shows anomalous noise that can be attributed to damage from soldering.

FET-Card Microphonics

A nonstandard type of microphonics was discovered in the FET cards, with the cause finally isolated and a fix found by Tom Shutt and Storn White. The platinum traces on the window must overlap copper pads at the end of the window to make electrical contact to the traces on the card. However, the platinum traces are thin (3 μm) and must run over a step of height 71 μm at the edge of the copper pad. This step occurs at the edge of the window, so motion of the window stresses this joint. It was shown that this movement gives rise to microphonics by modulating the resistance of the joint. This result is somewhat surprising because no gross change in the resistance is visible except under extreme stress. Small changes in the resistance, of order tens of $\mu\Omega$, can result in a noise signal on the *source* line because of the large ($\sim$ 2 mA) DC current flowing through the FET channel. A semi-permanent fix can be made by placing a cirlex spacer underneath the window at its edge under the joint and applying solder over the joint. This results in a slight decrease in the thermal impedance but provides a rigid support for the joint. This solution was applied to the FET cards used for Run 19 and was seen to work consistently: the Run 19 FET cards are free of microphonics. The FET cards designed for ZIP-detector operation remove this issue by moving the copper pad further back over the cirlex so the joint is always supported.

5.3.5 Striplines

The drain and source connections to the cryogenic FETs are carried to room temperature by copper-kapton striplines. These striplines also carry detector-biasing, detector-pulsing, and LED lines down from the warm electronics. The primary design criteria for the stripline were length, thermal conduction, and resistance. The Icebox's E-stem length is set by the size of the shield and veto and the need to provide a significant thermal impedance between the room-temperature electrical feedthrough and the LHe layer of the Icebox. This length sets the length of the striplines. The cross-sectional area and length of the traces are determined by the desired power flow from room temperature to the LHe layer. There is also a maximum resistance requirement on the

FET-source traces, $\approx 20\,\Omega$, above which the Johnson noise of the source trace dominates over the FET-gate noise.

The striplines are a natural extension of flexible-circuit technology used in printers and plotters. However, the extreme length and multiple layers push the technology. A stripline is 115 in (2.92 m) long, 1 in (2.54 cm) wide, and 18 mils (0.45 mm) thick [137]. The stripline carries 50 traces, etched out of 1/2-oz (0.0007 mils, or 0.018 mm) copper. The traces are 5 mils (0.13 mm) wide and separated by 10 mils (0.25 mm) for the bulk of their length. For the last 1 m on the warm end, seven traces for the front-end-ground and FET-source traces are widened to 35 mils (0.89 mm) to decrease their resistance. Two complete, conductive shield layers are fabricated by sputtering 1 μm copper onto kapton; they are overlaid with kapton for protection. The shield layer is very thin in order to maintain low thermal conductance. The traces meet connector patterns on each end; these connectors mate to the FET card and the room-temperature hermetic electrical feedthrough. In practice, about half of the 50 traces are assigned to front-end ground, interleaved with signal lines, to isolate the signal lines from each other and to decrease the front-end-ground resistance. Each stripline deposits a power load of about 12 mW on the LHe layer of the Icebox [104]. Finally, though they are called "striplines," the transmission-line properties are not of interest because the bandwidth of the detector signals is small compared to the signal travel time of tens of ns (the dielectric constant may be large, as for room-temperature coaxial cables). The stripline capacitance and inductance introduce some interesting effects in the ionization-amplifier feedback circuit, but the stripline may still be treated as a lumped element.

5.4 Warm Front-End Electronics

Because of the unique detector-biasing circuitry and noise requirements, CDMS uses custom-designed front-end-electronics boards. A 9U-format, 18-slot electronics crate with a custom backplane houses the front-end boards, one for each detector. Signals come in on 50-conductor, twisted-pair "Detector I/O" cables that connect from the Icebox hermetic electrical feedthrough to a connector on the rear of the board. Power-supply voltages are available on the backplane, carried in from a separate power-supply unit situated near the crate. The amplified analog signals are output onto a 25-pin connector that mates to a 50-foot-long, twisted-pair cable that runs to the Receiver/Trigger/Filter rack (discussed in Chapter 6). Latches and registers on each board are addressed and data read or written using a standard address/data-bus configuration on the backplane. The backplane digital bus is connected to a single "Digital I/O" board in slot 19 of the rack. The front of the Digital I/O board connects via a cable to a "GPIB Interface Box" in a nearby rack. This box contains a IOTech Digital-488 GPIB-controllable I/O register board. A computer communicates via this GPIB interface with the Digital I/O board to set the state of the front-end boards.

I briefly describe the amplifier and detector-biasing circuits as well as the other functionality of these boards. The amplifier design used to read out BLIP detectors has been discussed in great detail in a paper by Dominique Yvon [138] and in Tom Shutt's dissertation [106]. Since analysis of these circuits, made of a number of discrete components, is rather complicated and has been presented fully in these references, I do not discuss these issues here.

5.4.1 System Design Aspects

An interesting system design feature is the use of multiple power supplies and ground returns. There are three different power-supply/ground-return sets in the front-end-electronics system: front-end ground, analog ground, digital ground. Chassis ground is an additional ground system. All four grounds are tied together at the enclosure containing the front-end-electronics racks. The purpose in having multiple supplies is to prevent large, fast signals being driven on the board from disturbing the front-end amplifiers. The front-end amplifiers, are by design, extremely sensitive: the voltage amplifier is capable of sensing pulses with peak heights of less than 100 nV with signal-to-noise of better than 1. An example where multiple supplies is critical is the thermistor-bias circuit. It switches a 10-V-amplitude square wave at 1 kHz in the initial stages of generating the thermistor bias. The large currents involved can result in voltage drops across the power-supply and ground planes. If the amplifiers used the same supply, the fluctuating power-supply and ground voltages would inject noise into the amplifier circuits. This type of effect is prevented (or, at least, mitigated) by producing the large-amplitude square wave bias in a circuit connected to the analog power supplies and grounds and only shipping it to the front-end power-supply and ground system just before it goes into the stripline, after it has been filtered into a much "friendlier" sine wave. In spite of these precautions, there is still significant cross-talk of the initial 1-kHz square wave to the ionization amplifiers.

Another aspect is the use of digital control on these low-noise analog boards. While convenient for configuring the boards, care must be taken to ensure that the digital circuits do not inject noise into the analog circuits. No clock signal is used on the digital backplane; read and write strobes synchronize address/data-bus operations. Therefore, when no configuration changes are being made, there is no digital traffic on the backplane or boards. A separate digital power supply is used to prevent power-supply and ground currents from disturbing the amplifiers. The sharp edges of digital signals from the Digital I/O board are slowed before distribution on the backplane to prevent capacitive cross-talk. Data acquisition is always halted when configuration changes are made. The double-buffering of the GPIB-control commands through the GPIB Interface Box and Digital I/O board isolates the front-end boards from the GPIB. A fiber-optic GPIB extender connects the GPIB Interface Box to the data-acquisition computers, electrically isolating the front-end digital control.

5.4.2 Front-End Amplifiers

I discuss here the basic operating principles and noise response of the front-end voltage and ionization amplifiers. Detailed discussion of how these noises affect the respective measurements has been presented in Chapter 4.

The voltage amplifier is source-follower configuration which acts, in practice, as a non-inverting op-amp configuration; see Figure 5.14. The FET gate corresponds to the noninverting op-amp terminal, the source to the inverting terminal, and the amplifier output to the op-amp output. The feedback resistors in the simplified schematic perform their same function in the real amplifier. The action of a FET is to change the channel conductance in response to a change in $V_{GS} = V_G - V_S$, where V_G is the gate voltage and V_S is the source voltage [139]. The amplifier senses this conductance change as a change in the drain current; it amplifies this change and applies the amplified current to the feedback-resistor network shown in Figure 5.14. This changes the source

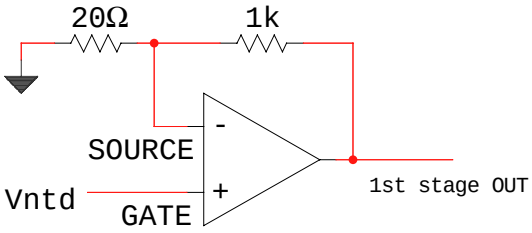

Figure 5.14: Op-amp analogue of voltage amplifier. The source and feedback resistors, are shown, $R_s = 20\,\Omega$ and $R_f = 1\,\mathrm{k}\Omega$. Figure taken from [133].

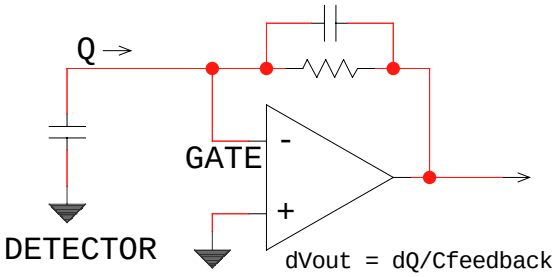

Figure 5.15: Op-amp analogue of ionization amplifier. $C_f \approx 1\,\mathrm{pF}$ and $R_f = 40\,\mathrm{M}\Omega$. Figure taken from [133].

voltage in the direction needed to make V_S track V_G, thereby returning V_{GS} to its quiescent value. Hence the name "source-follower." The amplification factor is $G = (R_s + R_f)/R_s \approx 50$. A second stage of amplification is provided by a standard inverting amplifier with $G = 10$. The main noise sources are the gate noise (the channel Johnson noise referred to the gate) and any noise on the source — Johnson noise of the source resistor, variations in the source resistance, etc. Noises at the gate and source can be modeled as noise sources at the noninverting and inverting inputs of the noninverting op-amp configuration.

The ionization amplifier is similar to an op-amp integrator; see Figure 5.15. The gate now corresponds to the inverting input of the op-amp. The amplifier output corresponds to the op-amp output. The source corresponds to the noninverting input and is grounded. Thus, gate-voltage changes cannot be nulled by moving the source. Instead, the amplifier provides current to the detector to prevent the gate from moving from ground. Viewed another way, an event puts charge on the detector (and FET-gate) capacitance, causing the gate voltage to move away from ground. This charge must be removed from the detector to keep the gate at ground. The amplifier senses the change in FET drain current when the gate moves away from ground; the output of the amplifier swings as is necessary to counter the movement of the gate. For a δ-function current pulse containing total charge Q, the voltage at the output of the circuit is $V = Q/C_f$. The feedback resistor R_f shorts the capacitor on long timescales, producing a decaying output pulse with time constant $R_f C_f$. A second stage amplifies this voltage pulse by a factor of 200. Noise analysis can

be done with the op-amp integrator model of the circuit and has been discussed in Chapter 4.

One interesting feature realized recently is that the speed and stability of the ionization-amplifier circuit is limited by the output impedance of the amplifier. The present configuration has an output impedance of approximately 5 kΩ. This forms a RC with the stripline capacitance ($\sim 100 - 200$ pF). If the amplifier is sped up by modifying its intrinsic response time, the system oscillates. It is believed that this is because of the phase shift due to the RC. Merle Haldeman and Sae Woo Nam have implemented a version of the amplifier with a lower output impedance and have observed that the new configuration can be sped up by a factor of a few without oscillating.

5.4.3 Detector Biasing

The primary elements of detector biasing have been discussed in Chapter 4. I discuss here how the voltage sources for the thermistor and ionization bias are implemented.

The ionization bias is straightforward. A computer-controlled DAC provides a reference value for the voltage. This voltage is buffered to switch it over to front-end ground; it is then filtered with a passive RC filter and an active op-amp integrator. As is discussed in Chapter 4, the ionization voltage bias is irrelevant on pulse timescales: the detector itself provides the bias. The purpose of the ionization bias is to recharge the detector capacitance between events. Thus, the filtering can be quite aggressive to ensure no noise is added by the bias; the RC time of the filter and the integrator is 100 ms.

The phonon bias is more complex. A clean sine wave is generated on the board as follows. A DAC is set to specify the voltage of the sine wave, up to a fixed numerical factor. The DAC output is sent to the analog input of an AD630 modulator/demodulator IC, which operates, essentially, by switching an inverter in and out of the signal path based on the zero crossings of its control input. A low-amplitude 1-kHz sine wave from the backplane provides these zero crossings. The DC output of the DAC is thus converted to a 1-kHz square wave. The square wave is filtered using a 4-pole low-pass Butterworth filter with a 3 dB point of 1 kHz to produce a 1-kHz sine wave. The sine wave is buffered and switched to the front-end ground circuit. Just before sending it to the stripline, the bias is divided using a passive resistive divider. The bias is kept large while on the board to make sure that the output noise of all the components in the production path remains small compared to the bias; the bias and noise are divided by the same factor by the passive output divider, maintaining the high signal-to-noise ratio. This voltage is applied to the top of a divider network consisting of a cold 40 MΩ resistor and the thermistor, producing a current bias for the thermistor. The biases for the two thermistors on a given detector are operated 180 degrees out of phase; the resulting cancellation reduces by a factor of 10 the cross-talk to the ionization circuits on the board, on the detector, and to adjacent detectors.

5.4.4 Phonon Lockin

The lockin method used for measurement of thermistor pulses was described in Chapter 4. The details of the implementation are as follows. The signal output by the voltage amplifier is still in modulated form. This signal is filtered using a 2-pole low-pass filter with a cutoff frequency of 1.85 kHz to remove the third and higher harmonics of 1 kHz. The resulting signal is demodulated by an AD630 modulator/demodulator IC, switching on the zero crossings of the signal itself. The DC component of this rectified signal is determined by an integrator circuit (with 66-s time constant)

and subtracted from the rectified signal. This provides a "DC-ref" signal that is proportional to the rms value of the modulated signal; in the limit that RC rolloff at the thermistor is unimportant, this is proportional to the thermistor resistance. The now AC-coupled signal is filtered by a 8-pole low-pass Butterworth filter with a cutoff frequency of 450 Hz. As discussed earlier, this filters out the components of the signal that are shifted to 2 kHz, 4 kHz, etc. by the simple demodulation scheme. The resulting 450-Hz bandwidth is sufficient to faithfully reproduce the phonon-pulse shape for interactions in the crystal; however, thermistor events are bandwidth-limited. The output of the Butterworth filter is the pulse that is digitized.

Because the DC-level of the output of the demodulator is set by the carrier amplitude, small variations in the carrier amplitude appear as baseline drifts. The DC ref is subtracted from the demodulator output with a 66-s time constant, so long-timescale drifts are taken out and the baseline remains constant. Short timescale drifts are not corrected. If the detector temperature is very unstable for some reason, the signal baseline may wander. Events for which the signal baseline is outside a reasonable range are discarded, as is discussed in Chapter 7.

5.4.5 IBAPACAP

The IBAPACAP, or Implant Biasing And Pulsing and ChArge Pulsing, circuits produce pulses to mimic events. As mentioned in Chapter 4, a small resistive heater ($\sim 100\,\Omega$) on the detector can be used for DC heating or to produce heat pulses. Pulser capacitors are placed at the gates of the ionization-amplifier FETs; step-function voltage pulses into these capacitors present δ-function current pulses to the ionization amplifiers. Both pulser circuits are initialized and enabled via the digital interface. It is undesirable to use this interface to trigger the IBAPACAP because the digital traffic would likely cross-talk to the analog circuits and could even cause heating in the detectors. Instead, a trigger signal sent out on the GPIB is transmitted from the GPIB Interface Box on a separate, twisted-pair cable and received differentially by an auxiliary board in the electronics crate, where it is converted to a square analog pulse with slow edges and transmitted onto the backplane. Once received on each board, it is reconditioned using a Schmitt trigger. The reason for the slow backplane pulse is that, with 18 boards receiving the signal, the large line capacitance would require high-speed currents to be driven, likely resulting in cross-talk. Once the signal is received on each board, it is possible to speed up the edges again because the trace lengths and capacitances on the boards are small.

The phonon pulser operates as follows. The pulser consists of a small patch of resistive material on the surface of the detector, electrically isolated from the ionization electrodes and thermistors. A DAC sets the amplitude of the pulses to be applied to the heater pads. A one-shot sets the length of the pulse to be 50 μs. The pulse must be short compared to the rise time of the phonon pulse to accurately simulate a δ-function energy deposition. The sharp edges of the one-shot pulse are filtered using a 4-pole low-pass filter with a cutoff frequency of 100 kHz to prevent capacitive cross-talk to the phonon and ionization circuits. The filtered pulse is sent differentially down the pulser lines after a final passive resistive-divider stage: pulses of the same size but opposite sign are applied to opposite ends of the pulser pad on the detector to minimize cross-talk to other lines.

The phonon pulser produces pulses of a fixed amplitude at known times. At the start of a run, this is useful for optimizing the phonon bias value. The optimal bias is a complex function of refrigerator temperature, sensor characteristics, and the other sensor's bias, and so is difficult

to calculate from first principles. The phonon pulser is used to map out the phonon pulse height versus bias empirically. During normal data-acquisition, pulser runs at intervals of a few hours calibrate the relationship between the thermistor DC ref and pulse height, allowing corrections for small drifts in refrigerator temperature.

The phonon-pulser circuit was designed with greater goals in mind. Since the voltage of the pulse and the heater-pad resistance are nominally known, it is possible to calculate the energy deposited and thus use the phonon pulser as an absolute energy calibration at low energies rather than relying on 511-keV positron-annihilation photons (see Chapter 6). The phonon pulser fails in this respect — the phonon pulser energies do not give the expected event energies. The pulse applied to the heater is not square and is probably partially shunted by parasitic capacitances in the lines down to the detector, so the incorrect calibration may not be surprising. Another goal had been to use the phonon-pulser pulses to calibrate and calculate the efficiency of the phonon-pulse-fit-χ^2 cut. Unfortunately, it appears that, above 50-keV phonon pulse height, phonon-pulser events deviate from the shape of events due to particle interactions. It had also been expected that phonon-pulser data could be used to measure the phonon-trigger efficiency. This turns out to be unnecessary because it can be measured using low-energy events that are acquired when another detector triggers.

The ionization pulser also has problems. For the same reasons given above, it is not used to determine trigger thresholds. It is not useful for calibrating drifts in the ionization amplifiers because the pulser itself is as susceptible to drifts as the ionization amplifiers. The ionization pulser can be used to check the resolution in ionization as a function of energy.

It is likely that these problems with the IBAPACAP could have been corrected with more work, but these issues were a low priority at the time. While the phonon pulsing scheme is not expected to be applicable to ZIP detectors, ionization pulsing may prove useful if the pulse calibration can be stabilized.

5.4.6 Output Stage

The output stage of the boards contain assorted conveniences. The ionization-signal polarities may be switched, which is necessary if the ionization-bias polarity is switched. Two multiplexers make various housekeeping signals available to the DAQ system: the inputs to the lockins (*i.e.*, before demodulation), the phonon biases, the ionization biases, etc. Every signal is buffered before making it available to to the outside world. Front-panel test points are buffered with OPA627's with $50\,\Omega$ output resistors to ensure they can drive coaxial cables to oscilloscopes. Signals are output onto the 50-foot-long, 25-conductor cables that run to the Receiver/Trigger/Filter rack using LM6321 line drivers, which are capable of driving the highly capacitive loads presented by the cables while maintaining sufficient bandwidth.

Chapter 6

Data Acquisition, Reduction, and Calibration

6.1 Introduction

In this chapter, I describe the triggering and data-acquisition system, fitting of pulses to determine energies, and calibration of the energy over extended periods.

6.2 Trigger and Data Acquitision

The trigger and DAQ system were developed by Eric Aubourg, Dan Bauer, Steve Eichblatt, Rick Gaitskell, Josef Jochum, Sae Woo Nam, Francois Queinnec, Dennis Seitz, Tom Shutt, Andrew Sonnenschein, and myself. The Receiver/Trigger/Filter 9U-format boards were produced at Fermilab under the direction of Mike Crisler and Steve Eichblatt.

6.2.1 Overview

A diagram of the data-acquisition system is shown in Figure 6.1. The data-acquisition system

- configures the detectors and their front-end boards by setting assorted bias values, polarities, gains, etc.;
- receives and digitizes the analog detector pulses output by the front-end electronics;
- receives, amplifies, and digitizes the integral of the pulses from the veto-counter phototubes;
- discriminates each analog detector pulse to form detector-trigger signals;
- forms a global trigger based on the detector-trigger signals;
- digitizes a time history of detector and veto-counter triggers;
- and ships the digitized information to the analysis computers.

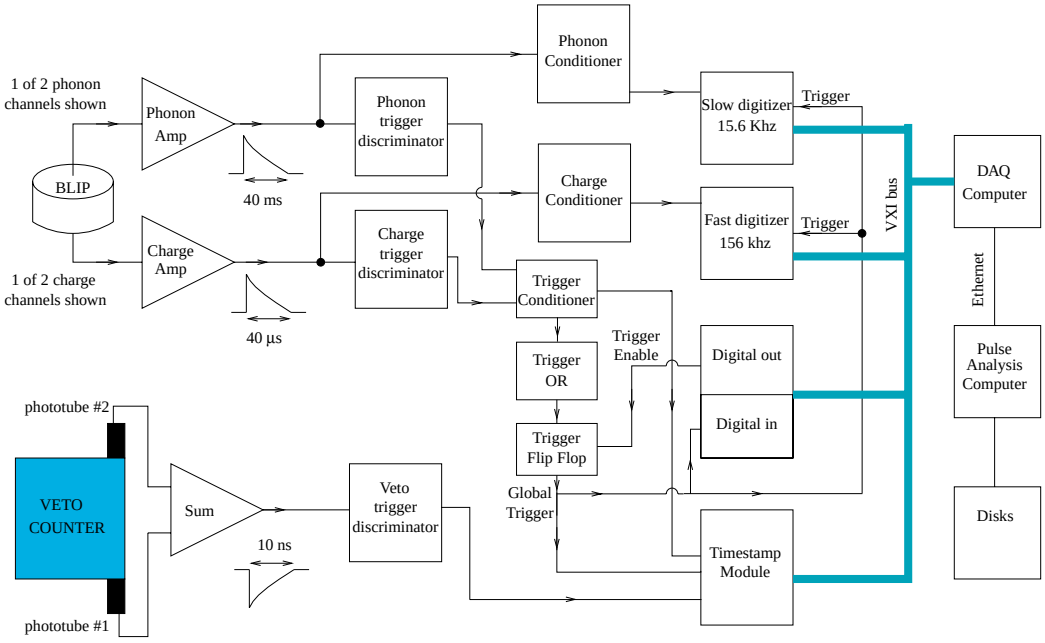

Figure 6.1: Schematic of data-acquisition system. Only one detector and one veto counter are shown. The modules that digitize the veto-counter pulse integrals are not shown. Figure taken from [8].

The system is controlled by Pilot, a Power Macintosh clone running MacOS, containing a PCI bus and a G3 upgrade card. Two auxiliary NuBus Power Macintoshes, Digitizer and Monitoring, perform slow-control and monitoring functions. The data-acquisition system is coded in LabView, in commercially-available LabView-callable drivers, and in custom-written C subroutines. The various programs running are:

- **Dark Control**: Runs on Pilot. This program starts, pauses, and stops all the other programs, including those on other computers.
- **Light Event**: Runs on Pilot. The main DAQ loop. When running and not paused, checks for triggers, performs readout of event information from assorted modules, and sends data to the analysis cluster.
- **Dark Script**: Runs on Digitizer. A scripted system for performing front-end board configuration and pausing/restarting the DAQ. Its primary role is to configure the front-end boards at the start of every data run, to pause the DAQ and ground the detectors briefly approximately once an hour, and to bias the detectors and restart the DAQ following grounding periods. Since IBAPACAP events are initiated by GPIB, Dark Script also performs phonon IBAPACAP runs every few hours. Dark Script also starts **Async IBAPACAP**, which sends GPIB triggers to fire the IBAPACAP asynchronously with Dark Script and the DAQ cycle. Dark Script ran on Digitizer during Run 19 because GPIB lockup occurred when both Async

IBAPACAP and Light Event tried to access GPIB simultaneously.

- **File Control**: Runs on Pilot. Checks disk space availability on the analysis computers and provides names and locations for data files being written.

- **Icebox Monitoring System**: Runs on Monitoring. Monitors temperatures, pressures, liquid levels, etc. in the Icebox and refrigerator. Interlocked with Dark Script so it only runs during grounding periods to avoid causing cross-talk to the detectors. Includes **Fridge Status**, which provides a resumé of refrigerator parameters.

- **Monitor HV**: Runs on Monitoring. Monitors veto-phototube high-voltage supply to provide warnings of tubes whose supply has failed.

- **Monitor DC**: Runs on Monitoring. Monitors DC-ref outputs of phonon lockins, which give sensitive histories of detector temperatures to allow for pulse-height drift corrections.

- **Monitor Scalers**: Runs on Monitoring. Measures veto-counter- and detector-trigger rates via CAMAC scalers. Also monitors hardware Live Time and Real Time clocks provided by Trigger Logic Board.

- **netServer**: Java-based file-transfer server running on an analysis computer in the Control Room, similar to ftp, optimized for transfer of large blocks of data over dedicated connections. netServer gives much larger transfer speeds ($\sim$ 1 MB/s between Pilot and Hadrian) than other programs in use (*e.g.*, Fetch, ftp, scp). Light Event deposits data on the analysis cluster via netServer. The Java platform allowed netServer to be moved from a Sun Ultra 1 to a PC running Linux without modification.

The monitoring programs send email if excursions in the refrigerator operation, veto high voltage, detector DC refs, or veto- or detector-trigger rates occur.

6.2.2 Slow Control and Monitoring

Dark Script configures the detectors and front-end boards at the beginning of each data run (once per day, beginning after cryogen transfers). It

- sets the phonon-bias value and type (1-kHz sine wave, 10-Hz square wave, or DC);

- sets the phonon-lockin mode: bypass or lockin, external or carrier reference, fast or slow AC coupling;

- puts the voltage amplifier in fast or slow autozero mode;

- sets the ionization-bias value and polarity and output-driver polarity;

- and sets pulse height, channel, and enables ionization IBAPACAP and initiates Async IBAPACAP routine.

During Run 19, the detectors were operated on a 50-minute-biased/5-minute-grounded cycle. Dark Script performs this cycling, also pausing the DAQ during grounding. Halfway through each 50-minute period, Dark Script momentarily pauses the DAQ to switch the ionization channel being pulsed via IBAPACAP. The ionization IBAPACAP pulse height is also modified at every grounding cycle, stepping through four different values in the 0-to-100-keV range. At the end of every four cycles, a phonon-IBAPACAP run is performed. This consists of 100 pulses at 1 Hz at each of four different pulse heights, again in the 0-to-100-keV range. As discussed in the previous chapter, the IBAPACAP is not completely successful in reproducing true pulses, especially in the case of attempts to perform ionization and phonon pulsing simultaneously. In the end, the phonon IBAPACAP pulses are useful for calibrating the DC-ref-based drift correction for phonons, while the ionization IBAPACAP pulses are useful for little besides checking pretrigger-cut efficiencies and energy resolution. Dark Script logs its operations to the analysis cluster.

The configurations of the Receiver/Trigger/Filter boards (discussed below) are not modified during normal running and so are set by the operator by hand at the start of each data run using a standalone LabView routine.

The refrigerator monitoring system is described in great detail in Peter Barnes's dissertation [92]. The other slow-control software is described adequately above.

6.2.3 Receiver/Trigger/Filter Electronics

Analog pulses are output by the front-end electronics described in the previous chapter. The gains have been chosen such that the pulse heights are approximately 5 mV/keV. The signals are driven onto twisted-pair cable. One cable, with 25 conductors, is allotted to each detector. The cable has foil shields surrounding all the conductors and surrounding each individual pair to ensure good isolation between channels and protection from electromagnetic and radiofrequency interference (EMI/RFI). The distance between the front-end and data-acquisition racks is approximately 30 feet; this large distance minimizes pickup by the front-end electronics of EMI/RFI emitted by the data-acquisition crates and computers.

The signals are recieved by Receiver/Trigger/Filter (RTF) boards in the RTF crate via 25-pin-D connectors at the rear of the boards. This is primarily an analog electronics crate, similar in architecture to the front-end crate, with one board per detector. In order to prevent ground loops, each RTF board differentially receives the analog signals from the corresponding front-end board. The signals are split into two paths. One path conditions the signals for shipment to the digitizers in a VXI crate. First, the ionization signals are notch-filtered to removed residual 1-kHz cross-talk from the thermistor bias, described in the previous chapter. The notch filter used is the Analog Devices 780RT-3, which provides a tunable notch center and width. The filter is set to have $Q = 10$. Because of component-value variations, the notch's center would be too far off from 1 kHz if Q were further increased. Furthermore, optimal filtering, which was not implemented until early 1999, renders the notch filter redundant. Second, both ionization and phonon signals are attenuated and offset to match the pulses to the $\pm$1-V range of the digitizers. The attenuation is chosen so the digitizers do not saturate until the pulse energy is $\gtrsim$ 1 MeV. Third, the pulses are anti-alias-filtered using 4-pole Butterworth filters constructed from discrete op-amps. This filter choice was made by comparing RC, Bessel, and Butterworth filters of 2 and 4 poles, constraining the filter parameters to yield the same 3 dB ($1/\sqrt{2}$) attenuation at the desired Nyquist frequency, f_{Ny}. The figure of merit used was the ratio of the integral of the power spectral density (*i.e.*, V^2/Hz)

of the noise aliased from the f_{Ny} to $2f_{Ny}$ range to the in-band noise. It was desired that this ratio be less than 0.1 to ensure negligible contribution from aliased noise. Of the filters considered, only the 4-pole Butterworth met this criterion. The details of the calculation are described in [140]. For $f_{Ny} = 78$ kHz, the filter poles are placed at 49.0 kHz with phases of ± 112 and ± 158 degrees. The phonon signals are filtered similarly, using $f_{Ny} = 7.8$ kHz. The filtered signals are output from the RTF boards to the digitizers by single-ended line-drivers via front-panel SMD connectors. There are $50\,\Omega$ resistors in series with these outputs to absorb any signal reflected from the high-impedance digitizer inputs.

The second signal path in the RTF crate discriminates the analog signals to provide detector-trigger signals. The inner- and outer-electrode ionization signals are summed to yield a single ionization signal, and the two phonon sensors' signals are similarly summed. It can be shown that, in the presence of different noise and signal heights in two channels, the optimal weighting (the one that yields the lowest threshold) is

$$w = \frac{\sigma_A^2}{\sigma_B^2}\frac{B}{A} \tag{6.1}$$

where A, B are the pulse heights in the two channels per keV, $\sigma_{A,B}^2$ are the rms noises in the two channels, and w is the factor by which channel B is multiplied when summing. In practice, the gains and noises of the two channels of the pairs are sufficiently close that the sums are done with equal weights for simplicity. The summed signals are amplified by factors of 100 (ionization) and 200 (phonon) and diode-clamped to prevent saturation of the op-amps in the chain (which can result in ringing). The signals are filtered using one high-pass and one low-pass single-pole RC filter. These simple filters are used instead of filters with more poles in order to prevent ringing that would result in multiple triggering of the following discriminator for large pulses. In addition, more complex causal filters do not appreciably improve the signal-to-noise. Appendix B provides a formula for the energy resolution of an arbitrary filter. For each detector channel, a template pulse and a real noise spectrum are obtained and a numerical search performed to find time constants for the filters that optimize the resulting energy resolution and thus minimize the trigger threshold. A LM311 discriminator provides a threshold trigger on the filtered pulse. The discriminator is operated with maximal positive hysteresis to prevent multiple leading-edge triggering. This hysteresis is especially important for the phonon pulses, which are much slower than the switching speed of the LM311. The trigger signal is thus active from the time of threshold crossing until the pulse crosses back through zero. Because of the high-pass filtering, the zero crossing occurs quickly compared to the decay of the pulse, and so the length of the trigger signal is not particularly useful. An additional high-threshold trigger is provided by a discriminator viewing the unamplified, unfiltered, summed pulses. This high trigger is used to veto high-energy events in calibration mode. The trigger signals are output via a 50-pin-D connector at the rear of the board.

The RTF boards are configurable via a backplane digital bus that is interfaced to GPIB, exactly as with the front-end electronics boards. Users can set thresholds, enable/disable trigger circuits, and choose whether to use the standard conditioning circuitry. The assorted filters, gains, attenuations, and offsets are configurable by modifying component values.

6.2.4 Veto Electronics

Because the phototube pulses provided by the veto counters are standard signals in high-energy physics, the veto electronics are built almost entirely from off-the-shelf NIM modules. Standard 50 Ω coaxial cables carry the negative-going pulses from the phototubes down the length of the tunnel to NIM racks adjacent to the RTF rack. The signals are amplified using LeCroy 612A Fixed-Gain Photomultiplier Amplifiers and the phototube signals from each veto counter summed using LeCroy 428F Linear Fan-In modules. A channel consisting of the sum of all the phototube signals is also constructed. LeCroy 623B Discriminator modules trigger on the summed signals and provide counter-hit signals to the trigger-history buffer, described below. The thresholds are set to ~ -30 mV to trigger on muons but not photons. The typical total veto-trigger rate is 6 kHz. To monitor veto performance, LeCroy 2249A Charge ADCs are gated by the -30 mV trigger signals to measure the summed pulse integral of each counter. The gate width is 75 ns. The spectrum consists of a tail of photon events near threshold and a wide muon bump at higher pulse integral. The pulse-integral spectrum thus provides a diagnostic of the veto's performance for muons, which is critical to the success of the experiment at the shallow site. To probe the possibility that sub-threshold events could be important, a second 2249A module is gated by a window trigger built from the logical OR of the -30 mV ("high") triggers and an additional -20 mV ("low") trigger on the *analog sum* of all the counter signals. The trigger rate of the low threshold is too large to allow use for vetoing.

Pulse-integral spectra for each veto side and the summed signal of all the counters are shown in Figure 6.2 for detector-triggered events. The veto-side signals are amplified to show the threshold region in detail, so the muon bump is not visible. The "High" spectra show events for which one or more counters has a high trigger: the dark line shows the spectrum of events for which the given veto side has a high trigger ("on-side" events), and the light line shows the spectrum of events for which a high trigger occurred, but the given veto side did not have a high trigger ("off-side" events). The two spectra thus give the above- and below-threshold spectra for the given veto side. The steep rise in the below-threshold spectra is due to photon interactions (noise is at the 1 mV level). The spike at zero pulse integral is due to events in which nothing, not even a photon, hit the counter during the ADC gate from the high trigger. The "Low" plots show the spectra for events for which the window trigger fires but none of the veto counters' high triggers fire. If the veto side has a window trigger, then a non-zero pulse integral appears; otherwise, the event appears in the noise distribution near zero pulse integral. The low threshold of the window trigger is seen in the rolloff of the "Low" spectra at low pulse integral, which differs from the steep rise seen in the off-side "High" spectra at low pulse integral. Finally, the "OR" plots give the sum of the pulse integrals of all the counters. The "OR" signal is not amplified as much as the individual counter signals, so the muon bump is visible in the spectrum. Because the veto-sum signal contains noise contributions from all the counters, its high and low thresholds are set much higher than the individual veto-side thresholds. The "ORHigh" plot shows the pulse-integral spectrum for events in which a veto-side high trigger occurred. The two spectra correspond to events in which the veto sum's high trigger did or did not fire. Because the "OR" high threshold is much higher than the individual counter thresholds, it is possible to have events in which one of the veto sides had a high trigger but the "OR" high threshold did not. Finally, the "ORLow" plot shows the corresponding window-trigger spectrum.

Grossly, the pulse-integral spectra are as expected. There is a muon bump in the "ORHigh"

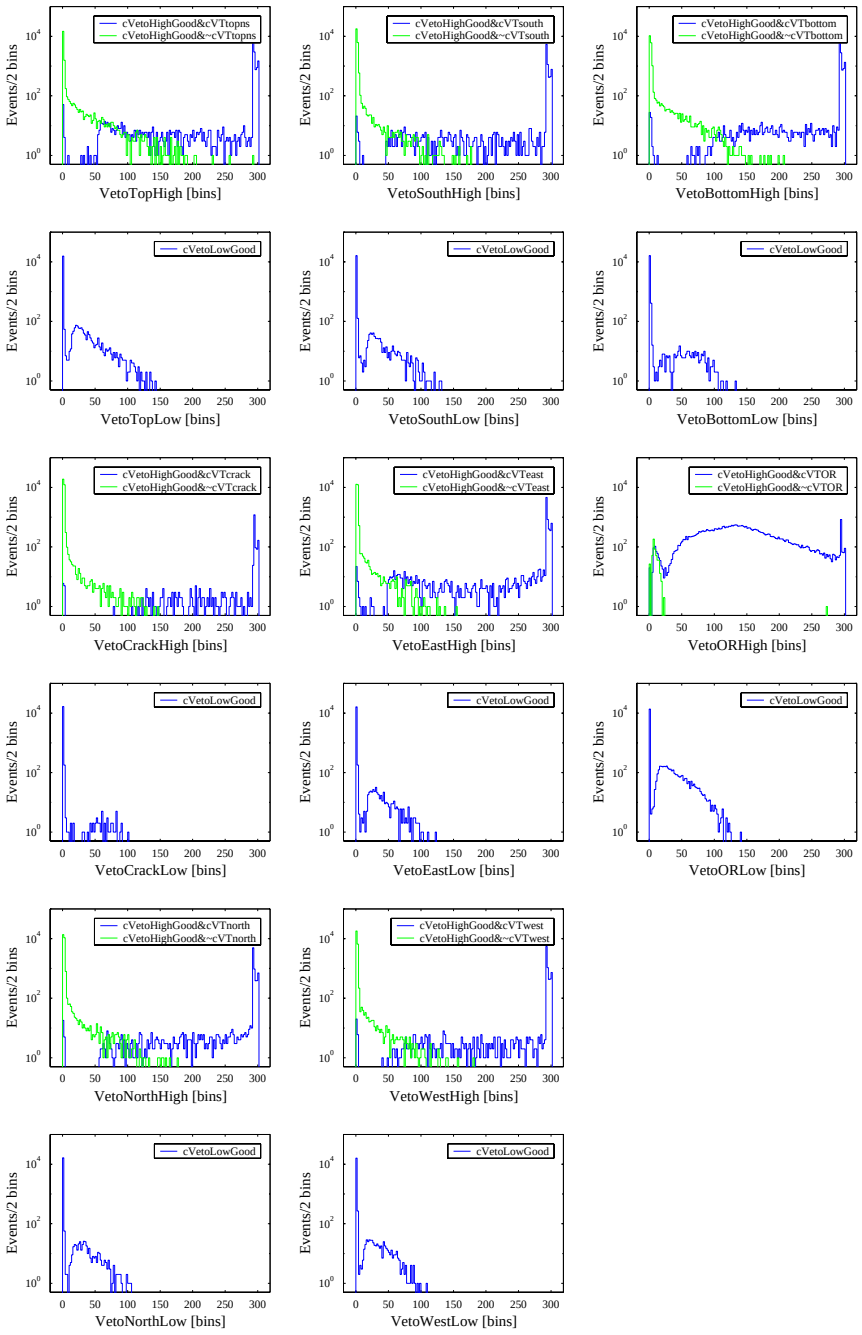

Figure 6.2: Veto pulse-integral spectra for detector triggers for 3 days of data. The spectra for all counters on a given veto side are coadded. The plots are described in the text. Note that the gain for the pulse integral of the summed signal (the "VetoOR" channels) is much lower than for the individual veto sides; the muon bump is thus visible in the VetoOR channel but occurs above saturation for the individual veto sides.

spectrum. The "High" spectra of the individual sides show that there is nothing obviously wrong when an off-side trigger occurs: there is indeed little or no pulse integral in the untriggered counter. Also, the "Low" spectra are as expected — when a window trigger occurs, the pulse integrals of all the counters are indeed low or vanishing. The high-pulse-integral portion of these spectra have the same shape as the off-side "High" spectra. Thus, as expected, the "Low" spectra are dominated by accidental coincidences with photon interactions. Nevertheless, the veto ADC spectra provide a good, continuous diagnostic of the veto's health by showing the pulse-integral spectra near threshold and the position of the muon bump.

6.2.5 Trigger-History Buffer

The trigger-history buffer is a VXI Technology VM1602 timestamp module. It accepts up to 96 inputs. Whenever a single input fires, the module records a timestamp and the status of all channels. The veto-counter, detector-ionization-, and detector-phonon-trigger signals are routed to this module. The buffer provides a history of all veto and detector activity around the time of the Global Trigger (defined below). All hits between 10 ms before the Global Trigger and 10 ms after are saved. The primary use of the history information is to determine whether a detector event is coincident with a veto event. Also, pre- and post-trigger pileup can be monitored by checking for additional ionization triggers near the trigger time. Finally, by recording post-trigger hits, the standard threshold trigger can be applied offline, allowing a uniform analysis threshold regardless of whether the detector in question is the source of the Global Trigger. By recording the veto history over a long period, a history of veto-counter-trigger rates can also be constructed offline.

6.2.6 Digitizers

Two types of digitizers were used during Run 19. OmniByte Comet digitizers were used to sample the phonon traces at 64 μs spacing for 2048 samples, or 131 ms. Joerger VTR1012 digitizers sampled the ionization traces at 6.4 μs for 2048 samples, or 13.1 ms. Both modules use 12-bit ADCs. The resulting event size is 16 kB/detector. The Comets are capable of digitizing at frequencies up to 2 MHz and the Joergers up to 10 MHz. This high-frequency capability was not necessary for Run 19.

It is useful to note here the positioning of the pulses in the digitized traces to clarify later discussion. For the case of ionization triggers, the situation is simple. The ionization pulse begins to rise slightly before the trigger time, crossing threshold at the trigger time. The phonon pulses are delayed by 1–2 ms because of the Butterworth filters used in the lockin demodulation circuit. The delay is very small compared to the 80 ms posttrigger period, is fixed, and is energy independent, so it introduces no difficulty. Since the pulses used for triggering are filtered heavily, the trigger time may appear later in the digitized pulse than expected; conversely, the digitized pulse appears slightly early, relative to the trigger time. In addition, there is a slewing error because a threshold trigger is used: the trigger time appears later in the pulse (relative to an energy-independent point, such as the pulse peak) for low-energy pulses than for high-energy pulses because low-energy pulses reach the trigger threshold later. This slewing error is important for two reasons. First, in fitting pulses, it is necessary to fit for a time offset since the template and data pulses may be of different energies (or may have been taken with different trigger thresholds). This is discussed in detail in Section 6.3. Second, the time of trigger may be delayed relative to the corresponding veto event

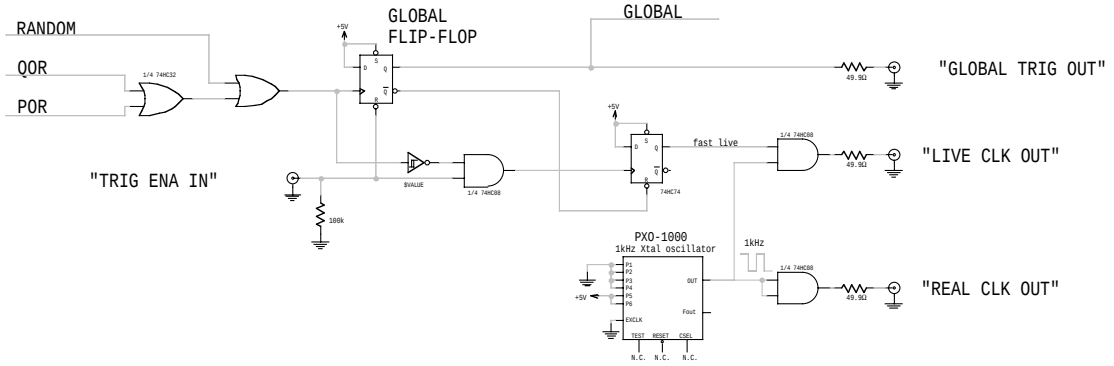

Figure 6.3: Global Trigger generation circuitry. Figure taken from [133].

that appears in the history buffer, which may affect the position of the window chosen to accept veto coincidences.

For phonon triggers without an ionization trigger, the situation is more complex. Since thermal-phonon pulses are so slow, the trigger time occurs up to many ms after the ionization pulse has occurred. Therefore, it is necessary to acquire ionization traces that are extremely long compared to the pulse length (12.8 ms compared to the 40 μs pulse decay time) so the ionization event appears in the large pretrigger region. Furthermore, because the trigger does not occur until the phonon pulse passes threshold, the trigger time is appreciably later (many ms) in both pulses for phonon triggers as compared to ionization triggers. Conversely, the phonon and ionization pulses appear earlier in the digitized traces, relative to the trigger time, as compared to ionization triggers. Therefore, when fitting, the phonon- and ionization-pulse templates must be shifted earlier in the trace by a significant amount to provide the correct time offset. The fitting procedure is discussed in Section 6.3.

6.2.7 Trigger Logic and DAQ Interlocking

It is necessary to provide a global trigger signal to inform the DAQ computer that an event has occurred. Each detector's RTF board outputs its high and low ionization- and phonon-trigger signals to the Trigger Distribution Board, a simple printed-circuit board attached to the output 50-pin-D connectors of the RTF boards. The Trigger Distribution Board routes the signals to the Trigger Logic Board, which resides in slot 10 of the RTF crate.

The primary function of the Trigger Logic Board is to take the OR of all the detector-trigger signals and to provide a global-trigger signal that can be sensed and enabled/disabled by the DAQ computer. The board is programmed to select one of four modes for each detector's ionization and phonon channels: disabled, low trigger, high trigger, and low trigger with high veto. The last mode delays the low trigger and blanks the low trigger using the high trigger, passing only triggers in a window between the low and high thresholds. This mode is used during photon-calibration runs, when a large fraction of the events are above 100 keV but are not useful.

The logical OR of the signals output by the above circuitry clocks a D flip-flop (the "Global

Flip-Flop" in Figure 6.3). The flip-flop outputs the "Global Trigger" signal to a front-panel BNC. This signal triggers the digitizers in the VXI crate and is also monitored by the DAQ computer. A flip-flop is used so that, once a Global Trigger occurs, the system cannot trigger again until the flip-flop is reset. The system is enabled/disabled via a "Trigger Enable" line from a front panel BNC that controls the flip-flop's RESET input. The DAQ computer uses this line to reset and reenable the flip-flop after a trigger has occurred and the digitizers are read out. The Global Trigger signal is also output to the history buffer to establish the time of trigger in the trigger history. Finally, internal to the Trigger Logic Board, the Global Trigger signal is used to establish a hardware live time for each event: Global Trigger gates the output of a 1-kHz oscillator; the oscillator is active when Global Trigger is inactive, and vice versa. Thus, by reading the gated 1-kHz signal with a scaler (via the "Live Clock" output), the time between the enabling and firing of Global Trigger is measured. The ungated clock signal is output from the board as "Real Clock" and is also monitored to measure the real time.

The DAQ computer communicates with the Trigger Logic Board via a HP 1330B VME digital bit-I/O module, which provides lines that can be configured as CMOS-level inputs or outputs. One line is configured as an input to accept the Global Trigger signal. The DAQ computer polls the bit-I/O module every 10 ms to check whether Global Trigger has gone active, indicating a detector trigger. Another line is configured as an output and is used to control the Trigger Logic Board's Trigger Enable line.

A secondary function of the Trigger Logic Board is to condition each of the detector-trigger signals for the trigger-history buffer. This consists of standardizing their lengths to 1 μs and differentially driving them onto twisted-pair cables to the history buffer.

The DAQ loop is as follows, beginning after the previous event has been fully downloaded:

1. The DAQ computer starts the digitizers running via the VXI interface to acquire pretrigger information. For the slow phonon traces, 48 ms of pretrigger information is recorded. The ionization digitizers require much less pretrigger information but are started simultaneously for simplicity. The DAQ computer also starts the history buffer.

2. In the intervening 48 ms, the previous event, which has been downloaded from the digitizers, is packaged and shipped via TCP/IP to the analysis cluster in the SUF Control Room. This pretrigger period may extend longer if computer or network slowdowns occur.

3. At the end of 48 ms, or when the DAQ computer is finished with shipping the previous event, the computer enables the Global Trigger by activating the Trigger Enable line via the bit-I/O module.

4. The DAQ computer sits in a loop, polling the bit-I/O module to check whether the Global Trigger line has gone active. The digitizers and history buffer acquire trace and trigger information continuously. The Live Clock output is running.

5. When an event occurs, Global Trigger goes active. This triggers the digitizers, instructing them to stop digitizing after a preconfigured posttrigger period. Global Trigger stops the Live Clock and is sent to the history buffer. Global Trigger is also shipped to the NIM crate to blank off clearing of the veto-counter ADCs. Since the total veto-trigger rate is 6 kHz, it is not possible to save ADC information for every veto trigger; rather, the ADCs are cleared

by a signal delayed from the original veto trigger. This signal is blanked off by the Global Trigger, thus saving veto information if there was a detector trigger.

6. 10 ms after the Global Trigger, the DAQ computer stops the history buffer and downloads its contents. The history is dominated by veto triggers. The last 200 hits are downloaded; the DAQ computer determines the time origin by finding the Global Trigger and retains 10 ms of pretrigger and 10 ms of posttrigger information. The CAMAC ADCs containing the veto pulse-integral information are also downloaded.

7. At 83 ms after the Global Trigger, the phonon digitizers stop recording posttrigger data. The DAQ computer downloads the digitized traces at this point. The ionization traces are not downloaded earlier because the high-frequency traffic on the VXI backplane can cross-talk to the phonon channels.

8. Once the traces are downloaded, the cycle begins again with the initiation of pretrigger sampling.

The Trigger Logic Board possesses a "Random" input that feeds directly into the large OR prior to the Global Flip-Flop. A line from the bit-I/O module is connected to this input and lets the DAQ computer force a trigger at any time. At the point in the DAQ loop where the computer is about to enable triggers, it checks to see whether it is time for a random trigger (based on a preset desired fraction of randoms). If so, the computer waits a random time between 500 ms and 1000 ms and then fires the random line. This type of random trigger is called a "hardware" random because it uses the triggering hardware in the standard way. "Software" randoms may also be generated by instructing the digitizers via VXI to begin posttrigger sampling, but this may result in extra noise due to VXI-backplane traffic during such a command. The use of random triggers in offline analysis is discussed in Section 6.3.

A better random-trigger scheme would be to decouple the generation of randoms from the DAQ loop by issuing the random trigger from a different computer, or at least from an independent process running on the DAQ computer. At present, their linkage to the DAQ cycle makes them more susceptible to pileup and thus of questionable use for establishing cut efficiencies. For this reason, ionization-pulser events, discussed earlier, are good substitutes.

6.2.8 Data Handling

Each detector possesses four, 2048-sample traces sampled at 12-bit precision. The resulting event size is 16 kB/detector. The typical VXI data-transfer speed during Run 19 was 2 MB/sec, yielding a download dead time of approximately 50 ms/event with 6 detectors operating. This dead time is nonnegligible but is smaller than the pretrigger/posttrigger sampling time and so does not dominate. The download of the history-buffer information is anomalously slow, 80 ms, due to the architecture of the module. This download time is hidden in the phonon posttrigger-sampling period and so is not important.

The trace data are not processed by the DAQ computer. The history-buffer data from the three submodules of the timestamp module are merged in chronological order and with times assigned relative to the Global Trigger time. The veto-ADC data are also not processed. The two latter data require negligible storage space. The traces, history-buffer, and veto pulse-integral

information are assembled into an elementary event structure and sent to a data-receiving computer over TCP/IP. The local area network consisting of the SUF Tunnel and Control Room computers is isolated from the rest of the world via a network switch, so high bandwidth is maintained on this connection. During the early part of the run, a Sun Ultra 1 (Tiberius) was used as the receiving computer, as it had been during Run 18. During the latter part of the run, a Pentium III PC running Linux (Hadrian) took over as receiving computer. The transfer speed is typically 1 MB/s, sufficiently fast that data transfer contributes no dead time during smooth running.

The slow-control and monitoring routines log their output to the analysis computer's disk via netServer also. At present, the only slow data that are merged into the data stream are the ionization-bias and phonon-bias settings (not the actual biases, which are not monitored) and the phonon-lockin DC refs. The first are necessary for calculation of recoil energy from the phonon and ionization energies; the latter are used for drift correcting the phonon pulse height. With the insensitivity of ZIP detectors to refrigerator conditions, it is unlikely that other monitoring information will be merged into the data stream in the near future.

Data are stored on large (20–30 GB) local disks on the analysis cluster. With six detectors running, the trigger rate was 0.6 Hz, resulting in a data rate of approximately 3 GB/day. Following first-pass processing (see next section), the data are gzipped, reducing the storage requirement by ~ 3. The raw data are archived to DLT once every couple of weeks and the raw data erased from disk. The entire Run 19 data set occupies about 2 TB uncompressed (that's right, terabytes!) and requires about 30 DLTs.

The file structure is in the form of "series" and "files." A series corresponds to a data run, the period between starting and stopping the DAQ. The DAQ is usually stopped once per day for cryogen transfers. A series can be shorter if work on the system is being done. Each series is assigned a 10-digit series number of the form YYMMDDHHMM corresponding to the start time of the series. Each series of data is written as separate files containing 1000 events each. The series is divided this way to give chunks of data that can be easily handled. This division also guards against data corruption — one file, or 1000 events, is the most that can be lost at one time. File division also permits parallel processing because single files can be processed on different CPUs.

6.3 Pulse-Energy-Fitting and Veto-Correlation Code: DarkPipe

The first-pass data-analysis code, DarkPipe, fits the acquired traces to templates to determine phonon and ionization energies and does some processing of the history-buffer information. In this section, I discuss the fitting algorithms used for both unsaturated and saturated events. The pulse-fitting code is derived from algorithms used by the Berkeley and Stanford groups for laboratory testing but has been rewritten in Matlab and largely automated. A number of people have contributed to DarkPipe, include Roland Clarke, Angela Da Silva, Rick Gaitskell, Sae Woo Nam, Richard Schnee, and myself. During Run 19, Richard Schnee took responsibility for keeping the data flowing smoothly through DarkPipe and also implemented the ionization optimal filter.

Logistically, DarkPipe processes individual 1000-event files independently and even non-sequentially — DarkPipe requires no "history" to operate. DarkPipe outputs an array of Matlab-format files, one for each detector ("DPBX.mat" files, X = 3–6), one containing general information, including detector-trigger information ("DPTG.mat" files) and one containing the veto-trigger history and pulse-integral information ("DPVT.mat" files). When a full series is finished, the Matlab

files for the 1000-event blocks are merged, producing a single file of each type for each series.

For BLIP detectors, a linear template fit is performed to the ionization and phonon pulses, minimizing the χ^2 defined by

$$\chi^2 = \sum_{i=1}^{N} \frac{|V_i - V_0\, s_i|^2}{\sigma_i^2} \tag{6.2}$$

where V_i are the data samples ($N = 2048$), s_i is the template (normalized in some reasonable way — *e.g.*, peak-normalized to 1), V_0 is the fitted pulse amplitude, and σ is the rms noise per sample. Minimization with respect to V_0 yields

$$V_0 = \frac{\sum_{i=1}^{N} \frac{V_i s_i}{\sigma^2}}{\sum_{i=1}^{N} \frac{s_i^2}{\sigma^2}} \tag{6.3}$$

In some cases, the fit is done in frequency space, but it remains linear. Many subtleties arise that I discuss in detail below.

For each channel, a template is built by averaging a number of ionization-triggered pulses. Pulses with energies of 100–200 keV are used — this ensures high signal-to-noise while being low enough in energy that the pulse shape is correct for lower-energy events. For the ionization template, slewing error is negated by using pulses in a small energy range. Slewing error is unimportant for phonon pulses since ionization-triggered pulses are used — the slewing error is of order μs. It is necessary to build different templates for each detector and channel because of pulse-shape variations. In the phonon channels, variations are caused by small differences in thermistor properties and detector heat sinking. Variations in the ionization pulse shape occur because of differences in feedback-component values and amplifier open-loop gains.

In addition to fitting energies, DarkPipe also merges the ionization-bias, phonon-bias, and phonon-lockin DC-ref information into the data stream. The quantities output by DarkPipe are known as "RQs," short for "Reduced Quantities."

6.3.1 Delay Estimation and Veto Correlation

Prior to performing the energy fits, it is necessary to determine the event "delay" — the position of the Global Trigger time relative to the particle interaction. In the vast majority of events, any multiple scattering occurs on ns timescales, much faster than the pulses, so it is reasonable to speak of a single particle-interaction time. (There are interesting cases where this does not hold). In practice, the particle-interaction time is not determined, but rather how much the templates must be shifted to match the observed pulse. The start of the template ionization pulse occurs within 1–2 μs of the particle interaction, which is sufficiently good precision to check for coincidence with the veto.

The delay is determined using the detector that gave rise to the Global Trigger, as determined by order in the history buffer. The validity of using this detector is easy to see. First, consider the case in which all detectors have the same pulse shapes and noise. The trigger thresholds are set commensurately in this case. Because a threshold trigger is used, the first detector to trigger is the one with the largest pulse height — it must rise most quickly to reach peak at the same time as the other pulses. Therefore, the first-triggering detector has the highest signal-to-noise and provides the best delay estimation. Variations in noise, thresholds, and small variations in pulse

shape complicate the situation, but the same result holds: the first-triggering detector has the best signal-to-noise. Consider the case of two detectors that have the same energy-to-voltage conversion but that have noises that differ by a factor of 2. The trigger threshold is set higher on the noisier detector. Therefore, for the same energy event, the noisier detector triggers later. One could search for a best estimator of the delay based on all the detectors' pulses, with weighting accounting for different pulse heights and noises, but this is not likely to help much since one detector usually has a much higher pulse height than the others.

Calculation of the delay is done using optimal filtering on the first-triggering detector. The numerical value of the delay is given by the time of the peak of the convolution of the "optimal filter" for the first-triggering detector channel with the digitized pulse. A detailed discussion of optimal filtering is given in Appendix B, and a short heuristic explanation is given below in the discussion of of pulse-energy fitting. If the event is an ionization trigger, the calculation is done on the ionization pulse summed over both electrodes. If the event is a phonon trigger, the summed (or average) phonon pulse is used first. Because the phonon pulse is sampled at 64 μs intervals, it does not provide a sufficiently precise time-offset estimate to allow correlation with the veto — the veto-trigger rate is 6 kHz, so there would be accidental coincidences 38% of the time, assuming the phonon-delay-estimation noise were only one time bin. The phonon delay is thus used as the center of a search in the ionization trace. The optimal-filter convolution is performed on the ionization pulse over the restricted search window. As is discussed below, it is advantageous to minimize the search-window size because the energy threshold for finding a pulse increases with the window size. If no above-threshold pulse exists, the search finds a noise excursion. In the case of a phonon trigger, the width of the search windows for the phonon and ionization signals are 14.4 ms and 1.6 ms, respectively.

The delay determined in the above way is used as the time offset in the fitting algorithm for the pulses in all the detectors. It is also used to determine the nearest veto trigger.

Practical considerations make the delay calculation more complicated. In the case that the phonon pulse is used to determine the delay, one such problem is that the phonon pulse is truncated (in time) by the sampling window. This is done intentionally because there is negligible signal-to-noise at low frequencies and because longer traces result in larger dead time (at 0.6 Hz trigger rate and 80 ms posttrigger sampling, the dead fraction due to posttrigger sampling alone is 5%). For technical reasons discussed in Appendix B, truncation makes it difficult to calculate the standard optimal filter and thus to perform the optimal delay estimation. Fortunately, since the noise in the phonon channel is almost white because of the use of the lockin, the optimal filter is well approximated by the time-reverse of the phonon template. Thus, a simple cross-correlation (not convolution) of the pulse with the template would be sufficient. However, the phonon template is the same length as the phonon traces. In calculating the cross-correlation, when the template is shifted by one sample early (left), there is no template sample against which to multiply the last trace sample, so this piece of the cross-correlation integral at this lag is lost, reducing the cross-correlation at this lag below its true value. This type of biasing can result in a cross-correlation that peaks at a different lag than it should, or in double peaks. To fix this, 200 samples at the end of the data trace are discarded, corresponding to the maximum lag allowed. The remaining 1848 samples are used in the cross-correlation. Thus, the number of cross-correlated samples remains fixed, and there are always enough template samples to calculate the cross-correlation in an unbiased way. This biasing does not occur for the ionization cross-correlation because the ionization pulses are

not truncated by the digitization window; the samples lost in the above way are noise and do not contribute to the cross-correlation integral.

A second issue is saturated traces. When a trace is saturated, the optimal-filter delay calculation fails because the digitized pulse does not fit the template. If a trace is saturated because the baseline is below the digitizer window, nothing can be done and the event is not fitted. For a trace with a good baseline that is saturated only because the pulse amplitude is large, a simplified delay algorithm is employed. The saturating pulse's time offset can be determined by comparing the time at which the pulse rises above noise and comparing to a similar point for the template. The point chosen is the last point before saturation that is within five standard deviations of the pulse baseline. Though it is not obvious that this method should work, it determines the delay sufficiently well that the calculated delay provides the correct time offset for other unsaturated traces.

Pileup in the phonon trace also causes the phonon-delay calculation to fail. Pileup affects less than 10% of phonon pulses, and the phonon-delay calculation is performed for a similarly small fraction of events, so this failure is unimportant. Ionization triggers with phonon pileup can be analyzed, as is discussed in the next section. Pileup in the ionization trace occurs very rarely (the single-detector trigger rate is 0.3 Hz and the ionization trace length is 13.1 ms, yielding a pileup fraction of 0.4%) and so such events are simply rejected.

6.3.2 Pulse-Energy Fitting

In principle, once the delay is determined, it is straightforward to use Equation 6.3 to fit for the pulse amplitude. This is done for phonon pulses, with the addition of a floating baseline offset and an arbitarily normalized exponential with a fixed time constant to allow for pileup — the tail of the previous pulse is well approximated by an exponential of known time constant. The fit remains linear. One minor problem is that the χ^2 of the fit is incorrectly normalized. When calculating the χ^2 in time, correlations in the noise between samples are not taken into account. Such correlations are discussed in detail in Appendix B. They can be understood heuristically as resulting from the finite bandwidth of the phonon-channel electronics — a noise impulse requires approximately 350 μs $\sim (2\pi \times 450\text{ Hz})^{-1}$ to decay, so noise "hangs around" for many samples. The thermal time constants do not result in significant noise correlation because the thermistor noise is dominated by Johnson noise, not thermodynamic fluctuations. The correct χ^2 could be calculated using the autocorrelation matrix of the noise. Since a hypothesis test is not performed based on phonon-pulse-fit χ^2, this normalization issue is unimportant and is neglected.

For the ionization traces, it is advantageous to use optimal filtering to calculate the fit energy because the noise of the ionization channels has significant frequency structure. Pickup of 1 kHz and harmonics from the thermistor bias occurs. The frequency region below 500 Hz is dominated by FET $1/f$ noise and 60 Hz. Optimal filtering calculates the pulse fit in frequency space, where frequency components with low signal-to-noise can be deweighted to minimize their effect on the fit. The technique is discussed in detail in Appendix B. In the DarkPipe implementation, the pulse and template are Fourier transformed, the optimal filter ($= s^*(f)/J(f)$, where $J(f)$ is the noise power spectral density in V^2/Hz) calculated, the optimal filter multiplied against the pulse transform, and the result transformed back to time. This is equivalent to calculating the optimal filter in Fourier space, transforming to time, and convolving it with the pulse; *i.e.*, filtering the pulse with the optimal filter. In either case, the optimal time-offset and energy estimators are given by

the time and the value of the peak of the filtered trace. The time offset provides the phase factor to apply to the template in frequency space to allow calculation of the χ^2 in frequency space, where it can be correctly normalized because noise components at different frequencies are uncorrelated.

This algorithm is complicated by the sparse ionization-trace sampling used in Run 19. Based on a simple two-pole model of the ionization pulse, I determined that there is negligible signal-to-noise for the delay measurement above the frequency corresponding to the pulse rise time, $f_r = (2\pi\tau_r)^{-1} \sim 50$ kHz. This argument is made in detail in Appendix B. Therefore, the traces were sampled at 156.25 kHz, corresponding to a Nyquist frequency of 78.125 kHz (100 kHz was not an available digitization frequency). However, in making the calculation of this effect, I neglected the fact that a template fit is done rather than a fit to an analytic form. For a template fit, it is not possible to calculate the convolution of the trace and the optimal filter for delays that are not integral multiples of the sampling interval Δt. Since the peak value of this convolution is the energy estimator, the energy is therefore underestimated. Essentially, the time jitter of the pulse relative to the digitizer sampling results in an energy jitter. The expected underestimation of the energy is calculated in Appendix B and is significant.

A method to correct for this problem was developed by Richard Schnee. DarkPipe uses five ionization templates constructed by interpolation of the original template on spacings of 0.2 bins. Since the actual behavior of the template in between the samples is not known, the resulting filters are not quite optimal, but are close. Because the time offset is already known to within one bin from the calculation described earlier, it is only necessary to calculate the filtered traces at that time bin for each template. In frequency space, this corresponds to multiplying the Fourier-transformed data trace by the appropriate phase factor corresponding to the gross time offset, performing a linear, no-time-shift fit to each of the five frequency-space templates using a formula similar to Equation 6.3, and choosing the best fit. Because of the nature of the time discretization, the best fit template is the one that yields the maximum value for the energy. This calculation is performed for each trace separately; it has been observed empirically that sub-integral time shift calculated for one detector is not necessarily the best value for other detectors. This is surprising, given that all the digitizers are slaved to the same master clock.

A second complication arises because of cross-talk between the inner and outer ionization channels of a single detector. Each trace is the sum of its own pulse and a cross-talk pulse whose amplitude is proportional to that of the pulse in the other channel. There is a simple matrix equation for the χ^2 in this case:

$$\chi^2 = \sum_{k=1}^{N} \left(\vec{\tilde{V}}_k - \underline{\tilde{s}}_k \vec{A}\right)^\dagger \underline{J}_k^{-1} \left(\vec{\tilde{V}}_k - \underline{\tilde{s}}_k \vec{A}\right) \tag{6.4}$$

with

$$\vec{\tilde{V}}_k = \begin{pmatrix} \tilde{V}_{i,k} \\ \tilde{V}_{o,k} \end{pmatrix} \tag{6.5}$$

$$\vec{A} = \begin{pmatrix} A_i \\ A_o \end{pmatrix} \tag{6.6}$$

$$\underline{\tilde{s}}_k = \begin{pmatrix} \tilde{s}_{i,k} & \tilde{x}_{o\to i,k} \\ \tilde{x}_{i\to o,k} & \tilde{s}_{o,k} \end{pmatrix} \tag{6.7}$$

$$\underline{J}_k^{-1} = \begin{pmatrix} \frac{1}{J_{i,k}} & 0 \\ 0 & \frac{1}{J_{o,k}} \end{pmatrix} \tag{6.8}$$

where $\widetilde{V}_{i,k}$ and $\widetilde{V}_{o,k}$ are the kth frequency components of the Fourier transforms of the inner- and outer-electrode traces, A_i and A_o are the fit amplitudes to be determined, $\tilde{s}_{i,k}$ and $\tilde{s}_{o,k}$ are the kth frequency components of the Fourier transforms of the peak-normalized templates, and $\tilde{x}_{o\to i,k}$ and $\tilde{x}_{i\to o,k}$ are the kth frequency components of the Fourier transforms of the outer-inner and inner-outer cross-talk templates. The x's are normalized such that that they give the correct cross-talk amplitude for a unit-amplitude pulse in the other electrode. $J_{i,k}$ and $J_{o,k}$ are the kth frequency components of the inner- and outer-electrode noise power spectral densities. Minimization of χ^2 results in a matrix equation for $\vec{A}$ with solution

$$\vec{A} = \left(\sum_{k=1}^{N} \Re \left(\underline{s}_k^\dagger \, \underline{J}_k^{-1} \, \underline{s}_k \right) \right)^{-1} \left(\sum_{k=1}^{N} \Re \left(\underline{s}_k^\dagger \, \underline{J}_k^{-1} \, \vec{\widetilde{V}}_k \right) \right) \tag{6.9}$$

This formula is analogous to the usual expression without cross-talk, Equation 6.3. In practice, the unshifted template is used to calculate the delay to within one bin without any consideration of cross-talk. Then, the above calculation for $\vec{A}$ is made for each shifted template. The time offset that maximizes the sum of the two components of $\vec{A}$ is the one chosen.

As noted earlier, optimal filtering cannot be applied to saturated pulses. However, once the delay is determined, a template fit in time can be done by using only the unsaturated samples. If the pulse is not much larger than the saturation threshold, this works well — the rising and falling edges of the pulse provide enough information. However, pulse-shape variations cause this to fail for very high-energy pulses. Even at only moderately high energy, $\sim$ 1 MeV, the phonon-pulse shape changes because of the intrinsic nonlinearity of the system — the standard Taylor-expansion solution becomes invalid. At the energies at which the amplifiers themselves saturate, the amplifiers misbehave as they recover from saturation, resulting in varying decay times. For these reasons, the energy fitting becomes both noisier and nonlinear at high energies.

Optimal filtering also fails for pulses with pileup because the combination of two pulses does not match the pulse template that was used to calculate the optimal filter. As noted earlier, pileup occurs rarely for ionization pulses and such pulses are simply rejected.

6.3.3 Additional Trace Calculations

DarkPipe calculates a number of other straightforward trace parameters, many of which are useful for later analysis. These are (where "XX" is a detector channel: QI, QO, P1, or P2):

- XXbs: mean of pretrigger region, in digitizer ADC bins;
- XXstd: standard deviation of pretrigger region, in digitizer ADC bins;
- XXrt: 10% to 90% rise time of pulse, no fitting or smoothing done, in digitizer time bins;
- XXst: a. k. a. "start time," the last point before the maximum of the trace that is within five standard deviations (as defined by XXstd) of the trace baseline. Essentially, gives the position of the rising edge of the largest pulse in the trace in digitizer time bins;

- XXsat: boolean indicating whether the trace contains saturated samples. The front-end ionization amplifiers saturate before the signals reach the top of the digitizer windows, so the saturation limit is hardwired.

6.3.4 Random Triggers

To establish the noise performance of the delay and fitting algorithms, random triggers are taken by the data-acquisition system. These are traces taken at pseudo-random times that ostensibly have no events in them. The detector-trigger rate is low enough that it is easy to tag and discard random triggers that contain accidental events.

One use of random triggers is to provide empty traces containing only noise for calculation of the noise power spectral density (PSD). This quantity is used to calculate the weighting factor for the optimal filter. The noise environment during Run 19 was stable enough that there was no need to calculate the noise PSD continually. Rather, the noise PSD is calculated and saved with the fitting templates, which are constructed only when necessitated by changes in operating conditions. Typical noise spectra are shown in Figures 4.7 and 4.12.

More important for later analysis are the energy fits done for random triggers. For both phonon and ionization channels, a fit is performed with no time shift applied to the template. The resulting distribution of fitted energies should be Gaussian. For the ionization channels, where the noise and baseline are extremely stable, this holds. For the phonon channels, thermal or "thermophonic" (microphonics depositing mechanical energy in the detectors, yielding a thermal signal) instabilities can result in non-Gaussian baseline noise. Typically, the phonon-channel noise histograms are almost Gaussian, but have tails to high energy. Figures 4.21, 4.22, 4.23 show noise histograms that illustrate these points.

For ionization, a "sliding" noise fit is also performed. A search window, equal in width to that used for the ionization search for phonon triggers, is used, centered near the nominal trigger time of the event. The standard optimal-filter delay search and energy fit are performed. The resulting energy distribution is approximately Gaussian but is offset positively from zero, is narrower than the zero-delay noise distribution, and has a non-Gaussian tail to high energy. The asymmetrically positive ionization optimal filter picks out the most positive noise excursion in the search window because the optimal-filter algorithm picks out the peak of the filtered trace. The expected distribution is related to the zero-delay noise distribution:

$$P(E_{max}) = M\left[\mathrm{erf}(E_{max}, \sigma_E)\right]^{M-1} \frac{1}{\sigma_E\sqrt{2\pi}} \exp\left(-\frac{E_{max}^2}{2\sigma_E^2}\right) \tag{6.10}$$

where M is the width of the search window (number of samples) and σ_E is the width of the zero-delay noise distribution. This result is derived in Appendix B. It is nontrivial to calculate the expected mean and standard deviation for the sliding noise fit from the above distribution, but it is clear that the mean is displaced to more positive E_{max} as M is increased. Also, the error function yields a tail to high E_{max}. One can show numerically that the width decreases as M is increased. Histograms of the sliding noise fits done for random triggers are shown in Figures 6.4, 6.5, and 6.6. Note that the inner- and outer-electrode energies are calculated in the same combined fit, as would be done for a standard phonon-triggered event.

The usefulness of the observed distribution is that it indicates at what ionization energy, for phonon triggers, events are indistinguishable from noise. Phonon triggers with "noise charge" must

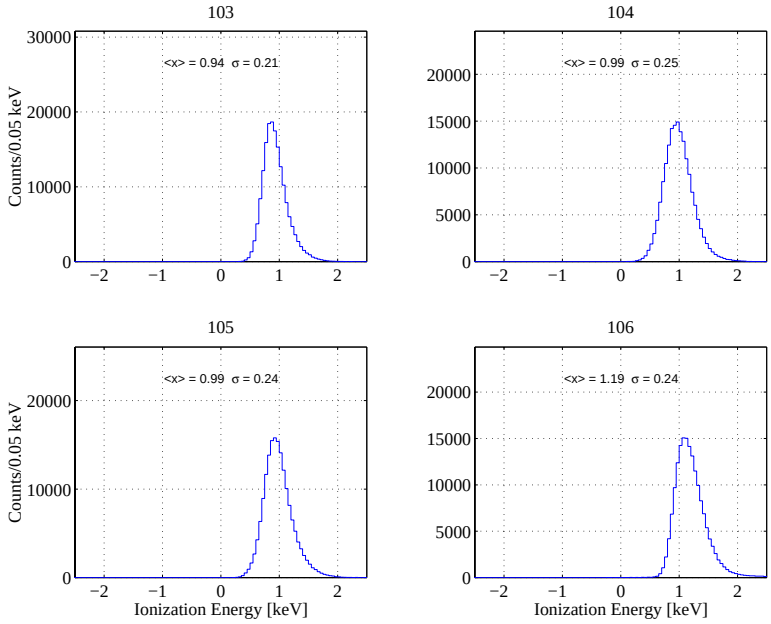

Figure 6.4: Distribution of summed ionization energy for random triggers as determined by "sliding" noise fit. The distributions are not expected to be Gaussian, so a histogram fit is not done.

be treated differently from events possessing ionization since only an upper limit on the ionization energy is established. Since no ionization is found in noise-charge events, no veto-coincidence determination can be made either. A sliding noise fit is not necessary for phonon triggers because: if the phonon trigger fires, there is something above noise in the trace to be found; analysis is not performed near the phonon-trigger threshold because there typically is no ionization information; and, the phonon-search window is small (225 samples, or 14.4 ms) relative to the phonon pulse length, so the noise distribution calculated from Equation 6.10 does not differ significantly from the usual noise distribution.

Table 4.3 gives baseline noises calculated using these random triggers and the zero-delay fits.

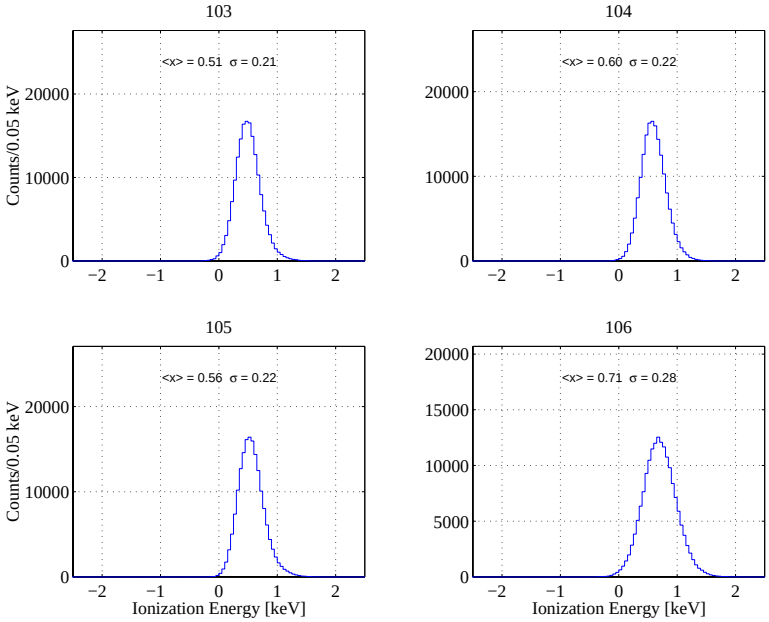

Figure 6.5: Distribution of inner-electrode ionization energy for random triggers as determined by "sliding" noise fit.

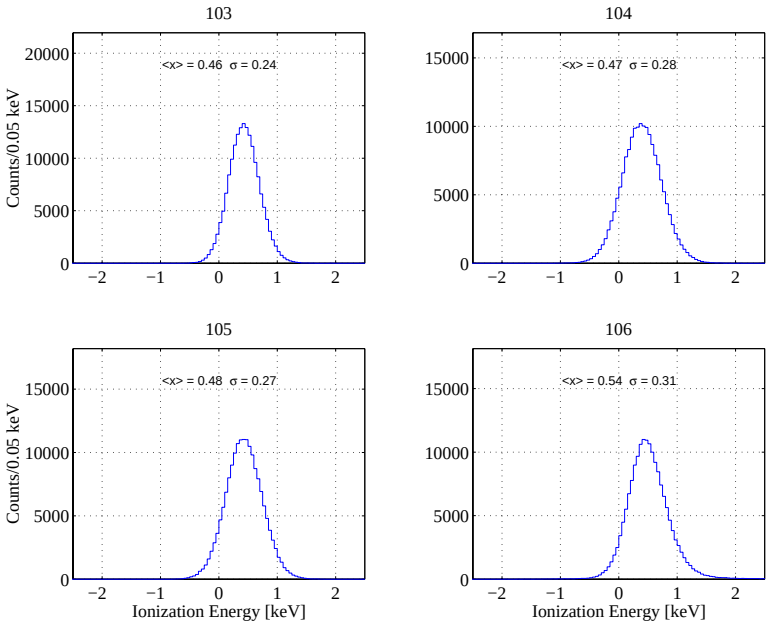

Figure 6.6: Distribution of outer-electrode ionization energy for random triggers as determined by "sliding" noise fit.

6.4 Second Pass Data Analysis: PipeCleaner

Due to drifts in both refrigerator base temperature and the electronics, the energies fitted by DarkPipe exhibit drifts. It is necessary to perform an absolute, time-dependent calibration to correct these drifts. PipeCleaner does this. PipeCleaner assumes that the ionization calibration is time-independent and calibrates the phonon energies against ionization using the prominent bulk electron-recoil band. PipeCleaner writes "datum" files that store this information for later use or modification. In addition, PipeCleaner calculates distributions of assorted event quantities that only make sense after the energies are correctly calibrated, such as recoil energy (by Neganov-Luke correction) and ionization partition, $(\mathrm{QI}-\mathrm{QO})/(\mathrm{QI}+\mathrm{QO})$, termed "Relational Reduced Quantities," or "RRQs." PipeCleaner also sets cut values to be used in later analysis. In order to accumulate enough data to perform these calculations, PipeCleaner operates on full data series ($\sim$ 1 day) rather than 1000-event files.

A number of people have contributed to PipeCleaner. Tom Shutt wrote the shell and developed the original calibration algorithms. Roland Clarke and Josef Jochum also contributed. More recently, PipeCleaner has been taken over by Thushara Perera. Richard Schnee and I contributed new code during Run 19.

The reader will no doubt conclude that PipeCleaner has the potential to perform many more analysis tasks automatically than it does now. I am certain this will happen during CDMS II as more thought and manpower are devoted to this stage of the analysis.

6.4.1 Phonon Calibration

PipeCleaner uses "prgamma" units, defined as follows:

$$\mathrm{prgamma} = \frac{E_P}{1 + (eV_b/3.0\ \mathrm{eV})} \tag{6.11}$$

where V_b is the ionization-bias voltage, 3.0 eV is the usual energy-per-charge-pair in germanium, and E_P is the phonon energy. For bulk electron-recoil events (γ's), $E_Q = E_R$ and this formula effectively performs the Neganov-Luke correction; hence the name prgamma. It is not correct for events with $E_Q \neq E_R$. The choice of whether to use phonon energy or prgamma in calculating calibrations is arbitrary; the correct recoil energy is always calculated during physics analysis. This choice is primarily historical, referring back to times when the ionization bias used was < 0.5 V and the difference between phonon energy and "prgamma" was small. In general, "phonon energy" denotes the true, physical phonon energy (recoil + drift heat), "prgamma" is the phonon energy scaled in the above way, and recoil energy refers to the correctly calculated physical recoil energy.

The first step in the phonon calibration is DC-ref correction. As described in the previous chapter, the phonon-lockin DC ref gives the thermistor's average resistance. To first order, the pulse height is linear in deviations of the thermistor resistance due to thermal drifts. A simple linear correction is made to the phonon pulse height based on the DC-ref measurement. The correction is calibrated using phonon IBAPACAP events. This procedure is especially useful for correcting the "bump in the night" thermal excursions observed during Run 19B, which were presumably due to a mixture-circulation instability. The correction slopes are fairly stable; only three different sets of corrections are needed for the entire run.

The second step is to calibrate the two phonon channels against each other — they should give the same energy for every event (except thermistor interactions — see Chapter 4). The "phonon partition" quantity, $(\mathrm{P1}-\mathrm{P2})/(\mathrm{P1}+\mathrm{P2})$, is calculated for the events in the range 10-to-100-keV prgamma. This region is chosen because the energy calibrations are most critical at low energy. The ratio $r = \mathrm{P1}/\mathrm{P2}$ is fitted and each channel is corrected by $\sqrt{r}$. The mean and standard deviation of phonon partition are calculated and events greater than three standard deviations from the mean are discarded. A new scaling factor is calculated and the process iterated until the standard deviation stabilizes.

The final step is calibration of phonons against ionization. prgamma is calculated from $\langle \mathrm{P1}, \mathrm{P2} \rangle$. A histogram of $R = E_Q/\mathrm{prgamma}$ is calculated and fitted. The phonon data are corrected to force $R = 1$. The fit and correction are iterated until the correction factor stabilizes. Code also exists to calibrate the phonon channels against the 511-keV positron-annihilation line, but this calibration is troublesome: there are usually very few events in this peak in one day's data; and, because of the 6-V ionization bias used for most of Run 19, the 511-keV peak lies at 1533-keV phonon energy, large enough that the phonon pulse height is no longer linear.

The above correction factors are saved to the datum file for the series.

In practice, the assumption of ionization stability is not valid. The ionization channels are still linear at 511 keV, so they are calibrated against this line. This calibration is usually performed for a relatively large block of time (days to weeks) because the energy resolution of the ionization signal is poor enough at 511 keV that a photon peak cannot be found in a single day's data. Moreover, changes in the ionization calibration can usually be traced to discrete events — cycling of power on the front-end electronics crate, for example. The calibration constants are calculated manually and the necessary correction factor inserted in the datum files. The phonon calibrator is then rerun to correct the phonon channels aginst the calibrated ionization energy. The quality of the calibrations were discussed in Chapter 4.

6.4.2 The April 3 Event: Ionization Linearization

On April 3, 1999, a Stanford-wide power outage occurred. The sudden loss of power damaged the electronics chain in a subtle way, introducing a nonlinearity in the ionization-energy calibration. The problem was, unfortunately, not noticed until much later. The nonlinearity occurs at approximately 10 keV in ionization. Because this is the region of particular interest for WIMP searches, it is necessary to correct the nonlinearity so the nuclear-recoil band behaves in a reasonable way. An empirical linearization was performed.

There is good evidence that nonlinearity of the electronics chain is the only reasonable explanation for the effect. First, the inner and outer ionization electrodes of a given detector exhibit different nonlinearities. It is hard to imagine detector physics that could do this. Second, photon-calibration data taken at two different ionization biases indicate that the nonlinearity appears at the same ionization energy but at different phonon energy for the two biases. As described in Chapter 4, the Neganov-Luke effect increases the observed phonon energy by an amount proportional to ionization bias and the ionization produced. A nonlinearity in the phonon-channel electronics should thus remain at the same phonon energy when the bias is changed; the observed nonlinearity moves in the way expected for an ionization nonlinearity. The suddenness of the change and the existence of an obvious cause rule out software problems.

The cause of the nonlinearity has not been investigated in detail — all elements of the

electronics chain are suspect. One possibility is that the nonlinearity occurs in the digitizers. My personal opinion is that, because of their demonstrated sensitivity to power cycling, the front-end boards are the more likely culprit.

The single-scatter veto-coincident photon data provide a well-defined band of bulk electron recoils that is used to parameterize the nonlinearity. For each detector, the events in this data set are binned in phonon energy (not prgamma, though this could be done instead) and the mean ionization energy in each bin calculated by a Gaussian fit to the distribution of the ratio of ionization energy to phonon energy. The bins are made in phonon energy rather than recoil energy because the phonon energy is the independent parameter (under the assumption that the nonlinearity is electronic, not physical) — the recoil energy is calculated from the phonon and ionization energies and thus would also contain a nonlinearity. The mean E_Q/E_P values thus derived are Savitzky-Golay filtered to yield a smooth parameterization. This parameterization is inverted and used to correct the observed ionization energies independently of the value of the phonon energy. I show a sample of pre- and post-linearization data in Figures 6.7, 6.8, 6.9, and 6.10 to indicate the nature of the problem and the quality of the fix.

In spite of performing this linearization, the nuclear-recoil band appears in a different position in the pre- and post-April 3 data sets. This is apparent in both the veto-coincident-neutron data and the second neutron calibration. Rather than attempt a correction that depends on both ionization and phonon energies, the nuclear-recoil band is redefined based on the second neutron calibration, as is discussed in Chapter 4. In Chapter 7, I discuss in detail the checks done to ensure the nuclear-recoil efficiency is known accurately over the entire data set.

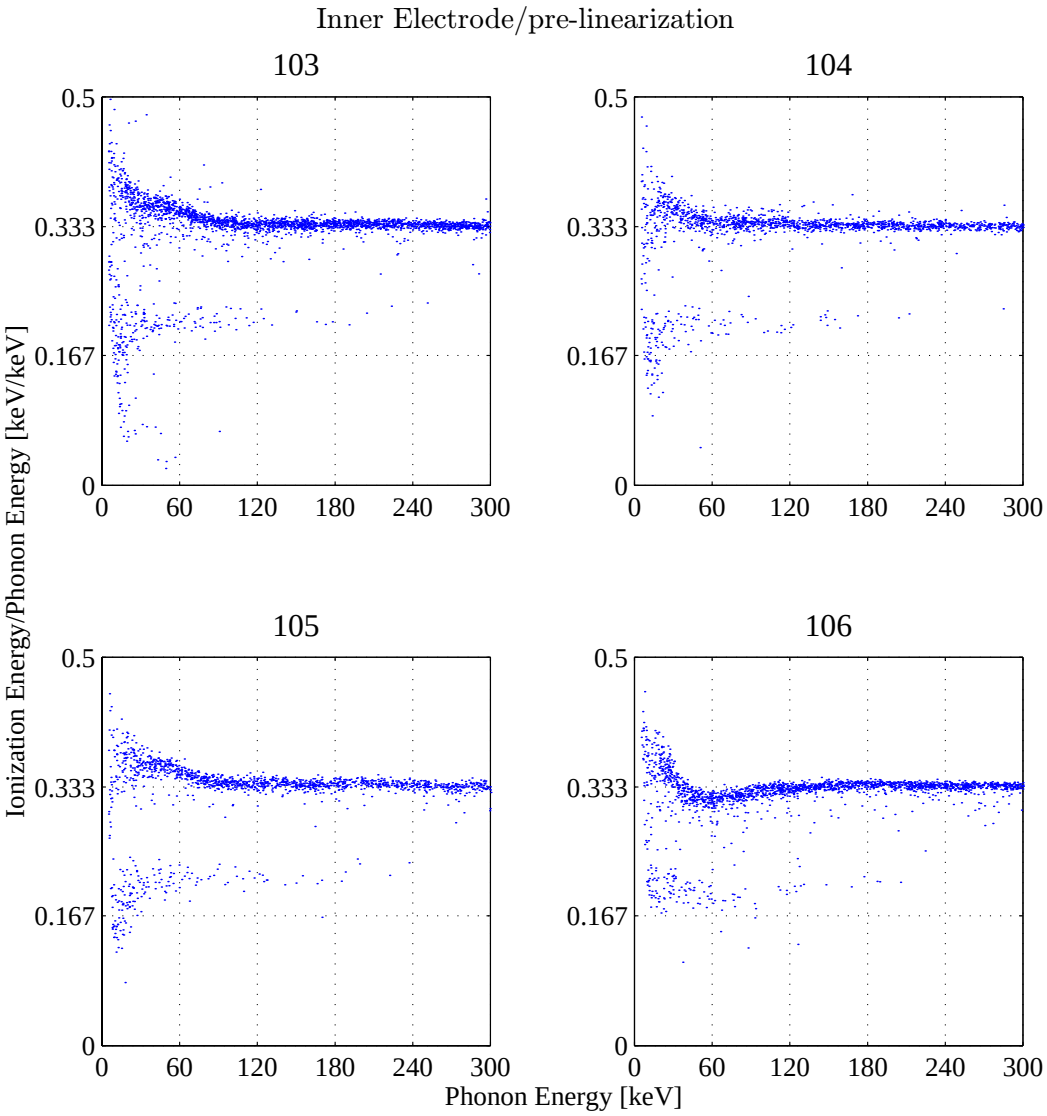

Figure 6.7: Inner-electrode ionization energy/phonon energy vs. phonon energy for veto-coincident cQinOnly single scatters for a subset of post-April 3 data before linearization. The events in the lower band on each plot are veto-coincident neutrons.

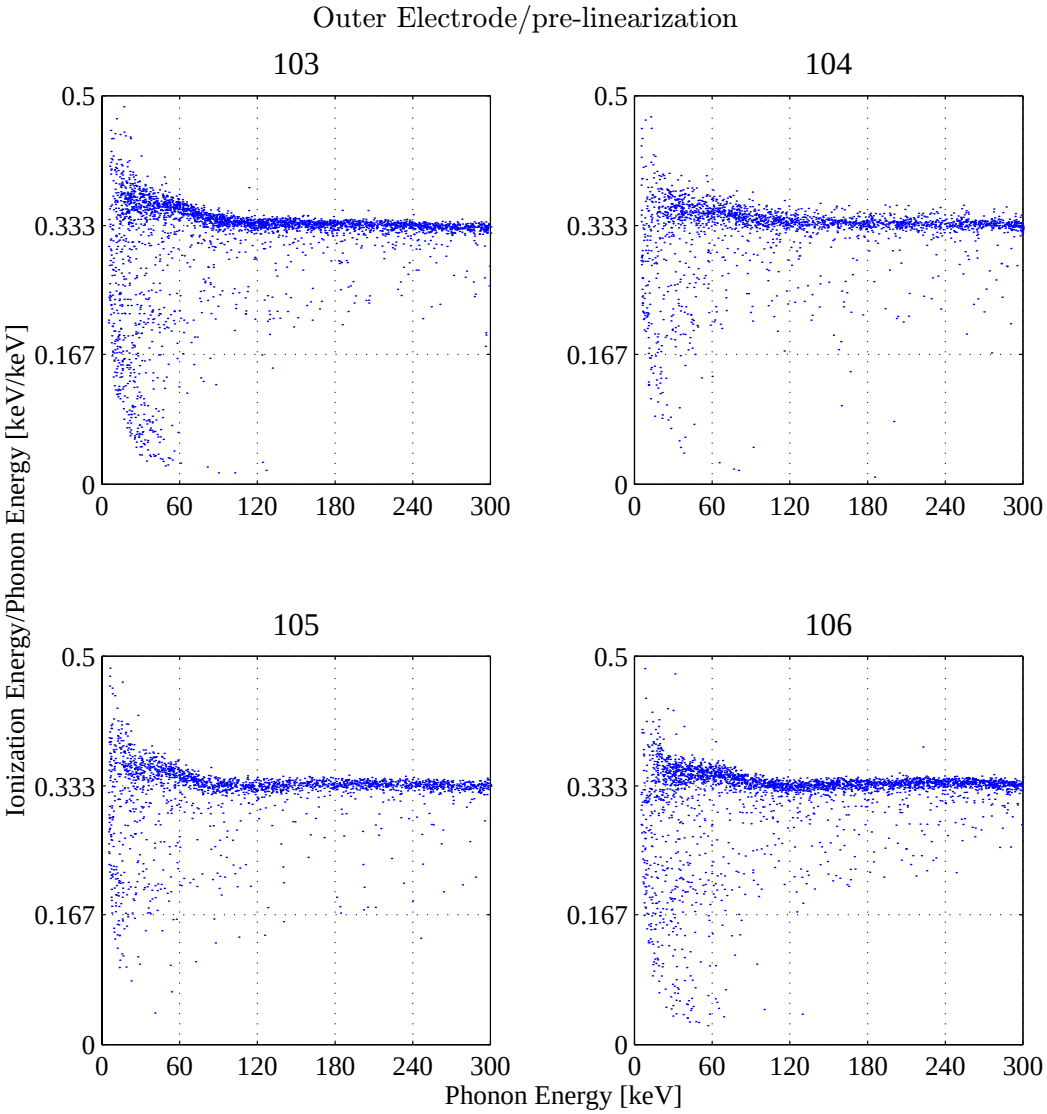

Figure 6.8: Outer-electrode ionization energy/phonon energy vs. phonon energy for veto-coincident cQoutOnly single scatters for a subset of post-April 3 data before linearization.

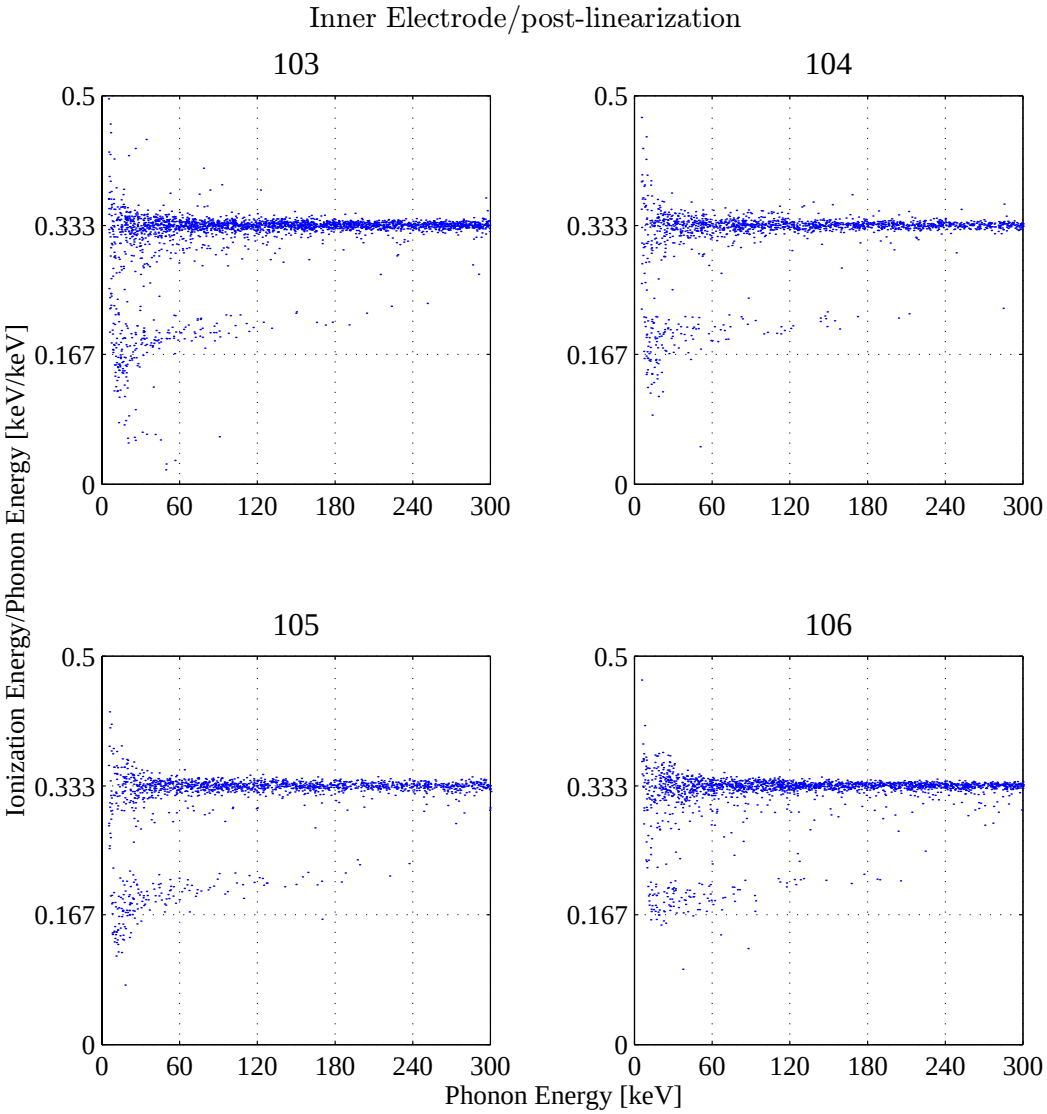

Figure 6.9: Inner-electrode ionization energy/phonon energy vs. phonon energy for veto-coincident cQinOnly single scatters for a subset of post-April 3 data after linearization.

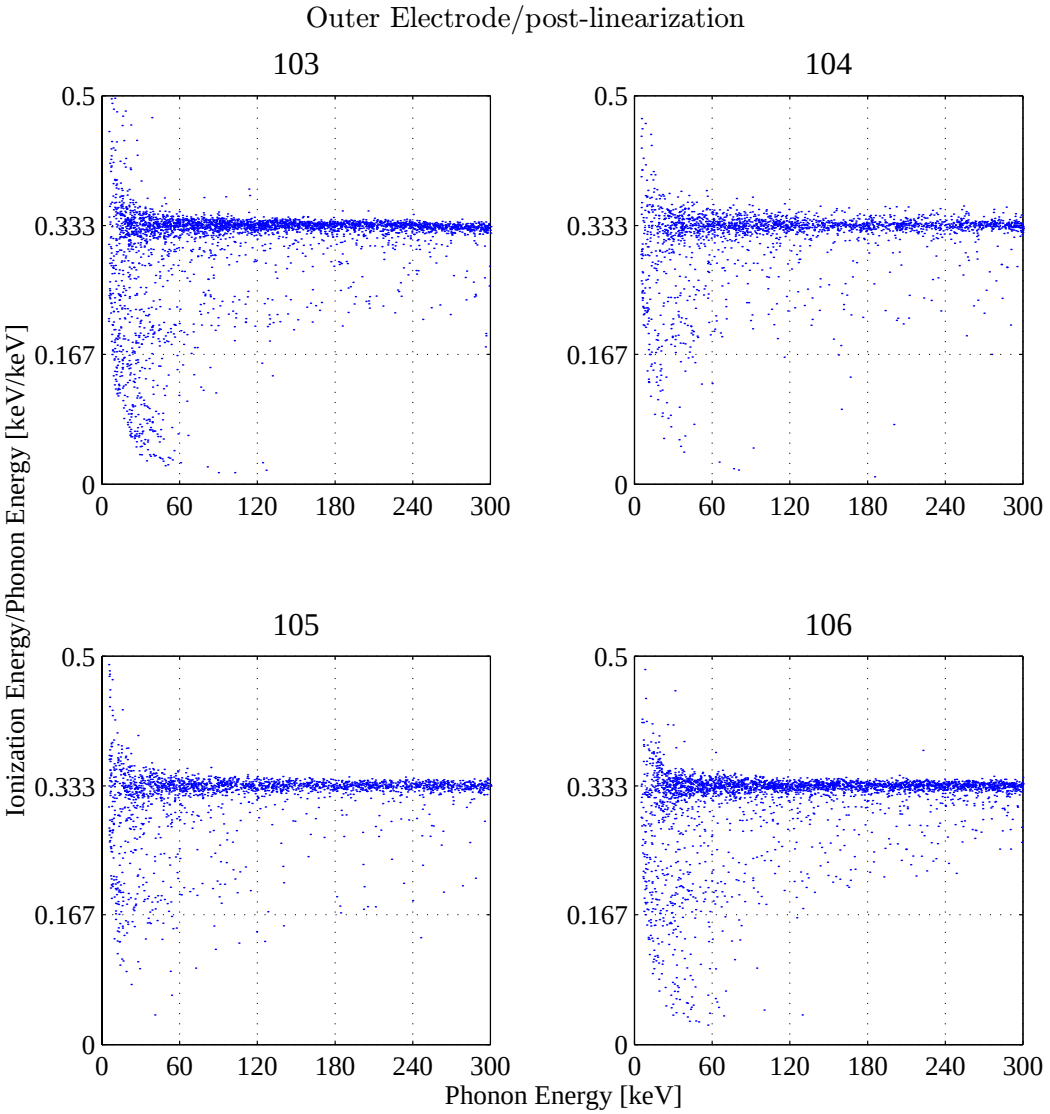

Figure 6.10: Outer-electrode ionization energy/phonon energy vs. phonon energy for veto-coincident cQoutOnly single scatters for a subset of post-April 3 data after linearization.

6.4.3 Automated Cut Determination

During physics analysis, it is necessary to make cuts on a number of trace parameters. These cuts are discussed in detail in the next chapter. A number of these cuts can be defined automatically using PipeCleaner:

- pretrigger-trace baseline (the XXbs RQ);
- pretrigger-trace standard deviation (the XXstd RQ);
- and χ^2 from ionization and phonon fits.

Pretrigger-Trace Cuts

The pretrigger baseline is the mean of the trace in the pretrigger region. The cut values were determined manually by looking at the distributions of the pretrigger-trace baselines, as is discussed in the next chapter. At present, these values are hardwired into PipeCleaner. PipeCleaner is used only to calculate the boolean mask for the cut and save it in the RRQ files, which saves substantial time during later analysis. It is unnecessary to make the pretrigger-trace-cut determination more adaptive.

To determine the pretrigger-trace standard-deviation-cut parameters, the mean and standard deviation of the pretrigger-trace standard-deviation values are calculated for events passing the pretrigger-baseline cut. Events more than 3 standard deviations from the mean are discarded and the process iterated until stable. The final cut is made at 3 standard deviations above the mean. There is no need to make a lower edge cut, as is shown in the following chapter.

The cut parameters are saved to the datum file.

Pulse-Fit-χ^2 Cuts

The pulse-fit-χ^2 cuts are nontrivial. Two issues arise: energy dependence of χ^2 and non-Gaussianity in the distribution of χ^2 at high energy. To illustrate this, Figures 6.11 and 6.12 show typical scatter plots of χ^2 vs. energy. Figures 6.13 and 6.14 show histograms of the χ^2 distributions for various energy ranges. At low energy, the distributions are well behaved. At high energy, though, pulse-shape changes result in severe deviation of χ^2 from its low-energy value: clearly, the templates are not good models for the pulses at higher energies. Andrew Sonnenschein discussed one model for this energy dependence, nonzero template noise [8]. Since templates are made from a finite number of pulses, they themselves are noisy. The signal-to-noise of the template is roughly $E_t\sqrt{N_t}/\sigma_E$, where E_t is the energy of the pulses averaged to make the template, N_t the number of pulses, and σ_E the zero-energy noise. This form holds under the assumption that the noise is energy independent. Pulses have signal-to-noise E_p/σ_E. Therefore, at a pulse energy of $E_p = E_t\sqrt{N_t}$, the pulse and template have the same signal-to-noise and the template noise contributes significantly to χ^2. This model predicts a dependence of the form $a[1 + (E/E_0)^2]$ for χ^2 as a function of energy.

This quadratic dependence appears to work well for the ionization channel. However, in practice, no cut is made on the optimal-filter-fit χ^2 for ionization pulses except to discard events where an optimal-filter fit could not be done, primarily saturated traces. Looking at Figure 6.11, it is clear that no cut is necessary – there is a well-defined distribution that contains essentially all

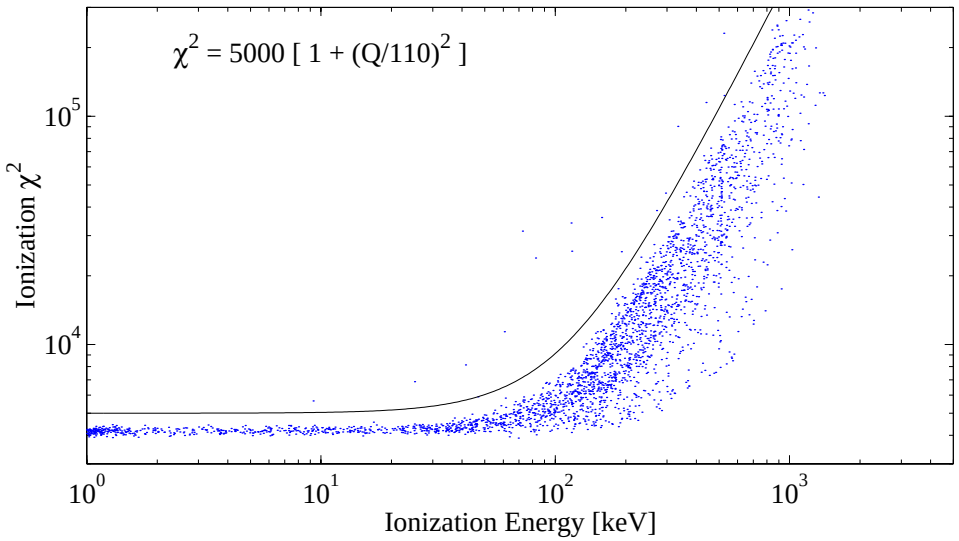

Figure 6.11: Typical ionization-pulse-fit χ^2 vs. energy. The ionization χ^2 is a true χ^2; at low energy, its mean value is about 4100, corresponding the sum of the number of samples in the inner- and outer-electrode ionization traces (4096) minus the number of fit parameters. The line has the quadratic functional dependence discussed in the text.

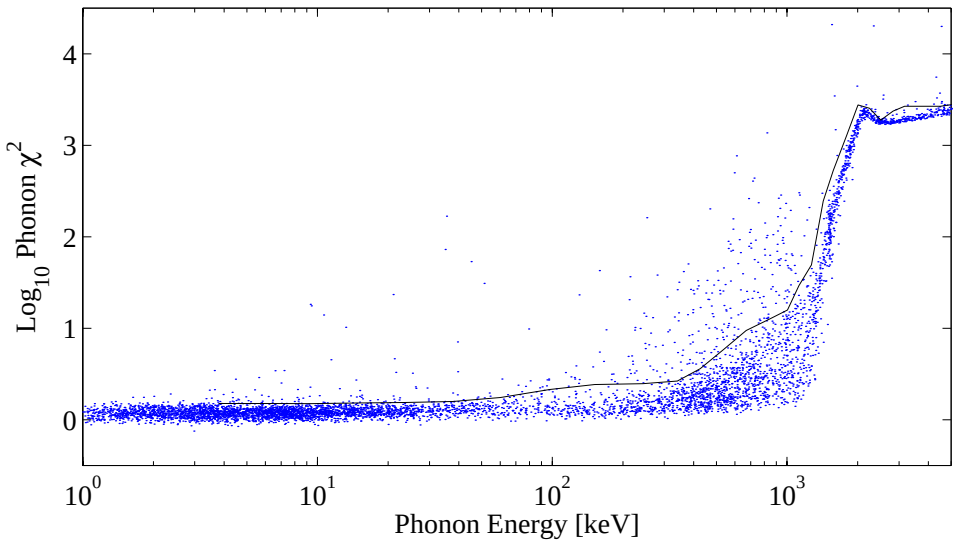

Figure 6.12: Typical phonon-pulse-fit χ^2 vs. energy. The phonon χ^2 is a reduced χ^2 for approximately 2000 samples, but, as noted in the text, it is not properly normalized. The line on the plot indicates the position of the cut calculated by PipeCleaner.

the events; there is no anomalous tail to high χ^2, though the distribution as a whole shifts to higher χ^2. For saturated events, the optimal-filter fit fails; however, at high energies, the signal-to-noise is good enough that an optimal-filter fit is not necessary and the standard template-fitted energy is used. Saturation can occur for low-energy pulses if the trace baseline is near the edges of the digitizer window, which can happen if the ionization amplifier is recovering from a muon event; such events are discarded in later analysis by pileup cuts, so it is not important to fit their energies well.

The quadratic model does not work well for phonon χ^2, as is clear from Figure 6.12. The slow rise away from the low-energy value is presumably due to minor pulse-shape nonlinearity as the energy is increased. The abrupt change at ~ 1 MeV coincides with the beginning of saturation. Furthermore, the χ^2 distributions change on timescales of one to a few days, indicating that the phonon χ^2 is sensitive to pulse-shape changes due to thermal drifts. Therefore, an empirical approach is taken in defining the phonon-χ^2 cut. The data are sliced into bins of equal logarithmic width in phonon energy. Near the sharp rise, the bins are narrowed. The χ^2 values for the data in each slice are histogrammed. The resulting histograms are rarely Gaussian; typically, they exhibit some reasonable smooth distribution with an obvious tail to high χ^2, as is seen in Figure 6.14. Usually, there is a discontinuity in the slope of the distribution. In principle, it is this point where the cut should be placed — the tail is due to anomalous events that have high χ^2 for possibly pathological reasons. In practice, it is difficult to take the second derivative of these distributions because there are not many events in each histogram. Empirically, making the cut at the point of most negative logarithmic derivative, $[1/P(\chi^2)]dP(\chi^2)/d\chi^2$, works well. It is not obvious that this is the correct place to make the cut. For a Gaussian, this point is at $+\infty$. But, intuitively, this makes sense. The purpose is to discard a tail of events that, by eye, do not seem to be part of the distribution. The reason such a tail is obvious to the eye is because there is sharply negative slope below the point and a more shallow slope above it. The algorithm based on the logarithmic derivative finds this feature. The point of most negative absolute slope is not used because it usually occurs near the peak of the distribution, where the number of events is largest. In practice, it is necessary to add a small positive number, *e.g.* 3, to the denominator to prevent a singularity at bins with no counts. The cut values thus determined are Savitzky-Golay filtered to yield a cut that is a smooth function of phonon energy. The energy bins and cut values are saved to the datum file. A typical cut determined by this automated technique is shown in Figure 6.12.

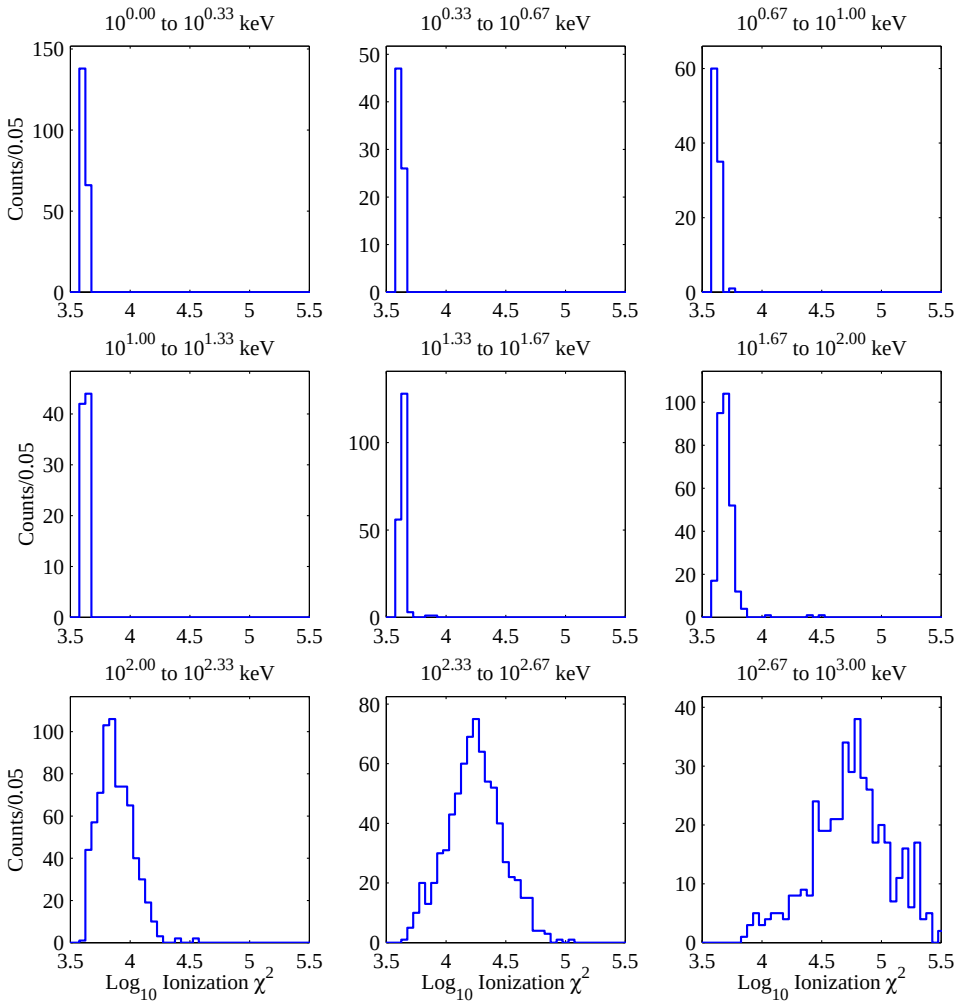

Figure 6.13: Typical ionization-pulse-fit-χ^2 distributions for various ionization-energy ranges. This is the same data set as shown in Figure 6.11. Note the extreme change in width of the distribution between 100 and 1000 keV.

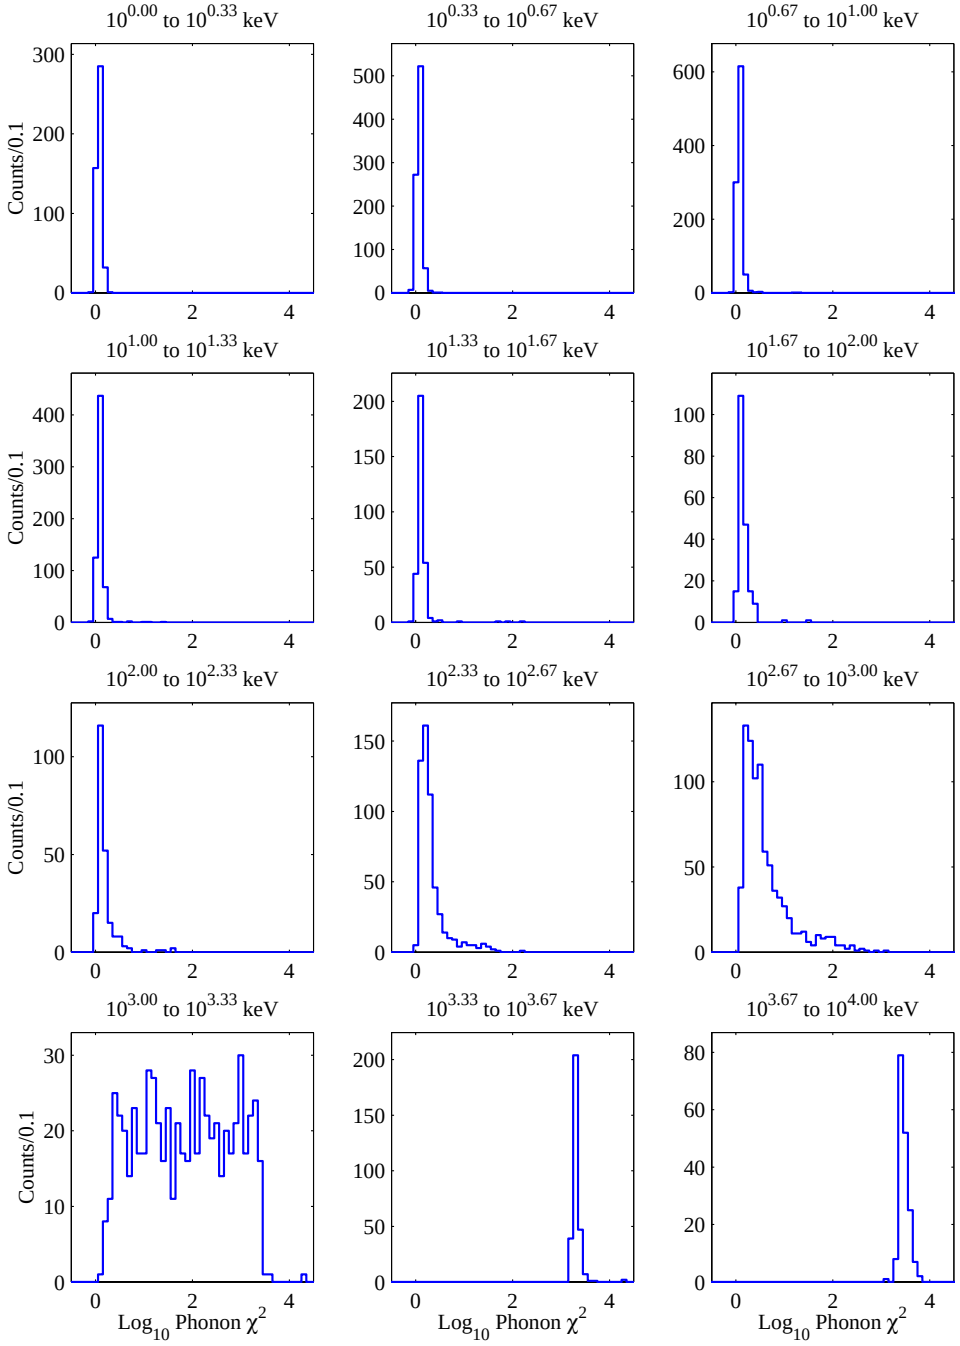

Figure 6.14: Typical phonon-pulse-fit-χ^2 distributions for various phonon-energy ranges. This is the same data set as shown in Figure 6.12. Note the tails on the distributions between 100 and 1000 keV.

6.4.4 Noise-Parameter Determination

PipeCleaner also histograms the energies for random triggers and performs Gaussian fits to the distributions. The fit parameters are recorded in the datum file. Fits are done for both the "sliding" and "zero-delay" energy distributions. The fit parameters are later used to set ionization thresholds using the sliding or zero-energy noise fits. The Gaussian fits are typically very good for the ionization channels because of the stability of the noise. The phonon channels' distributions tend to possess tails to positive energy, presumably due to thermal disturbances, but the phonon noise parameters are not used in later analysis.

6.4.5 Conditions

The datum files can be modified to include "detector-death"-type conditions. The main use of this feature is to intentionally "kill" a detector during periods where it was known to be malfunctioning — *e.g.*, ionization bias failed, trigger outburst, etc. These conditions are propagated via their inclusion in the "cGoodTime" cut mask generated during the "crunching" phase.

6.4.6 Crunching

PipeCleaner "crunches" the RQs to RRQs by saving new quantities that have been derived from the RQs using the calibration parameters discussed above. These consist of corrected phonon and ionization energies, recoil energy, phonon and ionization partition (defined earlier), ionization yield ($Y_R = E_R/E_Q$), etc. The noise parameters determined for the series are saved in arrays equal in length to the data set, selecting the sliding noise parameters if the given detector was the triggering detector and phonon triggered and otherwise using the zero-delay noise parameters.

In addition, PipeCleaner generates mask files that contain boolean cut masks calculated from the data using the cut parameters determined above. Included in this is the "cGoodTime" mask, which is discussed in detail in Chapter 7.

In general, conditions are not determined until the data has been passed through PipeCleaner once. Therefore, the crunching step is usually repeated. The data are crunched, conditions are determined and incorporated in the datum files, the data are recrunched and inspected again, etc. This cycle works well and is likely inevitable. For CDMS II, I expect this procedure can be formalized and coded so that fewer of the steps require oversight.

Chapter 7

Data Set, Cuts, and Efficiencies

7.1 Chronology and Raw Data Set

Run 19 took place from October, 1998, through September, 1999. A chronology is given in Tables 7.1, 7.2, and 7.3. A total of 96.4 raw live-days of data were acquired, corresponding to 63.6 raw kg d summed over BLIPs 3 through 6. Raw live-days denotes the live time of the DAQ system, before any cuts are made, excepting periods where the raw data are completely discarded due to obvious problems. Though BLIPs 1 and 2 were alive during the first half of the run, they are neglected when quoting exposures because they do not contribute to the WIMP sensitivity due to their electrodes. In Figure 7.1, the integrated real and live times for which the DAQ was on and taking dark-matter data (*i.e.*, excluding grounding, pulser runs) is shown. For the DAQ live time alone, the largest slope observed is ~ 0.6 live-day/real day, and there are of course many periods where that number is not achieved. In Table 7.4, the expectations for data-taking efficiencies are shown. The lowest efficiency expected during stable running is 67%, which is only about 10% higher than is observed: most of the dead time is understood. Dead-time reductions can be made in every category except for cryogen transfers. However, it should be realized that, even with 0.95 efficiencies individually for cryogen transfers, grounding, and DAQ, and without phonon-pulser runs, one still only expects a live fraction of $0.95^3 = 0.86$. In addition to the normal losses during running, there are significant blocks of time where the data rate drops to zero. Early in the run, the bulk of these were due to DAQ and electronics debugging and detector-operation optimization. From April onward, downtime occurred primarily because of assorted phonon-system pathologies and veto repairs and upgrades.

Pathologies in the phonon-channel behavior resulted in a nonnegligible loss of live time during Runs 19A and 19B. Such periods are easily cut out of the data set because they have obvious symptoms, such as phonon-trigger-rate excursions, obvious changes in phonon-pulse shape, or large excursions in the thermistor DC refs. The symptoms and possible explanations of this behavior are discussed in Appendix C.

Ionization- and phonon-trigger-rate distributions for the entire run for each detector are shown in Figure 7.3. The ionization trigger-rate distributions are very clean. The tails to low rate are an artifact of the way the distributions are calculated. The phonon trigger-rate distributions are not as clean, but clearly the periods with trigger-rate excursions constitute a very small part of the data set (note that both axes are logarithmic).

Table 7.1: Run 19A Timeline

Aug	–	Sep, 1998	Removal of original complement of 3 striplines, installation of full-size internal lead shield, 18 new striplines and 2 old striplines
			Debugging and testing of Run 19 detectors and tower
Oct 8			Detector testing at UCB test facility completed
Oct 9	–	Oct 16	Installation of BLIP1/2 and BLIP3–6 towers
			Modifications to monitoring system external wiring
Oct 17	–	Oct 21	Run 19A cooldown
Oct 22	–	Nov 3	Electronics debugging, phonon-bias optimization
Nov 3	–	Nov 11	Ionization-bias scan: ± 4 V, ± 6 V, ± 8 V
Nov 12	–	Dec 4	Low-background data at 6-V ionization bias on all detectors; live fraction low due to computer crashes
Dec 4	–	Dec 9	Crashing DAQ problem traced to bad memory chip on Pilot's G3 upgrade card
Dec 10	–	Dec 14	DAQ back online; "slow pulses" begin
Dec 14	–	Dec 27	Low-background running attempted at 4-V ionization bias on all detectors, but assorted problems and repairs cause significant dead time
Dec 27			IVC leak during LHe transfer causes warmup; Run 19A ends

Table 7.2: Run 19B Timeline

Jan 5	–	Jan 15, 1999	Run 19B cooldown
Jan 17	–	Jan 27	Modifications/repairs to front-end and RTF electronics
Jan 28	–	Mar 21	Low-background running at 4-V ionization bias with some electronics and detector-bias optimization
Mar 22	–	Mar 30	Stable low-background running at 6-V ionization bias
Mar 31	–	Apr 5	Neutron calibration
Apr 3			Campus-wide power outage; nonlinearity appears
Apr 5	–	Jun 4	Mostly stable running, but with phonon-trigger outbursts
May 28			BLIP1/2 turned off due to appearance of "slow pulses" and sufficient data acquired
Jun 4			Halt run to determine cause of phonon-channel problems (trigger outbursts and slow pulses)
Jun 14			Attempt > 4 K warmup to check helium-film hypothesis; refrigerator runs away due to IVC/OVC leak; end of Run 19B

Table 7.3: Run 19C Timeline

Jun 20	–	Jun 30	Run 19C cooldown
Jul 1	–	Jul 9	Low-background running at 1-V ionization bias for electrode study
Jul 10	–	Jul 15	Low-background running at 6-V ionization bias
Jul 15	–	Jul 20	Photon calibration (last day at 1-V ionization bias)
Jul 20	–	Sep 14	Low-background running at 6-V ionization bias
Jul 20	–	Jul 30	5% veto hole in West counter; repaired Jul 30
Aug 3			Veto ADCs debugged — good veto ADC data available from here on
Aug 27	–	Aug 31	Pot pump out of commission; circulation halted and refrigerator maintained at 4 K
Sep 15	–	Sep 20	Low-background running at 1-V ionization bias
Sep 21	–	Sep 22	Photon calibration at 1-V ionization bias
Sep 23			Neutron calibration
Sep 23	–	Sep 26	Data taken without top polyethylene
Sep 27	–	Sep 28	Debugging of IVC leaks while cold
Sep 28			Begin warmup; end of Run 19C

	Efficiency	
Loss Mechanism	min	max
Cryogen transfers	0.92	0.96
Grounding	0.91	
Phonon pulser runs	0.95	0.96
DAQ dead time	0.84	
Overall	0.67	0.70

Table 7.4: Run 19 data-taking-efficiency summary. Estimated minimum and maximum expected efficiencies for those that can vary or whose exact value is not known are indicated by the "min" and "max" columns.

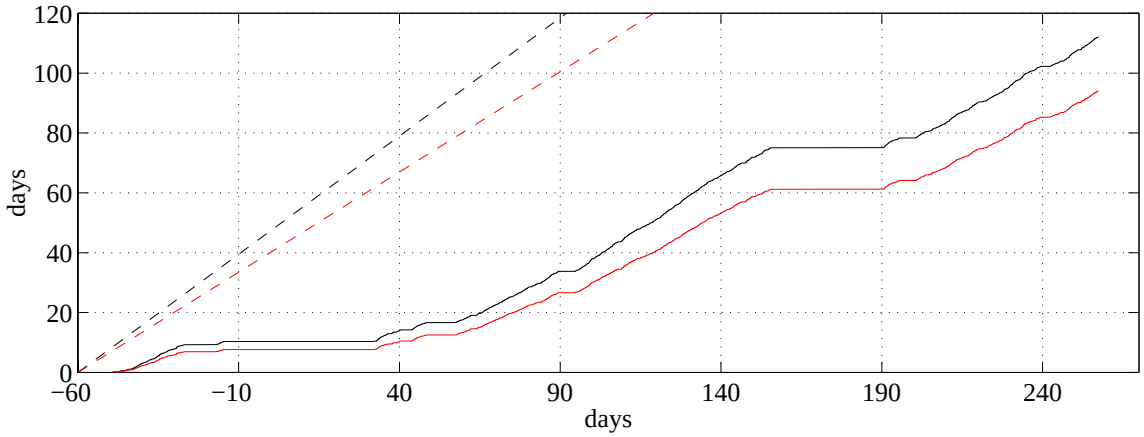

Figure 7.1: Cumulative data-acquisition real time and live time. The top curve indicates the actual time that the DAQ was operating; the lower curve indicates when it was "live" — *i.e.*, waiting for a trigger. The two dashed lines correspond to slopes of 0.70 and 0.67, the expected maximum and minimum slopes from Table 7.4. The origin of the horizontal axis is Jan 1, 1999.

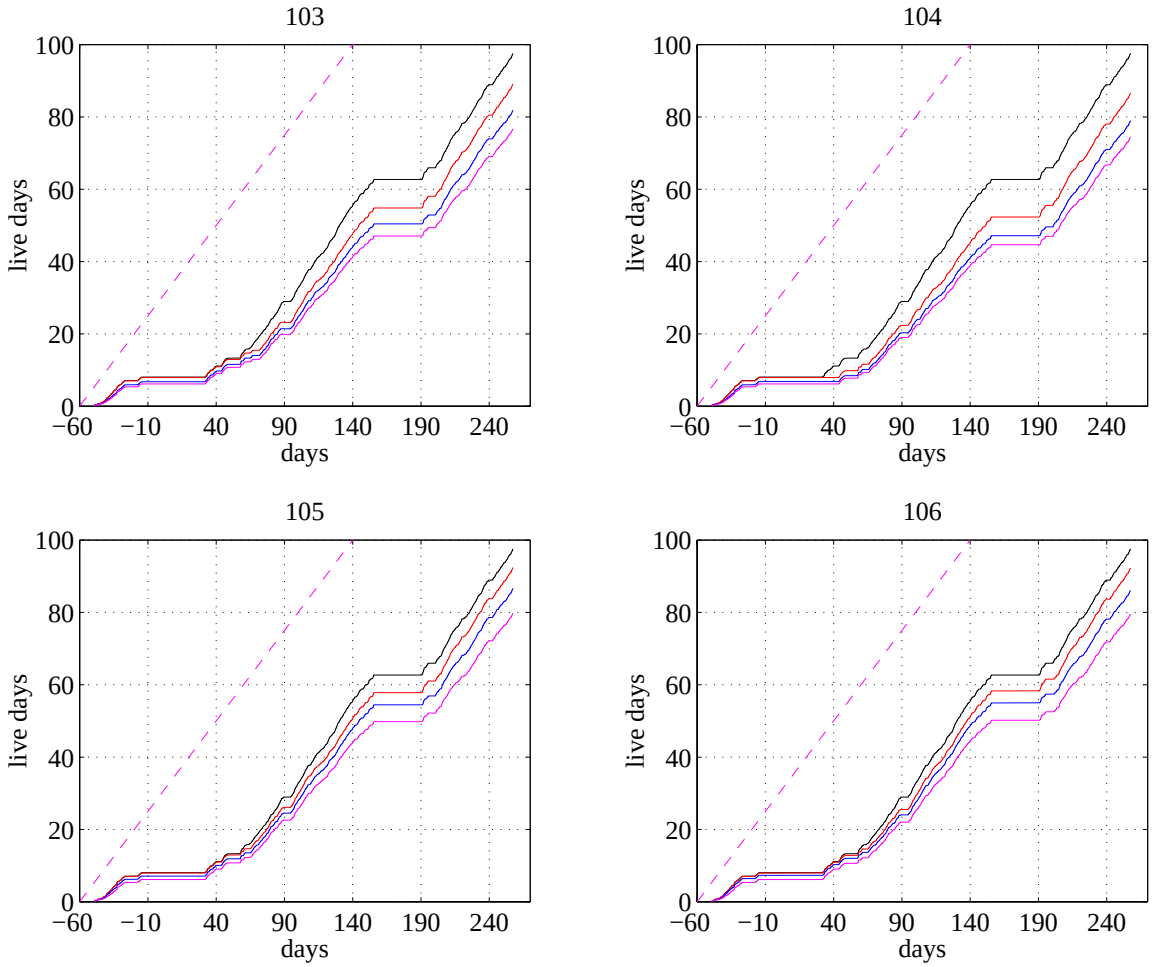

Figure 7.2: Cumulative live time vs. real time for BLIPs 3 through 6 and for various cuts. The data sets are, from top to bottom: all events that are not random and not detector-pulser events (same as bottom curve in Figure 7.1); and passing cGoodTime (detector is "alive"); and passing pretrigger cuts for the given detector (cPreQual for the detector); and passing pretrigger cuts for all detectors that are "alive" (cPreQual(1)). The dashed lines have slope of 0.5, which matches the lowest curves well. The origin of the horizontal axis is Jan 1, 1999. See Section 7.2 for definition of the cuts.

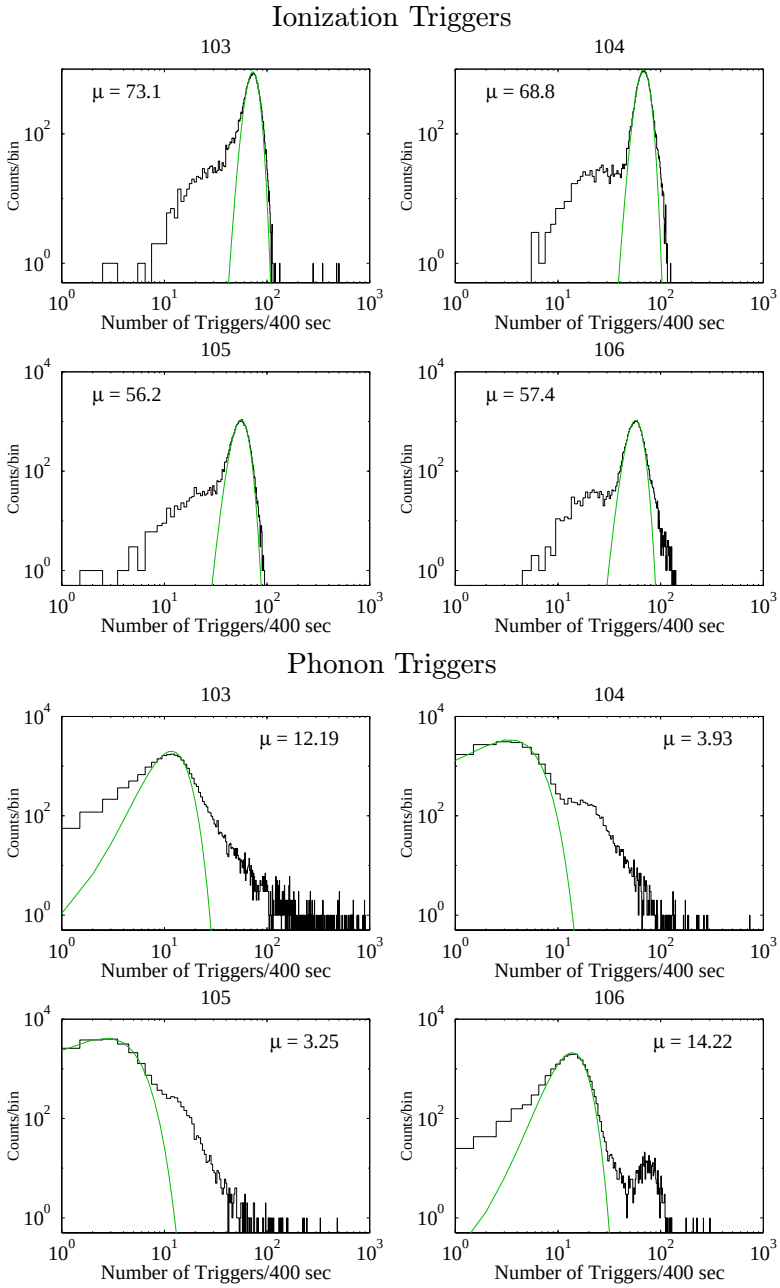

Figure 7.3: Histograms of the number of triggers observed per 400-s block of live time for the entire run. Note that both scales are logarithmic to clearly show the outlier behavior. The overlaid curves are Poisson fits, with the best-fit mean μ shown. Note also that an event is counted as a trigger for a given detector only if that detector was the first-triggering detector; this is why the apparent mean trigger rate is about 0.15 Hz instead of 0.33 Hz.

7.2 Data-Quality Cuts

To prepare the data for a search for WIMP-induced nuclear recoils, a number of data-quality cuts are made. The goal of these cuts is to remove pileup, periods of high noise or trace baseline wandering, and to select only those events where the pulse fits are of sufficient quality to ensure the accuracy of the energy estimate. The last point is critical: since the nuclear-recoil sensitivity of CDMS detectors is founded on high-resolution measurements of both phonon and ionization energies, degradation or blatant inaccuracy of these measurements must be avoided at all costs. Overestimation of the phonon energy is especially troublesome, as it pushes events to low ionization yield.

As Andrew Sonnenschein pointed out in his dissertation [8], cuts on the phonon- and ionization-pulse-fit χ^2 can be used to eliminate many different types of problem events at once, but one is left without any understanding of the causes of the losses. To better understand such losses, cuts are defined with specific rationales in mind and the effects of the individual cuts are checked. This set of cuts relies heavily on the experience of Andrew Sonnenschein and Richard Schnee from Run 18.

All cuts are named and begin with a lower case "c" to indicate "cut." In general, a cut on a given reduced quantity (RQ) has a name that is the RQ name prepended with a "c"; *e.g.*, cPstd. The plots shown in this section are made for a randomly selected 10% of the data unless otherwise indicated. Also, in general, it is unnecessary to show figures for both phonon sensors of a given detector because they behave similarly.

7.2.1 Pretrigger-Trace-Quality Cuts

As mentioned in the previous chapter, a number of cuts are made on pretrigger-trace quantities. These are made to ensure the traces are free of pileup, within the digitizer window, and that the noise environment is reasonable. In this section, I describe the justifications for these cuts.

- **cP1bs, cP2bs**: These are cuts on the mean pretrigger baseline in the phonon channels, P1bs and P2bs, in units of digitizer bins. The goal of these cuts is to ensure the pretrigger trace is in a reasonable range – not so close to the bottom of the digitizer window that noise could push samples below the edge, and not so close to the top of the window that saturation occurs at an unexpectedly low energy. As mentioned in Chapter 5, either case can occur if the thermal disturbances cause the thermistor carrier amplitude to fluctuate too quickly for the DC-ref subtraction to compensate ($\tau = 66$ s). Another case is events preceded by large events, *e.g.*, through-going muons. The event under consideration may be on the steeply falling tail of a muon event or on the recovery from undershoot due to AC coupling. The choice of the cut is made by comparing the distribution of the P1bs and P2bs parameters for unsaturated and saturated events: if the phonon baselines were absolutely stable, then there would be no correlation between whether an event is saturated (high energy) and the baseline. The distributions of P1bs are shown in Figure 7.4. As one can see, the saturated-event distribution deviates from the unsaturated-event distribution very close to the bottom of the window and above about 3000 digitizer bins. These distributions are stable over the run, so a fixed acceptance range of (15, 3000) is used for all detectors and all times. The logical AND of cP1bs and cP2bs is cPbs.

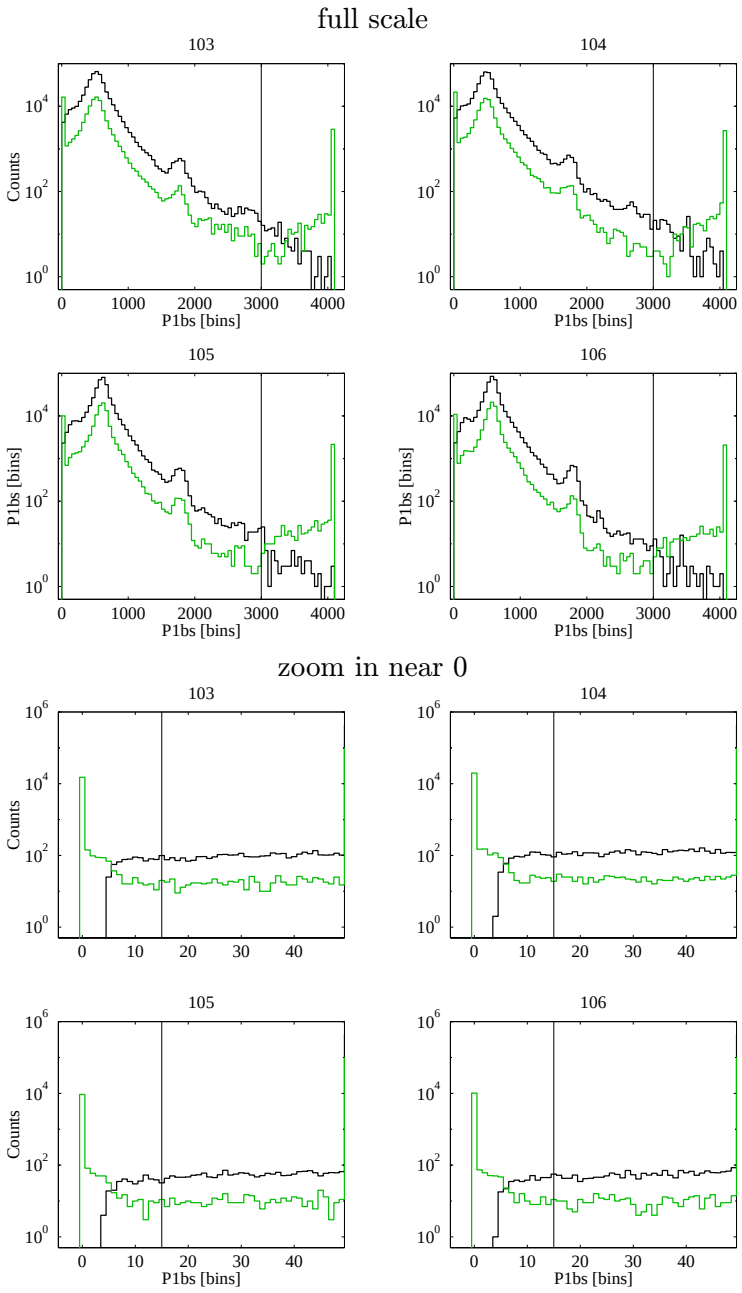

Figure 7.4: Distributions of mean phonon pretrigger-trace baseline for sensor 1, P1bs. Dark: unsaturated events for a random 10% of the data. Light: saturated events. The solid vertical lines show the cut position. The lower set of plots zooms in near 0 to show clearly the position of the lower cut. The main peak in the accepted region in the top plots corresponds to the nominal baseline position. The secondary bump corresponds to a short period when the baselines were artificially offset. Plots for P2 are not shown because they are identical.

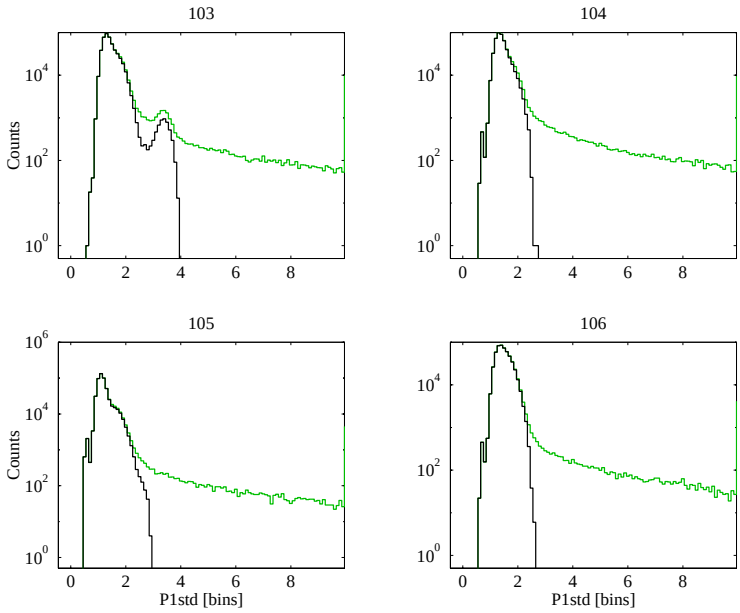

Figure 7.5: Distributions of P1std, the phonon pretrigger-trace standard deviation for sensor 1. Dashed: distribution before cut. Solid: distribution after cut. The bimodality in the BLIP3 plot occurs because of a glitch (presumably electronic) that appears occasionally in the traces — such events are accepted because the glitch has little effect on the pulse fit.

- **cQIbs, cQObs**: These cuts are the analogues of cP1bs and cP2bs for the ionization channels. They reject very few events: the ionization-trace baselines are intrinsically very stable because the DC level has no dependence on the detector behavior. Also, pileup occurs far less frequently because the ionization pulses are very short, $\tau_f = 40$ μs. The same range, (15, 3000), is used. cQbs is also defined analogously.

- **cP1std, cP2std**: These are cuts on the standard deviation of the pretrigger baseline, which is the standard deviation of the ADC values for the samples in the pretrigger region. These cuts remove low-level baseline wandering that increases the baseline noise, periods of high phonon-channel noise, and pretrigger pileup. Typical distributions for P1std and the effect of the cuts are shown in Figure 7.5. Only events passing cPbs have been plotted, as cPbs removes events with zero or anomalously low P1std and P2std due to part of the trace being below the digitizer window. The distributions have three main characteristics: a clear central peak corresponding to "normal" events, a hard lower edge, and an extended, exponentially decreasing tail. The absolute lower edge results from quantization noise. The exponential tail is due to pretrigger pileup. Events with any kind of pileup may be discarded while maintaining high efficiency because of the low trigger rate (0.33 Hz/detector).

  As discussed in the previous chapter, PipeCleaner is used to determine these cuts adaptively based on the data. This is necessary due to varying noise conditions — a fixed cut would be excessively harsh during some periods and excessively loose during others. The distributions

of events passing the combined cPstd cut are also shown in Figure 7.5. The smooth upper end of the post-cut distribution reflects the variation in the adaptive cut position.

- **cQIstd, cQOstd**: These cuts on the ionization pretrigger-trace standard deviation also reject very few events because the ionization-noise environment is stable. Moreover, the pretrigger-trace standard deviation is dominated by cross-talk of the 1-kHz phonon bias and so is not very sensitive to small changes in the ionization-noise environment. The cuts do appear to eliminate some pileup, though these events could probably have been eliminated by cQPileup (see below). These cuts are determined adaptively in the same way as the cPstd cuts. The distributions of events before and after the cut are shown in Figure 7.6. The distributions for QO are wider than for QI because the 1-kHz cross-talk to QO is larger: the capacitance to the neighboring detector's thermistors is larger for QO than QI.

- **cPdcref**: As noted earlier, there are periods during which, due to refrigerator instability, electronics anomalies, or other unknown causes, the phonon-lockin DC refs leave their normal range, indicating the detector temperatures have changed. Small variations in DC ref can be corrected by the "DC-ref correction" discussed in the previous chapter. However, large excursions leave the range where a linear correction can be done, so such times are discarded. Figure 7.7 illustrates the cut. The distributions are clearly non-Gaussian, and different ranges can be considered "normal" at different times. Too loose a cut degrades the phonon measurement resolution due to uncorrectable drifts; too tight a cut reduces the live time unnecessarily. Since neither the resolution nor the live time degrade quickly with the position of the cut, the exact position of the cut is not critical and it is determined manually.

- **cGoodTime**: This cut is used to discard periods of known low-quality data. For the early part of Run 19, its primary use is to discard a detector's data when its electronics were in need of repair. Another use is to cut periods of DAQ or veto malfunction. The cut also includes cPdcref.

- **cErrorMask**: cErrorMask requires the ErrorMask RQ to be 0. ErrorMask is a mask with bit assignments for a number of trigger-related problems. For example, the global trigger is sent to the three independent banks of the history-buffer module to allow synchronization; ErrorMask indicates if one of these synchronization bits is not found.

- **cPreQual** cPreQual is the logical AND of the above cuts. A general cPreQual, cPreQual(1), is defined to be the logical AND of cPreQual for those detectors out of BLIPs 3 through 6 that pass their own cGoodTime cuts. The goal of cPreQual(1) is to ensure that all live detectors have clean pretrigger traces. In principle, such a harsh cut is not necessary. For example, if one detector has pretrigger pileup and the others do not, the first detector can be discarded and any possible multiple-scatter events reclassified as single scatters. The rate of such misidentification is not significant and can be estimated. The main reason for the requirement is due to nonidealities in the electronics and detectors, cross-talk, etc. For example, if one detector has a pretrigger muon, its ionization channel can electrically cross-talk to other detectors as it comes out of saturation.

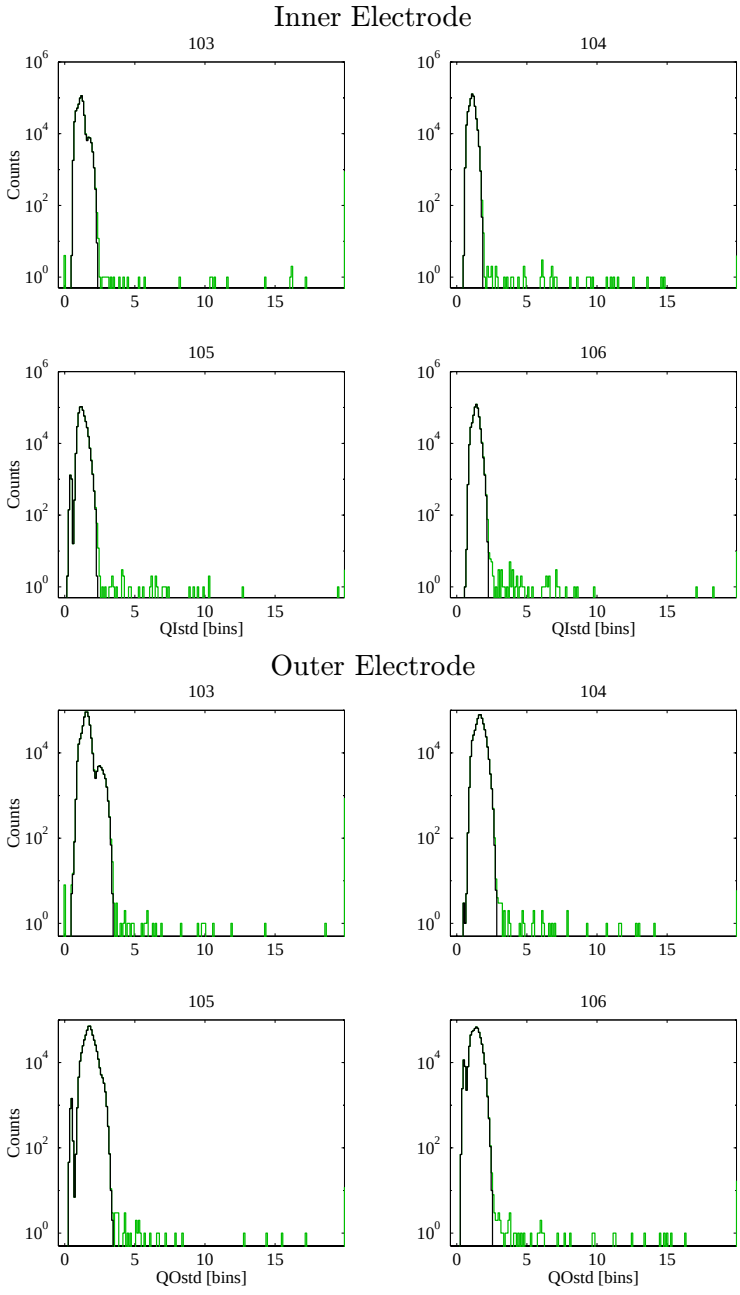

Figure 7.6: Distributions of ionization pretrigger-trace standard deviation for QI (top plots) and QO (bottom plots). Light: distribution before cut. Dark: distribution after cut.

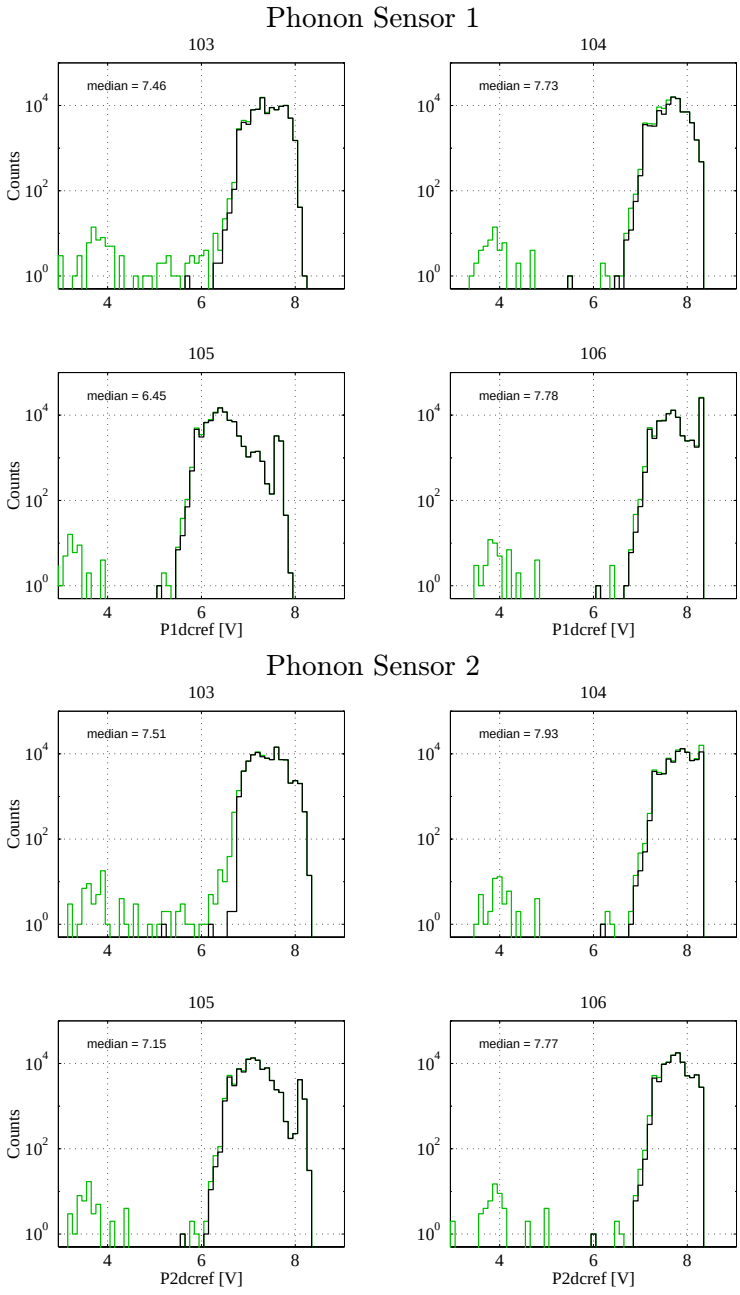

Figure 7.7: Distributions of thermistor DC-ref data. Light: before cut. Dark: after cut.

7.2.2 Posttrigger-Trace-Quality Cuts

Additional cuts are used to eliminate posttrigger pileup. In principle, the QPreTime cut should be included in the pretrigger cuts, but it is more convenient to combine it with the QPostTime cut.

- **cQPileup**: This cut checks the history buffer for pretrigger or posttrigger ionization triggers using QPreTime and QPostTime. These RQs contain the times of the last ionization trigger prior to and the first ionization trigger after the Global Trigger time. Distributions of QPreTime and QPostTime are displayed in Figure 7.8. The large peak at the origin corresponds to events coincident between detectors. For the majority of events, if there is no coincident trigger, then QPreTime and QPostTime time out — no additional triggers are found in the history buffer, which is consistent with the 20-ms width of the buffer and the typical $\sim$ 3-s time between events. An acceptance window extending from -50 μs to 300 μs is necessary to accept coincidences. The large posttrigger size is needed presumably because of double triggering in the electronics, not because of physical delays of this size. In the QPostTime distribution, in addition to the coincidence peak, a flat distribution of accidental triggers is observed. The distribution appears flat (instead of exponential) because the 10-ms region shown is small compared to the average time between events of $\sim$ 3 s. The QPreTime distribution is more complex. The flat distribution with QPreTime $\lesssim -9$ ms arises from accidentals. The bump between -5 ms and 0 ms arises from events in which the Global Trigger was enabled between the ionization and phonon pulses for an event. For such events, the system triggers on the phonon pulse and the ionization trigger appears in the pretrigger history. Such events are possible because the phonon pulse is intrinsically very slowly rising, the lockin low-pass filter introduces a time delay, and the phonon-trigger filter performs another low-pass-filtering step to optimize the energy threshold. From the distribution, it is seen that the delay can be as large as 5 ms. These events are also rejected, though it is possible they may be usable. The fraction of events rejected is very small, $\sim$ (5 ms)/(3 s) = 0.17%. Explicitly, the cQPileup cut is defined to retain two classes of events: first, events in which both QPreTime and QPostTime time out; and second, events in which a coincidence occurs, with QPreTime > -50 μs and QPostTime < 300 μs. All other events are rejected.

- **cPst**: As described in Chapter 6, the P1st and P2st RQs indicate the time of the last sample within 5 standard deviations of the trace baseline that occurs before the largest peak in the pulse. Thus, these RQs are sensitive to posttrigger pileup where the second event is larger than the first. Surprisingly, the distributions of these RQs are quite well behaved in spite of the lack of any filtering in their determination. The behavior of these parameters is dependent on the trigger type, as is shown in Figure 7.9. Posttrigger pileup can be seen clearly in the on-detector ionization-trigger plot as a tail of events to high P1st. The sloping at very high energies occurs because the pulses are large enough that the change of slope of the pulse rise becomes apparent above the noise. The flaring of the ionization-trigger and off-detector phonon-trigger distributions at low energy follows E^{-1}, as expected for a time-like quantity derived from a pulse (see Appendix B for a discussion of the expected noise of time-offset estimation). The strong energy dependence in the on-detector phonon-trigger case reflects the energy-trigger-time relationship for phonon-triggered events. The cuts are defined to accept the bulk of the distributions; they are indicated in the plots.

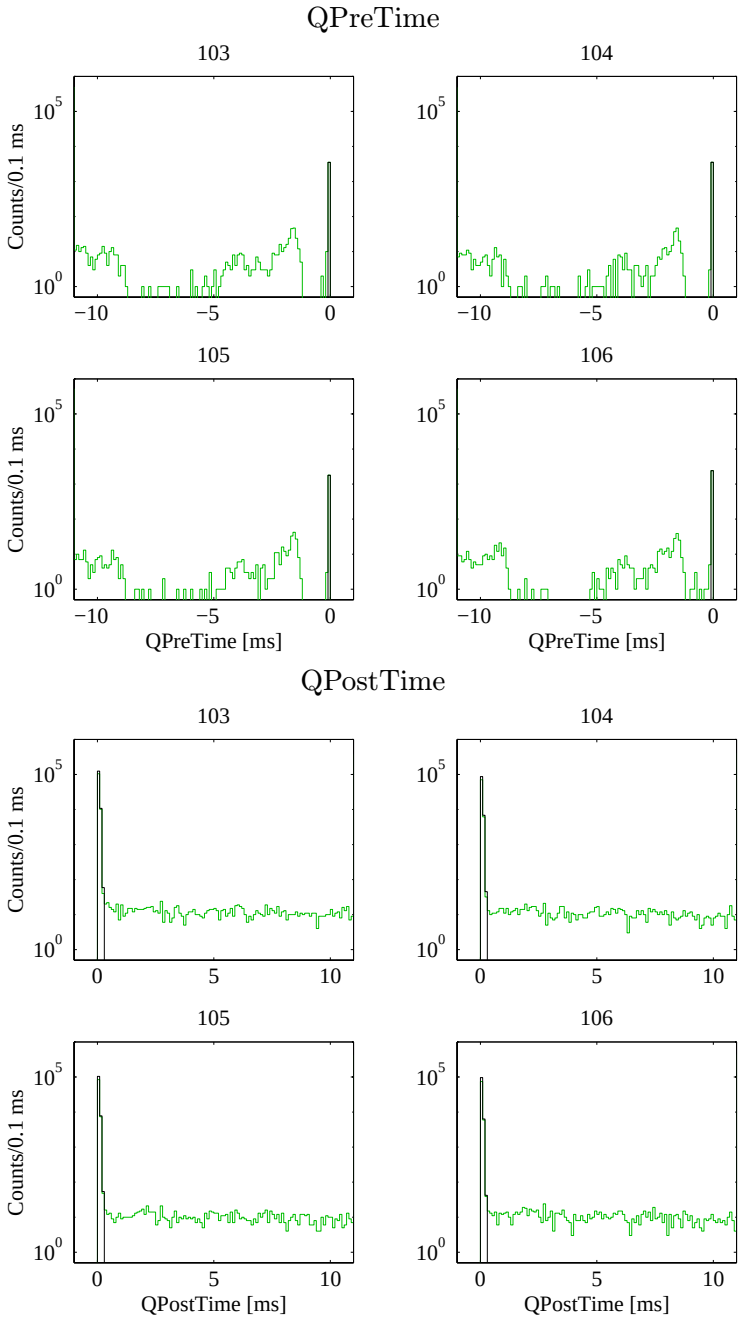

Figure 7.8: Distributions of QPreTime and QPostTime. Dark: events passing cut. Light: events failing cut. For most events, QPreTime and QPostTime time out and so the events do not appear. The distributions are explained in the text.

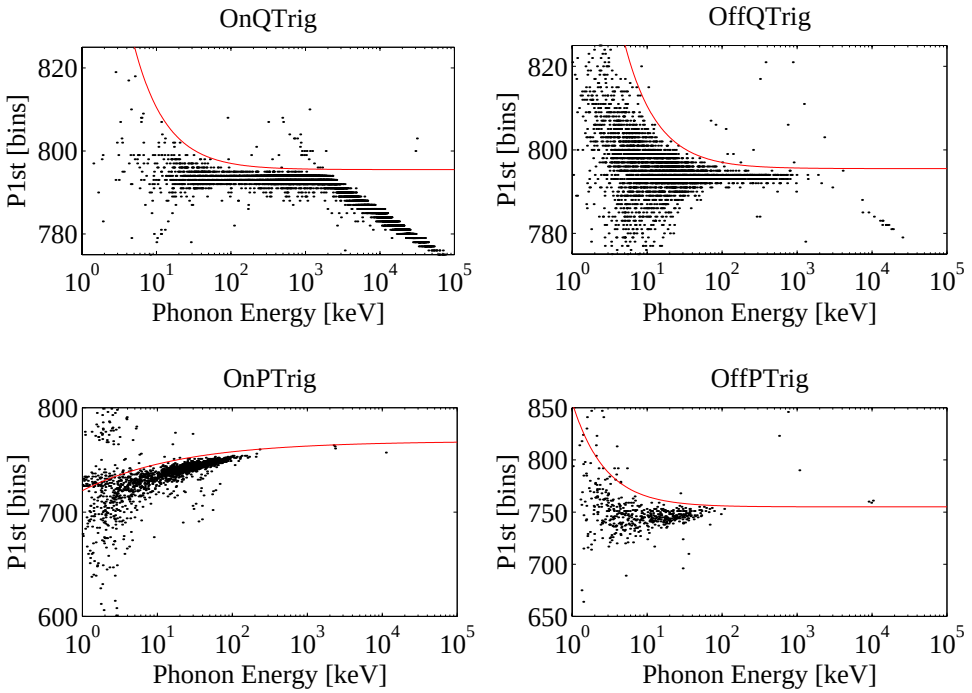

Figure 7.9: Typical plots of P1st vs. phonon energy for on-detector ionization triggers, off-detector ionization triggers, on-detector phonon triggers, and off-detector phonon triggers. The cuts are indicated in the plots; events above the lines are cut.

- **cBaseQual**: This cut is the logical AND of cPreQual, cQPileup, and cPst, so named because it ensures the quality of the entire trace baseline by removing pretrigger and posttrigger pileup.

Figure 7.10 illustrates the usefulness of the pretrigger- and posttrigger-trace cuts by comparing the distribution of phonon χ^2 before and after the cuts. Clearly, the cuts defined so far remove the great majority of pathological pulses.

7.2.3 Pulse-Quality Cuts

These cuts are intended specifically to check the pulses themselves. It is preferable to discard pathological pulses and accept the resulting efficiency loss rather than accept them with misestimated energies.

- **cpchisq, cqchisq**: Cuts on χ^2 of the phonon- and ionization-pulse fits. Recall that, while called a χ^2, the distribution of phonon χ^2 does not match a statistical χ^2 distribution. As discussed in Chapter 6, variations in χ^2 appear due primarily to pulse-shape changes. Therefore, standard, statistically rigorous techniques for determining the natural cut position are not applicable, and the adaptive algorithm described in Chapter 6 is used. Given the cleanliness of the phonon-χ^2 distribution after application of the previous cuts, it is questionable whether such a cut is necessary. At this point, the cut is historically embedded in the analysis. The

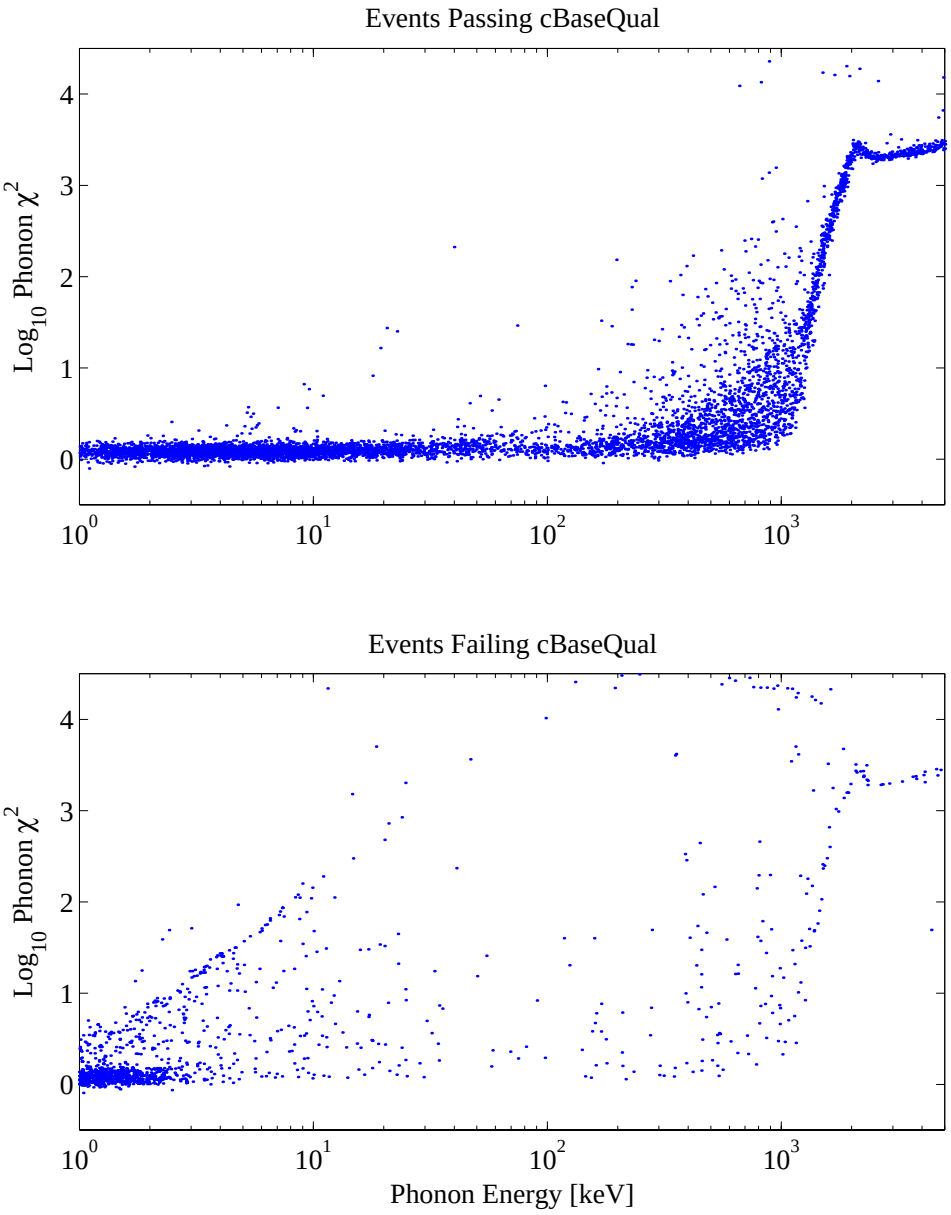

Figure 7.10: Typical phonon-pulse-fit χ^2 vs. energy for events passing (top) and failing (bottom) cBaseQual. A detector's hardware trigger must fire for it to pass cBaseQual; events without hardware triggers account for the band of normal events at 1–2 keV being discarded.

adaptive method for defining the cut was only developed during Run 19C and has not been applied to the Runs 19A and 19B data. The Run 19A and 19B cuts were determined by eye.

Because the ionization χ^2 is well behaved, a cut on ionization χ^2 is barely necessary. A very liberal cut is made, accepting all events for which an optimal-filter fit could be done. The adaptive algorithm is unable to find a reasonable cut point because the distributions are lacking in features the algorithm was designed to find.

The ionization-χ^2 cut presents one problem. Due to the high interdetector capacitance in the close-packed geometry, there is significant cross-talk between detectors. No attempt is made to correct for this. It would be possible to perform a common fit, including cross-talk, as is done for the inner and outer electrodes of single detectors, but this additional complication has not been tackled yet. Because no correction is made, a significant tail to high χ^2 for low-energy ionization pulses is seen when any adjacent detector exhibits saturation in the ionization channel. The ionization-χ^2 cut is therefore defined using a data set in which no detectors are saturated. The cut rejects a large number of events in which there are saturated detectors. However, as mentioned earlier, all multiple scatters can be discarded without any loss of sensitivity for the WIMP-search analysis. Another data set that is important for the final analysis is the veto-anticoincident double-scatter neutrons. However, the simulations show that such neutrons are rarely accompanied by a high-energy particle since the neutron is produced outside the shield. The neutrons themselves deposit small energies (as discussed in Chapter 3). Therefore, the cut's efficiency for these events is high. Another data set that may be biased by this effect is the veto-coincident neutrons. These neutrons are produced in the inner-lead shield and the copper of the cryostat, so accompanying high-energy photons have a better chance of interacting in the detectors simultaneously. As is discussed below, it is seen that the simulations, efficiency-corrected using the "single-scatter efficiency," match the data well, indicating that only a small fraction of multiple-scatter veto-coincident neutrons are rejected because of a coincident high-energy detector interaction.

- **cppart**: The phonon partition quantity is $(\mathrm{P1}-\mathrm{P2})/(\mathrm{P1}+\mathrm{P2})$. "Partition" is a misnomer, since the energy is not partitioned as the ionization is; the name is given because the expression is analogous to the ionization partition, $(\mathrm{QI}-\mathrm{QO})/(\mathrm{QI}+\mathrm{QO})$. As described in Chapter 4, the ppart quantity identifies thermistor interactions, so a cut is defined to discard these events. The distributions of phonon partition and cuts are shown in Figure 7.11.

- **cQDelay**: QDelay is the time offset of the ionization pulse from the standard template position as determined by the ionization search. The distribution of QDelay vs. phonon energy for *phonon triggers* is shown in Figure 7.12. Recall that the delay is large and negative for such events because the phonon pulse rises so slowly, yielding a Global Trigger time late relative to the pulse. For a fixed phonon energy, the events form a band. The width of the band is determined by two effects. First, for a given phonon energy, there is some jitter in the time of the hardware phonon trigger relative to the pulse peak due to noise on the pulse. Thus, there is jitter in the position of the pulse in the digitized trace, producing noise in the calculated phonon delay and thus in the center of the ionization-search window. Second, once the search window has been established, the search can find any value in the 1.6-ms-wide window if there is no real ionization pulse present above noise. The latter effect primarily determines the band width above 10 keV. The band moves to earlier QDelay as the phonon

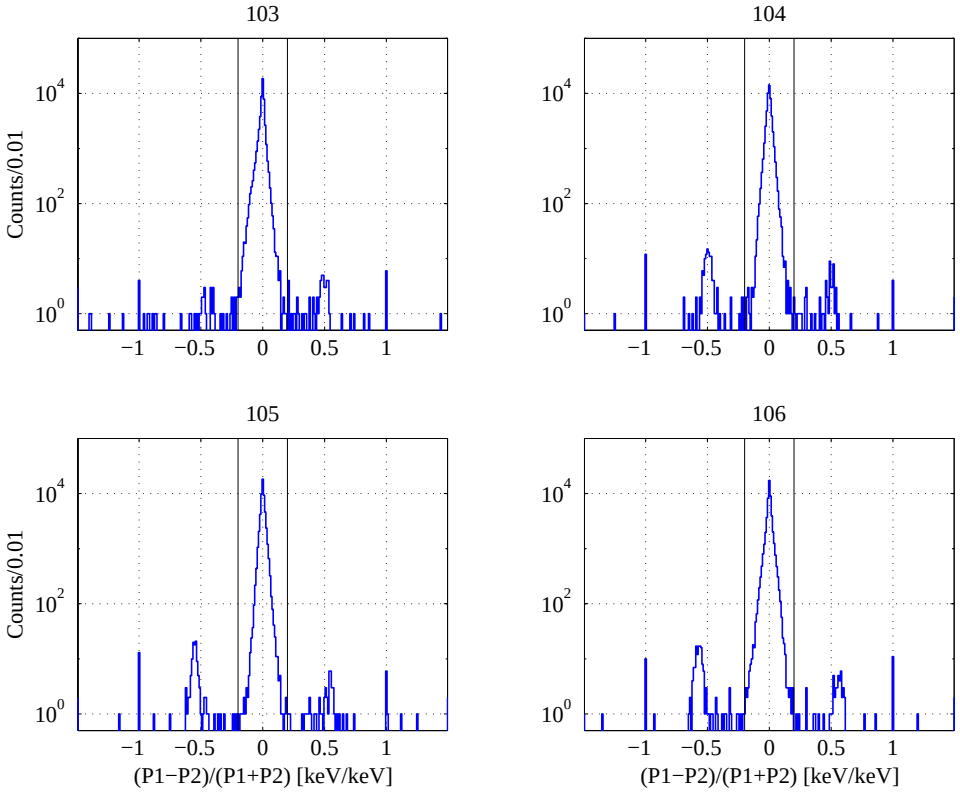

Figure 7.11: Histograms of phonon partition. The lines indicate the acceptance region; events failing the cut are thermistor interactions. cBaseQual and the phonon-χ^2 cuts remove such events due to their unusual time structure (the thermistor with the interaction has a very fast rise time), so these cuts have not been performed before making these histograms. Recall that only 10% of the data are shown.

energy is decreased because a constant-threshold hardware phonon trigger is used: as the phonon energy decreases, the trigger occurs later in the pulse and thus the calculated phonon delay is more negative.

For any phonon energy, events with "real charge" found are seen as an increase in the point density at the center of the band in QDelay. The underlying uniform distribution consists of "noise charge" events, where no true ionization event is found and a noise excursion is picked out by the search. A cut on QDelay is necessary for the following reason. As the phonon energy decreases, the Global Trigger time becomes later relative to the pulse, so the predicted ionization delay becomes more negative and approaches the early edge of the ionization trace for very low phonon energy. The ionization search tends to prefer to saturate at this edge, as is indicated by the collection of events at QDelay ~ -9 ms. The saturation produces a fitted energy that is higher than the typical noise-charge energy. These events are cut by cQDelay to prevent them from leaking through the ionization-threshold cuts (see below). A hard cut at a fixed value of QDelay is made. The length of the ionization pretrigger trace

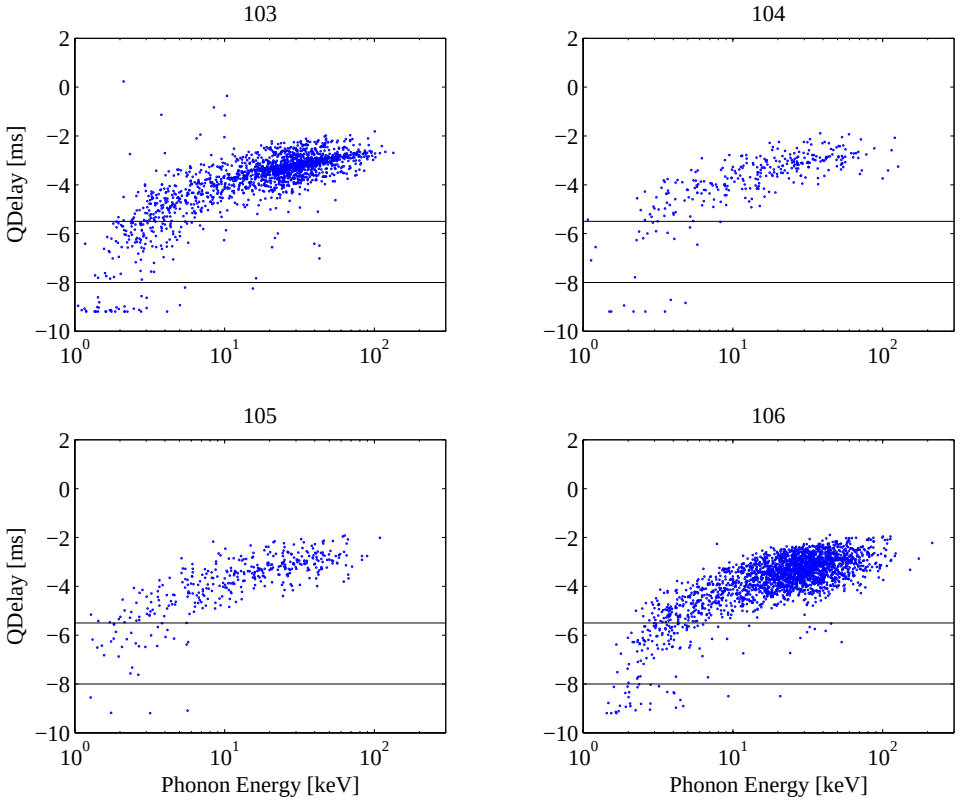

Figure 7.12: QDelay vs. phonon energy. The two lines indicate the position of the cQDelay cut; the cut at -5.5 ms is used for data with 6 ms of pretrigger information and the cut at -8 ms for data with 9 ms of pretrigger information. The BLIP3 and BLIP6 plots contain more points presumably because they shield BLIPs 4 and 5, accepting the brunt of the low-energy-particle flux. The distributions and cuts are described in the text.

was increased from about 6 ms to 9 ms midway through the data set when it was realized it was too short; therefore, two cut values are used, -5.5 ms for the 6 ms data and -8 ms for the 9 ms data. These two cut values are indicated in Figure 7.12.

- **cDataQual**: cDataQual is the logical AND of cBaseQual and the above three cuts, named in this way for obvious reasons.

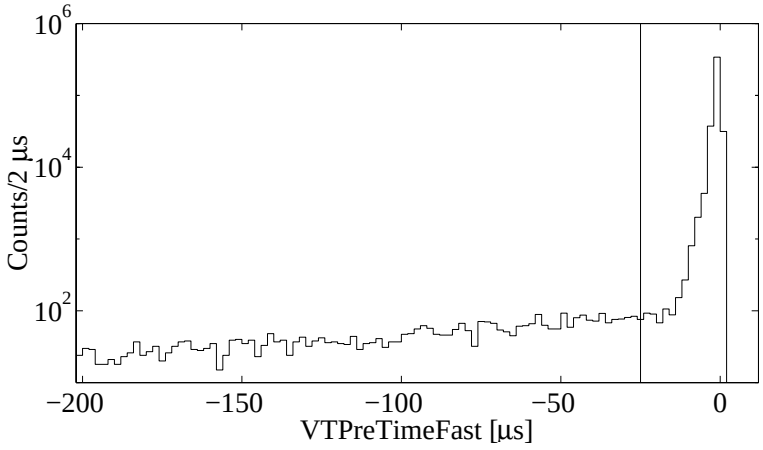

Figure 7.13: Distribution of last veto-trigger times for ionization-trigger events. The exponential accidental distribution has a slope corresponding to $\tau = 150$ μs. The 25-μs coincidence window is indicated.

7.3 "Physics" Cuts

In this section, I describe the cuts that help define, for physics reasons, the data set to be used for a search for WIMP-induced nuclear recoils. The previous cuts are essentially independent of the physical classification of the event — bulk electron recoil, surface electron recoil, or nuclear recoil, veto coincident or anticoincident — except indirectly through energy dependences, while the cuts discussed here select events based on "physics."

7.3.1 Veto-Anticoincidence Cut

In the final dark-matter analysis, a cut is made to remove events coincident with activity in the veto. Because of the high trigger rate of the veto, the rate of accidental coincidences is nonnegligible. For ionization triggers, a veto-coincidence window of 25 μs is used: if there were any veto events in the 25 μs before the detector trigger, the event is considered veto-coincident. This window size was determined by choosing the point where the distribution of last veto-trigger times deviates from an accidental-trigger exponential; see Figure 7.13. The value chosen is, perhaps, too conservative.

The presence of phonon triggers modifies the above. As described in Chapter 6, a search for a pulse in the ionization trace is performed for phonon triggers. If an ionization event is found, its time can be compared to the veto-trigger history. The distribution of nearest veto-trigger times for phonon triggers with an ionization pulse found is shown in Figure 7.14. Based on the points where the distribution deviates from an exponential accidental distribution, the cut is set; a window of ± 25 μs is used. For comparison, the distribution for phonon triggers with only "noise charge" is also shown.

For phonon triggers without ionization, the event time is established by the phonon pulse alone. Because the phonon pulse rises so slowly, the uncertainty on the event time is comparable

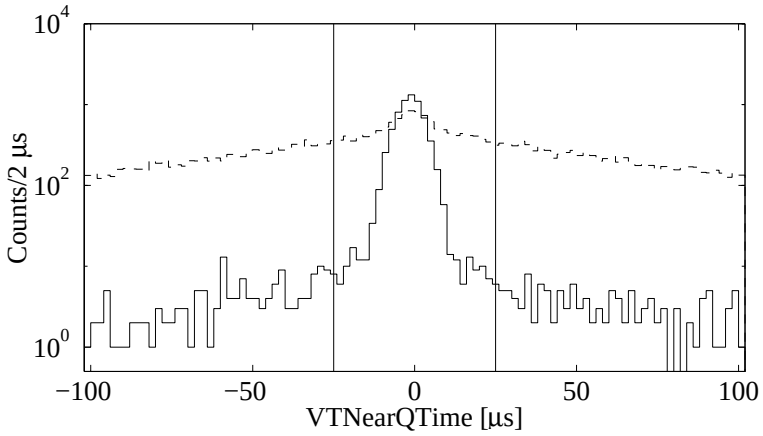

Figure 7.14: Distribution of nearest veto-trigger times for phonon-trigger events, relative to time of *ionization pulse*. Solid: with an ionization pulse above "noise charge" found. Dashed: no ionization pulse above noise charge found. The ±25-μs coincidence window is indicated. The appearance of coincident events in the noise-charge distribution reflects the conservativeness of the threshold dividing "real charge" from noise charge: some noise-charge events actually have real charge.

to the average time between veto events. Figure 7.15 shows the time of nearest veto trigger to the event time established by the phonon pulse alone. The absence of a peak at $t = 0$, even for the set of events shown in Figure 7.14 for which veto-coincidence was established, indicates essentially all phonon triggers with no ionization are accidentally coincident with veto triggers. Because veto coincidence cannot be established, all phonon triggers with only noise charge are discarded. The acceptance for WIMPs is accordingly reduced, as is discussed later.

7.3.2 Inner-Electrode-Containment and Nuclear-Recoil Cuts

To take advantage of the close-packed detector geometry implemented for Run 19, an inner-electrode-containment cut is defined. This consists of selecting events whose outer-electrode ionization signal is consistent with noise: the outer-electrode ionization energy is required to be within ±2 standard deviations of the noise mean, where the noise mean and standard deviation are given by the noise parameters calculated by PipeCleaner (see Chapter 6). Note that, when the detector under study was the triggering detector and phonon triggered, this requires an offset of the ionization signal from zero due to the search algorithm. This cut is named "cQinOnly." This cut is rather extreme, accepting only about 46% of the detector volume, as was shown in Chapter 4.

An alternate definition would be to accept events that have any signal in the inner electrode under the assumption that the shared region, while less shielded than the inner-electrode-contained volume, would still be sufficiently well shielded to increase the exposure without degrading the WIMP sensitivity. Such a definition consists of accepting all events having inner-electrode signal greater than some threshold. This threshold is taken to be 4 standard deviations above the noise mean in order to be sure of having true inner-electrode signal. For purely shared events, events must exceed this inner-electrode threshold, must have signal outside the outer-electrode noise definition

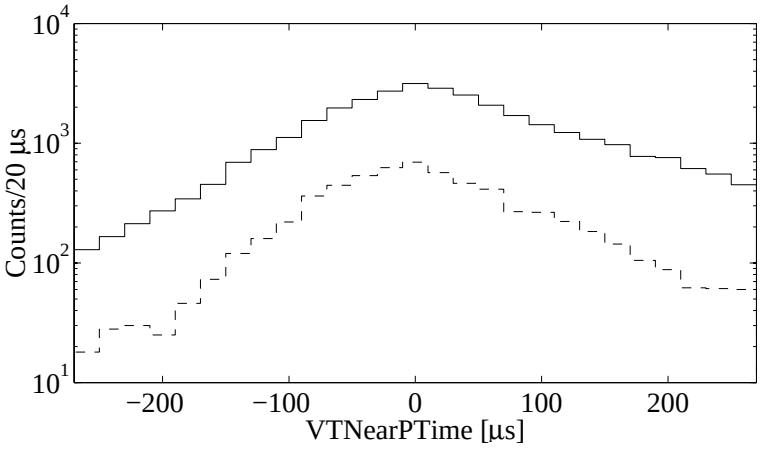

Figure 7.15: Distribution of nearest veto-trigger times for phonon-trigger events, relative to time of *phonon pulse*. Solid: with an ionization pulse above noise charge found (*c.f.* Figure 7.14, which showed the distribution relative to the time of the *ionization pulse*). Dashed: no ionization pulse above noise charge found. The lack of a coincidence peak in either distribution indicates the impossibility of using the phonon pulse alone for veto correlation at the normal veto-trigger rate (*i.e.*, threshold).

above, and must exceed a threshold on the summed-ionization signal. The latter threshold is also defined to be 4 standard deviations above the summed-ionization noise mean. The cut defining such shared events is "cQShare." The logical OR of cQinOnly and cQShare is "cQinShare." The "cQinShare" cut accepts approximately 65% of the detector volume. The three cuts are shown graphically in Figure 7.16.

The electron calibration discussed in Chapter 4 indicates that, for BLIP3, the electron rejection of the portion of the detector selected by the cQShare cut is no worse than that selected by cQinOnly. There are not enough cQShare events in BLIP4 to make a conclusion of this type. Based on the discussion of α-Si/Al-Schottky test devices in Chapter 4, BLIP3 should exhibit worse electron rejection than BLIP4 because BLIP3's negatively biased face is probed by the electron calibration. Naively, then, the cQinShare cut should be used to define the WIMP-search data set.

In practice, the cQinOnly cut is used for historical reasons — the WIMP-search analysis was initially performed with the cQinOnly cut because it is the most conservative choice. The possible use of the cQShare events is discussed in Chapter 8, where it is concluded that inclusion of such events would not degrade the WIMP sensitivity. However, due to time constraints, the analysis has not been redone with the addition of cQShare events. There is nothing, in principle, incorrect about this choice, as long as the efficiencies are appropriately corrected.

The nuclear-recoil cut was explained in Chapter 4. Later in this chapter, I discuss checks done to ensure the stability of the nuclear-recoil efficiency in the presence of both the ionization-nonlinearity correction (see Chapter 6) and the nuclear-recoil band shift (Chapter 4).

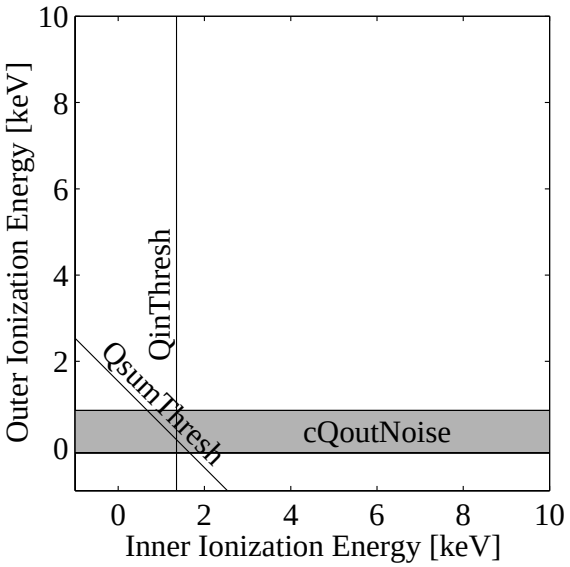

Figure 7.16: Definition of cQinOnly, cQShare, and cQinShare. The vertical line (QinThresh) is the inner-electrode threshold. The diagonal line (QsumThresh) is the summed-ionization threshold. The shaded band is the band in outer-electrode ionization energy accepted as noise (cQoutNoise). Therefore, cQinOnly corresponds to the area to the right of QinThresh and in the cQoutNoise band. cQShare corresponds to the region to the right of QinThresh, above QsumThresh (which has no effect here) and above the cQoutNoise band. cQinShare is the union of the two.

7.3.3 Multiple-Scatter Tagging

A significant fraction of events involve particle interactions in more than one detector. With respect to the WIMP search, these events may be rejected outright without any loss of efficiency — the probability of a single WIMP interacting twice is vanishingly low; if it were not, this experiment would be a lot easier! However, multiple-scatter events are useful for background studies and estimation. For example, the double-scatter-neutron data set is used to estimate the neutron background for statistical subtraction in the WIMP-search analysis.

Some care must be taken in the definition of multiple scatters because there are two minor ambiguities involved. First, since all detector traces are saved, even for detectors without "hardware" triggers, it is in principle possible to obtain a lower "software" threshold for secondary detectors in a multiple-scatter event because the intrinsic noise of the offline fitting algorithms is less than that of the hardware trigger filtering. This is not done because it introduces complications while adding very little interesting data. The effective detector threshold is not set by the phonon-energy threshold, but by the "noise-charge" threshold of the ionization search. Below 10-keV recoil energy, the nuclear-recoil band merges with the noise-charge distribution. If only noise charge is found, no ionization-yield measurement can be made, so neutrons cannot be distinguished from surface events, though an upper limit on the yield can be set that may allow discrimination from bulk electron recoils. Thus, data acquired by a software trigger below 10 keV are not of much use anyways.

Second, one must decide how to classify events where more than one detector triggered but only one detector passes data-quality cuts. In a world of perfectly operating detectors, it would be reasonable to discard these events completely: they certainly should not be included in the single-scatter set, and their loss from the multiple-scatter set can be accounted for by cut efficiencies. However, there are periods during which a subset of the detectors experienced trigger-rate outbursts. The rate of accidental coincidences is greatly increased; classifying such events as multiple scatters would be incorrect and would result in a difficult-to-correct-for loss of single-scatter efficiency. Therefore, such multiple scatters are reclassified as single scatters, conservatively treating the detector failing cuts as dead.

7.4 Cut Efficiencies

In this section, I discuss the efficiencies of the above cuts and the methods used to calculate them. I remind the reader of an elementary analysis point. Any given event may be removed by more than one cut; for example, an event with pretrigger phonon pileup may be removed by both cPbs and cPstd. When calculating cut efficiencies, it is therefore important to account for such correlations. This is trivial to do. One can calculate the efficiency of a set of "overlapping" cuts in one pass, treating them as a single cut. One can also calculate cut efficiencies "relatively:" calculate the efficiency of cut A relative to the entire data set, calculate the efficiency of cut B relative to the event set that passes cut A, etc., and then multiply the efficiencies together. One or the other method is used here, depending on which is more convenient.

7.4.1 Pretrigger-Cut Efficiency

The calculation of pretrigger-cut efficiency is straightforward because the cuts have no dependence on the event characteristics. All the pretrigger cuts are grouped together in calculating this efficiency. The efficiency is calculated in four ways:

- the fraction of "real" (*i.e.*, not random-trigger or pulser) events that pass the cut set;
- the ratio of the sum of the live time of the events passing the cut set to the sum of the live time of all events;
- the fraction of ionization-pulser (IBAPACAP, *c.f.* Chapter 5) events passing the cut set;
- and the ratio of the sum of the live time of the ionization-pulser events passing the cut set to the sum of the live time of all ionization-pulser events.

As was described in Chapter 5, the ionization pulser is completely asynchronous with the DAQ cycle and also appears to introduce no pretrigger noise. It thus is an excellent random-trigger source.

Calculating the efficiency in the above four different ways provides some cross-checks. There *should* be a difference between the fraction of events and the fraction of live time passing these cuts because the DAQ dead time is comparable to the recovery time of the phonon baseline — low live-time events should be cut preferentially. The live-time fraction is the correct efficiency to use because it represents the fraction of *time* the system is capable of detecting a WIMP-scattering event. The ionization-pulser data set provides an independent check.

Table 7.5 displays the efficiencies calculated in these four ways for the ensemble of cuts. The efficiencies for cGoodTime, cPreQual, and cPreQual(1) are all shown; the last cut is the one applied for the final analysis. Showing the other two indicates what fraction of the efficiency loss is due to "bad" data periods (cGoodTime) and what fraction is due to actual pretrigger-trace cuts (cPreQual). The cPreQual(1) requirement results in an additional efficiency loss of about 5% — even when a given detector has a clean pretrigger region, there is a 5% chance that another detector fails pretrigger cuts. This is small enough that is not worthwhile to try to recover such events. The systematic difference between efficiencies calculated using numbers of events and live times agrees with the expectation discussed above. There is good agreement between the data and the pulser set for each separate calculation method and cut (1% to 2% in most cases).

Cut	BLIP3	BLIP4	BLIP5	BLIP6
Live time (live-days)				
cGoodTime	87.6	85.3	91.1	91.3
cPreQual	80.8	78.4	85.6	85.8
cPreQual(1)	76.3	74.2	79.4	79.6
Fraction of data live time				
cGoodTime	0.91	0.88	0.95	0.95
cPreQual	0.84	0.81	0.89	0.89
cPreQual(1)	0.79	0.77	0.82	0.83
Fraction of data events				
cGoodTime	0.89	0.87	0.94	0.94
cPreQual	0.79	0.76	0.86	0.87
cPreQual(1)	0.71	0.69	0.75	0.75
Fraction of pulser live time				
cGoodTime	0.92	0.92	0.95	0.95
cPreQual	0.85	0.85	0.90	0.90
cPreQual(1)	0.81	0.81	0.84	0.84
Fraction of pulser events				
cGoodTime	0.92	0.92	0.96	0.96
cPreQual	0.82	0.81	0.89	0.90
cPreQual(1)	0.74	0.74	0.77	0.78

Table 7.5: Summary of live times and pretrigger-cut efficiencies. The total live time before any cuts is 96.4 live-days. "cGoodTime" denotes the live time and efficiencies of the cGoodTime cut for the given detector. "cPreQual" denotes these for the cPreQual cut for the given detector, which includes that detector's cGoodTime cut. "cPreQual(1)" denotes the efficiency of the cPreQual(1) cut, which requires any detector passing its own cGoodTime cut to also pass its cPreQual cut. cPreQual(1) differs among detectors due to periods when not all detectors pass their own cGoodTime cuts. As noted in the text, the estimates based on numbers of events are in general lower than the estimates by live time, while the data and pulser estimates agree well within each category.

7.4.2 Posttrigger-Cut Efficiency

The efficiency of the cBaseQual cut can be calculated directly from the trigger rate. The dominant efficiency loss is from cPst because phonon pulses are much longer than ionization pulses. cPst cuts events in which a second pulse appears in the phonon trace, but only when the peak height of the second event is larger than that of the first. A lower limit on the efficiency of the cPst cut can be calculated by assuming that the occurence of a second event of any energy causes an event to fail cPst. This is a reasonable estimate at low energies — if the first event is below 100 keV, the second event is likely to be of higher energy simply because most of the trigger rate comes from events above 100 keV. This efficiency is given by the accidental rate for a second event to appear in the 83 ms posttrigger period, which is just

$$1 - \epsilon_{cPst} = 0.083 \ \times \ R \tag{7.1}$$

where R is the single-detector trigger rate calculated from the data. The typical single-detector trigger rate is 0.33 Hz, so $f \approx 0.97$.

A small additional efficiency loss is imposed by the cQPileup cut, which discards any event for which an additional charge trigger occurs between –10 ms and 10 ms relative to the trigger time. The posttrigger portion of this dead time is already accounted for by the 83 ms dead time imposed by cPst. The pretrigger portion should be added, increasing the dead time from 83 ms to 93 ms. At the 1% precision being quoted here, this leaves $\epsilon_{cBaseQual}$ unchanged at 0.97.

The efficiency of cBaseQual is also calculated directly from the data when calculating the multiple-scatter efficiency matrix (see Section 7.4.7) and compares well with the above result.

7.4.3 Pulse-Quality-Cut Efficiency

A nonnegligible efficiency loss is caused by the phonon-χ^2 cut (hereafter, referred to as "the χ^2 cut" when it causes no confusion). It is believed that the events discarded by the cut are basically reasonable but have pulse-shape variations that would lead to an energy misestimate. Therefore, the cut reduces the overall efficiency. For more detailed discussion of the χ^2 distributions and the cut definiton, refer to Chapter 6.

The efficiency of the cut is simply the fraction of events that pass it. Figure 7.10 implies this efficiency is energy dependent. To determine this efficiency function, the data are binned logarithmically in phonon energy and the fraction of events in each bin that pass the cut is calculated. This method assumes that the energy calculated for failing events is approximately correct. The efficiency is seen to be a smooth function of energy, so small errors in the fitted energy have little effect. It is important to remember that the pre-cut data set must be the one that passes all the previous cuts in order to avoid grossly underestimating the efficiency of the χ^2 cut; this efficiency is then multiplied against the previous efficiencies to determine the overall efficiency.

The cut efficiency is certainly not independent of time. This arises mainly because the phonon-χ^2 distributions change sufficiently that it is impossible to set a cut that has the same efficiency at all times. The algorithm used to set the cut attempts to maintain a constant efficiency, but is not very successful. However, given that the trigger rates are fairly stable, the efficiency calculated from the data set as a whole should correctly incorporate the varying efficiency. For example, a period with a low cut efficiency is weighted according to the total number of events in the set before the χ^2 cut, which is proportional to the live time of the period, providing the correct

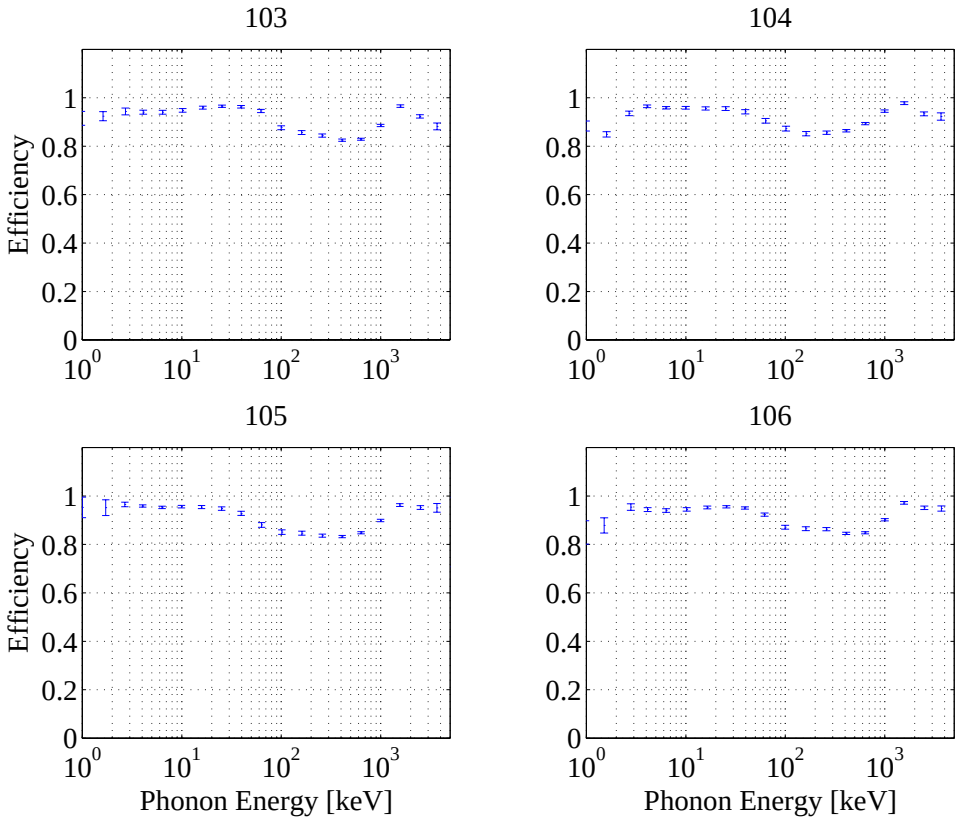

Figure 7.17: Efficiency of phonon-χ^2 cut vs. phonon energy. See also Figure 7.10.

weighting. The prior cuts remove extraordinary periods, so this is a valid procedure. Furthermore, the assumption is conservative in that it underestimates the efficiency. For example, if a trigger outburst is left in the data set from which the efficiency is calculated, then it is overweighted because it has too many events. The efficiency for such a period is lower than is typical because of the higher noise. Thus, the mean efficiency is decreased by such a period.

The efficiency of the phonon-χ^2 cut as a function of phonon energy is shown in Figure 7.17. The efficiency has artificial structure that arises mainly from the fact that, at around 200 to 300 keV, the χ^2 distribution broadens and exhibits a tail, as seen in Figures 6.12 and 7.10. While the shape of the efficiency function may appear strange, it is correct — a more stringent cut is made at higher energy, giving a lower efficiency.

As has been mentioned before, the ionization-χ^2 cut is very loose, with essentially 100% efficiency. However, as also previously noted, its efficiency for low-energy events in which another detector had a saturating ionization pulse is significantly lower. It was argued earlier that the data sets of primary interest for this analysis — veto-coincident and veto-anticoincident single-scatter and double-scatter nuclear recoils — do not suffer because they never or rarely contain a saturated detector. It is critical to realize that other data sets — for example, multiple-scatter photons or photon-electron events — are affected by this loss, and the 100%-efficiency assumption is completely

invalid. The efficiency function *must* be calculated independently for such analyses.

The cppart cut on phonon partition clearly is very liberal (see Figure 7.11) and thus has essentially unity efficiency.

It is difficult to calculate the efficiency of the cQDelay cut in a rigorous manner because the true ionization energy of an event discarded by cQDelay is unknown. An event can fail cQDelay for two reasons. If the event has a large phonon energy ($\gtrsim$ 10 keV), yet fails cQDelay because no "real-charge" pulse can be found, then the event is likely a dead-layer event — bulk nuclear and electron recoils at such energies produce enough ionization to be above noise charge. Because nuclear recoils are in general uniformly distributed through the detector, cutting such dead-layer events only reduces the efficiency for nuclear recoils by the ratio of the dead-layer volume to the total detector volume, which is less than 1% assuming a 30 μm dead layer. The other reason an event can fail cQDelay is that it has such low phonon energy that the ionization-search window overlaps the early edge of the ionization trace. This failure mode is independent of the event position in the detector and, in fact, is dependent mainly on the *phonon* energy. As is seen in Figure 7.12, the width of the QDelay band widens considerably below $\sim$ 5-keV phonon energy. Since the ionization-search-window width is fixed at 1.6 ms, this additional widening is due to increased noise in the *phonon*-delay estimation. Figure 7.18 shows the fraction of events passing cQDelay as a function of phonon energy. These plots indicate that the latter effect is indeed the dominant loss of efficiency due to cQDelay: the efficiency above 10 keV is almost unity. Therefore, it is justified to include these efficiency functions in the overall efficiency calculation in the same way that the hardware phonon-trigger efficiency is included; this is discussed in the next section. One caveat should be made: the April, 1999, neutron-calibration data do not satisfy the above criterion that the efficiency above 10 keV be unity; thus, the inclusion of the cQDelay efficiency in the above way is invalid. This point may partially explain the worse agreement of the data and simulation for this calibration set as compared to the September, 1999, neutron calibration and the muon-coincident neutron data (see Section 7.5).

7.4.4 Trigger Efficiency

The hardware-trigger efficiency can be measured using "off-detector" triggers. An "off-detector" trigger for detector A is an event in which any of the other detectors was the first to trigger. For all events, every detector channel is digitized and trace fits done. Detector A's trigger efficiency at energy E is given by the fraction of all detector-A off-detector events having energy E for which a detector-A hardware trigger is found in the posttrigger history. This is done separately for the phonon trigger as a function of phonon energy and for the ionization trigger as a function of ionization energy. This calculation is done on the event set passing all the prior cuts to ensure good energy estimates for detector A. The calculation is done for the whole data set at once — again, presumably periods of different trigger efficiency are automatically weighted by the number of events taken with a given trigger-efficiency curve. Figure 7.19 shows the ionization-trigger efficiency for all events and for inner-electrode-contained events. The horizontal axis is the summed or inner-electrode ionization energy, respectively. The trigger efficiencies are the same, as they should be because the hardware trigger cannot distinguish the two event classes (unless the pulse-height-to-energy conversions of the two channels differ appreciably; they do not). Figure 7.20 shows the phonon-trigger efficiency.

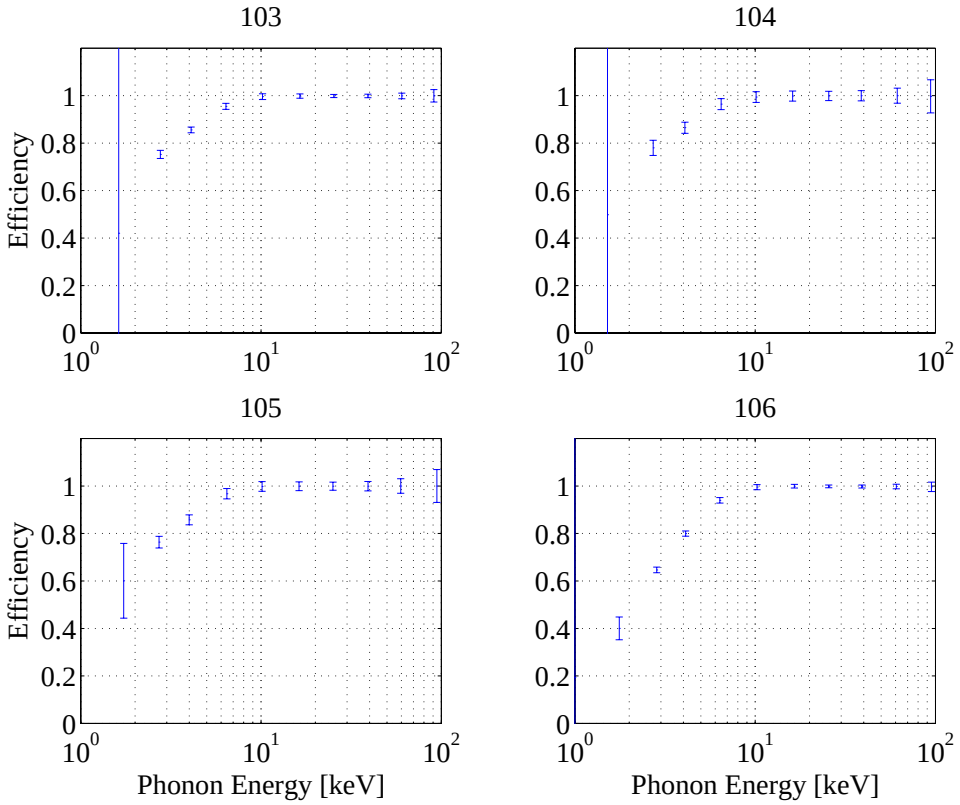

Figure 7.18: Efficiency of the cQDelay cut as a function of phonon energy.

7.4.5 Veto-Anticoincidence-Cut Efficiency

The efficiency of the veto-anticoincidence cut is determined by the rate of accidental detector-veto coincidences. The total veto-trigger rate about 6 kHz. For ionization triggers, the veto-coincidence window is 25 μs. Therefore, there is a 6 kHz $\times$ 25 μs $= 0.15$ probability that an accidental coincidence occurs, yielding an efficiency of 0.85. For phonon triggers with ionization found, the window is ± 25 μs, giving an efficiency of 0.70. For phonon triggers with only noise charge found, the efficiency vanishes because all such events are discarded.

7.4.6 Overall Nuclear-Recoil Efficiency

I discussed in detail the grid method used for calculation of the nuclear-recoil efficiency in Chapter 4. The "raw" grid, $\epsilon_{NR}^{raw}(E_R, E_Q)$, is defined as the grid including only the Gaussian distributions of nuclear recoils about the nuclear-recoil line. By definition, it satisfies the normalization condition $\int_{-\infty}^{\infty} dE_Q\, \epsilon_{NR}^{raw}(E_R, E_Q) = 1$ for all E_R. The nuclear-recoil efficiency grid, $\epsilon_{NR}(E_R, E_Q)$, incorporates the definition of the nuclear-recoil-acceptance region and the ionization-search threshold:

$$\epsilon_{NR}(E_R, E_Q) = \epsilon_{NR}^{raw}(E_R, E_Q)\, \theta(E_Q - E_Q^{thr})\, \theta(E_Q - E_Q^{lo}(E_R))\, \theta(E_Q^{hi}(E_R) - E_Q) \qquad (7.2)$$

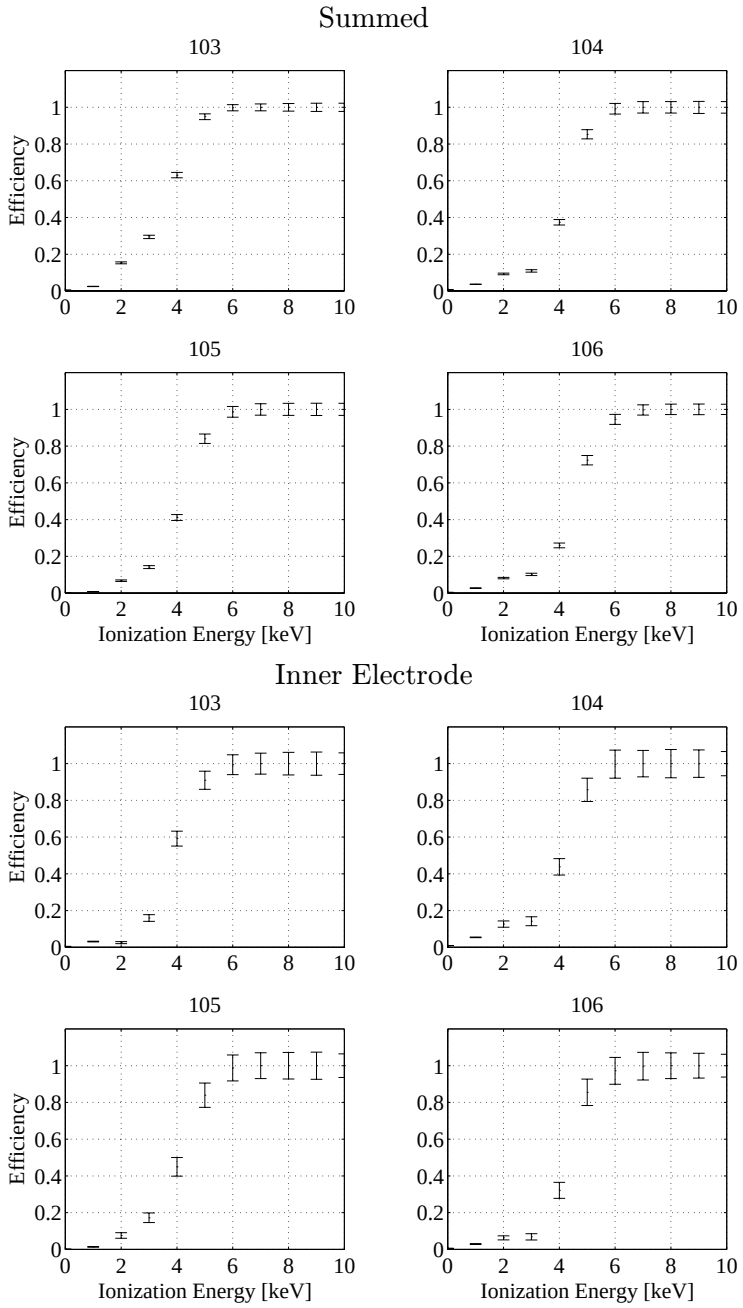

Figure 7.19: Efficiency of hardware ionization trigger vs. ionization energy. Separate plots are shown of all events and inner-electrode-contained events.

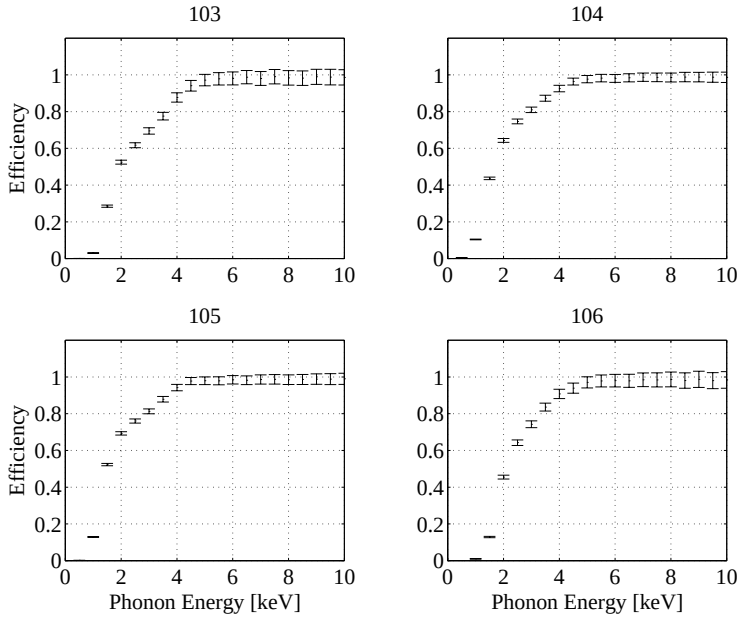

Figure 7.20: Efficiency of hardware phonon trigger vs. phonon energy.

where $E_Q^{lo}(E_R)$ and $E_Q^{hi}(E_R)$ define the nuclear-recoil-acceptance region and E_Q^{thr} is the ionization-search threshold (a constant). This second grid thus accounts for the efficiency lost by these two cuts. The nuclear-recoil efficiency at E_R is given by integrating $\epsilon_{NR}(E_R, E_Q)$ over E_Q.

This calculation does not include the effects of other cuts, including those with energy-dependent efficiencies, such as the phonon-χ^2 cut. The above procedure is generalized by calculating the efficiency of each cut on the grid of E_Q vs. E_R. For cuts that depend on phonon energy (E_P) rather than E_Q or E_R separately, the phonon energy at each point on the grid is calculated by inverting the Luke correction:

$$E_P = E_R + \frac{eV_b}{3.0\,\mathrm{eV}} E_Q \tag{7.3}$$

and the efficiency at that phonon energy applied. For example, the phonon-χ^2-cut efficiency on the grid is calculated this way, yielding an efficiency grid $\epsilon^P_{\chi^2}(E_R, E_Q)$. The trigger efficiency takes the form

$$\begin{aligned} \epsilon_{tr}(E_R, E_Q) &= 1 - \left[1 - \epsilon^Q_{tr}(E_Q)\right]\left[1 - \epsilon^P_{tr}(E_P)\right] \\ &= \epsilon^Q_{tr}(E_Q) + \epsilon^P_{tr}(E_P) - \epsilon^Q_{tr}(E_Q)\epsilon^P_{tr}(E_P) \end{aligned} \tag{7.4}$$

This form may not be obvious. Consider the first line. In the second term, the first factor is the probability that the ionization trigger fails to fire and the second factor is the probability that the phonon trigger fails to fire. Both must fail in order for the event to be lost, so the second term is the combined *inefficiency* of the hardware triggers. As argued above, the efficiency of the cQDelay cut can be absorbed into $\epsilon^P_{tr}(E_P)$. When applying a veto-anticoincidence cut, the veto-anticoincidence-cut and trigger efficiencies must be combined into one grid because of the dependence of the former

on the latter; the resulting efficiency grid takes the form:

$$\epsilon_{anti}(E_R, E_Q) = \epsilon^Q_{tr}(E_Q)\epsilon^Q_{anti} + \left[1 - \epsilon^Q_{tr}(E_Q)\right] \epsilon^P_{tr}(E_P)\epsilon^P_{anti}\theta(E_Q - E^Q_{thr}) \tag{7.5}$$

where $\epsilon^Q_{tr}(E_Q)$ is the ionization-trigger efficiency, $\epsilon^P_{tr}(E_P)$ is the phonon-trigger efficiency and $\epsilon^Q_{anti} = 0.85$ and $\epsilon^P_{anti} = 0.70$ are the veto-anticoincidence-cut efficiencies for ionization triggers and phonon triggers with charge found, respectively. The $\theta(E_Q - E^Q_{thr})$ enforces the requirement that the ionization found be above the ionization-search threshold E^Q_{thr}. The first term is the efficiency for ionization triggers. The second term accounts for the lost efficiency gained back by phonon triggers for which ionization is found.

The overall energy-dependent efficiency is calculated by multiplying all grids of the above form together, point-by-point, and integrating over ionization, yielding

$$\epsilon^{tot}_{anti}(E_R) = \epsilon_{pre}\,\epsilon_{post} \int_{-\infty}^{\infty} dE_Q\, \epsilon^P_{\chi^2}(E_R, E_Q)\, \epsilon_{anti}(E_R, E_Q)\, \epsilon_{NR}(E_R, E_Q) \tag{7.6}$$

where all E_P have been replaced using Equation 7.3, ϵ_{pre} is the pretrigger-cut efficiency, ϵ_{post} is the posttrigger-cut efficiency (both energy independent), $\epsilon^P_{\chi^2}(E_R, E_Q)$ is the phonon-χ^2-cut efficiency grid, $\epsilon_{NR}(E_R, E_Q)$ is the nuclear-recoil efficiency grid, and $\epsilon_{anti}(E_R, E_Q)$ from Equation 7.5 is the veto-anticoincidence efficiency grid.

Figure 7.21 shows the overall *exposure* as a function of recoil energy. The exposure is given by multiplying the overall efficiency function by the raw exposure, 96.4 live-days or 15.9 kg d. The softening of the veto-anticoincident efficiency function between 10 and 20 keV occurs because of the gradual change-over between phonon triggers, with veto efficiency 0.70, to ionization triggers, with veto efficiency 0.85. The high-energy ($>$ 20-keV) shape is entirely due to the energy dependence of the phonon-χ^2 cut. The change in the position of the nuclear-recoil band at the April 3 power outage has been accounted for by weighting the efficiencies calculated during the two epochs by their respective livetimes. It must be emphasized that the efficiency calculated here is only valid for nuclear recoils; for example, the efficiency for bulk electron recoils is different because they lie at a different position in the E_Q vs. E_R grid.

7.4.7 Multiple-Scatter Efficiency

I have so far concentrated on individual detector efficiencies. However, double-scatter neutron events play an important role in the final analysis. Naively, one calculates the efficiency for detection of a multiple-scatter event simply from the product of the individual detector efficiencies. This is not accurate because it is important to take into account correlations in the data-quality-cut efficiencies for different detectors. It is only the energy-independent data-quality-cut efficiencies that need to be considered in this respect. The χ^2-cut, cQDelay, and trigger efficiencies exhibit no correlations because their energy dependences are dominated by the individual detector noise and pulse-shape characteristics. The nuclear-recoil-cut efficiencies are also uncorrelated, aside from correlations introduced by real physics; *e.g.*, multiple scattering of a neutron. An example case of how data-quality cuts introduce correlations is pretrigger pileup. When detector A has pretrigger pileup, its neighbor has a higher-than-random chance of also having pretrigger pileup because the neighbor may be hit by the same particle or by particles produced by the same incident muon or

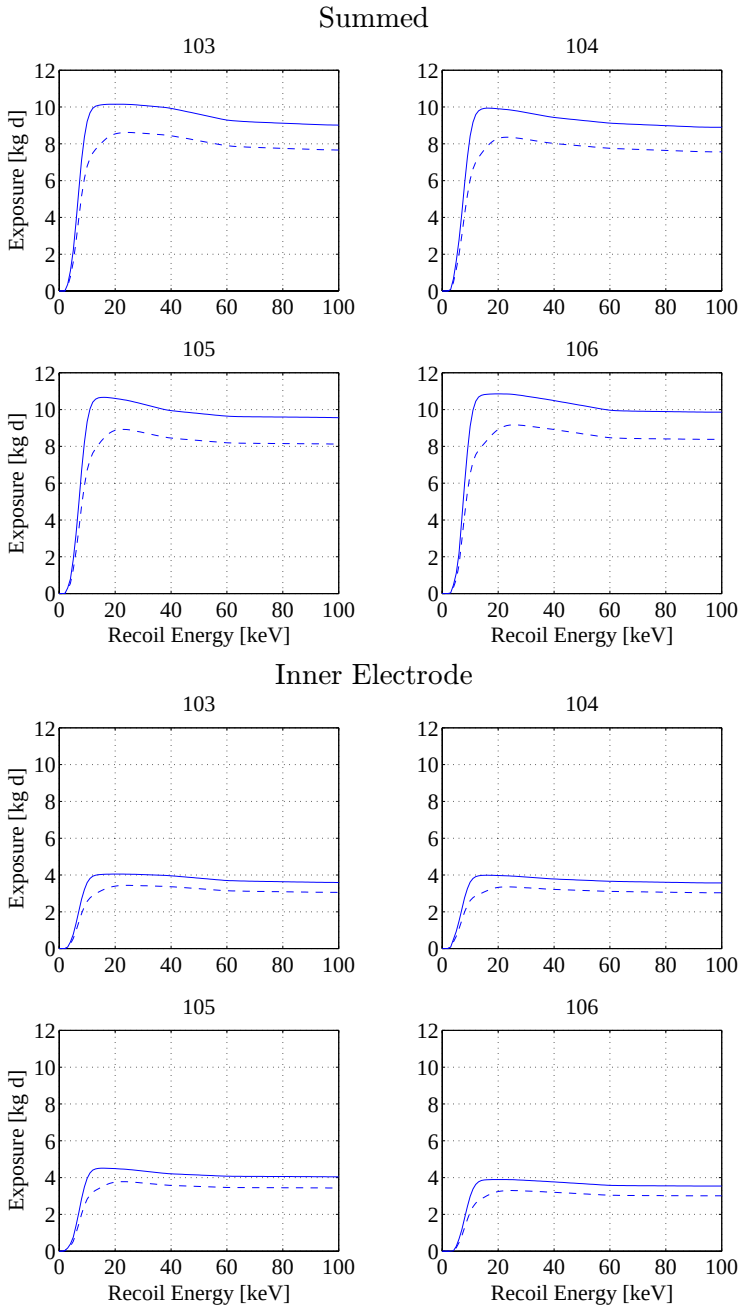

Figure 7.21: Overall nuclear-recoil *exposure* as a function of recoil energy. Top: no inner-electrode-containment cut. Bottom: with inner-electrode-containment cut. Solid: veto-coincident. Dashed: veto-anticoincident. The raw exposure for each detector is 96.4 live-days, or 15.9 kg d. Thus, the overall veto-anticoincident cut efficiency is about 50% when no inner-electrode-containment cut is made, dropping to about 20% when this cut is applied.

high-energy photon. Therefore, it is necessary to calculate a matrix of various joint efficiencies, which, for three detectors (four would not fit on the page!), has the form:

$$\mathbf{M} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \neg B3 & B3 & 0 & 0 & 0 & 0 & 0 & 0 \\ \neg B4 & 0 & B4 & 0 & 0 & 0 & 0 & 0 \\ \neg B3 \,\&\neg B4 & B3 \,\&\neg B4 & \neg B3 \,\&B4 & B3 \,\&B4 & 0 & 0 & 0 & 0 \\ \neg B5 & 0 & 0 & 0 & B5 & 0 & 0 & 0 \\ \neg B3 \,\&\neg B5 & B3 \,\&\neg B5 & 0 & 0 & \neg B3 \,\&B5 & B3 \,\&B5 & 0 & 0 \\ \neg B3 \,\&\neg B4 \,\&\neg B5 & B3 \,\&\neg B4 \,\&\neg B5 & \neg B3 \,\&B4 \,\&\neg B5 & B3 \,\&B4 \,\&\neg B5 & \neg B3 \,\&\neg B4 \,\&B5 & B3 \,\&\neg B4 \,\&B5 & \neg B3 \,\&B4 \,\&B5 & B3 \,\&B4 \,\&B5 \end{pmatrix}$$

where the entries are interpreted as follows: B3&B5 implies BLIP3 and BLIP5 both pass all their cuts and ¬B3&B5 indicates that BLIP3 fails its cuts while BLIP5 passes its own. If a detector is not explicitly listed in an entry, then its state is not considered. The matrix is read as follows. Each row corresponds to a particular combination of detectors. Entries in that row contain all possible configurations for that combination — all detectors passing cuts, some subset passing cuts, or none passing cuts. There is therefore a unitarity condition: each row must sum to 1. The matrix must be calculated from an event set in which no detectors are saturated; cross-talk from a detector whose state is nominally irrelevant can affect the probability of the *other* detectors passing their cuts. As discussed earlier, a cut requiring no saturated detectors is reasonable. This matrix is used in the obvious way for correcting the simulation for the efficiency of detecting multiple-scatter neutrons; the simulated and observed multiple-scatter-neutron spectra are discussed in the next section.

7.5 Checks on Accuracy and Stability of Cut Efficiencies

It is critical to check that the efficiency calculated by the above method is correct and that it is stable over time. This is especially important given the change in nuclear-recoil bands at April 3.

7.5.1 Accuracy

The absolute accuracy of the efficiency calculation can be checked against the simulation of the neutron calibration. The details of the simulation code are discussed in Chapter 3. A ^{252}Cf-fission neutron source is placed roughly at the center of the top face of the veto. The spectrum of neutrons emitted by the source is shown in Figure 7.22. Because these neutrons have such low energies, the top layers of polyethylene inside the shield are removed to permit the neutrons to penetrate to the cryostat. The data rate is also higher than normal to ensure the data set is dominated by neutrons. Otherwise, the data-taking conditions are as usual. The source activity is known, so the absolute normalization of the spectrum is fully determined. The overall cut efficiency determined by the method previously discussed is applied to this spectrum. It is necessary to recalculate the efficiencies from the neutron-calibration data set because the amount of pileup is significantly increased. A veto-anticoincidence cut is not applied to the neutron-calibration data

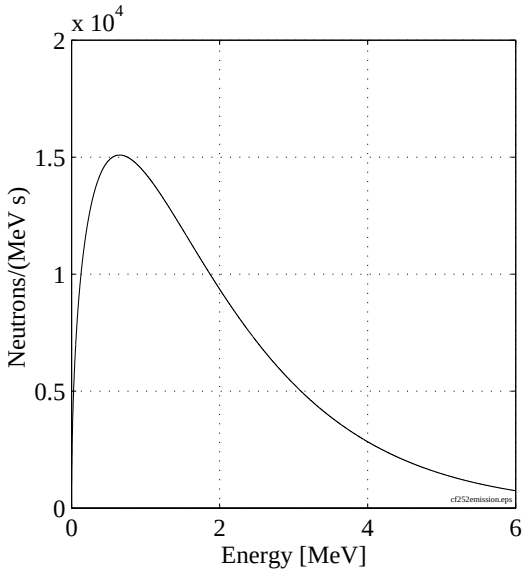

Figure 7.22: Energy spectrum of neutrons emitted by ^{252}Cf calibration source. Figure taken from [8].

because a significant fraction of neutrons trigger the veto as they pass through it. So, while the efficiencies themselves differ from those of the WIMP-search data set, the method is essentially the same.

Additionally, an artificial energy window of 10 to 100 keV is imposed. It is observed in the veto-anticoincident data set that leakage of low-energy electrons into the nuclear-recoil band becomes significant below 10 keV. This observation is corroborated by data taken with the α-Si/Al-Schottky test device, Figure 4.20. The neutron-calibration multiple-scatter-neutron data set also shows an excess rate at low energies — the spectra rise too steeply at low energy. (Note that this has little effect on the empirical nuclear-recoil-efficiency estimate because multiple scatters make up a small (10-20%) fraction of the total neutron rate.) Similar effects are seen in the muon-coincident double-scatter-neutron data. Application of a 10-keV threshold for both detectors in multiple-scatter events yields physically reasonable spectra. As mentioned in Section 4.5.1, a 100-keV upper limit is applied because the neutron-calibration data become so sparse that the nuclear-recoil band is not accurately determined above this energy. It is straightforward to apply these energy cuts to the simulation.

The observed and simulated spectra for the two neutron-calibration data sets are shown in Figure 7.23. The cQinOnly cut has been applied. There are no free parameters in the comparison; the simulation normalization is set by the source activity and the efficiencies calculated from the data. For both calibrations, the simulated spectra match the data to about 20 to 30% accuracy. The spectral shape, while not exactly right, is fairly close. Moreover, the significant change in the observed spectra at low energies, due to changes in cut efficiencies between the two calibrations, is reproduced by the simulation after application of the cut efficiencies. The discrepancy between the simulated and observed spectra can be used to estimate a systematic error for the cut-efficiency

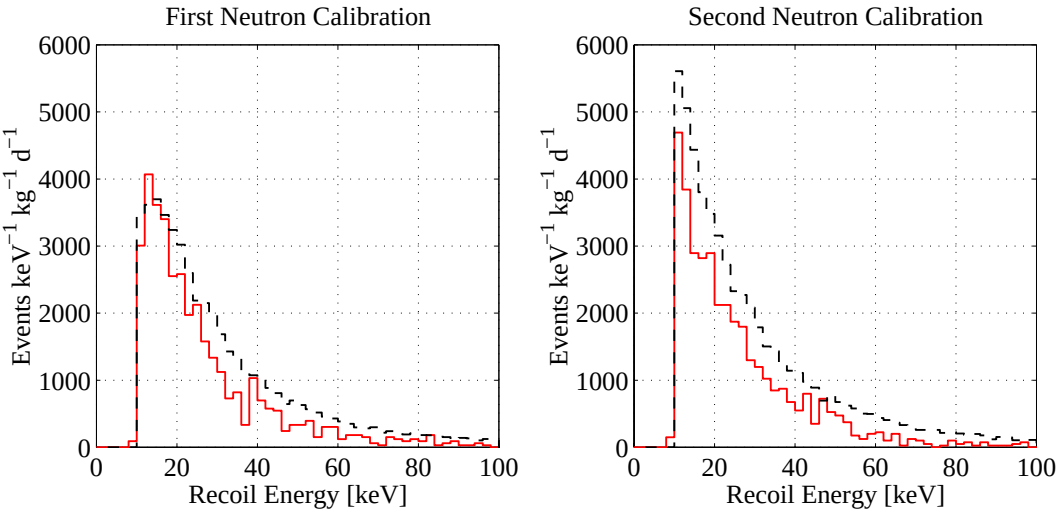

Figure 7.23: Observed and simulated neutron-calibration recoil-energy spectra, coadded over all four detectors, with no free parameters. Light solid line: observed spectra. Dashed line: simulated, with efficiency corrections applied. The spectra are for all cQinOnly nuclear recoils.

calculation by finding the scaling A that yields the best fit of the simulated spectral shape to the data and taking $|A-1|$ as the fractional systematic error. These numbers are listed in Table 7.6.

A second check of the accuracy of the efficiency calculation is provided by multiple-scatter neutrons. As discussed in the previous section, calculation of the efficiency for such events is nontrivial due to correlations in the cuts for detector combinations. The correlation matrix defined earlier is used in establishing this efficiency. The simulated and observed multiple-scatter-neutron spectra are shown in Figure 7.24. As noted above, all recoils of a multiple-scatter are required to be between 10 and 100 keV to appear in this histogram. The histogram is filled for each recoil of a multiple-scatter event; *e.g.*, a double scatter adds two entries to the histogram. (It should be noted that the number of multiple scatters with more than two recoils is negligible, though nowhere in this section is it assumed that there are no triple or quadruple scatters.) Plots are shown for two ionization-partition cuts. The first set of plots requires all recoils of a multiple scatter to pass cQinOnly ("Qin/Qin"). The second set only requires one cQinOnly recoil ("Qin/Qsum"). In all cases, the agreement of the data and simulation is good. The best-fit scaling factors are shown in Table 7.6. The different levels of agreement of the all-event and multiple-scatter spectra with the simulation are indicative of the systematic error in the efficiency calculation.

The accuracy of the nuclear-recoil efficiency can also be checked by comparing the simulated and observed spectra for muon-coincident neutrons. As discussed in detail in Chapter 3, these neutrons are produced by muons passing through the veto that interact in the copper cans of the cryostat or the internal lead shield. This data set offers the advantage that it is acquired at the same time as the WIMP-search data set, and thus the efficiencies are exactly the same, with the exception that no veto-anticoincidence cut is applied. Figure 7.25 shows the simulated and observed muon-coincident-neutron spectra for the same energy cuts and event categories as shown for the neutron-calibration data. There is good agreement. Best-fit scaling factors A are shown in

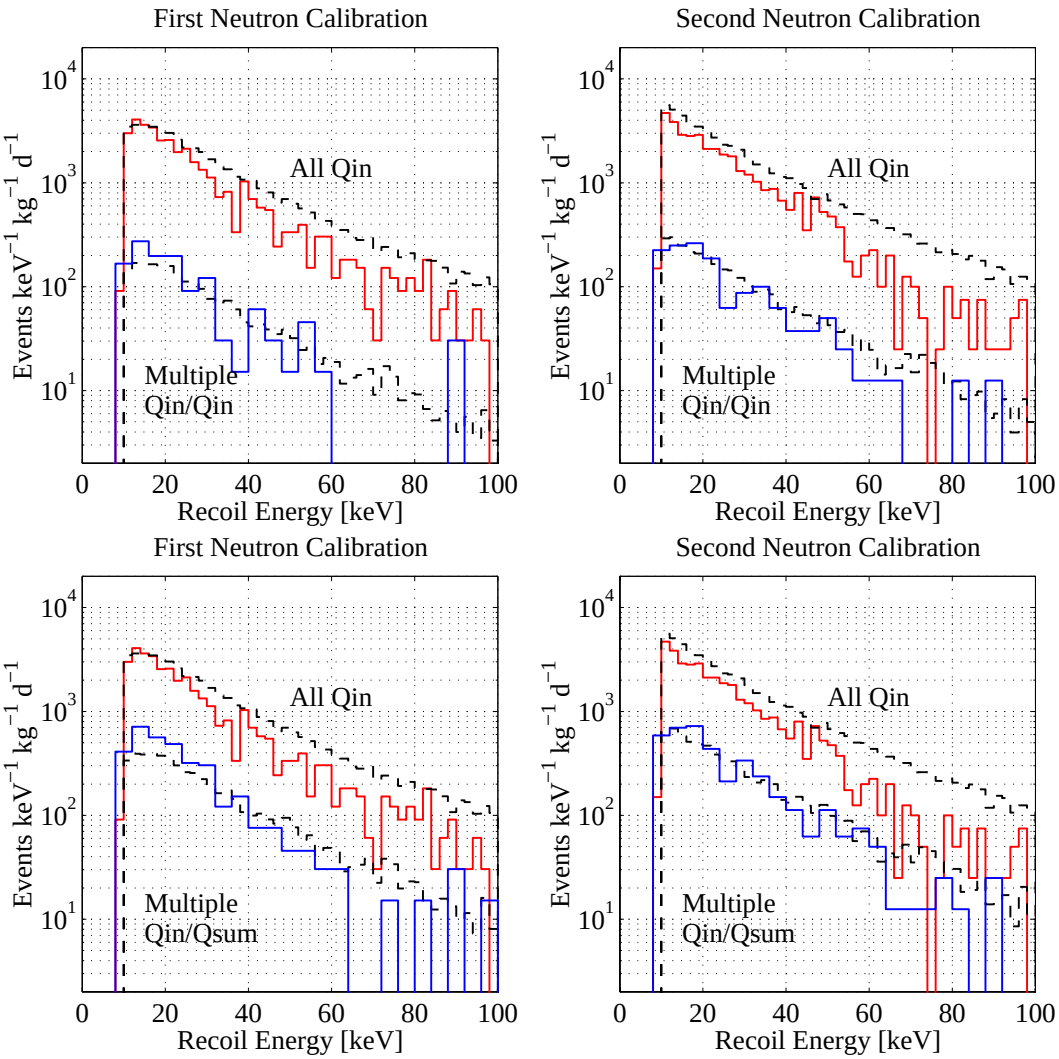

Figure 7.24: Observed and simulated multiple-scatter neutron-calibration spectra, coadded over all four detectors, with no free parameters. Solid: observed spectra. Dashed: simulated. Upper left: first calibration, all recoils of the multiple scatter are required to pass cQinOnly ("Qin/Qin"). Lower left: first calibration, only one recoil of the multiple scatter is required to pass cQinOnly ("Qin/Qsum"). Right: similar plots, second neutron calibration. The upper spectra in the plots are the same as were shown (with linear vertical scales) in Figure 7.23.

Event Set	First Neutron Calibration	Second Neutron Calibration	Muon-Coincident Neutrons
all Qin NRs	0.86	0.74	0.97
multiple Qin/Qin NRs	1.31	0.98	0.77
multiple Qin/Qsum NRs	1.51	1.18	1.33

Table 7.6: Scaling factors A that must be applied to the simulation to achieve agreement with the data. $|A-1|$ is a reasonable estimate of the fractional systematic error of the efficiency calculation.

	standard deviation, NRs/live-day		
Event Set	observed	expected	excess
single Qin NRs	2.2	2.1	0.7
multiple Qin/Qin NRs	0.7	0.6	0.4
multiple Qin/Qsum NRs	1.0	0.9	0.4

Table 7.7: Observed and expected dispersion of muon-coincident-neutron rates, in nuclear recoils/live-day. The excess dispersion is calculated by the quadrature difference of the observed and expected dispersions.

Table 7.6.

7.5.2 Stability

The stability of the nuclear-recoil acceptance over time is checked by Figure 7.26, which shows the muon-coincident-neutron rates, coadded over the four detectors, as a function of time in blocks of approximately 5 live-days. The rates of three different types of events are shown separately: cQinOnly single scatters, "Qin/Qin" multiple scatters, and "Qin/Qsum" multiple scatters. The standard 10-to-100-keV cut is made. While the stability is not as good as one would expect from random statistical fluctuations, the variations are not alarmingly large. There are two particularly critical points:, the April 3 power outage and the refrigerator warmup/cooldown cycle in June, which occurred at roughly 29 and 65 raw live-days, respectively. The rates show no statistically significant change at either of these points. Numerically, the dispersion of the data can be compared to the expected dispersion based on the data's own error bars. Each point corresponds to about 5 live-days, so the typical number of counts N yielding each point is given by the product of the mean rate and 5 live-days. The statistical uncertainty on N counts is $\sqrt{N}$, so the statistical uncertainty on each point is $\sqrt{N}/(5 \text{ live-days})$. This is the expected standard deviation of the data. The observed and expectated standard deviations are shown in Table 7.7 and the inferred excess standard deviation is calculated. The excess dispersions indicate what is already apparent: the rates vary more than are expected, but not egregiously. Furthermore, some excess dispersion is expected because variations in noise, amplifier gain, etc. have the effect of slightly modulating the acceptance of the nuclear-recoil cut.

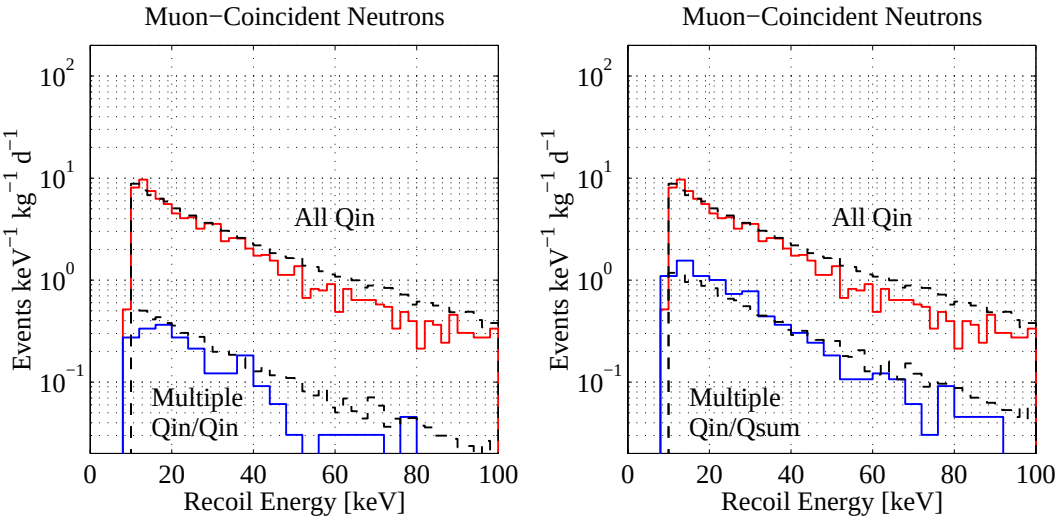

Figure 7.25: Muon-coincident-neutron recoil-energy spectra, coadded over the four detectors, for the entire run. Solid: observed spectra. Dashed: simulated. The top spectra in each plot are for all events; the bottom spectra are multiple-scatter events. Left: "Qin/Qin" multiple scatters. Right: "Qin/Qsum" multiple scatters.

7.5.3 Implications for WIMP-Search Analysis

It is seen that the efficiencies are accurate and stable at the 30% level. This is not as good as one might hope for. However, it is important to realize that, because the dark-matter analysis discussed in the next chapter is severely statistics limited — only 13 single-scatter and 4 multiple-scatter nuclear-recoil candidates are seen — this level of accuracy and stability is sufficient. For future data sets, as more nuclear-recoil candidates are accumulated and the statistical uncertainty associated with the data becomes small, the accuracy of the efficiency calculation (or, possibly, the simulation) will have to be improved.

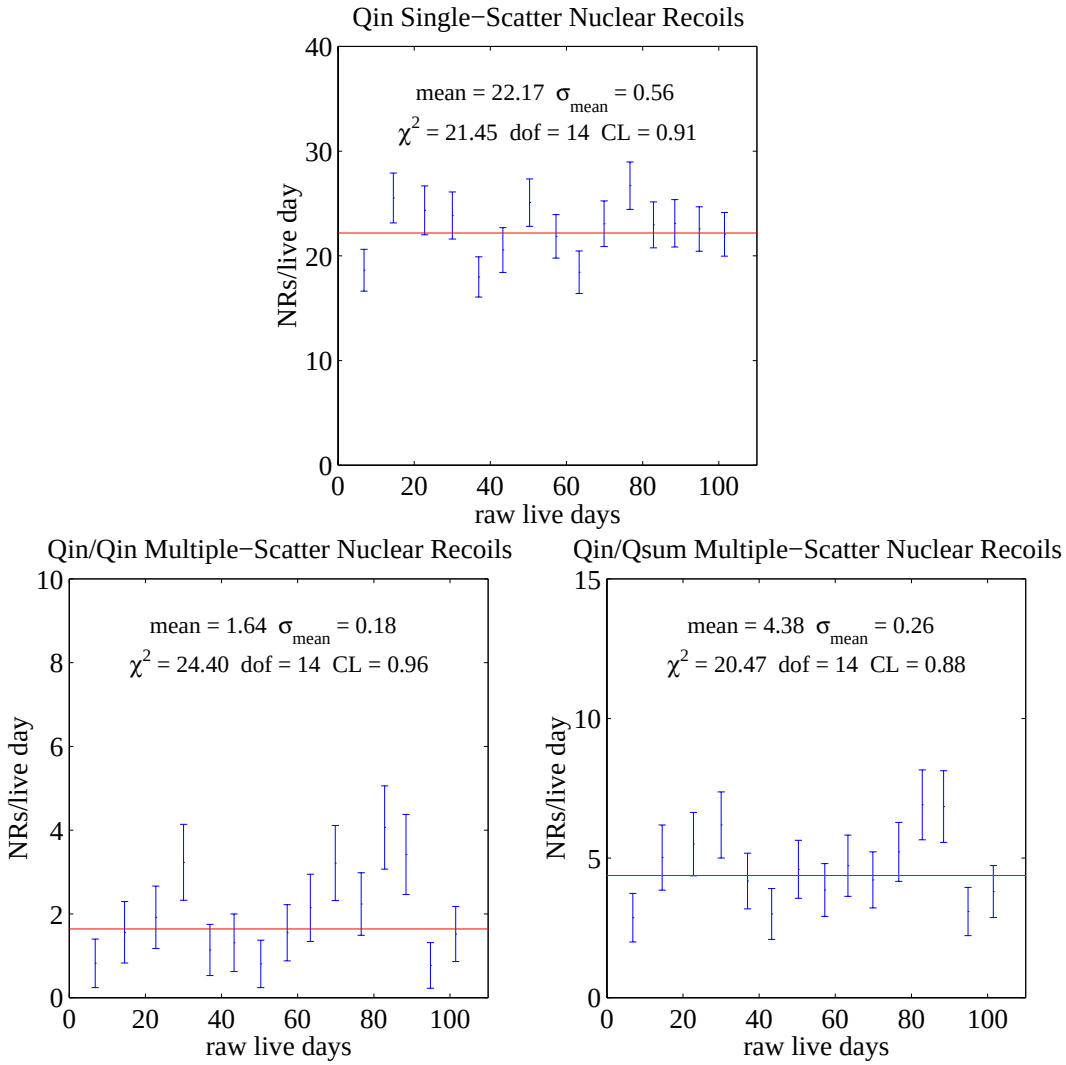

Figure 7.26: Muon-coincident-neutron rates vs. time, coadded over detectors. Top plot: single scatters. Bottom left: Qin/Qin multiple scatters. Bottom right: Qin/Qsum multiple scatters. Each point corresponds to approximately 5 live-days. The error bars are only statistical. The mean (calculated from the data) and the χ^2 of the data relative to the mean are shown.

Chapter 8

Background-Particle Rates and Constraints on the WIMP-Nucleon Cross Section

8.1 Introduction

In this chapter, I finally cut to the chase — the veto-anticoincident nuclear-recoil data set and its implications for the WIMP-nucleon cross section. Rates for photons, electrons, and nuclear-recoil candidates are presented. Using the double-scatter data set and data from CDMS Run 18, it is argued that the observed single-scatter nuclear recoils are due to scattering of "external" neutrons, neutrons produced by muon interactions outside the shield (see Chapter 3). The double-scatter neutron and Run 18 data sets are used to estimate the external-neutron background, which is then statistically subtracted from the single scatters to determine an exclusion limit on the spin-independent WIMP-nucleon elastic-scattering cross section. Rates of noise-charge events and of thermistor interactions due to tritium beta decay are also presented. Finally, some conclusions regarding residual photon and electron sources are made.

Recall the three ionization-partition cuts defined in Chapters 4 and 7: cQinOnly (inner-electrode-contained), cQShare (inner-outer-shared), and cQoutOnly (outer-electrode-contained). These cut names are used throughout this chapter for brevity.

8.2 Veto-Anticoincident Nuclear-Recoil Data Set

The veto-anticoincident nuclear-recoil data set is defined by the cuts presented in the previous chapter. This set is divided into single-scatter events, which are used to search for WIMP-induced nuclear recoils, and double-scatter events, which measure the external-neutron background.

8.2.1 Single-Scatter Data Set

Figures 8.1, 8.2, and 8.3 show plots of ionization energy vs. recoil energy for cQinOnly, cQShare, and cQoutOnly veto-anticoincident single-scatter events. Figures 8.4, 8.5, 8.6, 8.7, 8.8, and 8.9 show ionization yield vs. recoil energy plots for the same data. The ionization-yield plots are

split into two sets, before and after the April 3 change in the nuclear-recoil bands. The appropriate nuclear-recoil-acceptance regions are shown in the ionization-yield plots.

A number of features are clear:

- For all four detectors, the number of low-yield events (inside, above, and below the nuclear-recoil-acceptance regions) is significantly higher in the cQoutOnly data. This is expected, given the degraded photon rejection in the outer electrode observed in the photon-calibration data. It is likely that self-shielding also plays a role.
- The rate of low-ionization-yield events in BLIP3 is significantly higher than in the other detectors, even for the cQinOnly data set. As discussed in Section 4.3, BLIP3 suffered repeated processing steps during development of the new electrode-fabrication method, so its electrodes may have been contaminated or damaged during processing. These possibilities are discussed in Section 8.5.
- BLIP4 also shows an elevated low-yield event rate in its cQinOnly data in the form of a band in Y_R. As is discussed later, there is good evidence that these events are due to electrons emitted by the contaminant on BLIP3. There is good separation between this low-yield band and the nuclear-recoil-acceptance region.
- Most importantly, in both the cQinOnly and cQShare data sets, there appear to be populations of true nuclear-recoil events in BLIPs 4, 5, and 6.

The structure at about 10 keV is the 10.4-keV gallium line; the conversion from phonon energy and ionization energy to recoil energy and ionization yield makes it tilt. Barely visible are lines at 46 keV and 67 keV; the sources and implications of these lines are discussed in Section 8.5.

Photon and Electron Rates, Choice of Ionization-Partition Cut

Prior to performing the WIMP-search analysis, it is necessary to assess the bulk electron-recoil and surface-event rates, outside the nuclear-recoil band, to determine expected rates of misidentified particles. The primary goal is to determine whether an ionization-partition cut should be made to optimize the WIMP sensitivity. It is, of course, necessary that this decision be made on the basis of these data alone and not with any consideration of how many nuclear-recoil candidates are actually observed!

To assess the rates of bulk electron recoils and surface events, three "particle types" are defined: photons ($Y_R \sim 1$), nuclear recoils (anything in the nuclear-recoil-acceptance region), and electrons (all events between the photons and nuclear recoils). Figures 8.10, 8.11, and 8.12 show recoil-energy spectra for photons and electrons. Note that, because of the clear leakage of events into the nuclear-recoil band for cQoutOnly events, the cQoutOnly "electron" spectra underestimate the true rate of electron interactions; this is corrected in the discussion of residual background sources, Section 8.5. The average rates of photons and electrons between 10 and 100 keV are shown in Table 8.1. These rates are normalized using the volume fractions determined from the neutron-calibration data and simulation, as discussed in Section 4.5.1. It would be more physically reasonable to normalize the electron rates by surface area, rather than mass ($\sim$volume). However, volume-normalized rates indicate the impact of the electron rates on WIMP sensitivity. Electron

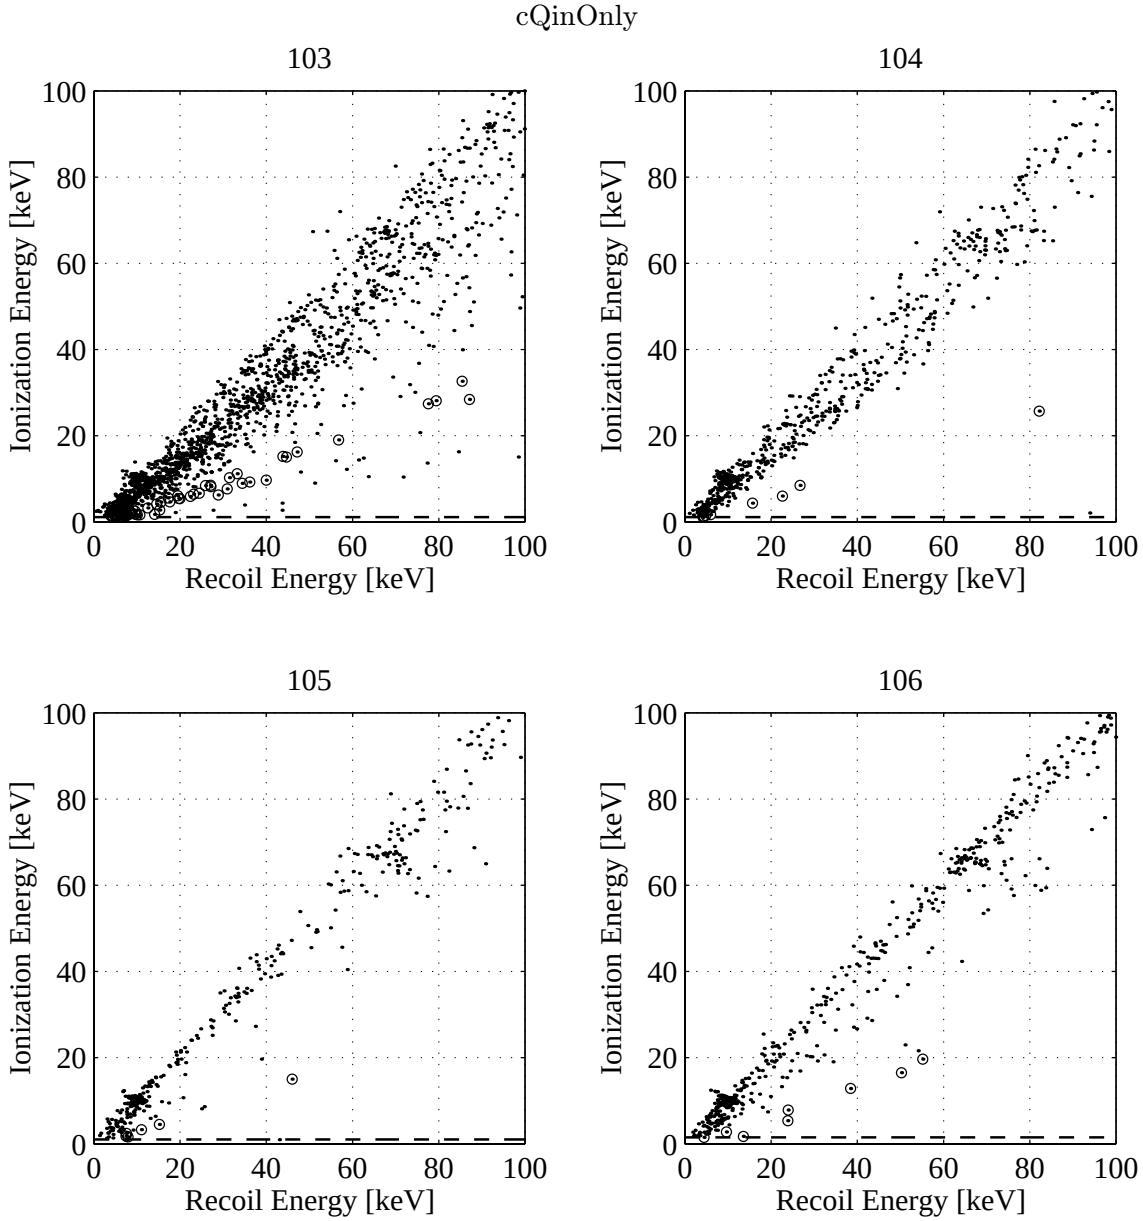

Figure 8.1: Ionization energy vs. recoil energy for veto-anticoincident cQinOnly single-scatter events. Nuclear-recoil candidates are circled. Note that, because the pre- and post-April 3 data sets are combined in these plots and the nuclear-recoil bands are different for the two data sets, uncircled events that appear to lie in the nuclear-recoil-acceptance region in fact do not. Dashed: ionization-search threshold.

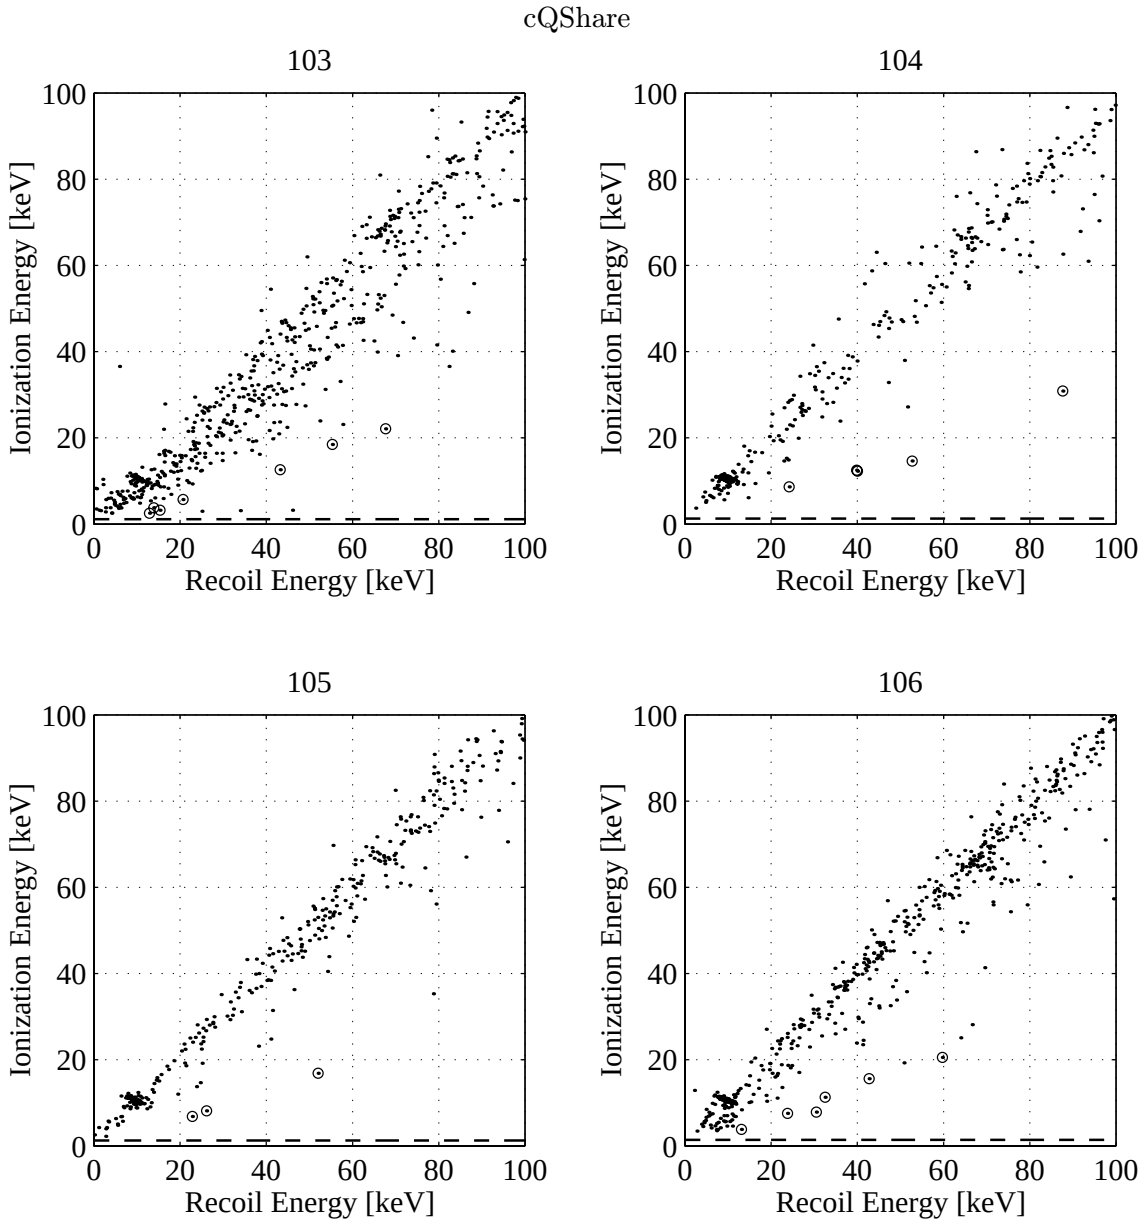

Figure 8.2: Ionization energy vs. recoil energy for veto-anticoincident cQShare single-scatter events. Legend as in Figure 8.1.

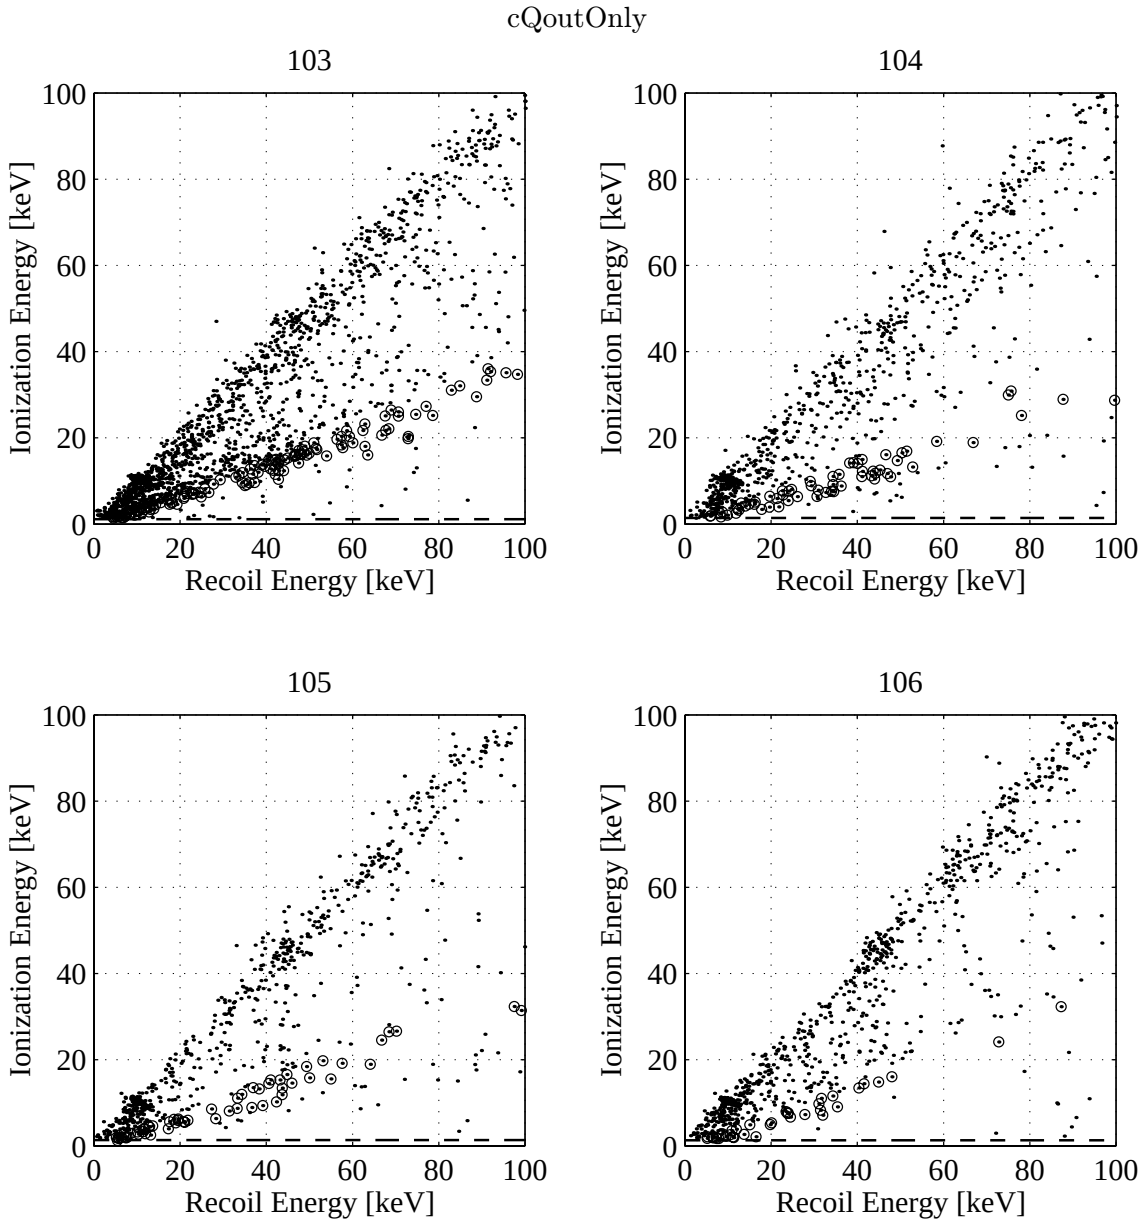

Figure 8.3: Ionization energy vs. recoil energy for veto-anticoincident cQoutOnly single-scatter events. Legend as in Figure 8.1.

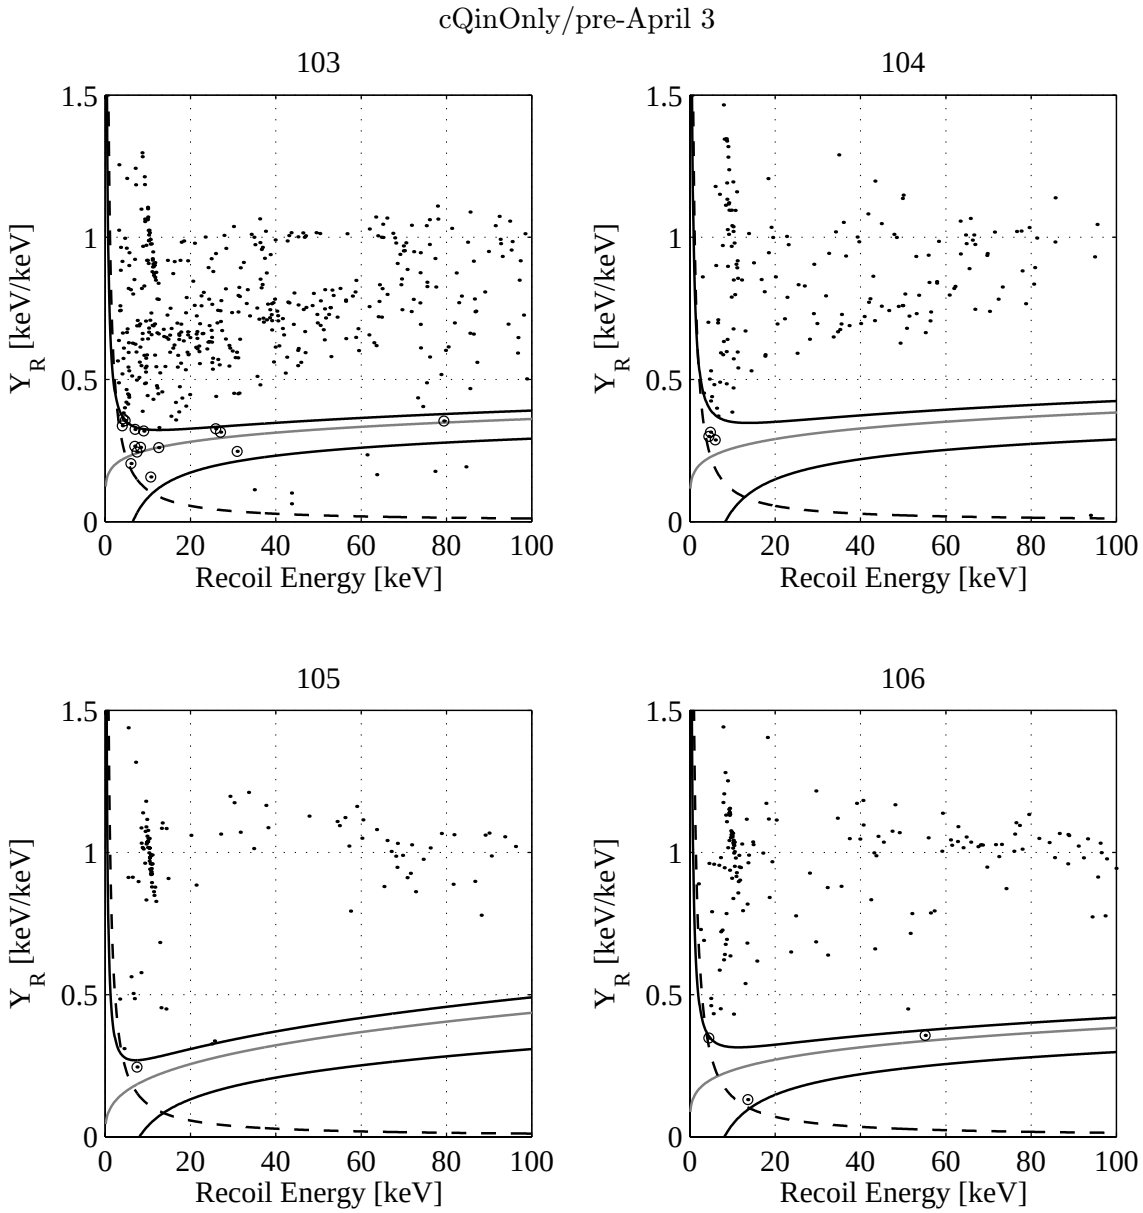

Figure 8.4: Ionization yield vs. recoil energy for veto-anticoincident cQinOnly single-scatter events in the pre-April 3 data set. Nuclear-recoil candidates are circled. Light line: center of nuclear-recoil band. Dark lines: nuclear-recoil acceptance region. Recall that the acceptance region is defined asymmetrically in Y_R. Dashed: ionization-search threshold.

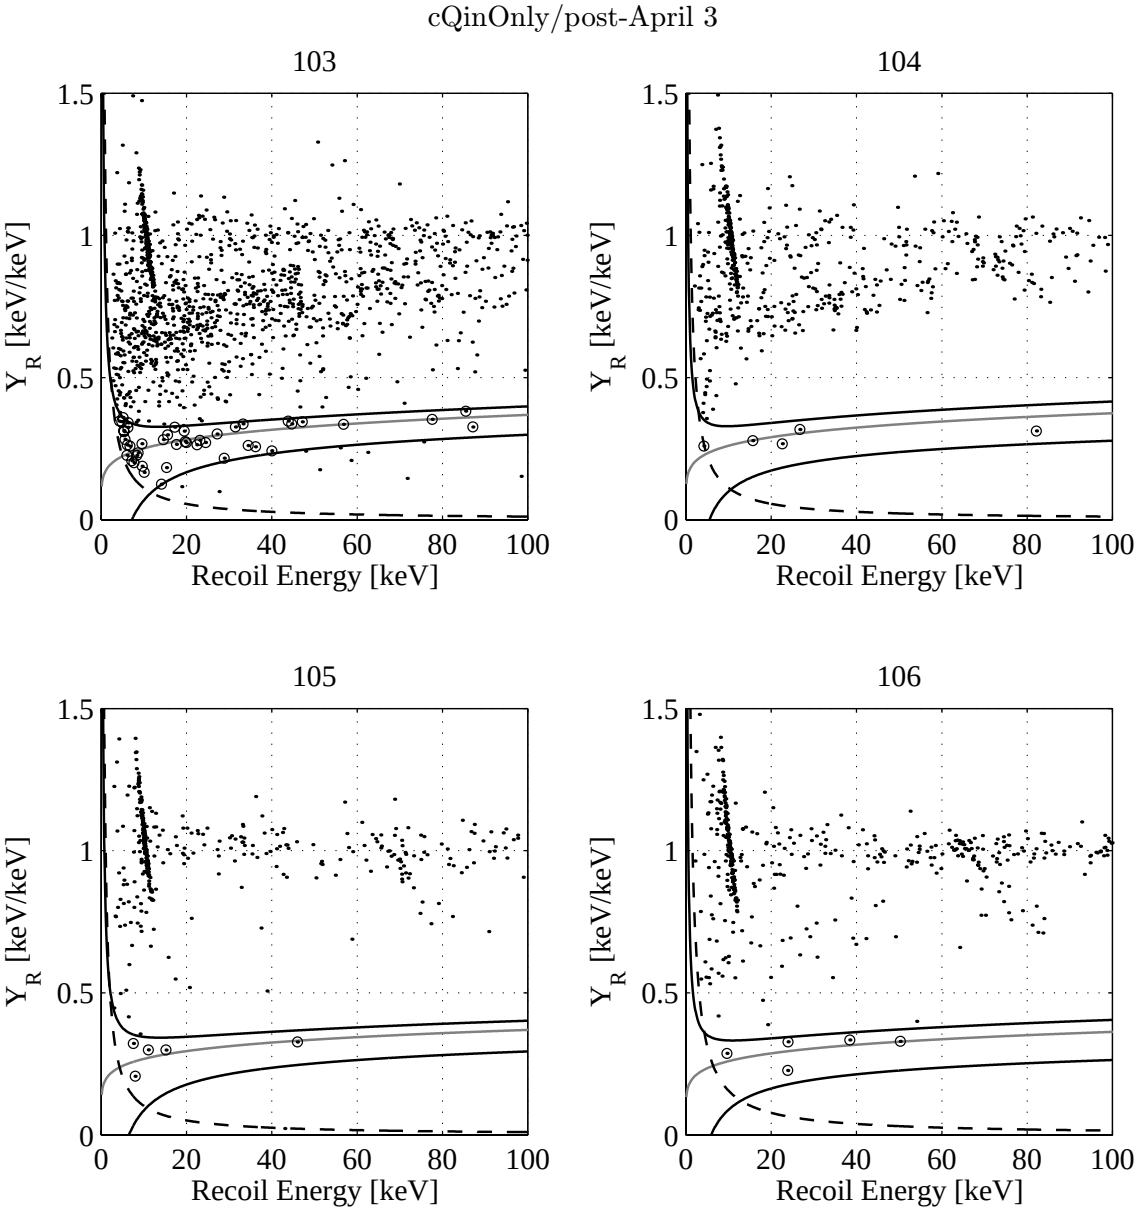

Figure 8.5: Ionization yield vs. recoil energy for veto-anticoincident cQinOnly single-scatter events in the post-April 3 data set. Legend as in Figure 8.4.

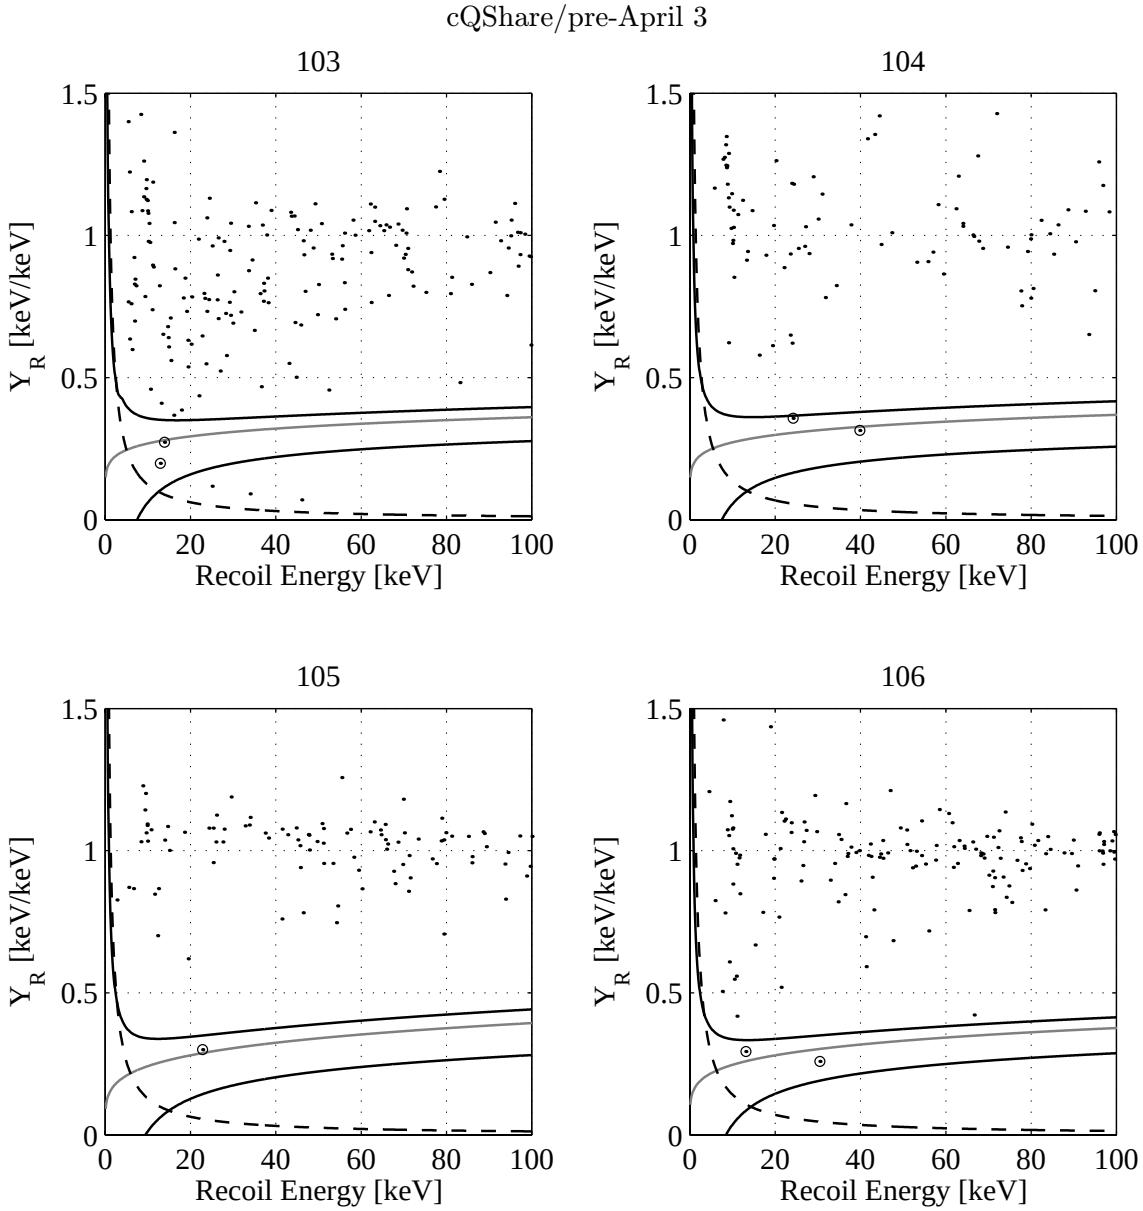

Figure 8.6: Ionization yield vs. recoil energy for veto-anticoincident cQShare single-scatter events in the pre-April 3 data set. Legend as in Figure 8.4.

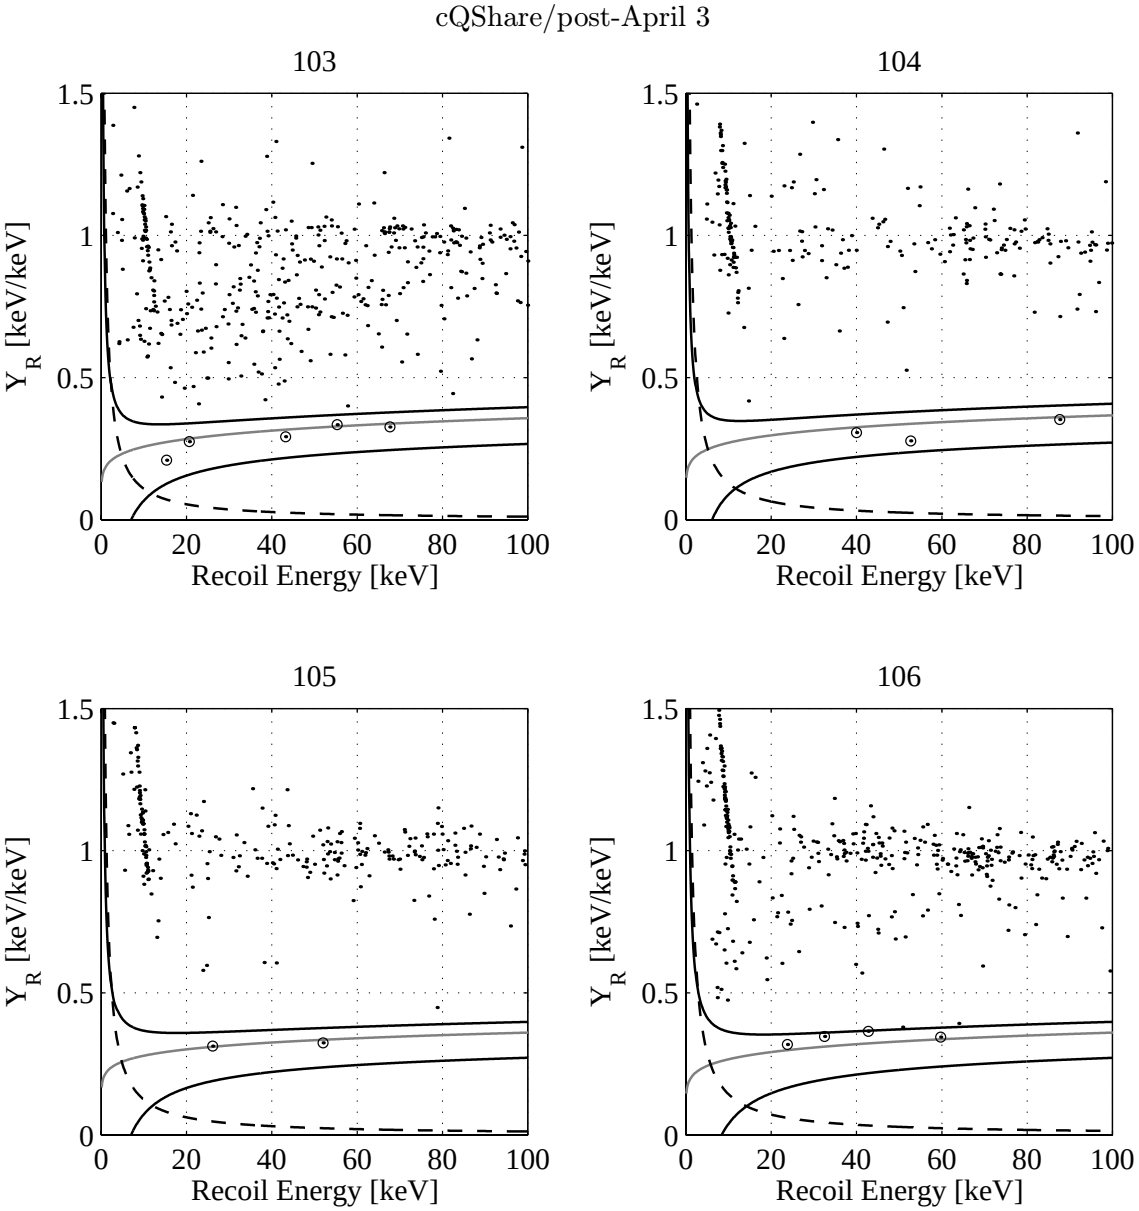

Figure 8.7: Ionization yield vs. recoil energy for veto-anticoincident cQShare single-scatter events in the post-April 3 data set. Legend as in Figure 8.4.

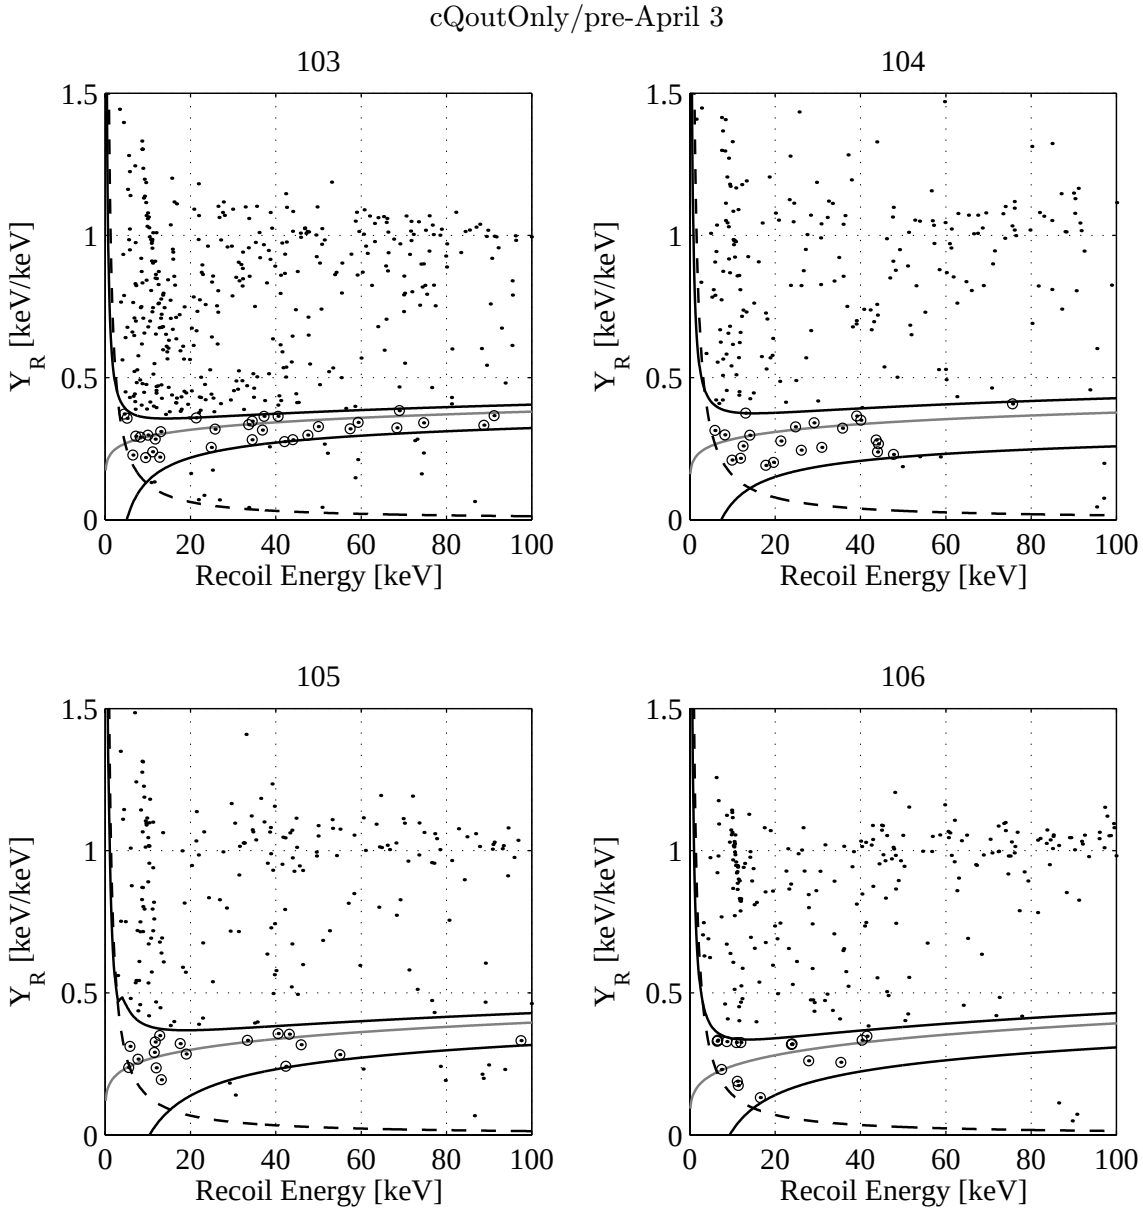

Figure 8.8: Ionization yield vs. recoil energy for veto-anticoincident cQoutOnly single-scatter events in the pre-April 3 data set. Legend as in Figure 8.4.

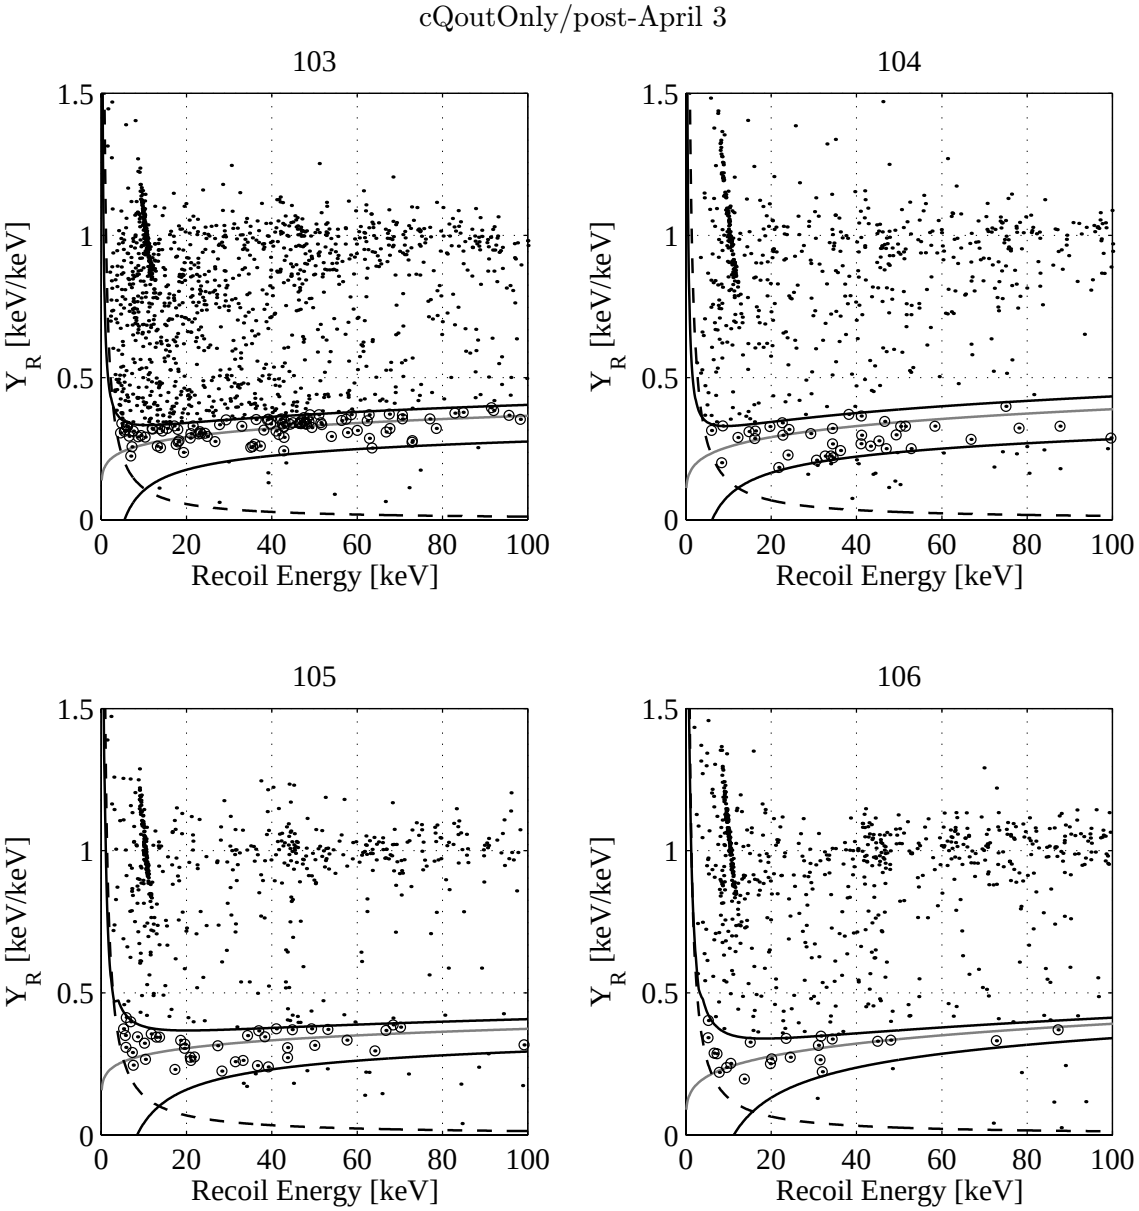

Figure 8.9: Ionization yield vs. recoil energy for veto-anticoincident cQoutOnly single-scatter events in the post-April 3 data set. Legend as in Figure 8.4.

Event set	BLIP3	BLIP4	BLIP5	BLIP6
Photons				
Qin	1.287±0.055	0.825±0.044	0.667±0.039	0.793±0.042
QShare	1.570±0.094	1.280±0.086	1.394±0.087	1.915±0.101
Qout	2.011±0.079	1.311±0.064	1.205±0.059	1.580±0.068
Electrons				
Qin	2.006±0.068	0.462±0.033	0.054±0.011	0.140±0.018
QShare	1.366±0.088	0.203±0.034	0.146±0.028	0.409±0.047
Qout	2.532±0.088	0.969±0.055	0.636±0.043	0.727±0.046

Table 8.1: Veto-anticoincident single-scatter-photon and -electron rates, averaged between 10 and 100 keV, in keV^{-1} kg^{-1} d^{-1}. The detector volume ratios used for the exposure normalization are 0.46:0.19:0.35, as calculated in Section 4.5.1. Uncertainties are statistical only; cut efficiencies add a systematic normalization error of about 10%, but the normalization scaling should be very similar for all the rates.

rates normalized by area are presented in Section 8.5. Note that, because the shared region appears to be located on the detector's flat faces, the ratio of cQinOnly to cQShare electron rates is the same for volume and area normalization; the relative rate of cQoutOnly electron events is overestimated by volume normalization (the denominator is too small). It should be noted that photons that scatter in both electrodes (and not necessarily in the nominal shared volume) contribute to the cQShare rate, elevating it above what is expected from the volume fraction; the higher cQShare photon rates are thus not unexpected.

A number of features are apparent in these rates. First, the cQinOnly photon rates among BLIPs 4, 5, and 6 are uniform at the 10% to 15% level. There is less uniformity among the cQShare event rates. There is also less uniformity among the cQoutOnly event rates. In the limit that photon events are dominated by Compton scattering of high-energy photons, the photon rates (in keV^{-1} kg^{-1} d^{-1}) should be the same for all ionization-partition cuts and detectors because the mean free path of high-energy photons is large compared to the detector thickness. (Note that the gallium-X-ray rates should also be proportional to volume.) The lack of uniformity among the detectors for the cQShare and cQoutOnly cuts may be due to varying exposures to 46-keV photons emitted by ^{210}Pb (presumably situated on the detector housing); these photons are apparent in the spectra. The implications of the variations in photon rates, and especially the differing intensities of the photon lines, are explored in Section 8.5. Second, though the electron rates observed for the cQoutOnly cut are similar to those observed for the cQinOnly or cQShare cuts, it is apparent from the ionization-yield plots that the cQoutOnly events appear to be distributed more uniformly in Y_R, while the cQinOnly and cQShare events form a band at relatively high Y_R with a tail. This is corroborated by the histograms of ionization yield, Figures 8.13, 8.14, and 8.15, which even indicate a peak in the Y_R distribution at fairly low Y_R for the cQoutOnly data. Finally, the cQinOnly and cQShare electron rates and Y_R histograms are quite similar. For a flux of electrons emitted by an external source (the detector package, for example), the inner-electrode regions are more shielded than the shared regions, so one would expect the latter to have significantly higher electron rates. However, if the residual electron source is surface contamination, which is likely to be uniformly

distributed over a detector's surface, it is reasonable that similar rates are observed. Implications for the residual electron-source location are discussed in Section 8.5.

Based on the ionization-yield plots and these rates, the following decisions can be made regarding definition of the WIMP-search data set. First, cQoutOnly events should be discarded, for a number of reasons. Though the cQoutOnly photon rates are not significantly higher than the cQinOnly or cQShare rates, the photon calibration indicates that the photon-misidentification parameter is 2 to 10 times higher for cQoutOnly than for cQinOnly or cQShare. This alone is good reason to discard the cQoutOnly data. Beyond this, the much flatter Y_R distributions seen for the cQoutOnly data indicate that, though the cQoutOnly electron rate is not significantly different from the rates seen for the cQinOnly and cQShare cuts, the electron-misidentification parameter is likely to be much worse. It would be preferable to use a separate electron calibration to demonstrate conclusively that the outer-electrode electron-misidentification parameter is higher. Unfortunately, the electron calibration presented in Section 4.5.3 does not probe the top-bottom break. However, physically, it is clear that if photon misidentification is worse in the outer electrode by a factor of 2 to 10, electron misidentification is likely to be drastically worse.

Second, there appears to be no reason to discard the cQShare data set. The electron and photon rates (modulo the overestimation issue discussed above) are not significantly higher than for the cQinOnly data set. The calibration data sets indicate that both the photon- and electron-misidentification parameters for the shared region appear to be no worse than for the inner-electrode region. The Y_R histograms for the background data corroborate this point. Thus, the expected rate of misidentified photons and electrons in the two regions should be about the same. It should be stressed that two pieces of data play roles in this decision: first, the observed similar rates; and second, the observed similar misidentification parameters β_γ and β_β. A counterexample would be a case in which, though the misidentification parameters were the same, the cQShare photon and electron rates were, for example, 10 times higher than the cQinOnly rates. The larger rates would yield 10 times larger expected rates of misidentified events, obviously indicating that the cQShare data set should be discarded in such a case.

When the WIMP-search analysis was initially performed, only cQinOnly events were used. This choice was based on an earlier analysis comparing the *veto-coincident*-electron rates for the cQinOnly and cQShare sets, which saw a significantly higher cQShare electron rate (of order 30 times) [141]. This was interpreted as evidence that the electron misidentification for shared events is higher, and thus the cQShare data set was discarded. However, the more recent analysis presented in Section 4.5.3 shows that, at the available statistical precision, the electron misidentification for cQShare events is no worse than for cQinOnly events. This latter result suggests that the higher veto-coincident cQShare electron rate is due primarily to a higher incident flux on the shared region, relative to the inner-electrode region, because the shared region is not as well shielded. Since the veto-anticoincident-electron event rates do not reflect such a difference, there is no reason to distinguish the cQinOnly and cQShare data sets.

Having said this, the cQShare data set is discarded from the WIMP-search analysis at this point simply because redoing the analysis with the inclusion of the cQShare data set is not practical on the timescale of this dissertation. There is, in principle, nothing wrong with this approach, as long as the WIMP-detection efficiency is appropriately calculated. The possible impact of inclusion of the cQShare events on the final results is mentioned in Section 8.4.

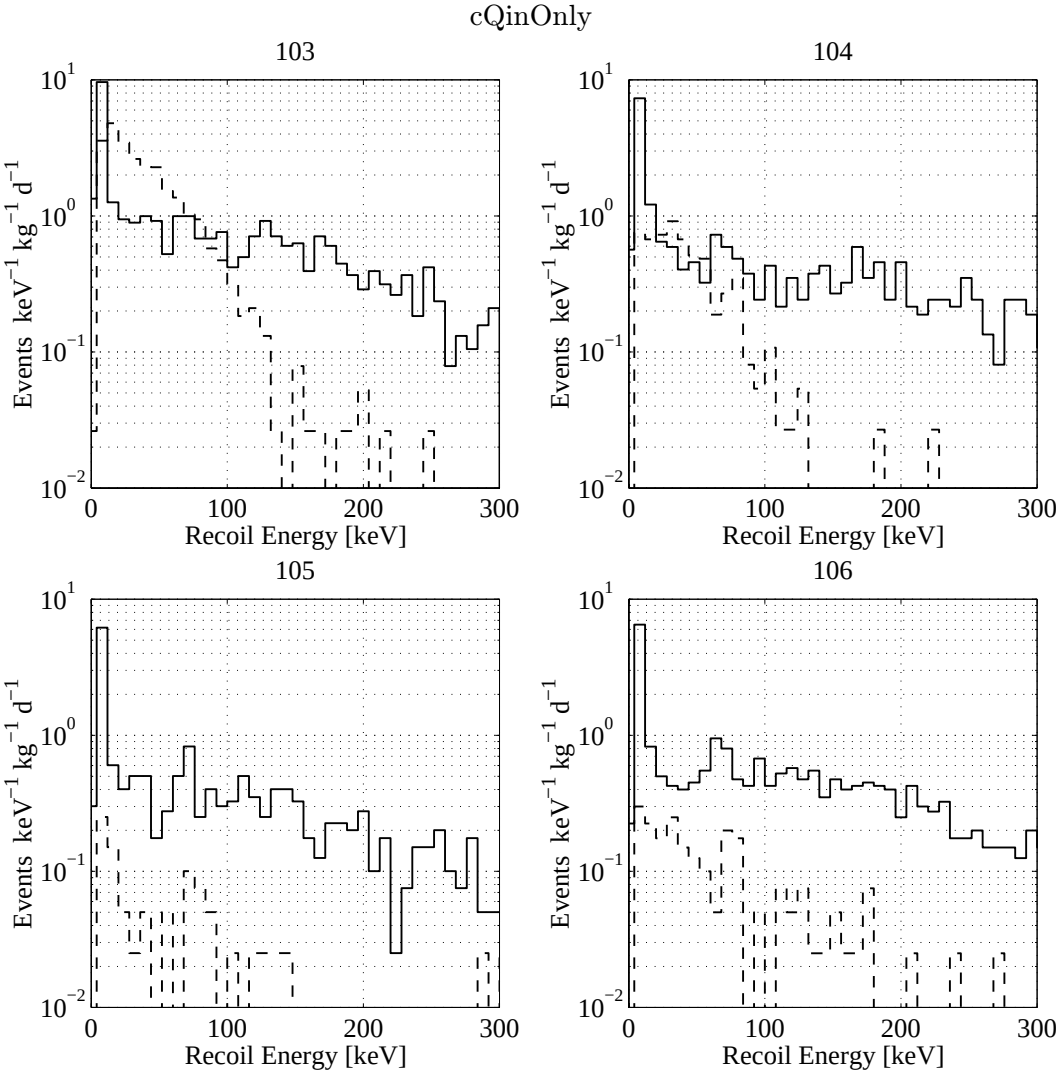

Figure 8.10: Single-scatter photon and electron recoil-energy spectra for veto-anticoincident cQinOnly events. Solid: photons. Dashed: electrons. Note that the energy range is 0 to 300 keV.

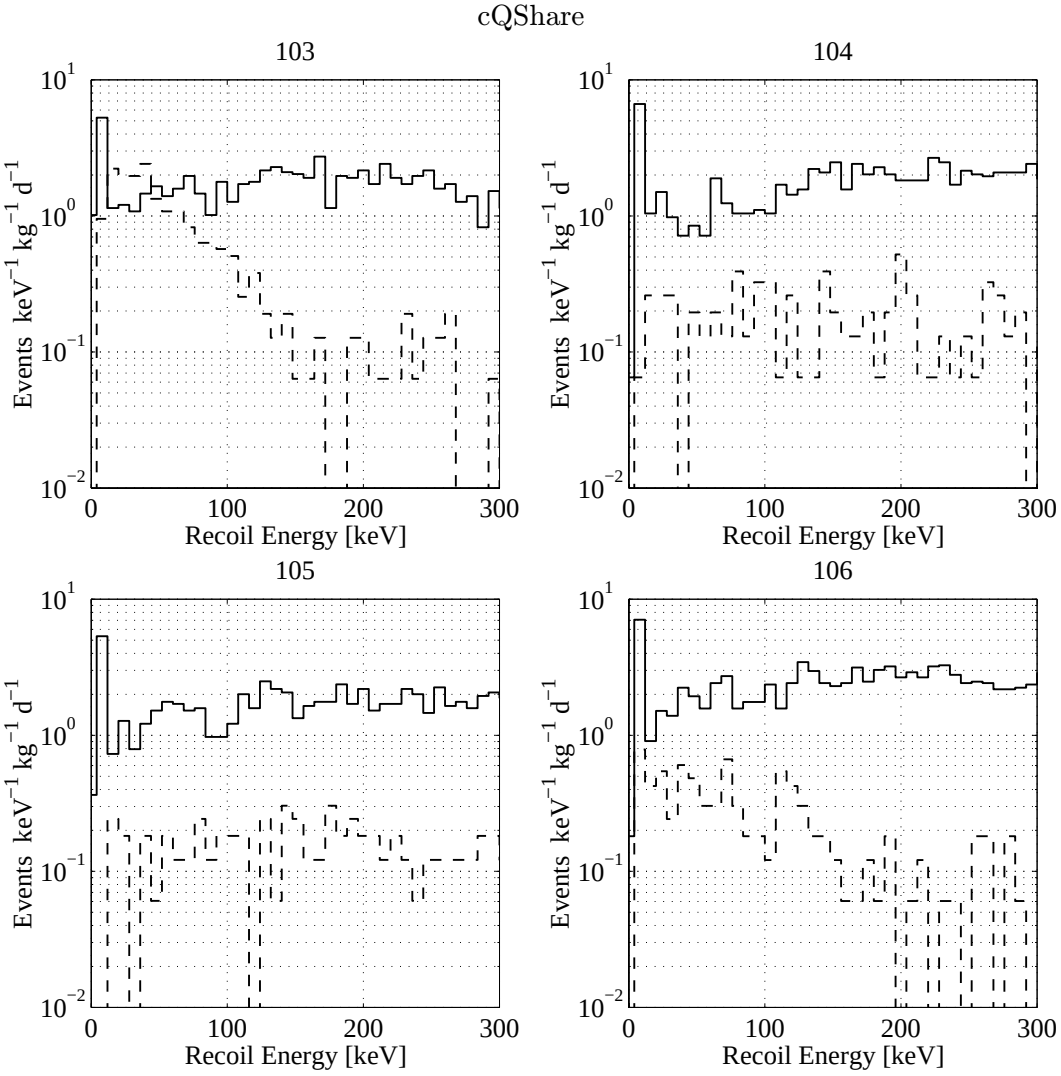

Figure 8.11: Single-scatter photon and electron recoil-energy spectra for veto-anticoincident cQShare events. Solid: photons. Dashed: electrons.

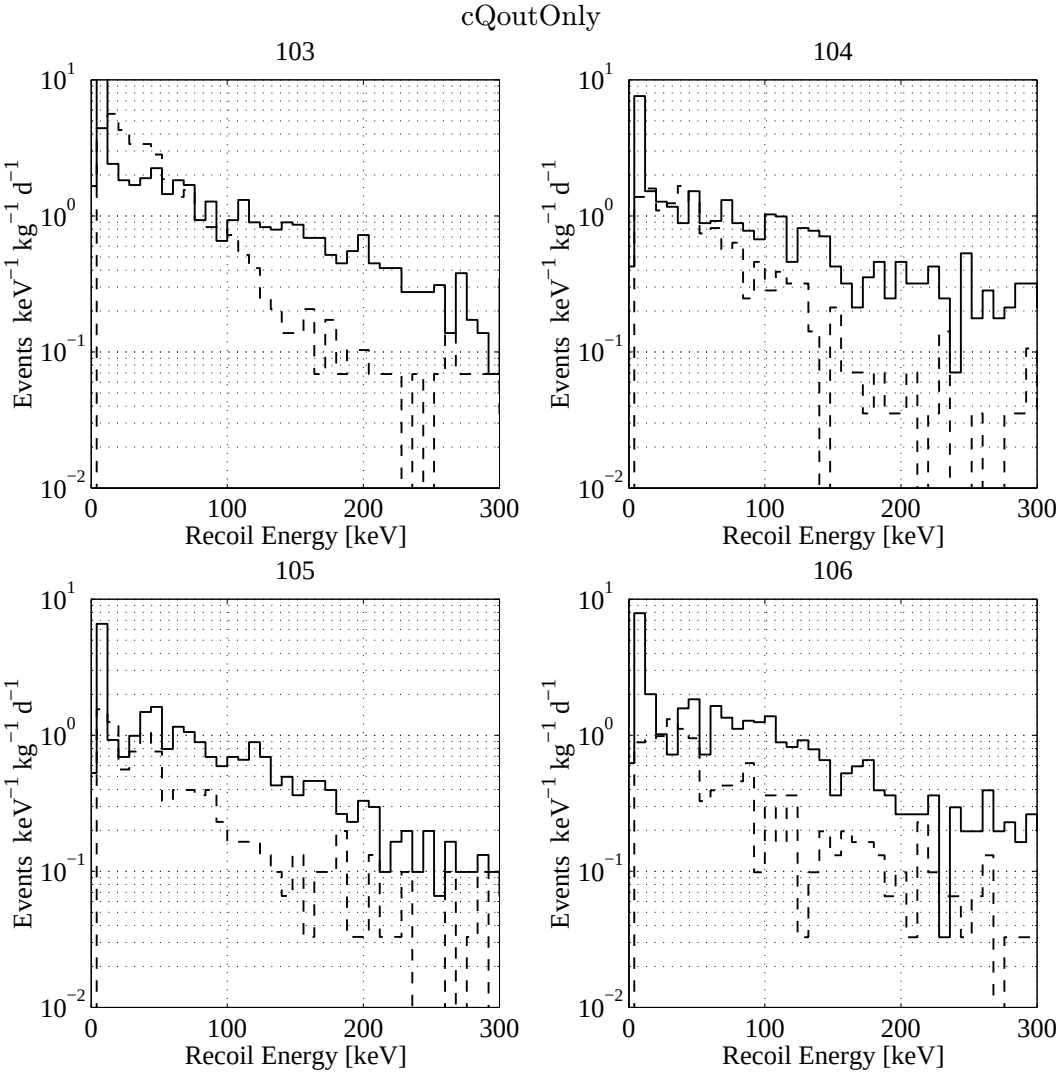

Figure 8.12: Single-scatter photon and electron recoil-energy spectra for veto-anticoincident cQoutOnly events. Solid: photons. Dashed: electrons.

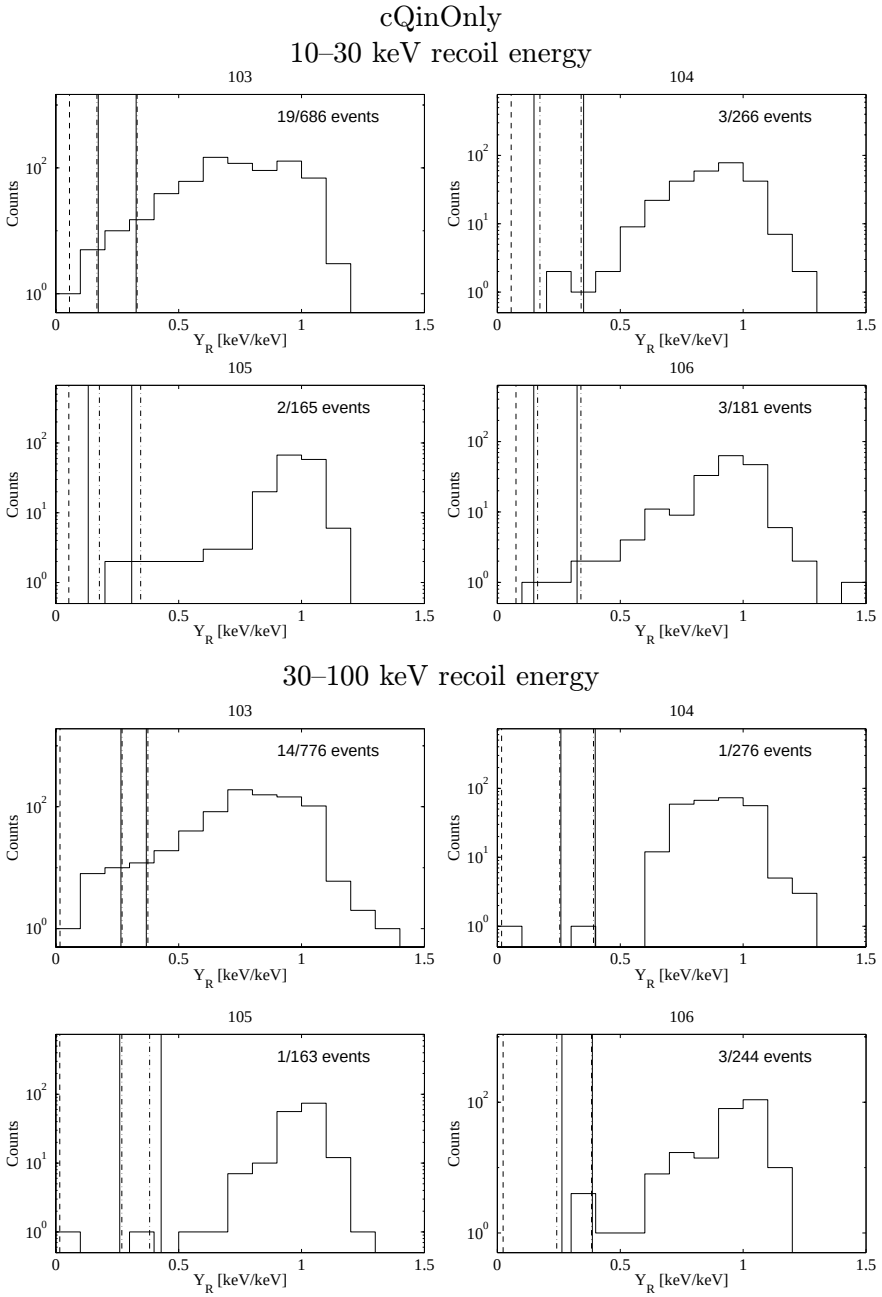

Figure 8.13: Distributions of ionization yield for veto-anticoincident cQinOnly single-scatter events in two recoil-energy bins. Dashed: ionization-search threshold. Dashed-dotted: nuclear-recoil-acceptance region for pre-April 3 data. Solid vertical lines: nuclear-recoil-acceptance region for post-April 3 data. The ionization-search thresholds and acceptance-region edges are averaged over the recoil-energy bins. Also shown are the number of nuclear-recoil candidates and the total number of events in each histogram, determined using the fully energy-dependent nuclear-recoil-acceptance regions.

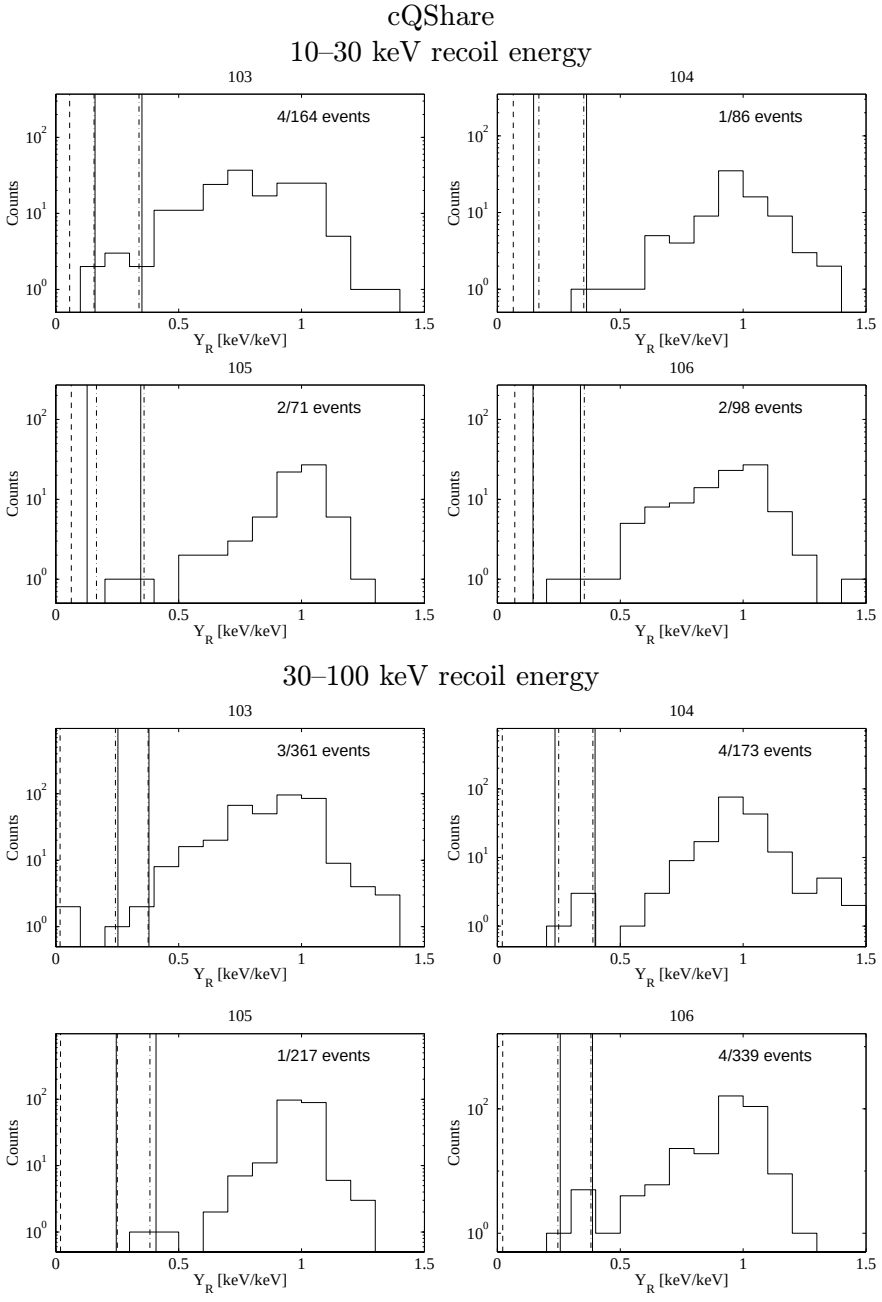

Figure 8.14: Distributions of ionization yield for veto-anticoincident cQShare single-scatter events. Legend as in Figure 8.13.

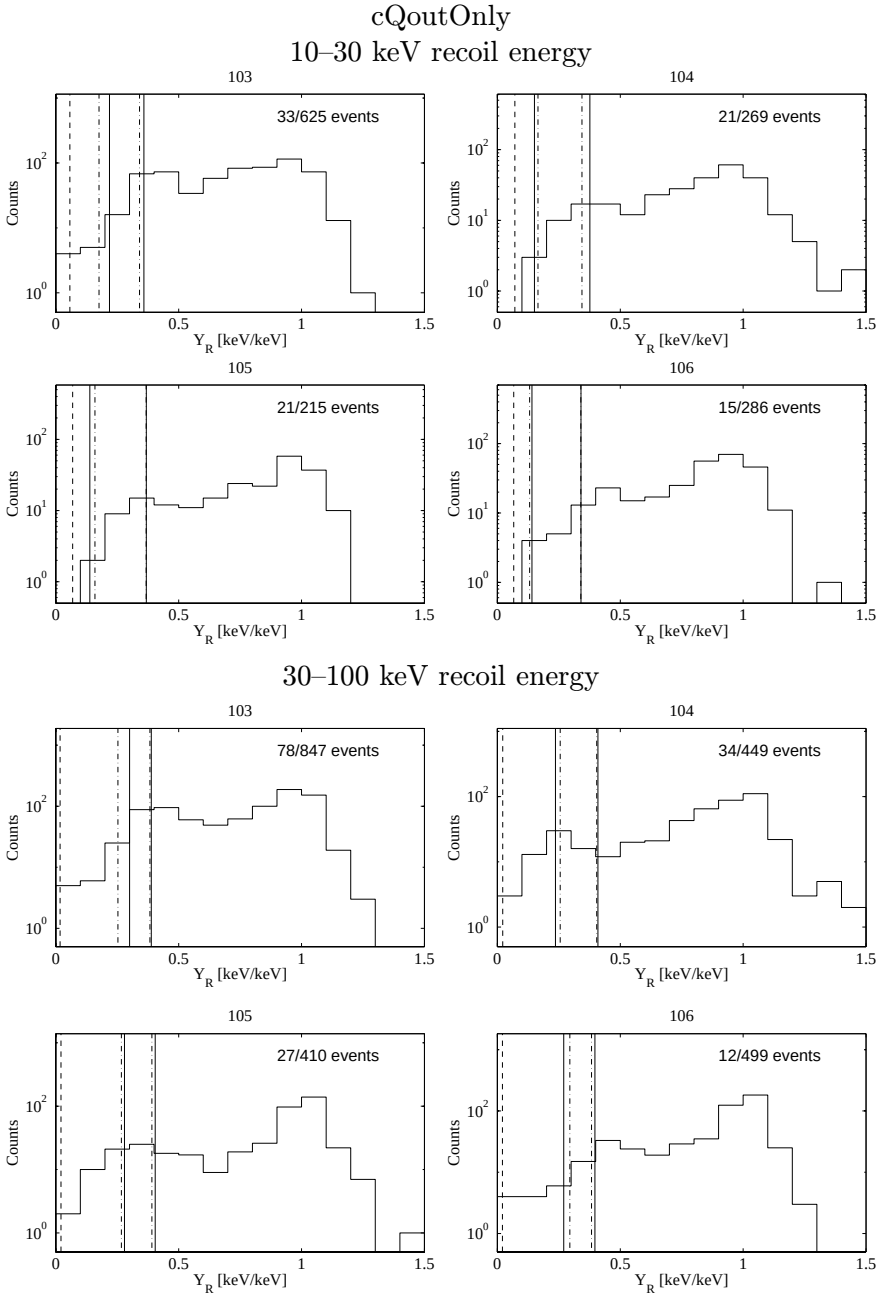

Figure 8.15: Distributions of ionization yield for veto-anticoincident cQoutOnly single-scatter events. Legend as in Figure 8.13.

Definition of WIMP-Search Data Set, Expected Numbers of Misidentified Photons and Electrons

The search for WIMP-induced nuclear recoils is restricted to cQinOnly events in BLIPs 4, 5, and 6 between 10 and 100 keV. The rationale for discarding events outside the 10-to-100-keV window was discussed in Section 7.5; it is also clear from the ionization-yield plots that the band of surface events intrudes into the nuclear-recoil-acceptance region below 10 keV. BLIP3 is discarded because of its clear contamination. The reasons for discarding cQoutOnly and cQShare events are discussed in the previous section.

Prior to considering the nuclear-recoil data set defined by the above cuts, it is useful to determine the expected numbers of misidentified photons and electrons passing these cuts. The observed photon and electron events rates can be combined with the photon- and electron-calibration data to set Bayesian 90% CL upper limits on the expected numbers of misidentified photons and electrons. Both the photon- and electron-calibration data are discussed in Chapter 4. Recall that the electron-calibration data set consists of BLIP3/BLIP4 double scatters arising from BLIP3's surface contamination; detailed justification for the definition of the electron-calibration set is discussed in Section 8.2.2. I describe the method for calculating the upper limit in detail for electrons; an analogous formalism holds for photons. This formalism is overkill for the single-scatter data, but it is useful to present it for the simple case of single-scatter misidentification to clarify the discussion of double-scatter misidentification later. Recall that $\beta_\beta \equiv N_l/N_\beta$, where N_l is the number of misidentified single-scatter electrons (the number that leak into the nuclear-recoil-acceptance region) and N_β is the number of correctly identified single-scatter events (the number that appear between the nuclear-recoil-acceptance region and the bulk electron-recoil band). Note that the denominator is *not* $N_{tot} = N_\beta + N_l$; these definitions of β_β and N_β make the formulae simpler. μ_l^b is defined to be the expected number of misidentified single-scatter electron events for the veto-anticoincident single-scatter ("background") data set. An upper limit will be set on μ_l^b. Let μ_l^c be the expected number of misidentified electron events in the calibration set. Define $\lambda = 1/\beta_\beta$. The expected values for the observables N_β^b, N_l^c, and N_β^c — the number of correctly identified single-scatter electrons in the background set, the number of misidentified electrons in the calibration set, and the number of correctly identified electrons in the calibration set, respectively — are

$$\begin{aligned} \langle N_\beta^b \rangle &= \lambda \mu_l^b \\ \langle N_l^c \rangle &= \mu_l^c \\ \langle N_\beta^c \rangle &= \lambda \mu_l^c \end{aligned} \tag{8.1}$$

The likelihood function for these data is given by

$$\mathcal{L}(N_\beta^b, N_l^c, N_\beta^c | \mu_l^b, \mu_l^c, \lambda) = \frac{(\lambda\mu_l^b)^{N_\beta^b} e^{-\lambda\mu_l^b}}{N_\beta^b!} \frac{(\mu_l^c)^{N_l^c} e^{-\mu_l^c}}{N_l^c!} \frac{(\lambda\mu_l^c)^{N_\beta^c} e^{-\lambda\mu_l^c}}{N_\beta^c!} \tag{8.2}$$

which is just the product of Poisson distribution functions with the appropriate expected values. To calculate a Bayesian upper limit on μ_l^b, it is necessary to integrate over the "nuisance" parameters μ_l^c and λ and calculate the confidence interval based on the resulting function of μ_l^b only. The 90%

Event set	N_β^b	N_l^c	N_β^c	$\mu_{l,90}^b$	$\mu_{l,10}^b$	$\mu_{l,50}^b$
10 – 30 keV	101	1	46	8.9	1.2	3.7
30 – 100 keV	180	0	39	11.1	0.5	3.2
10 – 100 keV	281	1	85	13.1	1.8	5.6

Table 8.2: Veto-anticoincident cQinOnly single-scatter-electron misidentification estimates, calculated by the method discussed in the text. The first column is the number of single-scatter-electron events observed in the background data for the given energy bin, coadded over BLIPs 4, 5, and 6. The second two columns list the calibration data used, which come from the BLIP3 column of Table 4.6. Recall that the calibration data set consists of the BLIP3/BLIP4 veto-anticoincident double-scatter events. The fourth column is the Bayesian 90% CL upper limit on the number of misidentified single-scatter electrons. Because of the large values of N_β^c and N_β^b, these values are very close to what one gets by calculating a Bayesian 90% CL upper limit on β_β from the electron calibration and multiplying by N_β^b. The fifth and sixth columns give the Bayesian 10% CL and 50% CL upper limits on μ_l^b. The large spread among the 10%, 50% and 90% CL values indicates the impact of the low statistical precision of the electron calibration.

CL upper limit on μ_l^b, termed $\mu_{l,90}^b$, is defined by

$$\frac{\int_0^{\mu_{l,90}^b} d\mu_l^b \int_0^\infty d\mu_l^c \int_0^\infty d\lambda\, \mathcal{L}(N_\beta^b, N_l^c, N_\beta^c | \mu_l^b, \mu_l^c, \lambda)}{\int_0^\infty d\mu_l^b \int_0^\infty d\mu_l^c \int_0^\infty d\lambda\, \mathcal{L}(N_\beta^b, N_l^c, N_\beta^c | \mu_l^b, \mu_l^c, \lambda)} = 0.9 \tag{8.3}$$

Note that a uniform integration measure is used for the integrals; in the Bayesian formalism, this choice is a matter of judgment. (A logarithmic measure for λ may be more appropriate, but, in general, the choice of measure does not significantly affect the result). It is necessary to invoke this machinery because β_β and β_γ are determined with significant statistical uncertainty: in both calibrations, the number of misidentified events is in the range 0 to 2. As a general point, this formalism is also necessary if N_β^b is small enough that it has significant statistical uncertainty, though, if β_β is determined with negligible uncertainty, the formalism simplifies greatly because the calibration pieces of the likelihood function become δ-function-like.

The denominator can be calculated analytically, but the numerator cannot: one of the integrals can be done analytically, but the remaining two are nontrivial. Instead, the numerator is calculated numerically. Tables 8.2 and 8.3 show the resulting upper limits on misidentified single-scatter electrons and photons, respectively. In the photon case, an approximation has been made: N_γ^b has been taken to be only the number of events in the photon band, while N_γ^c included both photons and electrons in the photon-calibration set. This introduces a negligible error because bulk electron recoils make up approximately 97% of photon-calibration events [132].

Based on the above analysis, photon misidentification contributes a negligible number of nuclear-recoil candidates. The upper limit that can be set on electron misidentification is not nearly so useful. This occurs for two reasons. First, the electron calibration is statistics limited: even if no misidentified events had been seen in the electron calibration, the 50% CL and 90% CL upper limits would still be nonnegligible. Second, the one event seen in the nuclear-recoil band in the electron calibration Y_R^2 plot (Figure 4.52) looks like a multiple-scatter neutron. The above analysis

Event set	N_γ^b	N_l^c	N_γ^c	$\mu_{l,90}^b$
BLIP4				
10 – 30 keV	187	0	1613	0.27
30 – 100 keV	154	0	1694	0.21
10 – 100 keV	341	0	3307	0.24
BLIP5				
10 – 30 keV	147	0	1623	0.21
30 – 100 keV	156	0	2556	0.14
10 – 100 keV	293	0	4179	0.16
BLIP6				
10 – 30 keV	156	2	1425	0.58
30 – 100 keV	198	0	1359	0.34
10 – 100 keV	354	2	2784	0.68
Coadded				
10 – 30 keV	490	2	4661	0.56
30 – 100 keV	498	0	5609	0.20
10 – 100 keV	988	2	10270	0.51

Table 8.3: Veto-anticoincident cQinOnly single-scatter-photon misidentification estimates, calculated by the method discussed in the text. The first column is the number of single-scatter-photon events observed in the background set in the given energy bin. The other two columns list the calibration data used, which are given by the histograms in Figure 4.49. As one would expect from the large values of N_γ^c and N_γ^b, the final upper limits are very close to what one gets by calculating a Bayesian 90% CL upper limit on β_γ from the calibration and multiplying by N_γ^b.

conservatively assumes it is a misidentified surface event.

WIMP-Search Data Set

The WIMP-search data set consists of the veto-anticoincident cQinOnly single-scatter nuclear-recoil candidates passing the cuts described in Chapter 7 and the previous section; it contains 13 events. From the ionization-yield plots for the cQinOnly data, it is clear on an intuitive level that these events are not electrons: there is clear separation of the nuclear-recoil candidates from the low-energy-electron band at suppressed ionization yield. The expected numbers of misidentified photons are calculated above, rigorously demonstrating that misidentified photons contribute negligibly to the 13 observed events. Such a strong statement cannot be made for electrons because of the small number of events in the electron calibration. However, an electron interpretation is strongly disfavored by the observation of four double-scatter nuclear recoils with negligible electron leakage, as is discussed in Section 8.2.2: this double-scatter rate predicts a single-scatter rate consistent with the observed number of single-scatter nuclear recoils, as is discussed in Section 8.3.2. Given these arguments, the single-scatter nuclear-recoil candidates are interpreted as real nuclear recoils rather than misidentified photons or electrons.

Figure 8.16 shows the recoil-energy spectrum of these 13 nuclear-recoil events, coadded

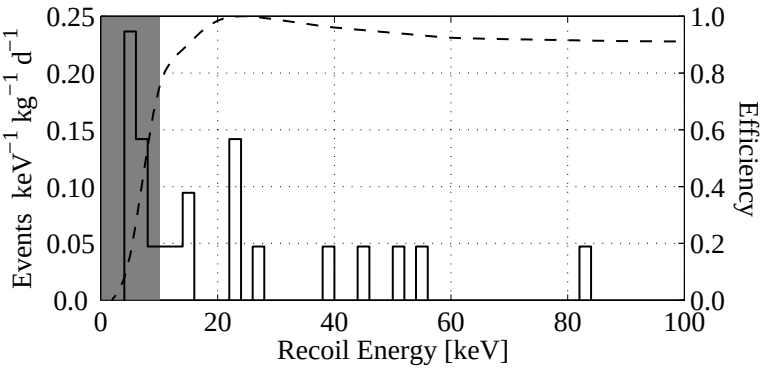

Figure 8.16: Recoil-energy spectrum for veto-anticoincident cQinOnly single-scatter nuclear recoils. The exposure used to normalize the spectrum is 10.6 kg d. The nuclear recoil efficiency, normalized to the 10.6 kg d exposure, is also shown. The grey shaded region indicates the 10-keV energy threshold applied. The spectrum is *not* corrected for the energy dependence of the efficiency function.

over BLIPs 4, 5, and 6. Also shown is the overall efficiency function for nuclear recoils as calculated in the previous chapter, normalized to have peak efficiency of 1. The exposure corresponding to the peak is 10.6 kg d. The exposure and efficiency function have been fully corrected for the cuts defining the data set. The spectrum has been normalized to the peak exposure, but has *not* been corrected for the energy dependence of the efficiency function. It is not helpful to do so here because of the small number of events; the energy dependence is correctly taken into account in later analysis. One conservative approximation has been made — the volume of the detector corresponding to cQinOnly events has not been corrected in the way discussed in Section 4.5.1. That is, the cQinOnly cut is taken to accept 40% of the detector volume, as is naively suggested by the neutron calibration, rather then the apparent 46% one derives after the two corrections discussed in that section. This decreases the exposure artificially by about 15%. Note that the use of 10.6 kg d is an approximation only for WIMPs; veto-anticoincident neutrons are subject to the same scattering physics as the calibration-source neutrons, so it is true that only 40% of veto-anticoincident neutron scatters are accepted by the cQinOnly cut.

Appendix D lists most of the useful parameters of these 13 nuclear-recoil events — their recoil energies and ionization yields, veto information, trigger-time information, etc. In this section, a number of Kolmogorov-Smirnov tests are performed on these events to establish that they exhibit no unusual characteristics. A Kolmogorov-Smirnov test compares the cumulative distribution of the events in a given event parameter to the expected distribution. It finds the point that deviates furthest from the expected cumulative distribution function (CDF) and assesses how many experiments would have observed events with larger deviations. For a detailed description of the formalism of the KS test, please consult [142] or [143].

The first, most obvious test is to check that the events exhibit no unexpected correlation with veto activity. A KS test of the last veto-trigger times for these events is shown in Figure 8.17. This test is only performed on the 10 ionization-trigger events. These times should follow an expo-

nential distribution corresponding to the last veto-trigger times being randomly distributed relative to the event time. The test indicates that 36% of experiments would have observed distributions that deviate further from the expected distribution. A different test, using the *nearest* veto-trigger time rather than the *last* veto-trigger time, could be performed for the 3 remaining events that are phonon triggers, but is not done because it would not be powerful with so few events. I also note that the events that occurred during the period that veto ADC information was available exhibit vanishing low and high veto pulse integrals for all counters, again indicating no simultaneous activity in the veto.

It is also possible to test the time distribution of the events. Each event has a live time and a number of live detectors (0 to 3) associated with it. (Fewer than three detectors may be alive at any given time because of downtime due to detector misbehavior or electronics problems). The coadded live time per event is the product of these two quantities and an average cut efficiency. Since the cut efficiencies are not strong functions of time or of detector, this last factor is taken to be a constant; since the test is performed on the cumulative coadded live time normalized to the total coadded live time, this constant factor cancels and can be ignored. The expected CDF is thus the cumulative fraction of events observed as a function of the cumulative fractional coadded live time. Note that this distribution includes events where BLIPs 4, 5, and 6 did not trigger; they were alive and could have triggered, so the live time of these events is included. The inverse of the slope of the CDF is the average fractional coadded live time per fraction of the number of events. The main contributions to variations in this number are discrete changes in the number of detectors alive at a given time (reasons for single detectors being dead are described in Chapter 7) or discrete changes in the trigger rate due to certain detectors' triggers being disabled. For example, the expected distribution shows a slope change at the point where BLIPs 1 and 2 were turned off: the average coadded live time per event increases because the overall trigger rate decreases. It is important to stress that the live times of BLIPs 1 and 2 (and BLIP3) are not included in the coadded live time per event; however, the BLIPs 4/5/6 live time for events in which BLIPs 1, 2, or 3 triggered and none of BLIPs 4, 5, or 6 triggered must be included to give the correct coadded live time. The KS test indicates 19% of experiments would have observed distributions that deviate further from the expected distribution. The observed distribution for the 13 events appears to systematically deviate at early live time. While troubling, the 19% probability indicates that this deviation is not cause for significant worry.

The distribution in ionization yield of the nuclear recoils can be compared to the expected distribution. The normalized deviation, termed Y_R^* in Andrew Sonnenschein's dissertation [8], is defined by:

$$Y_R^* = \frac{Y_R - Y_{NR}(E_R)}{\sigma_{NR}(E_R)} \tag{8.4}$$

where $Y_{NR}(E_R)$ is the position of the center of the nuclear-recoil band in ionization yield and $\sigma_{NR}(E_R)$ is the standard deviation in Y_R of the band, both functions of E_R. The usefulness of Y_R^* is that it puts nuclear recoils at different E_R on the same footing. In the absence of cuts in Y_R defining the acceptance region, the expected distribution is a simple Gaussian with mean 0 and standard deviation 1. The cuts that define the nuclear-recoil band truncate the distribution in an E_R-dependent manner. One ambiguity arises in calculating the effect of these cuts on the expected distribution. Without these cuts, the expected distribution of Y_R^* at all recoil energies is the same, so combining Y_R^* distributions from different recoil energies is trivial. With the cuts implemented,

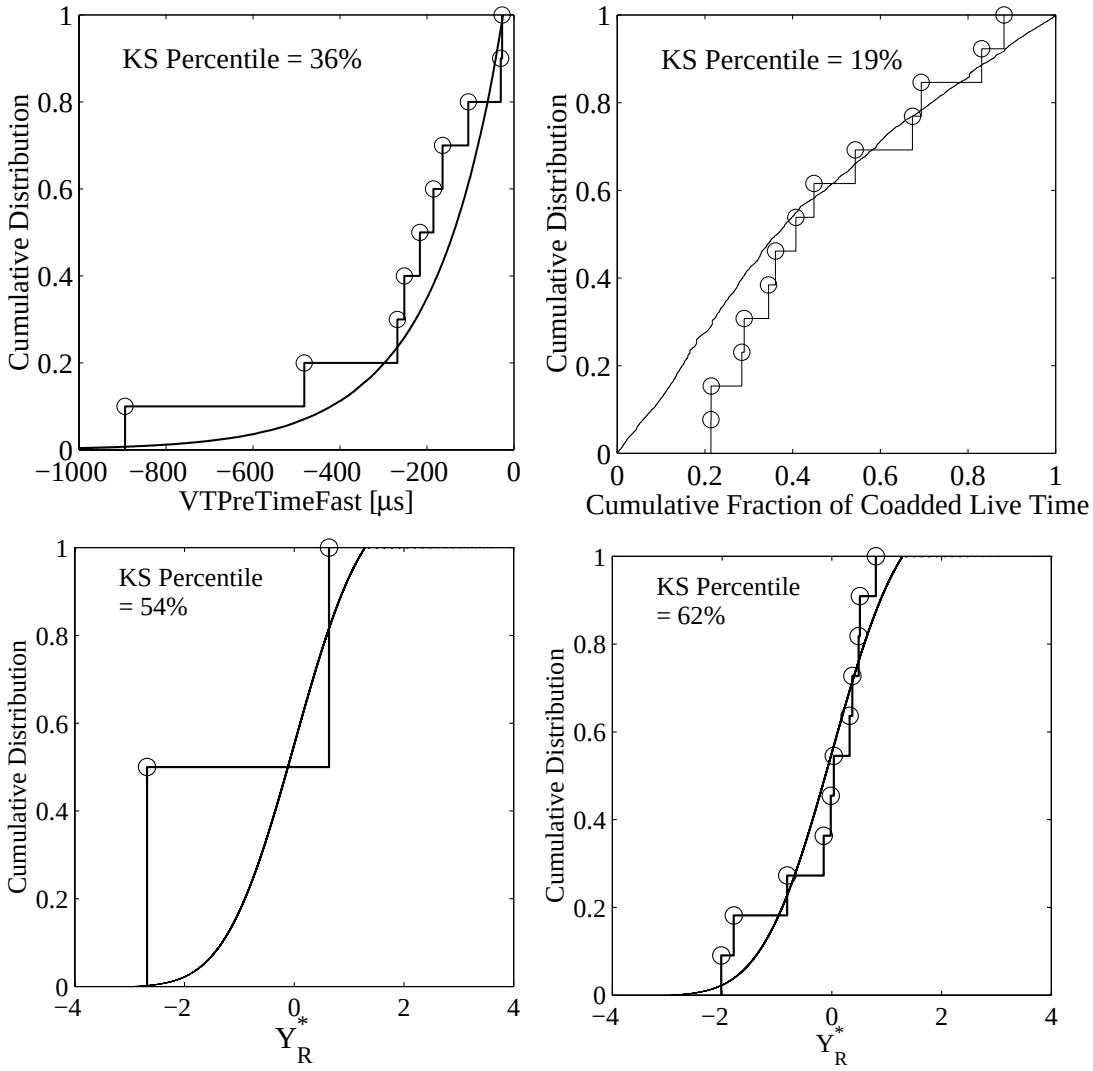

Figure 8.17: Kolmogorov-Smirnov tests for veto-anticoincident cQinOnly single-scatter nuclear-recoil candidates. "KS Percentile" indicates the percentage of experiments expected to observe larger deviation (larger KS statistic). Solid lines indicate expected distributions. Upper left: last veto-trigger-time distribution for ionization triggers. Upper right: cumulative coadded live time distribution. Lower left: pre-April 3 Y_R^* distribution. Lower right: post-April 3 Y_R^* distribution.

however, one must choose a weighting in E_R to apply when combining the expected distributions from different E_R. As is seen later, the recoil-energy spectrum fits the spectrum expected for external-neutron scattering well, so this spectrum, corrected by the overall efficiency function, is used as a weighting. The spectra of calibration-source, internal, and external neutrons are not different enough for it to matter which one is used. It is important, however, that a reasonable spectral weighting be used because it determines the relative importance of the cuts at low and high Y_R^* in modifying the distribution. Also, it is necessary to do the test separately for events before and after the April 3 power outage because the nuclear-recoil bands change at this point. Figure 8.17 shows this test for the 13 nuclear recoils. The data and distributions for BLIPs 4, 5, and 6, have been combined, weighting by each detector's exposure. The KS tests indicate that 54% and 62% (pre/post-April 3) of experiments would observe distributions that deviate further from the expected distribution. This level of agreement is important because misidentified electron events would be expected to have a distribution either flat in Y_R or weighted toward high Y_R.

Similar tests could be done for almost any event parameter. The above tests are most sensitive to the possibility that the 13 nuclear-recoil candidates are not veto-anticoincident nuclear recoils. These results indicate there is nothing unusual about these events at a statistically significant level.

8.2.2 Double-Scatter Data Set

In addition to the cuts described in the previous chapter, the veto-anticoincident double-scatter data set is defined by the following restrictions:

- at least one detector out of BLIPs 4, 5, and 6 has an inner-electrode event;
- exactly two detectors out of BLIPs 4, 5, and 6 trigger;
- and the recoil energies of both triggered detectors are between 10 and 100 keV.

The detailed rationale for these cuts was discussed in Section 7.5. Broadly, the purpose of these restrictive cuts is to discard events that are not double-scatter neutrons while maximizing the efficiency for observing these neutrons; a similar set of cuts has been made on the neutron-calibration and muon-coincident data sets and the resulting efficiency reductions are discussed in Chapter 7. These cuts discard triple-scatter neutrons, but the rate of such events relative to double scatters is only about 10%, so they contribute little to the analysis.

Ionization yield vs. recoil energy plots for this event set are shown in Figure 8.18. The events are distributed among the different ionization-partition cuts, so the summed-ionization signal is used for calculating the ionization yield to put all the events on the same footing; the added noise is not important here. BLIP3 is not shown because, by definition of the event set, it has no events. The cleanliness of these plots is remarkable. There are 8 scatters tagged as nuclear recoils, which appear in pairs. This last fact is quite important — there are no double-scatter events for which one recoil is a nuclear recoil and the other is not.

To better understand the physics of the double-scatter event set, plots of ionization yield vs. ionization yield ("Y_R^2" plots) for detector pairs are shown in Figures 8.19 and 8.20. The first figure shows double scatters among BLIPs 4, 5, and 6 only, with the aforementioned cuts. The second figure shows double scatters between BLIP3 and one of the other detectors, with the same

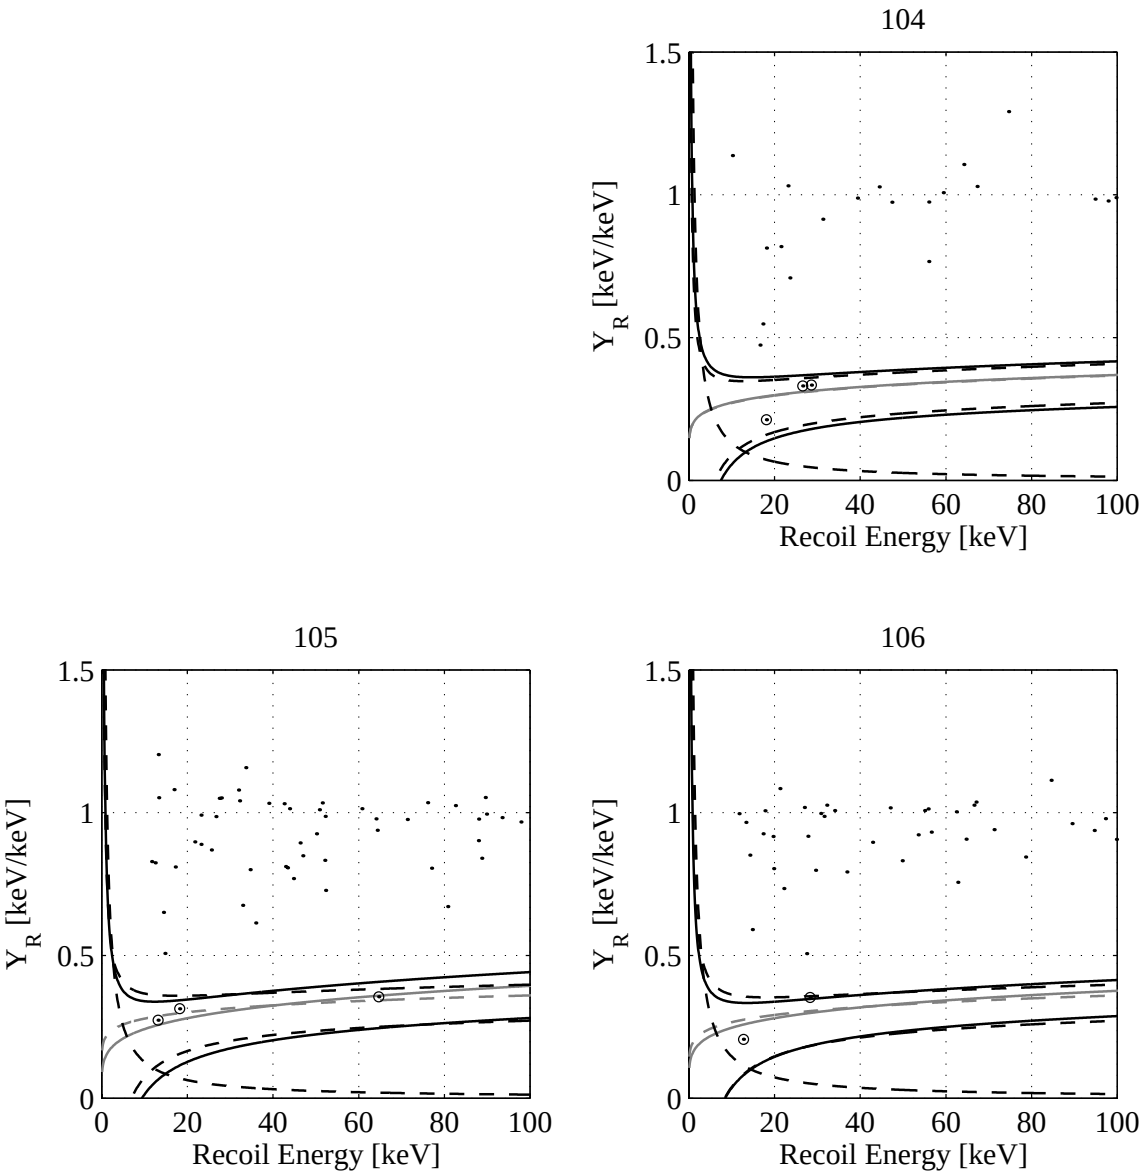

Figure 8.18: Ionization yield vs. recoil energy for veto-anticoincident BLIPs 4/5/6 double scatters in which at least one detector of BLIPs 4, 5, and 6 contains a cQinOnly event and both recoils are between 10 and 100 keV. Nuclear-recoil candidates are circled. Hyperbolic dashed line: ionization-search threshold. Light solid line: center of nuclear-recoil band, pre-April 3. Light dashed line: center of nuclear-recoil band, post-April 3. Dark solid line: nuclear-recoil-acceptance region, pre-April 3. Dark dashed line: nuclear-recoil-acceptance region, post-April 3.

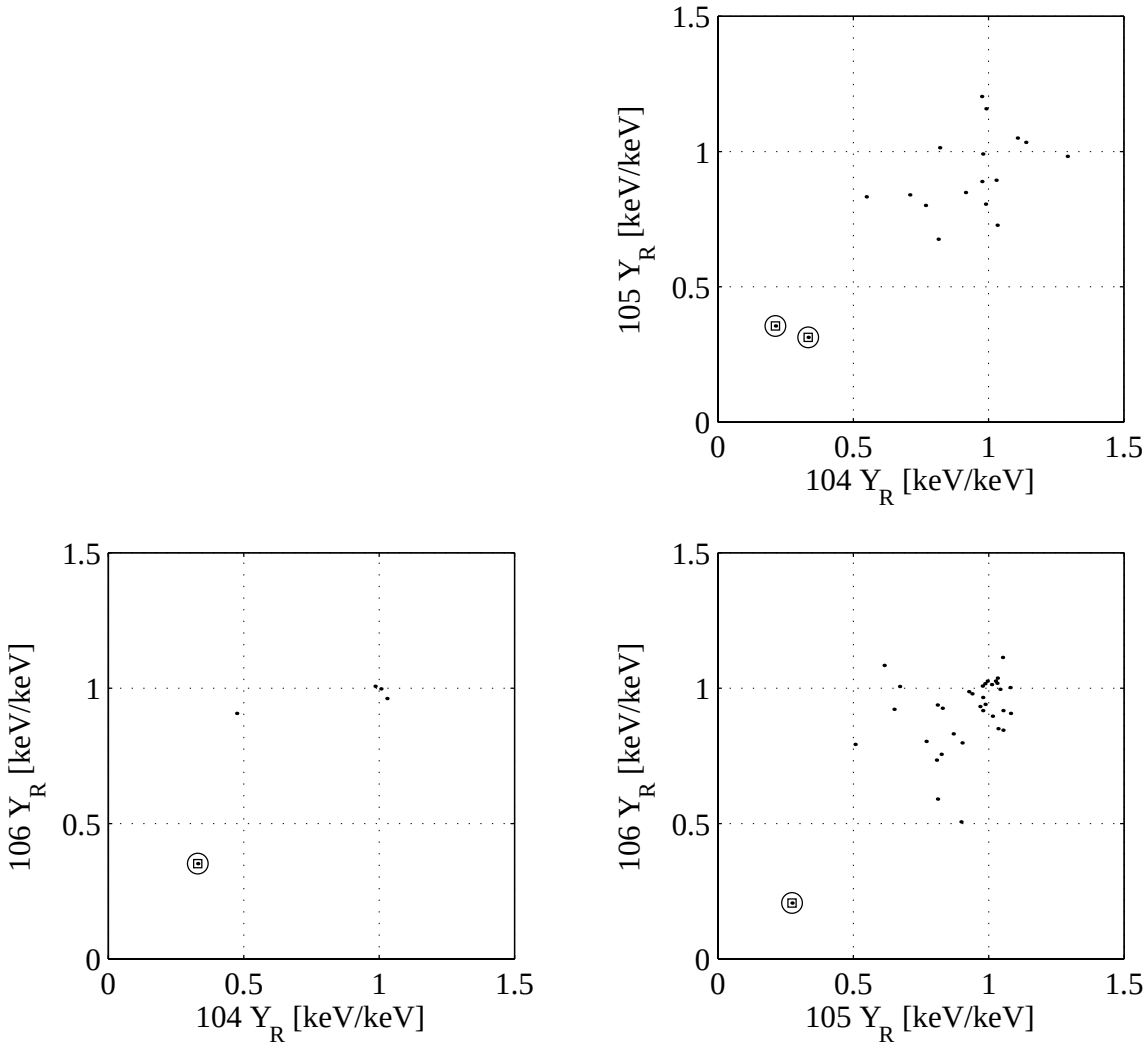

Figure 8.19: Ionization yield vs. ionization yield for veto-anticoincident BLIPs 4/5/6 double scatters in which at least one detector of BLIPs 4, 5, and 6 contains a cQinOnly event and both recoils are between 10 and 100 keV. Circles and squares indicate events identified as nuclear recoils by the x-axis and y-axis detectors, respectively.

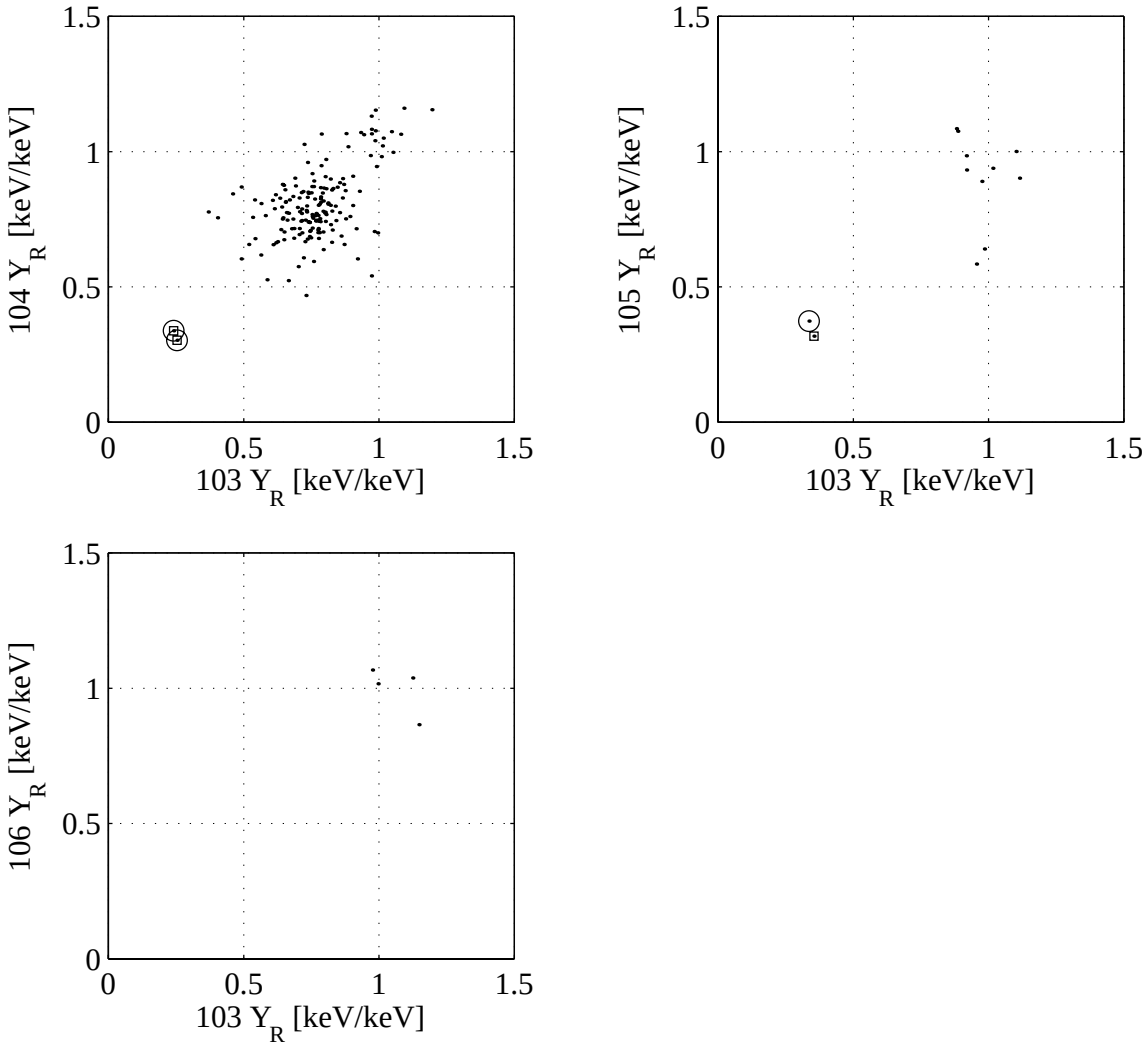

Figure 8.20: Ionization yield vs. ionization yield for veto-anticoincident double scatters in which BLIP 3 triggers, at least one detector (of all four) contains a cQinOnly event, and both recoils are between 10 and 100 keV. Circles and squares indicate events identified as nuclear recoils by the x-axis and y-axis detectors, respectively. The event marked by a circle and not a square is just above the nuclear-recoil-acceptance region of BLIP3, consistent with the less-than-unity acceptance of the cut. A similar explanation holds for the event marked by a square and not a circle. Note that the full electron-calibration data set discussed in Chapter 4 corresponds approximately to the BLIPs 3/4 plot, with the additional inclusion of events in which at least one detector has a cQShare event.

cuts except that BLIP3 is required to belong to each double scatter. (At least one detector is required to have a cQinOnly recoil, as above). The two sets of plots thus show all double-scatter events in which *any* detector has a cQinOnly recoil and both recoils are between 10 and 100 keV.

Consider first detector pairs that are not neighbors — BLIPs 3/5, BLIPs 4/6, etc. With the exception of 2 events in the BLIPs 3/5 pair, all such events are seen with similar Y_R in both detectors at either $Y_R \sim 1$ or $Y_R \sim 0.3$. These events are therefore unambiguously identified as either double-scatter penetrating photons (*i.e.*, bulk electron recoils) or double-scatter neutrons, respectively. When adjacent detector pairs are considered, low-yield events appear between the bulk electron recoils and neutrons. If an event appears at low yield in both detectors, it is most easily interpreted as a low-energy electron. If an event appears at low yield in one detector but at $Y_R \sim 1$ in the other detector, the event may be one in which a photon scatters very near the surface of one detector, ejecting a low-energy electron that hits the adjacent detector. This seems an unlikely explanation of the BLIPs 3/5 events of this type, though.

The BLIPs 3/4 plot is striking because it clearly shows a population of events that appear at low yield in both detectors and that is separated from both the double-scatter bulk electron recoils and the double-scatter neutrons. These events are interpreted as double-scatter low-energy electrons emitted from the contaminant on BLIP3 that scatter in both BLIPs 3 and 4. The absence of a similar population in the BLIPs 3/5 and BLIPs 3/6 plots supports this interpretation — low-energy electrons emitted by BLIP3 cannot traverse BLIP4 to reach BLIPs 5 and 6.

The most important conclusion to be drawn from these plots is that the identification of double-scatter nuclear-recoils, once the cuts listed previously are made, is umambiguous, with 4 double-scatter nuclear recoils appearing in the BLIPs 4/5/6 data set. To establish this point numerically, the BLIPs 3/4 "electron-calibration" data set is used. As is clear from the Y_R^2 plots, the BLIPs 3/4 sample contains an electron source. These events can be used as a low-statistics electron calibration. It is important to note that this electron-calibration data set is completely disjoint from the BLIPs 4/5/6 single- and double-scatter data sets. The analysis of this electron calibration is presented in Section 4.5.3 and the estimated value of the electron-misidentification parameter, β_β, for various ionization-partition cuts (cQinOnly, cQShare, and cQoutOnly) and energy bins are shown in Table 4.6. In using the electron calibration to estimate the number of double-scatter nuclear-recoil candidates arising from misidentified electrons, it is important to make use of the fact that, while the double-scatter electrons do cluster around $Y_R \sim 0.75$, there appears to be no correlation in the deviations from this central value of the ionization yields observed in the two detectors. This is clear from the BLIPs 3/4 plot in Figure 8.20 — the electron events do not form a line with slope 1. In order to be misidentified as a double-scatter neutron, a double-scatter electron must therefore be misidentified in *both* detectors; such misidentification is suppressed by a factor β_β^2 rather than only β_β.

This last point is nontrivial: one might have assumed that there would be significant correlation because the penetration depth of the electron would scale with its energy, and thus a very low-energy electron (say, 15 keV) would appear at lower yield in both interactions than a higher-energy electron (say 50 or 60 keV). There may be such a correlation, but the BLIPs 3/4 plot indicates that, when averaged between 10 and 100 keV using the energy spectrum of electrons emitted by the BLIP3 contaminant, the correlation becomes negligible. Furthermore, Figure 4.53 indicates that the bulk of the events in the electron-calibration data set are at low recoil energies, below 30 keV, and thus the electron calibration probes the possibility of correlation at low energies,

where it would be most likely. Clearly, correlation is not important.

There may be a straightforward explanation for the apparent lack of correlation. Because electrons are charged, they lose energy continuously rather than in a pointlike interaction, as is typical for photons. Furthermore, because electrons are so light, they experience significant deflection on each scattering — their path in the detector is not a straight line but is very jagged. A double-scatter electron is so deflected that it exits the crystal again. The deposited energy is therefore not a strong function of the electron energy — it depends on the track length in the crystal, which may be short for a high-energy electron if it is backscattered. The ionization yield is, however, likely to be well correlated with the track length: shorter tracks are also likely to be more shallow. Thus, for double-scatter electrons, the ionization yield for one scatter, while likely correlated with the deposited energy, may not be a good predictor of the actual electron energy, and thus may not be a good predictor of the ionization yield observed in the second recoil.

Numerically, the expected number of misidentified double-scatter electrons in the double-scatter nuclear-recoil sample is calculated as follows. Let N_l be the number of double-scatter electrons misidentified by both detectors and let N_β be the number of double-scatter electrons correctly identified by both detectors. The latter is the number of events that appear between the nuclear-recoil-acceptance region and the bulk electron-recoil band *in both triggered detectors*. A correction must be made because, as is seen in Figure 8.19, some events appear as electrons in one detector and as bulk electron recoils in the second detector. These events are artificially included in the double-scatter-electron set; this is conservative because photon misidentification is far less likely than electron misidentification.

The formalism developed for calculation of the 90% CL upper limit on the expected number of misidentified single-scatter electrons can be used, with the modifications

$$\langle N_\beta^b \rangle = \lambda^2 \mu_l^b \tag{8.5}$$

and

$$\mathcal{L}(N_\beta^b, N_l^c, N_\beta^c | \mu_l^b, \mu_l^c, \lambda) = \frac{(\lambda^2 \mu_l^b)^{N_\beta^b} e^{-\lambda^2 \mu_l^b}}{N_\beta^b!} \frac{(\mu_l^c)^{N_l^c} e^{-\mu_l^c}}{N_l^c!} \frac{(\lambda \mu_l^c)^{N_\beta^c} e^{-\lambda \mu_l^c}}{N_\beta^c!} \tag{8.6}$$

where, essentially, λ has been replaced by λ^2 everywhere in the background-data parts of the formulae. (Recall that $\lambda \equiv 1/\beta_\beta$.) This expresses the lack of ionization-yield correlation: if the expected number of double-scatter electrons misidentified as double-scatter neutrons is μ_l^b, then the expected number of correctly identified double-scatter electrons is $\lambda^2 \mu_l^b$.

In applying the above, one complication arises: the double-scatter electrons in the background set are distributed among the three ionization-partition cuts, so in principle one has to do a more complex analysis, dividing the data into different ionization-partition sets, using the appropriate calibration-data sets, and calculating the combined likelihood function. This results in many more nuisance parameters and becomes intractable. Instead, the following approximate approach is taken. As is shown in Table 8.4, the double-scatter background-data events are primarily cQinOnly events, as with the calibration data. They are as or more concentrated in the cQinOnly cut than the BLIP3 calibration data. Therefore, the limit is calculated using the background data and BLIP 3 calibration data summed over the three ionization-partition cuts. The BLIP4 calibration data are more concentrated in the cQinOnly cut and thus would yield too aggressive a limit. The BLIP3 calibration data are also expected to probe the bias polarity in which the ionization-collection tails are worse (see Section 4.2). A sum over energies is done — the background data appear to have a

Event set	N_β^b	N_l^c	N_β^c	$\mu_{l,90}^b$
10 – 30 keV				
Qin	$8/2 = 4$	1	46	0.08
QShare	$1/2 = 0.5$	1	35	0.04
Qout	$4/2 = 2$	0	9	1.00
Total	$13/2 = 6.5$	2	90	0.05
30 – 100 keV				
Qin	$14/2 = 7$	0	39	0.09
QShare	$3/2 = 1.5$	0	27	0.06
Qout	$2/2 = 1$	0	5	3.73
Total	$19/2 = 9.5$	0	71	0.03
10 – 100 keV				
Qin	$22/2 = 11$	1	85	0.05
QShare	$4/2 = 2$	1	62	0.02
Qout	$6/2 = 3$	0	14	0.44
Total	$32/2 = 16$	2	161	0.03

Table 8.4: Double-scatter-electron misidentification estimates, calculated by method discussed in text. The first column is the number of double-scatter-electron events observed for the given energy bin and ionization-partition cut. Since a given double scatter can be distributed among two different bins and cuts, this number is derived by dividing by 2 the number of *recoils* (not events) seen in the bin. Also listed are estimates combining the ionization-partition cuts for a given energy bin and combining energy bins for a given ionization-partition cut. The number on which to focus is the one at the lower right, which gives the expected number of double-scatter-electron events misidentified as double-scatter nuclear recoils, summing over all ionization-partition cuts and energy bins. Note that the estimates are, in many cases, significantly higher than one would find by calculating a Bayesian 90% CL upper limit on β_β from the calibration data and estimating the leakage by $N_\beta^b \beta_\beta^2$; this occurs because N_β^b is small and its fluctuations significantly weaken the final upper limits.

spectrum no softer than the calibration data. For completeness, limits are also calculated for each electrode cut and each energy bin separately. Since a given double-scatter event may appear in different energy bins and ionization-partition cuts in the two detectors, some prescription must be used in determining the "number of double scatters" in a given energy bin and ionization-partition cut. The prescription used here is to divide the number of recoils (not events) for each cut and energy bin by 2. This is exactly correct for the sum over all ionization-partition cuts and energy bins.

The double-scatter-electron misidentification limits are shown in Table 8.4. Clearly, the limits are on the whole very stringent except in cases where the number of electron-calibration events is low. To conclude, the double-scatter data set defined here shows that double-scatter neutrons are unambigously identified and that the expected misidentification of double-scatter electrons as double-scatter neutrons is negligible. Four double-scatter-neutron events are seen.

8.3 The Neutron Interpretation

Thirteen single-scatter and four double-scatter nuclear-recoil candidates are observed in the veto-anticoincident data set. It has been argued that these events are indeed nuclear recoils, not low-energy-electron surface events. The only possible particles that can cause these recoils are neutrons and WIMPs. Given the 10.6 kg d exposure, this number of single-scatter nuclear recoils corresponds to an integral event rate of about 1.3 kg^{-1} d^{-1} above 10 keV. This event rate approaches interesting MSSM WIMP models and is in the neighborhood of the rate expected from a WIMP with mass and cross section yielding the annual-modulation signal observed by the DAMA collaboration [88].

However, there are a number of reasons to believe that these events are due to neutrons, not WIMPs. First, the observation of 4 double scatters rules out WIMPs as the source of these events: since the WIMP single-scatter rate is so low ($\sim$1 kg^{-1} d^{-1} and lower), WIMPs, for all practical purposes, do not multiply scatter. Put in more practical terms, the mean free path of a WIMP with a 1 kg^{-1} d^{-1} event rate in germanium is about 10^{10} meters of germanium. Second, a WIMP interpretation is further disfavored by the observation of 4 nuclear recoils in the Run 18 data set obtained with a Si ZIP detector. If WIMPs caused these four events, then an unacceptably high number of WIMP events are expected in the Run 19 Ge data set. Finally, via the simulations discussed in Chapter 3, it is seen that the 13 single-scatter Ge events, 4 double-scatter Ge events, and 4 Si events are consistent (though not impressively so) with a neutron-background interpretation.

The most straightforward interpretation is therefore that these events are due to neutron scattering. These neutrons must be veto-anticoincident, so they are most probably neutrons (and their secondaries and tertiaries) generated by fast muons in the tunnel walls. The expected rates and recoil-energy spectra for these "external" neutrons are discussed in Chapter 3. In this section, I present in detail the argument for this neutron interpretation.

8.3.1 Run 18 Si Detector Nuclear-Recoil Event Rate

In CDMS Run 18, which occurred during 1998, 33 live days of data were taken with a 100-g Si ZIP detector ("Alex," after the first prime minister of Canada). The Si run yields a 1.6 kg d exposure after cuts. Four nuclear recoils are observed in the Si data set. Based on a separate electron calibration, the upper limit on the expected number of misidentified surface events is 0.76 events (90% CL). The Run 18 veto-anticoincident data are displayed in Figure 8.21.

The ZIP Detector and Data Set

A brief digression to discuss this detector and data set is warranted. Z-sensitive Ionization and Phonon-mediated (ZIP) detectors sense the *athermal* phonons produced by a recoiling particle. The Run 18 Si detector employs Schottky-barrier contacts for the ionization measurement, which are briefly discussed in Chapter 4. The detector geometry is similar to the BLIP geometry: the detectors are cylindrical, 7.5 cm in diameter and 1-cm thick. The ionization electrodes are segmented radially with a nominal 0.85:0.15 inner/outer volume ratio. The detector edges are not curved and the cylindrical detector wall is left bare. The athermal phonons are collected and sensed by Quasiparticle-Trap-Assisted Electro-Thermal-Feedback Transition-Edge Sensors (QTA-ETF-TES, or QET for short) deposited on one crystal face. Though the Schottky-barrier contacts

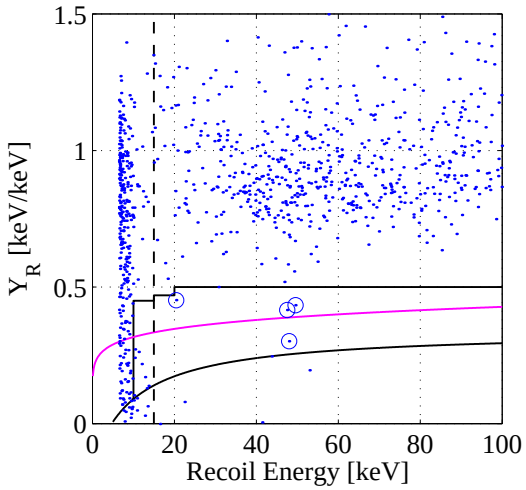

Figure 8.21: Run 18 Si ZIP detector veto-anticoincident data. Light solid line: center of nuclear-recoil band. Dark solid lines: nuclear-recoil-acceptance region. Dashed line: nuclear-recoil analysis threshold (15 keV). Figure provided by R. Clarke/R. Schnee.

also have a dead-layer problem (see Section 4.2), the athermal-phonon signal allows rejection of dead-layer events because this pulse shape is sensitive to the proximity of an interaction to the detector surface. A 90% CL lower limit on surface-event rejection of 98.8% above 15 keV was established by a laboratory exposure of the Run 18 Si ZIP detector to an electron source. Thus, this detector is capable of discriminating nuclear recoils from dead-layer events in the presence of a low-energy-electron surface-event rate of 60 kg^{-1} d^{-1} between 15 and 100 keV. An upper limit of 0.76 misidentified surface events is derived from this event rate and the surface-event-rejection lower limit. It should be noted that, above 20 keV, the surface-event-rejection lower limit is 99.7% and the expected number of misidentified surface events is only 0.26 events. A value of 0.76 events is used in calculation of exclusion limits to be conservative.

The ZIP detector technology has been developed by Blas Cabrera's group at Stanford over the last few years. The electron calibration and Run 18 data set is described in detail in Roland Clarke's dissertation [102]. (The detector is referred to as a Fast, Large Ionization- and Phonon-mediated (FLIP) detector in this reference.) The electron calibration and surface-event rejection are also described in [144]. A paper in preparation will present the Run 18 Si ZIP data and analysis [103]. The baseline detector for CDMS II is a ZIP detector using the above athermal-phonon sensing technology and the α-Si/Al-Schottky ionization contacts discussed in Section 4.2. Three silicon and three germanium ZIP detectors are being operated in CDMS Run 20, which began in March, 2000. Physics results from this run are not yet available.

Implications of the Silicon Detector Data

The four nuclear-recoil events observed in the Run 18 Si ZIP data cannot be WIMPs: whether their interactions with target nuclei are dominated by spin-independent or spin-dependent couplings, WIMPs yielding the observed Si nuclear-recoil rate would cause an unacceptably high

ratio	simulated	observed
$N_d/(N_s + N_d)$	0.06	0.24
$N_{Si}/(N_s + N_d)$	0.20	0.24

Table 8.5: Ratios of numbers of events predicted by external-neutron-background simulation.

number of nuclear recoils in the Ge data set. As discussed in Chapter 2, the WIMP-nucleus cross section scales as A^2 for WIMPs with spin-independent interactions. Expected recoil-energy spectra in germanium and silicon for a WIMP with spin-independent interactions are shown in Figure 2.10. Ge and Si differ by a factor of 5 to 7 in differential rate between 0 and 100 keV. One thus expects of order 30 to 50 times the number of WIMPs in the Ge data set as in the Si data set given the exposures and this rate ratio. The argument is more complicated for spin-dependent interactions, but it also holds that there should be many more nuclear recoils in the Run 19 Ge data set than are observed. Furthermore, the spin-dependent cross section corresponding to the observed Si event rate is significantly larger than expected from the MSSM. As is shown below, the observation of 4 Si nuclear recoils is consistent with the external-neutron-background interpretation of the 13 single-scatter and 4 double-scatter nuclear recoils in the Ge data set.

8.3.2 Comparison of Observed and Simulated Rate Ratios

The simulation of the external-neutron background is discussed in Chapter 3. As indicated there, a number of issues lead to a large normalization uncertainty, making it difficult to accurately predict the absolute flux of unvetoed external neutrons. However, normalization-independent predictions of the simulation, such as relative rates of single scatters and double scatters, relative rates in Si and Ge detectors, and the shapes of nuclear-recoil spectra, are accurately simulated. Therefore, only these normalization-independent quantities are used to test the consistency of the observed data with the external-neutron-background simulation. It is important to note that these predictions of the simulation are also insensitive to reasonable changes in the external-neutron energy spectrum. In Chapter 3, it is seen that the high-energy tail of the external-neutron spectrum has little effect on the spectrum of neutrons incident on the detectors that dominate the detector event rate — high-energy external neutrons produce many secondary and tertiary neutrons with energy spectra similar to internal neutrons, and the neutron-interaction cross section strongly favors interactions of low-energy neutrons.

The ratios of Ge double scatters N_d to Ge single scatters N_s and of Si neutron events N_{Si} to Ge single scatters for the data and the simulation are shown in Table 8.5. Schematically, the data and simulation can be compared in two ways: normalizing the simulation by the neutron-background rate that best fits N_s, N_d, and N_{Si} jointly, or normalizing by the neutron-background rate that best fits N_d and N_{Si} and predicting N_s. The latter is the intuitive interpretation of using the Ge doubles and Si events to predict the neutron background in the Ge singles set. These comparisons are shown in Figure 8.22.

More rigorously, a likelihood-ratio test can be used to compare the hypothesis that the three numbers of events are derived from a single neutron background with relative rates given by the simulation to the hypothesis that the three event sets arise from three different neutron-

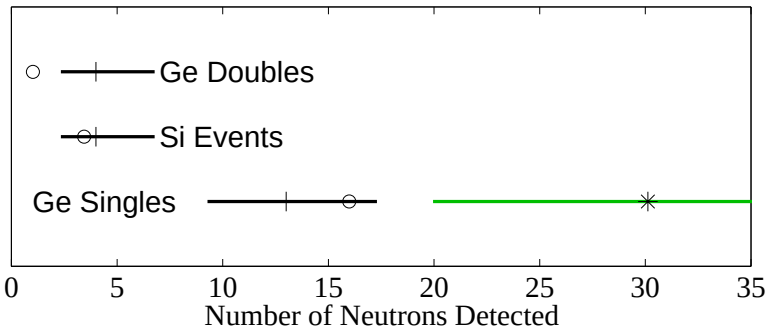

Figure 8.22: Schematic comparison of simulated and observed event ratios. The vertical marks and dark horizontal lines indicate the observed numbers of events and Feldman-Cousins 68% CL confidence intervals. The circles indicate the numbers of events expected based on normalizing the simulation by the neutron background that best fits N_s, N_d, and N_{Si} jointly. The $*$ and light line indicate the expected N_s and corresponding 68% CL confidence interval based on the neutron background that best fits N_d and N_{Si} jointly. Figure provided by R. Schnee.

background sources. Effectively, the latter hypothesis corresponds to three arbitrary background sources for the three event types, the most general possible hypothesis. The likelihood ratio of the two hypotheses is

$$R = \frac{\mathcal{L}(N_s, N_d, N_{Si}|\widehat{n})}{\mathcal{L}(N_s|\widehat{n}_1)\mathcal{L}(N_d|\widehat{n}_2)\mathcal{L}(N_{Si}|\widehat{n}_3)} \tag{8.7}$$

where $\mathcal{L}(N_s, N_d, N_{Si}|\widehat{n})$ is the likelihood of the three observed values calculated at the number of Ge (single + double) nuclear recoils $\widehat{n}$ that maximizes their likelihood, and $\mathcal{L}(N_s|\widehat{n}_1)$, $\mathcal{L}(N_d|\widehat{n}_2)$, and $\mathcal{L}(N_{Si}|\widehat{n}_3)$ are the individual likelihoods of the three observed values separately, calculated at the values of the number of Ge events that maximize each likelihood separately. R is limited from above by 1, which occurs when the best fit values $\widehat{n}_1$, $\widehat{n}_2$, $\widehat{n}_3$, and $\widehat{n}$ are all equal. The larger R is, the more consistent the data are with the simulation. In the limit of large numbers of events N_s, N_d, and N_{Si}, $-2\log R$ is distributed as χ^2 for 2 degrees of freedom. With the small numbers of events observed, a Monte Carlo must be performed to determine the expected distribution of R (or $-2\log R$), which then yields the fraction of experiments that would have yielded lower (worse) values of R. Since the true neutron background is unknown, the Monte Carlo must be done for different values of the "true" number of Ge events n that is assumed for the Monte Carlo. In doing the Monte Carlo, the assumed value n is used to calculate expected values μ_s, μ_d, and μ_{Si} for the numbers of each type of event. The Monte Carlo draws observed numbers of events N_s, N_d, and N_{Si} from Poisson distributions for each quantity given the expected values μ_s, μ_d, and μ_{Si}. The ratio R is calculated for each Monte Carlo outcome, yielding the distribution of R for that assumed n. Note that $\mathcal{L}$ is calculated at $\widehat{n}$, not at the assumed value n, for reasons explained below.

Table 8.6 shows the fraction of experiments with worse values of $-2\log R$ as a function of the assumed value of n. The advantage of testing the above likelihood *ratio* is that the calculated confidence level of the observed data is only weakly dependent on the assumed neutron background n — the test does not depend strongly on what normalization for the external-neutron background is assumed for the Monte Carlo. This characteristic results partially from the use of a likelihood ratio

assumed n	$\log(\mathcal{L}_d(n))$	$P(\mathcal{L}(n) < \mathcal{L}_d(n))$	$P(-2\log R > -2\log R_d)$
9.0	-12.06	0.001	0.072
14.0	-8.78	0.029	0.060
17.5	-8.30	0.062	0.061
21.6	-8.80	0.054	0.064
25.0	-9.81	0.026	0.059
29.0	-11.98	0.006	0.062

Table 8.6: Likelihood-ratio test of the consistency of the observed nuclear-recoil event sets with the simulation. The first column is the likelihood of the data, $\mathcal{L}_d = \mathcal{L}_d(N_s, N_d, N_{Si}|n)$, given the assumed neutron background n. The second column shows the fraction of simulated experiments at the assumed value of n that yield likelihoods lower than the data. The last column shows the fraction of simulated experiments with worse values of the likelihood ratio R than the value given by the data R_d (*i.e.*, with $-2\log R > -2\log R_d$). This last column indicates that the likelihood-ratio test has only weak dependence on the assumed value n. Thus, the result, that about 6% of simulated experiments yield worse agreement with the neutron-background simulation than the data, is robust against uncertainty in n.

rather than, for example, just the likelihood function in the numerator. The lack of dependence on n also arises from the calculation of the likelihoods at the "best-fit" neutron backgrounds $\widehat{n}_1$, $\widehat{n}_2$, $\widehat{n}_3$, and $\widehat{n}$ rather than using the assumed value of n. The conclusion is that the fraction of experiments that would have yielded worse (higher) values of $-2\log R$ is 6%. This level of consistency with the simulation is unpleasantly low. The poor agreement is due to the high number of multiples: the numbers of Ge singles and Si events yield an expectation of only about 1 Ge double scatter, while 4 are observed. However, based on the arguments made in Section 8.2.2, it is extremely unlikely that the number of Ge doubles is somehow overestimated. Little more can be said on the subject, except that maybe we just got lucky (or unlucky, depending on one's point of view).

Finally, the $\log\mathcal{L}$ data shown in Table 8.6 indicate the background normalization that best fits the observed N_s, N_d, and N_{Si} is $\widehat{n} = 17.5$ events. The exposure is 10.6 kg d, so this corresponds to a total (single + double) 10-to-100-keV neutron-background rate of approximately 1.6 kg^{-1} d^{-1}. However, the maximum in $\log(\mathcal{L})$ is wide. The values at which $\log(\mathcal{L})$ drops by 0.5, corresponding asymptotically to the 68% CL allowed region, are 14 and 21.6 events, so the neutron background is determined to a precision of about $\pm 20\%$.

8.3.3 Comparison of Observed and Simulated Spectral Shapes

The other normalization-independent quantity available from the simulation is the shape of the recoil-energy spectrum. A Kolmogorov-Smirnov test comparing the integral observed and simulated nuclear-recoil spectra is shown in Figure 8.23. Being a KS test, the test is independent of normalizations. The simulated spectrum has been corrected by the energy dependence of the overall efficiency function. The KS tests for the Ge single-scatters alone and for all Ge events are shown. The fraction of experiments expected to give larger deviations is 86% and 28% in the two cases, respectively. These levels of agreement are generally considered reasonable and

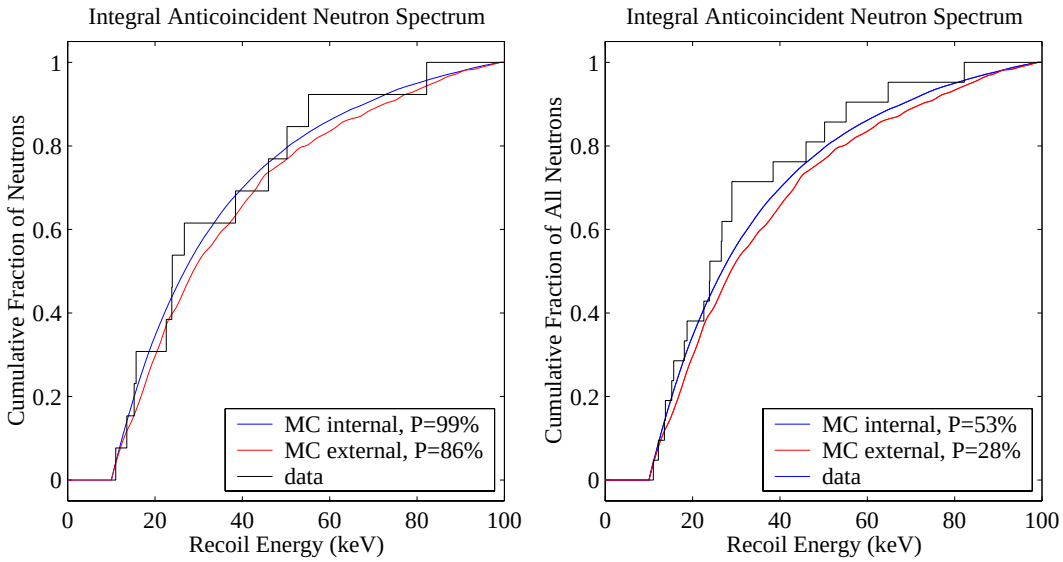

Figure 8.23: Kolomogorov-Smirnov tests comparing observed Ge single-scatter (left) and Ge all-event (right) integral recoil-energy spectra to efficiency-corrected expected spectra from the external-neutron (bottom smooth lines) and internal-neutron (top smooth lines) simulations. Figures provided by R. Schnee.

indicate the observed and simulated spectral shapes are consistent. These results should only be taken as support for the consistency of the data with the simulation; they are also consistent with interactions of a WIMP having mass in the 50 to 100 GeV c^{-2} range and so do not alone disfavor a WIMP interpretation. Also shown are the KS tests comparing the data to the *internal*-neutron simulation. It is again clear that the recoil-energy spectrum expected in the detectors is fairly independent of the high-energy tail of the external-neutron spectrum. The consistency of the data with the recoil-energy spectrum expected for external neutrons can also be seen in a differential recoil-energy spectrum, shown in Figure 8.24.

8.3.4 On the Possibility of Veto Inefficiency

The possibility that the observed neutron background is due to internal neutrons that appear as veto-anticoincident because of veto inefficiency — *i.e.*, the muon was missed by the veto — should be considered. As discussed in Chapter 3, it is not possible to directly measure the veto efficiency for muons that generate neutrons, but it is possible to measure the efficiency for detector-tagged muons. Figure 3.10 demonstrates that the veto pulse-integral spectra for muons that generate neutrons and for detector-tagged muons are very similar. The veto efficiency for detector-tagged muons had an average value of 99.90% over the run, with a worst excursion to 99.69%, as is shown in Figure 8.25. Using the single-scatter cQinOnly muon-coincident-neutron rate of 22 per live day (coadded over BLIPs 3, 4, 5, and 6) shown in Figure 7.26, the total number of internal neutrons seen in BLIPs 4, 5, and 6 over the run is about 1200. The expected number of "missed" internal neutrons is thus about 1, much less than the observed number of Ge nuclear

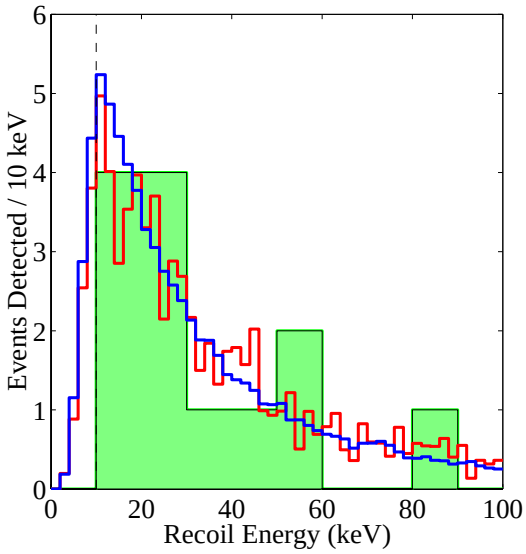

Figure 8.24: Comparison of observed single-scatter nuclear-recoil energy spectrum and expectations from simulation of external and internal neutrons. Data: wide bins. Light: external-neutron simulation. Dark: internal-neutron simulation. The simulated spectra have been efficiency-corrected. Dashed vertical line: 10-keV analysis threshold. Figure provided by R. Schnee.

recoils. However, one may invoke the possibility that internal neutrons may be produced by a subset of muons that pass through particular regions of the veto for which the efficiency is much less than shown above. Such a possibility cannot be ruled out.

However, in practice, it does not matter whether the neutrons are internal or external. It is seen in Chapter 3 that energy spectra of internal and external neutrons incident on the detectors are very similar. Therefore the expected recoil-energy spectra are essentially the same. This is corroborated by the fact that the KS test comparing the expected internal- and external-neutron recoil-energy spectra cannot distinguish the two cases. Furthermore, because the double-scatter/single-scatter and Ge/Si ratios are determined completely by the incident-neutron spectrum, they also are the same to good precision for internal and external neutrons. Thus, because of this degeneracy between the two types of neutrons, it is almost irrelevant whether the neutrons causing the observed nuclear recoils are internal or external. One issue arises, which is that the veto efficiency may have been different for the Run 18 Si data. However, the veto efficiency for detector-tagged muons during Run 18 was $> 99.99\%$ [8], higher than during Run 19, so the Si data presumably can only underestimate the neutron background, making the joint estimate lower. This yields an overly conservative exclusion limit on the WIMP-nucleon cross section.

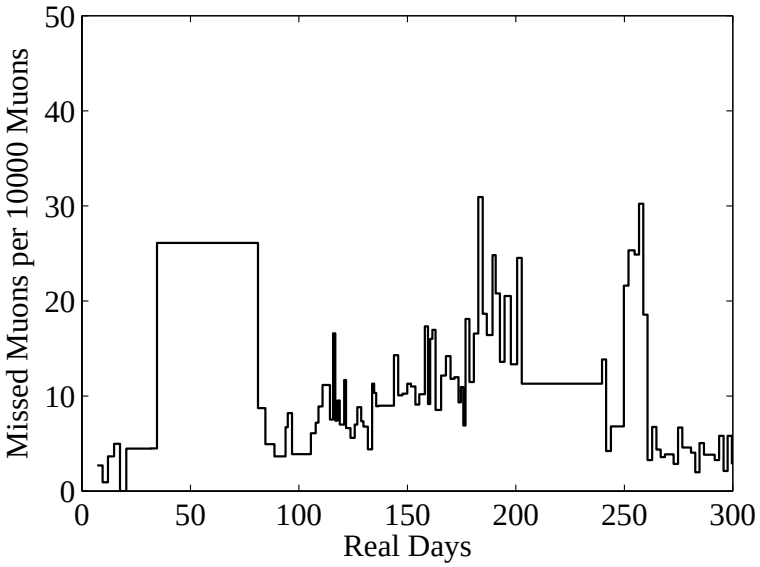

Figure 8.25: Veto inefficiency for detector-tagged muons over Run 19. The number of detector-tagged muons missed by the veto per 10000 detector-tagged muons is shown. The time bins do not all have the same width because they are defined by containing 10000 muons, not a fixed amount of real time. Variations in the bin width (especially the large ones) are due to intervals when the experiment was not operating. The mean veto efficiency is 99.90% and lowest veto efficiency observed is 99.69%. Figure provided by R. Schnee.

8.4 Exclusion Limits on the WIMP-Nucleon Cross Section

A 90% CL exclusion region in the plane of WIMP mass M_δ and the WIMP-nucleon spin-independent elastic-scattering cross section $\sigma_{\delta \mathrm{n}}$ is determined under the assumption that the neutron-background interpretation of the data is correct. Heuristically, this permits the neutron background predicted by the Ge double-scatter and Si event rates to be statistically subtracted from the Ge single-scatter data, yielding sensitivity to lower WIMP-nucleon cross sections by accounting for the component of the Ge single-scatter data set that is due to neutron scattering.

In practice, a more rigorous method, involving a modification of the Feldman-Cousins "unified" approach [83], is used to calculate the allowed region. The Feldman-Cousins method is most easily understood by specific examples, as are given in the original paper [83]. The Feldman-Cousins approach offers a number of desirable features, whose validity will not necessarily be obvious to the reader immediately, but that will be reiterated later in the context of its application here:

- The first feature is that it is a classical method — in calculating an excluded region of parameter space at, for example, 90% CL, it answers the question "What are the values of the input parameters such that, for those input parameters, the observed data in at least 90% of experiments would have been more consistent with the parameters than the observed experimental data?" This question is independent of the integration measure assumed on the parameter space being constrained and so is usually considered the best question to ask when presenting experimental results. What is meant by "consistent" is ambiguous; Feldman and Cousins use a likelihood ratio,

$$R = \frac{\mathcal{L}(X|\mu)}{\mathcal{L}(X|\widehat{\mu})} \tag{8.8}$$

  where μ are the set of input parameters being tested, X are the experimental data, and $\widehat{\mu}$ are the set of parameters that maximize the likelihood of ("best fit") the experimental data. Use of the likelihood ratio is appropriate because the absolute likelihood of the data has little meaning. For example, consider two experiments that measure the rate of some occurrence, but allow the first experiment to have a much larger amount of data; thus, the statistical uncertainty of the result measured by the first experiment is much smaller than that of the second experiment. Even if both experiments yield the same measured rate and the measured rate coincides exactly with the rate expected for some μ, the first experiment has a much larger $\mathcal{L}$ because its statistical uncertainty is smaller and so its likelihood function is narrower and has a larger peak value. The two experiments are equally valid, but the latter has a much smaller $\mathcal{L}$. Use of a likelihood ratio normalizes the likelihood by the highest possible likelihood the experiment could yield, thus putting the two experiments on the same footing.

- A second advantage of the Feldman-Cousins approach is that it is "unified" — its results smoothly transition between one-sided upper limits and two-sided confidence intervals. This eliminates the ambiguity as to whether to quote an upper limit or a confidence interval.

- A third advantage is that it ensures full coverage, and possibly overcoverage. This feature is particularly important when considering cases involving background subtraction. Various methods previously in use could yield upper limits or confidence intervals that do not meet their advertised confidence levels in a classical sense. In extreme cases, such methods could yield an empty set for the allowed region.

The analysis presented here is somewhat more complicated than the examples discussed in [83] because the prediction of the neutron background is subject to significant uncertainty. The normalization of the neutron background is thus taken as an additional free parameter, n, the total number of Ge neutron scatters, that is constrained by the observed relative rates of Ge single scatters, Ge double scatters, and Si events. The Feldman-Cousins method has been extended in order to project out the neutron-background "nuisance" parameter and give a final confidence region in the $(M_\delta, \sigma_{\delta \mathrm{n}})$ plane. This extension was developed and implemented by Bernard Sadoulet, Richard Schnee, and Steve Yellin and is described in detail in an internal note by Richard Schnee [145]. I describe the formalism briefly, but the reader should refer to this reference for a comprehensive discussion, including important practical aspects of the implementation. I describe the method in two steps: first, the naive application of the Feldman-Cousins method to the CDMS data to make clear the link to [83]; and then the modifications made to improve the robustness of the method for analyses with nuisance parameters.

The observables are the energies of the Ge single scatters E_i, the number of Ge double scatters N_d, and the number of Si events N_{Si}. The energies of the Ge doubles and Si events are not used because, given the small numbers of such events, including them would significantly increase the computational requirements without significantly improving the accuracy of the result. I do not discuss how the 90% CL upper limit of 0.76 misidentified electrons in the Si event set is included; it can be put into the likelihood function, as is described in [145]. The free parameters for which a confidence region will be calculated are the WIMP mass M_δ, the WIMP-nucleon cross section $\sigma_{\delta \mathrm{n}}$, and the expected total number of neutron-scattering events in Ge, n. Naive application of the Feldman-Cousins method involves calculation of the likelihood ratio

$$R = \frac{\mathcal{L}(E_i, N_d, N_{Si} | M_\delta, \sigma_{\delta \mathrm{n}}, n)}{\mathcal{L}(E_i, N_d, N_{Si} | \widehat{M_\delta}, \widehat{\sigma_{\delta \mathrm{n}}}, \widehat{n})} \tag{8.9}$$

for a grid of *physically allowed* values of M_δ, $\sigma_{\delta \mathrm{n}}$, and n, where the numerator is the likelihood of the data given the values taken from the grid and the denominator is the likelihood of the data evaluated at the values $(\widehat{M_\delta}, \widehat{\sigma_{\delta \mathrm{n}}}, \widehat{n})$ that maximize the likelihood of the data. Because of the small numbers numbers of events and the physical boundaries, the expected distribution of R cannot be determined analytically (*i.e.*, it is not just distributed as a χ^2), so a Monte Carlo is performed at each grid point to determine this distribution. Of particular importance is the value $R_{90}(M_\delta, \sigma_{\delta \mathrm{n}}, n)$ such that 90% of simulated experiments at the grid point give $R > R_{90}$. The bigger R is, the more "likely" the particular outcome, so R_{90} is the lower limit of the 90% most "likely" outcomes. The 90% CL allowed region in the space of triples $(M_\delta, \sigma_{\delta \mathrm{n}}, n)$ is the set of all points for which $R_{data} > R_{90}$. The WIMP contribution to the expected spectrum is calculated using the formalism outlined in Chapter 2, normalized to the 10.6 kg d exposure, and corrected for the energy-dependent efficiency function shown in Figure 8.16.

The above procedure yields a confidence region in the three-dimensional space of $(M_\delta, \sigma_{\delta \mathrm{n}}, n)$. However, a confidence region in $(M_\delta, \sigma_{\delta \mathrm{n}})$-space alone is desired. Therefore, n must be projected out. The most obvious way to do this is to allow a grid point $(M_\delta, \sigma_{\delta \mathrm{n}})$ if there is any value of n for which the triple $(M_\delta, \sigma_{\delta \mathrm{n}}, n)$ is allowed. This makes the allowed region as large as it can possibly be, thereby providing the most conservative excluded region.

The primary disadvantage of the above naive method is that the statistic R clearly depends on the value of the nuisance parameter n used. This is undesirable since the nuisance parameter is projected out. The modification designed for this analysis, and more generally applicable to analyses

with nuisance parameters, is to use a different statistic and thereby reduce the dependence of R on n. The likelihood ratio is redefined to be

$$R = \frac{\mathcal{L}(E_i, N_d, N_{Si}|M_\delta, \sigma_{\delta\mathrm{n}}, \widehat{n}^*)}{\mathcal{L}(E_i, N_d, N_{Si}|\widehat{M_\delta}, \widehat{\sigma_{\delta\mathrm{n}}}, \widehat{n})} \tag{8.10}$$

where $\widehat{n}^*$ is the value of n that maximizes the likelihood of the data for the particular parameter values $(M_\delta, \sigma_{\delta\mathrm{n}})$. The change in the statistic R is the only modification to the naive approach. The advantage of this method is that R_{data} now has only weak dependence on the value of n. The 90% CL region allowed by the observed data set consists of, as before, all parameter space for which $R_{\mathrm{data}} > R_{90}$ for some value of n.

As stated earlier, there are a number reasons for using the above method. First, it is a classical, or frequentist, statistical approach. Since the distribution of the likelihood-ratio statistic R is calculated by Monte Carlo, the edge of the allowed region answers the question "At what values of M_δ and $\sigma_{\delta\mathrm{n}}$ would 90% of experiments have yielded values of R greater than that observed for the data, given M_δ and $\sigma_{\delta\mathrm{n}}$?" This question pertains to the reproducibility of the experiment, which can be answered without any assumption of "prior" distributions for the parameters to be constrained. Furthermore, the n-projection approach taken here is somewhat conservative because the neutron background is allowed to take the value that yields the most conservative limit; conversely, if the neutron background were known precisely, the limit would presumably be more aggressive because some points $(M_\delta, \sigma_{\delta\mathrm{n}})$ allowed by taking n to a value disfavored by the data would be disallowed. Another reason for using the Feldman-Cousins approach is that it deals correctly with the physical boundaries. This is especially important in the case of a background-subtraction scenario such as this, where a naive Bayesian approach can yield a confidence region that extends to nonphysical parameter space (*e.g.*, $\sigma_{\delta\mathrm{n}} < 0$ or $n < 0$) that must be somehow renormalized. Finally, since the number of events is so small, even a Bayesian approach would not satisfy the asymptotic limit — the likelihood function is not Gaussian — and the exclusion limit would have to be calculated numerically anyways, using a method similar to the one discussed in Section 8.2.1.

Figure 8.26 displays the the border of the excluded region. Because all the nuclear recoils may be neutron scatters, the allowed region includes $\sigma_{\delta\mathrm{n}} = 0$. Hence the excluded region takes the form of an upper limit. For completeness, the full likelihood function for the data is shown in Figure 8.27. The neutron background is normalized to $n = 17$ events. This is slightly different than the best-fit value $n = 17.5$, but n is not well determined by the data and using $n = 17$ introduces negligible error.

The limit is most sensitive to the number of Ge doubles because their unexpectedly high number pushes the expected number of Ge singles up, improving the background subtraction. The limit is fairly insensitive to the number of Si events. However, even with the high number of Ge doubles, the 90% CL exclusion limit corresponds to an expectation of ~ 8 WIMPs in the Ge single-scatter data set for WIMP masses where the shape of the expected recoil spectrum from WIMPs matches the observed spectrum well. In an absolute sense, then, the high number of Ge doubles does not have that significant an effect: even if *no* neutron-background subtraction had been performed, the limit would only have been approximately a factor of two worse (90% CL upper limit on 13 events is 17.6 events, Bayesian or classical). The high expected number of WIMPs at the 90% CL exclusion limit (more than half the single scatters) gives an indication of the Poisson penalty exacted by the use of such small numbers of Ge multiples and Si singles to constrain the neutron background.

To gauge how likely the calculated limit is, one can simulate an ensemble of experiments assuming the most likely neutron-background normalization $n = 17.5$ and calculate limits for these simulated experiments. The exclusion limit for which 50% of simulated experiments yield less sensitive limits is shown on the exclusion plot (the "median expected 90% CL exclusion limit"). Clearly, we got lucky. However, it is important to note there is nothing incorrect about the calculated limit; the point is that most realizations of this experiment would have set less sensitive exclusion limits.

Regarding the effect of modifying the ionization-partition cut: redoing this analysis with inclusion of the cQShare event sample would yield a more conservative limit since the number of single scatters increases from 13 to 27 (as can be read off of Figures 8.6 and 8.7), the number of double-scatter neutrons is unchanged, and the exposure only increases by about 1/3. Future analysis will include the cQShare events [147].

The limit thus calculated probes the MSSM more deeply than previous limits. In particular, it begins to exclude a significant number of the Gondolo *et al.* pure gaugino models. The limit also begins to approach the more constrained Corsetti and Nath mSUGRA models, though does not exclude any. Perhaps more importantly, this exclusion limit tests the annual-modulation signal claimed by the DAMA collaboration [88]. The compatibility of this limit and the annual-modulation signal is discussed in detail in Appendix E.

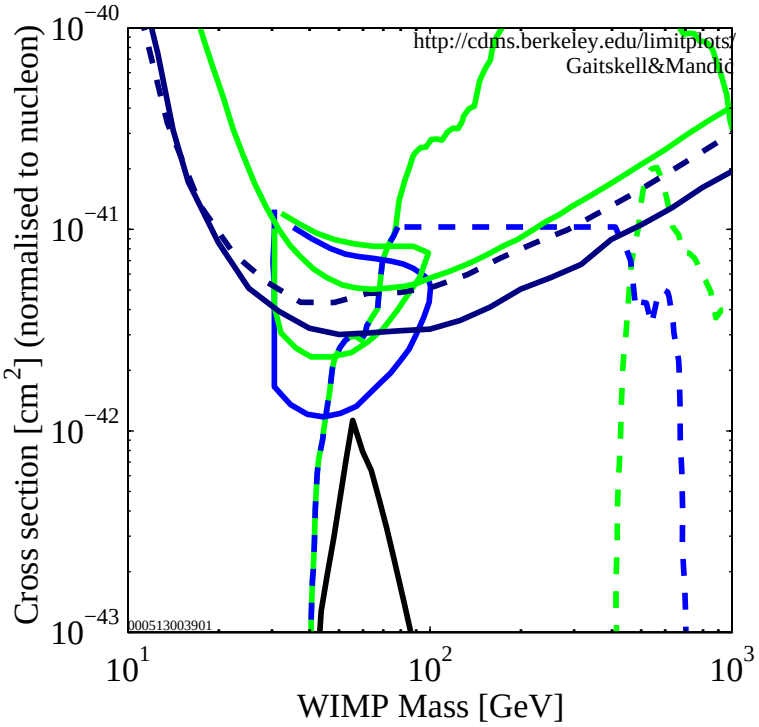

Figure 8.26: Exclusion limit on the spin-independent WIMP-nucleon elastic-scattering cross section based on combined analysis of Run 19 Ge single-scatter and double-scatter nuclear recoils and Run 18 Si nuclear recoils. Dark solid U-shaped curve: 90% CL exclusion limit from these data. Dark dashed U-shaped curve: median expected 90% CL exclusion limit assuming a neutron background $n = 17.5$. Light solid U-shaped curve: DAMA 1996 pulse-shape-analysis 90% CL exclusion limit [146]. Light heart-shaped region: DAMA 2000 annual-modulation 3σ CL allowed region without application of DAMA 1996 limit [88]. Dark heart-shaped region: DAMA 2000 annual-modulation 3σ CL allowed region from combined analysis with DAMA 1996 limit [88]. The other curves are the MSSM predictions shown in Figure 2.9. Light solid: Gondolo *et al.* mixed models [78]. Dark dashed: Gondolo *et al.* gaugino models. Light dashed: Gondolo *et al.* Higgsino models. Dark solid: Corsetti and Nath mSUGRA models [79]. Figure from [80].

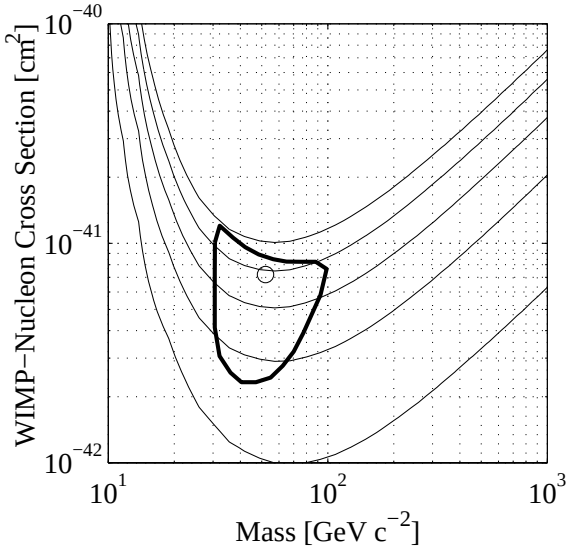

Figure 8.27: Likelihood function in the $(M_\delta, \sigma_{\delta \mathrm{n}})$ plane for combined Ge single-scatter, Ge double-scatter, and Si nuclear-recoil data. The neutron background is normalized to $n = 17$ events. The likelihood is maximized at $\sigma_{\delta \mathrm{n}} = 0$ (any M_δ), below the lower edge of the plot. The contours indicate $\Delta \log \mathcal{L} = -0.5, -2, -4.5, -8,$ and -12.5, with -0.5 the lowermost. The DAMA best-fit $(M_\delta, \sigma_{\delta \mathrm{n}}) = (52 \text{ GeV c}^{-2}, 7.2 \times 10^{-42} \text{ cm}^2)$ and 3σ CL contour (not using the 1996 exclusion limit) are shown. Figure provided by R. Schnee.

8.5 Residual Background-Particle-Interaction Rates

It is important to assess the interaction rates of all types of background particles in order to determine which sources need to be addressed to achieve the background-rate goals for both CDMS I and CDMS II. Classification by ionization partition and particle type provide a wealth of information in this respect. While only single-scatter events impact WIMP sensitivity, multiple-scatter events can illuminate the sources of the single-scatter events. In this section, I summarize the rates of single-scatter and double-scatter photon and electron events. Triple- and quadruple-scatter events are neglected because such events presumably arise from very energetic particles; such events shed little light on the low-energy processes that dominate single and double scatters. Similarly, events with saturated detectors are neglected. A comprehensive analysis of the data presented here is beyond the scope of this dissertation, but I hope this section provides a starting point.

Ionization-yield plots for the veto-anticoincident double-scatter data set are shown in Figures 8.28, 8.29, and 8.30. Saturated events are excluded, as mentioned above. It should be remembered that the cuts defining the double-scatter data set shown in Figure 8.18 are far more restrictive than those used to define the data plotted in these figures; the elevated number of events in these plots should not be cause for concern.

8.5.1 Exposure Normalization and Particle Identification

It is important to determine the correct normalizations for the various particle rates. As will be seen, the photon rates are dominated by Compton scattering of penetrating high-energy photons and thus volume (or mass) normalization is appropriate. For electrons, which do not penetrate deeply, area normalization is more useful. Table 8.7 lists the assumed volumes and areas of the three ionization-partition cuts. The volume fractions are determined from the neutron-calibration data and simulation as discussed in Section 4.5.1. The cQinOnly and cQShare cuts select regions on the detector flats and thus the accepted areas can be determined trivially. The cQoutOnly area requires an estimate of the total detector area. The curved edge of the detector is assumed to be a circle of radius 0.6 cm, corresponding to half the detector thickness of 1.2 cm, in accordance with the schematic shown in Figure 4.1. The total surface area is given by

$$A = 2\pi r_{flat}^2 + 2\pi^2 r_{eq} r_c - 2\pi(\pi - 2) r_c^2; \tag{8.11}$$

The first term is the area of the two flats, with $r_{flat} = 2.55$ cm. The second two terms give the area of the curved edge, with $r_{eq} = 3.00$ cm being the detector radius at the equator and $r_c = 0.6$ cm being the edge curvature radius. This form is an approximation because the curved edges are not perfectly circular, but the formula should be correct to a few percent. The area of the cQoutOnly cut is given by subtracting the cQinOnly and cQShare areas from the total area. One can infer radii of 1.95 cm and 2.31 cm for the boundaries separating the cQinOnly, cQShare, and cQoutOnly regions, but these should not be considered exact because of the indentation in the electrode break shown in Figure 4.1. The areas, on the other hand, account for the indentation.

The mass and area exposures for the data presented in this section are listed in Table 8.8. These exposures do not take into account the efficiencies of the energy-dependent cuts, which have negligible effect at low energy and a 10%–15% effect near 100 keV. Because this error is common to all the rates, rate comparisons are subject to much smaller errors.

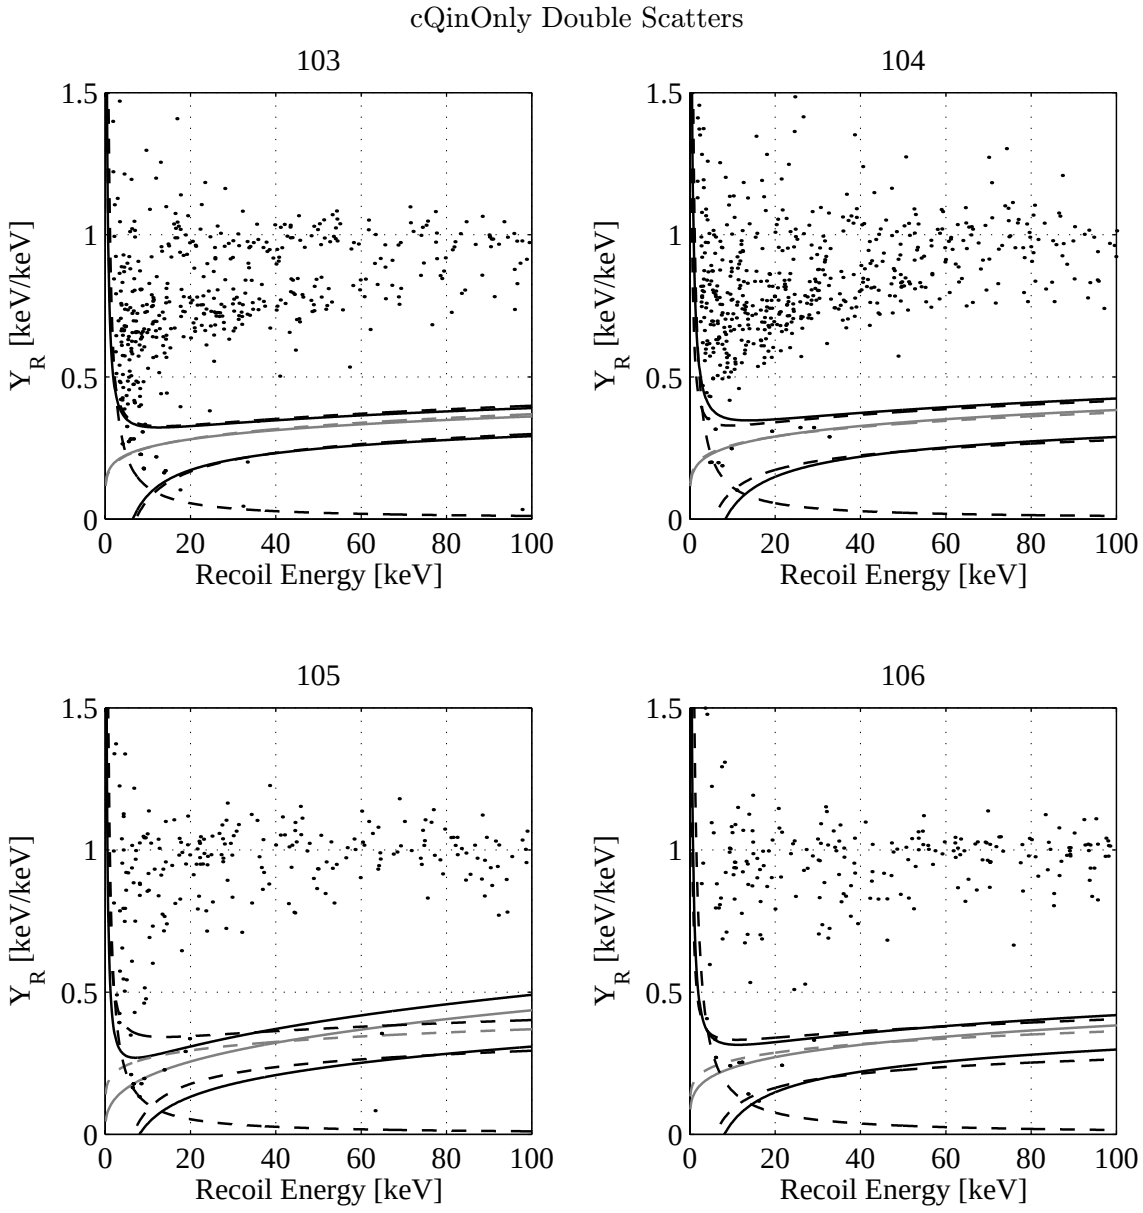

Figure 8.28: Ionization yield vs. recoil energy for veto-anticoincident cQinOnly recoils belonging to double scatters. Saturated events are excluded. Hyperbolic dashed line: ionization-search threshold. Light solid line: center of nuclear-recoil band, pre-April 3. Light dashed line: center of nuclear-recoil band, post-April 3. Dark solid line: nuclear-recoil-acceptance region, pre-April 3. Dark dashed line: nuclear-recoil-acceptance region, post-April 3.

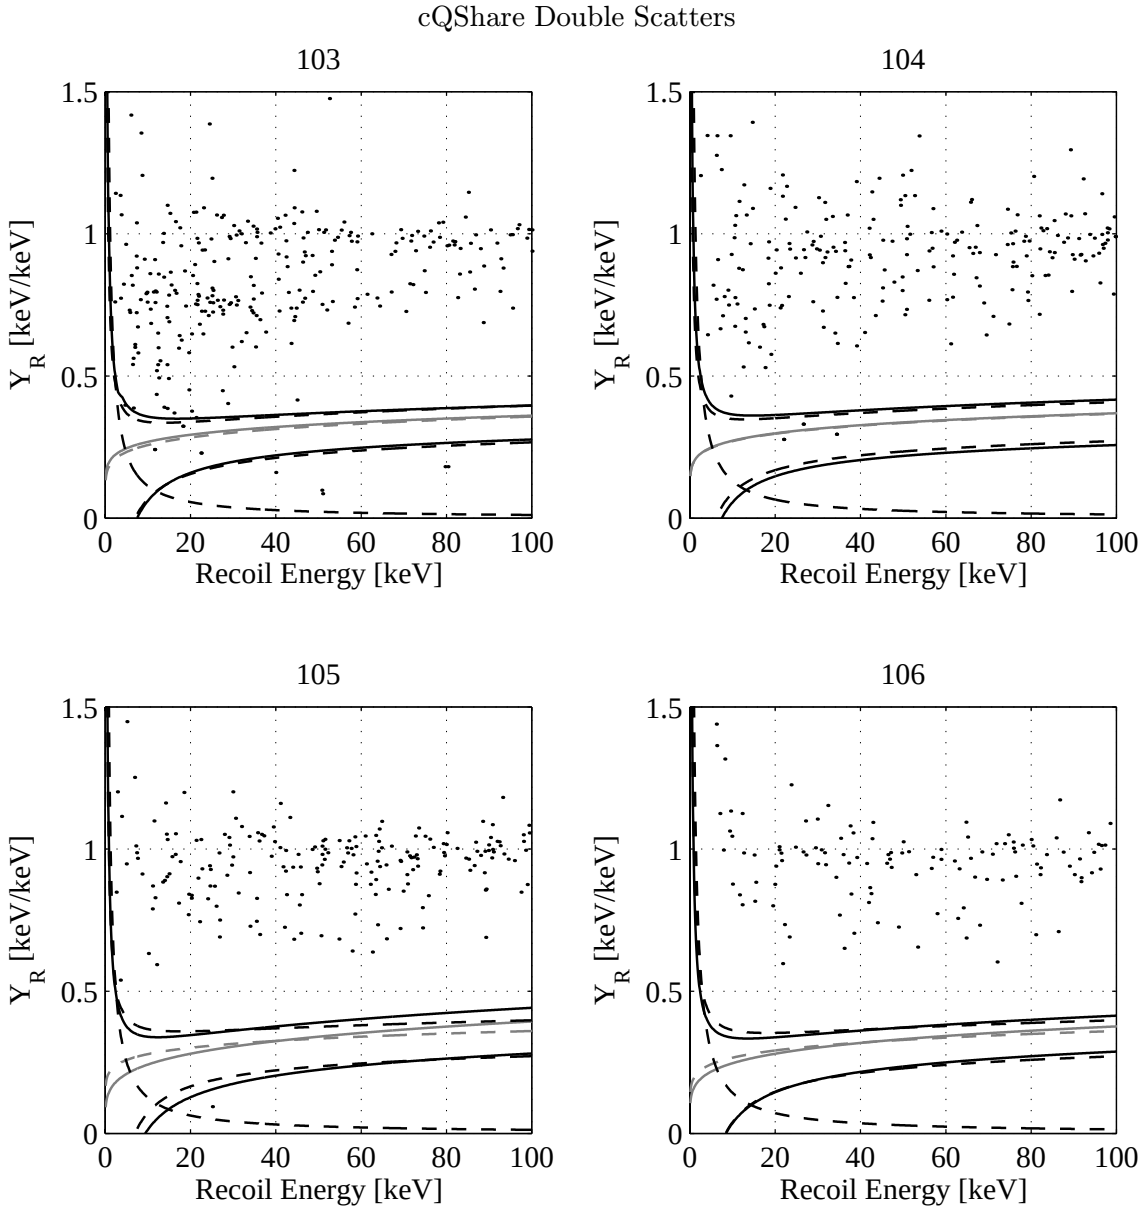

Figure 8.29: Ionization yield vs. recoil energy for veto-anticoincident cQShare recoils belonging to double scatters. Saturated events are excluded. Legend as in Figure 8.28.

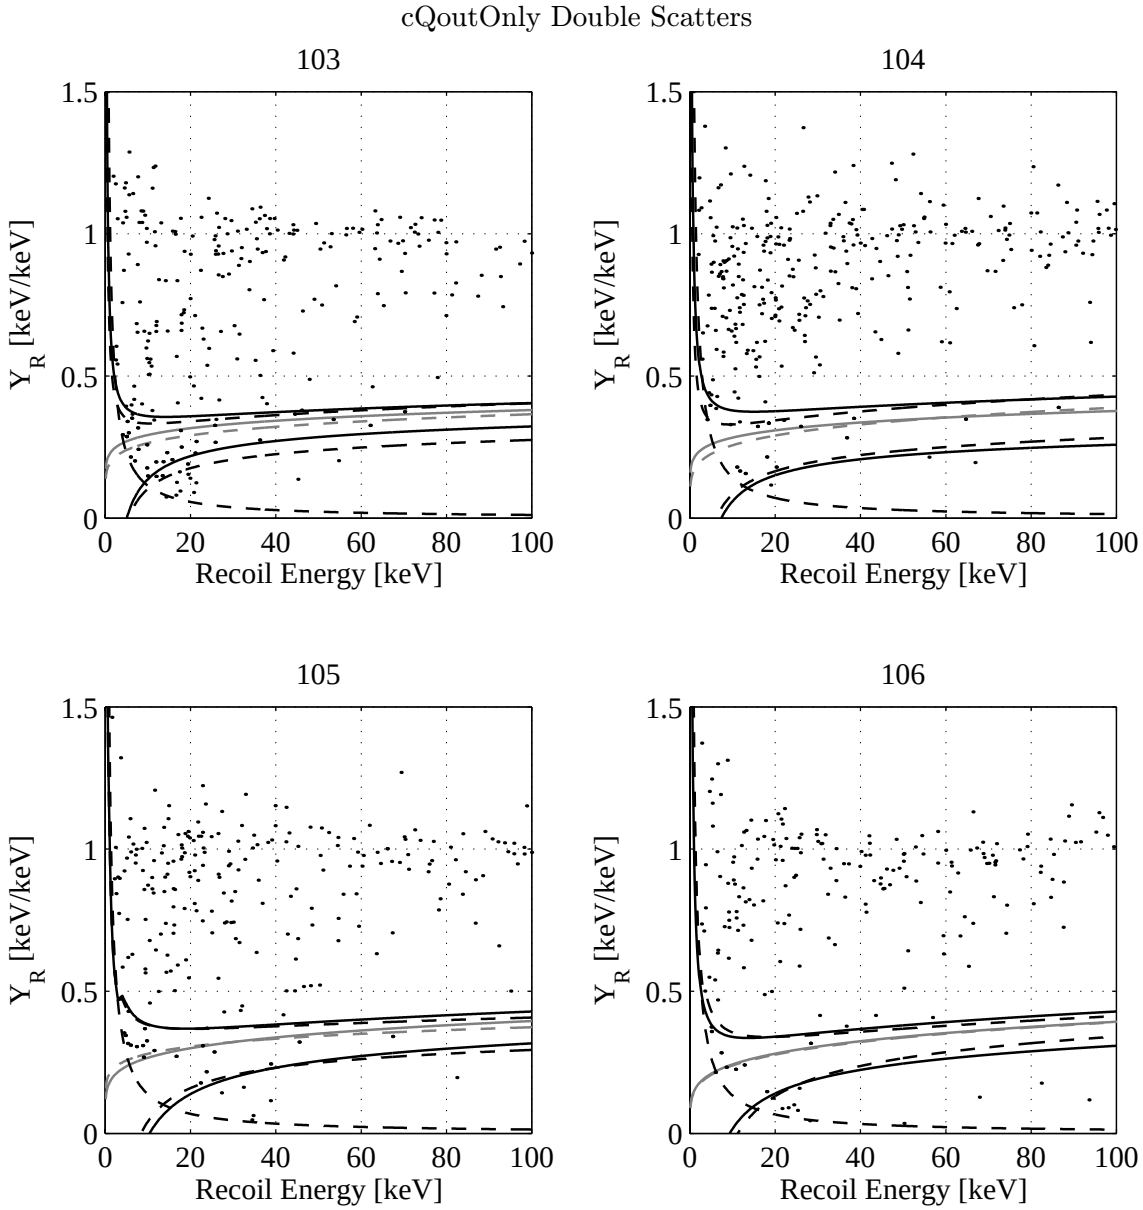

Figure 8.30: Ionization yield vs. recoil energy for veto-anticoincident cQoutOnly recoils belonging to double scatters. Saturated events are excluded. Legend as in Figure 8.28.

Event set	Mass [g]	Volume [cm^3]	Volume Fraction	Area [cm^2]	Area Fraction
Qin	75.9	14.3	0.46	23.8	0.32
QShare	31.4	5.9	0.19	9.8	0.13
Qout	57.7	10.8	0.35	40.2	0.55
Total	165.0	31.0	1.00	73.8	1.00

Table 8.7: Masses, volumes, and areas of ionization-partition cuts.

Event set	BLIP3	BLIP4	BLIP5	BLIP6
Volume Exposure [kg d]				
Qin	4.77	4.64	4.98	5.00
QShare	1.97	1.92	2.06	2.07
Qout	3.62	3.54	3.79	3.81
Area Exposure [cm^2 d]				
Qin	1500	1450	1560	1570
QShare	610	600	640	650
Qout	2520	2470	2640	2650

Table 8.8: Volume and area exposures for photons and electrons. No energy-dependent efficiency corrections have been made, though the veto-anticoincidence cut efficiency has been approximately taken into account using an energy-independent value of 0.85, the value for ionization triggers.

As was discussed in Section 8.2.1, particle identification is based on ionization yield. A straightforward cut enclosing the photon band defines photons; the cut flares out at low energy to allow for the widening of the band. Any event above 300-keV recoil energy is also considered a photon for the purposes of double scatters. Electrons are defined to be all events between the photon and nuclear-recoil bands. For the cQoutOnly cut, leakage of electrons into the nuclear-recoil band dominates over true nuclear recoils, so all events below the photon band are taken to be electrons. Finally, for double-scatters, if the second scatter is below 10 keV, the event is separately classified as "unidentifiable" because the bands merge.

8.5.2 Photon Rates

Single-scatter photon spectra are shown in Figures 8.10, 8.11 and 8.12. Spectra for photons belonging to double scatters are shown in Figures 8.31, 8.32, and 8.33. The cQinOnly and cQShare spectra are subdivided by the particle identification of the *other* scatter. No ionization-partition cut is made on the second scatter. The average rates between 10 and 100 keV are summarized in Table 8.9.

The cQinOnly photon spectra are fairly flat, as expected for Compton scatters. The event set is dominated by events in which the second scatter is also a photon, again consistent with Compton scattering. There is an extremely low rate of events with the second scatter being an

electron, at the level of 5–10% of all double scatters. Approximately 10% of double scatters have second hits below 10 keV. BLIP4 shows uniformly higher rates of all multiples. As is evident in many of the ionization-yield plots, BLIP4's photon band is less well defined, so this may just reflect acceptance of electrons into the photon data set. BLIP4 has a large rate of double-scatter electrons in coincidence with BLIP3 because of BLIP3's surface contamination, so such misidentification is especially problematic for BLIP4. Such misidentification may also explain the generally higher photon rates seen for BLIP3. Finally, there is some self-shielding against photons, as seen by the variations in overall rates, but it is only at the 20–30% level (assuming electron misidentification is the cause of the elevated BLIP3 and BLIP4 rates). This is also expected, as the penetration length of high-energy photons is large compared to the detector thickness.

The cQShare spectra are also fairly flat. The rates are elevated above both the cQinOnly and cQoutOnly rates, presumably because internal double scatters cause events to be classified as cQShare events though the interactions may not have occurred in the cQShare volume. The cQShare double scatters are again dominated by Compton scatters. Events in which the second scatter is an electron are a small fraction, as for cQinOnly events. It is interesting that the fraction of events in which the second scatter is below 10 keV is elevated. One possible explanation is that these photons also produce a low-energy-photon and -electron continuum in the surface of the detector housing; the cQShare region is not as well shielded as the cQinOnly region. Another geometrical effect may also be important. If a high-energy photon Compton scatters in the cQinOnly volume of a detector, resulting in the emission of low-energy particles from the detector surface that hit the adjacent detector, the incident photon is not severely deflected (because it did not lose much energy), so it is likely to be collimated with the low-energy particle. The scattered photon may thus scatter in the same detector as the low-energy particle, increasing the second scatter's energy. This effect is presumably less important for cQShare events.

The cQoutOnly photon spectra show a surprising rise below 20 keV. The cause is not clear, though electron confusion is likely. The low-energy region of the photon spectra is explored in Section 8.5.4.

The above discussion is admittedly fairly speculative. To understand in detail these photon rates, and especially the correlations between particle types, a full photon simulation is necessary. However, photon rates are not a particular worry because of the detectors' excellent photon discrimination. Even at the 1 keV^{-1} kg^{-1} d^{-1} rate seen, the expected photon misidentification is at the 0.001 keV^{-1} kg^{-1} d^{-1} level, which meets the CDMS II goal [148]. Therefore, further analysis of the photon rates is a low priority.

Photon Rates [keV^{-1} kg^{-1} d^{-1}]

Event set	BLIP3	BLIP4	BLIP5	BLIP6
cQinOnly				
single	1.287±0.055	0.825±0.044	0.667±0.039	0.793±0.042
multi w/photon	0.282±0.026	0.371±0.030	0.312±0.026	0.302±0.026
multi w/electron	0.047±0.010	0.089±0.015	0.016±0.006	0.016±0.006
multi w/< 10 keV	0.021±0.007	0.091±0.015	0.038±0.009	0.033±0.009
multi all	0.350±0.029	0.550±0.036	0.366±0.029	0.351±0.028
all	1.637±0.062	1.376±0.057	1.033±0.048	1.144±0.050
cQShare				
single	1.570±0.094	1.280±0.086	1.394±0.087	1.915±0.101
multi w/photon	0.638±0.060	0.817±0.069	0.708±0.062	0.446±0.049
multi w/electron	0.034±0.014	0.041±0.015	0.059±0.018	0.016±0.009
multi w/< 10 keV	0.164±0.030	0.324±0.043	0.297±0.040	0.113±0.025
multi all	0.836±0.069	1.182±0.083	1.064±0.076	0.575±0.056
all	2.405±0.117	2.462±0.119	2.458±0.115	2.490±0.116
cQoutOnly				
single	2.011±0.079	1.311±0.064	1.205±0.059	1.580±0.068
multi all	0.359±0.033	0.591±0.043	0.460±0.037	0.385±0.034
all	2.369±0.085	1.902±0.077	1.666±0.070	1.965±0.076

Table 8.9: Veto-anticoincident-photon rates, averaged between 10 and 100 keV, in keV^{-1} kg^{-1} d^{-1}.

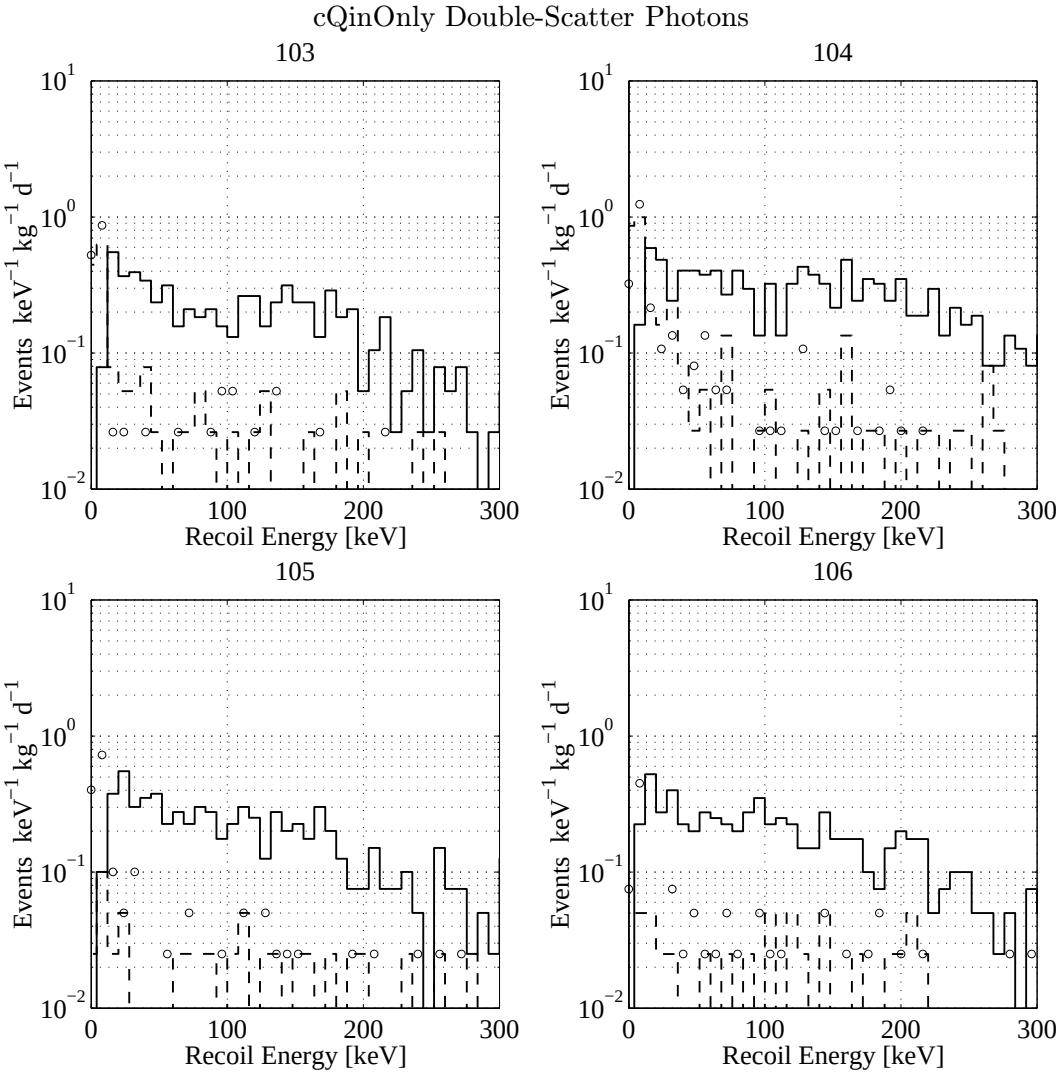

Figure 8.31: Recoil-energy spectra for cQinOnly photons belonging to veto-anticoincident double-scatter events. Solid: other scatter is a photon above 10 keV. Dashed: other scatter is an electron above 10 keV. Circles: other scatter is below 10 keV (photon/electron discrimination not possible).

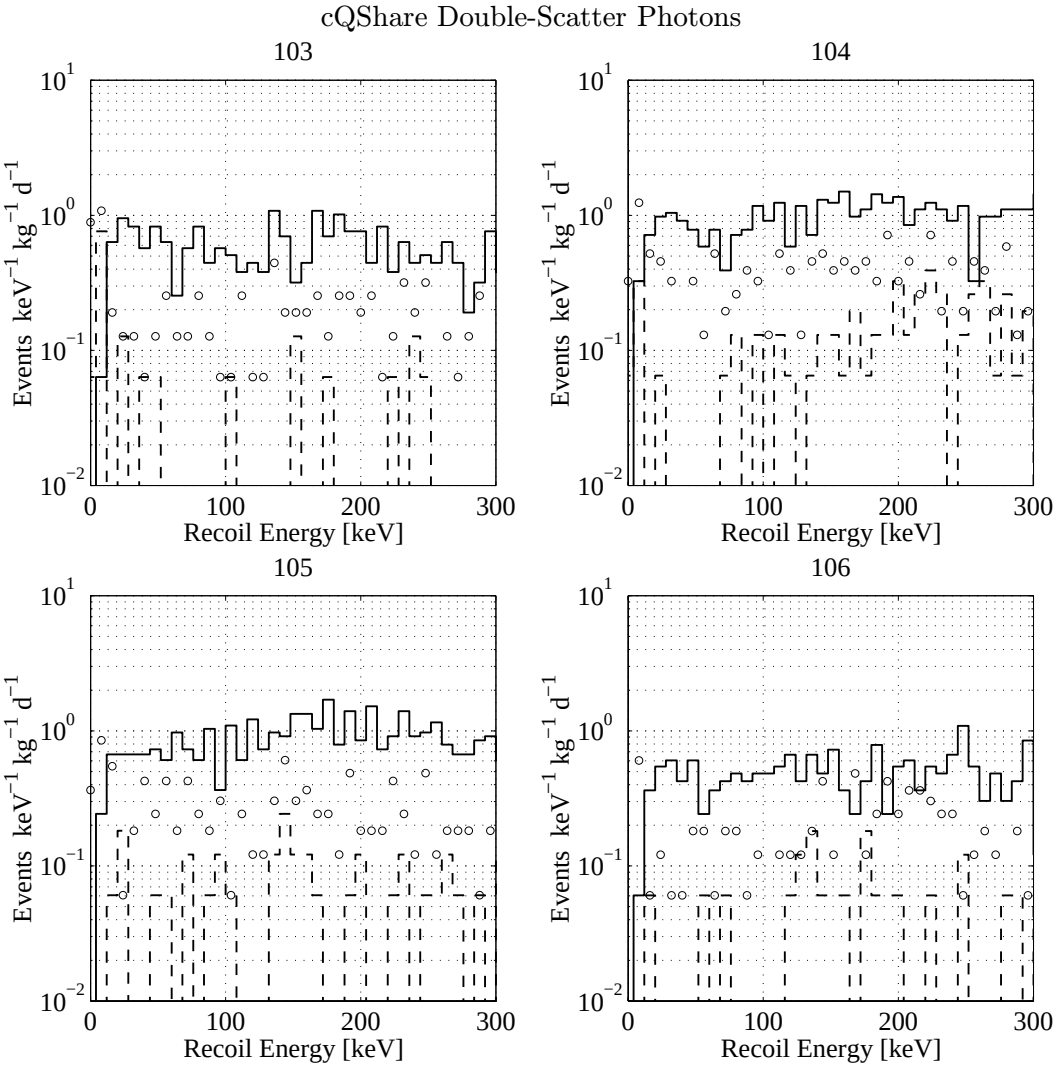

Figure 8.32: Recoil-energy spectra for cQShare photons belonging to veto-anticoincident double-scatter events. Legend as in Figure 8.31.

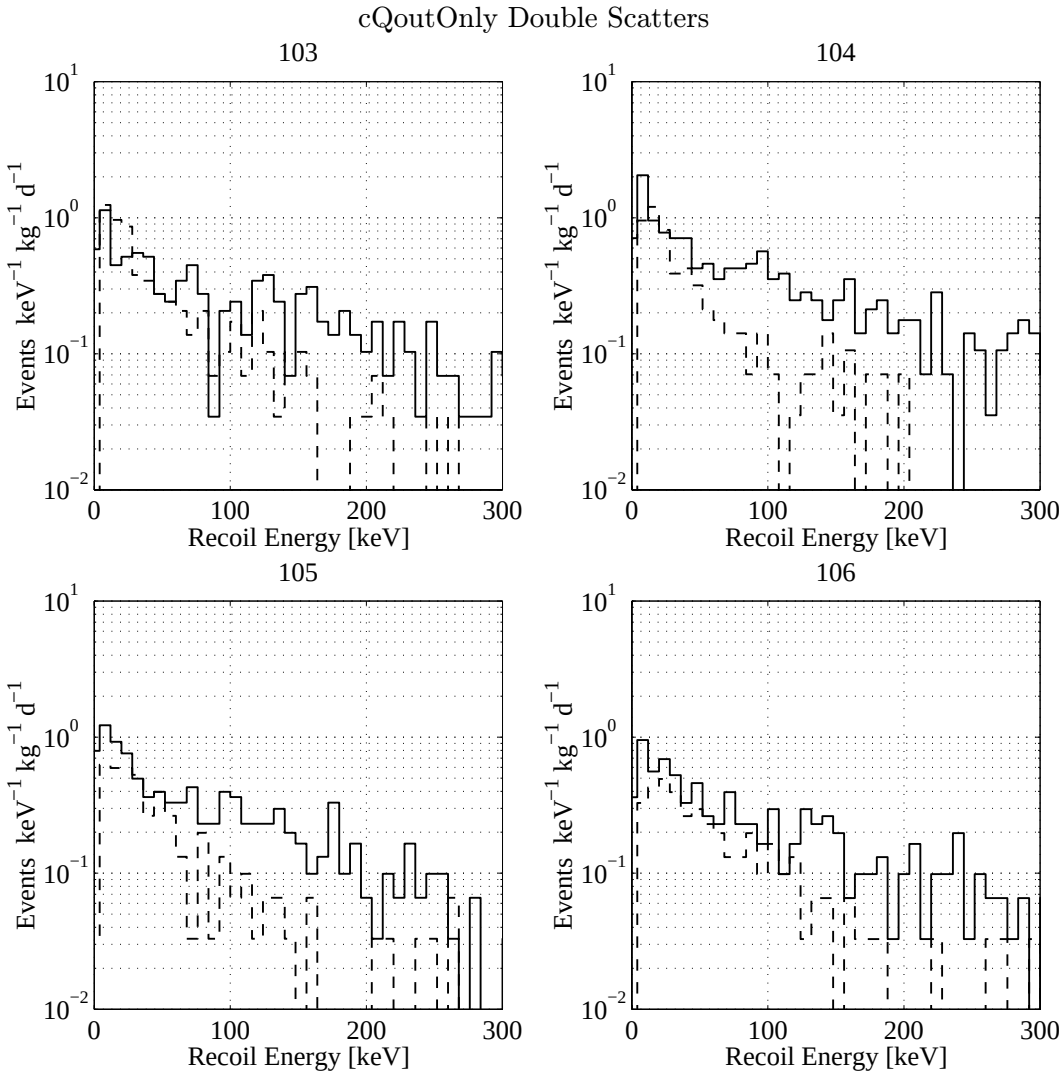

Figure 8.33: Recoil-energy spectra for cQoutOnly photons and electrons belonging to veto-anticoincident double-scatter events. Solid: photons. Dashed: electrons. Note that these spectra are *not* analogous to Figures 8.31, 8.32, 8.34, and 8.35 because the events are not subdivided by the identity of the second scatter.

Electron Rates [10^{-3} cm^{-2} d^{-1}]

Event set	BLIP3	BLIP4	BLIP5	BLIP6
cQinOnly				
single	576.5±19.7	132.7± 9.6	15.4± 3.1	40.2± 5.1
multi w/photon	16.1± 3.3	14.4± 3.2	7.7± 2.2	17.9± 3.4
multi w/electron	46.9± 5.6	70.1± 6.9	5.1± 1.8	1.9± 1.1
multi w/< 10 keV	52.3± 5.9	50.9± 5.9	2.6± 1.3	3.2± 1.4
multi all	115.3± 8.8	135.5± 9.7	15.4± 3.1	23.0± 3.8
all	691.8±21.5	268.2±13.6	30.8± 4.4	63.2± 6.4
cQShare				
single	392.7±25.2	58.3± 9.9	41.9± 8.1	117.5±13.5
multi w/photon	37.3± 7.8	38.3± 8.0	40.4± 7.9	32.5± 7.1
multi w/electron	77.9±11.2	40.0± 8.2	7.8± 3.5	3.1± 2.2
multi w/< 10 keV	79.5±11.4	30.0± 7.1	28.0± 6.6	13.9± 4.6
multi all	194.8±17.8	108.2±13.4	76.1±10.9	49.5± 8.7
all	587.5±30.9	166.5±16.7	118.0±13.5	167.0±16.1
cQoutOnly				
single	327.1±11.4	125.1± 7.1	82.2± 5.6	93.9± 6.0
multi all	49.1± 4.4	50.4± 4.5	38.3± 3.8	34.0± 3.6
all	376.2±12.2	175.5± 8.4	120.5± 6.8	127.9± 6.9

Table 8.10: Veto-anticoincident-electron rates, *summed* between 10 and 100 keV, in 10^{-3} cm^{-2} d^{-1}.

8.5.3 Electron Rates

Single-scatter-electron spectra are shown in Figures 8.10, 8.11 and 8.12. Spectra for electrons belonging to double scatters, analogous to the spectra shown for photons, are shown in Figures 8.34, 8.35, and 8.33. The total areal rates between 10 and 100 keV are summarized in Table 8.10. Note that the spectra are in keV^{-1} kg^{-1} d^{-1} while the table shows rates in cm^{-2} d^{-1}. The single-scatter rates in keV^{-1} kg^{-1} d^{-1} are shown in Table 8.1.

The contamination of BLIP3 is clearly visible in the BLIP3 and BLIP4 spectra. This contamination dominates these detectors' electron rates, so they are not useful for constraining residual electron sources. I therefore concentrate on BLIP5 and BLIP6. The lower single-scatter rate in BLIP5 reflects self-shielding. The ratio of double scatters to all events is *higher* for BLIP5 than BLIP6, as one would expect because one face of BLIP6 is not shielded by other detectors.

A simple model can be constructed to attempt to interpret these rates. Three different fluxes of electrons are assumed:

- F_γ: the flux of electrons ejected by photons ("photon-induced flux")
- $F_{i\beta}$: the flux of electrons due to beta emission by a contaminant on the surface of the detector housing, yielding an isotropic flux of electrons ("isotropic-electron flux")
- $F_{s\beta}$: the flux of electrons due to beta emission by a surface contaminant on the detector surfaces ("surface-contaminant flux")

F_γ is assumed to be emitted isotropically from any surface and to be independent of the emitting surface: the housing and detector surfaces emit the same flux F_γ. This is a valid approximation as long as the photons that eject electrons are of sufficiently high energy ($\gtrsim$ 100 keV) that their penetration length is large compared to the detector thickness. $F_{i\beta}$ is also assumed to be emitted isotropically by the detector-housing surface and to be homogeneous over the housing surface. Finally, $F_{s\beta}$ is assumed to be emitted isotropically by any detector surface and also to be homogenous. To characterize multiple-scattering and self-shielding for the inner electrodes, three parameters are assumed:

- m_γ: the double-scatter fraction for F_γ electrons
- m_β: the double-scatter fraction for $F_{i\beta}$ or $F_{s\beta}$ electrons
- s_s: the flux-reduction factor for electrons due to self-shielding

Two different double-scatter fractions are assumed because the physical sources of F_γ and of $F_{i\beta}$ and $F_{s\beta}$ electrons are quite different. A single double-scatter fraction for $F_{i\beta}$ and $F_{s\beta}$ electrons is assumed for simplicity. This may be invalid because $F_{i\beta}$ electrons are likely to have shallow incident angles on the inner electrodes while $F_{s\beta}$ electrons will have more evenly distributed incident angles. Finally, s_s describes the self-shielding for $F_{i\beta}$ electrons. F_γ electrons are not subject to self-shielding because *any* surface emits F_γ electrons. $F_{s\beta}$ electrons are also not subject to self-shielding. Various observable electron rates are defined as follows:

- $R^s_{\rm BX}$: BLIPX's single-scatter-electron rate
- $R^\gamma_{\rm BX}$: BLIPX's double-scatter electron rate for events where the second scatter is a photon
- $R^\beta_{\rm BX}$: BLIPX's double-scatter electron rate for events where the second scatter is an electron or is below 10 keV

These rates are given by the first four rows of Table 8.10. $R^\beta_{\rm BX}$ lumps double scatters with electrons or with $<$ 10-keV events together under the assumption that the latter is dominated by electrons. Photon-electron coincidences are presumably due primarily to Compton scattering of high-energy photons, so the spectrum of energies deposited by the photon should be flat, while the spectrum due to electrons should naturally rise at low energies.

The model consists of the following relations among the fluxes and observed rates:

$$R^s_{\rm B5} = F_\gamma\left[1-m_\gamma\right] + F_{i\beta}s_s\left[1-m_\beta\right] + F_{s\beta}\left[1-m_\beta\right] \quad (8.12)$$

$$R^s_{\rm B6} = F_\gamma\left[1-\frac{1}{2}m_\gamma\right] + F_{i\beta}\left[\frac{1}{2}+\frac{1}{2}s_s(1-m_\beta)\right] + F_{s\beta}\left[1-\frac{1}{2}m_\beta\right] \quad (8.13)$$

$$R^\gamma_{\rm B5} = F_\gamma m_\gamma \quad (8.14)$$

$$R^\gamma_{\rm B6} = \frac{1}{2}F_\gamma m_\gamma \quad (8.15)$$

$$R^\beta_{\rm B5} = F_{i\beta}s_s m_\beta + F_{s\beta}m_\beta \quad (8.16)$$

$$R^\beta_{\rm B6} = \frac{1}{2}F_{i\beta}s_s m_\beta + \frac{1}{2}F_{s\beta}m_\beta \quad (8.17)$$

The first line is BLIP5's single-scatter rate, consisting of the single-scatter photon-induced flux, the self-shielded single-scatter isotropic-electron flux, and the single-scatter surface-contaminant flux. The second line is BLIP6's single-scatter rate. Because BLIP6 is at the bottom of the stack, only one face is subject to multiple scattering and self-shielding, reflected by the various 1/2's in the expression. The third and fourth lines are the multiple-scatter photon-induced fluxes; again, BLIP6's multiple-scatter fraction is reduced by a factor of 2. The fifth and sixth lines are the multiple-scatter electron rates, including self-shielding, with BLIP6's rate again lower by a factor of 2.

It is clear that the model does not fit the observed rates perfectly; for example, $R^s_{\rm B5} = 2R^\beta_{\rm B5}$ does not hold. To rigorously assess the ability of the model to match the data, a standard χ^2 quantity can be calculated:

$$\begin{aligned}\chi^2 \;=\; & \frac{(R^s_{\rm B5} - \widehat{R^s_{\rm B5}})^2}{(\delta R^s_{\rm B5})^2} + \frac{(R^s_{\rm B6} - \widehat{R^s_{\rm B6}})^2}{(\delta R^s_{\rm B6})^2} \\ & + \frac{(R^\gamma_{\rm B5} - \widehat{R^\gamma_{\rm B5}})^2}{(\delta R^\gamma_{\rm B5})^2} + \frac{(R^\gamma_{\rm B6} - \widehat{R^\gamma_{\rm B6}})^2}{(\delta R^\gamma_{\rm B6})^2} \\ & + \frac{(R^\beta_{\rm B5} - \widehat{R^\beta_{\rm B5}})^2}{(\delta R^\beta_{\rm B5})^2} + \frac{(R^\beta_{\rm B6} - \widehat{R^\beta_{\rm B6}})^2}{(\delta R^\beta_{\rm B6})^2}\end{aligned} \tag{8.18}$$

where hats indicate the model rates and the denominators are the statistical uncertainties given in Table 8.10. This is not a χ^2 in the usual sense because the number of degrees of freedom equals the number of data points. Regardless, it can be minimized to yield the best-fit parameters $\widehat{F_\gamma}$, $\widehat{F_{i\beta}}$, $\widehat{F_{s\beta}}$, $\widehat{m_\gamma}$, $\widehat{m_\beta}$, and $\hat{s}_s$. The best-fit χ^2 is 15.8. The best-fit parameter values and resulting expected values for the six observable rates are listed in Table 8.11. Since the expected distribution of χ^2 is not known, the probability of observing a higher χ^2 cannot be determined without a simulation. However, it is clear from the discrepancies between the expected and observed rates that the model does not fit the data.

The poor fit arises because the model predicts $R^\gamma_{\rm B5}/R^\gamma_{\rm B6} = 2$ while the data are statistically precise and have $R^\gamma_{\rm B5}/R^\gamma_{\rm B6} = 0.5$. The photon-induced flux model, where only one flux is assumed for all surfaces, may be wrong because of the complex geometry: the convex curved surfaces may shield more than they emit and the concave detector housing surface may emit more than a flat surface that occupies the same solid angle. To check this, $R^\gamma_{\rm B6}$ can be discarded and the minimization redone, yielding $\chi^2 = 0.35$. The best-fit parameters and the expected values of the observables are again listed in Table 8.11. Clearly, the model fits the data reasonably well now. The decrease in χ^2 is significantly more than what one would expect from simply discarding a single data point, corroborating the fact that $R^\gamma_{\rm B6}$ is the crux of the inconsistency.

It is nontrivial to determine the probability of observing a higher value of χ^2 because the expected distribution is not analytic (there are –1 degrees of freedom!). The expected distribution can be determined by simulation. The simulated observables are drawn from Gaussian distributions with means determined by the best-fit parameters $\widehat{X_i}$ and widths given by the observed statistical errors. For the nth simulated outcome, the model parameters $\widehat{X_{i,n}}$ that best fit the simulated observables $N_{i,n}$ are determined and the quantity $\chi^2_n = \chi^2(N_{i,n}|\widehat{X_{i,n}})$ calculated. The distribution of χ^2_n is the expected distribution. Note that it is necessary to calculate $\chi^2(N_{i,n}|\widehat{X_{i,n}})$ rather than $\chi^2(N_{i,n}|\widehat{X_i})$ to minimize the dependence of the expected distribution on the particular values $\widehat{X_i}$

Best-Fit Model Parameters and Results

Quantity	all data	wo/ R^{γ}_{B6}	wo/ R^{γ}_{B5}
$\widehat{F_{\gamma}}$	11.3	16.4	43.7
$\widehat{F_{i\beta}}$	52.9	42.7	10.3
$\widehat{F_{s\beta}}$	0.8	5.9	11.0
$\widehat{m_{\gamma}}$	0.92	0.47	0.82
$\widehat{m_{\beta}}$	0.37	0.56	0.53
$\widehat{s_s}$	0.42	0.22	0.48
χ^2	15.8	0.35	0.35
$P(>\chi^2)$	0.001	0.54	0.68

Observed and Best-Fit Model Rates

Rate	Observed	Best-Fit Model		
		all data	wo/ R^{γ}_{B6}	wo/ R^{γ}_{B5}
R^{s}_{B5}	15.4±3.1	15.4	15.4	15.4
R^{s}_{B6}	40.2±5.1	40.2	40.2	40.2
R^{γ}_{B5}	7.7±2.2	10.4	7.7	35.8
R^{γ}_{B6}	17.9±3.4	5.2	3.8	17.9
R^{β}_{B5}	7.7±2.2	8.4	8.4	8.4
R^{β}_{B6}	5.1±1.8	4.2	4.2	4.2

Table 8.11: Electron-rate model best-fit parameters and expected rates. All rates and fluxes are in 10^{-3} cm^{-2} d^{-1}. m_{γ}, m_{β}, and s_s are unitless. The "all data" values correspond to fits including both R^{γ}_{B6} and R^{γ}_{B5}. For the "wo/ R^{γ}_{B6}" and "wo/ R^{γ}_{B5}" fits, R^{γ}_{B6} and R^{γ}_{B5}, respectively, are not included in the χ^2. $P(>\chi^2)$ is determined by simulation, as described in the text.

assumed for the simulation. Performing the test using the former statistic is also more conservative because $\chi^2(N_{i,n}|\widehat{X_{i,n}}) \leq \chi^2(N_{i,n}|\widehat{X_i})$. Based on the simulated distribution, the percentage of outcomes with higher χ^2 than observed is 54%, indicating the fit is reasonable. Furthermore, the best-fit value $\widehat{s_s}$ is close to the expected value of 0.2 (see Figure 5.5) in this case.

For completeness, the case where R^{γ}_{B5} is discarded is considered. The resulting fit has χ^2= 0.35 also. In this case, 68% of experiments would have had higher χ^2. $P(>\chi^2)$ thus does not strongly discriminate between the two models. It is perhaps more reasonable to discard R^{γ}_{B6} because doing so yields $s_s = 0.22$, close to expectations. Also, the simulated distribution of $\widehat{s_s}$ is flat from 0 to 1 for the case where R^{γ}_{B5} is discarded, indicating there is no sensitivity to s_s. However, $F_{i\beta}$ is a factor of 7 larger than $F_{s\beta}$ when R^{γ}_{B6} is discarded. One would expect $F_{i\beta}$ and $F_{s\beta}$ to be comparable since the detector housing is covered in germanium shielding that was handled as or more carefully than the detectors. Furthermore, m_{γ} is is surprisingly low for this case. Nevertheless, it may be necessary to accept this conclusion: it is corroborated by a simpler argument made below, the high value of F_{γ} in the case where R^{γ}_{B5} is discarded seems more unreasonable, and ^{210}Pb on the Detector Interface Boards and IR blockers or in the brass screws for the germanium shielding may serve as the source of the high $F_{i\beta}$ flux.

One could also extend the above model to include cQShare events by adding three new

	Inferred value		Contribution to $R^s_{\rm B5}$	
s_s	$F_{i\beta}$	$F_{s\beta}$	$F_{i\beta}$	$F_{s\beta}$
0.0	44.4	23.1	0.0	23.1
0.1	49.0	18.2	4.9	18.0
0.2	55.5	12.0	11.1	12.0
0.3	63.3	4.1	19.0	4.1

Table 8.12: $F_{i\beta}$ and $F_{s\beta}$ via robust method. The values of $F_{i\beta}$ and $F_{s\beta}$ and their contribution to $R^s_{\rm B5}$ as a function of s_s are shown. All fluxes and rates in 10^{-3} cm^{-2} d^{-1}.

parameters, m^S_γ, m^S_β and s^S_s, to allow for different multiple-scatter fractions and self-shielding reduction. It would not be self-consistent to add new flux parameters also. This extension is not performed here.

Because of the degeneracy between discarding $R^\gamma_{\rm B5}$ or $R^\gamma_{\rm B6}$, it is useful to step back and ask what robust conclusions can be made independent of the above fitting procedure:

- First, F_γ does not contribute significantly to the singles rates. On physical and geometrical grounds, m_γ is expected to be large, certainly bigger than ~ 0.5. Therefore, most of F_γ appears in $R^\gamma_{\rm B5}$ and $R^\gamma_{\rm B6}$ and the contribution of F_γ to $R^s_{\rm B5}$ and $R^s_{\rm B6}$ is small compared to $R^\gamma_{\rm B5}$ and $R^\gamma_{\rm B6}$. But, $R^\gamma_{\rm B5}$, and even $R^\gamma_{\rm B6}$, are comparable to or smaller than $R^s_{\rm B5}$ and $R^s_{\rm B6}$, so the contribution of F_γ to $R^s_{\rm B5}$ and $R^s_{\rm B6}$ is small.

- Second, F_γ is ~ 0.010 cm^{-2} d^{-1}, since most of F_γ appears in $R^\gamma_{\rm B5}$ and $R^\gamma_{\rm B6}$.

- Third, $F_{i\beta}$ is large compared to $F_{s\beta}$. Consider only the total rates

$$R^{s+\beta}_{\rm B5} \equiv R^s_{\rm B5} + R^\beta_{\rm B5} = F_{s\beta} + F_{i\beta} s_s \tag{8.19}$$

$$R^{s+\beta}_{\rm B6} \equiv R^s_{\rm B6} + R^\beta_{\rm B6} = F_{s\beta} + F_{i\beta}\left[\frac{1}{2} + \frac{1}{2} s_s\right] \tag{8.20}$$

  These rates are insensitive to the multiple-scatter fractions and are determined only by s_s, which is robustly estimated to be $s_s \sim 0.2$ purely by geometry (see Chapter 5). They are also insensitive to the problematic photon-induced flux. These rates are $R^{s+\beta}_{\rm B5} = 0.0231 \pm 0.0038$ cm^{-2} d^{-1} and $R^{s+\beta}_{\rm B6} = 0.0453 \pm 0.0054$ cm^{-2} d^{-1}. Solving for the fluxes with $s_s = 0.2$ gives $F_{s\beta} = 0.0120$ cm^{-2} d^{-1} and $F_{i\beta} = 0.0555$ cm^{-2} d^{-1}. The resulting contributions to $R^s_{\rm B5}$ are 0.0120 cm^{-2} d^{-1} and 0.0111 cm^{-2} d^{-1}, respectively. This result depends somewhat on s_s, as is illustrated in Table 8.12. The table shows that it generally holds that $F_{i\beta}$ is large compared to $F_{s\beta}$, but which one dominates $R^s_{\rm B5}$ is not well constrained.

So, what can be concluded? First, it is clear that self-shielding has a significant effect. The BLIP5 total electron rate is lower than the BLIP6 rate by about a factor of 2. Second, a significant fraction of electrons, of order 50% for BLIP5, can be vetoed by multiple scattering. These two effects yield a nontrivial factor-of-four rate reduction in going from the BLIP6 total electron rate of 0.063 cm^{-2} d^{-1} to the BLIP5 single-scatter-electron rate of 0.015 cm^{-2} d^{-1}. Third, surface

contamination of the detectors and the housing appear to be the most important residual electron sources. While the fit discarding R^{γ}_{B6} yields a F_{γ} contribution of 0.0077 cm^{-2} d^{-1} to the total $R^{s}_{\mathrm{B5}} = 0.0154$ cm^{-2} d^{-1} rate, it is hard to believe that m_{γ} is as low as 0.5. With $m_{\gamma} \sim 0.8 - 0.9$, the contribution of photon-induced electrons is negligible. Unfortunately, the relative importance of $F_{i\beta}$ and $F_{s\beta}$ to R^{s}_{B5} is not clearly determined. However, s_s will be increased in the future because ZIP detectors do not have curved edges and can be packed with separations as small as 1 mm. Thus, the contribution of $F_{i\beta}$ will be decreased and the most important electron source to address will be detector surface contamination.

To further evaluate what has been learned, it is useful to compare to rates observed in previous runs. The electron rate observed by the Run 18 Ge BLIP detectors, averaged between 20 and 100 keV, is 0.27 keV^{-1} kg^{-1} d^{-1} [8]. Converting to an areal rate summed between 20 and 100 keV gives 0.048 cm^{-2} d^{-1}. It is not clear from [8] whether this rate includes multiple-scatter events. However, for the Run 18 geometry, multiple scattering of non-penetrating particles was prevented by the individual copper detector housings; furthermore, the detectors were much further apart than in the Run 19 geometry. It is assumed here that the Run 18 0.048 cm^{-2} d^{-1} rate should be compared to the total Run 19 electron rates. The total Run 19 BLIP5 and BLIP6 areal electron rates, summed between 20 and 100 keV (rather than 10 and 100 keV), are 0.0212 ± 0.0037 cm^{-2} d^{-1} and 0.0524 ± 0.0058 cm^{-2} d^{-1}, respectively. The BLIP6 rate is thus comparable to and the BLIP5 rate about 40% of the Run 18 rate. Clearly, self-shielding has a significant effect.

However, it is surprising that the BLIP6 total rate is the same as the Run 18 rate in spite of its top side being self-shielded. Moreoever, the cQShare and cQoutOnly rates are substantially higher. These elevated rates are mysterious, since a number of precautions were taken for Run 19 to reduce contamination of the detectors. Furthermore, the Ge shielding lining the detector package was etched and maintained in at least as clean an environment as the Run 18 detectors. Because of the large difference between the cQinOnly rate and the cQShare and cQoutOnly rates, it is likely the new electron source is not surface contamination of the detectors, which should be uniformly distributed. One possibility is that ^{210}Pb is present on the circuit boards providing electrical connections to the detectors. This possibility is discussed later in the context of observation of a 46-keV photon line in the cQoutOnly data.

Because of this possible new electron source, it is not possible to determine the effect of the Ge shielding. It is possible that, though the shielding may have attenuated the $F_{i\beta}$ flux from the copper housing, the effect would be masked by the rate due to the new source. It is also possible that the shielding had no effect. As noted above, the sensitivity of the detectors to electrons emitted by the housing will be further decreased in the future by improved self-shielding, so the effect of the Ge shielding has little bearing on the future.

Finally, it is important to realize that the present BLIP5 single-scatter-electron rate roughly meets the CDMS II goals. The R^{s}_{B5} rate corresponds to 0.05 keV^{-1} kg^{-1} d^{-1} and appears to be fairly flat in energy down to about 20 keV (see Figure 8.10). The CDMS II electron-rate goal is 0.02 keV^{-1} kg^{-1} d^{-1} at 15 keV with 95% electron rejection [148]. As mentioned earlier, ZIP detectors have demonstrated approximately 99% electron rejection above 15 keV, so the electron-rate goal can be relaxed by a factor of 5 to 0.1 keV^{-1} kg^{-1} d^{-1}. The relaxed rate goal is certainly met above 20 keV and is almost met in the 12–20-keV bin (see Figure 8.10). Furthermore, it is expected that some improvement in surface contamination will be achieved in switching from BLIPs to ZIPs: the ZIP fabrication environment is much cleaner and only photoresist is used for masking. Analysis

of the type presented here should certainly be pursued to solidify these conclusions and, hopefully, to indicate how to push the electron rates well below the CDMS II goals, but significant reduction of surface contamination does not appear to be necessary.

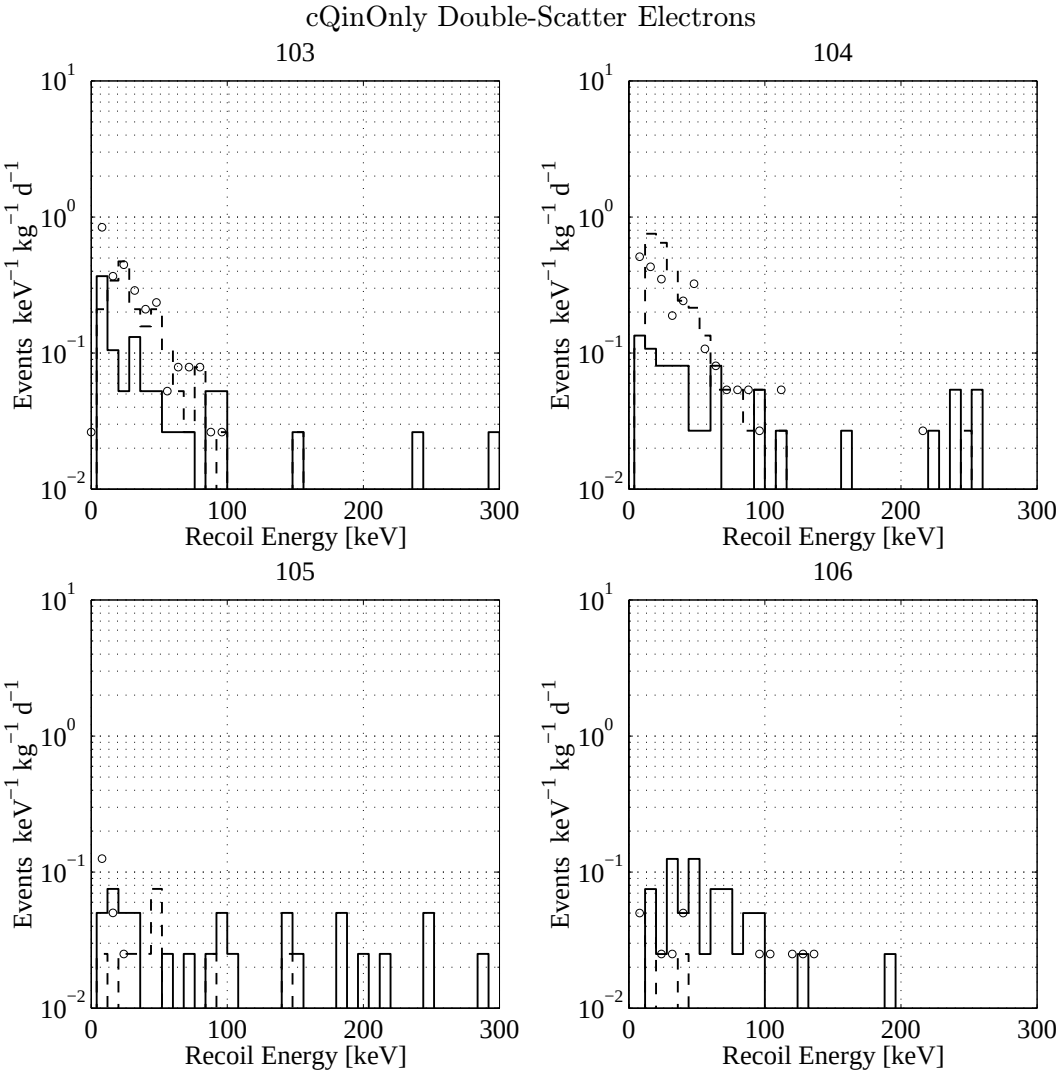

Figure 8.34: Recoil-energy spectra for cQinOnly electrons belonging to veto-anticoincident double-scatter events. Solid: other scatter is a photon above 10 keV. Dashed: other scatter is an electron above 10 keV. Circles: other scatter is below 10 keV (photon/electron discrimination not possible).

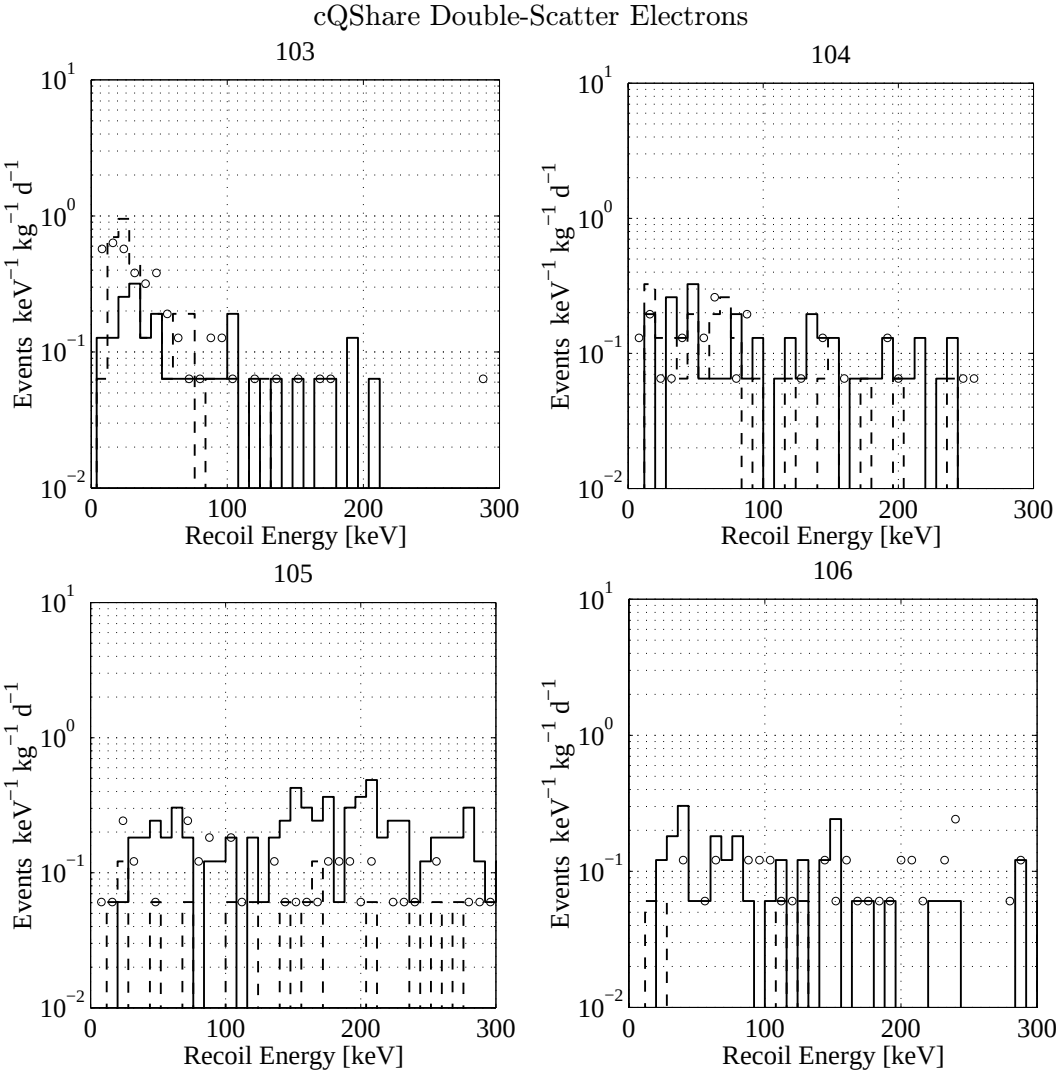

Figure 8.35: Recoil-energy spectra for cQShare electrons belonging to veto-anticoincident double-scatter events. Legend as in Figure 8.35.

8.5.4 Spectral Lines

Photon lines can be used to constrain the presence of radioactive contaminants. Figures 8.36, 8.37, and 8.38 show spectra of "prgamma" for veto-anticoincident single-scatter-photon events. "prgamma," defined in Chapter 6, is the phonon energy divided by $1 + V_b/(3\text{ eV})$; it gives the recoil energy for photons, for which the ionization energy and recoil energy are equal. The advantage of using prgamma is that it is less noisy than the calculated recoil energy. Spectra for double-scatter events are also shown, though they exhibit no photon lines. The double-scatter spectra exhibit large increases in rate below 20 keV. The reason for this is not known.

All the spectra exhibit the 10.4-keV Ga X-ray. This X-ray arises from two sources. Cosmogenic activation while the detectors are above ground creates ^{68}Ge, which decays by electron capture to ^{68}Ga with a 270.8 d half-life. Electron capture removes a K-shell electron; the X-ray is emitted when the empty state is filled. Thermal-neutron capture during neutron calibrations produces ^{71}Ge, which decays by electron capture with a 11.4 d half-life, producing a high rate of 10.4-keV Ga X-rays immediately following the neutron calibrations. For both types of X-rays, the source is distributed uniformly through detectors and so the ionization yield is not suppressed. Some work has been done by Richard Schnee to measure the gallium-X-ray rate as a function of time, but is not presented here. It would be interesting to determine the cosmogenic-activation rate using the rate of the 270.8 d decay, but this is also not done here. Though not visible due to the axis ranges of the spectra, BLIP3 has a gallium rate approximately 30% higher than the other detectors. This is probably due to enhanced activation during the neutron calibration: though the neutron directions are highly randomized by scattering, some anisotropy may remain, yielding a higher thermal-neutron flux on BLIP3 than the other detectors. The neutron-capture cross section rises at low energy, so thermal neutrons have much shorter penetration lengths than do the higher-energy neutrons responsible for nuclear-recoil events. Finally, the rise in rate below the gallium line in BLIPs 3 and 4 should not be ascribed any significance; as mentioned earlier, discimination between photons and electrons becomes difficult below 10 keV.

Also visible in all the spectra is a line at 67 keV. This line was first seen in Run 18 data [8]. As first noted by Rick Gaitskell, this line arises from deexcitation of ^{73m}Ge produced by thermal-neutron capture on ^{72}Ge, giving an energy of 66.7 keV for the line [149]. In Andrew Sonnenschein's dissertation [8], the rate of this line is used to measure the thermal-neutron background.

Of most interest is the 46-keV line visible in the cQoutOnly spectra. It may also be present in the BLIP6 cQinOnly spectrum, but at a much lower rate and significance. This line arises from the beta decay of ^{210}Pb. The decay has Q = 63.1 keV, but 89% of the time the electron is emitted with 16.6-keV endpoint and the daughter ^{210}Bi nucleus is left in an excited state. Most of the time, the excited ^{210}Bi nucleus decays by internal conversion, ejecting electrons from atomic states. A small fraction of the time (4.05% of all ^{210}Pb decays), the ^{210}Bi nucleus decays by emission of a 46.5-keV photon. This 46-keV line thus indicates the amount of ^{210}Pb present in the surroundings. The absence of these photons, or presence at a much lower rate, in the BLIP3 and BLIP6 cQinOnly spectra suggests that they are emitted by ^{210}Pb present on the sides, but not the endcaps, of the detector housing. One possibility is the Detector Interface Boards (DIBs) and IR blockers: there are a number of electronic components and a good deal of solder on these assemblies. Another possibility is the brass screws holding the germanium shielding in place. The faces of BLIPs 3 and 6 would not see these sources. If ^{210}Pb is directly visible to the detectors, then a large associated electron flux is expected directly from the beta decay. Such a rate may explain the elevated electron

rates in the outer electrodes. The 46-keV line has an intensity of roughly 10 kg^{-1} d^{-1} in each of the detectors. This implies a ^{210}Pb decay rate of at least 40 d^{-1}; of course, a significant fraction of the 46-keV photons may not be emitted toward the detectors, so the actual rate of decays may be higher. Every decay yields at least one electron. An areal event rate of 0.1 cm^{-2} d^{-1} in electrons corresponds to a total event rate of 4 d^{-1} for cQoutOnly, so less than 10% of the emitted electrons would have to reach the detectors and deposit more than 10 keV, which is plausible. To conclusively test this hypothesis, a simulation of a ^{210}Pb source on the DIBs or in the brass screws would be useful. Such work is beyond the scope of this dissertation.

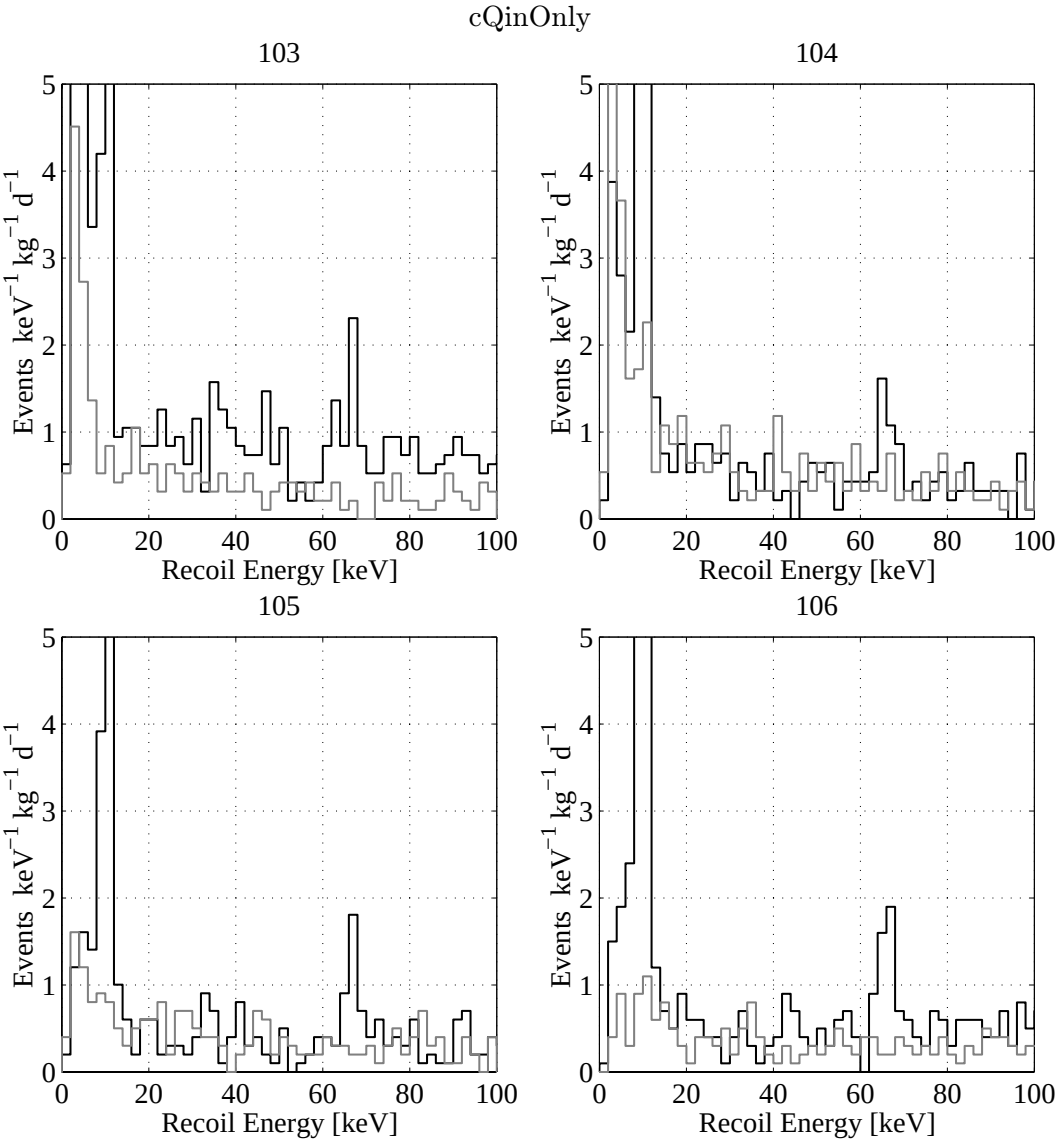

Figure 8.36: Photon "prgamma" spectra for veto-anticoincident cQinOnly photon events. Solid: single scatters. Light: double scatters.

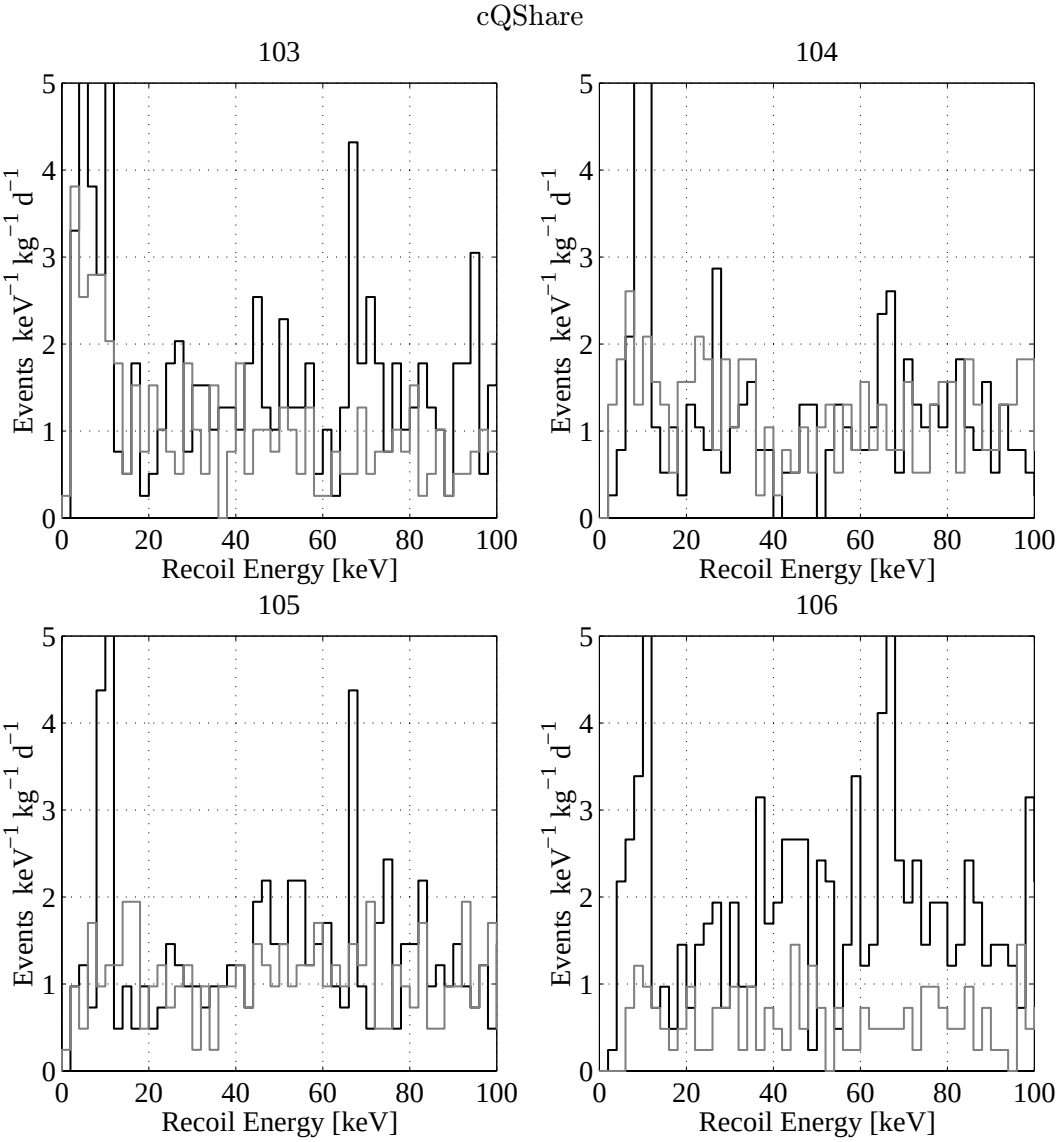

Figure 8.37: Photon "prgamma" spectra for veto-anticoincident cQShare photon events. Solid: single scatters. Light: double scatters.

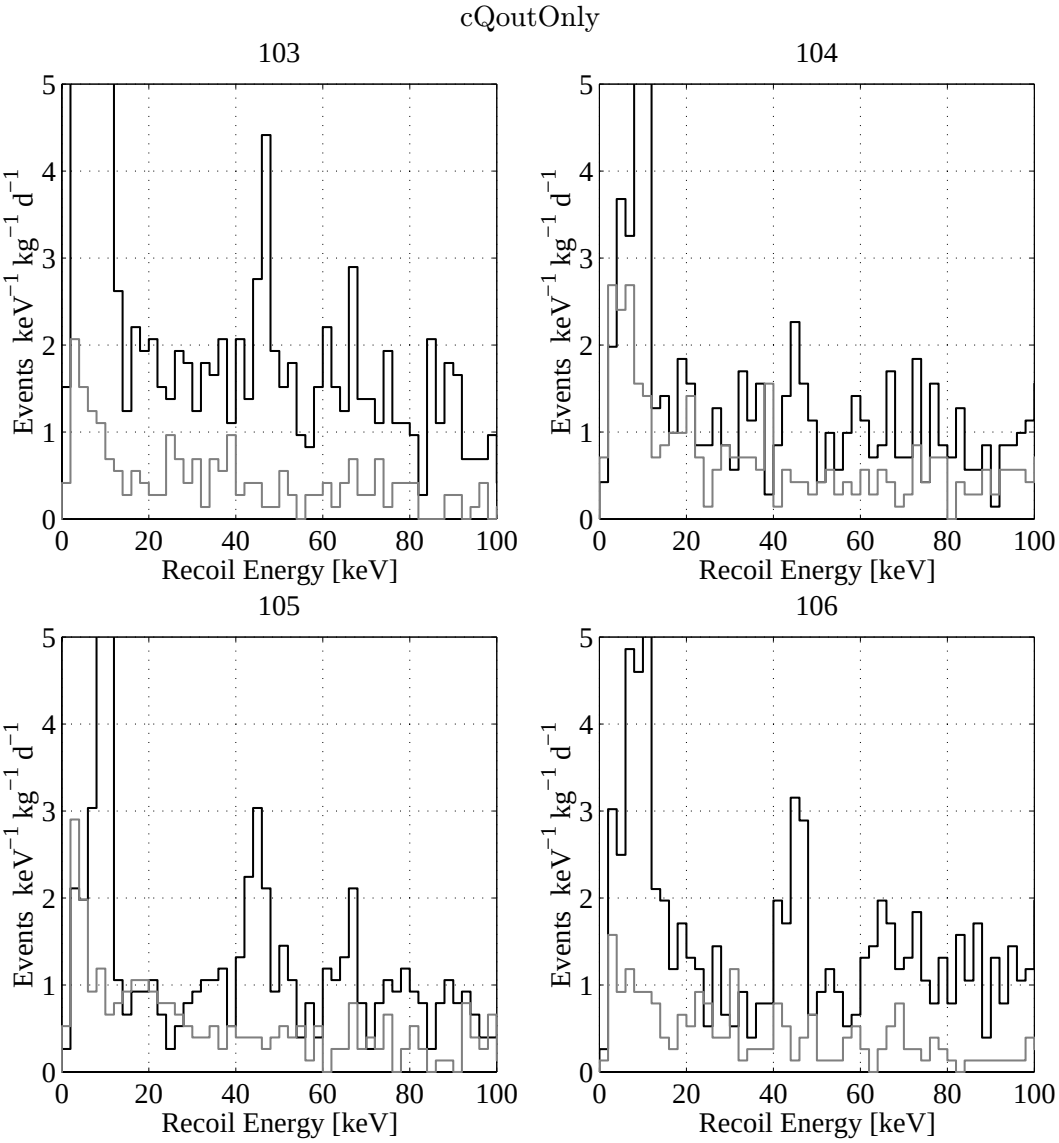

Figure 8.38: Photon "prgamma" spectra for veto-anticoincident cQoutOnly photon events. Solid: single scatters. Light: double scatters.

8.5.5 Tritium

The story of tritium contamination observed in previous runs in both the thermistors and detectors is somewhat of a footnote to the results presented in this dissertation; however, it is useful to summarize the situation and indicate what light Run 19 can shed on the issue.

Concurrent with discovery of tritium in the thermistors (see Chapter 4), a tritium beta-decay spectrum was also observed in the BLIP1 crystal-interaction data set with an overall rate of 120 d^{-1}, or 730 kg^{-1} d^{-1} [111]. These events suffer from suppressed ionization collection and thus could be due to an external source, detector surface contamination, or tritium in the dead layer itself. The observation of tritium in both the thermistor-interaction and crystal-interaction data suggested the possibility that the tritium either diffused from the thermistors into the detector dead layer or evaporated out of the thermistors and plated out on the detector surface during eutectic bonding of the thermistors. The excellent fit to the energy spectrum also suggested the source was not external — one would expect some distortion due to backscattering, though this was not investigated in detail at the time. Another possibility is that tritium was present in the ion beam used to implant boron into the detector surfaces for creation of the ionization electrodes.

A number of tests made it clear that, in fact, the tritium source was external, residing on or in the surfaces of the detector housing. However, the series of tests that finally led to this conclusion was not direct and the chronology is somewhat strange. First, to test the possibility that the tritium was either deposited on the detector surface or had diffused into the dead layer during eutectic bonding, a deep (30-μm) etch was performed on BLIP1. This removes the entire dead layer. The detector was reimplanted, so the implantation hypothesis would not be tested. BLIP2, a detector not previous operated at SUF, but fabricated in the same way as BLIP1, also was subjected to this procedure. BLIPs 1 and 2 were operated at SUF during Runs 18 and 19. These data sets are discussed in detail in Andrew Sonnenschein's dissertation [8]. In Run 18, the event rates due to tritium in BLIPs 1 and 2 were 120 d^{-1} and 40 d^{-1}, equal to and one-third of the Runs 15/16 BLIP1 rates, respectively. This test demonstrated that the thermistor hypothesis could not explain the data. Contamination by reimplantation, detector handling, or in the detector housings remained viable. Between Runs 18 and 19, the two detectors' housings were exchanged. In Run 19, the ratio of rates reversed itself, thus indicating the source of tritium was the detector housings.

Concurrently, a second test was done to see if it was possible to remove the tritium from the thermistors prior to eutectic bonding. Two thermisors were heated to 550 C and 650 C, respectively, for 2 hours to evaporate the tritium [130]. These thermistors were measured at SUF in Run 17, yielding tritium reductions of 99% and 50% for the 650 C and 550 C thermistors, respectively [117].

Thus, while there is indeed tritium in the thermistors that can be removed by baking, this contamination is not related to the source of the crystal-interaction tritium spectrum seen in BLIPs 1 and 2. The thermistor-baking result was available in October, 1997, while the disproval of the thermistor hypothesis was not available until May, 1998. Fabrication of the Run 19 detectors began in early 1998, so the precaution of baking the thermistors *prior* to bonding to the detectors was taken. An entire wafer of NTD Ge was baked at 650 C in the same way as the test device above and diced into thermistors. These thermistors were eutectically bonded to the detector crystals in the standard way. A second precaution taken for Run 19, against both tritium and higher-energy electrons, was to cover the inside of the detector package with high-purity germanium, as has been described earlier.

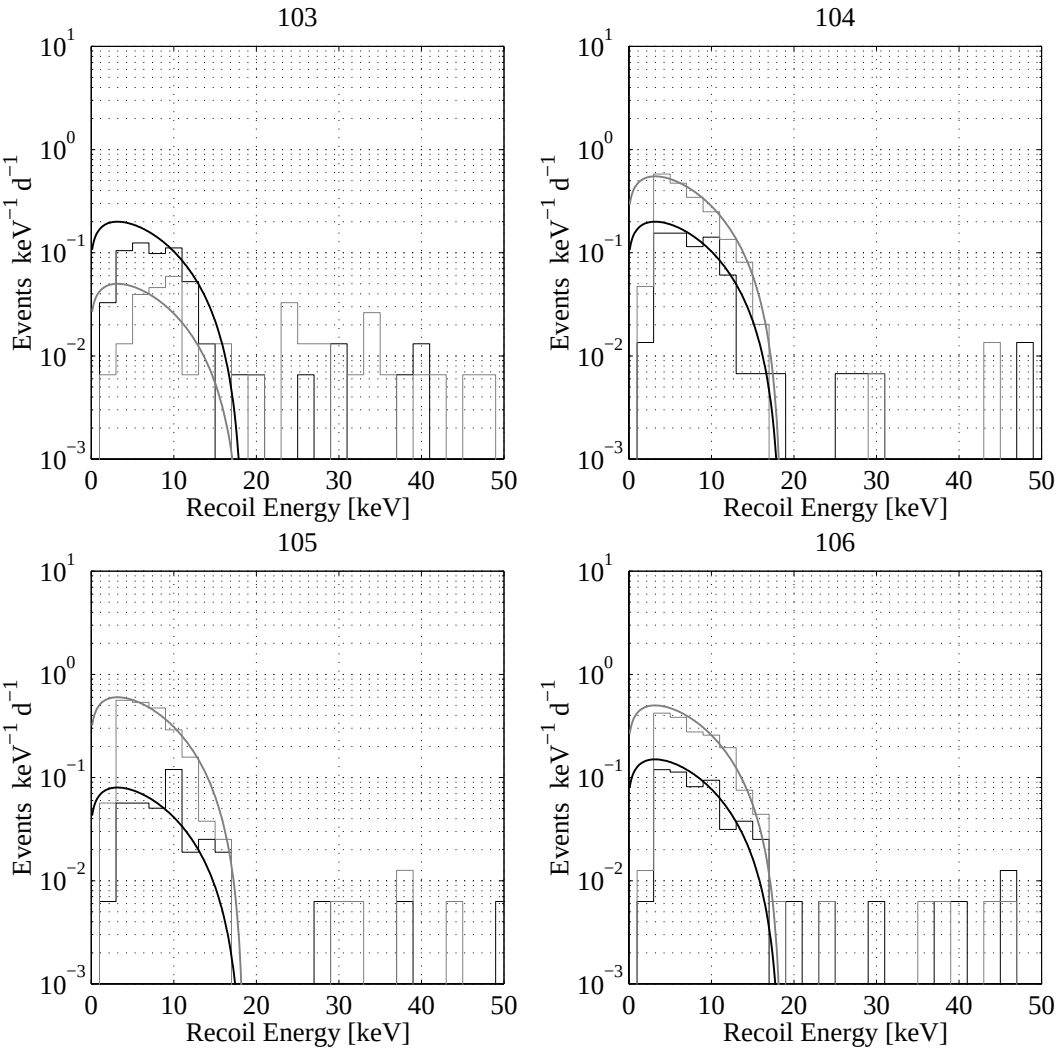

Figure 8.39: Thermistor event spectra. Dark: P1. Light: P2. The fitted phonon energies have been scaled down by an empirical factor of 2.7. The curves indicate tritium beta-decay spectra with amplitudes chosen to roughly match the spectra.

Run 19 can add two pieces of information to the story. First, the Run 17 tritium reduction is confirmed. Energy spectra for thermistor events are shown in Figure 8.39; the typical inferred total tritium beta-decay rate is in the range 0.5 d^{-1} to 6 d^{-1}, with large variations among thermistors. The tritium rate observed in Runs 15 and 16 was 570 d^{-1}, so a reduction factor of at least 100 is confirmed. Second, no tritium is apparent in the Run 19 crystal-interaction data set either; a limit of order 1000 times lower than the 120 d^{-1} observed in BLIP1 can be placed. Given the use of close packing and germanium shielding in Run 19, it may be possible for the Run 19 detector housings to be contaminated with tritium without producing an observable rate in the detectors. The Run 20 detector package eliminates the germanium shielding, so it should be possible to test this hypothesis. However, it is certainly possible that tritium contamination is not a general problem and that, with the the BLIPs 1 and 2 detector packages, we simply were unlucky.

8.6 BLIP3 Electron Misidentification

In terms of systematics, it is important to determine whether the BLIP3 single-scatter data are consistent with the electron-misidentification expectations from the electron calibration. This checks whether the ionization yield for the double-scatter events used in the electron calibration is systematically different from the single-scatter data to which the calibration is applied, clearly an important issue. Expected electron-misidentification numbers for the BLIP3 single-scatter data are calculated in the same way as is done for the BLIPs 4/5/6 data in Section 8.2.1 and are shown in Table 8.13. The observed numbers of nuclear-recoil candidates are also shown. Only about 4 to 5 true nuclear recoils are expected in the cQinOnly sample, small compared to the observed numbers. The numbers are consistent with the expected misidentification. One can also make the important conclusion that, had the BLIP3 single-scatter data set been used as the electron calibration, the expected numbers of misidentified events calculated in Sections 8.2.1 and 8.2.2 would not be significantly changed.

Event set	N_β^b	N_l^c	N_β^c	$\mu_{l,90}^b$	$\mu_{l,10}^b$	$\mu_{l,50}^b$	N_{NR}^b
cQinOnly							
10 – 30 keV	389	1	46	35.2	4.5	14.2	19
30 – 100 keV	475	0	39	32.0	1.2	8.4	14
10 – 100 keV	864	1	85	42.3	5.4	17.1	33
cQShare							
10 – 30 keV	87	1	35	10.3	1.3	4.2	4
30 – 100 keV	155	0	27	14.9	0.6	4.0	3
10 – 100 keV	242	1	62	15.7	2.1	6.6	7

Table 8.13: BLIP3 cQinOnly and cQShare single-scatter-electron misidentification estimates, calculated by method discussed in Section 8.2.1. The first column is the number of single-scatter-electron events observed for the given energy bin. The second two columns list the calibration data used, which come from the BLIP3 column of Table 4.6. The fourth column is the rigorously calculated Bayesian 90% CL upper limit on the number of misidentified single-scatter electrons. Because of the large values of N_β^c and N_β^b, these values are very close to what one gets by calculating a Bayesian 90% CL upper limit on β_β from the electron calibration and multiplying by N_β^b. The fifth and sixth columns give 10% CL and 50% CL upper limits on the number of misidentified electrons to give an idea of the expected range. The final column lists the observed numbers of nuclear-recoil candidates. No attempt has been made to subtract the expected numbers of single-scatter neutrons from these data, but, roughly, there should be 4 to 5 and 2 to 3 neutrons between 10 and 100 keV in the cQinOnly and cQShare data, respectively.

8.7 Veto-Coincident Data

For completeness, I present the veto-coincident data in this section. A detailed discussion of the veto-coincident-neutron spectra is presented in Section 7.5. Otherwise, these data have no bearing on the WIMP-search analysis and on residual background-particle rates, so no serious analysis is attempted. However, some features deserve brief discussion.

Ionization yield plots for cQinOnly, cQShare, and cQoutOnly events for single and multiple scatters are shown in Figures 8.40, 8.41, 8.42, 8.43, 8.44, and 8.45. The most disconcerting feature of these plots is the distribution of low ionization-yield events that follow the shape of the ionization-search threshold. In the single-scatter data, these events are only apparent in the BLIP3 data. In the multiple-scatter data, the rate of such events is progressively higher in the cQShare and cQoutOnly data sets. Events with saturated detectors have been removed, so electrical cross-talk is ruled out as the cause of these events. Cross-talk is also disfavored by the fact that the events have significant phonon energy. The variation of the rate of such events with ionization-partition cut and detector (with BLIP6 generally having the fewest such events, except possibly in the multiple-scatter cQoutOnly data) strongly suggests a physical origin. The absence of such distributions in the veto-anticoincident data corroborates a physical origin for the events. At present, no explanation has been found.

It must be emphasized that the presence of such events does not invalidate the analyis of the muon-coincident nuclear recoils presented in Section 7.5. First, the cQinOnly single-scatter data do not exhibit such events (with the possible exception of BLIP3), so the single-scatter nuclear recoils are free of contamination. Second, recall that extremely restrictive cuts were made to select double-scatter nuclear recoils; these cuts remove these low-yield populations in the double-scatter data set.

Recoil-energy spectra for the veto-coincident data are shown in Figures 8.46, 8.47, and 8.48. Since the low-yield distributions seen in the ionization-yield plots merge into the ionization-search threshold, spectra of the phonon energies for events with only noise charge are also presented, in Figure 8.49. Since these events have only noise charge, the observed ionization energy is not accurate. The phonon energy presumably gives the correct recoil energy for such events. The entire detector volume is used for normalization of the noise-charge spectra. Note that, because these events have only noise charge, it is not possible to determine if they are coincident or anticoincident with the veto. It is likely, based on the rates and the merging of the above-noise-charge distribution into the noise-charge distribution, that the noise-charge events are dominantly veto-coincident.

The cQinOnly and cQoutOnly spectra and, especially, the noise-charge event spectra, exhibit steep rises at low energy due to the low-yield events. The relative single- and double-scatter rates reflect the geometry, with BLIPs 4 and 5 exhibiting higher double-scatter photon fractions than BLIPs 3 and 6. Also, compared to the veto-anticoincident data, the electron double-scatter fractions are quite high, indicating most veto-coincident electrons are produced in showers or are ejected from the detectors and surroundings. This is especially true for the cQoutOnly data. The rise in cQinOnly and cQoutOnly photons and the drop in cQShare photons at low energy, as compared to their fairly flat behavior in the veto-anticoincident data, may indicate that a large part of the low-energy-photon flux is due to fluorescence photons — such photons have energies $\lesssim$ 100 keV and thus are unlikely to scatter twice in a single detector, causing the cQShare fraction to drop at low energy.

Little more can be said at this point, except to note that, since the low-yield events appear to be purely veto-coincident and do not contaminate the single-scatter cQinOnly and double-scatter neutron samples, it is not a high priority to understand their origin.

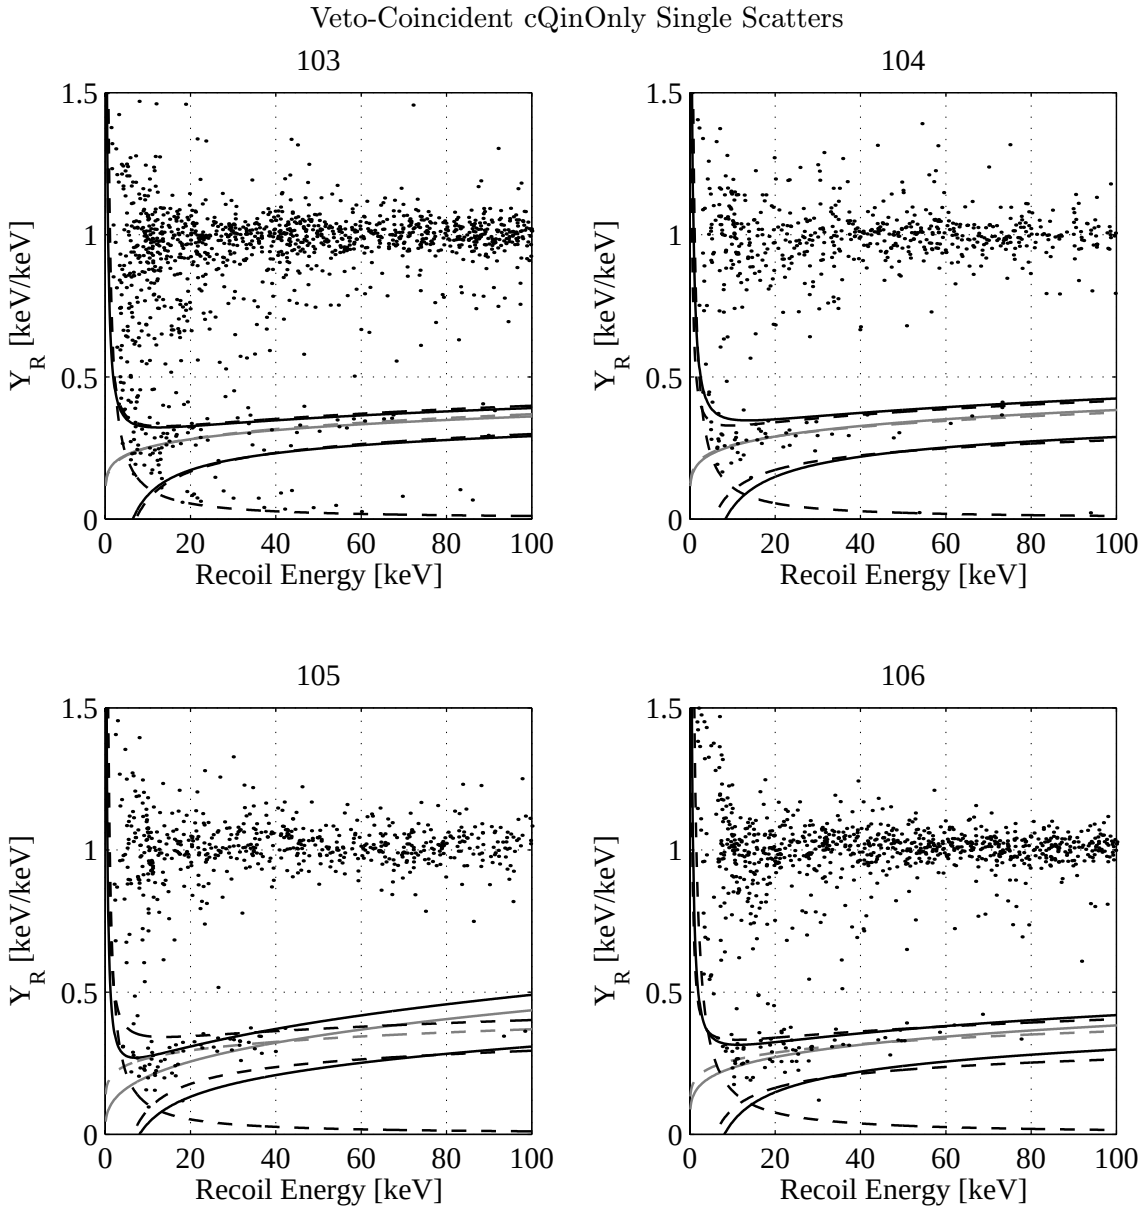

Figure 8.40: Ionization yield vs. recoil energy for veto-coincident cQinOnly single-scatter events. A random 10% of the data is shown. Hyperbolic dashed line: ionization-search threshold. Light solid line: center of nuclear-recoil band, pre-April 3. Light dashed line: center of nuclear-recoil band, post-April 3. Dark solid line: nuclear-recoil-acceptance region, pre-April 3. Dark dashed line: nuclear-recoil-acceptance region, post-April 3.

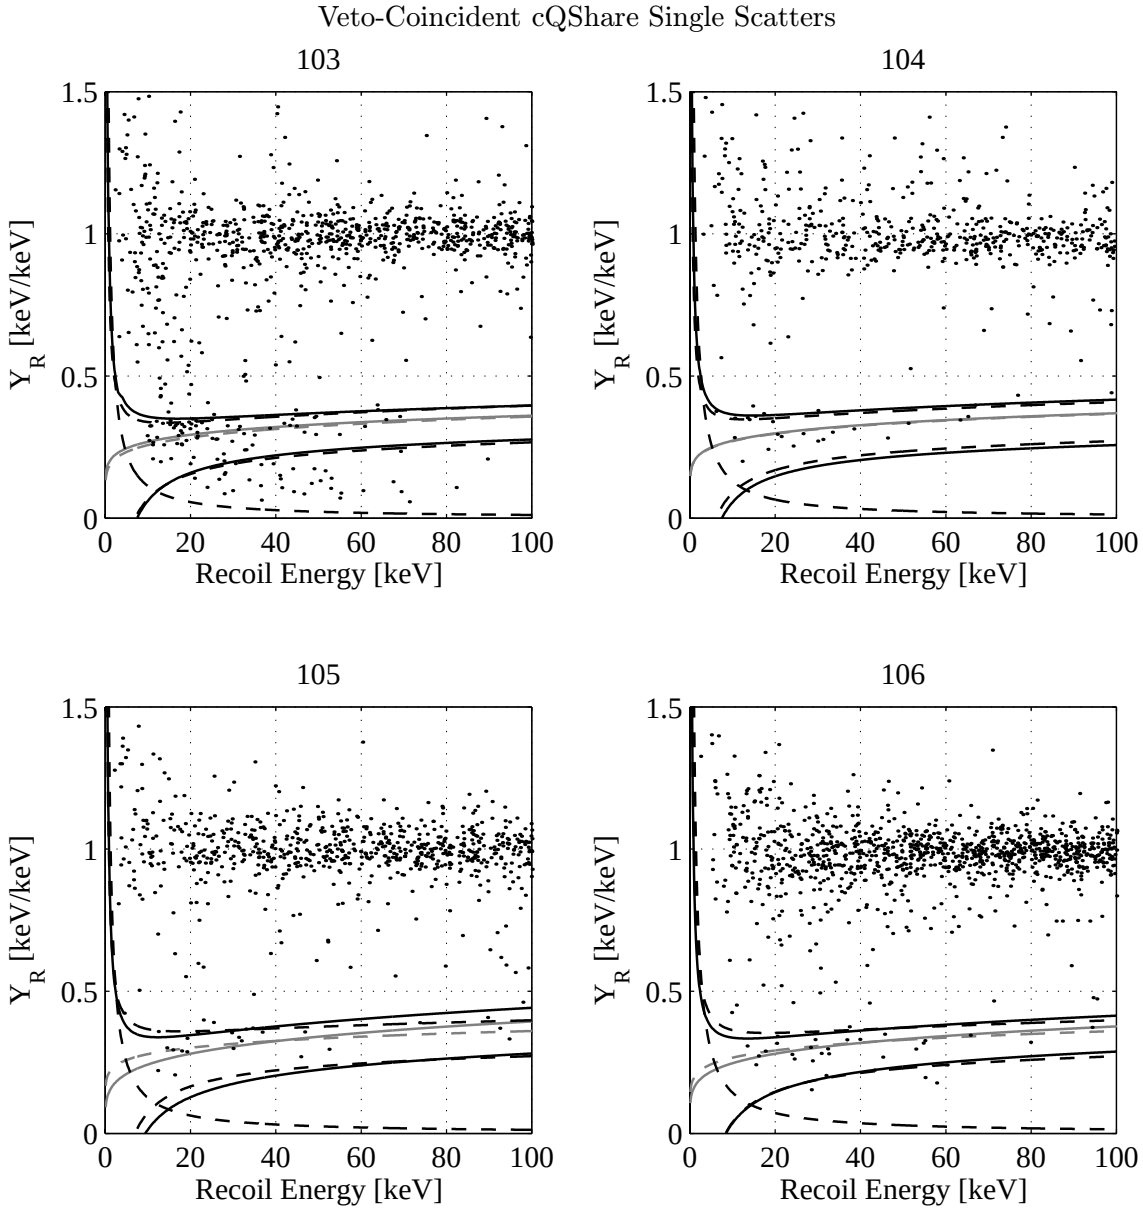

Figure 8.41: Ionization yield vs. recoil energy for veto-coincident cQShare single-scatter events. A random 10% of the data is shown. Legend as in Figure 8.40.

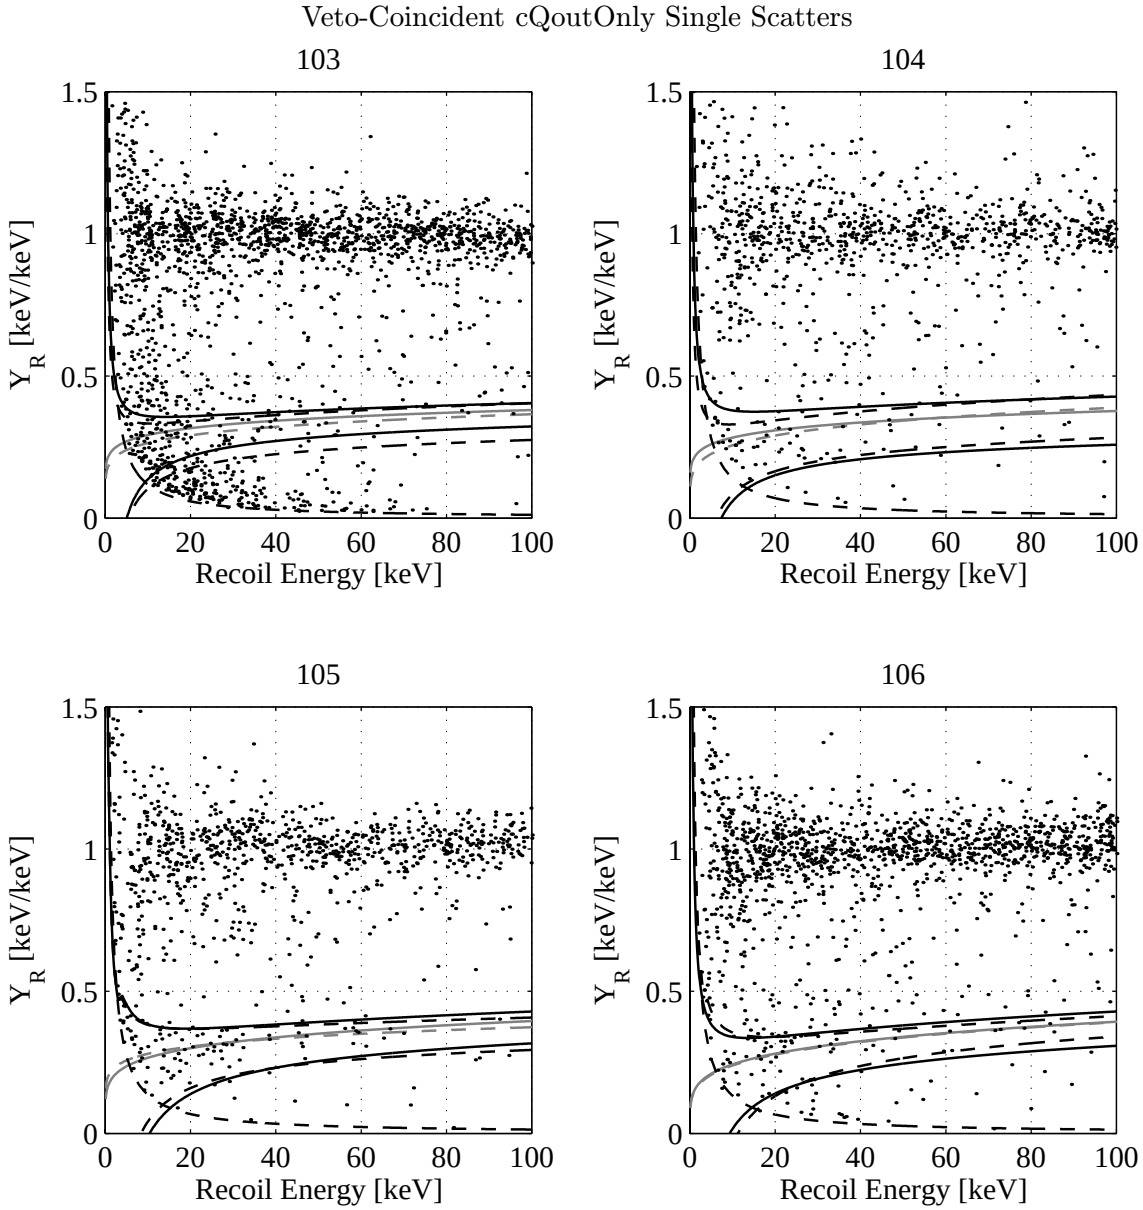

Figure 8.42: Ionization yield vs. recoil energy for veto-coincident cQoutOnly single-scatter events. A random 10% of the data is shown. Legend as in Figure 8.40.

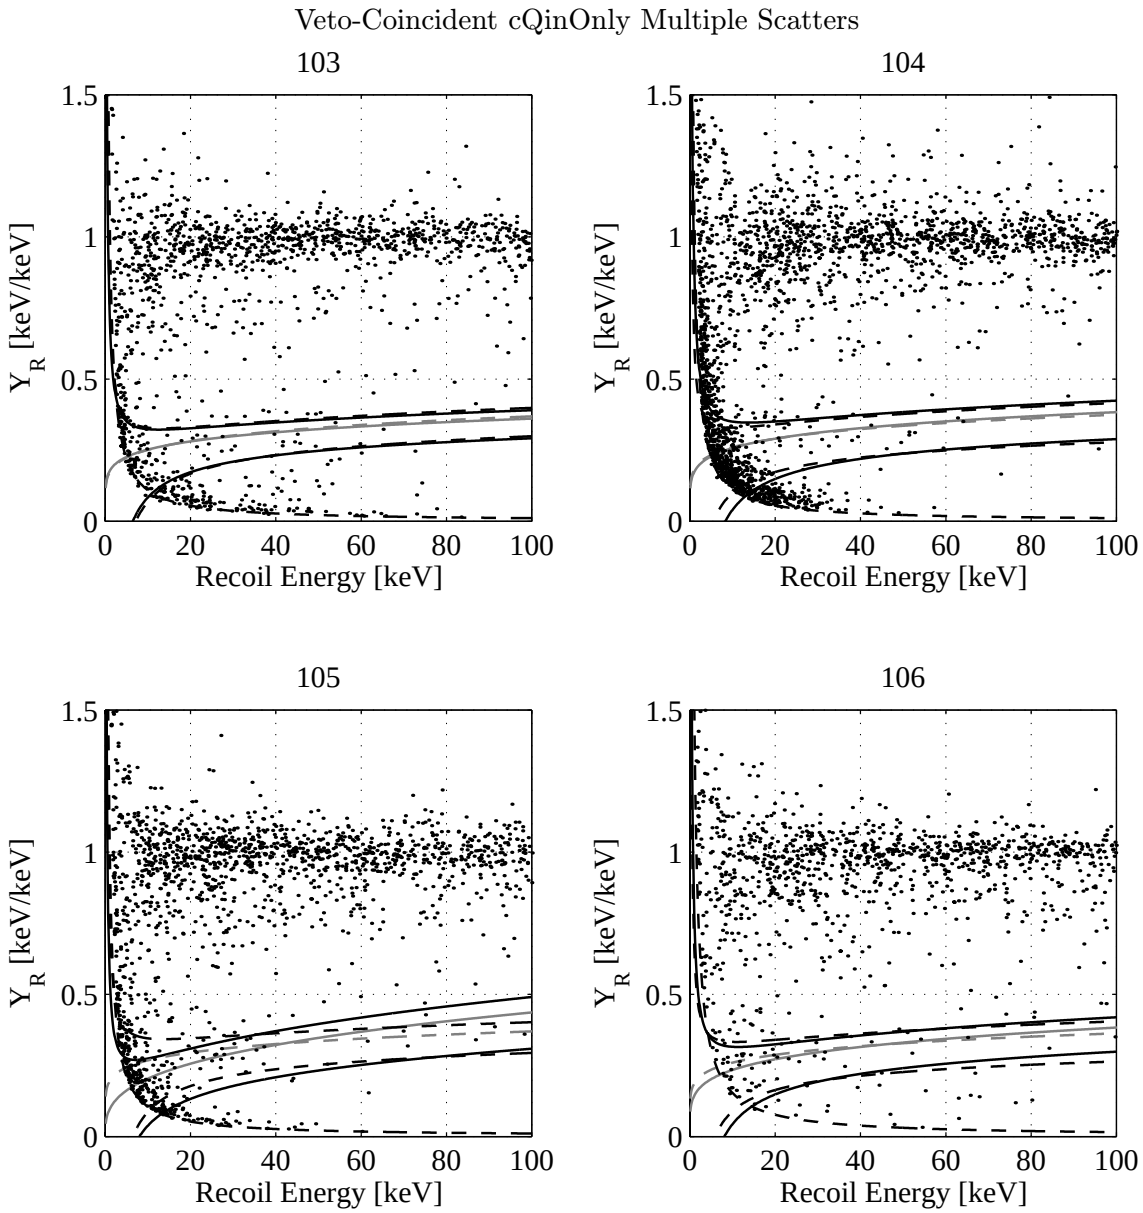

Figure 8.43: Ionization yield vs. recoil energy for veto-coincident cQinOnly multiple-scatter events. Events in which any detector is saturated have been removed. A random 10% of the data is shown. Legend as in Figure 8.40.

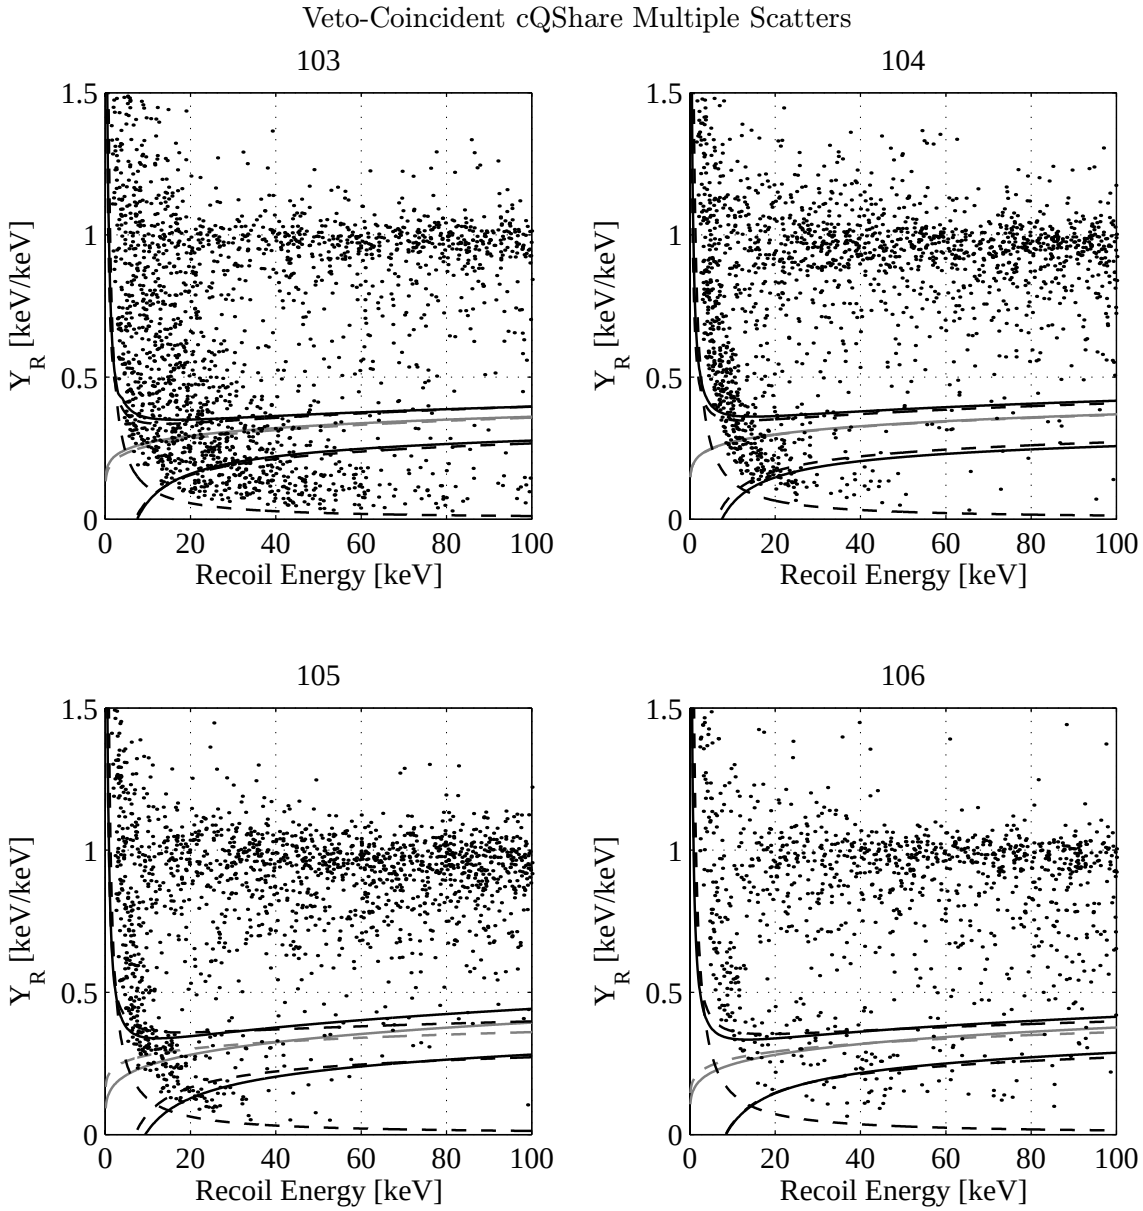

Figure 8.44: Ionization yield vs. recoil energy for veto-coincident cQShare multiple-scatter events. Events in which any detector is saturated have been removed. A random 10% of the data is shown. Legend as in Figure 8.40.

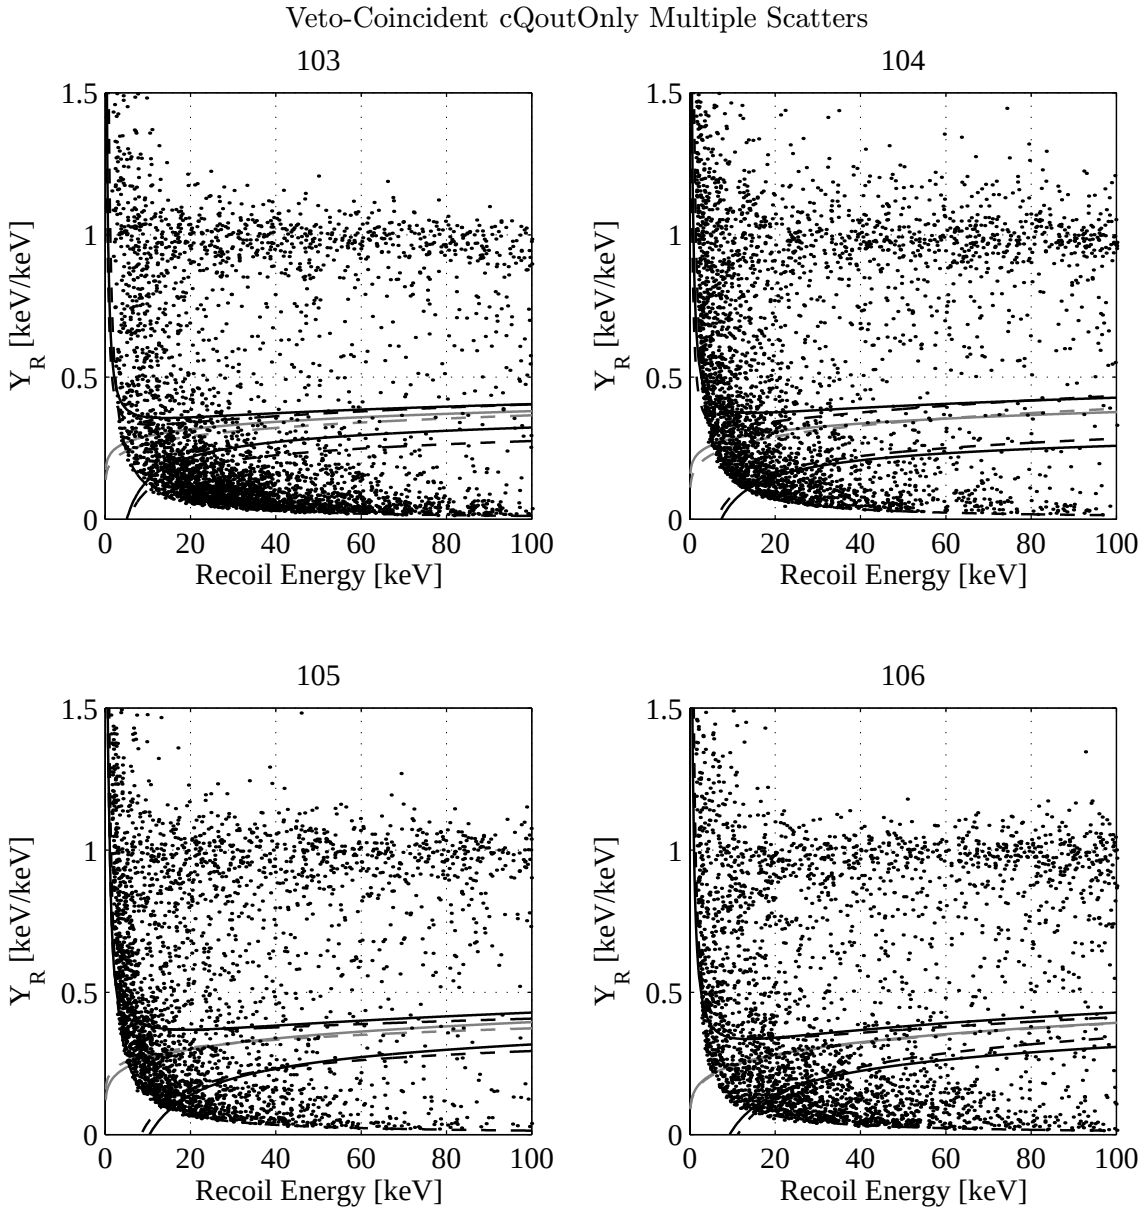

Figure 8.45: Ionization yield vs. recoil energy for veto-coincident cQoutOnly multiple-scatter events. Events in which any detector is saturated have been removed. A random 10% of the data is shown. Legend as in Figure 8.40.

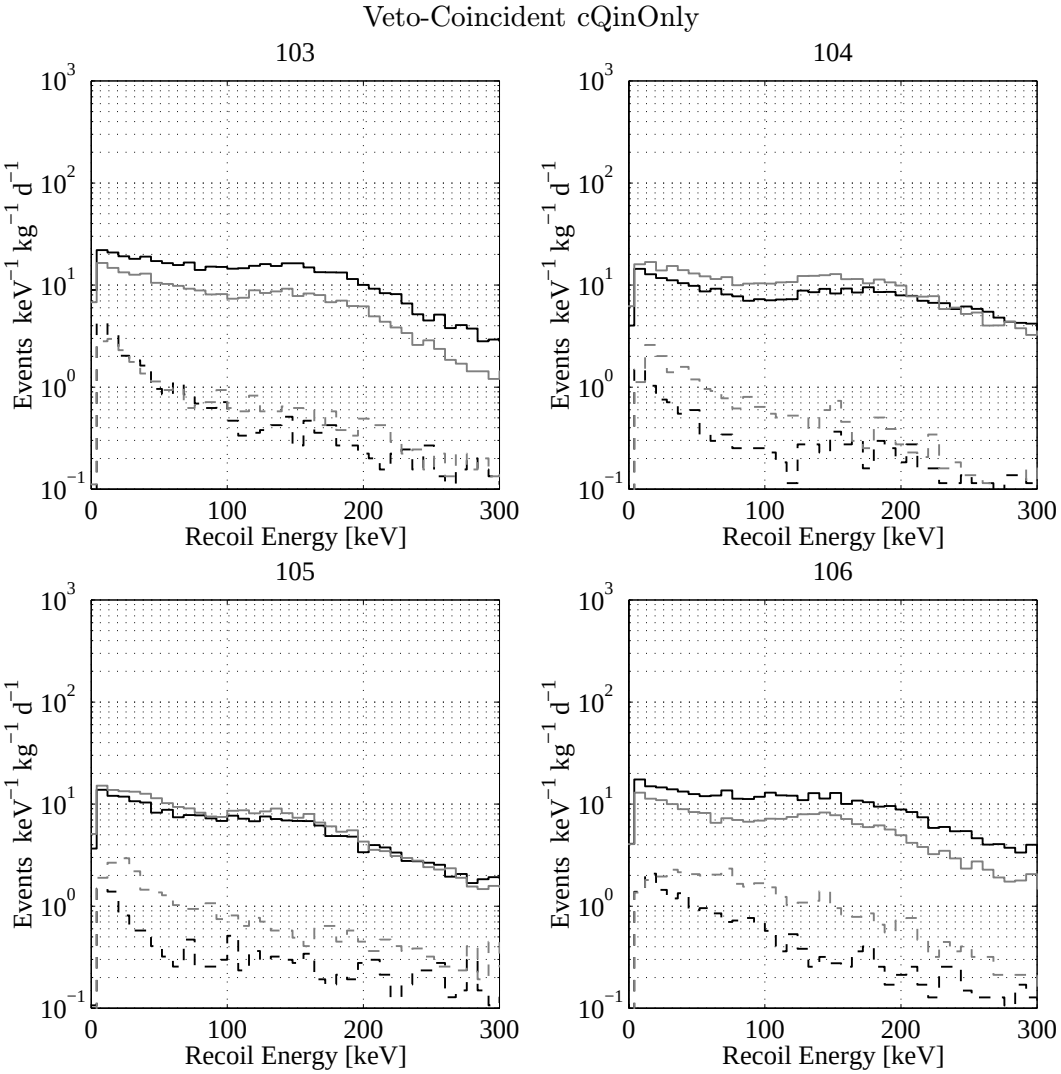

Figure 8.46: Recoil-energy spectra for veto-coincident cQinOnly events. Dark solid: single-scatter photons. Dark dashed: single-scatter electrons. Light solid: photons belonging to double scatters. Light dashed: electrons belonging to double scatters.

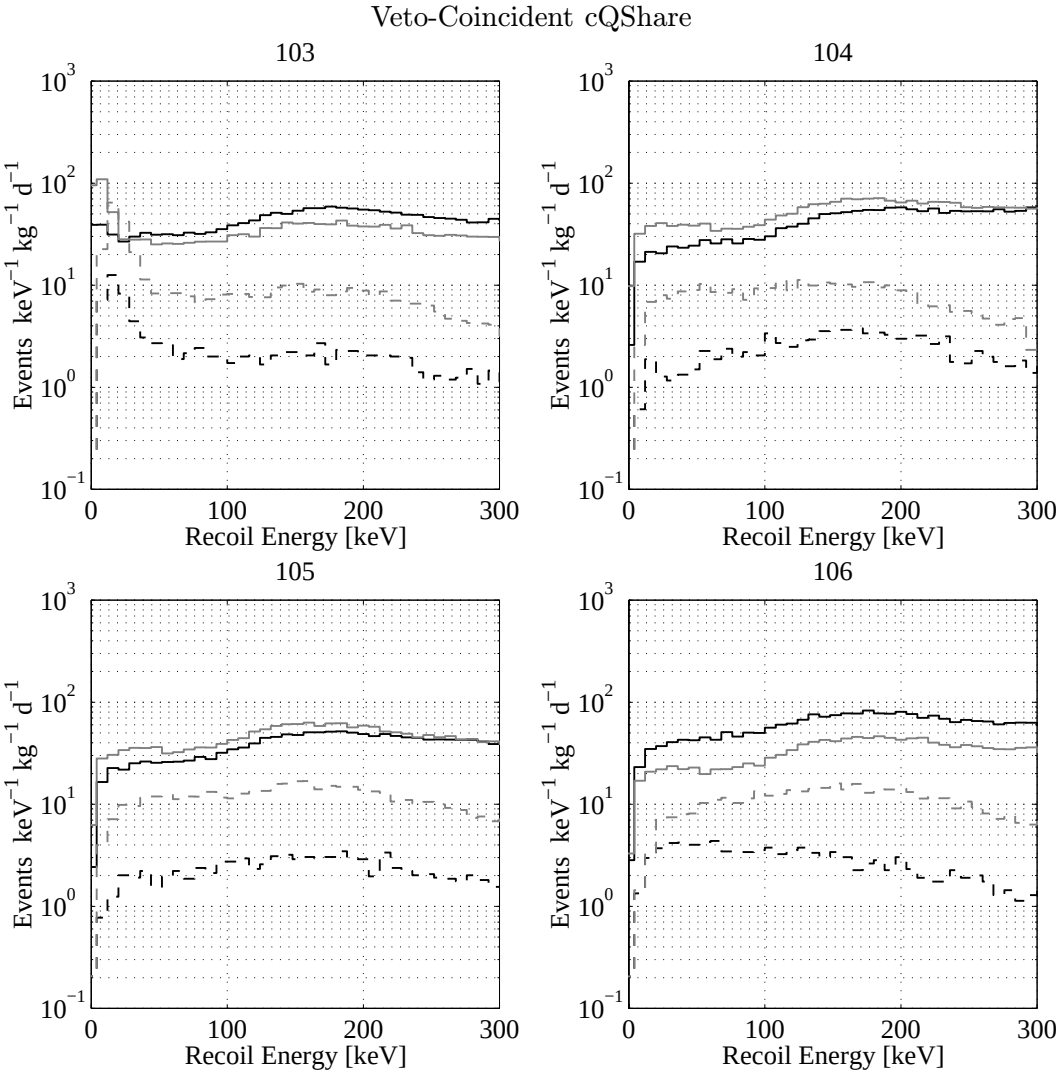

Figure 8.47: Recoil-energy spectra for veto-coincident cQShare events. Legend as in Figure 8.46.

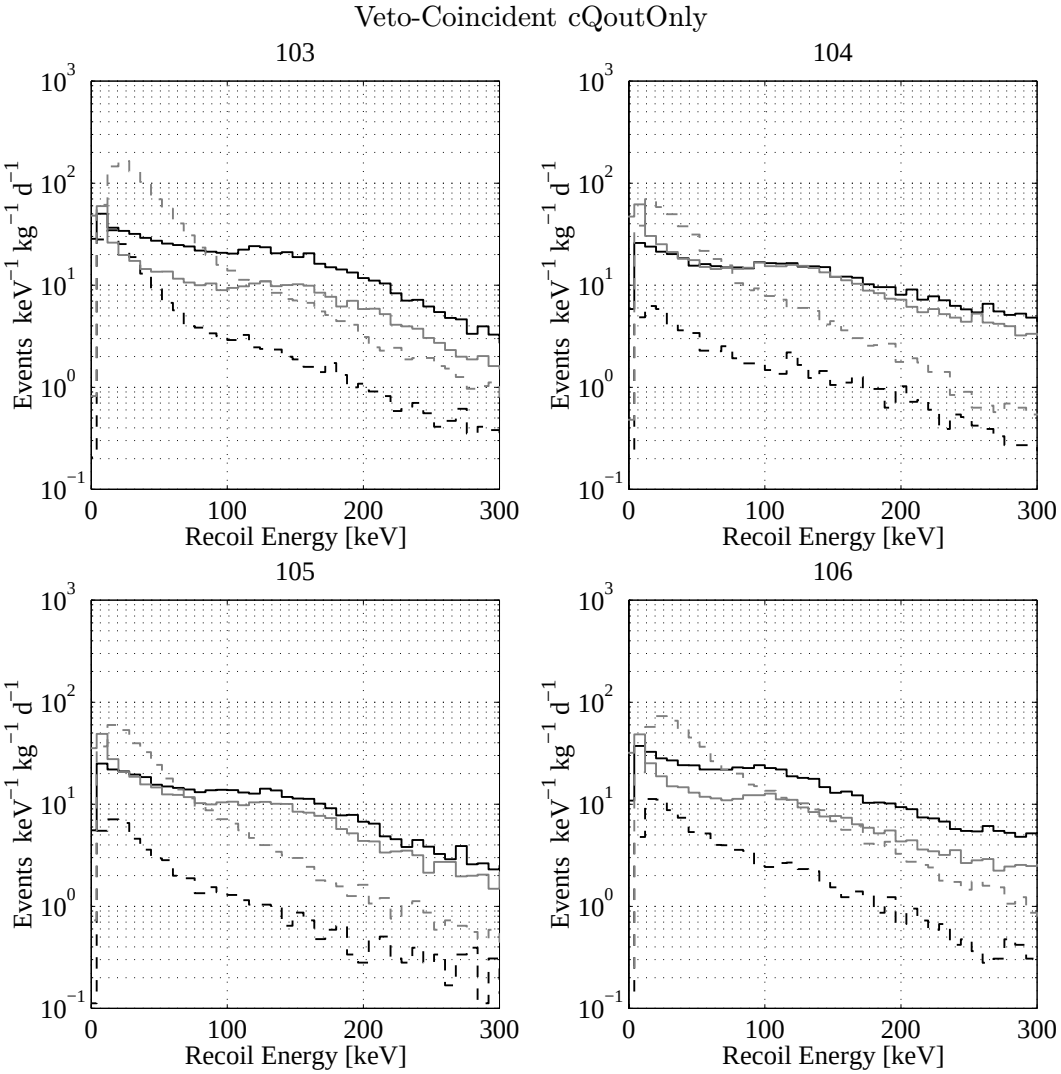

Figure 8.48: Recoil-energy spectra for veto-coincident cQoutOnly events. Legend as in Figure 8.46.

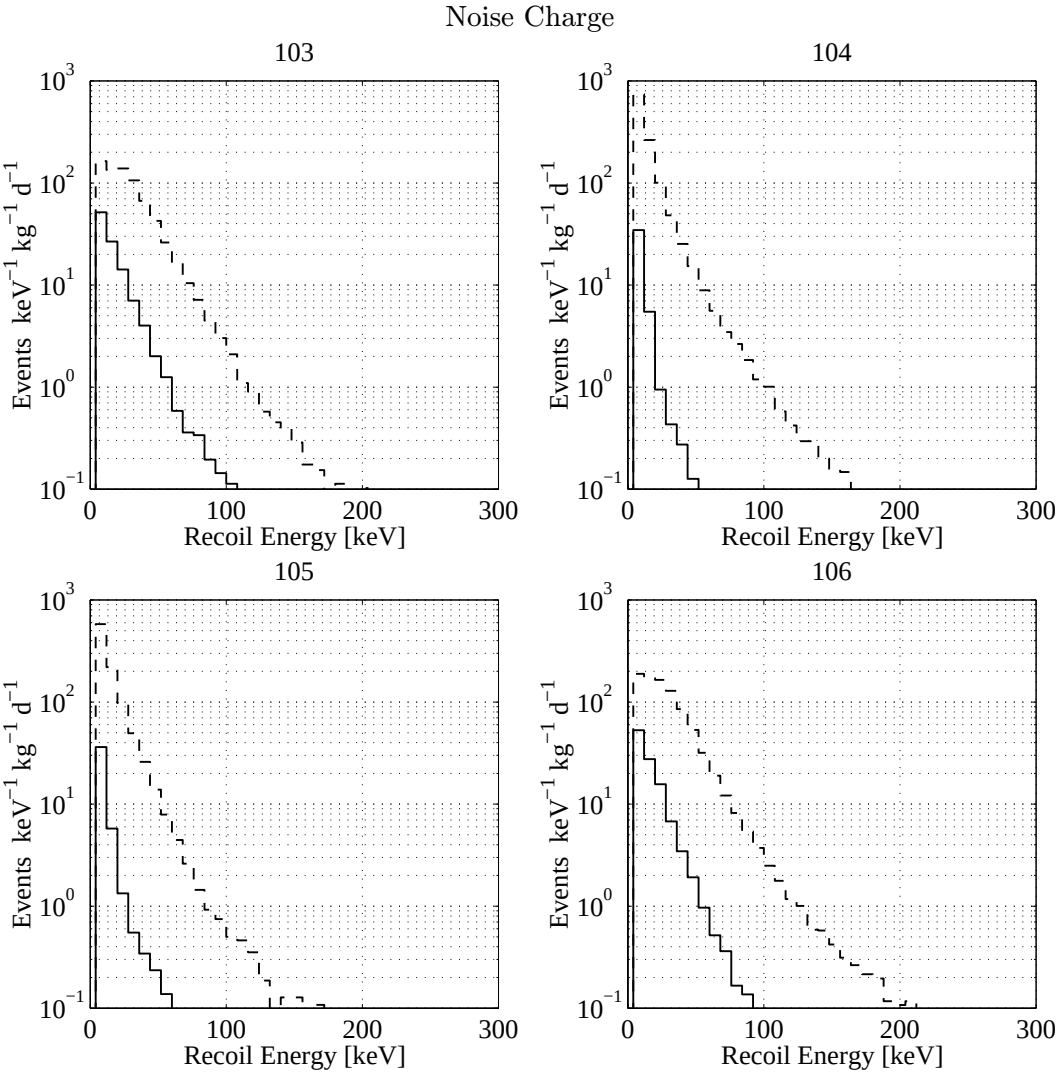

Figure 8.49: Recoil-energy spectra for noise-charge events. Solid: single scatters. Dashed: double scatters. Note that, because these events have only noise charge, no determination of coincidence or anticoincidence with veto can be made.

Chapter 9

Conclusion and Outlook

The data presented here yield the first significant exclusion limits on the WIMP-nucleon cross section from the Cryogenic Dark Matter Search. Interesting regions of minimal supersymmetric parameter space are being probed. The claimed annual-modulation detection by the DAMA collaboration is being tested, and, perhaps, disproved. CDMS is within sight of its SUF sensitivity goal, shown in Figure 9.1.

The CDMS philosophy of developing new technology has paid off handsomely. The 15-year investment made in understanding the solid-state physics of phonons and electron-hole pairs at low temperature has led to detectors that are making significant impact on one of the most pressing scientific questions of our time. Man-decades of engineering effort have gone into making this experiment, which many considered on the edge of cryogenic feasibility, run smoothly for extended periods.

The analysis presented here also points to open issues. Statistically precise electron and photon calibrations are needed to fully characterize the ionization yield distributions arising from interactions of these particles. More stringent tests and checks of the efficiency calculation are required. An explanation of the 20–30% systematic disagreement between the internal- and calibration-source-neutron data and simulations is necessary. And a more comprehensive understanding of the source of the residual electron background is important. To continue to make progress at the SUF site, and to reach the CDMS II goal at Soudan, all these systematic issues must be addressed.

But the future is bright. CDMS II will deploy 10 times as much detector mass. The α-Si/Al-Schottky ionization electrodes employed here and the QET athermal-phonon-sensing technology are being combined to give the CDMS II ZIP detectors an unprecedented ability to identify and reject bulk and surface electron recoils. Moreover, the ability of these detectors to diagnose background sources using position and ionization yield is only beginning to be exploited. With these tools, CDMS II will explore a significant portion of the parameter space available for minimal supersymmetric WIMPs, as shown in Figure 9.1. I personally look forward to seeing this dissertation quickly become obsolete!

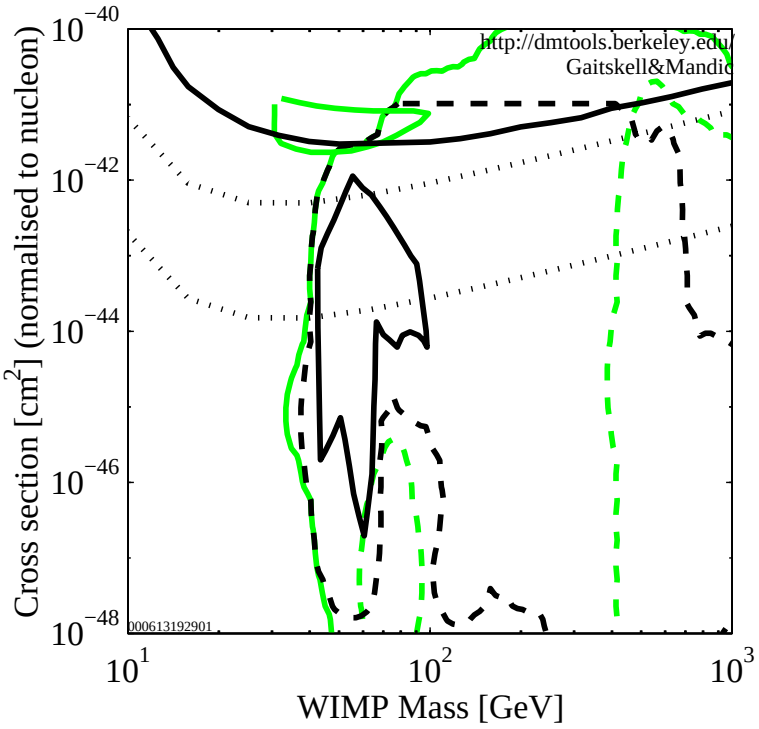

Figure 9.1: CDMS SUF and Soudan sensitivity goals. Also shown are the exclusion limit presented in this dissertation, the DAMA annual-modulation claim, and various MSSM models. Dark solid limit curve: exclusion limit from these data. Upper dotted limit curve: CDMS SUF sensitivity goal. Lower dotted limit curve: CDMS Soudan sensitivity goal. Light heart-shaped region: DAMA 2000 annual-modulation 3σ allowed region without application of DAMA 1996 limit. MSSM models: Light solid: Gondolo *et al.* mixed models [78]. Dark dashed: Gondolo *et al.* gaugino models. Light dashed: Gondolo *et al.* Higgsino models. Dark solid: Corsetti and Nath mSUGRA models [79]. Figure from [80].

Appendix A

Free Parameters and Frameworks in Minimal Supersymmetry

It is useful to discuss what parameters of supersymmetric models are presently constrained or unconstrained by experiment and to delineate the various schemes that exist for reducing the set of free parameters to a manageable number. The information presented here is derived from discussions between Paolo Gondolo and Rick Gaitskell, Vuk Mandic, and myself.

The parameters fixed by experiment are the following:

- **h**: The quark and lepton mass and mixing matrices, also known as the Yukawa couplings because they define a Yukawa-type coupling to the Higgs, whose vacuum expectation value turns these into mass terms. Fixed by observed Standard Model particle masses and mixing.
- g_1, g_2, g_3: The gauge couplings for the $U(1)$ hypercharge, $SU(2)$ electroweak, and $SU(3)$ strong interactions. They are fixed by the observed strengths of these interactions.
- $v_1^2 + v_2^2$: This quadrature sum of the vacuum expectation values of the two Higgs fields sets the mass of the W and Z weak bosons, and so is fixed by experiment.

The free parameters consist mainly of the coefficients of the "soft" supersymmetry-breaking terms in the Lagrangian. Obviously, supersymmetry must be broken to give the world we see, where no superpartners have yet been observed. However, the breaking must be done in such a way that the desirable traits of supersymmetry — in particular, cancellation of quadratic divergences above the SUSY-breaking scale — are preserved. This condition restricts the allowed breaking terms to a small set, first determined by Girardello and Grisaru [53]. Their coefficients are as follows:

- $\mathbf{M}^2_{\widetilde{f}}$: The sfermion masses. These are 3×3 hermitian matrices, one for each type of sfermion: $\widetilde{q_L}$, $\widetilde{u_R^c}$, $\widetilde{d_R^c}$, $\widetilde{l_L}$, $\widetilde{e_R^c}$, and $\widetilde{\nu_R^c}$ (assuming neutrinos are massive and thus that the normal particle ν_R exists!). Each matrix has 9 free parameters, giving a total of 54 free parameters. Note that these may be nondiagonal and therefore may themselves contribute to sfermion mixing, though this freedom has been accounted for in the above count.
- **A**: The sfermion-mixing matrices. These are three 3×3 matrices with no hermiticity requirement. There are therefore 18 free parameters per matrix, yielding a total of 54 free

parameters. There is a loose requirement that the minima of all the sfermions' scalar potentials (*i.e.*, the sum of all terms containing the powers of a given sfermion field alone, such as mass terms) be minimized at 0, but this constraint only puts some allowed range on the matrix components; it does not yield specific relations between components.

- M_1, M_2, M_3: The gaugino masses. Three free parameters.
- $m_1^2, m_2^2, B, \tan\beta \equiv v_2/v_1, \mu$: Higgs parameters (supersymmetry and electroweak-symmetry breaking). m_1^2 and m_2^2 are real scalars. B is a scalar and may be complex. $\tan\beta$ is limited by experiment from below and by m_t/m_b from above. Of these, three can be taken to be free and the others determined via the electroweak-symmetry breaking relations. Choices of parameter sets are
  - m_1^2, m_2^2, $\tan\beta$ free, B and μ dependent ([51], Equation 69).
  - B, μ, and $\tan\beta$ free, m_1^2 and m_2^2 dependent ([51], Equation 68).
  - $m_A^2 = -2B\mu/\sin 2\beta$, $\tan\beta$, μ free, B, m_1^2, and m_2^2 dependent. (m_A^2 is the mass of the pseudoscalar Higgs boson). A new non-independent quantity, m_A^2, has been introduced, so there are now three dependent parameters instead of two. The reason for this choice is that m_A^2 is a more physical parameter. This is what is used in most numerical codes. It should be noted that these conditions hold even for radiative electroweak-symmetry breaking — radiative EWSB imposes an additional condition.

Summing up the above, there are a total of 115 free parameters! There are different schemes for reducing the above freedom; these different schemes are called "frameworks," to distinguish them from a "model," which consists of a specific choice for all the parameters.

- **mSUGRA**: Short for *minimal supergravity* — this framework assumes that, above the Planck scale (10^{19} GeV), supersymmetry can be gauged: the invariance under global supersymmetry transformations can be promoted to invariance under local supersymmetry transformations. This yields general relativity, since supersymmetry transformations are related to Lorentz transformations. Some mechanism breaks this local invariance at energies below the Planck scale. However, since the GUT scale (10^{16} GeV) is so much lower, there is no reason to expect this mechanism to distinguish between the different gauge groups and particle generations. Thus, mSUGRA yields the maximal set of assumptions:
  - Gaugino-mass unification: $M_1 = M_2 = M_3 \equiv m_{1/2}$ at M_{GUT}.
  - Fermion Yukawa-coupling unification: $\mathbf{h}_E = \mathbf{h}_U = \mathbf{h}_D \equiv \mathbf{h}$ at M_{GUT}.
  - Sfermion-mass and -mixing unification: $\mathbf{M}_{\tilde{l}}^2 = \mathbf{M}_{\tilde{q}}^2 = \mathbf{M}_{\tilde{e}}^2 = \mathbf{M}^{\mathbf{2}}{}_{\tilde{\nu}}^2 = \mathbf{M}_{\tilde{u}}^2 = \mathbf{M}_{\tilde{d}}^2 = m_1^2\mathbf{1} = m_2^2\mathbf{1} \equiv m_0^2\mathbf{1}$ and $\mathbf{A}_e = \mathbf{A}_u = \mathbf{A}_d \equiv A_0\mathbf{1}$ at M_{GUT}.
  - $B = A_0 + m_0$ at M_{GUT}.
  - Some specification of the GUT gauge group. This is necessary to determine the running of the couplings from the Planck scale to the GUT scale.
  - Radiative electroweak-symmetry breaking: the electroweak-symmetry breaking parameters are generated by higher-order diagrams that can be summed into an effective Higgs potential. The parameters of the potential are thus set by other parameters and are not

free. This provides another relationship between μ^2 and the other Higgs sector parameters, thereby adding another condition and reducing the number of free parameters in the Higgs sector to two (usually taken to be m_A^2, the pseudoscalar mass, and $\tan\beta$).

This framework is the most constraining one available. It is not used, however, because some of its consequences violate observations. These will be noted below as more free frameworks are defined.

- **mSUGRA$_{\rm LHC}$/mSUGRA$_{\rm LEP}$**: This is what is meant when LEP data or LHC expectations are discussed in the context of "minimal supergravity." A few assumptions are relaxed:
  - Yukawa-coupling unification is not required. It is impossible to unify the first and second generation couplings at the GUT scale and still yield reasonable masses at observed energies. The first and second generation coupling unification is thus released immediately. mSUGRA$_{\rm LHC}$/mSUGRA$_{\rm LEP}$ also release the 3rd generation — this is also known as "no b/τ mass unification" for obvious reasons.
  - Radiative electroweak-symmetry breaking is employed, but, when picking models, no checks are made on the color/charge neutrality of the vacuum for the given set of parameters.
- **Generic SUGRA**: Releases assumptions further:
  - No Higgs-sfermion-mass unification. The two Higgs masses are assumed to be unified, $m_1^2 = m_2^2$. It is usually m_A^2 that is specified, which is related to these by EWSB. All the sfermion masses are assumed to be unified, $\mathbf{M}_{\tilde{l}}^2 = \mathbf{M}_{\tilde{q}}^2 = \mathbf{M}_{\tilde{e}}^2 = \mathbf{M}_{\tilde{\nu}}^2 = \mathbf{M}_{\tilde{u}}^2 = \mathbf{M}_{\tilde{d}}^2 \equiv m_0^2\mathbf{1}$. Lack of Higgs-sfermion-mass unification allows $m_{1,2}^2 \neq m_0^2$ at M_{GUT}.
  - No radiative electroweak-symmetry breaking — standard spontaneous electroweak-symmetry breaking is assumed, so there are three free Higgs parameters, as discussed earlier.

  The release of Higgs-sfermion-mass unification is done for three reasons. First, it allows one to separate the Higgs and sfermion sectors (the sleptons and squarks are linked by the sfermion-mass unification relation at the GUT scale), sending one or the other to very high mass to simplify analysis of models. Second, the gauge couplings do not unify properly at M_{GUT} if Higgs-sfermion-mass unification is enforced. Finally, superstring theories suggest there really is no connection between the Higgs and sfermion masses.
- **Constrained MSSM**: Makes some parameters "quasi-free:"
  - $\mathbf{A}$ and $\mathbf{m}_{\tilde{f}}$ matrices are diagonal (sfermion mixing same as fermion mixing) but otherwise unconstrained; *i.e.*, sfermion-mass unification is not imposed. Authors frequently forget to specify in which basis (mass-diagonal or weak-interaction-diagonal) this constraint is made, leading to much confusion. Frequently, they take only A_t and A_b nonzero since these dominate the sfermion-mass effects.
  - Higgs-sfermion-mass unification not required (as is necessary if sfermion-mass unification not required.).

- **MSSM**: No constraints, not even gaugino-mass unification.

A general comment that can be made about different frameworks is that, obviously, the more restrictive a framework is, the smaller will be the range of predictions for the WIMP-quark and WIMP-gluon scattering cross sections. As noted above, the most restrictive frameworks can be viewed as too restrictive, and possibly even in contradiction with experiment. However, an overly loose framework can make it artificially appear that a large fraction of SUSY WIMP models can never be detected. Where to draw the line is a judgment call.

Appendix B

Noise and Optimal Filtering

B.1 Introduction

The formalism of optimal filtering has traditionally been handed down via oral tradition, with each generation rediscovering truths lost in the shrouded mists of the past through careful study of the ancient writings of Bernard Sadoulet and various others, coupled with deep thought, usually under duress. This appendix is an effort to present some of the basic results of optimal filtering in a pedagogical manner, with care taken to get all the factors of 2 and π correct. I also make the correspondence between the continuous case and the discrete case.

Many of the manipulations performed below (non-convergent integrals, exchanging order of integrations for ill-defined integrals, treating delta functions like functions, etc.) are not rigorously justified. All calculations can always be carried over to the discrete case, where such problems vanish, to check that the results are correct.

B.2 Preliminaries: Fourier Transform and Power Spectral Density Conventions

B.2.1 Fourier Transforms — Continous Case

First, I define the conventions I use for Fourier transforms. The transform pair $g(t)$ and $\widetilde{g}(f)$ are defined by the forward and reverse transforms

$$\begin{aligned} \widetilde{g}(f) &= \int_{-\infty}^{\infty} dt\, g(t)\, e^{-j\omega t} \\ g(t) &= \int_{-\infty}^{\infty} df\, \widetilde{g}(f)\, e^{j\omega t} \end{aligned}$$

where $\omega \equiv 2\pi f$ is used for convenience. This sign convention for the argument of the exponential is used so one can make the correspondence $\frac{d}{dt} \leftrightarrow j\omega$, which is expected both from electrical circuit analysis and quantum mechanics. Some confusion may arise about the placement of the 2π; here, by choosing df as the measure for the frequency integral, the 2π is placed in the denominator of the right side of the reverse transform: $df = \frac{d\omega}{2\pi}$. This choice is made because it make sense in terms of units: if $g(t)$ carries some physical units (*e.g.*, volts), then $\widetilde{g}(f)$ carries units of 1/Hz times that

physical unit (*e.g.*, volts/Hz) because the measure for the integrand is df; using $d\omega$ would give $\tilde{g}$ units of volts/(radians/sec)). Some authors (*e.g.*, the *CRC Standard Mathematical Tables*) use a symmetric convention (where each transform gets a $\frac{1}{\sqrt{2\pi}}$); this would confound the units.

For the above reasons, it is important to remember that f is the argument of the transform $\tilde{g}$, not ω. There are many reasons to think otherwise. First, ω is what one thinks of as the conjugate variable to t because these two appear in the exponential. Second, using ω saves writing many 2π's, especially when solving linear differential equations by transforming to frequency space. Or, it may just appear that ω is the natural variable; for example, when $g(t) = e^{-t/\tau}$, then $\widetilde{g}(f) = \tau/(1 + 2\pi j f\tau)$, but one is tempted to write $\widetilde{g}(\omega) = \tau/(1 + j\omega\tau)$. Another reason is that, in actually calculating some of these integrals, one does a contour integral in the complex ω plane, not the f plane. Despite all these reasons to use ω, it is misleading to make ω the argument of the function $\tilde{g}$ because this implies that the correct measure for integration of $\tilde{g}$ is $d\omega$; however, one finds that one always has to use $d\omega/2\pi$, and thus the true measure for the function is df.

A number of useful properties make later manipulations and calculations easier. First, it is easy to show

$$\begin{aligned} g(-t) &\stackrel{\mathcal{FT}}{\longleftrightarrow} \widetilde{g}^*(f) \\ g(t) \text{ real} &\iff \widetilde{g}(-f) = \widetilde{g}^*(f) \end{aligned}$$

The δ-functions are defined by

$$\begin{aligned} \delta(f) &= \int_{-\infty}^{\infty} dt\, e^{-j\omega t} = 2\pi\delta(\omega) \\ \delta(t) &= \int_{-\infty}^{\infty} df\, e^{j\omega t} = \int_{-\infty}^{\infty} \frac{d\omega}{2\pi}\, e^{j\omega t} \end{aligned}$$

It is important to be careful with the factors of 2π.

A very useful result is the *convolution theorem*:

$$\widetilde{g}(f)\,\widetilde{h}(f) \stackrel{\mathcal{FT}}{\longleftrightarrow} [g * h](t) \equiv \int_{-\infty}^{\infty} dt_1\, g(t_1)\, h(t - t_1)$$

where $[g * h](t)$ signifies the convolution of g and h. Note the sign of the argument of h. Note also the strange units of the convolution: if g and h are in volts, the convolution carries volts^2 sec, or volts^2/Hz. This is rectified in the discrete case. The convolution of g and h at time t corresponds to the following sequence of operations: time reverse h; shift the time-reversed h later in time (to the right) by t; multiply the time-reversed, shifted h against g; and integrate this product. It is instructive to prove the convolution theorem explicitly because it illustrates one place where it is important to keep track of factors of 2π:

$$\begin{aligned} \mathcal{FT}\{[g * h](t)\} &= \int_{-\infty}^{\infty} dt\, e^{-j\omega t} \int_{-\infty}^{\infty} dt_1\, g(t_1)\, h(t - t_1) \\ &= \int_{-\infty}^{\infty} dt\, e^{-j\omega t} \int_{-\infty}^{\infty} dt_1 \int_{-\infty}^{\infty} df_1\, \widetilde{g}(f_1)\, e^{j\omega_1 t_1} \int_{-\infty}^{\infty} df_2\, \widetilde{h}(f_2)\, e^{j\omega_2 (t - t_1)} \\ &= \int_{-\infty}^{\infty} dt\, e^{-j\omega t} \int_{-\infty}^{\infty} df_1\, \widetilde{g}(f_1) \int_{-\infty}^{\infty} df_2\, \widetilde{h}(f_2)\, e^{j\omega_2 t} \int_{-\infty}^{\infty} dt_1\, e^{j(\omega_1 - \omega_2)t_1} \end{aligned}$$

where $g(t_1)$ and $h(t - t_1)$ are expanded in terms of their transforms and integration order is exchanged. It holds that $\int_{-\infty}^{\infty} dt_1\, e^{j(\omega_1-\omega_2)t_1} = 2\pi\delta(\omega_1 - \omega_2) = \delta(f_1 - f_2)$. Therefore,

$$\begin{aligned}\mathcal{FT}\{[g * h](t)\} &= \int_{-\infty}^{\infty} dt\, e^{-j\omega t} \int_{-\infty}^{\infty} df_1\, \widetilde{g}(f_1) \int_{-\infty}^{\infty} df_2\, \widetilde{h}(f_2)\, e^{j\omega_2 t}\delta(f_1 - f_2) \\ &= \int_{-\infty}^{\infty} dt\, e^{-j\omega t} \int_{-\infty}^{\infty} df_1\, \widetilde{g}(f_1)\, \widetilde{h}(f_1)\, e^{j\omega_1 t}\end{aligned}$$

Interchanging the order of integration and using $\int_{-\infty}^{\infty} dt\, e^{j(\omega_1-\omega)t} = \delta(f_1 - f)$ again yields

$$\mathcal{FT}\{[g * h](t)\} = \widetilde{g}(f)\, \widetilde{h}(f)$$

The convolution has the property (provable by changing variables from t_1 to $t_2 = t - t_1$)

$$[g * h](t) = [h * g](t)$$

The convolution theorem is very useful for linear systems, such as filters: typically, the effect of a linear system is easily represented in Fourier space as multiplication by a transfer function. The convolution theorem thus relates the filter's transfer function in Fourier space to its effect in time. By using a δ-function for g, the convolution theorem implies that the transform of the transfer function is just the impulse response of the system, its response to a δ-function input. One can see intuitively that, for a causal filter (one whose impulse response is zero for negative time), the convolution of the signal g with the impulse response h calculates the output at time t by summing up the impulse responses of the filter to the input over earlier times.

A corollary of the convolution theorem is

$$\begin{aligned}\widetilde{g}^*(f)\, \widetilde{h}(f) \stackrel{\mathcal{FT}}{\longleftrightarrow} [g \otimes h](t) &\equiv \int_{-\infty}^{\infty} dt_1\, g(t_1)\, h(t + t_1) \\ &= \int_{-\infty}^{\infty} dt_1\, g(t_1 - t)\, h(t_1)\end{aligned}$$

where $[g \otimes h](t)$ is the cross-correlation of g and h (again, with weird units). I call this the *cross-correlation theorem*, though this is not general practice. The cross-correlation is easy to visualize: in the first form, shift h earlier in time by t, multiply against g, and integrate the product; in the second form, shift g later in time by t, multiply against h, and integrate the product. This corollary is easy to prove by applying the convolution theorem to the time reverse of g; this gives $\tilde{g}^*$ on the left side and yields the sign change needed to produce the right side of the second line. Note that the cross-correlation, unlike the convolution, is not commutative:

$$[h \otimes g](t) = [g \otimes h](-t)$$

The cross-correlation and convolution can be related by a time reversal:

$$\begin{aligned}[g \otimes h](t) &\stackrel{\mathcal{FT}}{\longleftrightarrow} \widetilde{g}^*(f)\widetilde{h}(f) \\ &= \widetilde{g_r}(f)\widetilde{h}(f) \\ &\stackrel{\mathcal{FT}}{\longleftrightarrow} [g_r * h](t)\end{aligned}$$

where $g_r(t) = g(-t)$ is the time-reverse of g.

Finally, a corollary of the corollary is *Parseval's Theorem*:

$$\int_{-\infty}^{\infty} dt\,[g(t)]^2 \;=\; \int_{-\infty}^{\infty} df\,|\widetilde{g}(f)|^2$$

This can be proven by integrating both sides of the cross-correlation theorem over f:

$$\begin{aligned}
\int_{-\infty}^{\infty} df\,\widetilde{g}^*(f)\,\widetilde{g}(f) &= \int_{-\infty}^{\infty} df \int_{-\infty}^{\infty} dt\,e^{-j\omega t} \int_{-\infty}^{\infty} dt_1 g(t_1)\,g_1(t+t_1) \\
&= \int_{-\infty}^{\infty} dt_1\,g(t_1) \int_{-\infty}^{\infty} dt\,g(t+t_1) \int_{-\infty}^{\infty} df\,e^{-j\omega t} \\
&= \int_{-\infty}^{\infty} dt_1\,g(t_1) \int_{-\infty}^{\infty} dt\,g(t+t_1)\,\delta(t) \\
&= \int_{-\infty}^{\infty} dt_1\,g(t_1)\,g(t_1)
\end{aligned}$$

B.2.2 Fourier Transforms — Discrete Case

In reality, signals are not sampled continuously and over infinite time (though sometimes this experiment makes it look like the second statement may hold...). First consider the reduction of the sampling period to a finite time, $2T$. The proper transform pair is now

$$\begin{aligned}
\tilde{g}_n &= \frac{1}{2T}\int_{-T}^{T} dt\,g(t)\,e^{-j\omega_n t} \\
g(t) &= \sum_{n=-\infty}^{\infty} \tilde{g}_n\,e^{j\omega_n t}
\end{aligned}$$

with $f_n = \frac{n}{2T}$ and $\omega_n = 2\pi f_n$ as usual. One can justify this retroactively by identifying $df = \frac{1}{2T}$ and realizing the continuous case corresponds to $T \to \infty$; here, df has been shifted into the forward transform. Another way to check the transforms is to just substitute the first formula into the second one and see that one has an identity (recognizing $\sum_{n=-\infty}^{\infty} e^{j\omega_n t} = 2T\delta(t)$: one needs the $2T$ to yield the correct time units). The use of frequency bins of a given width can be interpreted as a version of the uncertainty principle — by sampling for only a finite time, only a nonzero frequency resolution can be achieved, given by the bin width. Note that the bandwidth is still infinite — there is information up to arbitrarily high frequency because the time signal is continuous. The units are different also: $\tilde{g}_n$ now carries the same units as $g(t)$. One can make the correspondence

$$\widetilde{g}(f) \longleftrightarrow \frac{\tilde{g}_n}{df} = 2T\,\tilde{g}_n$$

This will be useful below to convert from the continuous case to the discrete case.

This intermediate step to discreteness is not useful, so I immediately move to the fully discrete case. The transform pair is defined by

$$\tilde{g}_n \;=\; \frac{1}{N}\sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} g_k\,e^{-j\omega_n t_k}$$

$$g_k = \sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \tilde{g}_n \, e^{j\omega_n t_k}$$

where $g_k \equiv g(t_k) \equiv g(k\Delta t)$ and $2T = N\Delta t$. That is, dt has been replaced by the sampling interval Δt and has been cancelled everywhere. The frequency binning is the same as for the intermediate case; that is, it is only the limited sampling period that gives a finite frequency resolution. The feature introduced by sampling at finite frequency is reduction of the bandwidth. This is the Nyquist theorem: by sampling at a finite frequency, all information above half the sampling frequency is lost. The information is not actually lost; rather, it is aliased into the frequency range $\frac{N}{2}\frac{1}{2T}, \cdots, (\frac{N}{2}-1)\frac{1}{2T}$. (This is easy to show by trying to transform a signal $g_k = e^{j\omega t_k}$ where $f > \frac{N}{2}\frac{1}{2T} = \frac{1}{2\Delta t}$.) The correspondence

$$\tilde{g}(f) \longleftrightarrow \frac{\tilde{g}_n}{df} = 2T\,\tilde{g}_n$$

continues to hold. The units of g_k are the same as $g(t)$.

The δ-functions take the form

$$\delta(f_n) = \frac{1}{N} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} e^{-j\omega_n t_k}$$

$$\delta(t_n) = \frac{1}{N} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} e^{j\omega_k t_n}$$

Note that these functions now take on the values 0 or 1, rather than 0 or ∞. The factors of 2π do not appear in the discrete case.

In the discrete case, the convolution and cross-correlation are defined with more sensible units:

$$[g * h]_n = \frac{1}{N} \sum_{m=-\frac{N}{2}}^{\frac{N}{2}-1} g_m \, h_{n-m}$$

$$[g \otimes h]_n = \frac{1}{N} \sum_{m=-\frac{N}{2}}^{\frac{N}{2}-1} g_m \, h_{n+m}$$

That is, the cross-correlation and convolution carry units of volts^2. The $\frac{1}{N}$ factor is applied for two reasons: first, correspondence with the intermediate case gives $\frac{1}{N} \leftrightarrow \frac{\Delta t}{2T}$: this would give a sensible definition for the convolution and cross-correlation in the intermediate case; second, with this normalization, both the convolution and cross-correlation at zero shift give $\langle [g(t)]^2 \rangle$ for $h = g$, which is certainly more reasonable than giving $N\langle [g(t)]^2 \rangle$.

In the discrete case, there is the minor issue of the edges: since the traces do not extend in each direction to ∞, when one shifts a trace during convolution or cross-correlation, samples are lost off the edge because there is nothing against which to multiply them. One might try

correcting the normalization of $[g * h]_n$ by dividing by $N - |m|$ rather than N. How well this works depends on the traces themselves: if the traces go to zero at the ends of the window, then the "uncorrected" formula is in fact the best estimator of the true value; however, if the traces tend to some nonzero value at the edges of the window, then the "uncorrected" formula is quite a bad estimator. However, in practice, the "corrected" form does not fix this. One might hope that the use of a Hanning window would fix these edge effects; it also does not. The fundamental problem is that, by using traces that end before the pulse, one artificially changes the pulse shape, and thus its Fourier transform, in an unrecoverable way. The information past the edge of the digitization window is lost. In practice, for the pulse-analysis case where the trace and filter both are zero for many samples at the start of the window but are nonzero at the end, the best solution is to slightly shorten the data trace so that, in effect, the filter trace has more samples at its end. In the pretrigger direction, there is no edge effect because both traces are very close to zero. In the posttrigger direction, the longer filter trace eliminates the edge effect as long as the shift value is not larger than the number of extra samples in the filter trace; this typically holds.

With the above definitions, the convolution and cross-correlation theorems hold as before:

$$\begin{aligned} \tilde{g}_n \, \tilde{h}_n &\stackrel{\mathcal{FT}}{\longleftrightarrow} [g * h](t) \\ \tilde{g}_n^* \, \tilde{h}_n &\stackrel{\mathcal{FT}}{\longleftrightarrow} [g \otimes h](t) \end{aligned}$$

Parseval's theorem is

$$\frac{1}{N} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} g_k^2 = \sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} |\tilde{g}_n|^2$$

Note the position of the $\frac{1}{N}$.

B.3 Noise

The formalism for noise is confusing at first because the units that noise is usually quoted in, $\mathrm{V}/\sqrt{\mathrm{Hz}}$, do not correspond directly to any physical quantity. One becomes so accustomed to the convenience of having noises in $\mathrm{V}/\sqrt{\mathrm{Hz}}$ rather than V^2/Hz that one tends to forget the units are nonsensical. In this section, I explain the formalism for noise and why it carries units of $\mathrm{V}/\sqrt{\mathrm{Hz}}$.

The fluctuations of a physical voltage are assumed to be a Gaussian-distributed random variable; that is, $v(t)$ is chosen from a Gaussian distribution whose variance is $\langle [v(t)]^2 \rangle$ ($\langle v(t) \rangle = 0$ is assumed for simplicity). However, $\langle [v(t)]^2 \rangle$ does not fully determine the noise characteristics of the system — it does not contain information about correlations of $v(t)$ with itself in time. Physically, it is important to consider such correlations — for example, for a temperature fluctuation at time t, it takes a time τ for the extra energy associated with the fluctuation to flow out of the system, so $T(t)$ contains some information about $T(t+\tau)$. This information is contained in the autocorrelation function $R(\tau)$:

$$\begin{aligned} R(\tau) &= \langle v(t) v(t+\tau) \rangle \\ &= \lim_{T \to \infty} \frac{1}{2T} [v \otimes v](\tau) \\ &= \lim_{T \to \infty} \frac{1}{2T} \int_{-T}^{T} dt \, v(t) \, v(\tau + t) \end{aligned}$$

Note that this differs from the definition of the cross-correlation for the continuous case by a factor $2T$; the $2T$ is inserted here to give the autocorrelation function reasonable units. Also, since $v(t)$ extends infinitely into the past and future, the integral is infinite; the $2T$ makes it finite.

Such correlations are also coded in the frequency spectrum of the noise. The Fourier transform of the autocorrelation function is the noise power spectral density, $J(f)$:

$$J(f) = \lim_{T\to\infty} \int_{-T}^{T} dt\, R(t)\, e^{-j\omega t}$$

which has units of V^2/Hz. One typical case is $R(t) = \sigma^2 e^{-|t|/\tau}$, which holds for linear systems with a single decay time. The resulting power spectral density is $J(f) = 2\tau\sigma^2[1 + (\omega\tau)^2]^{-1}$. The noise PSD rolls off at frequencies above $(2\pi\tau)^{-1}$, expressing the strong autocorrelation of the noise for time separations smaller than τ. Note that all noises must be rolled off at some frequency; otherwise, $J(f)$ does not tend to zero as $f \to \infty$ and $R(t)$ becomes infinite. Perfect "white noise," with $J(f)$ a constant, does not exist. However, any noise of the first form can be considered white for $f \ll (2\pi\tau)^{-1}$; this expresses the fact that, for $t \gg \tau$, the autocorrelation function tends to zero: the noise appears uncorrelated in time for $t \gg \tau$.

A corollary of the definition of the power spectral density is

$$R(t) = \int_{-\infty}^{\infty} df\, J(f)\, e^{j\omega t}$$

and, specifically,

$$\langle [v(t)]^2 \rangle = R(0) = \int_{-\infty}^{\infty} df\, J(f)$$

This expression ultimately justifies the use of $J(f)$ to represent the frequency behavior of the noise: the integral of the noise power spectral density gives the noise variance in time. The units of $J(f)$, V^2/Hz, are thus sensible. The quantity usually discussed or given in noise spectra is $\sqrt{J(f)}$, which thus has units of $\mathrm{V}/\sqrt{\mathrm{Hz}}$. In many places, this second quantity is incorrectly referred to as the power spectral density (including in earlier sections of this dissertation). Usually, the context or units indicate whether it is $J(f)$ or $\sqrt{J(f)}$ that is implied by "power spectral density." The above relation can also be used as a check that calculational routines have been coded correctly.

The advantage of moving to frequency space for noise is that, for linear systems, noise components at different frequencies are independent. This results simply from the fact that, for linear systems, the behavior of the system at frequency f is independent of its behavior at other frequencies. Consider, for example, a thermal system with heat capacity C coupled to a thermal bath by a coupling G and at temperature $T = \delta T(t) + T_0$. In time, energy conservation gives

$$C\,\frac{d\delta T}{dt} = -GP(t)$$

$P(t)$ is just the noise power exchanged between the system and the thermal bath; though the system and bath are at the same DC temperature and there is no net power flow, there is power flowing back and forth to maintain the equilibrium condition. There are correlations in time according to the above equation. However, its transform is

$$j\omega\, C\widetilde{\delta T}(f) = -G\widetilde{P}(f)$$

where C and G are taken to be independent of frequency. $P(t)$ is the driving force and can be specified in either time or frequency space. The solution for $\widetilde{\delta T}(f)$ is

$$\widetilde{\delta T}(f) = \frac{1}{G}\frac{\widetilde{P}(f)}{1+j\omega\tau}$$

where $\tau = C/G$ is the thermal time constant. Thus, the behavior of the temperature fluctuations at frequency f depends only on $\widetilde{P}(f)$ and not on $\widetilde{\delta T}(f)$ at other frequencies. Typically, $P(f)$ (or its analogue) can be traced back to a physical process for which there should be no correlations among frequencies, at least for timescales long compared to the timescale of the physical processes involved, which are usually quite short, 10^{-12} s or less. This lack of correlation is very much unlike the behavior of $\delta T(t)$, where there are strong correlations between $\delta T(t)$ and $\delta T(t+\delta t)$ for $|\delta t| \lesssim \tau$. Thus, frequency space is a basis in which the noise is diagonal, which is in general very convenient. It is also important for calculating a properly normalized χ^2 when fitting a pulse-shape model to data.

Although the noise power spectral density can be determined from the autocorrelation function calculated from many noise traces (traces that have no events in them), it is common practice to calculate it directly from the Fourier transform of the noise traces. This can be seen as follows:

$$\begin{aligned}
J(f) &= \lim_{T\to\infty}\int_{-T}^{T} dt\, R(t)\, e^{-j\omega t} \\
&= \lim_{T\to\infty}\int_{-T}^{T} dt\, e^{-j\omega t} \lim_{T\to\infty}\frac{1}{2T}[v\otimes v](t) \\
&= \lim_{T\to\infty}\int_{-T}^{T} dt\, e^{-j\omega t} \lim_{T\to\infty}\frac{1}{2T}\int_{-\infty}^{\infty} df_1\, e^{j\omega_1 t}\, \widetilde{v}^*(f_1)\, \widetilde{v}(f_1) \\
&= \lim_{T\to\infty}\frac{1}{2T}\int_{-\infty}^{\infty} df_1\, e^{j\omega_1 t}\, |\widetilde{v}(f_1)|^2 \delta(f-f_1) \\
&= \lim_{T\to\infty}\frac{1}{2T}|\widetilde{v}(f)|^2
\end{aligned}$$

where limit and integration orders have been freely exchanged and the definition of $\delta(f)$ has been used. That is, the power spectral density is just the squared amplitude of the Fourier transform divided by the sampling period.

It is necessary to convert the above relations to the discrete case for application. The autocorrelation function is

$$\begin{aligned}
R(t_n) = R_n &= [v\otimes v]_n \\
&= \frac{1}{N}\sum_{m=-\frac{N}{2}}^{\frac{N}{2}-1} v_m\, v_{n+m}
\end{aligned}$$

where the $2T$ denominator is no longer needed because the cross-correlation includes it for the discrete case. The $\frac{1}{N}$ factor plays the normalizing role of $\frac{1}{2T}$. The power spectral density can be written in two forms:

$$J_p = J(f_p) = \mathcal{FT}[R(t_k)]$$

$$= \frac{1}{N} \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} e^{-j\omega_k t_n} \sum_{m=-\frac{N}{2}}^{\frac{N}{2}-1} v_m \, v_{n+m}$$

corresponds to the definition of $J(f)$ in terms of the autocorrelation function. The more useful form is

$$\begin{aligned} J(f_p) &= \lim_{T\to\infty} \frac{1}{2T} |\widetilde{v}(f_p)|^2 \\ &= \lim_{T\to\infty} \frac{1}{2T} |2T\tilde{v}_p|^2 \\ &= 2T|\tilde{v}_p|^2 \end{aligned}$$

where the limit has been discarded in the discrete case. This is the form that is usually used for calculating $J(f)$: a number of noise traces are acquired, $|\tilde{v}_p|^2$ calculated for each trace, and the mean value of $|\tilde{v}_p|^2$ calculated, multiplying by $2T$. This final expression is what one would expect: the $\tilde{v}_p$ are in units of volts and $\frac{1}{2T}$ is the frequency bin width, so the expression gives V^2/Hz.

It is also important to understand the relation between an individual noise realization $v(t)$ (or $\widetilde{v}(f)$) and $J(f)$. $J_p/(2T)$ gives the mean square absolute value of the frequency component $\tilde{v}_p$. Under the assumption of Gaussianity, $J_p/(2T)$ thus gives the variance of the Gaussian distribution from which $\tilde{v}_p$ is drawn. Since $\tilde{v}_p$ is complex, $J_p/(2T)$ only specifies the distribution for $|\tilde{v}_p|$. However, again because of linearity, the real and imaginary components of $\tilde{v}_p$ are independent. One can choose one of two prescriptions. One is to pick $\Re[\tilde{v}_p]$ and $\Im[\tilde{v}_p]$ independently from a Gaussian distribution of variance $(J_p/2)/(2T)$; then $\langle|\tilde{v}_p|^2\rangle = J_p/(2T)$. The other is to pick $|\tilde{v}_p|$ from a Gaussian of variance $J_p/(2T)$, rectifying negative values, and to pick the phase of $\tilde{v}_p$ from a uniform distribution over the interval $[0, 2\pi)$. Finally, for most applications, $v(t)$ is real, and thus $\widetilde{v}(-f_p) = \widetilde{v}^*(f_p)$: the negative frequency components are not independent of the positive ones.

B.4 Optimal Pulse-Height and Time-Offset Estimators

With the above formalism, the problem of finding optimal pulse-height and time-offset estimators for a pulse can be formulated rigorously. The procedure described here is termed "optimal filtering" for reasons that will become apparent. A real pulse has the form

$$v(t) = As(t) + n(t)$$

where $s(t)$ is the expected pulse shape (normalized to have peak height of 1), A is the amplitude to be estimated, and $n(t)$ is a particular noise realization. Let the noise power spectral density be given by $J(f)$. It is desired to find the best estimator for A. To pose the question clearly, one needs a criterion for "best." This is defined by the χ^2 of the fit of the event to the expected pulse shape:

$$\chi^2 = \int_{-\infty}^{\infty} df \, \frac{|\widetilde{v}(f) - A\widetilde{s}(f)|^2}{J(f)}$$

As justified above, the χ^2 is calculated in frequency space because different frequency components are independent. The above χ^2 is properly normalized because

$$df \, \frac{\langle|\widetilde{v}(f) - A\widetilde{s}(f)|^2\rangle}{J(f)} = df \, \frac{\langle|\widetilde{n}(f)|^2\rangle}{J(f)}$$

$$\begin{aligned} &= df \lim_{T\to\infty} 2T \\ &= 1 \end{aligned}$$

That is, $\langle\chi^2\rangle$ equals the sum of the number of frequency components, which is the number of degrees of freedom and is formally infinite. It becomes finite in the discrete case. Note that though the negative frequency components are not independent of the positive ones, the positive components each have two degrees of freedom, amplitude and phase, or real and imaginary components, so the χ^2 is indeed properly normalized. It is straightforward to find the best estimator $\widehat{A}$ by minimizing χ^2 with respect to A:

$$\widehat{A} = \frac{\int_{-\infty}^{\infty} df \frac{\widetilde{s}^*(f)\widetilde{v}(f)}{J(f)}}{\int_{-\infty}^{\infty} df \frac{|\widetilde{s}(f)|^2}{J(f)}}$$

One can show that this estimator is unbiased; *i.e.*, $\langle\widehat{A}\rangle = A$. It is interesting to calculate the expected variance of A, $\sigma_A^2 = \langle[\widehat{A} - A]^2\rangle$, which gives the expected resolution. It is an interesting excercise in Fourier algebra to do this directly by calculating $\langle[\widehat{A}]^2\rangle$ and using $\langle[\widehat{A} - A]^2\rangle = \langle[\widehat{A}]^2\rangle - A^2$; one has to take special care with the correlation between $n(f)$ and $n(-f)$. However, in the interest of saving space, I use the general result $\sigma_A^2 = \left[\frac{1}{2}\frac{\partial^2\chi^2}{\partial A^2}\right]^{-1}$, which yields

$$\sigma_A^2 = \left[\int_{-\infty}^{\infty} df \, \frac{|s(f)|^2}{J(f)}\right]^{-1} \tag{B.1}$$

This last formula is extremely useful. It indicates the best possible resolution achievable given the assumed pulse shape and noise power spectral density, so it can be used to gauge whether the fitting procedure is meeting its expected performance. In theoretical comparisons of different types of particle-detection schemes, the above formula can be used to estimate the optimal performance of each scheme.

The above formalism can be extended to fit for an arbitrary time offset. As described in Chapter 6, such a time-offset search is necessary in cases where the pulse is not in the expected position in the digitizer window. A time-offset search is performed to find an ionization pulse in the digitized trace for phonon-trigger events. More generally, it is used to correct for variation with pulse amplitude of the position of the pulse relative to the trigger time, though the search range is much shorter in the latter case. The χ^2 is defined as before, but this time allowing for a time offset:

$$\chi^2 = \int_{-\infty}^{\infty} df \, \frac{|\widetilde{v}(f) - Ae^{-j\omega t_0}\widetilde{s}(f)|^2}{J(f)}$$

If the template pulse normally starts at $t = 0$, the shifted pulse $e^{-j\omega t_0}\widetilde{s}(f)$ starts at time $t = t_0$. Minimizing χ^2 with respect to t_0 does not yield a simple formula for $\widehat{t_0}$ because of the more complex dependence of χ^2 on t_0. Minimizing χ^2 yields

$$\begin{aligned} -\int_{-\infty}^{\infty} df \, \frac{\widetilde{v}^*(f)Ae^{-j\omega t_0}\widetilde{s}(f) + \widetilde{v}(f)Ae^{j\omega t_0}\widetilde{s}^*(f)}{J(f)}\bigg|_{t_0=\widehat{t_0}} &= 0 \\ -2\int_{-\infty}^{\infty} df \, \frac{\widetilde{v}^*(f)Ae^{-j\omega t_0}\widetilde{s}(f)}{J(f)}\bigg|_{t_0=\widehat{t_0}} &= \end{aligned}$$

where $\widehat{t_0}$ is the value of t_0 that yields the minimization condition. The two quadratic terms vanish because they are independent of t_0. The second line is derived by using the relation

$$\int_{-\infty}^{\infty} df\, \frac{\widetilde{v}(f) A e^{j\omega t_0} \widetilde{s}^*(f)}{J(f)} = \int_{-\infty}^{\infty} df\, \frac{\widetilde{v}^*(f) A e^{-j\omega t_0} \widetilde{s}(f)}{J(f)}$$

which can be demonstrated using the convolution theorem and the fact that $\mathcal{FT}[\widetilde{v}(f)]$ and $\mathcal{FT}[e^{-j\omega t_0}\widetilde{s}(f)/J(f)]$ are real. The minimization condition can be written in a more illustrative form. Define the *optimal filter* by

$$\widetilde{\phi}(f) = \frac{\widetilde{s}^*(f)}{J(f)}$$

with Fourier transform $\phi(t)$. $\widetilde{\phi}(f)$ does not have the correct units, but this is rectified later. The minimization condition is then (dropping the -2 and the constant A):

$$\frac{\partial}{\partial t_0}[\phi * v](t_0)\Big|_{t_0=\widehat{t_0}} = 0$$

The prescription is thus to apply the filter $\phi(t)$ to the trace and find the time of the peak. The best estimator for the amplitude can be determined as before, treating $e^{-j\omega\widehat{t_0}}\widetilde{s}(f)$ as $\widetilde{s}(f)$ was treated before; effectively, $e^{-j\omega\widehat{t_0}}\widetilde{s}(f)$ gives the correct expected pulse shape because it incorporates the time offset. Thus the best estimator for the amplitude is given by

$$\widehat{A} = \frac{\int_{-\infty}^{\infty} df \frac{e^{j\omega\widehat{t_0}}\widetilde{s}^*(f)\,\widetilde{v}(f)}{J(f)}}{\int_{-\infty}^{\infty} df \frac{|\widetilde{s}(f)|^2}{J(f)}} \tag{B.2}$$

The numerator is again the convolution of the trace $v(t)$ with the optimal filter $\phi(t)$, evaluated at the time offset $\widehat{t_0}$. Since $\widehat{t_0}$ corresponds to the time of the peak of the filtered trace, $\widehat{A}$ is just proportional to the peak amplitude of the filtered trace. Hence, the procedure is called "optimal filtering." The denominator is the peak amplitude of the trace given by optimally filtering the expected pulse shape $s(t)$ (with peak amplitude 1); it is a normalization factor that corrects for the fact that the optimal filter changes the peak amplitude. It also rectifies the issue with the units of $\widetilde{\phi}(f)$. The expected energy resolution, σ_A^2 is unchanged since it is determined from the second quadratic term in χ^2 and thus is independent of the time offset. However, in the language of filters, it can be seen to have a more intuitive meaning by rewriting:

$$\sigma_A^2 = \frac{\int_{-\infty}^{\infty} df\, J(f) \left|\frac{\widetilde{s}(f)}{J(f)}\right|^2}{\left[\int_{-\infty}^{\infty} df\, \frac{|\widetilde{s}(f)|^2}{J(f)}\right]^2} \tag{B.3}$$

The numerator is the variance in time of the filtered trace: $J(f)\,|\widetilde{s}(f)/J(f)|^2$ is the filtered power spectral density. The denominator is the square of the peak amplitude of the trace given by filtering the expected pulse shape (whose peak height is 1) with the optimal filter. The expression is thus the squared noise-to-signal ratio, $(N/S)^2$, of the filtered expected pulse shape. It is instructive to

write the signal-to-noise ratio explicitly for a pulse of amplitude A:

$$\frac{\widehat{A}}{\sigma_A} = \frac{\int_{-\infty}^{\infty} df \frac{e^{j\omega\widehat{t_0}}\widetilde{s}^*(f)\widetilde{v}(f)}{J(f)}}{\left[\int_{-\infty}^{\infty} df\, J(f) \left|\frac{\widetilde{s}(f)}{J(f)}\right|^2\right]^{1/2}} \tag{B.4}$$

That is, the signal-to-noise ratio is just the ratio of the peak height of the filtered pulse to the filtered noise in time.

Viewed as an impulse response function, the optimal filter is always acausal since it is the time reverse of $\mathcal{FT}[\widetilde{s}(f)/J(f)]$, which is causal because $s(t)$ is causal and $J(f)$ is real. In practice, this means optimal filtering cannot be implemented in hardware: to construct the optimally filtered pulse at a given time, one has to know the value of the incoming pulse in the future!

In the case that the noise is white, the formulae become more intuitive. If the noise is white, the optimal filter is just $\widetilde{\phi}(f) = \widetilde{s}^*(f)$, which yields $\phi(t) = s(-t)$. Since $g_r * h = g \otimes h$ where $g_r(t) = g(-t)$, optimal filtering corresponds to cross-correlating the pulse with the expected shape. The best time offset is given by the time offset that maximizes the cross-correlation, as one would expect.

The expected time resolution can also be calculated using the second derivative of the χ^2; in fact, this is the only way to calculate it, given that there is no explicit form for $\widehat{t_0}$. This yields

$$\sigma_{t_0}^2 = \left[A^2 \int_{-\infty}^{\infty} df\, \omega^2 \frac{|s(f)|^2}{J(f)}\right]^{-1} \tag{B.5}$$

As one might expect, an additional frequency-weighting factor enters, expressing the fact that time-offset estimation requires more high-frequency information than amplitude estimation. The time resolution decreases as $1/A$: the more energetic the pulse is, the better the time estimation, which is also reasonable.

Finally, it is useful to convert the above formulae to the discrete case. The correspondence is made using

$$\begin{aligned}
\widetilde{s}^*(f) &\longleftrightarrow \frac{\tilde{s}_n^*}{\Delta f} \\
\widetilde{v}(f) &\longleftrightarrow \frac{v_n}{\Delta f} \\
J(f) &\longleftrightarrow J(f_n) \\
\int_{-\infty}^{\infty} df &\longleftrightarrow \sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \Delta f
\end{aligned}$$

which yields

$$\begin{aligned}
\widehat{A} &= \frac{\sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \Delta f e^{j\omega_n \widehat{t_0}} \frac{\tilde{s}_n^*}{\Delta f} \frac{v_n}{\Delta f} \frac{1}{J(f_n)}}{\sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \Delta f \frac{|\tilde{s}_n|^2}{[\Delta f]^2} \frac{1}{J(f_n)}} \\
&= \frac{\sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} e^{j\omega_n \widehat{t_0}} \frac{\tilde{s}_n^* v_n}{J(f_n)}}{\sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \frac{|\tilde{s}_n|^2}{J(f_n)}}
\end{aligned}$$

and

$$\begin{aligned}\sigma_A^2 &= \left[\sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \Delta f \frac{|\tilde{s}_n|^2}{[\Delta f]^2}\frac{1}{J(f_n)}\right]^{-1} \\ &= \left[2T\sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \frac{|\tilde{s}_n|^2}{J(f_n)}\right]^{-1}\end{aligned}$$

Note the factor of $2T$: this arises because $J(f_n)$ is in V^2/Hz.

B.5 Non-Optimal Pulse-Height Estimation

As mentioned above, optimal filtering cannot be implemented in an analog circuit in hardware. Therefore, it is of interest to determine the expected resolution expression for causal filters. An amplitude estimator can be calculated by analogy to the optimal filter case, giving, for a non-optimal filter $\widetilde{g}(f)$:

$$\widehat{A} = \frac{\int_{-\infty}^{\infty} df\, e^{j\omega(\widehat{t_0}+t^*)}\widetilde{g}(f)\widetilde{v}(f)}{\int_{-\infty}^{\infty} df\, e^{j\omega t^*}\widetilde{g}(f)\widetilde{s}(f)} \tag{B.6}$$

where the optimal filter $\widetilde{s}^*(f)/J(f)$ has been replaced by the non-optimal filter $\widetilde{g}(f)$. The extra exponential factor $e^{j\omega t^*}$ requires some explanation. The denominator of the optimal-filter expression, Equation B.2, is the peak of the trace given by filtering the expected pulse shape $s(t)$ with the optimal filter. Effectively, this is a correction factor for the fact that filtering changes the pulse amplitude. When filtering the expected shape by the optimal filter, the peak automatically occurs at zero time offset because the optimal filter and expected pulse shape are already aligned in time. This does not necessarily hold for the non-optimal filter, so the peak has to be found empirically. Furthermore, the time offset of the peak, t^*, must be used to correct the offset found for the data trace: even a pulse with $t_0 = 0$ has its filtered peak offset to t^*. One can shown that the above estimator is also unbiased.

The expected energy resolution of the non-optimal filter must be calculated from $\sigma_A^2 = \langle[\widehat{A}]^2\rangle - A^2$ since $\widetilde{g}(f)$ does not minimize χ^2. Again, this is an exercise in Fourier transform algebra that I do not perform explicitly; the result is

$$\sigma_A^2 = \frac{\int_{-\infty}^{\infty} df\, J(f)|\widetilde{g}(f)|^2}{\left[\int_{-\infty}^{\infty} df\, e^{j\omega t^*}\, \widetilde{g}(f)\, \widetilde{s}(f)\right]^2} \tag{B.7}$$

One can easily see that this reduces to the optimal-filter expression, Equation B.1, in the case $\widetilde{g}(f) = \widetilde{s}^*(f)/J(f)$ and $t^* = 0$. Furthermore, the expression is analogous to the alternate expression for the optimal-filter energy resolution given by Equation B.3. The numerator is the variance in time of the noise after filtering: $J(f)|\widetilde{g}(f)|^2$ is the filtered noise power spectral density. The denominator is the square of the peak amplitude of the trace given by filtering the expected shape $s(t)$ by the non-optimal filter, which is a normalizing factor of the type seen in Equation B.3. The

signal-to-noise ratio for a pulse of amplitude A in the non-optimal-filtering case is

$$\frac{\widehat{A}}{\sigma_A} = \frac{\int_{-\infty}^{\infty} df e^{j\omega(\widehat{t_0}+t^*)}\widetilde{g}(f)\widetilde{v}(f)}{\left[\int_{-\infty}^{\infty} df\, J(f)|\widetilde{g}(f)|^2\right]^{-1/2}} \tag{B.8}$$

As with the optimal-filter case, Equation B.4, the signal-to-noise ratio is just the ratio of the peak height of the filtered pulse to the filtered noise in time.

To illustrate the non-optimal case further, consider the simplest possible non-optimal filter: no filtering, $\widetilde{g}(f) = 1$. The pulse-height estimator is just the peak amplitude of the unfiltered pulse because the denominator of Equation B.6 reduces to 1 and the numerator is just the pulse itself. The energy resolution expression, Equation B.7, also reduces: the denominator is 1 and the numerator is $\langle[v(t)]^2\rangle = \int_{-\infty}^{\infty} df\, J(f)$. Since one uses the peak of the unfiltered pulse as the estimator, the noise in time gives the resolution directly.

The expressions in the discrete case for the non-optimal amplitude estimator and resolution are

$$\begin{aligned}
\widehat{A} &= \frac{\sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} e^{j\omega_n \widehat{t_0}}\tilde{g}_n\tilde{v}_n}{\sum_{n=-\frac{N}{2}}^{\frac{N}{2}-1} \tilde{g}_n\tilde{s}_n} \\
\sigma_A^2 &= \frac{1}{2T}\frac{\sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} J(f_n)|\tilde{g}_n|^2}{\left[\sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} e^{j\omega t^*}\,\tilde{g}_n\tilde{s}_n\right]^2}
\end{aligned}$$

B.6 Applications

The obvious application of the optimal-filtering formalism is to optimally fit the energy of acquired pulses. Optimal filtering can be implemented fully on digitized traces of sufficient length. For Run 19, optimal filtering was particularly necessary because of the large amount of 1-kHz cross-talk — optimal filtering notches out the frequencies with large cross-talk, leaving the resolution essentially unchanged from the case in which there is no cross-talk. Optimal filtering is implemented in a fairly straightforward way. The trace is Fourier transformed and filtered by the optimal filter, which is calculated from noise traces and the pulse template. The filtered pulse is transformed back to time. The time offset and amplitude are given by the time offset of the peak and the peak amplitude. Other aspects of the implementation are discussed in Chapter 6.

An important application of both the amplitude resolution expression, Equation B.1, and the time resolution expression, Equation B.5, is to determine the maximum digitization frequency needed. The ionization pulses can be modelled approximately as 2-pole exponentials, with one pole at the fall time, τ_f, and a second at the rise time, τ_r, with

$$s(f) = \frac{\tau_f}{1+j\omega\tau_f}\frac{\tau_r}{1+j\omega\tau_r}$$

$s(f)$ is flat up to $(2\pi\tau_f)^{-1}$, falls as $1/f$ between $(2\pi\tau_f)^{-1}$ and $(2\pi\tau_r)^{-1}$, and falls as $1/f^2$ above $(2\pi\tau_r)^{-1}$. Assuming white noise, which is not a bad approximation for the ionization channels,

Channel	BLIP3	BLIP4	BLIP5	BLIP6
Optimal Filter				
QI [μV]	10	11	12	12
QO [μV]	10	13	15	17
P1 [nV]	23	32	32	32
P2 [nV]	22	32	32	34
Non-Optimal Filter				
QI [μV]	24	22	28	25
QO [μV]	28	31	33	37
P1 [nV]	64	76	61	73
P2 [nV]	60	79	65	111

Table B.1: Resolutions for optimal and non-optimal filtering for Run 19 ionization and phonon channels. Resolutions are in pulse height at the first-stage ionization amplifier output and the phonon amplifier input. These are *calculated*, not measured, values.

the interesting quantity, $|s(f)|^2/J(f)$, is flat, falls as $1/f^2$, and falls as $1/f^4$ in the three regimes. Thus, the integral of $\frac{|s(f)|^2}{J(f)}$ up to a frequency f_c, which is the denominator of the amplitude resolution formula, Equation B.1, is linear in f_c up to $(2\pi\tau_f)^{-1}$, goes as $1/f_c$ between $(2\pi\tau_f)^{-1}$ and $(2\pi\tau_r)^{-1}$, and as $1/f_c^3$ above $(2\pi\tau_r)^{-1}$. Since τ_f and τ_r are typically separated by about an order of magnitude, little improvement in energy resolution is gained above $(2\pi\tau_f)^{-1}$. For the time resolution, two additional powers of frequency enter, so the scalings are as f_c^2, f_c^0, and $1/f_c$ in the three regimes. Little improvement in time resolution is gained above $(2\pi\tau_r)^{-1}$. As is described below, this result only holds for fitting to an analytic form; it is modified for template fitting.

The main application of non-optimal filtering is to find the optimal filters for the trigger circuitry. The filter that gives the best resolution also gives the lowest energy threshold because it maximizes the signal-to-noise ratio. The resolution expression, Equation B.7, can be used to compare different types of filters and to optimize the parameters of a given type of filter. The RTF trigger circuits use 1-pole low-pass and high-pass filters whose time constants have been optimized using Equation B.7. It is interesting to compare the expected optimal and non-optimal resolutions to see how much worse the resolution is for this particular non-optimal filter. This is done in Table B.1 for the Run 19 ionization and phonon channels. In general, the non-optimally filtered resolutions are a factor of 2 to 3 worse than the optimally filtered resolutions.

B.7 Real-World Effects

B.7.1 Sampling Period

If there are sharp features in the noise spectrum, such as 1-kHz cross-talk, it is important to use a long sampling period, $2T$, so that the frequency bin width, $\frac{1}{2T}$, is very small compared to the total bandwidth. This allows the bins containing spectral features to be discarded with negligible loss of signal-to-noise. For Run 19, very long ionization traces, 12.8 ms, were taken, yielding $\Delta f = 78.125$ Hz. The fall time of the ionization pulses is about 40 μs, corresponding to

about 4 kHz, so less than 10% of the bandwidth useful for energy estimation was lost by discarding the 1, 2, 3, and 4 kHz bins.

B.7.2 Expected Amplitude Distribution for Fit with Time Search

The amplitude resolution given by Equation B.1 is only valid if the time offset is not strongly correlated with the amplitude estimate. This holds if the pulse is of sufficient amplitude that the time offset is well established, independent of fluctuations around the expected amplitude; recall that the time-offset resolution decreases as $1/A$. However, near threshold or for noise traces, the time offset and amplitude are strongly correlated. An expression for the expected distribution of fitted amplitudes can be calculated in this case. The following argument was first suggested to me by Rick Gaitskell.

The case of a pulse well above threshold can be mimicked for noise traces by fixing the time offset and calculating the amplitude estimator in the standard way. The expected distribution of amplitudes is Gaussian with width σ_V, as calculated by optimal or non-optimal filtering. That is,

$$P(V) = \frac{1}{\sigma_V\sqrt{2\pi}}\,\exp\left(-\frac{V^2}{2\sigma_V^2}\right)$$

The Gaussianity of the distribution is corroborated by the baseline resolution distributions shown in Chapter 4.

Now, let t_0 be a free parameter that can take values over a range $M\Delta t$, where Δt is the sampling interval. M indicates the width of the search window in samples. Also, assume that the values of t_0 are restricted so the pulse does not approach the edges of the digitized trace; *i.e.*, no edge effects occur. Each value of t_0 yields a value $V(t_0)$ drawn from the above Gaussian distribution. The optimal-filter prescription is to take the largest such value of V as the best pulse-height estimator. The value of t_0 yielding this value of V is the best time estimator. Thus, the largest V, V_{max}, out of a sample of M values of V drawn from the above distribution is taken. The probability of this procedure giving a value V_{max} is given by the probability that there is value of t_0 giving $V = V_{max}$ and that the other $M-1$ samples give $V < V_{max}$. There are M ways for this to happen since any of the M samples may take the value V_{max}. Thus, the distribution for V_{max} is

$$\begin{aligned} P(V_{max}) &= M\left[\int_{-\infty}^{V_{max}} dV\,P(V)\right]^{M-1} P(V_{max}) \\ &= M\left[\mathrm{erf}(V_{max},\sigma_V)\right]^{M-1}\frac{1}{\sigma_V\sqrt{2\pi}}\exp\left(-\frac{V_{max}^2}{2\sigma_V^2}\right) \end{aligned}$$

B.7.3 Jitter

As shown above, for a 2-pole exponential pulse, the bulk of the signal-to-noise for the t_0 estimation comes from frequencies below $(2\pi\tau_r)^{-1}$. Therefore, in principle, it is unnecessary to digitize with a Nyquist frequency larger than $(2\pi\tau_r)^{-1}$. However, this result only holds in the case that the pulses are fit to an analytic form that can be calculated for time shifts that are not integral values of the sample interval Δt. If one is performing a template fit, the template can only be shifted by integral multiples of Δt. Essentially, the convolution of the optimal filter with

the pulse cannot be calculated between time bins. This degrades both the time and amplitude estimation; the phenomenon is a form of time jitter. The time shift can only be calculated to a precision of Δt. The effect on the amplitude is slightly more subtle: since the peak of the convolution gives the amplitude estimation, and one does not have the convolution for non-integral time values, the peak of the convolution is underestimated when one takes the largest of the two values that bracket the peak. The expected underestimation of the amplitude is

$$\begin{aligned}
\delta E(\delta t) &= E(0) - E(\delta t) \\
&= E(0) \frac{\sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \frac{|\tilde{s}_k|^2}{J_k} \left[1 - e^{j\omega_k \delta t}\right]}{\sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \frac{|\tilde{s}_k|^2}{J_k}} \\
&= E(0) \frac{\sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \frac{|\tilde{s}_k|^2}{J_k} \left[1 - \cos(2\pi f_k \delta t)\right]}{\sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \frac{|\tilde{s}_k|^2}{J_k}}
\end{aligned}$$

which is just the difference between the peak amplitude of the convolution and the amplitude a time δt away from the peak. (The third line is derived from the second by shuffling the negative and positive frequency components around.) The expression vanishes if $\delta t = 0$; *i.e.*, if the peak of the convolution coincides with a time sample. Richard Schnee has shown that one can correct for this effect by using multiple templates, shifted by time offsets smaller than Δt. The implementation of this for Run 19 is discussed in Chapter 6.

Appendix C

Phonon-System Misbehavior

A number of possible manifestations of a problem with the BLIP phonon channels were seen during Runs 19A and 19B, as indicated in the run chronologies presented in Chapter 7. The symptoms became progressively worse during May, 1999, until it was necessary to stop taking data on June 4 and study the problem intensively. Two different symptoms were encountered:

- The primary symptom, which began appearing on April 13, was outbursts in the phonon-trigger rate due to significantly increased noise in one or more detectors' phonon channels. There were two forms of this symptom:
  - Excess low-frequency noise, typically only increasing the noise below about 100 Hz.
  - Periodic events, with slow rise times suggesting "thermophonic" (microphonics depositing mechanical energy in the detectors, yielding a thermal signal) origin. The episodes were transient and the pulses varying in amplitude during an episode, with pulses correlated between detectors. The events were strikingly periodic, though the actual frequency was seen to vary. There appear to have been two modes, one around 10 Hz and another around 1 Hz.
- On or around December 4, 1998, the phonon pulses of all the detectors became "slow" — all the rise and fall times increased by a factor of a few. However, the pulses generated by the IBAPACAP appeared unchanged from the standard shape. The DC refs were normal, indicating the thermistor temperatures were no different from normal, and did not change when the IBAPACAP was enabled. This symptom recurred in BLIPs 1 and 2 on May 28, 1999; they were turned off in response. It recurred again during the studies that began June 4, 1999.

Furthermore, when any of the above symptoms occurred, they tended to occur in more than one detector at a time, though typically an onset was seen in one detector, then another, and so on.

It was found that, in all cases, a "LED bake" (turning on the LEDs used for neutralization for 20 minutes) ended the episodes, though, over time, the quiescent period thereby induced decreased in length from many days to less than an hour. A LED bake not only deposits an enormous number of photons into the detectors, but also deposits of order 1 mW for 20 minutes into the

detectors and their mounting hardware. This heating raises the entire mixing-chamber layer of the Icebox up to hundreds of mK. Similar quiescent periods, though less long-lived, were obtained by cycling the AC power on the front-end electronics, which also dumps a large amount of power into the refrigerator as the amplifiers find their quiescent state. This type of recovery was only seen if the electronics were cycled with the Detector I/O cables (the cables connecting the electronics to the detector striplines) connected, indicating that the symptoms were not simply due to the electronics boards finding some anomalous state. Even more convincingly, it was found that applying power directly to the detectors by applying a DC current to the IBAPACAP pulser on the detector, which results in a heating of the detector alone, momentarily alleviated the slow pulse problem, but the symptoms reappeared as the detector cooled down to its original temperature. This test was not tried during a noise episode.

In parallel with the above observations, a number of other possible causes were explored. In Run 18, ice in the liquid nitrogen bath caused bubbling that gave rise to thermophonic events. The bubbling could be prevented by allowing the LN bath to pressurize. This was tried and had no effect on the outbursts. Furthermore, LED baking, which did work, dumps a negligible power into the LN layer and so would be unable to affect a bubbling problem.

Another frequent microphonics suspect in dilution refrigerators is the pot needle valve. This valve permits a minute flow of LHe from the bath into the pot to replenish the helium that is pumped off. The flow of the liquid through the valve is known to give rise to microphonics seen in both the test facility and the Icebox previously via noise in the ionization circuits. Closing the pot needle valve temporarily relieves such microphonics by stopping the liquid flow, and this was attempted, to no effect. However, the pot cannot be completely ruled out as a cause because LED baking dumps enough power into the refrigerator to raise the pot pressure from its usual 0.8–0.9 mm Hg to a few mm Hg. Such a drastic change in the pot state could easily affect the flow of liquid from the bath.

Given that the observed symptoms could also be caused by electronics problems, some detectors were switched back to the 3U-format front-end electronics used in previous runs. This was done in the midst of a periodic pulse episode. By turning on the AC power with the Detector I/O cables disconnected and then plugging in the cables, one can avoid dumping significant power into the detectors. This was done to avoid killing the symptoms via the power cycle. The periodic pulsing remained with the 3U electronics. Furthermore, after a power cycle to alleviate the periodic pulsing, the detectors fell into the "slow pulse" state and slow pulses were seen from detectors connected to both 3U and 9U electronics.

To check the thermal state of the detectors, a 10-Hz square-wave bias, offset so as to apply zero power during half of the cycle, was applied to one thermistor of one detector and the other thermistor was observed with the standard 1-kHz sine wave and lockin setup. The slow square-wave bias provides a periodic power input that cycles the crystal temperature with sharp transitions. The rise and fall times observed in the second thermistor were typical of standard pulse time constants, indicating the heat capacities and conductances of the thermal circuit were at the nominal values. This result is reasonable, as it is analogous to the observation of normally shaped events from IBAPACAP pulses.

Finally, there was already an expectation of large amounts of He gas in the IVC. During both the Run 19A and 19B LHe cooldowns, helium was allowed to condense in the bath while the IVC pressure of He exchange gas was quite high, tens to 100 mT. Once liquid has condensed,

the He exchange gas becomes so cold that a standard turbopump cannot remove it. This was indeed observed — the IVC pressure remained in the vicinity of 3 mT for all of Run 19B in spite of continual turbopumping. Furthermore, a leak detector connected to the IVC observed a sharp ($> 10\times$) increase in the ^{4}He leak rate during He transfers (though no change in IVC pressure was observed). Thus, additional He was being allowed into the IVC over time. In addition, sudden IVC pressure excursions (going up to tens of mT from the 3 mT baseline) were observed. The reality of the excursions was confirmed by correlated increases in base temperature due to the increased thermal conductance between the mixing-chamber and higher-temperature layers.

The most plausible explanation so far found for these symptoms is that the high ^{4}He pressure in the IVC led to condensation of He onto the detectors. For this explanation to work, the resulting superfluid helium film must be decoupled from the thermal phonons generated by any of the various electrical heating mechanisms, but must be well coupled to the athermal phonons produced initially by particle interactions. Presumably, the He film could absorb athermal phonons and convert their energy into long-lived excitations. If the decay of these excitations occurs over tens of ms, they would return the absorbed energy to the detector very slowly, resulting in slow pulses. IBAPACAP pulsing and square-wave thermistor biasing would not produce athermal phonons and thus would exhibit their normal behaviors. The noise and periodic pulsing observed might be explained by evaporation and recondensation of He on the detectors, or even macroscopic dripping. The assorted heating impulses (LED bakes, power cycles, DC bias of IBAPACAP) would deposit enough power to temporarily evaporate the films, but He would recondense on the cold, clean detector surfaces soon — an empty surface in the presence of so much He gas is not an equilibrium situation.

Based on these arguments, the Icebox was cycled to room temperature during June and recooled, being careful to achieve a lower He pressure in the IVC before condensation of LHe. The residual exchange-gas pressure was approximately 10 mT when liquid was condensed during the cooldown, about an order of magnitude less than the previous cooldown. Only one phonon-trigger outburst was observed during the three months of Run 19C.

Appendix D

Veto-Anticoincident Nuclear-Recoil Event Parameters

This appendix contains a list of event parameters for each of the 13 single-scatter and 4 multiple-scatter nuclear recoils. The parameter lists also indicate whether the event passes various cuts (defined in Chapter 7) by a 0 or 1 for event parameters beginning with "c." The definitions of the event parameters are as listed in Table D.1.

SeriesNumber	number of file series, in the form YYMMDDHHMM
EventNumber	event number in series
EventTime	time of event, in Macintosh seconds (seconds since Jan 1, 1904)
LiveTime	time between enabling of trigger and event trigger [ms]
TimeBetween	time between event trigger of this event and the previous event [ms]
LastISRTime	time since last front-end subrack command control [s]
VTPreTimeFast	time between event trigger and last veto trigger [μs]
VTNearQTime	difference between event time, as determined from ionization delay, and nearest veto trigger [μs]
VTNearPTime	difference between event time, as determined from phonon delay, and nearest veto trigger [μs]
VetoXXXHigh	XXX veto side pulse integral in the ADC triggered by veto high trigger (not available for Runs 19A/B)
VetoXXXLow	XXX veto side pulse integral in the ADC triggered by veto window trigger (not available for Runs 19A/B)
pri	recoil energy calculated using inner-electrode ionization signal only; valid for cQinOnly events [keV]
qi	inner-electrode ionization energy [keV]
pr	recoil energy calculated using summed ionization signal [keV]
qsum	summed ionization signal [keV]
qo	outer-electrode ionization energy [keV]
qpart	ionization partition, (qi - qo)/qsum
pheat	phonon energy [keV]
P1bs, P2bs	mean of phonon pretrigger baseline [digitizer bins]
P1std, P2std	standard deviation of phonon pretrigger baseline [digitizer bins]
cQTrig	1 if ionization trigger occurred for the given detector at trigger time
cPTrig	1 if phonon trigger occurred for the given detector at trigger time
QPreTime	time of last ionization trigger prior to event trigger [μs]. Can be small (-1 μs) and negative for triggers coincident with event trigger. -99999 indicates timeout.
QPostTime	time of first ionization trigger after event trigger [μs]. Can be small and positive for coincident events. 99999 indicates timeout.
PPostTime	time of first phonon trigger after event trigger [μs]. For ionization triggers, this is typically 2000–4000 μs because of the slow rise of the phonon pulse. Becomes later as the phonon energy decreases (pulse slewing effect). 99999 indicates timeout.

Table D.1: Definitions of event parameters.

single scatter 1

SeriesNumber = 9903232041 EventNumber = 130628
EventTime = 3005118847
LiveTime = 1753 TimeBetween = 1985 LastISRTime = 402
VTPreTimeFast = -268 VTNearQTime = 201 VTNearPTime = 3889
VetoTopHigh -999999
VetoTopLow -999999
VetoCrackHigh -999999
VetoCrackLow -999999
VetoNorthHigh -999999
VetoNorthLow -999999
VetoSouthHigh -999999
VetoSouthLow -999999
VetoEastHigh -999999
VetoEastLow -999999
VetoWestHigh -999999
VetoWestLow -999999
VetoBottomHigh -999999
VetoBottomLow -999999
VetoORHigh -999999
VetoORLow -999999

quantity	103	104	105	106
pri	0.01	-0.12	-0.48	55.16
qi	-0.09	0.01	0.27	19.68
pr	0.53	0.79	-0.21	55.13
qsum	-0.35	-0.45	0.14	19.70
qo	-0.26	-0.46	-0.14	0.02
qpart	-0.47	-1.06	2.99	1.00
pheat	-0.17	-0.10	0.07	94.52
cQinThresh	0	0	0	1
cQoutNoise	1	1	1	1
cQsumThresh	0	0	0	1
cQinOnly	0	0	0	1
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	0	0	1
cDataQual	0	0	0	1
P1bs	580.87	605.50	743.36	685.13
P2bs	561.35	593.28	687.19	696.33
P1std	1.22	1.43	1.27	1.60
P2std	1.43	1.11	1.33	1.82
cQTrig	0	0	0	1
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	99999	2543

single scatter 2

SeriesNumber = 9903232041 EventNumber = 160980
EventTime = 3005125243
LiveTime = 1075 TimeBetween = 1302 LastISRTime = 1130
VTPreTimeFast = -388 VTNearQTime = -45.2 VTNearPTime = 167
VetoTopHigh -999999
VetoTopLow -999999
VetoCrackHigh -999999
VetoCrackLow -999999
VetoNorthHigh -999999
VetoNorthLow -999999
VetoSouthHigh -999999
VetoSouthLow -999999
VetoEastHigh -999999
VetoEastLow -999999
VetoWestHigh -999999
VetoWestLow -999999
VetoBottomHigh -999999
VetoBottomLow -999999
VetoORHigh -999999
VetoORLow -999999

quantity	103	104	105	106
pri	0.56	1.22	0.91	13.55
qi	-0.02	0.04	-0.44	1.77
pr	0.03	-0.21	0.86	12.18
qsum	0.24	0.75	-0.42	2.46
qo	0.26	0.71	0.02	0.69
qpart	-1.17	-0.90	1.11	0.44
pheat	0.52	1.29	0.02	17.10
cQinThresh	0	0	0	1
cQoutNoise	1	0	1	1
cQsumThresh	0	0	0	1
cQinOnly	0	0	0	1
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	0	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	0	0	1
cDataQual	0	0	0	1
P1bs	677.66	705.00	691.85	706.39
P2bs	647.10	693.58	626.60	711.18
P1std	1.43	1.30	1.14	1.59
P2std	1.48	1.49	1.38	2.56
cQTrig	0	0	0	0
cPTrig	0	0	0	1
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	99999	99999

single scatter 3

```
SeriesNumber = 9904141258   EventNumber = 240254
EventTime = 3007002813
LiveTime = 336   TimeBetween = 569   LastISRTime = 1448
VTPreTimeFast = -105   VTNearQTime = 66.8   VTNearPTime = -118
VetoTopHigh        -999999
VetoTopLow         -999999
VetoCrackHigh      -999999
VetoCrackLow       -999999
VetoNorthHigh      -999999
VetoNorthLow       -999999
VetoSouthHigh      -999999
VetoSouthLow       -999999
VetoEastHigh       -999999
VetoEastLow        -999999
VetoWestHigh       -999999
VetoWestLow        -999999
VetoBottomHigh     -999999
VetoBottomLow      -999999
VetoORHigh         -999999
VetoORLow          -999999
```

quantity	103	104	105	106
pri	0.30	-0.20	0.10	38.43
qi	-0.09	0.10	-0.08	12.85
pr	0.08	0.52	-0.36	39.34
qsum	0.02	-0.26	0.15	12.39
qo	0.11	-0.36	0.23	-0.45
qpart	-9.54	-1.76	-1.99	1.07
pheat	0.12	-0.00	-0.05	64.12
cQinThresh	0	0	0	1
cQoutNoise	1	1	1	1
cQsumThresh	0	0	0	1
cQinOnly	0	0	0	1
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	0	0	1
cDataQual	0	0	0	1
P1bs	502.63	327.15	469.72	557.09
P2bs	466.53	320.74	397.12	562.97
P1std	1.39	1.48	1.26	1.39
P2std	1.29	1.45	1.26	1.22
cQTrig	0	0	0	1
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	99999	3503

single scatter 4

```
SeriesNumber = 9904151325   EventNumber = 410598
EventTime = 3007107389
LiveTime = 1459   TimeBetween = 1723   LastISRTime = 259
VTPreTimeFast = -894   VTNearQTime = 193.8   VTNearPTime = 1027
VetoTopHigh        -999999
VetoTopLow         -999999
VetoCrackHigh      -999999
VetoCrackLow       -999999
VetoNorthHigh      -999999
VetoNorthLow       -999999
VetoSouthHigh      -999999
VetoSouthLow       -999999
VetoEastHigh       -999999
VetoEastLow        -999999
VetoWestHigh       -999999
VetoWestLow        -999999
VetoBottomHigh     -999999
VetoBottomLow      -999999
VetoORHigh         -999999
VetoORLow          -999999
```

quantity	103	104	105	106
pri	0.72	0.45	0.65	23.95
qi	0.66	-0.04	-0.19	7.86
pr	1.24	0.76	0.66	23.83
qsum	0.39	-0.19	-0.19	7.92
qo	-0.26	-0.15	-0.00	0.06
qpart	2.34	-0.58	0.95	0.98
pheat	2.03	0.37	0.28	39.67
cQinThresh	0	0	0	1
cQoutNoise	1	1	1	1
cQsumThresh	0	0	0	1
cQinOnly	0	0	0	1
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	0	0	1
cDataQual	0	0	0	1
P1bs	80.13	21.45	258.19	186.55
P2bs	41.34	29.83	161.38	182.39
P1std	1.42	1.08	1.16	1.42
P2std	0.98	1.37	2.53	1.36
cQTrig	0	0	0	1
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	7098	99999	99999	3680

single scatter 5

SeriesNumber = 9904291555 EventNumber = 50264
EventTime = 3008279046
LiveTime = 558 TimeBetween = 800 LastISRTime = 737
VTPreTimeFast = -252 VTNearQTime = 75 VTNearPTime = 1013
VetoTopHigh -999999
VetoTopLow -999999
VetoCrackHigh -999999
VetoCrackLow -999999
VetoNorthHigh -999999
VetoNorthLow -999999
VetoSouthHigh -999999
VetoSouthLow -999999
VetoEastHigh -999999
VetoEastLow -999999
VetoWestHigh -999999
VetoWestLow -999999
VetoBottomHigh -999999
VetoBottomLow -999999
VetoORHigh -999999
VetoORLow -999999

quantity	103	104	105	106
pri	0.11	26.72	-0.27	-1.27
qi	0.06	8.50	0.09	0.59
pr	0.66	27.86	0.37	0.05
qsum	-0.22	7.92	-0.22	-0.06
qo	-0.28	-0.57	-0.32	-0.66
qpart	-1.55	1.14	-1.82	-20.01
pheat	0.23	43.71	-0.08	-0.08
cQinThresh	0	1	0	0
cQoutNoise	1	1	1	1
cQsumThresh	0	1	0	0
cQinOnly	0	1	0	0
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	0
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	1	0	0
cDataQual	0	1	0	0
P1bs	452.50	380.22	451.17	460.75
P2bs	415.28	375.79	376.10	463.69
P1std	1.31	1.28	1.11	1.38
P2std	1.28	1.50	1.35	1.38
cQTrig	0	1	0	0
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	3023	99999	99999

single scatter 6

```
SeriesNumber = 9905021512   EventNumber = 290639
EventTime = 3008577524
LiveTime = 1220   TimeBetween = 1448   LastISRTime = 293
VTPreTimeFast = -216   VTNearQTime = 67.8   VTNearPTime = 4174
VetoTopHigh        -999999
VetoTopLow         -999999
VetoCrackHigh      -999999
VetoCrackLow       -999999
VetoNorthHigh      -999999
VetoNorthLow       -999999
VetoSouthHigh      -999999
VetoSouthLow       -999999
VetoEastHigh       -999999
VetoEastLow        -999999
VetoWestHigh       -999999
VetoWestLow        -999999
VetoBottomHigh     -999999
VetoBottomLow      -999999
VetoORHigh         -999999
VetoORLow          -999999
```

quantity	103	104	105	106
pri	-0.02	0.23	-0.12	23.86
qi	-0.18	0.17	0.26	5.42
pr	-1.08	0.32	-0.15	23.16
qsum	0.35	0.12	0.28	5.77
qo	0.53	-0.05	0.02	0.35
qpart	-2.06	1.76	0.88	0.88
pheat	-0.39	0.56	0.41	34.70
cQinThresh	0	0	0	1
cQoutNoise	1	1	1	1
cQsumThresh	0	0	0	1
cQinOnly	0	0	0	1
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	0	0	1
cDataQual	0	0	0	1
P1bs	968.01	896.64	920.31	929.51
P2bs	935.87	871.38	874.76	937.61
P1std	1.36	1.61	1.05	1.85
P2std	1.46	1.78	1.05	1.86
cQTrig	0	0	0	1
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	99999	3109

single scatter 7

```
SeriesNumber = 9905131333   EventNumber = 90702
EventTime = 3009487068
LiveTime = 570   TimeBetween = 833   LastISRTime = 185
VTPreTimeFast = -164   VTNearQTime = 134.4   VTNearPTime = 2621
VetoTopHigh        -999999
VetoTopLow         -999999
VetoCrackHigh      -999999
VetoCrackLow       -999999
VetoNorthHigh      -999999
VetoNorthLow       -999999
VetoSouthHigh      -999999
VetoSouthLow       -999999
VetoEastHigh       -999999
VetoEastLow        -999999
VetoWestHigh       -999999
VetoWestLow        -999999
VetoBottomHigh     -999999
VetoBottomLow      -999999
VetoORHigh         -999999
VetoORLow          -999999
```

quantity	103	104	105	106
pri	0.34	15.67	0.23	0.61
qi	-0.09	4.37	-0.24	-0.32
pr	0.55	15.15	1.16	0.03
qsum	-0.19	4.62	-0.70	-0.03
qo	-0.10	0.26	-0.46	0.29
qpart	-0.05	0.89	-0.32	19.09
pheat	0.16	24.40	-0.24	-0.04
cQinThresh	0	1	0	0
cQoutNoise	1	1	1	1
cQsumThresh	0	1	0	0
cQinOnly	0	1	0	0
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	1	0	0
cDataQual	0	1	0	0
P1bs	538.61	417.18	585.89	560.04
P2bs	508.15	414.48	508.78	559.02
P1std	1.25	1.41	1.27	1.45
P2std	1.24	1.56	1.23	1.35
cQTrig	0	1	0	0
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	3385	99999	99999

single scatter 8

SeriesNumber = 9905191346 EventNumber = 140408
EventTime = 3010021639
LiveTime = 2322 TimeBetween = 2518 LastISRTime = 240
VTPreTimeFast = -30 VTNearQTime = -31.6 VTNearPTime = -163
VetoTopHigh -999999
VetoTopLow -999999
VetoCrackHigh -999999
VetoCrackLow -999999
VetoNorthHigh -999999
VetoNorthLow -999999
VetoSouthHigh -999999
VetoSouthLow -999999
VetoEastHigh -999999
VetoEastLow -999999
VetoWestHigh -999999
VetoWestLow -999999
VetoBottomHigh -999999
VetoBottomLow -999999
VetoORHigh -999999
VetoORLow -999999

quantity	103	104	105	106
pri	-0.73	82.19	0.66	-0.54
qi	0.30	25.72	-0.33	0.45
pr	-0.44	81.56	-0.98	-1.29
qsum	0.15	26.03	0.49	0.82
qo	-0.15	0.32	0.82	0.37
qpart	2.99	0.98	-2.36	0.09
pheat	-0.14	133.61	-0.01	0.35
cQinThresh	0	1	0	0
cQoutNoise	1	1	0	1
cQsumThresh	0	1	0	0
cQinOnly	0	1	0	0
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	0
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	1	0	0
cDataQual	0	1	0	0
P1bs	501.75	487.00	604.33	541.75
P2bs	463.96	475.84	536.40	550.81
P1std	1.43	1.39	1.02	1.78
P2std	1.24	1.19	1.57	1.59
cQTrig	0	1	0	0
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	2652	99999	99999

single scatter 9

SeriesNumber = 9906021605 EventNumber = 140863
EventTime = 3011241504
LiveTime = 6266 TimeBetween = 6529 LastISRTime = 278
VTPreTimeFast = -185 VTNearQTime = 176.8 VTNearPTime = -490
VetoTopHigh -999999
VetoTopLow -999999
VetoCrackHigh -999999
VetoCrackLow -999999
VetoNorthHigh -999999
VetoNorthLow -999999
VetoSouthHigh -999999
VetoSouthLow -999999
VetoEastHigh -999999
VetoEastLow -999999
VetoWestHigh -999999
VetoWestLow -999999
VetoBottomHigh -999999
VetoBottomLow -999999
VetoORHigh -999999
VetoORLow -999999

quantity	103	104	105	106
pri	0.03	22.58	0.37	-0.10
qi	0.00	6.04	-0.30	0.09
pr	1.02	23.20	0.89	0.37
qsum	-0.49	5.73	-0.56	-0.14
qo	-0.49	-0.31	-0.26	-0.23
qpart	-1.01	1.11	0.07	-2.31
pheat	0.04	34.66	-0.22	0.09
cQinThresh	0	1	0	0
cQoutNoise	1	1	1	1
cQsumThresh	0	1	0	0
cQinOnly	0	1	0	0
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	1	0	0
cDataQual	0	1	0	0
P1bs	602.31	611.24	772.85	794.50
P2bs	568.85	595.42	716.03	799.68
P1std	1.25	1.30	1.25	1.19
P2std	1.33	1.66	1.23	1.72
cQTrig	0	1	0	0
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	3036	99999	99999

single scatter 10

SeriesNumber = 9907281402 EventNumber = 340427
EventTime = 3016119623
LiveTime = 4818 TimeBetween = 4987 LastISRTime = 45922
VTPreTimeFast = -324 VTNearQTime = 438.6 VTNearPTime = 413

VetoTopHigh	0
VetoTopLow	0
VetoCrackHigh	0
VetoCrackLow	0
VetoNorthHigh	0
VetoNorthLow	0
VetoSouthHigh	0
VetoSouthLow	0
VetoEastHigh	0
VetoEastLow	0
VetoWestHigh	0
VetoWestLow	0
VetoBottomHigh	0
VetoBottomLow	0
VetoORHigh	0
VetoORLow	0

quantity	103	104	105	106
pri	-0.36	0.12	15.22	0.82
qi	0.02	-0.06	4.55	-0.60
pr	-0.85	1.14	15.31	0.24
qsum	0.26	-0.57	4.50	-0.31
qo	0.25	-0.51	-0.05	0.29
qpart	-0.87	-0.78	1.02	2.90
pheat	-0.33	-0.00	24.31	-0.37
cQinThresh	0	0	1	0
cQoutNoise	1	1	1	1
cQsumThresh	0	0	1	0
cQinOnly	0	0	1	0
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	0	1	0
cDataQual	0	0	1	0
P1bs	672.62	670.82	682.23	667.68
P2bs	631.10	659.00	616.42	671.39
P1std	1.80	1.35	1.11	1.24
P2std	1.27	1.59	1.18	2.03
cQTrig	0	0	0	0
cPTrig	0	0	1	0
QPreTime	-99999	-99999	-99999	-4813862
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	99999	99999

single scatter 11

SeriesNumber = 9907311450 EventNumber = 190309
EventTime = 3016345204
LiveTime = 806 TimeBetween = 1041 LastISRTime = 257
VTPreTimeFast = -482 VTNearQTime = 93.4 VTNearPTime = -482
VetoTopHigh 0
VetoTopLow 0
VetoCrackHigh 0
VetoCrackLow 0
VetoNorthHigh 0
VetoNorthLow 0
VetoSouthHigh 0
VetoSouthLow 0
VetoEastHigh 0
VetoEastLow 0
VetoWestHigh 0
VetoWestLow 0
VetoBottomHigh 0
VetoBottomLow 0
VetoORHigh 0
VetoORLow 0

quantity	103	104	105	106
pri	0.32	-0.33	45.98	0.86
qi	-0.17	0.27	15.05	-0.41
pr	0.24	-0.29	45.66	0.72
qsum	-0.13	0.25	15.21	-0.34
qo	0.04	-0.02	0.16	0.07
qpart	1.58	1.16	0.98	1.40
pheat	-0.01	0.21	76.08	0.03
cQinThresh	0	0	1	0
cQoutNoise	1	1	1	1
cQsumThresh	0	0	1	0
cQinOnly	0	0	1	0
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	0	1	0
cDataQual	0	0	1	0
P1bs	377.43	417.69	523.47	536.32
P2bs	338.59	406.63	445.97	542.27
P1std	1.57	1.22	1.12	1.44
P2std	1.26	1.25	1.59	1.48
cQTrig	0	0	1	0
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	2627	99999

single scatter 12

SeriesNumber = 9908201410 EventNumber = 110537
EventTime = 3018053934
LiveTime = 4543 TimeBetween = 4752 LastISRTime = 1289
VTPreTimeFast = -27 VTNearQTime = -27 VTNearPTime = -479

VetoTopHigh	0
VetoTopLow	0
VetoCrackHigh	0
VetoCrackLow	0
VetoNorthHigh	0
VetoNorthLow	0
VetoSouthHigh	0
VetoSouthLow	0
VetoEastHigh	0
VetoEastLow	0
VetoWestHigh	0
VetoWestLow	0
VetoBottomHigh	0
VetoBottomLow	0
VetoORHigh	0
VetoORLow	0

quantity	103	104	105	106
pri	0.24	-0.62	-0.25	50.23
qi	-0.12	0.24	0.19	16.55
pr	0.13	-1.27	-0.23	49.60
qsum	-0.06	0.56	0.18	16.86
qo	0.05	0.33	-0.01	0.31
qpart	2.74	-0.16	1.15	0.96
pheat	0.01	-0.15	0.13	83.32
cQinThresh	0	0	0	1
cQoutNoise	1	1	1	1
cQsumThresh	0	0	0	1
cQinOnly	0	0	0	1
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	0	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	0	0	1
cDataQual	0	0	0	1
P1bs	405.54	465.90	586.38	599.55
P2bs	366.02	454.03	513.95	606.73
P1std	1.34	1.50	1.04	1.32
P2std	1.32	1.37	1.21	1.42
cQTrig	0	0	0	1
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	99999	2822

single scatter 13

SeriesNumber = 9908261350 EventNumber = 240974
EventTime = 3018600625
LiveTime = 2896 TimeBetween = 3105 LastISRTime = 1089
VTPreTimeFast = -242 VTNearQTime = 38.6 VTNearPTime = 77
VetoTopHigh 0
VetoTopLow 0
VetoCrackHigh 0
VetoCrackLow 0
VetoNorthHigh 0
VetoNorthLow 0
VetoSouthHigh 0
VetoSouthLow 0
VetoEastHigh 0
VetoEastLow 0
VetoWestHigh 0
VetoWestLow 0
VetoBottomHigh 0
VetoBottomLow 0
VetoORHigh 0
VetoORLow 0

quantity	103	104	105	106
pri	-0.57	1.31	11.02	-0.24
qi	0.14	-0.67	3.30	0.05
pr	0.60	0.99	10.73	-0.04
qsum	-0.44	-0.51	3.45	-0.06
qo	-0.58	0.16	0.15	-0.10
qpart	-1.64	1.62	0.92	-2.61
pheat	-0.29	-0.03	17.63	-0.15
cQinThresh	0	0	1	0
cQoutNoise	0	1	1	1
cQsumThresh	0	0	1	0
cQinOnly	0	0	1	0
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	0	1	1	1
cPreQual	0	1	1	1
cBaseQual	0	0	1	0
cDataQual	0	0	1	0
P1bs	481.98	459.98	549.63	565.47
P2bs	440.38	450.29	474.16	570.25
P1std	1.18	1.31	1.19	1.10
P2std	1.11	1.36	1.13	1.54
cQTrig	0	0	0	0
cPTrig	0	0	1	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	99999	99999

multiple scatter 1

SeriesNumber = 9903081411 EventNumber = 90278
EventTime = 3003791803
LiveTime = 2730 TimeBetween = 2941 LastISRTime = 294
VTPreTimeFast = -141 VTNearQTime = -134.6 VTNearPTime = 1355
VetoTopHigh -999999
VetoTopLow -999999
VetoCrackHigh -999999
VetoCrackLow -999999
VetoNorthHigh -999999
VetoNorthLow -999999
VetoSouthHigh -999999
VetoSouthLow -999999
VetoEastHigh -999999
VetoEastLow -999999
VetoWestHigh -999999
VetoWestLow -999999
VetoBottomHigh -999999
VetoBottomLow -999999
VetoORHigh -999999
VetoORLow -999999

quantity	103	104	105	106
pri	0.73	0.05	13.78	12.15
qi	-0.49	-0.19	3.13	3.08
pr	0.46	0.12	13.15	12.74
qsum	-0.28	-0.24	3.60	2.63
qo	0.20	-0.05	0.47	-0.45
qpart	2.42	0.60	0.74	1.34
pheat	0.08	-0.21	17.95	16.25
cQinThresh	0	0	1	1
cQoutNoise	1	1	1	1
cQsumThresh	0	0	1	1
cQinOnly	0	0	1	1
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	0	1	1	1
cPreQual	0	1	1	1
cBaseQual	0	0	1	1
cDataQual	0	0	1	1
P1bs	1173.42	1173.48	1168.58	1124.60
P2bs	1086.16	1130.70	1142.16	1142.36
P1std	1.24	1.44	1.16	1.55
P2std	1.39	1.44	0.90	1.55
cQTrig	0	0	1	0
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	99999	3479	3839

multiple scatter 2

SeriesNumber = 9904141258 EventNumber = 180759
EventTime = 3006993310
LiveTime = 2587 TimeBetween = 2828 LastISRTime = 291
VTPreTimeFast = -120 VTNearQTime = 106.8 VTNearPTime = -1139
VetoTopHigh -999999
VetoTopLow -999999
VetoCrackHigh -999999
VetoCrackLow -999999
VetoNorthHigh -999999
VetoNorthLow -999999
VetoSouthHigh -999999
VetoSouthLow -999999
VetoEastHigh -999999
VetoEastLow -999999
VetoWestHigh -999999
VetoWestLow -999999
VetoBottomHigh -999999
VetoBottomLow -999999
VetoORHigh -999999
VetoORLow -999999

quantity	103	104	105	106
pri	-0.13	31.22	-0.55	29.00
qi	0.07	6.46	0.08	9.62
pr	-0.63	26.57	-1.77	28.30
qsum	0.32	8.78	0.69	9.97
qo	0.25	2.33	0.61	0.35
qpart	-0.55	0.47	-0.77	0.93
pheat	0.02	44.13	-0.39	48.23
cQinThresh	0	1	0	1
cQoutNoise	1	0	1	1
cQsumThresh	0	1	0	1
cQinOnly	0	0	0	1
cQoutOnly	0	0	0	0
cQShare	0	1	0	0
cQinNR	1	0	1	1
cQNR	1	1	0	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	1	0	1
cDataQual	0	1	0	1
P1bs	317.10	353.58	472.75	425.61
P2bs	276.46	359.62	382.44	418.80
P1std	1.30	1.28	1.22	1.65
P2std	1.24	1.42	1.17	2.14
cQTrig	0	1	0	0
cPTrig	0	0	0	0
QPreTime	-99999	-1	-99999	-99999
QPostTime	99999	99999	99999	1
PPostTime	99999	3041	99999	3311

multiple scatter 3

```
SeriesNumber = 9904271059   EventNumber = 200939
EventTime = 3008114831
LiveTime = 242   TimeBetween = 528   LastISRTime = 22206
VTPreTimeFast = -39   VTNearQTime = 31.4   VTNearPTime = 731
VetoTopHigh        -999999
VetoTopLow         -999999
VetoCrackHigh      -999999
VetoCrackLow       -999999
VetoNorthHigh      -999999
VetoNorthLow       -999999
VetoSouthHigh      -999999
VetoSouthLow       -999999
VetoEastHigh       -999999
VetoEastLow        -999999
VetoWestHigh       -999999
VetoWestLow        -999999
VetoBottomHigh     -999999
VetoBottomLow      -999999
VetoORHigh         -999999
VetoORLow          -999999
```

quantity	103	104	105	106
pri	0.53	29.00	18.72	1.37
qi	-0.44	9.33	5.46	-0.62
pr	-0.54	28.59	18.22	0.95
qsum	0.09	9.53	5.71	-0.41
qo	0.54	0.20	0.25	0.21
qpart	-10.52	0.96	0.91	2.03
pheat	-0.36	47.65	29.64	0.14
cQinThresh	0	1	1	0
cQoutNoise	1	1	1	1
cQsumThresh	0	1	1	0
cQinOnly	0	1	1	0
cQoutOnly	0	0	0	0
cQShare	0	0	0	0
cQinNR	1	1	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	1	1	0
cDataQual	0	1	1	0
P1bs	379.52	361.99	518.08	517.65
P2bs	341.32	358.25	443.89	515.55
P1std	1.25	1.32	1.12	1.67
P2std	1.18	1.74	1.34	1.42
cQTrig	0	1	0	0
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	5	99999
PPostTime	99999	2865	3135	99999

multiple scatter 4

SeriesNumber = 9905011832 EventNumber = 150696
EventTime = 3008478565
LiveTime = 4707 TimeBetween = 4967 LastISRTime = 588
VTPreTimeFast = -42 VTNearQTime = -34 VTNearPTime = -1817
VetoTopHigh -999999
VetoTopLow -999999
VetoCrackHigh -999999
VetoCrackLow -999999
VetoNorthHigh -999999
VetoNorthLow -999999
VetoSouthHigh -999999
VetoSouthLow -999999
VetoEastHigh -999999
VetoEastLow -999999
VetoWestHigh -999999
VetoWestLow -999999
VetoBottomHigh -999999
VetoBottomLow -999999
VetoORHigh -999999
VetoORLow -999999

quantity	103	104	105	106
pri	1.58	25.85	64.76	0.89
qi	-0.72	-0.03	22.95	-0.42
pr	0.41	18.10	64.69	0.22
qsum	-0.13	3.85	22.99	-0.09
qo	0.58	3.88	0.04	0.33
qpart	9.86	-1.02	1.00	8.37
pheat	0.14	25.79	110.66	0.04
cQinThresh	0	0	1	0
cQoutNoise	0	0	1	1
cQsumThresh	0	1	1	0
cQinOnly	0	0	1	0
cQoutOnly	0	1	0	0
cQShare	0	0	0	0
cQinNR	1	0	1	1
cQNR	1	1	1	1
cGoodTime	1	1	1	1
cPreQual	1	1	1	1
cBaseQual	0	1	1	0
cDataQual	0	1	1	0
P1bs	291.20	245.46	373.66	423.04
P2bs	251.03	250.65	286.01	417.95
P1std	1.32	1.19	1.02	1.65
P2std	1.37	1.36	1.21	1.38
cQTrig	0	0	1	0
cPTrig	0	0	0	0
QPreTime	-99999	-99999	-99999	-99999
QPostTime	99999	99999	99999	99999
PPostTime	99999	3292	2523	99999

Appendix E

The DAMA Claim and Compatibility of DAMA and CDMS

E.1 Introduction

The DAMA collaboration has recently presented analysis of a four-year data set in which they claim to observe an annually modulated event rate characteristic of WIMP interactions. In this appendix, I discuss their claim in some detail and present an analysis of the compatibility of the DAMA and CDMS results. This is done within the framework of spin-independent WIMP-nucleon interactions. I comment on the possibility of spin-dependent interactions at the end. The work described here has been done by Rick Gaitksell, Bernard Sadoulet, Richard Schnee, and myself. Unless otherwise indicated, all information on the DAMA experiment is derived from their recent preprint on the four-year data set [88] or their 1996 exclusion-limit paper [146].

E.2 The DAMA Experiment and Annual-Modulation Claim

The DAMA (DArk MAtter) experiment uses high-purity NaI(Tl) scintillation-mediated detectors to search for an annually modulated event rate characteristic of WIMP interactions. While these detectors have no intrinsic ability at low energy to discriminate WIMP-induced nuclear recoils from interactions of background particles, their large mass (~ 100 kg) and low intrinsic background-event rates make it possible to achieve the statistical precision needed to look for the unique annual-modulation signature of WIMP interactions. Annual modulation is discussed in Chapter 2. While the solar system possesses a velocity of approximately 232 km s^{-1} in the direction of the galactic disk's revolution about the center of the galaxy, the WIMPs making up the galactic halo have isotropic velocities with no net revolution. The WIMPs obey a Maxwell-Boltzmann distribution with most probable velocity $v_0 = 220$ km s^{-1}. Therefore, for a terrestrial target, the WIMPs come primarily from the direction in which the solar system is moving. In the laboratory frame, the WIMP velocity distribution is modulated by the revolution of the Earth around the Sun, which adds or subtracts a sinusoidally varying velocity with amplitude 15 km s^{-1} to the 232 km s^{-1} solar velocity in the direction of the Sun's motion. This varying component causes the recoil-energy spectrum to be modulated with a period of 1 year. The modulation changes the average energy and overall normalization of the recoil-energy spectrum. In any given energy bin, the modulation

is well approximated by a sine wave. The phase is fixed by the orbit of the Earth, yielding maxima and minima on days 152.5 and 335, corresponding to June 1 or 2 and December 3 or 4 [77].

The DAMA experiment employs nine 9.70 kg NaI crystals, shaped as rectangular prisms, each with two low-radioactivity photomultiplier tubes viewing the crystal ends via light guides. The crystals are contained in individual low-radioactivity copper housings; the assembled detectors are housed in a low-radioactivity copper box that is surrounded by shielding consisting of concentric layers of lead, cadmium foil, and paraffin/polyethylene (innermost to outermost). The lead and copper stop photons and electrons; the paraffin and polyethylene moderate neutrons. Cadmium has a high neutron-capture cross section, so the cadmium foils also attenute the neutron flux. The apparatus is continuously purged with high-purity N_2 gas to prevent radon from entering. The experiment is situated in the Gran Sasso National Laboratory, at a depth of 4000 mwe. At this depth, the muon flux is very small, correspondingly reducing the muon-induced electromagnetic and neutron backgrounds.

A given detector triggers if coincident pulses are seen in its two photomultiplier tubes. This yields an energy threshold in the range 1 to 2 keV "electron-equivalent," which is proportional to the scintillation light observed. For bulk electron recoils, this energy is defined to be equal to recoil energy. Since nuclear recoils are "quenched" in NaI, the recoil energy for nuclear recoils is given by $E_R = E_{ee}/\mathrm{QF}$ where QF is the quenching factor, 0.30 for Na recoils and 0.09 for I recoils, and E_{ee} is the electron-equivalent energy, also termed "keVee."

Under the assumption of spin-independent interactions and a WIMP mass of tens of GeV c^{-2} or higher, observable WIMP interactions occur primarily with iodine. Interactions of such a WIMP with Na are suppressed in cross section because of its low A, as was discussed in Chapter 2. Kinematic considerations also reduce the average recoil energy for Na recoils, pushing them below the threshold. For low-mass WIMPs, $\lesssim$ 10 GeV c^{-2}, iodine is not as favored, but MSSM WIMPs below about 30 GeV c^{-2} are already excluded. Furthermore, the WIMP mass that best fits the DAMA annual-modulation signal is about 50 GeV c^{-2}. Therefore, interactions on iodine and QF = 0.09 are assumed in this discussion and the recoil-energy threshold is thus between about 10 and 20 keV.

The data set published in [88] consists of a 57,986 kg d exposure obtained over a four-year period. The *efficiency-corrected* "DC" background-particle spectrum observed is shown in Figure E.1. By "efficiency-corrected," I mean the spectrum that would be seen if the detectors had unity efficiency at all energies — the spectrum is divided by the efficiency function given in Figure E.2. Note that the spectrum is in keVee. The background-particle interaction rate is 1 to 2 keV^{-1} kg^{-1} d^{-1}. Conversion to recoil-energy units for nuclear recoils reduces the effective background rate by a factor $1/0.09 \approx 11$, so clearly the intrinsic background rate in these detectors is quite low, making a sensitive annual-modulation search feasible.

The "AC" signal observed in the 2-to-6-keVee energy range is shown in Figure E.3. This figure is the same as Figure 2 of [88]. I refer to the data presented in Figure E.3 as the "Figure 2" data. The AC signal is calculated as follows. The data are divided into energy bins of width 1 keVee. In each bin, the mean counting rate over the data set is determined (separate means are calculated for the first year and the remaining three years [87]). For each time bin, the difference between the observed number of counts and the number of counts expected from the mean count rate is determined. This residual is divided by the exposure for the given time bin (mass×time). The resulting count rate is corrected for efficiencies by *dividing* by the experiment's efficiency function,

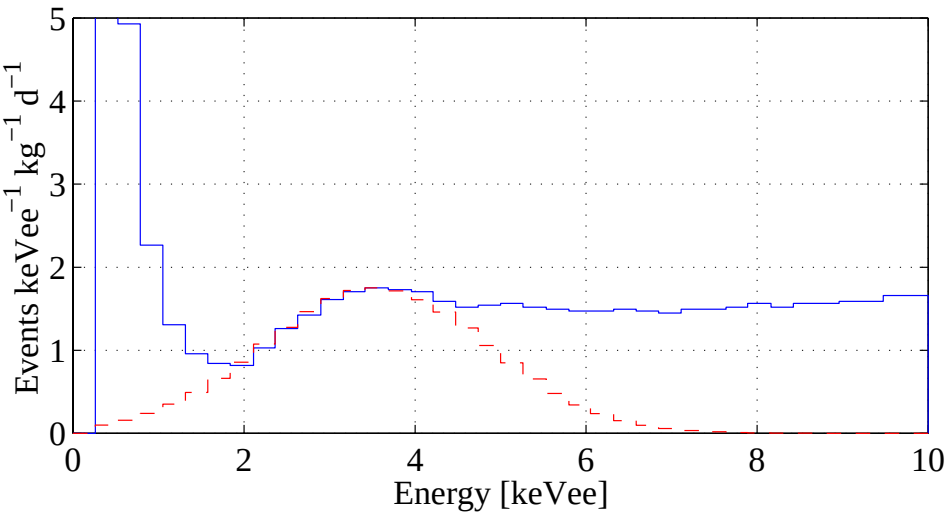

Figure E.1: Efficiency-corrected "DC" spectrum observed by the DAMA experiment with resolution function overlaid. The energy axis is in "keVee," which corresponds to recoil energy for photons but requires division by the quenching factor (0.09) to convert to recoil-energy for nuclear recoils. The energy resolution function is $\sigma_{E_{ee}} = 0.70\sqrt{E_{ee}}$ (E_{ee} is keVee) and is derived from Figure 13 of [150].

which is shown in Figure E.2. The resulting count rates for the four energy bins between 2 and 6 keVee are *averaged* to yield Figure E.3.

This interpretation of their procedure is confirmed by checking the size of their error bars. The complete DAMA data are not available, so it is not possible to calculate the expected error bars exactly. However, the modulation amplitude is very small compared to the mean count rate shown in Figure E.1 (of order 3%), and the number of counts in a given time and energy bin is large, so the statistical uncertainty on the number of counts is given accurately by $\sqrt{\langle N \rangle}$ where $\langle N \rangle$ is the number of counts expected in the given time and energy bin from the mean counting rate. The error bars calculated in this way match those given by DAMA in Figure E.3.

Clearly, there is an annual modulation of the DAMA counting rate. To determine the best fit WIMP mass and cross section and to set a confidence region in the $(M_\delta, \sigma_{\delta \mathrm{n}})$ plane, an approximate maximum-likelihood analysis is performed. The data are divided into 1-keVee energy bins and 1-day time bins and are separated by detector. The expected count rate due to both the DC background rate and possible WIMP interactions in the ith time interval for the jth detector in the kth energy bin is

$$\mu_{ijk} = \left[b_{jk} + S_{0,k} + S_{m,k} \cos\left(\frac{2\pi(t_i - 152.5)}{365.25} \right) \right] \Delta t_i \, M_j \, \Delta E \, \epsilon_{jk} \tag{E.1}$$

b_{jk} is the time-independent background-particle event rate (in keV^{-1} kg^{-1} d^{-1}) in the jth detector and kth energy bin. $S_{0,k}$ is the mean event rate in the kth energy bin expected for a WIMP with a given M_δ and $\sigma_{\delta \mathrm{n}}$. $S_{m,k}$ is the modulation amplitude of the WIMP event rate in the kth energy bin for the WIMP model. t_i is the time of the ith time bin. Δt_i is the width in time of the ith

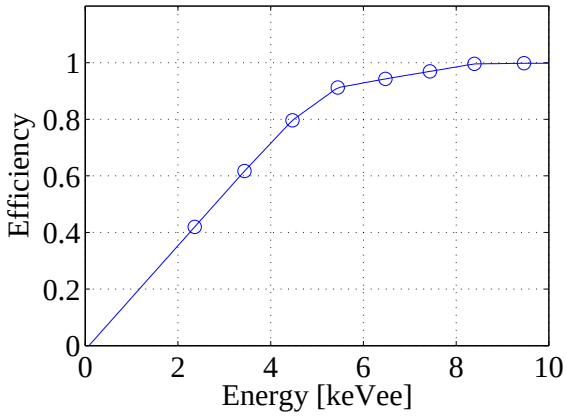

Figure E.2: DAMA efficiency function, taken from [150]. The circles indicate the data available from this reference; the line is an interpolation extrapolated to yield zero efficiency at $E_{ee} = 0$.

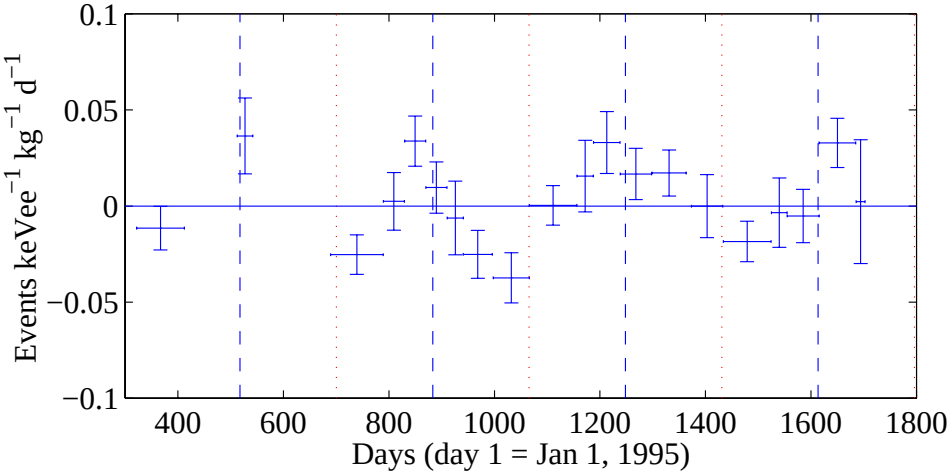

Figure E.3: Average event rate vs. time observed in the 2-to-6-keVee energy range. The horizontal axis is time in days. The error bars are DAMA's. The dashed and dotted vertical indicate the maxima and minima of the expected annual-modulation signal at days 152.5 and 335. This is Figure 2 of [88].

time bin. M_j is the mass of the jth detector. $\Delta E = 1$ keVee is the width of the energy bins. ϵ_{jk} is the efficiency of the jth detector in the kth energy bin. The cosine, of course, gives the sinusoidal modulation with the appropriate phase and period. The data are given by the ensemble of N_{ijk}, the number of counts observed in the ith time bin, in the jth detector, and in the kth energy bin. The experiment's likelihood function is given by

$$\mathcal{L}_0(N_{ijk}|b_{jk}, M_\delta, \sigma_{\delta\mathrm{n}}) = \Pi_{ijk} \frac{\mu_{ijk}^{N_{ijk}} \exp(-\mu_{ijk})}{N_{ijk}!} \tag{E.2}$$

which is just a product of Poisson distribution functions with the appropriate expected numbers of counts μ_{ijk}. The function $y = -2 \log \mathcal{L}_0$ is distributed about its minimum value as χ^2 with $\nu = 2 + N_{det} N_E$ degrees of freedom, where N_{det} is the number of detectors (the j index), N_E is the number energy bins (the k index), and an additional 2 degrees of freedom arise from the WIMP mass M_δ and cross section $\sigma_{\delta\mathrm{n}}$. During minimization of the y-function, the additional constraints $M_\delta \geq 30$ GeV c^{-2} and $b_{jk} \geq 0$ are applied. The first constraint applies present lower limits on the WIMP mass. The second constraint requires that no detector have a negative background rate, which seems a reasonable requirement.

The $(M_\delta, \sigma_{\delta\mathrm{n}})$ value that minimizes y is (52 GeV c^{-2}, 7.2×10^{-42} cm^2). The b_{jk} are projected out by evaluating them at the values $\widehat{b_{jk}}$ that minimize y for the given $(M_\delta, \sigma_{\delta\mathrm{n}})$. A 3σ CL allowed region is calculated in the standard way by finding the contour in the $(M_\delta, \sigma_{\delta\mathrm{n}})$ plane that yields $y = \widehat{y} + 9$ where $\widehat{y}$ is the value of the y-function at the minimum-y point $(\widehat{M_\delta}, \widehat{\sigma_{\delta\mathrm{n}}})$. If the constraints on M_δ and b_{jk} were not applied, this 3σ CL contour would correspond to 98.9% CL because, after projecting out the b_{jk}, y is distributed about its minimum value as χ^2 with two degrees of freedom, the WIMP mass and cross section. However, the additional constraints may invalidate this correspondence. A Monte Carlo must be performed to determine exact confidence levels. DAMA do not quote an explicit percentage confidence level. The best-fit $(M_\delta, \sigma_{\delta\mathrm{n}})$ and the 3σ CL contour are shown in Figure E.4.

It is interesting to plot the expected energy spectrum of the *modulation signal* (not the DC WIMP-interaction recoil-energy spectrum) for the best fit value of $(M_\delta, \sigma_{\delta\mathrm{n}})$. It is straightforward to calculate the annual modulation using the formalism presented in Chapter 2. It is important, however, that the DAMA energy resolution be convolved with the expected recoil-energy spectrum because the width of the resolution function is large compared to the energies of interest, as is shown in Figure E.1. The resolution function is given in Figure 13 of [150] and takes the form $\sigma_{E_{ee}} = 0.70\sqrt{E_{ee}}$ where E_{ee} is electron-equivalent energy. The expected recoil-energy spectrum of annual modulation is converted to E_{ee} and convolved with this resolution function. The result is shown in Figure E.5.

In the same preprint that presents the four-year data set, DAMA also perform a combined analysis of the annual-modulation data and their 1996 pulse-shape-analysis exclusion limit. The intent is to use the exclusion limit to better constrain the parameters of the WIMP yielding the annual-modulation signal. They do this by adding to the y-function a term

$$\Delta y = \frac{(S_{0,k} - r_k)^2}{\sigma_k^2} \theta(S_{0,k} - r_k) \tag{E.3}$$

where r_k is "the lowest counting rate in the most selective energy bin" of the data in the 1996 paper. That is, for a given M_δ and $\sigma_{\delta\mathrm{n}}$, they pick the detector and energy bin that most severely constrains

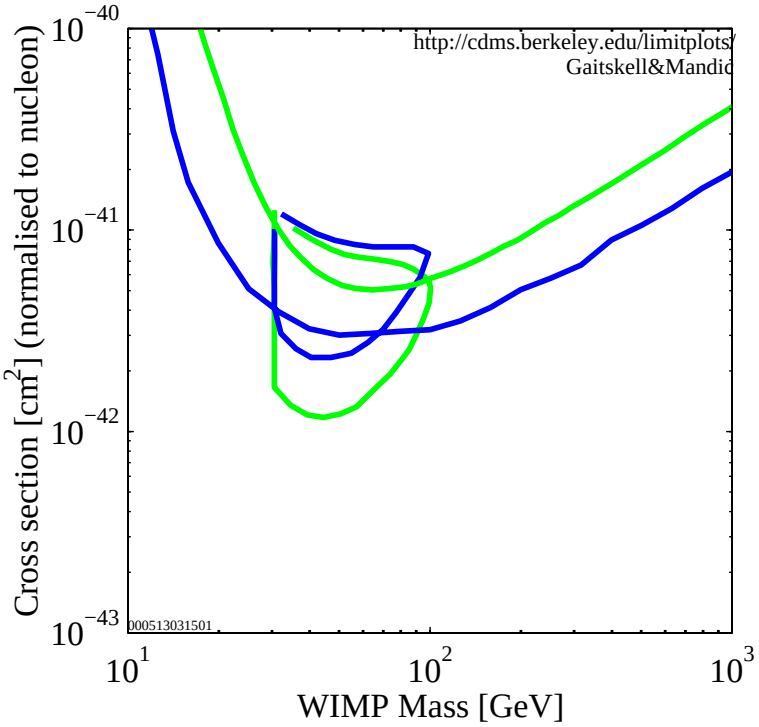

Figure E.4: DAMA allowed regions, exclusion limit, and the CDMS exclusion limit. Light solid limit line: DAMA 1996 exclusion limit. Dark solid limit line: CDMS limit. Dark heart-shaped contour: 3σ CL contour calculated from the annual-modulation analysis alone. The minimum-y point is $(M_\delta, \sigma_{\delta \mathrm{n}}) = (52\ \mathrm{GeV\ c^{-2}}, 7.2 \times 10^{-42}\ \mathrm{cm^2})$. Light heart-shaped contour: 3σ CL contour from combined analysis of the DAMA annual-modulation data and the 1996 pulse-shape-analysis exclusion limit, yielding a minimum-y point at $(M_\delta, \sigma_{\delta \mathrm{n}}) = (44\ \mathrm{GeV\ c^{-2}}, 5.4 \times 10^{-42}\ \mathrm{cm^2})$.

the cross section of the WIMP. σ_k is the uncertainty associated with r_k. The Heaviside function enforces the constraint only if the DC WIMP interaction rate exceeds r_k. The behavior of the likelihood function is presumably made more abnormal by this modifiction, so the 3σ CL contour derived from the modified y-function is certainly not a 98.9% CL confidence contour. Regardless, the analysis can be done this way, and the resulting minimum-y WIMP mass and cross section is $(M_\delta, \sigma_{\delta \mathrm{n}}) = (44\ \mathrm{GeV\ c^{-2}}, 5.4 \times 10^{-42}\ \mathrm{cm^2})$. The 3σ CL contour from this analysis is also shown in Figure E.4.

E.3 Critique of the Annual-Modulation Claim

A number of issues arise regarding the above annual-modulation claim.

First, the DC background-event spectrum has an unexpected shape. Typically, one expects a flat background due to Compton-scattering of high-energy photons with peaks due to X-rays and a possible continuum due to electrons (from various sources) superimposed. Such a flat spectrum

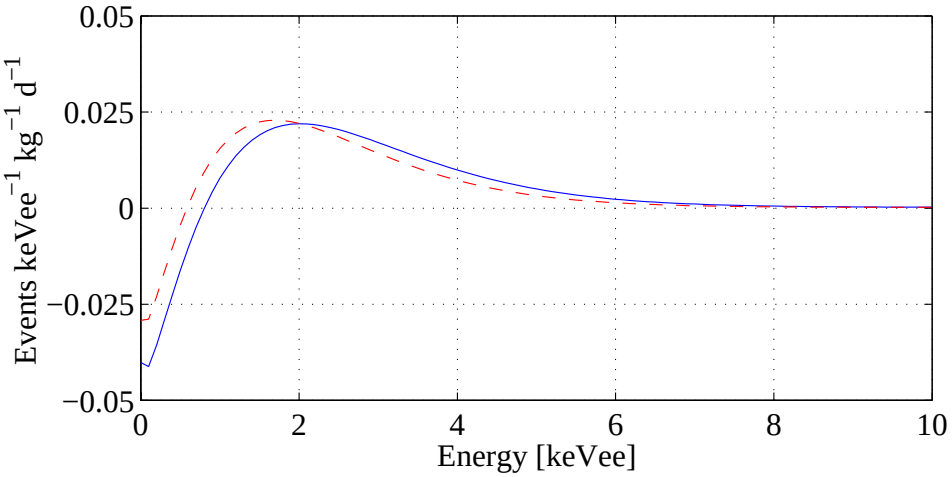

Figure E.5: Expected energy spectrum of annual-modulation signal: the difference between the energy spectrum at day 152.5 and the mean energy spectrum, giving the amplitude of the modulation as a function of energy. Note that electron-equivalent energy units are used. Solid: $(M_\delta, \sigma_{\delta \mathrm{n}}) = (52\ \mathrm{GeV\ c^{-2}}, 7.2 \times 10^{-42}\ \mathrm{cm^2})$, best fit to annual-modulation data alone. Dashed: $(M_\delta, \sigma_{\delta \mathrm{n}}) = (44\ \mathrm{GeV\ c^{-2}}, 5.4 \times 10^{-42}\ \mathrm{cm^2})$, best fit to combined annual-modulation and 1996 exclusion-limit analysis. The energy resolution function has been convolved with the expected spectra.

is seen in calibration data taken with a ^{137}Cs source, which emits 662-keV photons (see Figure 11 of [150]). Multiple-scatter data are included in this calibration spectrum, but it is not expected that the single-scatter data shown in Figure E.1 should exhibit a significantly different spectrum. The efficiency-corrected background spectrum, Figure E.1, displays a flat counting rate above 5 keVee, a bump at 3.6 keVee that presumably is an X-ray peak, and then a falloff below 3.6 keVee that has the shape of the 3.6-keVee resolution function. No candidate for this peak is given. However, regardless of the source of the peak, the falloff below 3.6 keVee suggests a sharp dropoff in the true energy spectrum (before smearing by the detector-resolution function) at 3.6 keVee, which would be unphysical. The dropoff may indicate that the efficiency function shown in Figure E.2 is incorrect below 3.6 keVee. As is seen in Figure E.5, the majority of the annual-modulation signal is in the 2-to-4-keVee energy range, so such an error could be significant.

Second, the minimum-y WIMP mass and cross section derived by the above analyses do not appear to fit the "Figure 2" data reasonably. Figure E.6 shows the Figure 2 data with the expected annual-modulation signal for $(M_\delta, \sigma_{\delta \mathrm{n}}) = (52\ \mathrm{GeV\ c^{-2}}, 7.2 \times 10^{-42}\ \mathrm{cm^2})$, which has amplitude $A = 0.0111\ \mathrm{keV^{-1}\ kg^{-1}\ d^{-1}}$. Recall that the Figure 2 data is the modulated counting rate averaged between 2 and 6 keVee. Figure E.6 also shows the expected signal for $(M_\delta, \sigma_{\delta \mathrm{n}}) = (44\ \mathrm{GeV\ c^{-2}}, 5.4 \times 10^{-42}\ \mathrm{cm^2})$, which has $A = 0.0038\ \mathrm{keV^{-1}\ kg^{-1}\ d^{-1}}$. It is clear that neither of of these amplitudes fits the Figure 2 data well. The best-fit amplitude for the Figure 2 data is is $A = 0.0221\ \mathrm{keV^{-1}\ kg^{-1}\ d^{-1}}$ and is shown in Figure E.6. Formally, the χ^2 of the two annual-modulation amplitudes $A = 0.0111\ \mathrm{keV^{-1}\ kg^{-1}\ d^{-1}}$ and $A = 0.0038\ \mathrm{keV^{-1}\ kg^{-1}\ d^{-1}}$ are 29.0 and 31.6, respectively. There are 20 data points, but, the modulation amplitude and mean count

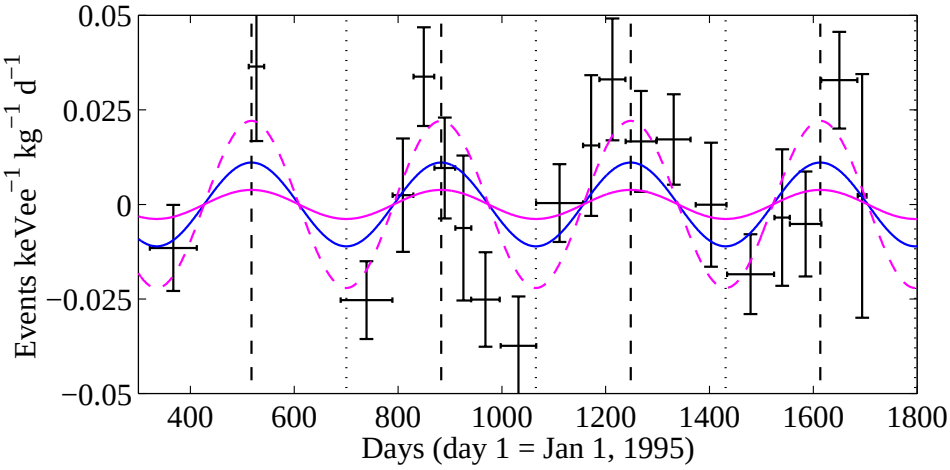

Figure E.6: DAMA modulated event rate vs. time with expected annual-modulation signals superimposed. The data are the same as Figure E.3, though the vertical axis limits have been changed. Light dashed (top): modulation amplitude that best fits these data alone, $A = 0.0221$ keV^{-1} kg^{-1} d^{-1}. Dark solid (middle): modulation amplitude corresponding to $(M_\delta, \sigma_{\delta\mathrm{n}}) = (52\ \mathrm{GeV\ c^{-2}}, 7.2 \times 10^{-42}\ \mathrm{cm^2})$, the minimum-$y$ WIMP mass and cross section for the analysis of the annual-modulation data alone, with $A = 0.0111$ keV^{-1} kg^{-1} d^{-1}. Light solid (bottom): modulation amplitude corresponding to $(M_\delta, \sigma_{\delta\mathrm{n}}) = (44\ \mathrm{GeV\ c^{-2}}, 5.4 \times 10^{-42}\ \mathrm{cm^2})$, the minimum-$y$ WIMP mass and cross section from combined analysis of the annual-modulation data and the 1996 pulse-shape-analysis exclusion limit, with $A = 0.0038$ keV^{-1} kg^{-1} d^{-1}.

rate are nominally determined from the data. Two separate means are calculated, one for the first year and a second for the remaining three years [87]. Thus, there are 17 degrees of freedom. The resulting χ^2 percentiles are 96.6% and 98.3%. This level of disagreement is troubling. Furthermore, in the paper, the χ^2 percentiles are calculated erroneously by assuming 20 degrees of freedom.

Presumably, the discrepancy between the annual-modulation amplitude expected from the minimum-y analyses and the simple fit to the Figure 2 data arises from the $b_{jk} \geq 0$ constraint: a WIMP with cross section that fits the Figure 2 data well yields DC WIMP event rates larger than the observed total event rates. This reinforces the criticism of the shape of DC background spectrum made above — it drops too quickly below 3 keV, which is presumably where the constraint has the largest effect.

Third, there appears to be a problem with the energy spectrum of the annual modulation. Though no energy spectrum of the modulated signal is available in [88], it is noted that, if the 5–6-keV bin is discarded in calculating the Figure 2 data, the χ^2 of the Figure 2 data with respect to the expected annual-modulation signals improve to 11.5 and 12.3 for the analyses with and without the 1996 limit, respectively. The corresponding χ^2 percentiles, assuming 17 degrees of freedom, are 17.1% and 21.8%, respectively. Though a plot analogous to Figure E.3 without the 5–6-keV data is not given, one can easily infer that the modulation amplitude of the data must decrease significantly when the 5–6-keV bin is discarded in order to reduce the χ^2 by these large amounts. This is disturbing — a significant fraction of the annual-modulation signal appears to be in the

5–6-keV bin, where Figure E.5 indicates there should be negligible signal.

E.4 Compatibility Analysis of the DAMA and CDMS Results

Though there are a number of reasons to doubt the DAMA claim, it is also useful to perform an analysis of the compatibility of the CDMS and DAMA results, taking the latter at face value. It is intuitively clear from Figure E.4 that the CDMS data are not compatible with the DAMA annual-modulation claim. Though the CDMS data do not exclude the entire DAMA contour at 90% CL, it must be remembered that the DAMA contour is a 3σ CL contour, corresponding asymptotically to a 98.9% CL allowed region. The two experiments are clearly out on the tails of each other's likelihood functions. In this section, the compatibility of the two experiments is evaluated in both heuristic and rigorous ways.

In making the comparison, it is important to compare to the appropriate version of the DAMA result. The CDMS data are fully compatible with the DAMA 1996 exclusion limit, as is clear from Figure E.4. There is no need to test this. Therefore, the comparison should be made with the DAMA annual-modulation analysis alone. Clearly, comparison with the combined DAMA analysis would yield a higher level of compatibility, but this is because the inclusion of the 1996 exclusion limit dilutes the test.

Heuristically, the compatibility of the two data sets can be checked by simply plotting the recoil-energy spectrum expected for various points in the DAMA 3σ CL allowed region over the CDMS recoil-energy spectrum. This is done in Figure E.7, correcting the WIMP spectra for the CDMS efficiency function. The left plot shows the spectrum expected for the DAMA minimum-y point at (52 GeV $\mathrm{c}^{-2}, 7.2 \times 10^{-42}$ cm^2). Clearly, it is very inconsistent, and, as indicated Figure E.4, it is ruled at $> 90\%$ CL. The second spectrum shown in the left plot of Figure E.7 corresponds to (40 GeV $\mathrm{c}^{-2}, 2.3 \times 10^{-42}$ cm^2), the lowest cross-section point of the DAMA 3σ CL region. It is reasonably consistent with the observed spectrum.

However, comparison to the CDMS recoil-energy spectrum does not account for the neutron component of CDMS nuclear recoils. The right plot of Figure E.7 shows a "background-subtracted" spectrum, where the neutron background has been normalized to the value that best fits the CDMS data, $n = 17.5$ events (see Section 8.3), and subtracted from the observed spectrum. This background-subtracted spectrum is fairly inconsistent with even the (40 GeV $\mathrm{c}^{-2}, 2.3 \times 10^{-42}$ cm^2) point, though not at 90% CL. CDMS excludes this point at only 75%. Finally, for completeness, the right plot of Figure E.7 also shows a WIMP spectrum correponding to the lowest cross-section point of the DAMA combined analysis 3σ CL contour, at about (40 GeV $\mathrm{c}^{-2}, 1.2 \times 10^{-42}$ cm^2). This point is consistent with the CDMS data. This is not surprising, given that this point is well below the CDMS 90% CL exclusion limit.

The compatibility of the CDMS and DAMA results can be tested rigorously by a likelihood-ratio test. This is similar the likelihood-ratio tests done in Chapter 8. The two hypotheses are that the CDMS and DAMA data are derived from the same WIMP mass and cross section or from different WIMP masses and cross sections. The likelihood ratio is

$$R = \frac{\mathcal{L}_C(X_C|\widehat{M_\delta}, \widehat{\sigma_{\delta \mathrm{n}}})\mathcal{L}_D(X_D|\widehat{M_\delta}, \widehat{\sigma_{\delta \mathrm{n}}})}{\mathcal{L}_C(X_C|\widehat{M_{\delta C}}, \widehat{\sigma_{\delta \mathrm{n} C}})\mathcal{L}_D(X_D|\widehat{M_{\delta D}}, \widehat{\sigma_{\delta \mathrm{n} D}})} \tag{E.4}$$

where $\mathcal{L}_C$ is the CDMS likelihood function evaluated for the CDMS data X_C and $\mathcal{L}_D$ is the DAMA

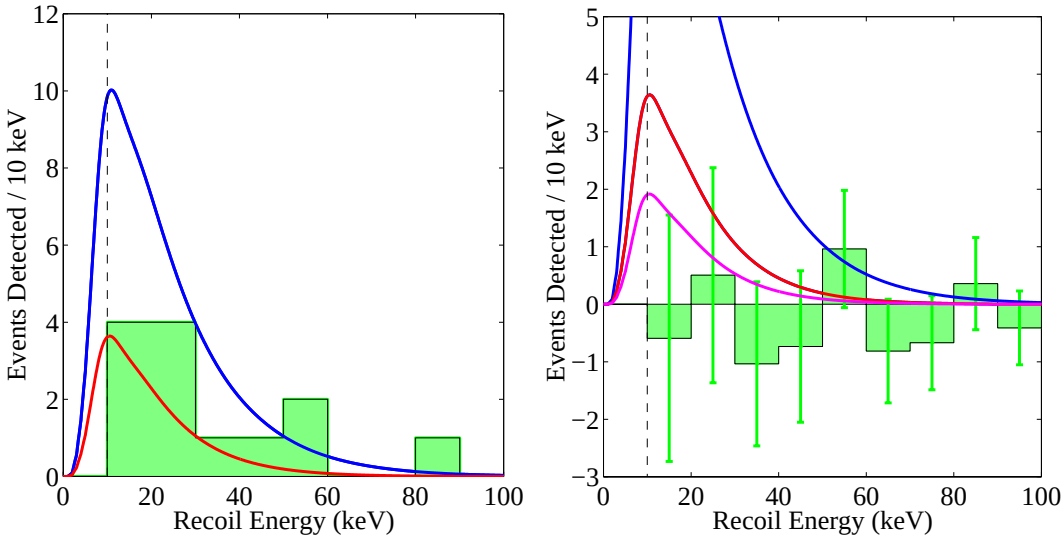

Figure E.7: Comparison of CDMS nuclear-recoil energy spectrum with that expected for the WIMP model that best fits the DAMA data. Left: comparison to full CDMS nuclear-recoil spectrum (13 events). Top curve: DAMA annual-modulation analysis minimum-y point at (52 GeV c^{-2}, 7.2×10^{-42} cm^2). Bottom curve: lowest cross-section point of DAMA annual-modulation 3σ CL allowed region. Right: "background-subtracted" CDMS nuclear-recoil energy spectrum. Top two curves are the same as in the left plot. The bottom curve corresponds to (40 GeV c^{-2}, 1.2×10^{-42} cm^2), the lowest cross-section of the 3σ CL contour of the DAMA combined analysis. Figures provided by R. Schnee.

likelihood function evaluated for the DAMA data X_D. The numerator evaluates the combined likelihood function at the WIMP mass and cross section $(\widehat{M_\delta}, \widehat{\sigma_{\delta \mathrm{n}}})$ that maximizes the combined likelihood. The denominator evaluates the likelihood functions at their separate maximum-likelihood masses and cross sections $(\widehat{M_{\delta C}}, \widehat{\sigma_{\delta \mathrm{n} C}})$ and $(\widehat{M_{\delta D}}, \widehat{\sigma_{\delta \mathrm{n} D}})$ and takes the product. R is bounded above by 1. The larger R is, the more compatible the two experiments are. In the asymptotic limit, where both likelihood functions are products of Gaussian probability functions, $-2 \log R$ is distributed as χ^2 for two degrees of freedom. The asymptotic limit does not hold here, especially for the CDMS data, but Monte Carlo calculation of the expected distribution of R is computationally intractable because both experiments must be simulated over a large area in the plane of WIMP mass and cross section. Therefore, the asymptotic limit is used an indication of the level of compatibility.

The full DAMA likelihood function is not available. However, it is straightforward to calculate a likelihood function using the Figure 2 data alone — the error bars are known. This disregards the $b_{jk} \geq 0$ constraint that forces the cross section down, but, as noted above, the large effect of this constraint is suspicious. Contours for the DAMA likelihood function calculated in this way are shown in Figure E.8. Clearly, the likelihood function moves higher in cross section because the $b_{jk} \geq 0$ constraint is not applied. The contours are not well localized in mass because no spectral information is available from the Figure 2 data.

A contour plot for the combined likelihood function (the numerator of R) is shown in

Figure E.9. The combined likelihood is maximized at (42 GeV $c^{-2}, 4.9 \times 10^{-42}$ cm^2). In the asymptotic limit, the CDMS data and the DAMA Figure 2 data are incompatible at $> 99.98\%$ CL: only 0.02% of experiments would have found lower values of R than the two data sets give. The actual confidence level should not be taken too seriously because the conditions for application of the asymptotic limit are not met. But, clearly the two experiments are very incompatible. This can be seen graphically with two plots: a comparison of the $(\widehat{M_\delta}, \widehat{\sigma_{\delta n}})$ annual-modulation amplitude to the DAMA Figure 2 data and the $(\widehat{M_\delta}, \widehat{\sigma_{\delta n}})$ recoil-energy spectrum to the background-subtracted CDMS data, shown in Figures E.10 and E.11. Clearly, the WIMP parameters that maximize the joint likelihood fit each experiment quite badly.

The incompatibility result is fairly robust. To be more generous, the likelihood function calculated by the above method can be scaled down in cross section so it approximately overlays the published DAMA 3σ CL contour (annual-modulation only), and this function used as an approximation to the true DAMA likelihood function (*i.e.*, including the $b_{jk} \geq 0$ constraint). When this is done, the new joint likelihood is maximized at (42 GeV $c^{-2}, 4.2 \times 10^{-42}$ cm^2). The incompatibility confidence level drops to 99.8% CL. Also, a reasonable criticism of the CDMS analysis is that too many Ge double-scatter nuclear recoils are seen as compared to the simulation. Discarding these events makes the CDMS data fully consistent with the simulation but also weakens the CDMS exclusion limit because the predicted neutron background is decreased. However, even after both discarding the multiples and shifting the likelihood contours down, the two experiments are still incompatible at 96.8% CL.

E.5 On Spin-Dependent WIMP Interactions

A final note — it may be possible to reconcile the CDMS and DAMA results if the signal seen by DAMA is due to WIMP scattering via spin-dependent interactions. Iodine and sodium both are mono-isotopic and odd-p nuclei, while the natural germanium used in the CDMS detectors has only a 7% abundance of the odd-n isotope ^{73}Ge (see Chapter 2 for a brief discussion of spin-dependent interactions). Thus, DAMA certainly would be more sensitive to a WIMP interacting primarily through spin-dependent interactions. It is difficult to compare the CDMS and DAMA data precisely for the spin-dependent case because nuclear-physics modeling is needed to relate the Na and I cross sections to a Ge cross section. Spin-dependent interactions may even render the DAMA data more internally consistent by modifying the annual-modulation spectrum relative to the DC WIMP-interaction spectrum. However, the spin-dependent cross section implied by the DAMA data is approximately a factor of 50 larger than is expected for MSSM WIMPs [151].

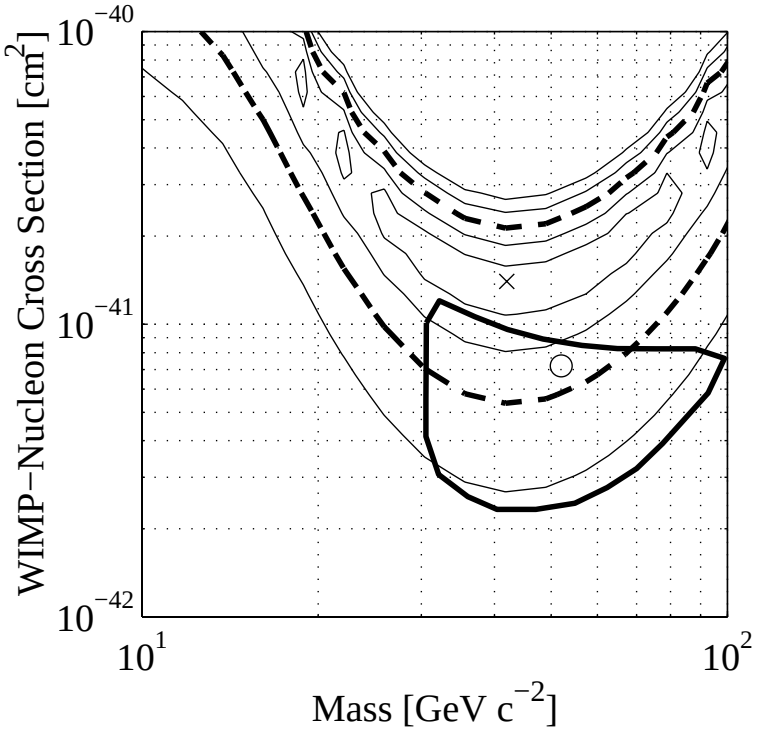

Figure E.8: DAMA log likelihood function calculated from Figure 2 data (*i.e.*, no spectral information, $b_{jk} \geq 0$ constraint not applied). The contours indicate $\Delta \log \mathcal{L}_D = -0.5, -2, -4.5, -8,$ and -12.5, corresponding to 1σ, 2σ, 3σ, 4σ, and 5σ contours. (The lower portion of the 5σ contour is not visible on the plot.) The thick dashed contour is $\Delta \log \mathcal{L}_D = -4.5$. The maximum likelihood point is (42 GeV c^{-2}, 1.4×10^{-41} cm^2), marked with a $\times$. The minimum-y point and 3σ CL contour determined by the DAMA annual-modulation analysis (without inclusion of the 1996 exclusion limit) are shown by the circle and the thick solid contour. The likelihood contours are not localized in WIMP mass because no spectral information is available from DAMA's Figure 2; the contours' shapes reflect the sensitivity of the 2-to-6-keVee bin as a function of M_δ.

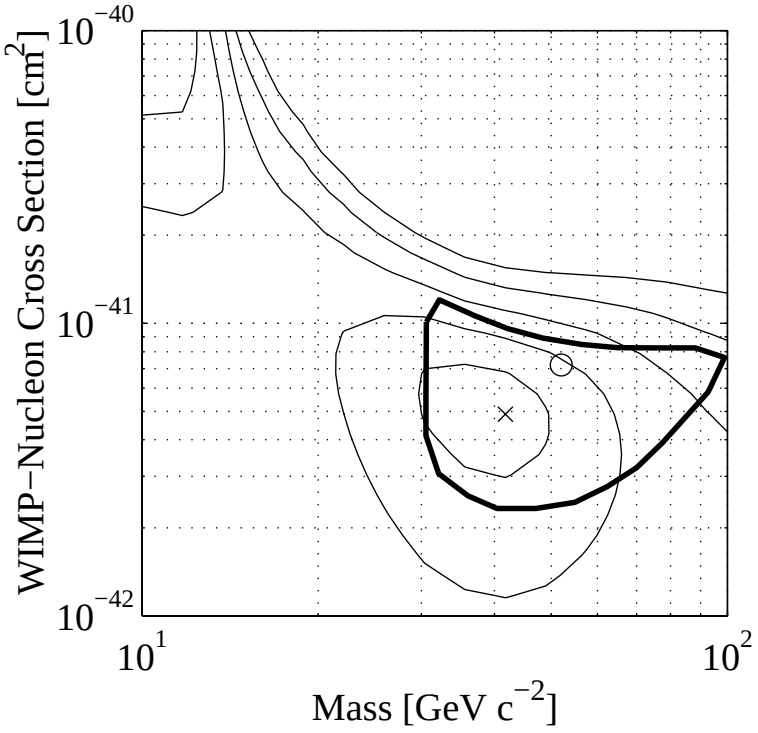

Figure E.9: Combined CDMS/DAMA log likelihood function $\log(\mathcal{L}_C\mathcal{L}_D)$ in the $(M_\delta, \sigma_{\delta\mathrm{n}})$ plane. The contours indicate $\Delta\log(\mathcal{L}_C\mathcal{L}_D) = -0.5, -2, -4.5, -8,$ and -12.5, corresponding to 1σ, 2σ, 3σ, 4σ, and 5σ contours (note that these contours have nothing to do with the compatibility CL!). The $\times$ indicates the joint maximum-likelihood $(\widehat{M_\delta}, \widehat{\sigma_{\delta\mathrm{n}}})$ at (42 GeV c^{-2}, 4.9×10^{-42} cm^2). The minimum-y point and 3σ CL contour determined from by the DAMA annual-modulation analysis (without inclusion of the 1996 exclusion limit) are shown by the circle and the thick solid contour.

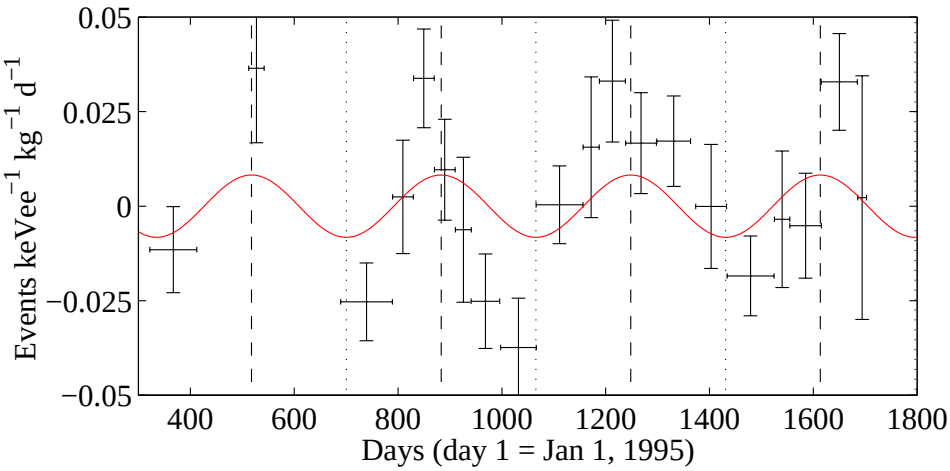

Figure E.10: Comparison of DAMA Figure 2 data to the annual-modulation signal expected for the WIMP that maximizes the joint CDMS/DAMA likelihood. The annual-modulation amplitude $A = 0.0082$ keV^{-1} kg^{-1} d^{-1} is shown for a WIMP with mass and cross section $(\widehat{M_\delta}, \widehat{\sigma_{\delta n}}) = (42 \text{ GeV c}^{-2}, 4.9 \times 10^{-42} \text{ cm}^2)$.

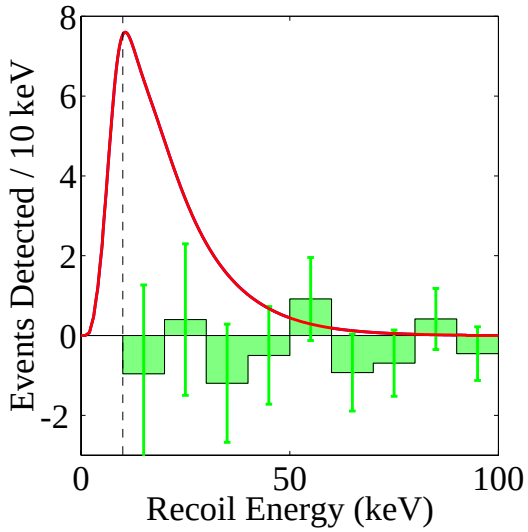

Figure E.11: Comparison of CDMS recoil-energy spectrum to that expected for the WIMP that maximizes the joint CDMS/DAMA likelihood. The background-subtracted CDMS spectrum is shown with the spectrum expected for a WIMP with mass and cross section $(\widehat{M_\delta}, \widehat{\sigma_{\delta n}}) = (42 \text{ GeV c}^{-2}, 4.9 \times 10^{-42} \text{ cm}^2)$. Figure provided by R. Schnee.

Bibliography

[1] P. J. E. Peebles, *Principles of Physical Cosmology* (Princeton University Press, Princeton, NJ, 1993).

[2] E. W. Kolb and M. S. Turner, *The Early Universe* (Addison-Wesley, Reading, MA, 1990).

[3] M. S. Turner and J. A. Tyson, Rev. Mod. Phys. **71**, 145 (1999).

[4] P. Salucci and M. Persic, in *ASP Conf. Ser. 117: Dark and Visible Matter in Galaxies and Cosmological Implications* (Brigham Young University, Provo, UT, 1997), pp. 1–27.

[5] D. S. Davis and R. E. White, Astrophys. J. Lett. **470**, L35 (1996).

[6] M. Loewenstein and R. E. White, Astrophys. J. **518**, 50 (1999).

[7] M. Persic, P. Salucci, and F. Stel, Mon. Not. Roy. Astron. Soc. **281**, 27 (1996).

[8] A. H. Sonnenschein, Ph.D. thesis, University of California, Santa Barbara, 1999.

[9] J. Binney and S. Tremaine, *Galactic Dynamics* (Princeton University Press, Princeton, NJ, 1987).

[10] M. Girardi *et al.*, Astrophys. J. **505**, 74 (1998).

[11] M. Girardi *et al.*, Astrophys. J. **530**, 62 (2000).

[12] D. A. White, C. Jones, and W. Forman, Mon. Not. Roy. Astron. Soc. **292**, 419 (1997).

[13] M. Bartelmann and P. Schneider, astro-ph/9912508, to appear in Phys. Rep.

[14] J. E. Carlstrom *et al.*, astro-ph/9905255.

[15] B. S. Mason and S. T. Myers, astro-ph/9910438.

[16] T. Padmanabhan, *Structure Formation in the Universe* (Cambridge University Press, Cambridge, 1993).

[17] E. Hivon, private communication.

[18] E. Torbet *et al.*, Astrophys. J. Lett. **521**, L79 (1999).

[19] A. D. Miller *et al.*, Astrophys. J. Lett. **524**, L1 (1999).

[20] P. D. Mauskopf *et al.*, astro-ph/9911444.

[21] K. Coble *et al.*, Astrophys. J. Lett. **519**, L5 (1999).

[22] J. B. Peterson *et al.*, Astrophys. J. Lett. **532**, L83 (2000).

[23] S. Dodelson and L. Knox, Phys. Rev. Lett. **84**, 3523 (2000).

[24] P. de Bernardis *et al.*, Nature **404**, 955 (2000).

[25] S. Hanany *et al.*, astro-ph/0005123.

[26] L. Knox and L. Page, astro-ph/0002162.

[27] M. Harwit, *Astrophysical Concepts* (Springer-Verlag, New York, 1988).

[28] M. Hamuy *et al.*, Astron. J. **112**, 2391 (1996).

[29] S. Perlmutter *et al.*, in *Thermonuclear Supernovae, NATO ASI* (Kluwer Academic Publishers, Dordrecht, 1997), p. 749.

[30] A. G. Riess, W. H. Press, and R. P. Kirshner, Astrophys. J. **473**, 88 (1996).

[31] S. Perlmutter *et al.*, Astrophys. J. **517**, 565 (1999).

[32] A. G. Riess *et al.*, Astron. J. **116**, 1009 (1998).

[33] A. Aguirre, Astrophys. J. **525**, 583 (1999).

[34] N. A. Bahcall, X. Fan, and R. Cen, Astrophys. J. Lett. **485**, L53 (1997).

[35] A. Dekel, astro-ph/9911501.

[36] D. Tytler, J. M. O'Meara, N. Suzuki, and D. Lubin, astro-ph/0001318, to appear in Phys. Scr.

[37] E. Gawiser and J. Silk, Science **280**, 1405 (1998).

[38] D. H. Weinberg, J. Miralda-Escudé, L. Hernquist, and N. Katz, Astrophys. J. **490**, 564 (1997).

[39] C. Flynn, A. Gould, and J. N. Bahcall, Astrophys. J. Lett. **466**, L55 (1996).

[40] C. Alcock *et al.*, astro-ph/0001272.

[41] B. D. Fields, K. Freese, and D. S. Graff, New Astron. **3**, 347 (1998).

[42] D. S. Graff, K. Freese, T. P. Walker, and M. H. Pinsonneault, Astrophys. J. Lett. **523**, L77 (1999).

[43] B. D. Fields, K. Freese, and D. S. Graff, Astrophys. J. **534**, 265 (2000).

[44] R. Cen and J. P. Ostriker, Astrophys. J. **514**, 1 (1999).

[45] B. J. Carr and J. H. MacGibbon, Phys. Rep. **307**, 141 (1998).

[46] K. Jedamzik, Phys. Rep. **307**, 155 (1998).

[47] C. Hagmann *et al.*, Phys. Rev. Lett. **80**, 2043 (1998).

[48] C. Athanassopoulos *et al.*, Phys. Rev. Lett. **81**, 1774 (1998).

[49] K. Eitel, New J. Phys. **2**, 1 (2000).

[50] A. O. Bazarko, hep-ex/9906003.

[51] M. Drees, hep-ph/9911409.

[52] F. Wilczek, hep-ph/0002045.

[53] L. Girardello and M. Grisaru, Nucl. Phys. B **194**, 65 (1982).

[54] G. Jungman, M. Kamionkowski, and K. Griest, Phys. Rep. **267**, 195 (1996).

[55] L. Bergström, P. Ullio, and J. H. Buckley, Astropart. Phys. **9**, 137 (1998).

[56] D. D. Dixon *et al.*, New Astron. **3**, 539 (1998).

[57] P. Gondolo, astro-ph/9807347.

[58] J. Ellis, T. Falk, G. Ganis, and K. A. Olive, hep-ph/0004169.

[59] L. Bergström, J. Edsjö, and P. Gondolo, Phys. Rev. D **58**, 103519/1 (1998).

[60] J. M. Losecco *et al.*, Phys. Lett. B **188**, 388 (1987).

[61] M. M. Boliev *et al.*, Nucl. Phys. B Proc. Supp. **48**, 83 (1996).

[62] N. Sato *et al.*, Phys. Rev. D **44**, 2220 (1991).

[63] M. Mori *et al.*, Phys. Lett. B **270**, 89 (1991).

[64] M. Mori *et al.*, Phys. Lett. B **289**, 463 (1992).

[65] M. Ambrosio *et al.*, astro-ph/0002492.

[66] P. Gondolo and J. Silk, Phys. Rev. Lett. **83**, 1719 (1999).

[67] A. M. Ghez, B. L. Klein, M. Morris, and E. E. Becklin, Astrophys. J. **509**, 678 (1998).

[68] R. Genzel, A. Eckart, T. Ott, and F. Eisenhauer, Mon. Not. Roy. Astron. Soc. **291**, 219 (1997).

[69] P. Gondolo, hep-ph/0002226.

[70] S. Orito *et al.*, Phys. Rev. Lett. **84**, 1078 (2000).

[71] L. Bergström, J. Edsjö, and P. Ullio, Astrophys. J. **526**, 215 (1999).

[72] S. Coutu *et al.*, Astropart. Phys. **11**, 429 (1999).

[73] S. W. Barwick *et al.*, Phys. Rev. Lett. **75**, 390 (1995).

[74] M. W. Goodman and E. Witten, Phys. Rev. D **31**, 3059 (1985).

[75] A. Drukier and L. Stodolsky, Phys. Rev. D **30**, 2295 (1984).

[76] H.-Y. Cheng, hep-ph/0002157.

[77] J. D. Lewin and P. F. Smith, Astropart. Phys. **6**, 87 (1996).

[78] J. Edsjö and P. Gondolo, Phys. Rev. D **56**, 1879 (1997).

[79] A. Corsetti and P. Nath, hep-ph/0003186.

[80] R. J. Gaitskell and V. Mandic, Dark Matter Tools Working Group limit plotter, http://dmtools.berkeley.edu/limitplots.

[81] D. O. Caldwell *et al.*, Phys. Rev. Lett. **61**, 510 (1988).

[82] R. J. Gaitskell *et al.*, Nucl. Phys. B Proc. Supp. **51**, 279 (1996).

[83] G. J. Feldman and R. Cousins, Phys. Rev. D **57**, 3873 (1998).

[84] L. Baudis *et al.*, Phys. Rev. D **59**, 022001/1 (1999).

[85] S. Ahlen *et al.*, Phys. Lett. B **195**, 603 (1987).

[86] A. Morales *et al.*, hep-ex/0002053, submitted to Phys. Lett. B.

[87] R. J. Gaitskell, private communication.

[88] R. Bernabei *et al.*, INFN/AE-00/01.

[89] N. J. C. Spooner, Phys. Rep. **307**, 253 (1998).

[90] D. Snowden-Ifft, C. Martoff, and J. Burwell, Phys. Rev. D **61**, 101301/1 (2000).

[91] A. Da Silva, Ph.D. thesis, University of British Columbia, 1996.

[92] P. D. Barnes, Jr., Ph.D. thesis, University of California, Berkeley, 1996.

[93] T. A. Perera, Ph.D. thesis, Case Western Reserve University, in preparation.

[94] R. P. Kokoulin and A. A. Petrukhin, Acta Phys. Acad. Sc. Hung. **29**, 277 (1970).

[95] T. Murota, A. Ueda, and H. Tanaka, Prog. Theor. Phys. **16**, 482 (1956).

[96] T. A. Perera, private communication.

[97] E. Segré, *Nuclei and Particles: An Introduction to Nuclear and Subnuclear Physics* (W. A. Benjamin, Reading, MA, 1980).

[98] F. Khalchukov *et al.*, Il Nuovo Cim. **18C**, 517 (1995).

[99] J. C. Barton, in *19th Intern. Cosmic Ray Conf.* (Scientific and Technical Information Branch, NASA, Washington, D.C., 1985), Vol. 8, pp. 98–101.

[100] T. Florkowski, L. Morawska, and K. Rozanski, Nucl. Geophys. **2**, 1 (1988).

[101] F. Khalchukov, A. Mal'gin, V. Ryassny, and O. Ryazhskaya, Il Nuovo Cim. **6C**, 320 (1983).

[102] R. M. Clarke, Ph.D. thesis, Stanford University, 1999.

[103] R. Abusaidi *et al.*, in preparation.

[104] CDMS Collaboration, "Cold Hardware Review", CDMSNote/0004002.

[105] N. Wang, Ph.D. thesis, University of California, Berkeley, 1991.

[106] T. A. Shutt, Ph.D. thesis, University of California, Berkeley, 1993.

[107] A. K. Cummings, Ph.D. thesis, University of California, Berkeley, 1997.

[108] W. K. Stockwell, Ph.D. thesis, University of California, Berkeley, 1996.

[109] T. Shutt, "Selected Plots from the CDMS Presentation at the December 1996 Texas Astrophysics Meeting", CDMS internal note, Jan. 1997.

[110] R. J. Gaitskell, CDMSNote/9703002.

[111] S. R. Golwala, CDMSNote/9708003.

[112] D. S. Akerib *et al.*, Nucl. Phys. B Proc. Supp. **70**, 64 (1999).

[113] F. C. Wellstood, Ph.D. thesis, University of California, Berkeley, 1988.

[114] B. Sadoulet *et al.*, Physica B **219**, 741 (1996).

[115] InterFET Corp., "Junction Field Effect Transistors", device IF4500, process NJ450L.

[116] SUF Detector Log, Vol. 10, p. 88.

[117] R. J. Gaitskell, CDMSNote/9711001.

[118] BLIPs 1–2 UCB Test Facility Log, Vol. 3, p. 7.

[119] B. Cabrera, "Electron-Phonon Scattering", CDMS internal note, Jan. 1993.

[120] BLIPs 3–6 UCB Test Facility Log, Vol. 2, pp. 74–98.

[121] T. Shutt *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **444**, 340 (2000).

[122] B. S. Neganov and V. N. Trofimov, JETP Lett. **28**, 328 (1978).

[123] P. N. Luke, J. Appl. Phys. **64**, 6858 (1988).

[124] T. Shutt *et al.*, in *Proceedings of the Seventh International Workshop on Low Temperature Detectors* (Max Planck Institute of Physics, Munich, 1997), pp. 224–226.

[125] R. Diaz, Undergraduate honors thesis, University of California, Berkeley, 1995.

[126] T. Shutt, "Ionization contact studies update", electronic mail, May 21, 1998.

[127] P. N. Luke, private communication.

[128] S. M. Sze, *Physics of Semiconductor Devices* (John Wiley and Sons, New York, 1981).

[129] Edwards High Vacuum International, Apiezon Wax W40.

[130] J. Emes, private communication.

[131] SUF Detector Log, Vol. 10, pp. 45–84.

[132] T. A. Perera and D. S. Akerib, CDMSNote/9808001.

[133] CDMS electronics web site, http://cdms.berkeley.edu/cdms_restricted/electronics.

[134] W. D. Greason, *Electrostatic Damage in Electronics: Devices and Systems* (John Wiley and Sons, New York, 1987).

[135] D. S. Akerib *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **400**, 181 (1997).

[136] J. J. Bock, Ph.D. thesis, University of California, Berkeley, 1994.

[137] S. White, "Request for Quote for Fabrication of CDMS Striplines", submitted to Basic Electronics, Inc.

[138] D. Yvon *et al.*, Nucl. Instrum. Methods Phys. Res., Sect. A **368**, 778 (1996).

[139] P. Horowitz and W. Hill, *The Art of Electronics* (Cambridge University Press, Cambridge, 1989).

[140] S. R. Golwala, "Anti-alias Filter Calculations", CDMS electronics web site, http://cdms.berkeley.edu/cdms_restricted/electronics.

[141] R. Schnee, private communication.

[142] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C: The Art of Scientific Computing* (Cambridge University Press, Cambridge, 1992).

[143] W. Eadie *et al.*, *Statistical Methods in Experimental Physics* (North Holland, Amsterdam, 1971).

[144] R. M. Clarke *et al.*, Appl. Phys. Lett. **76**, 2958 (2000).

[145] R. Schnee, CDMSNote in preparation.

[146] R. Bernabei *et al.*, Phys. Lett. B **389**, 757 (1996).

[147] R. Abusaidi *et al.*, in preparation.

[148] CDMS Collaboration, CDMS II Proposal, 1999.

[149] R. J. Gaitskell, CDMSNote/9805001.

[150] R. Bernabei *et al.*, Il Nuovo Cim. **112A**, 545 (1999).

[151] V. Mandic, private communication.