A Cryogenic Search for WIMP Dark Matter

by

Walter Kenneth Stockwell

B. A. (University of California, Berkeley) 1990

M. A. (University of California, Berkeley) 1992

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, BERKELEY

Committee in charge:

Professor Bernard Sadoulet, Chair

Professor Richard Packard

Professor Eugene Haller

Fall 1996

UMI Number: 9723199

UMI
300 North Zeeb Road
Ann Arbor, MI 48103

The dissertation of Walter K. Stockwell is approved:

11/19/96

Professor Bernard Sadoulet, Chair — Date

Richard Packard — 11/19/96

Professor Richard Packard — Date

11/26/96

Professor Eugene Haller — Date

University of California, Berkeley

Fall 1996

Table of Contents

List of Figures and Tables

Acknowledgements

CDMS is (to me at least) a large project and requires the efforts of a number of people, from undergraduates to professors, machinists to engineers. I cannot possibly mention everyone who, as a part of this project, has made my own contribution not only possible, but interesting and enjoyable. A sincere "Thank you" to all of you -- if you are reading this, you are probably one of the people I am referring to.

In particular, I wish to thank my advisor Bernard Sadoulet for providing the opportunity to work on this project. His guidance and leadership have been crucial to the present success of CDMS. Best wishes for continued success.

The post-docs I have worked with have provided an help, encouragement, and guidance. Dan Akerib, Tom Shutt, and Betty Young have been especially helpful, both in actual work and in providing a good role model for an impressionable graduate student.

Of course, no experiment gets very far without graduate students. All of the graduate students have been great, interested in the experiment and physics, and not interested in games and power struggles. Tom Shutt, Angela DaSilva, Peter Barnes, Jr., Alan Cummings, Sunil Golwala, Roy Therrien -- without their hard work, curiosity, and drive CDMS would still be just a timeline. Storn White, though not a graduate student, has worked as hard on CDMS as the rest of us.

Garth Smith has been invaluable in designing the tower and made the production drawings. John Emes constructred all of our detectors (except P2.)

I must also thank Eugene Haller's students Bill Knowlton and Kohei Itoh. Bill built P2 and helped me learn how to run a detector. Kohei worked on characterizing NTD Ge.

Even with all of these great people to work with, I don't think I could have finished graduate school without the support and encouragement of my (non-physics) friends in Berkeley. They've given me much practice in explaining just what that dark matter is supposed to be.

My family has been unwavering in their confidence in my abilities, and generous with their love and support.

I saved the best for last. Emma, I still don't know how or why we managed to find each other, but I do know my life would be much duller without you. I would not have finished graduate school without you. Thank you thank you thank you!

1. Dark Matter

When we look at the night sky, what do we see? Stars, galaxies, nebulae -- the universe. Or do we? There is much evidence today to suggest that much of the universe is dark, and cannot be seen by electromagnetic radiation at any frequency. We see evidence for the existence of this unknown matter, but cannot see it directly. This is the dark matter problem; what is this stuff we cannot see and how much of it is there?

The existence of the dark matter has profound effects upon the nature and structure of all of what we see in the universe. Most of the galaxies we can see are held together by dark matter. The way galaxies clump together into clusters and clusters into larger still super clusters is certainly affected by the nature and amount of dark matter in the universe. Dark matter could even decide the fate of the universe, whether the universe is open and will expand forever, or closed and will recondense in a "big crunch." The Big Bang is the standard model of cosmology, and one of its important features is an expanding universe. We observe this expansion; most of the stars and galaxies we see have their light red shifted before it reaches us. We can use the expanding universe as a convenient measure of how much dark matter exists, by comparing measurements of the mass density of systems to the critical density ρ_c, the density necessary to close the universe.

1.1 The Friedmann Equation

On large scales, the universe appears spatially isotropic and homogeneous. This is seen by galaxy counts and redshift surveys, and also by the extreme isotropy and homogeneity of the cosmic microwave background (CMB). The extreme smoothness in the CMB reflects the smoothness in the mass density of the universe at the time the

CMB decoupled from matter. The metric that describes a spatially isotropic and homogeneous universe is the Robertson-Walker metric[1]

$$ds^2 = d\tau^2 - R^2(\tau)\left(\frac{dr^2}{1-kr^2} + r^2 d\theta^2 + r^2 \sin^2\theta d\phi^2\right), \tag{1.1}$$

where the r coordinate is dimensionless, τ is the proper time, k is a constant which describes the curvature of the space, and $R(\tau)$, the scale factor, has dimension of length. The constant k is limited to the values -1, 0, +1, corresponding to a space with negative, zero, or positive curvature. Because the metric is symmetric, all of the dynamics are wrapped up into one quantity, the scale factor. To discover how $R(\tau)$ evolves, we turn to the Einstein field equations. In particular, the 0-0 component of the Einstein equations yields the Friedmann equation

$$\frac{\dot{R}^2}{R^2} + \frac{k}{R^2} = \frac{8\pi G}{3}\rho, \tag{1.2}$$

where R is again the scale factor $R(\tau)$, k = -1, 0, or +1, and ρ is the mass-energy density of the universe. The factor $\dot{R}/R$ is more commonly called the Hubble "constant", $H(\tau)$ -- it is not really a constant. The Friedmann equation can be rewritten as

$$\frac{k}{H^2R^2} = \frac{3H^2\rho}{8\pi G} - 1 \equiv \Omega - 1, \tag{1.3a}$$

where

$$\Omega = \frac{\rho}{\rho_c} \text{ and } \rho_c = \frac{3H^2}{8\pi G}. \tag{1.3)b}$$

The quantity ρ_c is known as the critical density; this is the density that separates an open universe from a closed universe. Ω, ρ, and ρ_c are all are time dependent. We will label their present day values with a subscript "0". We can see that k sets the sign of (Ω-1). If $\Omega < 1$, k = -1, and we live in an open universe, a universe with negative spatial curvature. Similarly, if $\Omega = 0$, we live in a flat universe; and for $\Omega > 1$, we live in a closed universe. Because k is a constant, and $H^2R^2 > 0$, the sign of $(\Omega - 1)$ cannot change, e.g. an open universe cannot evolve into a flat or a closed universe.

We want to measure the amount of dark matter in the universe. We can start by measuring the current value of Ω, Ω_0, and comparing this to the density of luminous matter in the universe, Ω_{lum}. To do this we must measure ρ_0, the current mass density of the universe, and we must know ρ_{c0}, the current value of the critical density. The value of ρ_{c0} (and many other cosmological quantities such as distances) depends on the present-day value of the Hubble constant, H_0. Unfortunately, our knowledge of H_0 is not very exact. Measurements of H_0 vary from about 50 km/s/Mpc to about 80 km/s/Mpc. It will be useful to explicitly remove our uncertainty in quantities due to their dependence on the Hubble constant by parameterizing H_0 as $H_0 = 100$ h km/s/Mpc, where h = [0.5 , 0.8].

1.2. The Dark Matter Problem

We will show the existence of dark matter by comparing the amount of matter we can see (Ω_{lum}) to the amount of matter we infer exists in various systems such as galaxies, galaxy clusters, and the universe itself.

1.2.1 Luminous Matter

A first attempt to measure Ω_0 might be to look into the universe and add up all the mass we can see -- stars and gas. We can do this by measuring mass-to-light ratios for many galaxies and using this data to make an average mass-to-light ratio for galaxies, Υ_{lum}. We then measure the luminosity density of the sky, j_0, and calculate the average mass density of luminous matter as

$$\rho_{lum} = j_0 \cdot \Upsilon_{lum} . \qquad (1.4)$$

Using this method, ρ_c corresponds to a mass-to-light ratio of $\Upsilon_c = 1200\ h\ \Upsilon_*$, where Υ_* is the mass-to-light ratio of the sun. We find $\Omega_{lum} \approx 0.003\ h^{-1}$ for the solar neighborhood using this method.[2]

1.2.2 Galaxies

We can make measurements of the mass of spiral galaxies in a particularly easy way. Consider looking at a spiral galaxy edge on. Because the galaxy is rotating, we will see one side of the galaxy has a blue shift as it rotates towards us; the other side will have a red shift as it rotates away from us. By measuring the speed of the rotation, we can estimate the mass of the galaxy by equating the force needed to hold the galaxy together to the force from the total mass inside where we measure the velocity:

$$\frac{GM(r)}{r} = v(r)^2, \qquad (1.5)$$

where v(r) is the velocity at a radius r from the center of the galaxy, and M(r) is the mass within that radius. The measurement of v(r) is called the rotation curve of the galaxy. We can measure the rotation curve for many galaxies well beyond the visible disk of the galaxy using the hydrogen 21 cm emission line. In almost all of the rotation curves for spiral galaxies measured, we find a rotation curve that rises at small radii, then flattens to a constant velocity for as far out as the rotation curve can be measured. We see from equation 1.5 that a constant v(r) implies $M(r) \propto r$, so that the mass of the galaxy increases with radius even though we see no visible component of the galaxy. This is in contrast to what we would expect if the disk was the dominant mass of the galaxy. We would then expect the rotation curve as measured by the gas to fall $\propto r^{-1/2}$ for r greater than the visible disk.

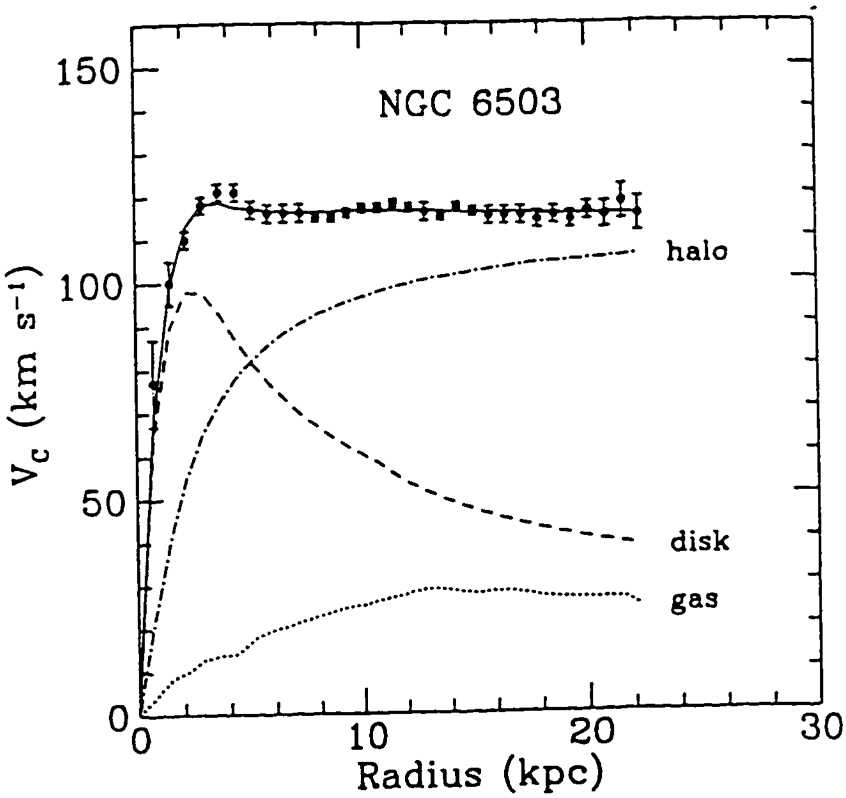

Figure 1.1 Rotation Curve for Spiral Galaxy NGC6503
Shown is the measured rotational velocity for galaxy NGC6503 as a function of radial distance from the galactic center. The labeled curves are the calculated contributions to the measured curve from gas, the visible disk, and a dark halo. The visible disk extends to r ≈ 5 kpc.

Figure 1.1 shows the rotation curve for a typical spiral galaxy, NGC6503.[3] The visible galactic disk extends to r ≈ 5 kpc, and the rotation curve is measured out to r ≈ 22 kpc. The rotation curve is flat well past the visible disk, and shows no sign of falling at the limits of the measurement, showing that we have not found the total mass associated with this galaxy.

We can make some estimates of the mass of our galaxy and other spirals using bound clusters and dwarf galaxies as test particles to extend our knowledge of the rotation curve past the measurable 21 cm emissions. Modeling of the orbits of the satellite clusters and dwarf galaxies of the Milky Way suggest the Milky Way could have a dark halo extending to as much as 200 kpc from the galactic center, with a total mass of a $3 - 5 \cdot 10^{12}$ M_{sun}.[4] If our galaxy is typical, this would imply $\Omega_{gal} \approx 0.1$ h^{-1}.

1.2.3 The Local Group

We can estimate the amount of mass contained in the local group by noticing that the Milky Way and Andromeda are approaching each other despite the expansion of the universe. The two galaxies must have sufficient mass and time to overcome their initial Hubble expansion velocity and fall towards each other with the observed velocity of Andromeda relative to the Milky Way. There is some uncertainty in the actual approach speed of Andromeda; we see Andromeda approaching us at 300 km/s, but somewhere between 65 km/s - 150 km/s of that speed is due to the rotation of the solar system within the Milky Way. Trimble[5] reviews some of the attempts to measure Ω for the Local Group, and finds values of $\Omega_{LG} \approx 0.1$ h^{-1}.

1.2.4 Galaxy Clusters and Superclusters

Measurements of the mass of galaxy clusters have been made using three basic strategies, dynamical measurements using the motions of the individual galaxies within the cluster, modeling of hot gas temperatures in clusters, and gravitational lensing of faint galaxies behind clusters. All three methods agree reasonably well. A well studied example is the Coma cluster. The mass of the Coma cluster was measured two ways in an analysis by White et al.[6] First by measuring the velocity dispersion of the individual galaxies and using the virial theorem to estimate the binding energy of the cluster. Second, by measuring the X-ray temperature of the gas bound to the cluster;

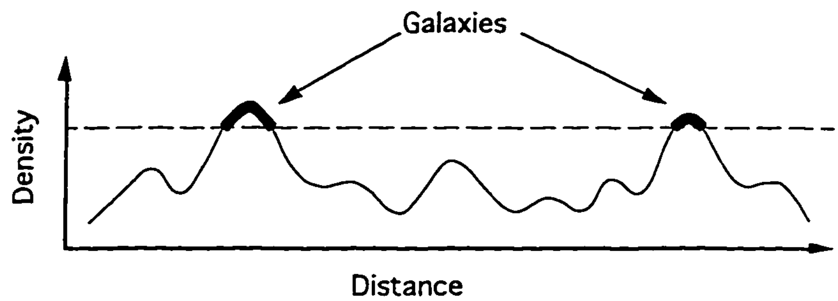

Figure 1.2 Galaxy Biasing.
The figure illustrates the problem of biasing in galaxy formation. If galaxies only form at the highest peaks in the density distribution, we will miss significant amounts of mass by only looking at galaxies.

the x-ray flux gives us information about how much gas is bound, and the actual temperature tells us how deep the cluster potential well is. The two methods agree with each other, finding the mass of the cluster to be $5.7 - 11 \cdot 10^{14}\ h^{-1}\ M_{sun}$ within a radius of 1.5 Mpc, corresponding to an $\Omega_{cluster} = 0.2 - 0.4$.

1.2.5 Large Scale Flows

Ideally, we want to measure Ω on the largest scale possible, so that our measured value of Ω will represent Ω_0, the true value of Ω for the universe as a whole. Attempts have been made to use measurements of large scale flows of galaxies and clusters to measure Ω. These calculations are complicated by the unknown effect of biasing of the mass-to-light distribution. We do not know that galaxies are accurate tracers of mass; galaxies could, for example, form only at the highest peaks of the density distribution. (figure 1.2) We then miss significant amounts of mass by only tracing the mass distribution of the visible galaxies. The relationship between mass and light is often assumed to be linear with a biasing parameter b, such that

$$b = \frac{\delta n / n}{\delta \rho / \rho}, \tag{1.6}$$

where $\delta n/n$ is the measured excess in galaxy count and $\delta\rho/\rho$ is the underlying excess in density.

Dekel[7] reviews the various methods of estimating Ω_0 from large scale flows. Most of these methods measure $\beta \equiv \Omega^{0.6}/b$, so we will only have a limit on Ω that depends on our knowledge of the biasing parameter b. As an example of these methods, people have measured β by comparing velocity inferred from the CMB dipole with galaxy counts. The Local Group is moving with a speed of $\approx$ 630 km/s with respect to the CMBR, which must represent our peculiar velocity with respect to the Hubble flow. If this peculiar velocity arises from gravitational interaction of the Local Group with a nearby mass concentration, then the dipole velocity should be aligned with the direction of this excess mass, as evidenced by an excess in the galaxy count. In fact, the galaxy distribution of the IRAS catalog shows an excess of galaxies aligned within 20° of the CMBR dipole, and suggests $\beta=0.9\pm0.2$.[7]

The POTENT method measures β in a different way. POTENT is a procedure of reconstructing the actual three-dimensional galaxy velocity field from the radial velocities by assuming a curl-free field.[8] We can reconstruct the mass field by assuming the velocities arise from gravitational instability. Again, we cannot find Ω directly, but β:

$$\nabla \cdot \mathbf{v} = -\beta \frac{\delta n}{n}. \tag{1.7}$$

These methods find values for β in the range 0.6 - 1[7], suggesting a firm lower bound $\Omega_0 > 0.3$.

1.2.6 Theoretical Considerations

Theoretical considerations lead us to believe Ω_0 should be very close to 1. We see from the Friedmann equation (1.3) that because H and R are time dependent, $(\Omega-1)$ is also. However, if the universe is flat, that is if $k = 0$ in the Friedmann equation, we find $\Omega = 1$ for all time. In a flat universe, Ω is constant. If $k \neq 0$, we can calculate the time evolution by relating the Hubble parameter to the density, ρ. At early times, the curvature term in the Friedmann equation can be neglected, and we find

$$H^2 \propto \rho. \tag{1.8}$$

This follows directly from (1.3) by neglecting k/R^2. The density, ρ, is proportional to R^{-3} when the universe is matter dominated, and proportional to R^{-4} when the universe is radiation dominated. Putting this together, we can write down the scaling of $|\Omega-1|$ as the universe expands,

$$|\Omega-1| \propto \left(\frac{R_{EQ}}{R_0}\right)\left(\frac{R}{R_{EQ}}\right)^2, \tag{1.9}$$

where R_{EQ} is the scale factor at the transition between a matter dominated universe and a radiation dominated universe.[1] $R_{EQ} \approx 10^{-4} R_0$. Because we know the quantity $|\Omega-1|$ is at most of order 1 today, we can infer that $|\Omega-1| \leq 10^{-60}$ at the Planck time. Faced with this apparent fine tuning of the mass density of the universe, it seems that $\Omega = 1$ exactly is the natural value for Ω.

Inflationary theories of the early universe force $\Omega \rightarrow 1$. The theory of inflation supposes that the early universe went through a period of exponential growth. Inflation provides an explanation for the observed isotropy of the universe, provides an method of producing the initial spectrum of density fluctuations, and dilutes away unwanted relics such as magnetic monopoles that would otherwise be observed. Inflation provides for Ω =1 by making the curvature radius of the universe very large, thus ensuring we live in a flat universe. Future space missions such as MAP, scheduled to be launched in 2001, and COBRAS/SAMBA, scheduled to be launched in 2004, should provide data about the CMBR with enough resolution and accuracy to seriously test inflationary models.[9]

1.3 Dark Matter Candidates

There is strong evidence that dark matter exists. What could this dark matter be? It is convenient to divide potential dark matter candidates into broad categories, and examine the qualities these categories have before focusing on specific candidates. The first such division will be between baryonic and non-baryonic dark matter candidates.

1.3.1 Baryonic Dark Matter

The most obvious candidate for dark matter might be simply normal matter that is somehow too cold or too diffuse to be noticed. We know Jupiter-sized objects form; why can't the dark matter be condensed objects that were too small to start fusion burning? We can get some help on this question from Big Bang Nucleosynthesis (BBNS.)

Big Bang Nucleosynthesis traces the formation of the light elements in the big bang.[1] When the universe is hot ($T \geq 1$ MeV) the population of neutrons and protons are kept in thermal equilibrium by weak interactions. Below this temperature, weak interaction times are greater than the Hubble time, so are no longer effective in keeping equilibrium. At temperatures below about $T \approx 0.3$ MeV, neutrons and protons will start to stick together and form the light elements. Virtually all of the free neutrons end up in

^4He nuclei, with some trace amounts ending up in Deuterium, ^3He, and ^7Li. Heavier elements are formed in stars and super-novae. The relative abundances of all these light elements can be predicted as a function of one parameter, which can be taken as $\eta \equiv n_b/n_\gamma$, where n_b is the number density of baryons and n_γ is the number density of photons in the universe. We can then use BBNS to measure η by measuring the primordial abundances of the light elements and relating measurements of these abundances to predictions of these abundances vs. η. Current measurements constrain $\eta = [4{\cdot}10^{-10}, 7{\cdot}10^{-10}]$, which translates into $\Omega_B h^2 = [0.015, 0.026]$. With our uncertainty in h, this means we know $\Omega_B = [0.015, 0.10]$.

This determination of Ω_B leaves us with two dark matter problems. First, $\Omega_B > 0.01 \geq \Omega_{LUM}$, which means a sizable portion of the baryons in the universe are dark. Second, if $\Omega_0 > 0.10$, which seems very likely from large scale measurements, then much of the dark matter must be non-baryonic. What could the baryonic dark matter be? If it were dust, it would heat up from starlight and radiate in the infra-red. Atomic hydrogen should radiate at 21 cm. Hydrogen in any form would absorb light from distant quasars. All of these are constrained by observation to be much less than what is needed to account for the baryonic dark matter. One viable alternative is to lock the baryons into massive compact halo objects -- MACHOs. MACHOs avoid the previous observational limits by clumping the baryons into small, cold objects that would be too dim to be seen directly, and too small to obscure background sources such as quasars. MACHOs could be brown dwarf sized objects that were not massive enough to ignite ($M < 0.08M_{sun}$), old white dwarf stars that have cooled enough that they cannot be seen today, or black holes.

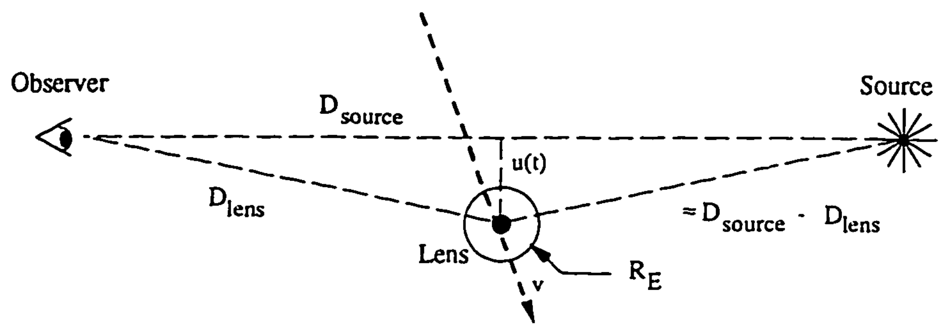

Figure 1.3 Gravitational Microlensing
A compact object moves with speed v through our line of sight to a distant source. R_E is the Einstein radius of the object, and u(t) is the distance between the object and the line of sight to the source scaled by R_E.

Three groups have mounted efforts to look for MACHOs in the halo of our Milky Way galaxy.[10-12] These groups utilize gravitational microlensing of stars to observe the MACHOs. Gravitational microlensing refers to gravitational lensing when the source splitting is too small to be resolved. In this case, the observer sees a brightening of the source star as the lens effectively focuses more light towards the observer. Because of the relative motion of the observer here on Earth, the lensing object, and the source, the amplification will be time dependent. The time dependent amplification can be written

$$A(t) = \frac{u^2 + 2}{u\sqrt{u^2 + 4}}, \tag{1.10a}$$

where

$$u(t) = \left[u_{min}^2 + \left(\frac{2(t - t_{max})}{\hat{t}} \right) \right]^{1/2}, \tag{1.10b}$$

$\hat{t} = 2 \cdot R_E / v$ is the time for the lens to move across the line-of-sight, R_E, the Einstein radius for the lens, is $R_E = \sqrt{4GmD/c^2}$, m is the lens mass, D = $D_{lens}(D_{source}$-$D_{lens})/D_{source}$, and u(t) is the distance between the lens and the line-of-sight scaled by R_E.[13] Figure 1.3 illustrates the geometry.

The MACHO group monitors ≈ 8.6 million stars in the LMC and the galactic bulge. When a MACHO orbiting in the halo intersects a line of sight to a distant star, they see an amplification of the star's luminosity as the gravitational field of the MACHO acts as a lens to gather more of the stars light than they would normally see. Microlensing events can be distinguished from periodic stars by the fact that they should be symmetric in time, amplify by the same amount in all colors, and only occur once.

All three groups have seen microlensing events. Unfortunately, observation of the microlensing events cannot give one enough information about the microlensing objects and their distribution in the galactic halo to determine uniquely their contribution to the dark matter in the halo. An assumption must be made about the distribution of the MACHOs in the galaxy; usually, the assumption is of a spherical halo made of a single mass type MACHO consistent with galactic constraints. Current analyses find it unlikely that MACHOs could make up the entirety of the dark matter in the Milky Way, though they could reasonably make up as much as 30% of the dark halo.[14] These analyses typically find MACHO masses above the hydrogen burning limit, suggesting that these MACHOs might be seen visually.

1.3.2 Non-Baryonic Dark Matter Candidates

If $\Omega_0 \geq 0.10 \geq \Omega_B$, then we have non-baryonic dark matter as well as baryonic dark matter. When examining non-baryonic dark matter candidates, it is useful to divide the candidates into hot dark matter and cold dark matter by their effect on structure formation in the early universe. Non-baryonic dark matter candidates are generally new particles that are extensions to the standard model postulated to resolve some problem with the model.

1.3.2.1 Hot Dark Matter

Hot dark matter refers to particle candidates that are relativistic at the start of galaxy formation. This does not mean that the particles must still be relativistic today. The best candidate for hot dark matter would be a neutrino with a non-zero mass, or a family of neutrinos with non-zero mass. One can show for light neutrinos that [1]

$$\Omega_{\nu 0} h^2 = 1.1 \cdot 10^{-2} \frac{m_\nu}{eV}, \tag{1.11}$$

where $\Omega_{\nu 0}$ is the sum of contributions from all light neutrinos with mass. If $\Omega_{\nu 0} h^2$ is to be of order 1, then $m_\nu \approx 92$ eV.

Hot dark matter is not the favored constituent for most of the dark matter however, due to its effect on structure formation. Hot dark matter has the effect of damping small wavelength perturbations in the power spectrum through free streaming. Free streaming is the process whereby the relativistic neutrinos would migrate away from over dense regions of the universe into underdense regions, thus smoothing perturbations before they become Jeans unstable. The first structures to form in a universe dominated by hot dark matter are supercluster-size objects. Structure formation then proceeds in a "top-down" manner, with galaxy-sized objects being formed fairly late compared to observation. Hot dark matter helps form the large voids and filaments seen, but fails to match observation on smaller scales. A universe with all of the dark matter being hot seems ruled out by this incompatibility with observation.

1.3.2.2 Cold Dark Matter

Cold dark matter refers to particle candidates that are non-relativistic during structure formation. Two such candidates are Axions and WIMPs.

Axions

Axions were first proposed as a solution to the strong CP problem in quantum chromodynamics. Axions arise from an assumed spontaneous breaking of Peccei-Quinn symmetry at a some energy scale f_{PQ}, which guarantees that CP symmetry is not violated in the strong interaction. Axions would acquire a mass that is inversely proportional to this energy scale f_{PQ}, and can couple to nucleons, electrons, and photons with a coupling strength that is directly proportional to their mass. The P-Q energy scale is not set by QCD or theory, which *a priori* leads to a very large range of possible mass for the axion. The very large range of masses can be greatly reduced by lab measurements and consideration of astrophysical processes. Axions with a mass less than about 10^{-6} eV would overclose the universe; axions with a mass greater than about 10^{-2} eV would noticeably affect the evolution of stars. Axions, though light, do not act like hot dark matter as they would never actually have been in thermal equilibrium with the universe; axions form as a zero momentum condensate from the spontaneous symmetry breaking of the P-Q scalar field. Interestingly, the allowed range of mass for the axion coincides with what mass would make it a cosmologically interesting particle.

Axions have the property of coupling to the EM field through $\vec{E}\cdot\vec{B}$. Groups are searching for axions using large, cryogenically cooled high-Q microwave cavities immersed in strong magnetic field. The axions should convert to microwave photons in the magnetic field, thus depositing a power in the cavity at the resonant frequency corresponding to the mass of the axion.

Weakly Interacting Massive Particles

WIMPs, Weakly Interacting Massive Particles, are the other major class of cold dark matter. WIMPs (we will refer to them as χ) are massive particles that fall out of thermal equilibrium with the universe when they are non-relativistic. WIMPs are kept

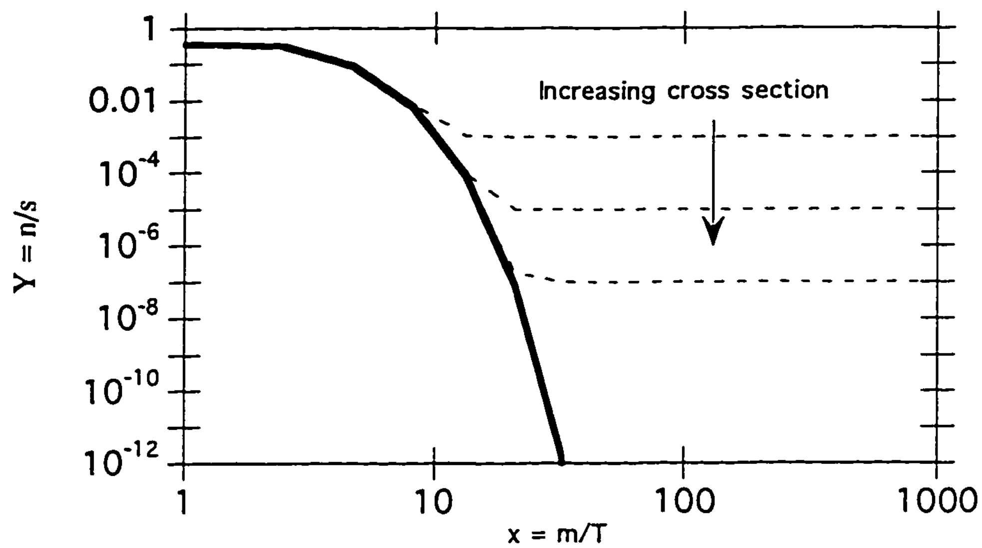

Figure 1.4 Relic WIMP Abundance
Y is the number density of WIMPs per comoving volume; x is the mass of the WIMP scaled by the temperature of the universe. The heavy line shows the abundance of a species that remains in thermal equilibrium; the dashed lines show the relic abundance of a species that decouples from thermal equilibrium. As the interaction cross section of the species increase, the WIMP remains in thermal equilibrium longer, thus decreasing its relic abundance.

in thermal equilibrium with the early universe through reactions like $\chi\bar{\chi}\leftrightarrow q\bar{q},l\bar{l}$. These reactions keep the WIMP coupled to the rest of the universe as long as the reaction rate, Γ is larger than the Hubble expansion rate, H. As the universe cools below m_χ, the above reaction becomes biased away from the χ side, and the number density of WIMPs will fall as $(m_\chi/T)^{3/2}\exp(-m_\chi/T)$. If, however, the reaction rate falls below the expansion rate, the WIMPs will "freeze out" -- they will be unable to annihilate any further, and their abundance will be fixed at some non-thermal value. Figure 1.4 shows the evolution of a species in thermal equilibrium with the universe, and the evolution of a species that falls out of thermal equilibrium. One can show from calculations of Big Bang freezeout that the relic abundance of WIMPs is inversely proportional to $\langle\sigma|v|\rangle$, their average interaction cross section. Assuming that these relic particles are stable, we can use this to calculate their abundance today. A calculation of the WIMP abundance today gives[2]

$$\Omega_\chi h^2 \approx \frac{3\cdot 10^{-27}\,\mathrm{cm}^2\mathrm{s}^{-2}}{\langle\sigma|v|\rangle}. \tag{1.12}$$

This result is roughly independent of m_χ (the result depends on the number of degrees of freedom present in the universe at the time the WIMPs freeze out), and is inversely proportional to the interaction cross section. If $\Omega_\chi h^2 = O(1)$, then the χ interaction cross section is reminiscent of a weak scale interaction. This is suggestive that physics at the weak scale might contribute to a solution to the dark matter problem.

We know that the Standard Model is not the complete picture for particle physics. In particular, new physics of some sort must be introduced at the weak scale to solve the hierarchy problem in the Standard Model. Supersymmetric (SUSY) extensions to the Standard Model are one proposed way of solving the hierarchy problem. SUSY theories postulate an additional symmetry in particle physics that would provide a fermionic supersymmetric partner for every boson, and a bosonic supersymmetric partner for every fermion. The Standard Model lagrangian is specified up to some energy, say Λ_{GUT}, above which we don't understand how the particle physics should be specified. If we use the Higgs mechanism to provide the spontaneous symmetry breaking of the electroweak sector, we run into problems. When we examine radiative corrections to the masses of the scalar Higgs, we find the masses are of the order of our limiting energy scale $\Lambda >> \Lambda_{weak}$. This destroys the observed hierarchy between the electroweak and GUT scales. This can be avoided within the standard model only by fine tuning. Supersymmetry avoids this by providing a fermion partner for every boson, and a boson partner for every fermion. These partners exactly cancel the loop corrections to the masses, removing the radiative mass correction altogether.

SUSY theories predict the existence of new particles with masses of order the weak scale to ensure SUSY is not broken at the weak scale. Most forms of SUSY include R parity, a global symmetry that ensures that the lightest neutral SUSY particle (sparticle) must be stable. Simply stated, R parity ensures that any state with a single sparticle cannot evolve to a state with no sparticle. The lightest sparticle (LSP) is expected to be a superposition of the SUSY partners to the gauge bosons -- the photino, Higgsino, Wino, and Bino. The exact coupling of the LSP to ordinary matter depends somewhat on its exact composition, but in all cases the LSP is predicted to have a mass in the approximate range 10 GeV -- 1000 GeV with a coupling strength of order the weak coupling, just what we want for a cosmologically interesting WIMP.[2] Interestingly, supersymmetric (SUSY) extensions to the standard model of particle physics provide a natural WIMP candidate.

2. The CfPA Cryogenic Dark Matter Search

The Center for Particle Astrophysics, with funding from the DOE and NSF, is coordinating a large experiment to search for WIMP dark matter, the Cryogenic Dark Matter Search, or CDMS. As implied by the name, the experiment will consist of cryogenic detectors which will look for direct interactions of WIMPs with our detectors. CDMS is a medium-sized collaboration involving scientists from UC Berkeley, UC Santa Barbara, Stanford University, and the Lawrence Berkeley National Laboratory.

This author was a graduate student in the Direct Detection group at U. C. Berkeley working on CDMS primarily in thermometry, detector development, and tower development.

2.1. Experimental Strategy

As described in the previous chapter, WIMP dark matter interacts with baryonic matter with a scattering cross section of a weak force type strength, and of course, gravitationally. Although it is through its gravitational interaction that we are able to deduce the existence of dark matter, we can not hope to directly sense dark matter through gravitational interactions in the lab. CDMS will attempt to detect the WIMP dark matter directly through a scattering interaction with a detector. From freeze out arguments we have a prediction of the scattering cross section, $\approx 10^{-36}$ cm^2, if the WIMP dark matter is cosmologically significant. If we assume the WIMPs arise from a SUSY theory, we can guess at the probable mass; SUSY plausibly predicts a the lightest supersymmetric partner (LSP) to be in the range of $\approx$ 10 GeV to $\approx$ 1 TeV. Although it is difficult to directly measure the rotation curve of the Milky Way since we are inside it, estimates have been made of the mass density of our galaxy in our solar neighborhood, $\approx$ 0.3 GeV cm^{-3}. Calculations have been made of the probable velocity dispersion of the hypothesized WIMP halo, giving a mean velocity of the WIMPs ≈ 270 km s^{-1}. Using these numbers we can plan our experiment.

We are looking for the WIMPs that should reside in the halo of our Milky Way galaxy. WIMPs cannot clump gravitationally in the galaxy because they have no efficient means of dissipating energy. Our galaxy should be immersed in a roughly spherical halo of WIMPs. The WIMPs will have a Maxwellian velocity distribution in the halo. We expect the WIMPs to scatter off our detector with the vector sum of their halo velocity and our orbital velocity around the Milky Way. The event rate per unit detector mass should be

$$R = n_\chi \sigma v \frac{1}{M}, \tag{2.1}$$

where n_χ is the number density of WIMPs in our galactic neighborhood, σ is the elastic scattering cross section per nucleus, v is the relative velocity of the WIMPs to our detectors, and M is the target nucleus mass. If we use a Ge target, a 10 GeV WIMP with a scattering cross section of $5 \cdot 10^{-36}$ cm^2 with a halo velocity of 270 km/s will give us a rate ≈ 3 events kg^{-1} day^{-1}. This low rate means we will need large mass detectors and will have to count for long periods of time.

The momentum transfer for each scattering events in our detector is simply

$$|\vec{p}|^2 = 2m_r^2 v^2 (1-\cos\theta), \tag{2.2}$$

where m_r is the reduced mass $m_\chi m_N/(m_\chi+m_N)$, v is the relative speed between the WIMP and our detector, and θ is the scattering angle in the center of momentum frame. The energy deposited in the detector is then

$$E_{dep} = \frac{|\vec{p}|^2}{2m_N} = \frac{m_r^2 v^2 (1-\cos\theta)}{m_N}. \tag{2.3}$$

If we take the case of a 10 GeV WIMP impacting on a Ge nucleus at 270 km/s, we find a maximum $E_{dep} \approx 1$ keV. This is a very small energy deposition; we will need detectors with small thresholds to be able to detect the WIMP events.

A more detailed and accurate calculation of scattering rates and energy deposition can be found in Jungman, et al.[2] In our rate calculation, we ignored important details such as spin vs. scalar interactions, the effect of coherence in the interaction between the WIMP and nuclei, and the possible SUSY composition of the WIMPs. Our simple rate calculation turns out to be high by at least an order of magnitude for realistic SUSY WIMP particles. A careful look at the energy deposited in scattering events allows WIMPs to deposit substantially more energy than 1 keV per scatter; this is primarily due to the actual velocity distribution of the WIMPs in halo having a high energy tail.

2.2. Overview of CDMS

Because the experiment must run in a very low temperature, radio-pure environment, CDMS is much more than a particle detector. CDMS includes an underground facility, a complex shielding system, a unique cryogenic system, and innovative cryogenic detectors and electronics. We will describe the major components of CDMS, starting with its detectors.

2.2.1 Detectors

CDMS is planning on first using two types of detectors. The Stanford group is developing detectors based on silicon absorbers with transition-edge phonon sensors (TES) and ionization collection; the Berkeley group is developing detectors based on germanium absorbers with neutron-transmutation doped (NTD) germanium phonon sensors and ionization collection. Both of these detectors operate at cryogenic temperatures, ≈ 20 mK. These detectors measure both the recoil energy from the phonons recoils produce, and the number of electron-hole pairs produced by particles scattered in the detector. These detectors use different technologies to measure the recoil energy of events, and are based on absorbers of different materials. The ionization measurement is similar in both detector, but different from conventional 77 K ionization detectors. Measuring two independent quantities from each event will enable us to discriminate between our dominant background, photons, and a WIMP signal. It is expected that these different detectors will complement each other. The use of different technologies and absorbers constitute an element of redundancy in identifying a WIMP signal. A WIMP signal seen in one type of detector should give a predictable response in the other type.

2.2.1.1 Berkeley Detectors

The Berkeley group has been developing large cryogenic germanium detectors that simultaneously measure the recoil energy and ionization produced by events in the detector.[15] We use a short cylinder of ultra-pure germanium as our target crystal. We measure the recoil energy caloriometrically with NTD germanium thermistors. The surface of the germanium crystal is implanted to make it conducting; we apply a drift field across the detector and measure the charge produced in the detector by particle recoils. These detectors are operated in a dilution refrigerator at ≈ 20 mK. We have built and tested a 60 g detector, E5, and are in the process of constructing six 170 g detectors, the Berkeley Large Ionization and Phonon detectors, or BLIPs.

Phonon Measurement

We measure the recoil energy of the particle interaction with neutron-transmutation doped (NTD) germanium thermistors that have been doped very close to the metal-insulator transition. We use a specific type of NTD Ge, NTD #29, that has a useful resistivity at 20 mK. Electrical conduction in NTD Ge occurs through thermally assisted hopping. The performance of our thermistors is described by a hot electron model, in which the conduction electrons are not well coupled to the lattice of the thermistor, and can heat above the lattice temperature. We have shown that the behavior of our thermistors at 20 mK and low bias is not well described by an electric field effect.[16] This is troublesome experimentally. To measure the resistance of a thermistor accurately, we would like to use as much current as possible. A large current will heat the electrons in the thermistor, and our resistance measurement will no longer reflect the temperature of the phonons in the thermistor. A "large" current for our thermistors at 20 mK is a few nA. The resistivity of our thermistors is a functi.›n of the temperature of the conduction electrons,

$$\rho = \rho_0 \exp\left(\sqrt{\Delta / T_e}\right), \tag{2.4}$$

where ρ_0 and Δ are parameters that describe the performance of each thermistor, and T_e is the temperature of the conduction electrons. Typical values of ρ_0 for NTD 29 are $\rho_0 \approx 0.1\ \Omega$-cm; typical values for Δ for NTD 29 are $\Delta \approx 7$ K. We can adjust the actual resistance of our thermistor at its operating temperature by changing its geometry.

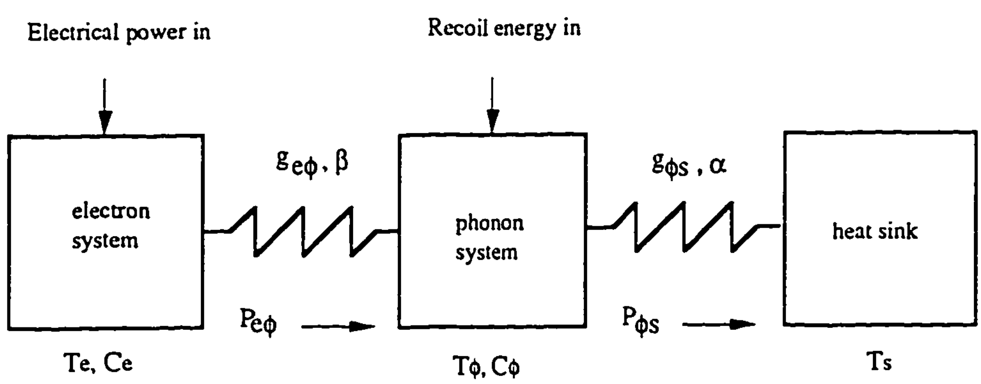

Figure 2.1 The Hot Electron Model.
The hot electron model assumes the electrons in our thermistor are thermally decoupled from the phonons. The electrons have a temperature, T_e, and heat capacity, C_e. The phonon system has a temperature, T_ϕ, and heat capacity, C_ϕ. The electrons are coupled to the phonons through $\kappa_{e\phi}$, and the phonons to the sink through $\kappa_{\phi s}$. These couplings allow power to flow between the electrons and phonons, $P_{e\phi}$, and phonons and sink, $P_{\phi s}$. The bias power dissipated in the thermistor enters as heat in the electron system; particle interactions inject their recoil energy into the phonon system.

At cryogenic temperatures the electrons in the thermistor decouple thermally from the phonons in the thermistor. In the hot electron model, as shown pictorially in Figure 2.1 the electrons have a temperature, T_e, a heat capacity, C_e, and some conductance to the phonon system, $g_{e\phi}$. The phonon system (we assume the phonons in the thermistor are well coupled to the phonons in the detector substrate)[17] also has some temperature, T_ϕ, heat capacity, C_ϕ, and conductance to the heat sink, $g_{\phi s}$, as well as to the electrons. The electrical resistance of the thermistor is assumed to be a function only of the electron temperature, T_e. The hot electron model effect makes the I-V curve of a thermistor highly non-linear. As we push more current through the electron system, the electrons heat up, which lowers the resistivity of the NTD Ge. The effect can be quite dramatic; thermistor I-V curves often "saturate" at some voltage; increasing current does not increase the voltage across the thermistor, and in some cases even lowers the voltage. We usually bias our thermistors near the top of the linear region as this corresponds to our optimal pulse height.

The decoupling between the electron and phonon systems has a profound effect on the thermistors sensitivity. We measure an electrical signal, and are thus sensitive to changes in T_e. Recoils in the detector, however, heat the phonon system. The electron-phonon coupling can be described as a thermal conductance much like thermal conduction in a solid:

$$\kappa_{e\phi} = \beta g_{e\phi} T^{\beta-1}, \tag{2.5}$$

where $g_{e\phi}$ and β are parameters we must measure for each NTD Ge. Experimentally, $g_{e\phi}$ is a function of the thermistor volume. This is because the electrons and phonons are in "thermal contact" throughout the volume of the thermistor. β should be a material parameter, and independent of the thermistor geometry. With this conductance, we can calculate the thermal power that flows between the electron and phonon systems when the thermistor is biased:

$$P_{e\phi} = \int_{T_\phi}^{T_e} \kappa_{e\phi}(T)dT = g_{e\phi}\left(T_e^\beta - T_\phi^\beta\right). \tag{2.6}$$

There is a similar relation between the phonon system and the heat sink for the system:

$$P_\phi s = \int_{T_s}^{T_\phi} \kappa_{\phi s}(T)dT = g_{\phi s}\left(T_\phi^\alpha - T_s^\alpha\right). \tag{2.7}$$

For our NTD 29 thermistors, we find $\beta = 6$ and $g_{e\phi}/V = 0.05$ W K^{-6} mm^{-3}, where V refers to the thermistor volume. The values of α and $g_{\phi s}$ are determined by the geometry of the detector heat sinking; we sink the detector through a gold pad sputtered onto the crystal. The conductance between the crystal and sink is dominated by the acoustic mismatch resistance of the gold-germanium interface. We have measured this resistance, and find $\alpha = 4$ and $g_{\phi s}/A = 2.9 \cdot 10^{-5}$ W K^{-4} mm^{-2}, where A refers to the area of the heat sink. We will discuss the hot electron model in more depth in Chapter 3.

The thermistors are approximately current biased through a large bias resistance. When a particle scatters in the detector, the interaction will heat the germanium absorber. As the thermistor heats up, its resistance drops, and because it is current biased, the voltage across it drops as well. The voltage is read out using a low-noise FET in the refrigerator as the front end of a low-noise voltage amplifier. The resulting pulse is relatively slow. The rise time is ≈ 1 ms, and the fall time is ≈ 20 - 50 ms. The fall time is set by the heat-sinking of the crystal to the refrigerator. With E5 we have demonstrated a baseline electronics noise of 600 eV FWHM and a resolution of 775 eV FWHM at 60 keV. Figure 2.2 is a diagram of the phonon measurement. Figure 2.3 shows the phonon spectrum obtained by E5.

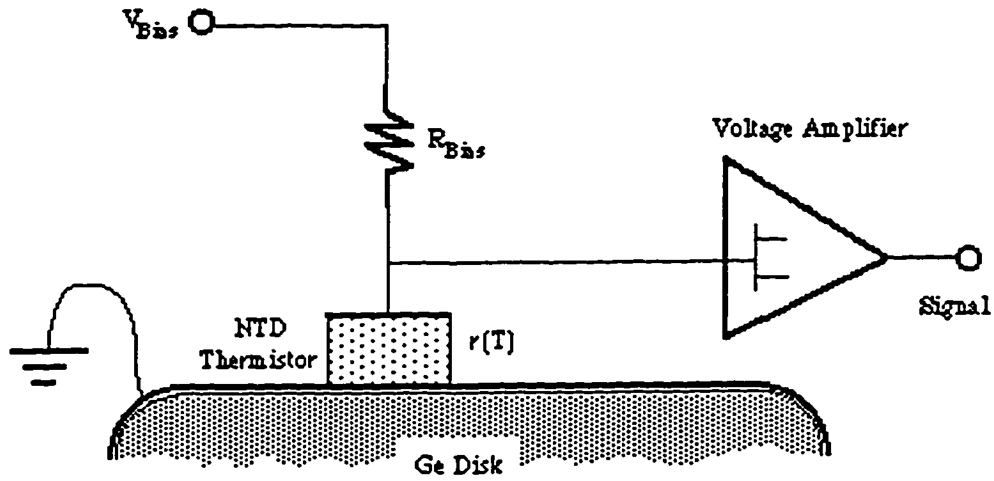

Figure 2.2 The Phonon Measurement Circuit.
The NTD thermistor is bonded to the surface of the detector. The surface is doped to form a conducting layer, and is grounded. We apply a bias voltage at the top of a bias resistor, $R_{BIAS} >> r(T)$. This makes the sensor roughly current biased. The voltage at the top of the thermistor is sensed by the front-end FET, and amplified by the voltage amplifier.

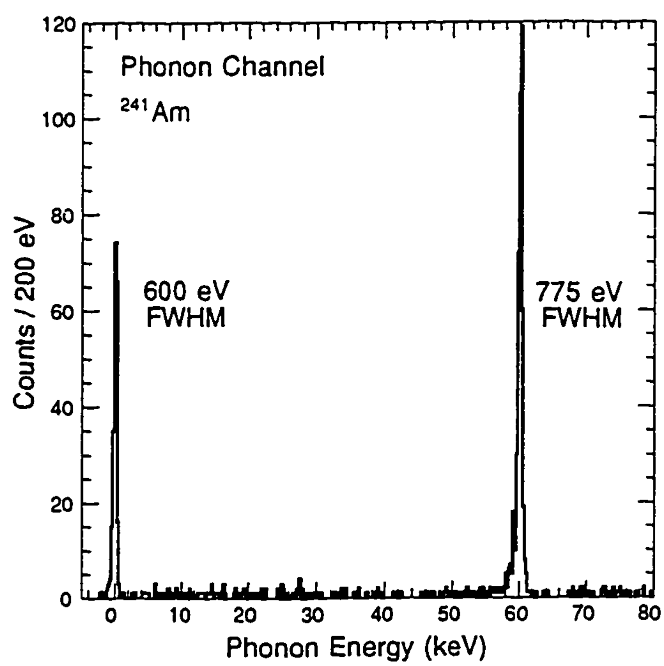

Figure 2.3 Phonon Spectrum in E5.
E5 is a 60 g detector made by the Berkeley group. Shown is the phonon spectrum of E5 when exposed to a 60 keV gammas from a collimated ^{241}Am source. The baseline noise is 600 eV FWHM and the resolution is 775 eV FWHM at 60 keV.

Ionization Measurement

The surfaces of the crystal are ion implanted beyond the metallic limit to form a conducting layer. The surface is divided at the equator to form two halves. The half with the thermistors is called the "phonon" side; the other half is called the "ionization" or "charge" side. The charge side is further segmented into an inner disk and an outer ring. These sections can be biased separately, and can measure charge separately. The segmentation is designed so that the two channels collect charge from roughly equal detector volumes. The phonon side is grounded; we apply a bias voltage to the other contact, typically ≤ 0.5 V. The detector is capacitively coupled to a current amplifier using a front end FET located inside the refrigerator. Figure 2.4 shows the measurement circuit for the ionization measurement.

Charge collection in these detectors is very different from conventional germanium ionization detectors run at 77 K. 77 K semiconductor detectors require a large field to be applied to deplete the detector of free charge so that one can detect the charge created by particle interactions in the detector. Typical drift fields are on the order of a few kV cm^{-1}. A drift field of this magnitude would be disastrous for our cryogenic detector. As we will see, a large drift field will destroy the sensitivity of our phonon measurement.

For a given amount of energy deposited in the detector by a recoil event, E_{dep}, εE_{dep} will initially go into creating electron-hole pairs. The remaining $(1-\varepsilon)E_{dep}$ goes into phonons directly. The parameter ε is $\approx 1/3$ for electron recoils and $\approx 1/9$ for nuclear recoils. The exact value is slightly energy dependant, with events being generally less ionizing for less recoil energy. The number of charge pairs created will be the initial ionization energy divided by the bandgap, or $n_Q = \varepsilon E_{dep}/E_g$.

When we bias the detector with a voltage V, we apply an electric field V/t across the detector, where t is the detector thickness. The electrons will drift towards the

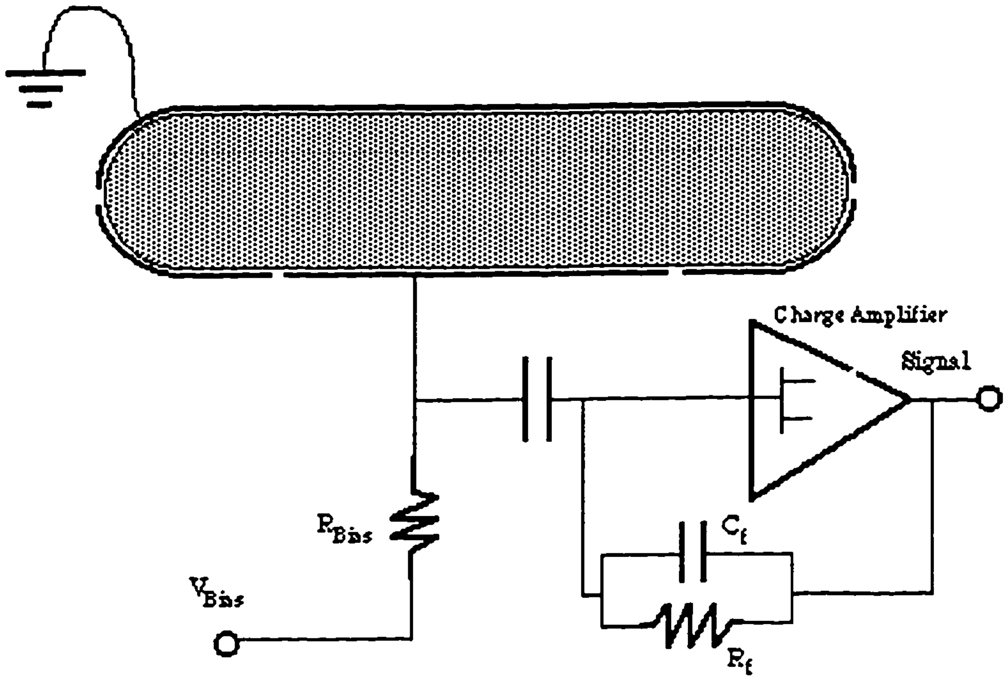

Figure 2.4 The Ionization Measurement Circuit
The surface of the detector is implanted to form a conducting layer, but the implant is broken on the equator to electrically isolate the two halves. The top half is grounded; the bottom half is biased to drift charge across the detector. The bottom surface is segmented radially to form two independent charge channels. The front-end FET is capacitively coupled to the detector. For clarity, only the inner charge amplifier circuit is shown.

positively biased contact and the holes towards the negatively biased contact. The charge carriers gain energy as they move across the detector,

$$\Delta E = eV\frac{l}{t}, \tag{2.8}$$

where l is the distance the charge drifts. We should consider electrons and holes separately; in the absense of trapping, however, an electron-hole pair produced in the detector will gain energy $\Delta E = eV$, just as if a single charge had drifted entirely across the detector. This energy is radiated in phonons and adds to the phonon signal measured. The total energy added is $n_Q\Delta E = en_QV$. The thermal amplification of an ionization signal is known as the Luke effect.[18] This effect is present in our devices and conventional p-i-n devices.

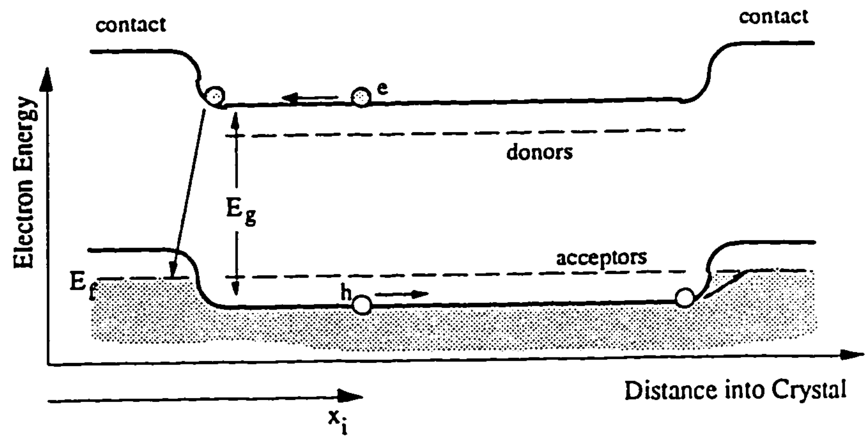

Figure 2.5 Charge Collection in a p^+-p-p^+ Device.
The contacts in a p^+-p-p^+ device are doped with acceptors so that the Fermi level in the contacts equals the energy of the acceptor impurity sites above the electron valence band. An event occurs at a distance x_i into the crystal producing e^--h^+ pairs. The electrons drift to the positive electrode (left in the figure), the holes to the negative electrode (right). At the contact, the electrons fall down to the fermi level. The holes relax up to the fermi level in the opposite contact. The initial energy E_g needed to create the charge pair is released in phonons when both the e^- and h^+ relax at the contacts.

Our detectors are made from high-purity p-type germanium. We use boron degenerately doped ionization contacts on both sides of the detector, so that the detectors are p^+-p-p^+ devices. Electrons that reach the contact will recombine at the contact, as the contact is a resevoir for holes. The electron releases the gap energy in phonons in this process. A hole that reaches the contact will fall down to the Fermi level, releasing E_f in phonons in this process. ($E_f << E_g$, so we are not too concerned with this process.) The net effect is to transfer the entire energy E_g that went into creating the charge pair into phonon energy. Figure 2.5 shows a diagram of the charge collection process in E5. This is in marked contrast to a p-i-n device, where there is no recombination at the contacts, and the energy used in creating the electron-jole pairs will not be seen in a phonon measurement.

These two processes (the Luke effect and charge recombination at the contacts) are very fast compared to the phonon measurement and together will add an energy E_Q

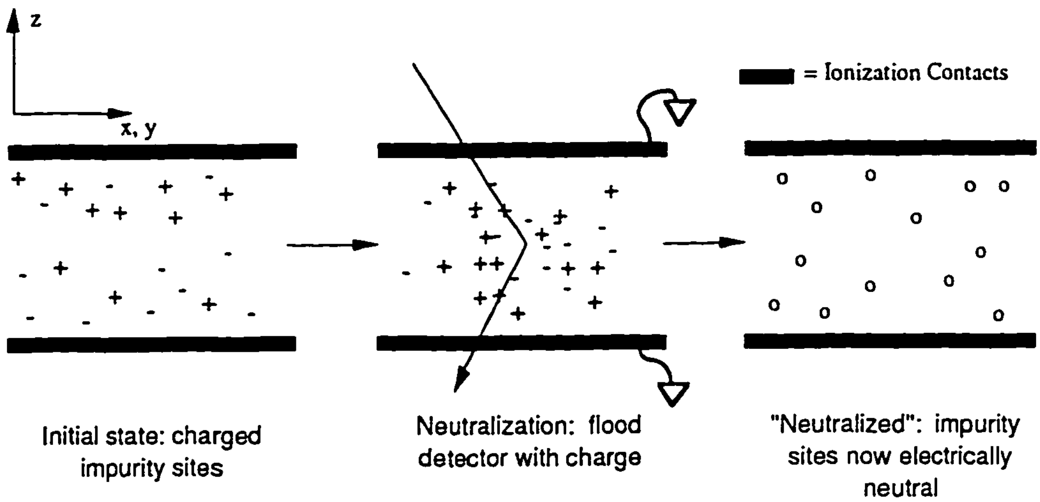

Figure 2.6 Neutralizing the Berkeley Detector.
When first cooled own, the detector will have many charged impurity sites that will act as charge traps and interfere with the charge measurement. The z-axis is the thickness of the detector. We neutralize the crystal by grounding both ionization contacts and irradiating the detector with an IR LED to flood the crystal with free charge. The free charge flows to the charged sites and become trapped, neutralizing the crystal.

to the phonon measurement,

$$E_Q = n_Q E_g + e n_Q V = \varepsilon E_{dep}\left(1 + \frac{eV}{E_g}\right). \tag{2.9}$$

The recombination of the electrons in the contact returns the energy that originally went into creating the electron-hole pairs, and the drifing electrons and holes add the Luke effect energy to the energy that will be seen in the phonon measurement.[19] The bandgap in germanium at 20 mK is ≈0.7 eV; if we do not want our thermal measurement to be dominated by the Luke effect, it is clear we must keep the drift voltage low. A drift voltage of 0.3 V will give a 10% increase to the total thermal signal measured. For this reason, we are compelled to use a very low bias across the detector.

Using such a low drift voltage is possible because the detector is run at 20 mK. The thermal energy available at 20 mK, k_BT, is 17 μeV << 1 eV, the bandgap energy in

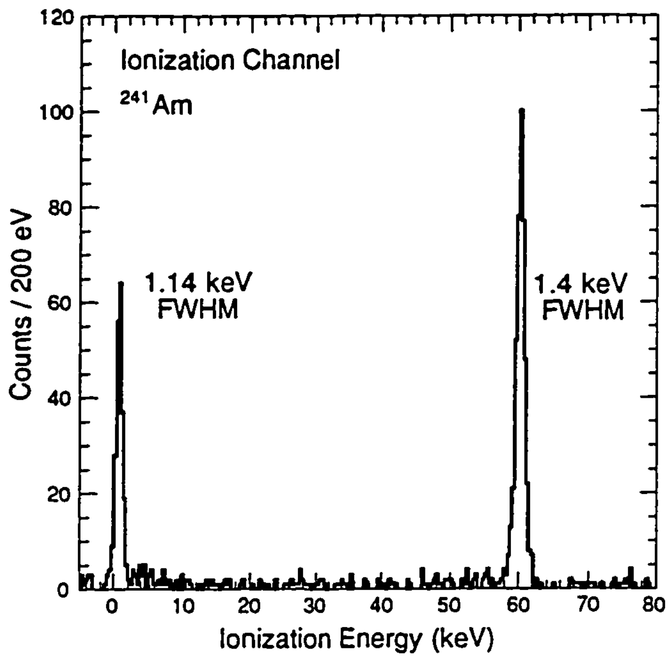

Figure 2.7 Ionization Spectrum in E5.

Shown is the ionization spectrum measured by E5 when exposed to 60 keV gammas from a collimated 241Am source. The electronics baseline noise is 1.14 keV FWHM and the resolution at 60 keV is 1.4 keV FWHM.

germanium. No free charge can be thermally excited in the detector, so we do not need a large drift voltage to deplete the detector. We do have a problem with trapping on charged impurity sites (A^-, D^+). Charge trapping directly affects our ionization measurement, and indirectly affects our phonon measurement through the recombination and the Luke effect, as shown in equation 2.9. The acceptor impurity sites sit about 10 meV above the valence band; the donor impurity sites sit about 10 meV below the conduction band. We would like to have the detector in an electrically neutral state, with no space charge and all the impurity sites filled (A^0,D^0). We achieve this by grounding both ionization contacts and irradiating the crystal with an infrared light emitting diode or γ rays. The photons produce electron-hole pairs in the crystal which drift to and trap on the charged sites, neutralizing them. This non-equilibrium state is stable at 20 mK because the thermally available energy is much less than the impurity gap. This process is illustrated in figure 2.6.

The ionization measurement degrades over time. Events in the biased detector will tend to reionize the impurity sites; in Berkeley, with an event rate of ≈ 5 Hz, this necessitates resetting the crystal after ≈ 30 minutes. The reset time will be longer underground as we will have a much lower event rate. We reset the crystal by again grounding both contacts and turning on the LED for a brief time.

Figure 2.7 is a measured ionization spectrum from E5 of the 60 keV line from a ^{241}Am source. We see a width of 1.4 keV FWHM at 60 keV and 1.41 keV FWHM on the baseline.

Background Discrimination

Measuring two quantities for each event (charge pairs produced and recoil energy) allows us to discriminate between WIMP events and our dominant background events. WIMPs will scatter off of nuclei in the detector, while photons will scatter from electrons in the detector. Electron recoils are more efficient at charge production than nuclear recoils. By measuring the charge pairs produced and the recoil energy of each event, we can calculate the ionization yield of each event, and use the ionization yield to discriminate between nuclear and electronic scatters. Figure 2.7 shows the ionization yield, η, measured for events in the Berkeley detector E5 under two conditions. We define η as

$$\eta = \frac{E_Q}{E_\phi}, \tag{2.10}$$

where EQ is the energy measured in the charge measurement and Eϕ is the energy measured in the phonon measurement. We have normalized the ionization axis so that a a 60 keV γ event is 60 keV on both energy axes. First we took data with a ^{241}Am source, shielded and collimated to only let the 60 keV gammas interact with the

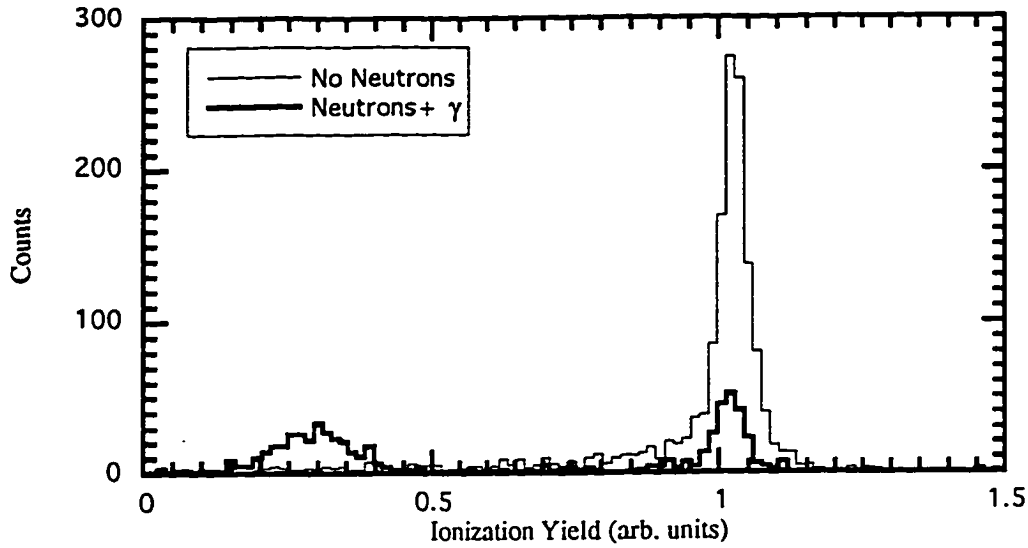

Figure 2.8 Background Discrimination with E5.
Shown is a histogram of the ionization yield, η, measured for E5 under two conditions. The "No Neutrons" line shows events measured with a ^{241}Am source shielded to only put 60 keV photons into the detector. The "Neutrons" line shows the distribution of events when we add a strong neutron source near the detector. The neutron data shows a clear separation between the photon events with $\eta \approx 1$ and the neutron events with $\eta \approx 0.3$.

detector. We see a distribution of events centered around $\eta = 1$. Second, we placed a strong ^{252}Cf source near the detector, outside the refrigerator. ^{252}Cf decays by alpha emission and spontaneous fission, producing both photons and neutrons. The neutron spectrum has a mean energy of $\approx$ 2 MeV with a high energy tail up to $\approx$ 8 MeV. We surrounded the source with a lead shield to block the photons. MeV neutrons will scatter elastically off nuclei in the detector, much like WIMPs. We see a second class of events with $\eta \approx 0.3$, well separated from the photon events, from the neutrons scattering in the detector. We note the "No Neutron" data shows some events in the range η = [0, 0.5]. These events are either real neutron scatters from background neutrons in the room or misidentified photon events with poor charge collection. We can estimate our background discrimination capabilities using the data shown in figure 2.8. We use the data with neutrons to define a nuclear recoil window, from η = 0.13 - 0.45. We then sum the number of events in this window in the no neutron data and

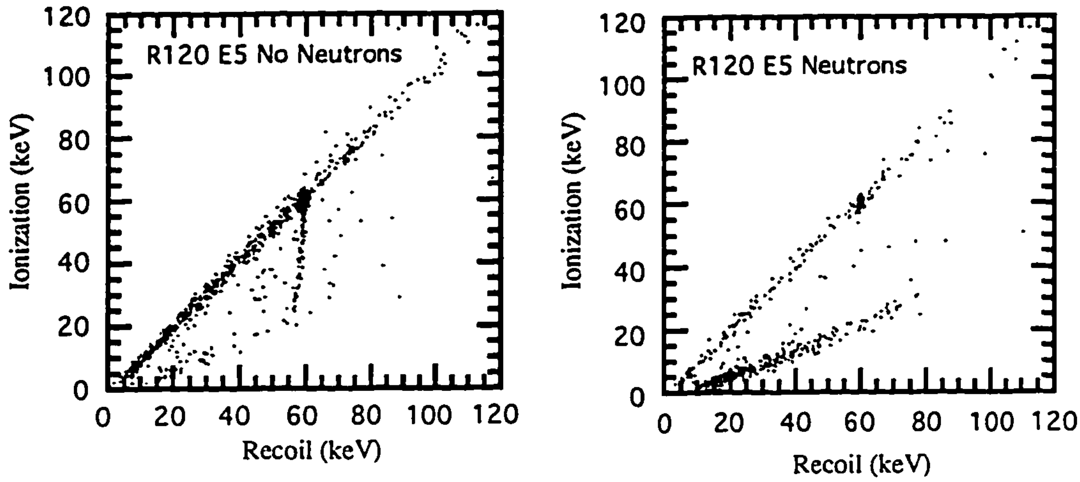

Figure 2.9 E5 Spectra: Without Neutrons and with Neutrons.
This is the data on which figure 2.7 is based, the comparison of ionization yield for nuclear and electronic recoils. The ionization measurement is calibrated to the charge produced by a 60 keV gamma ray. We can plainly see the existence of a second class of events, nuclear recoils, in the "Neutron" data. We have made no cuts in the data.

compare to the total number of events in the data, from $\eta = 0$ - 1.5. Let us define the function R,

$$R = 1 - \frac{N_n}{N_T}, \tag{2.10}$$

where N_n is the number of counts in the nuclear recoil window as defined above, and N_T is the total number of counts in the data set. For this data from E5, we find R = 0.96. Figure 2.9 shows the actual data we used to make figure 2.8. In the first data set, without neutrons, we see a single class of events lying on a line of slope 1, corresponding to electron recoils from the 60 keV gammas from the ^{241}Am source, and a continuum of gammas from background in the room. The second data set shows a second class of events lying on a slope of ≈ 3, corresponding to nuclear recoils from the neutron source. (We have normalized the ionization axis so that the 60 keV gammas are 60 keV in both measurements.) We see now that the data on which we base our measurement of R probably contained some neutrons at low energy, as well as some misidentified photons. We might be able to improve our R by use of cuts, rejecting low energy events where our resolution does not completely separate the two classes of events. We will be rejecting signal events as well as background events. P. D. Barnes, Jr. has made an analysis of our background rejection as a function of signal acceptance.[20] The quantity of interest in our dark matter search is the background rate multiplied by an effective rejection factor, Q(E):

$$Q(E) = \frac{\beta(1-\beta)}{(\alpha-\beta)^2}, \tag{2.11}$$

where α is the signal acceptance and β is the background acceptance when one is cutting on events above E. Smaller Q gives a more sensitive dark matter limit. For the above data, with $\alpha = 1$, we find Q(0) = 0.04. Barnes shows that with Q(2 keV) = 0.02, and a very efficient muon veto, the BLIP detectors should be capable of achieving a sensitivity to heavy WIMPs of 0.3 events/(kg day).

Berkeley Large Ionization and Phonon Detectors

For CDMS, the Berkeley group is building 170 g germanium detectors, the Berkeley Large Ionization and Phonon (BLIPs) detectors. We plan to initially install six of these detectors, thus giving us 1 kg of detector. The first BLIPs will be made of natural germanium. Because of the possibility that the WIMPs could have a spin coupling, we have also obtained enriched ^{76}Ge (spin 0) and enriched ^{73}Ge (spin 1) from which to construct additional BLIPs.

2.1.1.2 Stanford Detectors

The Stanford group is working on a different type of cryogenic detector. Whereas the Berkeley germanium detectors measure the recoil energy of an event by measuring a temperature rise resulting from the thermalized phonons from the interaction, the Stanford detectors measure the initial athermal phonon pulse, before the

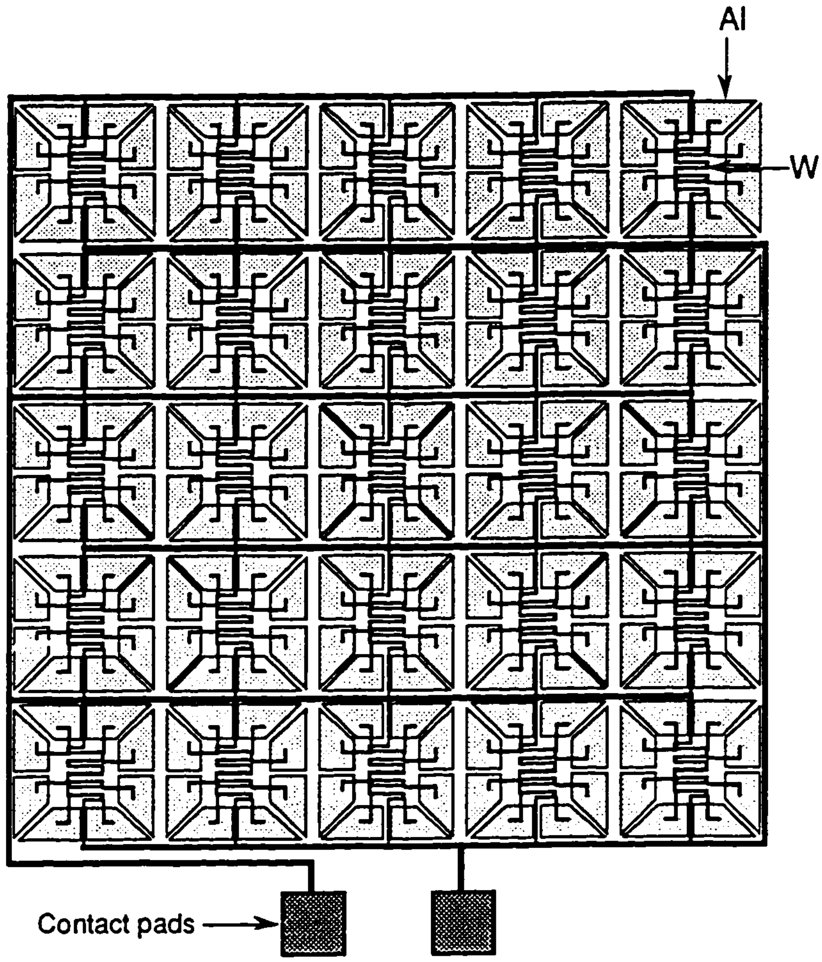

Figure 2.10 The Stanford W/Al QET Sensor Layout.
Shown is the Stanford W/Al QET phonon sensor layout. Each square in the layout contains a tungsten TES meander connected to aluminum phonon collection pads. The TES meanders are connected in parallel and are voltage biased.

phonons can thermalize. A recoiling nucleus should produce a "hot spot" of phonons with energies of order the Debye temperature for the material.[21] These phonons propagate ballistically through the crystal, and thermalize in the detector. These athermal phonons possibly carry information about the type of event (nuclear or electronic recoil), the direction of recoil, and the position of the event in the detector. The Stanford detectors combine a measurement of the athermal phonon pulse with an ionization measurement. The phonon measurement uses tungsten/aluminum quasiparticle trap assisted electrothermal feedback transition edge sensors (W/Al

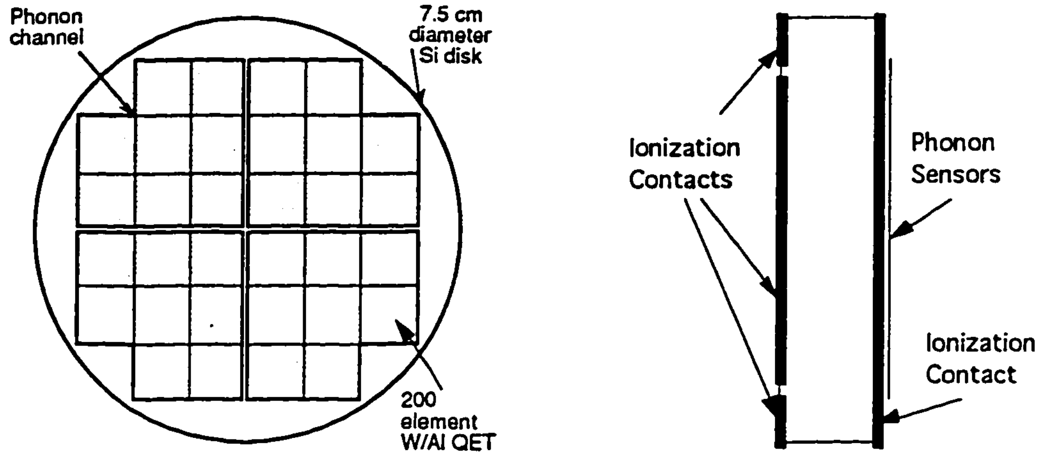

Figure 2.11 The Stanford Large Ionization and Phonon Detectors
One face of the SLIP is covered with 32 W/Al QET sensors grouped into four phonon channels. In addition, both faces of the disk are implanted to form ionization contacts. The face opposite the phonon sensors is segmented into a central disk and guard ring.

QET).[22] These W/Al QET sensors allow us to instrument large surface areas without a large heat capacity, thus allowing us to make fast sensors on large crystals. The current Stanford detectors are based on silicon absorbers, but we are developing W/Al QET sensors on germanium absorbers.

The W/Al QET sensors work using a large area aluminum film to absorb the athermal phonons and tungsten transition edge sensors to measure the energy absorbed by the aluminum pads. A recoil event in a crystal produces a hot spot of high energy ($\approx$ 2 meV in silicon) phonons. When these phonons reach the surface of the crystal, they will have enough energy to be absorbed in the aluminum film ($2\Delta_{Al} \approx 0.3$ meV) by breaking Cooper pairs and creating quasiparticles. The quasiparticles diffuse through the aluminum film until they reach a tungsten transition edge sensor (TES). Tungsten has a much lower gap energy than the aluminum ($2\Delta_W \approx 0.02$ meV), so the quasiparticles relax to the lower tungsten gap. As the quasiparticles relax, they emit phonons which are absorbed by the tungsten. The net effect has been to collect the energy incident on the crystal surface from the athermal phonon pulse, and transfer it to the phonon system of the tungsten TES. Figure 2.9 shows the physical layout of the Stanford W/Al QET phonon sensors.

A transition edge sensor operates by maintaining a superconductor at its superconducting transition. Power into the superconductor will heat it and tend to drive it normal. The TES could be either current or voltage biased. By using voltage biasing, it is possible to take advantage of electrothermal feedback in the sensor[22]. In this method, the TES is voltage biased in the middle of its transition region, and the current is monitored with a SQUID. The TES is maintained at its bias point by its Joule self-heating. Because the sensors are voltage biased, if the sensor resistance decreases, the Joule power dissipated increases, and heats the tungsten; if its resistance increases, the power dissipated decreases, and the tungsten cools. The power dissipated in the TES is simply $I \cdot V$. When an event heats the tungsten TES, the current drops to decrease the Joule heating. The power incident on the TES is $\Delta I \cdot V$, and the total energy is the integral of $\Delta I \cdot V$. Because the tungsten film is the only part of the system that heats, and its heat is actively removed by the electrothermal feedback, these sensors are very fast. Pulses in the W/Al QET are of order 100 μs.

We are able to cover large areas of the crystal surface with the phonon sensors by connecting individual sensors in parallel. Our current detectors are made with a 7.6 cm diameter Si absorber. We are able to cover most of the surface area with four separate phonon channels, each channel being made up of eight of the QET arrays shown in figure 2.10. Figure 2.11 is a diagram of our large Si detectors utilizing both ionization and phonon measurements. We are able to discriminate between nuclear and electronic recoils in the same way as the Berkeley Ge detectors, using a measurement of the total recoil energy and a measurement of the charge pairs produced. In addition, a comparison of phonon pulses in the four channels can give some x-y position

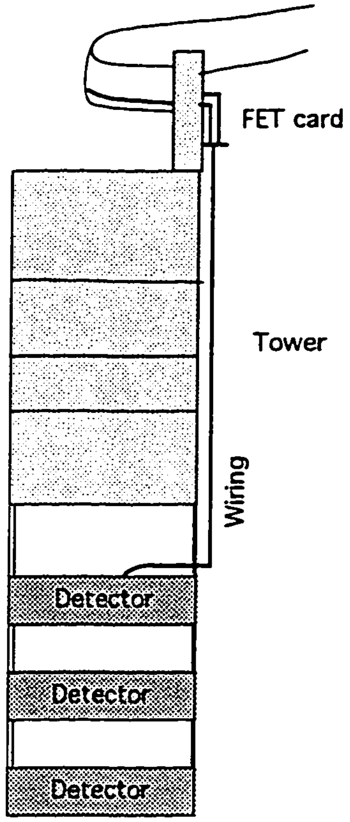

Figure 2.12 The Tower and FET Card.
The tower inserts into the icebox cans. Each tower can hold up to six detectors, including cold electronics, wiring from detector to 4 K, and FET cards for each detector. Each FET card holds four JFETs for the front end of our amplifiers. The FETs self heat to 120 K, their optimum operating temperature.

information. We are also working on developing these W/Al QET sensors for use with Ge crystals.

2.2.2 The Tower and FET card

Holding a fair number of detectors in a dilution refrigerator and connecting them to the outside world is not a trivial matter. We developed what we call "towers" to accomplish this. Each tower will hold six detectors, complete with wiring and cold electronics components. The towers insert into the cold environment independently of each other. Initially we will run at least two towers -- one for the BLIPs, and one for the Stanford detectors. We also plan to run another tower with detectors made from enriched ^{73}Ge and ^{76}Ge.

The tower holds the detectors at 10-20 mK. Wires go up the sides of the tower from the detectors to 4 K, and then through the FET card to the stripline. The FET card holds the front-end JFETs for our amplifiers. The FETs are mounted on a thin Kapton window, and self-heat to their optimum operating temperature of ≈ 120 K. The tower system (including the FET card) will be discussed in detail in Chapter 4. Figure 2.12 shows a rough outline of the tower system.

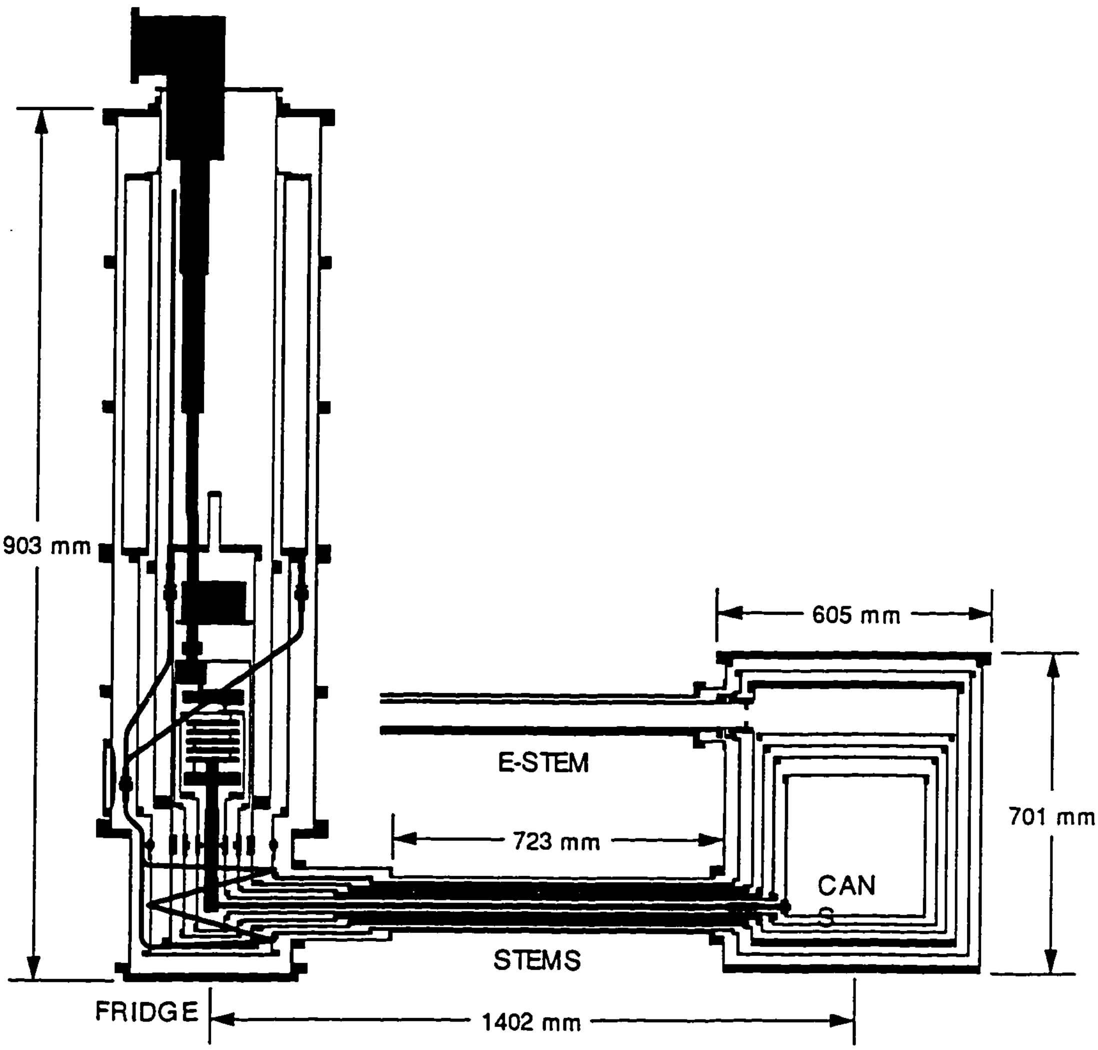

Figure 2.13 The Ice Box and Dilution Refrigerator.
We use an Oxford 400 dilution refrigerator with modified tails to cool the Icebox. The E-stem is rotated into the page ≈30 degrees to clear the refrigerator.

2.2.3 The Icebox

The detectors need to be operated in a low radioactive background, cryogenic environment. Dilution refrigerators are not easily made radio-pure. Common materials used in making dilution refrigerators and dewars such as stainless steel, aluminum, and fiberglass are not suitable materials from which to construct low background apparatus. We decided to therefore separate the dilution refrigerator, which is radio-impure, from our detector volume, which should be radio-pure. The Icebox accomplishes this goal.

Figure 2.13 is a diagram of the Icebox and dilution refrigerator. The Icebox is a set of nested cans made of radio-pure copper. The cans are connected to the refrigerator though a horizontal cold finger, the C-stem. Each temperature layer in the refrigerator is connected to a corresponding can in the Icebox. The icebox provides us with a cold volume big enough to hold seven towers with no internal shielding. When used at a shallow site, we will sacrifice some of the cold volume to hold a neutron moderator. In this configuration, we can hold three towers in the Icebox.

The Icebox also has a horizontal stem, the E-stem, that will carry striplines out from the 4K layer to room temperature. All electrical connections to detectors will be through the E-stem.

The Icebox/dilution refrigerator was first successfully run in September, 1994. After testing at Berkeley, the Icebox was moved to the Stanford Underground Facility and reassembled. The Icebox/refrigerator was again tested in July 1995 and demonstrated the capability to cool the heat loads necessary to run detectors in the Icebox.

2.2.4 The Stanford Underground Facility

To shield the detectors from cosmic rays, we will place our experiment underground. For the first phase of the experiment, we decided we needed an

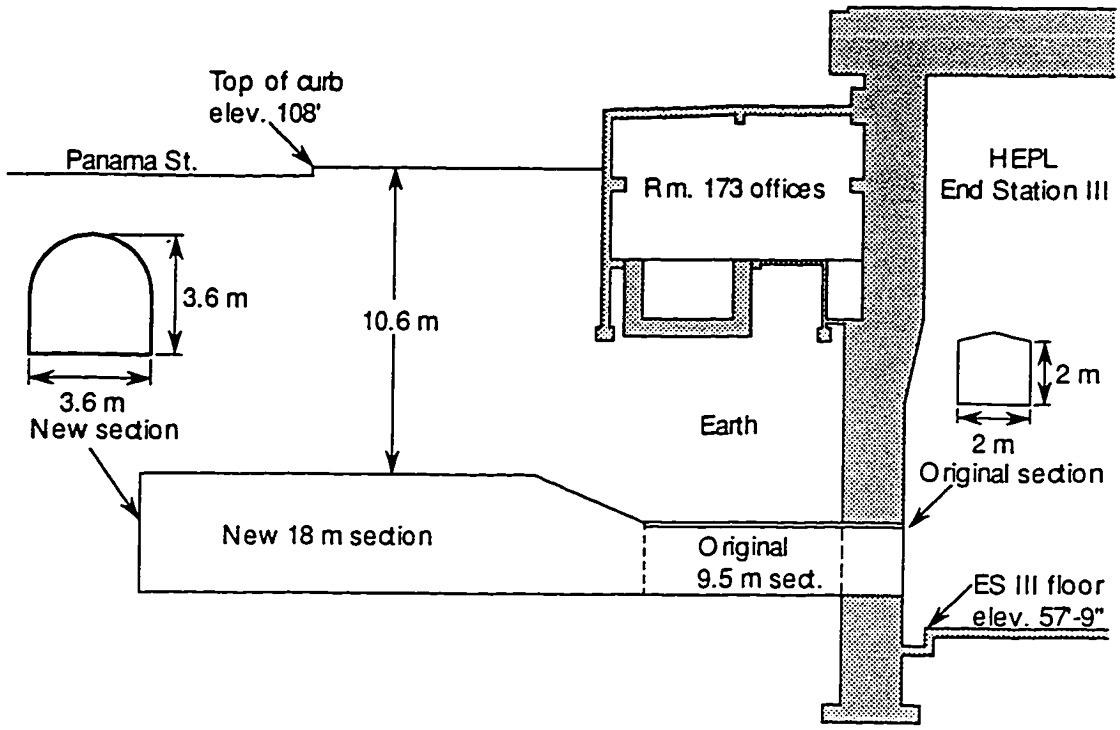

Figure 2.14 The Stanford Underground Facility (SUF).
We have extended one of the HEPL end stations to use as our underground laboratory.
SUF has an overburden of 16 m water-equivalent.

underground site that was close to our home laboratory to ease the inevitable period of troubleshooting. For this purpose, we constructed the Stanford Underground Facility by extending an existing tunnel at HEPL end station III. Figure 2.14 shows the layout of the SUF. The SUF has an overburden of 16 m water-equivalent. This is deep enough to shield the experiment from the hadronic component of the cosmic rays, but not deep enough to stop the cosmic ray induced muons.

The muon flux in SUF has been measured.[23] The muons will have two effects on the experiment. First, they will interact with the detectors directly, creating a background that is easy to identify, but still must be dealt with. Second, the muons will create neutrons near the experiment by spallation processes with nuclei in the tunnel. Unfortunately, neutrons interact with our detectors in the same way that WIMPS do, by nuclear recoil. Thus being at a shallow site has the disadvantage of subjecting our experiment to a neutron background that may sabotage the background discrimination

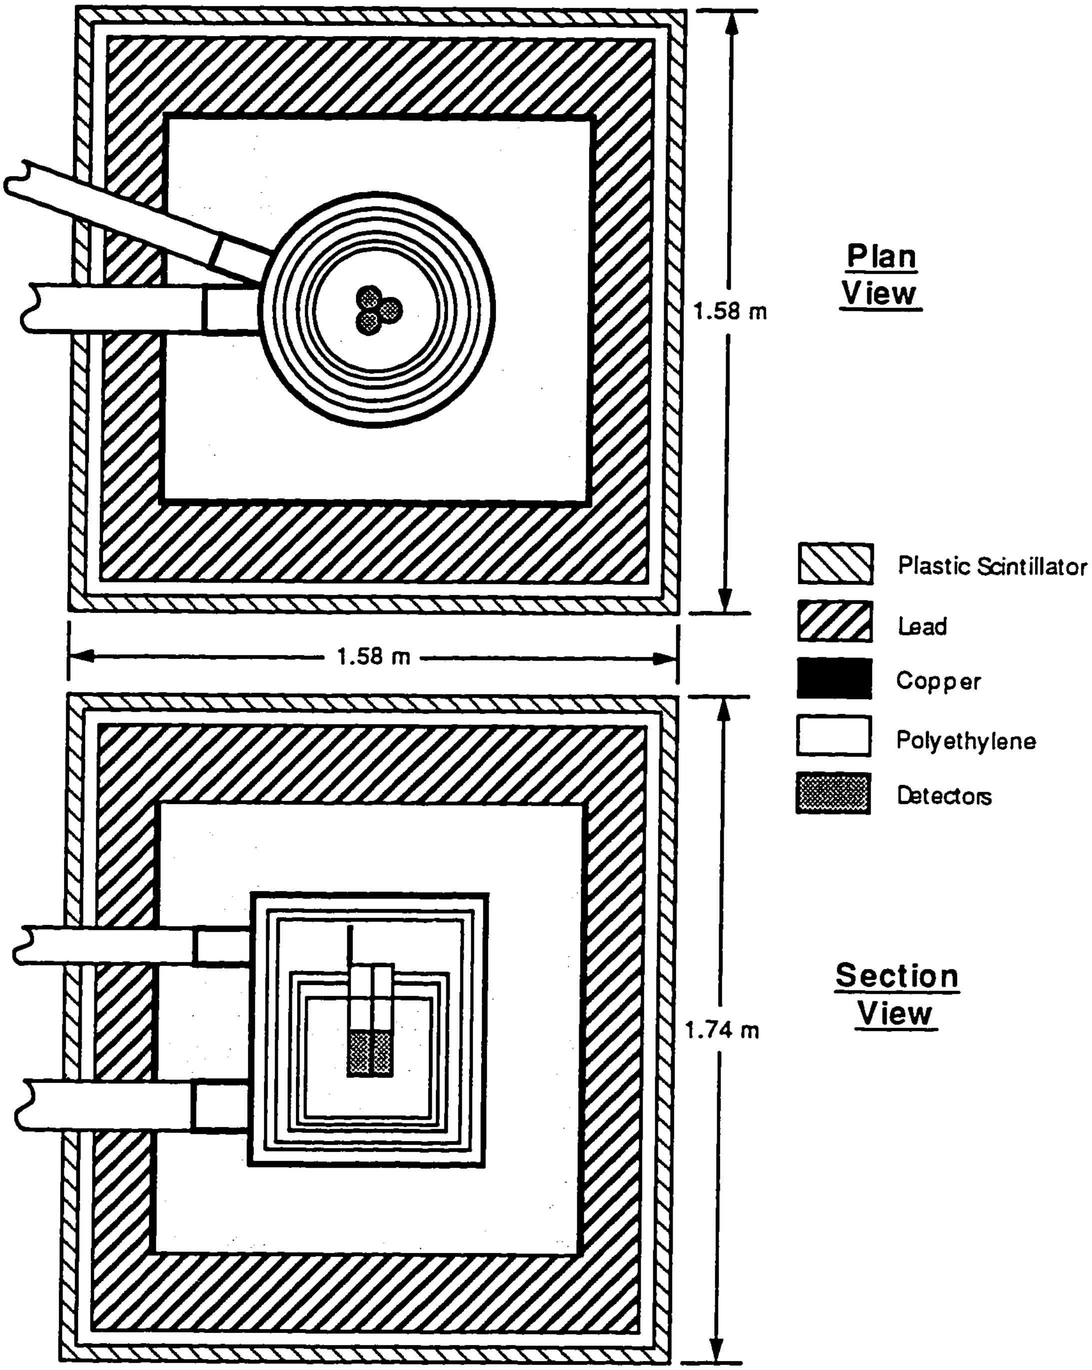

Figure 2.15 The Shield at the Stanford Underground Facility.
Plan and section views of the shielding system surrounding the cryostat. Inside the cryostat and surrounding three detector towers is ~10 cm of polyethylene to moderate the neutrons produced in the copper. Outside the cryostat is ~20–25 cm of polyethylene to moderate the neutrons produced in the 15 cm lead shield, which is surrounded by the 4 cm thick muon veto scintillators. The thick-lined region of the cryostat is surrounded by a magnetic shield.

capabilities of our detectors. The γ background is also much higher than we would have at a deep site.

An integral part of the SUF is the shield system designed to surround the icebox. Figure 2.15 is a diagram of our shielding strategy. The shield contains both an active muon veto and passive elements to shield the experiment from photons and moderate neutrons. The first element of the passive shield is a box of clean lead bricks enclosing that the Icebox. This shields the detectors from the gamma and x-ray background in the tunnel. Because of the muon flux, however, the lead bricks become a neutron source. The second passive element of the shield is polyethylene to moderate the neutrons that are created in and around the shield. We do not need to capture the neutrons for the simple kinematic reason that they don't transfer energy very effectively in a collision with the much larger germanium nucleus. A 20 keV neutron can deposit at most $\approx$1 keV in a single scattering event. Our goal with the polyethylene is to slow the neutrons so that the events they produce are below threshold in the detectors. To combat the effects of muon-produced neutrons, the lead box is covered on all sides by a scintillator veto. This allows us to mark events that occur in the detectors in coincidence with a muon passing through the experiment. We are in the process of exploring the efficiency of the muon veto, but we expect to greatly reduce the neutron background using this veto. The shield system has been extensively modeled and is described more completely in the Ph.D. thesis of Angela DaSilva.[23]

3. Detector Optimization: Interface Studies

The Berkeley group has built and tested a number of devices in an attempt to optimize both charge and phonon measurements with the germanium disk/NTD germanium thermistor detectors. We will focus on efforts to understand and improve the phonon measurement using NTD Ge thermistors. My own contribution in this effort has been an examination of the role of the eutectic interface between the thermistor and detector using the device P2. W. B. Knowlton built P2; he and I ran the detector to take the data in this thesis. This analysis of the data is my own. Unfortunately, the set of measurements presented here are somewhat incomplete. Time pressures from the larger experiment (CDMS) prevented us from following up on these measurements before the completion of this thesis.

3.1 NTD Germanium Thermistor Studies

The Berkeley group bases its phonon measurement on a change in resistance with a change in temperature of an NTD Ge thermistor. This measurement can in practice be very simple; we correlate a pulse height in our measurement with an amount of energy deposited in the detector. Optimizing this system, however, proves not to be so simple. We would like our thermistors to have the highest sensitivity to energy deposited in the absorber crystal possible. We want to increase the signal/noise ratio of our phonon measurement, which includes the detector, thermistor and amplifier chain. We are making a resistance measurement. Increasing the bias current will allow us to make a more accurate resistance measurement. Increasing the bias power will increase the temperature of our thermistors, which makes our thermistors less sensitve to the small temperature changes we are trying to measure. Our thermistor performance is partially described by the simple hot electron model, which postulates that the

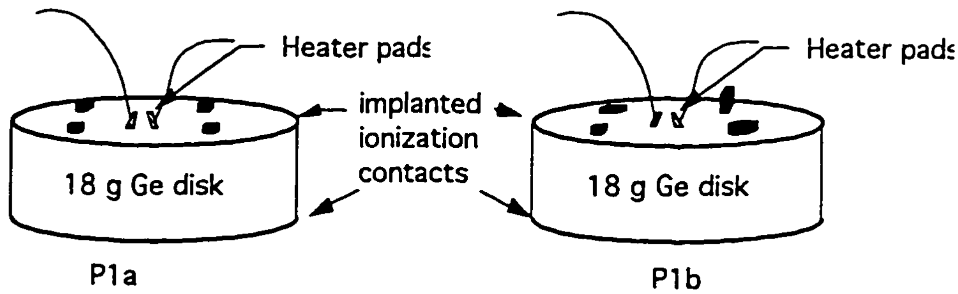

Figure 3.1 P1a and P1b.
P1a was designed to explore the effect of different amounts of doping on NTD Ge sensitivity. P1a was made with four thermistors, identical except they were made of different types of NTD Ge. P1b was designed to measure the effect of sensor geometry on sensitivity. The four thermistors on P1b were all NTD 29, but were of different volume and shape. Both detectors had implanted ionization contacts and gold heater pads.

electrons in our thermistor are not well thermally coupled to the phonon system of the thermistor and absorber lattice. The heating of the electrons with bias power is affected by the doping, size , and shape of the thermistor, and interface between the thermistor and detector.

We have built two devices, P1a and P1b, to better understand what parameters most affect thermistor performance. Figure 3.1 is a diagram of these devices. Both P1a and P1b have four NTD thermistors mounted on a 2 cm diameter by 1 cm thick single crystal disk of high purity Ge. The thermistors on P1a were all the same geometry, 1.72 mm^2 x 0.88 mm thick, and were similarly mounted, but were of different NTD Ge types (NTD 28, NTD 29, NTD 12, and NTD 23). Different NTD types correspond to different initial neutron doses and thus different dopant concentrations. The thermistors on P1b were all of NTD 29, but were very different geometries (2.23 mm^2 x 0.31 mm, 3.53 mm^2 x 2.85 mm, 4.35 mm^2 x 0.49 mm, and 8.50 mm^2 x 1.17 mm). In addition, both P1a and P1b were operated over a series of refrigerator runs with different amounts of heat sinking. All of the NTD Ge thermistors were attached to the crystal with our standard 1000 Å eutectic.[16] Both sides of the detector crystal were boron implanted

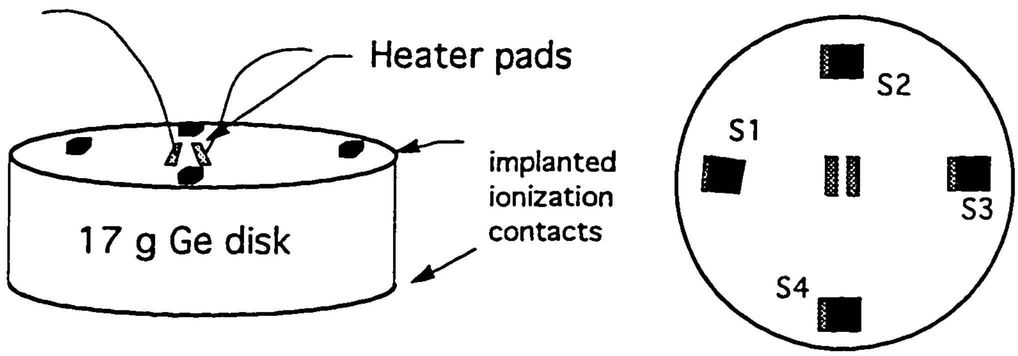

Figure 3.2 P2
P2 was constructed on the same size Ge crystal as P1a and P1b. P2 has four identical NTD-29 Ge thermistors only differing in the thickness of the eutectic used to bond the thermistors to the disk. S1 was deliberately misaligned by ≈8° from the disk crystal axis.

for ionization collection contacts. Both crystals had small gold pads (2.5 mm x 0.5 mm) deposited in the center of the top surface to use as heater contacts. The resistance between the pads was ≈ 100 Ω; by electrically pulsing the pads, we could inject a known amount of heat into the crystal. These detectors allowed us to explore the dependence of our thermistors' performance on doping levels, volume, and geometry.

In addition to optimizing the type of NTD Ge and the geometry of the thermistor, we can change the sensitivity of the phonon measurement by changing the coupling between the crystal and the heat sink. For a given device, less heat sinking leads to a larger pulse height at a given r_{bias}. This is also seen in our model. Less heat sinking has the disadvantage of creating longer pulses which leads to practical problems in using the detector in an experiment (such as pile-up). The phonon pulses typically have fall times from 10 ms to 100 ms depending on the exact heat sinking configuration of the device.

One last factor in our thermistor performance is the interface between the thermistor and the crystal. The interface can have two important effects on the performance of our detector. If the thermistor phonon system is not well coupled to the crystal phonon system, the pulse height will be reduced. More interesting is the possibility of coupling athermal phonons directly into the thermistor electron system.

For this to be possible we need an interface that is transparent to the athermal phonons. Coupling the energy of an event directly into the thermistor electrons, rather than the detector phonons, will increase our pulse height by heating only the heat capacity of the electrons initially. In addition, we may see differences that would allow us to distinguish between electron and nuclear recoils. Knowlton[24] has shown that the Ge-Au-Ge eutectic interface can be transparent to ballistic phonons. A 1000 Å eutectic showed 4% transmission of ballistic phonons, but a 500 Å eutectic demonstrated ≈50 % transmission and a 300 Å eutectic ≈ 70 % transmission using ≈ meV phonons.

We constructed a device, which we call P2 (figure 3.2), to study the effect of the interface between the germanium absorber and the NTD germanium thermistor. P2 is similar to the P1a and P1b detectors. The Ge crystal was the same size as those used for P1a and P1b, ≈ 17 g, 2.2 cm diam. x 1.0 cm thick. Four NTD 29 thermistors were attached, all 2.0 x 2.0 x 0.5 mm, the same size as sensor S3 on P1b (P1b:S3). The sensors were attached using a eutectic method developed by W. Knowlton, *et al.*, at Lawrence Berkeley National Laboratory following the initial tests of N. Wang[16].[25] An initial film of Pd is electron beam deposited on both mating surfaces, the thermistor and the Ge disk. Then the gold film is deposited on both the Ge crystal and the NTD Ge thermistor over the Pd. The Pd thickness is ≈10% the thickness of the Au film, and is used to help the Au film stick to the Ge. The thermistor is then pressed against the corresponding gold pad on the crystal, the whole assembly is heated above the Au-Ge eutectic temperature of 361 °C, and then slowly cooled over a period of 10 hours to promote epitaxial growth between the thermistor and Ge disk. Small test pieces of Ge were bonded first to test this method. The bonded pieces were cut across the joint, and the bond was examined with high resolution transmission electron microscopy. It was found that the Ge lattice was continuous across parts of the bond. The gold came out of solution below the eutectic temperature and formed gold precipitates in the region of the

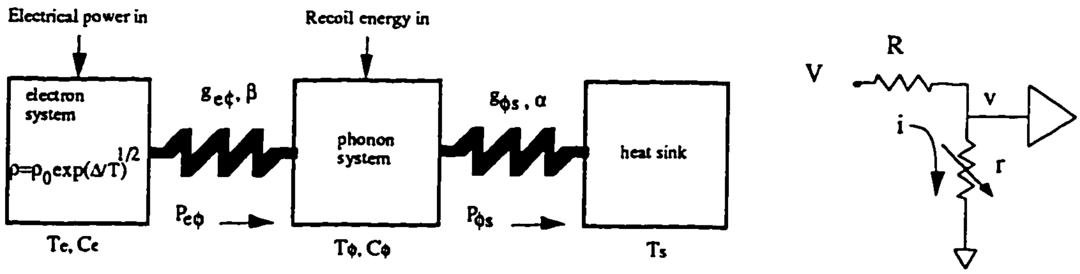

Figure 3.3 The Hot Electron Model
The left side of the figure shows how the hot electron model divides the thermistor thermal system into separate electron and phonon systems. The right side shows the thermistor bias circuit.

joint. For P2, two thermistors were attached using a 300 Å eutectic, one thermistor with a 500 Å eutectic, and one thermistor with a 1000 Å eutectic. (The "thickness" of the eutectic refers to the thickness of the gold films that are deposited onto both the crystal and the thermistor.) Eutectics made with 1000 Å gold on each surface show some penetration of the Ge lattice through the Au interface, but in the 300 Å eutectics the Au precipitates out in drops rather than forms a film. This leaves a considerable portion of the area of the interface as continuous Ge lattice.[24] One of the 300 Å eutectic thermistors had its crystal axes deliberately misaligned by 8° from the crystal axis of the Ge crystal. Misaligning the crystal axes between the thermistor and crystal has the effect of preventing epitaxial growth of Ge across the interface. A cut through a similarly rotated test piece showed a region of amorphous Ge ≈300 Å thick at the interface. For reference, we named the thermistors as follows: S1 = 300 Å eutectic, misaligned; S2 = 300 Å eutectic, aligned; S3 = 500 Å eutectic, aligned; S4 = 1000 Å eutectic, aligned.

3.2 The Hot Electron Model and P2

We describe the performance of our thermistors with the hot electron model. The hot electron model proposes a substantial thermal impedance between the electrons in the thermistor and the detector phonons.

3.2.1 Modeling the Thermistors with the Hot Electron Model

We can use the simple hot electron model (SHEM) to calculate the pulse height response of the thermistors to events in the detector. We start with a biased thermistor in an equilibrium state (see figure 3.3), so that

$$P_{ele} = P_{e\phi} = P_{\phi s} \tag{3.1}$$

where

$$P_{ele} = i^2 r, \tag{3.2}$$

i is the bias current and r is the biased thermistor resistance,

$$P_{e\phi} = g_{e\phi}\left(T_e^{\beta} - T_{\phi}^{\beta}\right), \tag{3.3}$$

and

$$P_{\phi s} = g_{\phi s}\left(T_{\phi}^{\alpha} - T_s^{\alpha}\right). \tag{3.4}$$

An event in the crystal changes the power balance in the electron-phonon system. We can calculate the response of the thermistor using the above equations and the relationship between energy and temperature in the detector,

$$E = CT \tag{3.5}$$

and its time derivative

$$P = C\frac{dT}{dt}. \tag{3.6}$$

These equations are highly non-linear; we solve the linearized case of the response to small changes δP around the equilibrium solution. The equations have been worked out in detail for pulses in Shutt.[26] The calculation is long, but fairly straight forward. We show the result for reference for recoil energies deposited entirely into the phonon system:

$$\delta v(t) = i\frac{R_{\parallel}}{r}\frac{\partial r}{\partial T}\frac{E_{dep}}{C_e}\frac{G_{\phi e}}{C_{\phi}}\frac{1}{s_1 - s_2}\left(e^{-s_2 t} - e^{-s_1 t}\right), \tag{3.7}$$

where $\partial v(t)$ is the voltage pulse we measure, E_{dep} is the event energy deposited into the phonon system, $R_{||}$ is the parallel resistance of the thermistor and bias resistor, and R is the bias resistance. We define the conductances $G_{e\phi}$ and $G_{\phi e}$ as

$$G_{e\phi} = \alpha g_{e\phi} T_e^{\alpha-1} \tag{3.8a}$$

and

$$G_{\phi e} = \alpha g_{e\phi} T_\phi^{\alpha-1}. \tag{3.9}$$

Similarly, we define $G_{\phi s}$ as

$$G_{\phi s} = \beta g_{\phi s} T_\phi^{\beta-1}. \tag{3.10}$$

For P2 (and most of our detectors), s_1 and s_2 at optimum bias are time constants close to our naive guesses for the time constants based on the heat capacities and conductances in the system,

$$s_1 \approx \frac{G_{e\phi}}{C_e}\left(1 - \frac{R-r}{R+r}\frac{P_{ele}}{G_{e\phi}}\frac{1}{r}\frac{\partial r}{\partial T}\right), \tag{3.11}$$

and

$$s_2 \approx \frac{G_{\phi s}}{C_\phi}. \tag{3.12}$$

When s1 >> s2 (this is the case for our detectors), equation 3.13 simplifies a bit to:

$$\delta v(t) = i\frac{R_{||}}{r}\frac{\partial r}{\partial T}\frac{E_{dep}}{C_\phi}\frac{G_{\phi e}}{G_{e\phi}}a\left(e^{-s_2 t} - e^{-s_1 t}\right). \tag{3.13}$$

If we examine equation 3.16 we can make a few observation about the pulse height. Increasing the bias current will increase the pulse height. A competing effect, however, can be seen in both $\partial r/\partial T$ and $G_{\phi e}/G_{e\phi}$. Increasing the bias current heats the electron system which decreases the signal by $G_{\phi e}/G_{e\phi} \propto (T_\phi/T_e)^5$. As important is the decrease in $\partial r/\partial T_e$, which can dominate at high bias. As the electron system heats significantly above the phonon system it becomes insensitive to changes in the phonon system. So we expect the pulse height of a given energy event to first increase with bias, and then to decrease with increasing bias as the electron system decouples. This simple model will give us pulses with two time constants, s_1 and s_2. In fact, our pulses have three time constants. We will consider a few extensions to the SHEM that might better describe our thermistors.

3.2.2 Extensions to the Hot Electron Model

3.2.2.1 Low Conductivity Interface

If the phonon systems in the thermistor and detector are not well coupled, we will have an extra heat capacity and coupling that will give us an additional time constant. Figure 3.4 is a representation of this model. In general, this will lower our pulse height for a given energy deposited in the detector phonon system. This is because the electron system will be hotter than in the SHEM and because any temperature change in the detector phonons must be filtered through the thermistor phonons before the electrons can respond. We usually simplify this model by recognizing that the heat capacity of the thermistor phonons is much less than the heat capacity of the detector phonons. In this limit, the low conductivity interface slows the rise time s_1, speeds the fall time s_2, and reduces the pulse height, all compared to the SHEM.

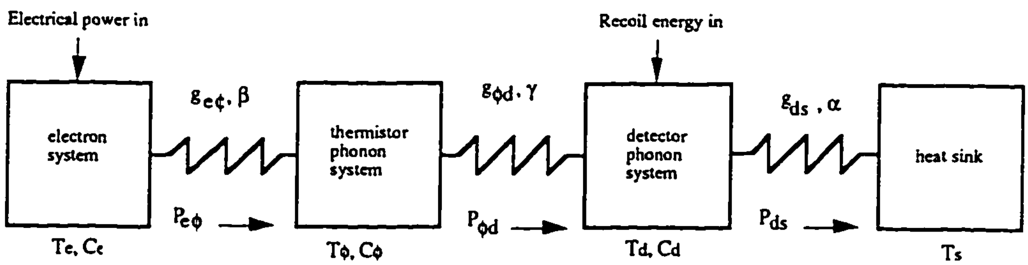

Figure 3.4 Low Conductivity Interface Extension to the Hot Electron Model.
This extension to the SHEM proposes a significant thermal impedance between the thermistor phonon system and the detector phonon system.

3.2.2.2 Hanging Heat Capacity

If we have a significant heat capacity on our detector weakly coupled to the detector phonons but not the electrons or heat sink, we will introduce another time constant into out pulse. Figure 3.5 shows schematically this model. This extra heat capacity could be, for instance, a region of gold evaporated on the surface of the crystal. The hanging heat capacity will have the effect of decreasing the pulse height, as some of the energy of the pulse must go into heating the extra heat capacity instead of the electron system. The time constants that actually appear in the pulse will be similar to s_1^{-1} and s_2^{-1} from the SHEM, and $s_x^{-1} \approx C_x/G_{xd}$ if the three times are well separated. We can simplify this model by assuming that the extra heat capacity and the phonon system are well coupled compared to the phonon sink coupling. In this case, once the hanging heat capacity heats up to the phonon temperature, the two systems will cool together.

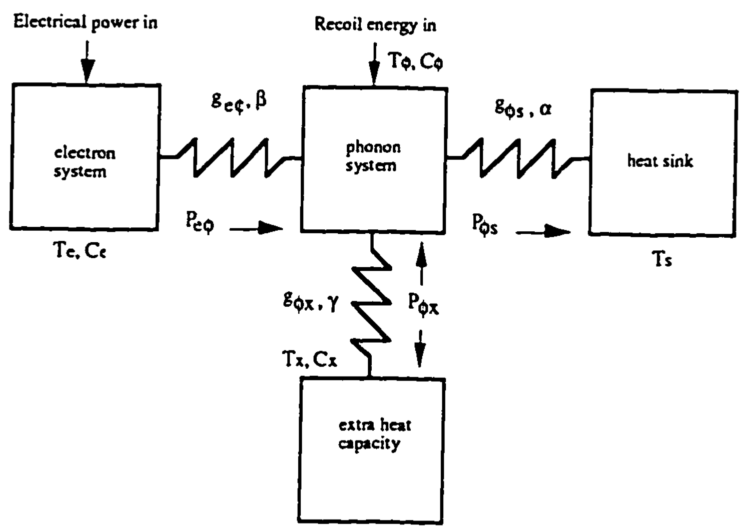

Figure 3.5 Hanging Heat Capacity Extension to the Hot Electron Model.
This extension to the SHEM proposes an extra heat capacity on the detector that is coupled to the phonon system.

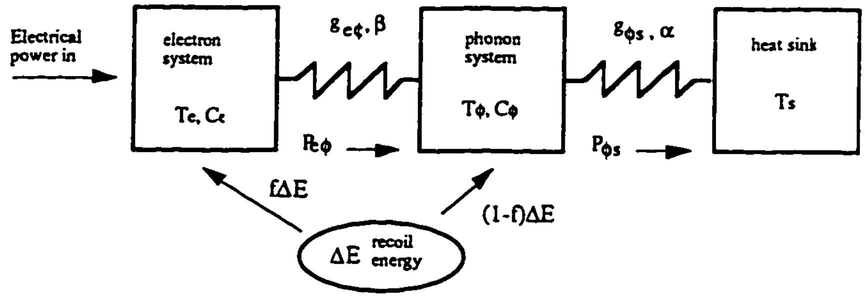

Figure 3.6 Collection of Athermal Phonons.
In this extension to the SHEM, a fraction fΔE of the initial recoil energy ΔE couples directly to the electron system. The rest of the recoil energy heats the phonon system.

3.2.2.3 Collection of Athermal Phonons

The previous hot electron model extensions have been concerned with changes in the detector/thermistor system itself, but all assumed the initial event energy is absorbed by the phonon system. Figure 3.6 represents an extension to the SHEM where we include the possibility of the thermistor electrons directly absorbing energy from the recoil phonons. An event in the detector will produce a spectrum of phonons with energies well above the distribution of the thermal phonons in the detector. Thin Ge-Au eutectic interfaces (< 500 Å Au per side) have been shown to have a transmission coefficient ≥ 50% for ballistic phonons. P2 was constructed with these thin eutectic interfaces, so we hope to couple these ballistic phonons into the thermistor directly. If we could couple the energy carried by athermal phonons directly into the electron system, we would gain a factor of $\approx (C_\phi/C_e)\cdot f \approx 4.5\cdot f$ in pulse height, where f is the fraction of the event energy that is absorbed directly by the electron system. Unlike the previous models discussed, this model will increase the pulse height above the pulse height predicted by the SHEM. This energy couples to the electron system with some time constant, s_r^{-1}, that will appear directly in the pulse shape. The fastest of the three time constants, s_1^{-1}, s_2^{-1} and s_r^{-1}, will be the rise time of the pulse. The slower two will be fall times.

The time constant s_r^{-1} is some characteristic time for the phonons to couple to the electron system. This might be a "reverberation" time for the phonons in the crystal, where the time constant is determined by the rate at which the athermal phonons "leak" out of the detector crystal and into the thermistor. In this case, we would expect to see

$$s_r \equiv \tau_r^{-1} \approx \frac{1}{4}\eta\frac{Av_s}{V}, \tag{3.14}$$

where η is the transmission coefficient of the thermistor/detector interface, A is the surface area covered by the thermistors, vs. is the phonon sound speed, and V is the detector volume. This results from considering the rate at which the energy density in the detector can leak into the thermistor. In the case that the phonons are not well thermalized by the NTD Ge thermistors, η must be multiplied by the efficiency of thermalization.

3.2.3 Phonon Sources

One of our goals in this study is determining if our thermistors can detect an athermal component to events in the detector. This convolves the initial spectrum produced by an event with any evolution of that spectrum (both spatial and in frequency) with the response of the detector. We use two different phonon sources in this study, heater pulses and photon events, and we should consider the possible differences in phonon spectrum produced by these methods.

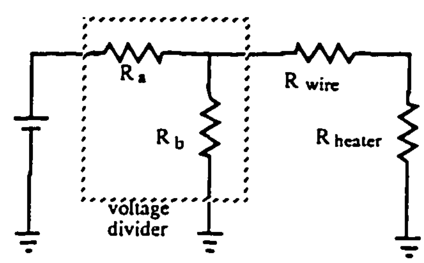

Figure 3.7 P2 Heater Bias Circuit.
We bias the heater through a voltage divider, R_a and R_b. The wire resistance was 12.5 Ω and the heater resistance was 39.5 Ω.

We inject heat pulses into the detector using the implanted face of the detector as a resistive heating element. We need to understand the bias circuit that we use to power the heater pads in order to accurately calculate the power deposited in the detector. We measured the cold resistance of the implanted face between the heater pads using a resistance bridge and a 4-wire measurement. We find the resistance is 39.5 Ω. We often used a voltage divider so that we could set a measurable voltage and divide it down before the heater pad. Figure 3.7 shows the heater bias circuit, and includes the voltage divider, wire resistances and heater resistance. We measure the voltage at the top of the divider; knowing all the circuit elements, we can calculate the power deposited in the heater pad. We used two 25 Ω wires in parallel to bias the heater (R_{wire} = 12.5 Ω.) Because the pulser pad and wire are low resistance (total resistance of 52 Ω), the voltage divider will not act as a stiff source. We used 10 kΩ and 100 Ω in the divider for a nominal 100:1 division. When we account for the low impedance of the heater however, we find the voltage at the heater is

$$v_{heater} = v_{bias} \frac{r_{heater}}{r_{wire} + r_{heater}} \cdot \frac{r_b || (r_{wire} + r_{heater})}{r_a + r_b || (r_{wire} + r_{heater})}, \tag{3.15}$$

where v_{bias} is the voltage applied at the top of the divider, v_{heater} is the voltage across the heater, and the resistances are as labeled in figure 3.7. We use the symbol "||" to denote adding two resistances in parallel. The power in the heater is then simply $(v_{heater})^2/r_{heater}$.

The top and bottom surfaces of the crystal are implanted with Boron ions to a depth of ≈ 2000 Å to create a degenerately-doped layer that we use for ionization measurement contacts. When we push a current through the heater, we are effectively creating a "black body" source of phonons as the power is dissipated in the surface layer. We can estimate the heater temperature using the hot electron effect. We have measured the electron-phonon coupling constant for NTD Ge, and we can use this to estimate the coupling constant for the heater by scaling the coupling by the ratio of the number of carriers in the heater to the number of carriers in the NTD Ge,

$$g_h = g_{e\phi}\frac{n_h}{n_{th}}, \tag{3.16}$$

where n_h is the number of carriers in the heater and n_{th} is the number of carriers in one of our thermistors. The number of carriers in the heater is the area of the heater multiplied by the boron ion dose used in implanting the contacts. The number of carriers in the thermistor is the volume of the thermistor multiplied by the dopant density. We then use equation 3.3 to relate the power dissipated and the effective heater temperature, just like the hot electrons in the thermistor. We find an effective g_h = $2.5 \cdot 10^{-3}$ W K^{-6}. We typically used 256 μV pulses across the heater, so that the power applied was 1.64 nW in a 10 μs pulse. With this power the heater should reach ≈ 93 mK. For a black body radiator at this temperature, less than 1% of the total heater pulse energy is radiated in phonons with temperatures greater than 1K.

A particle recoil event in the crystal should initially produce a very different phonon spectrum compared to the heater events. The interaction will initially produce phonons with a characteristic energy of order the Debye temperature of the lattice. This is a result of relaxation of excitations of both the lattice itself and electron-hole pairs formed by the interaction. These high frequency phonons quickly down convert to lower frequency phonons by scattering off defects in the lattice and by spontaneous decay. Cabrera, et al,[21] present both a Monte Carlo calculation and data from silicon based detectors demonstrating that a significant portion of the initial event energy can reach the surface of crystals in the form of ballistic phonons, at least for events close to the surface (25 μm) or in small devices (0.3 mm thick).

We can estimate the evolution of the phonon spectrum in the detector by considering the dominant scattering process for phonons as a function of phonon frequency. We will consider three processes for scattering and ultimately thermalizing the phonon distribution of an event (heater or particle.) These are spontaneous decay of phonons, defect scattering, and boundary scattering.

Spontaneous decay is a process where a high energy phonon decays into two phonons. This process is dominated by events where the initial phonon decays into two phonons with equal energy. The time constant for spontaneous decay of a phonon with frequency ω is given by[27]

$$\tau_{SD}^{-1} = \gamma\omega_D\left(\frac{\omega}{\omega_D}\right)^5, \quad \gamma \approx \frac{\hbar\omega_D}{M_0 v_s^2} \approx 10^{-3} \text{ for Ge.} \tag{3.17}$$

M_0 is the mass of the unit cell, vs. is the sound speed in the crystal, and ω_D is the Debye frequency of the crystal. Other phonon-phonon scattering processes will be slower than this because our detector is very cold; the other terms are proportional to $T^n\omega^{5-n}$, where T refers to the detector temperature.

Defect scattering will be dominated by isotopic scattering in our crystal. The rate for isotopic scattering can be estimated by[27]

$$\tau_i^{-1} = \eta\omega_D\left(\frac{\omega}{\omega_D}\right)^4, \quad \eta = \frac{3}{2}\frac{\langle \Delta M^2\rangle}{\langle M^2\rangle} \approx 3\cdot 10^{-3} \text{ for Ge,} \tag{3.18}$$

where $\Delta M = M - \langle M\rangle$, and brackets refer to averaging over isotopes.

Our detector is a finite size, and it may well be the case that phonons will encounter the surface of the detector before they have time to scatter by the other two processes. Boundary scattering can be described by[28]

$$\tau_b^{-1} = \frac{v_s}{L}, \quad L = \frac{1+p}{1-p}L_0, \tag{3.19}$$

where L_0 is some characteristic dimension of the crystal (in our case the diameter) and p describes the reflectivity of the surface of the crystal. If p = 1, then the surface is perfectly reflective, and the phonons never decay by boundary scattering. If p =0, then

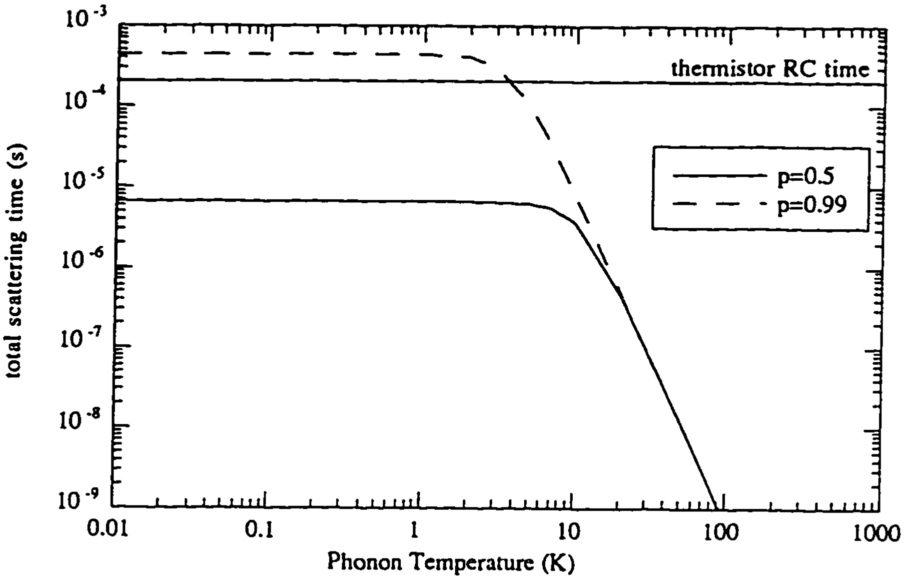

Figure 3.8. Total Phonon Scattering Time in P2.
Shown is the total scattering time in P2 as a function of phonon temperature. The p=0.5 curve refers to boundary scattering off the doped ionization contacts; p=0.99 refers to scattering off the thermistors only. The horizontal line is the approximate RC time of the thermistors when they are biased.

the phonons are completely absorbed at the surface and reemitted as thermal phonons. We don't know what p is for our detector. If we assume that the thermistors are the dominant thermalizers in the detector, we can estimate $p \approx 1\text{-}a/A$, where a is the surface area covered by thermistors and A is the total detector surface area. This gives us $p \approx 0.99$. If the implanted contacts are the dominant thermalizer, then $p = 1 - a/A = 0.5$.

We can calculate the total rate of decay as a function of phonon frequency by adding these time constants,

$$\tau_{TOT}^{-1} = \tau_b^{-1} + \tau_i^{-1} + \tau_{SD}^{-1}. \qquad (3.20)$$

We can see that total scattering rate will be dominated by the fastest rate. From the ratio of (3.3) to (3.4) we can see that $\tau_{SD}/\tau_i \approx 1000/T_\phi$. For phonons below $\theta_D = 320$ K, $\tau_{SD} >> \tau_i$, and we will not notice the effect of spontaneous decay. Figure 3.8 shows the total scattering time as a function of phonon temperature (where we use temperature as a label of phonon energy.) We also show the approximate RC time of the thermistors at their bias point, ≈ few 100 μs. We see that if $p = 0.5$ (thermalizing on the implanted contacts) that the entire phonon spectrum will be efficiently thermalized before our detector can respond. If $1 - p < 0.01$, then our detector has a chance at seeing an athermal component to the phonon spectrum.

3.3 Characteristics of P2

P2 was run before we modified the Oxford 75 at Berkeley, so P2 was mounted in the old-style detector holder in a light-tight can screwed directly into the mixing chamber. P2 was held in the detector holder using spring loaded sapphire tipped screws. The device was heat sunk through the top of each thermistor. The top surface of each thermistor was covered with a gold film, and a 2 mil diameter Au wire was bonded to the film. The other end of each wire was bonded to a heat sink Cu-Ka pad on the detector holder. We expect the conductance $G_{\phi s}$ to be dominated by the Au-Ge interface on the top of the thermistors. This method of heat sinking the crystal through the thermistors is different from the method we use on our later detectors, where we heat sink through a gold pad sputtered directly onto the detector crystal.

The data presented has a number of limitations. We were limited to a single refrigerator run due to severe time constraints on available refrigerator time. The detector was heat sunk in a non-standard way, through the thermistor instead of the crystal. Ideally, we would change the heat sinking to our standard method and compare with this data to see if this has an effect. More importantly, the thermistor bias wire for S4 was open. S4 is the thermistor with the eutectic interface identical to the interface we use on the rest of our detectors. Without S4 we cannot directly compare the results with the thinner interfaces to our standard 1000 Å eutectic.

Before we can compare the sensors on P2 to our models, we must know the material parameters contained in the models. The SHEM contains a number of thermistor and detector parameters, all of which we can measure. The resistivity of the NTD is characterized by two parameters, ρ_0 and Δ; the power flow between the electrons and phonons by $g_{e\phi}$ and β; the power flow between the phonon system and the heat sink by $g_{\phi s}$ and α. These parameters are enough to describe the NTD Ge I-V curves and any other static measurement. To calculate the NTD Ge response to pulses, we need in addition the electron and phonon heat capacities, C_e and C_ϕ, and the electrical characteristics of the thermistor and bias circuit.

3.3.1 Δ and ρ_0

The first useful measurement of the thermistors is calibrating their zero-bias resistance vs. temperature. This will allow us to measure the hot electron model parameters ρ_0 and Δ that describe the variation of the thermistor resistivity with electron temperature; also, we then can use the thermistors as thermometers on the detector

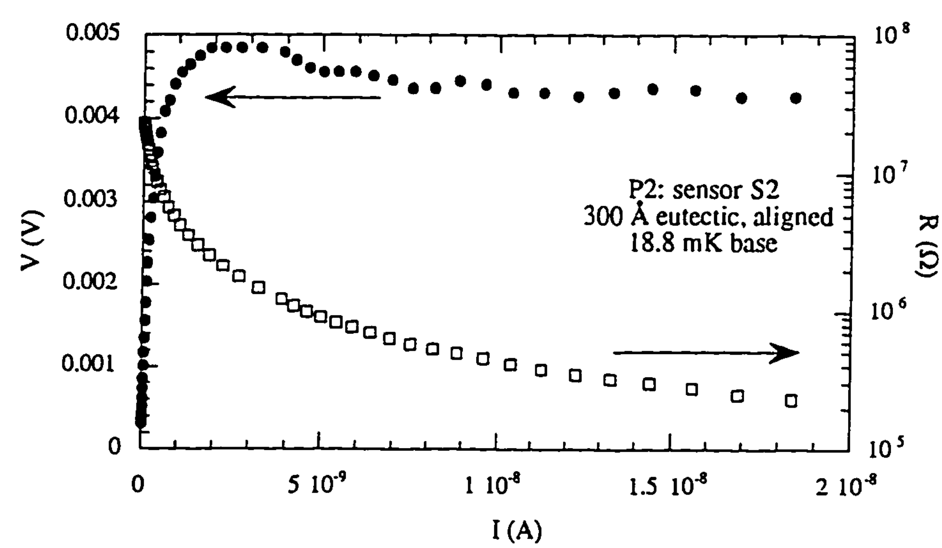

Figure 3.9 I-V Characteristic for S2 on P2 at 18.8 mK.
The left axis shows the voltage across S2 as a function of bias current; the right axis is the resistance, R = V/I or the thermistor at each bias point. NTD Ge thermistor have extremely non-linear I-V curves as the electron system heats above the base temperature.

crystal. Figure 3.9 shows the I-V characteristic of sensor S2 at the refrigerator base temperature of 18.8 mK. We can calculate the resistance of each point on the I-V curve as $R = V/I$. We notice that R is greatly dependent on the bias current -- this is because the electron system heats with increased bias power. To calibrate the sensors we vary the refrigerator temperature, T_s, and measure the zero-bias resistance of the thermistor. We are aiming to measure the resistance of the thermistors with a low enough bias power that the assumption $T_e = T_s$ is good. In practice, we try to set the measurement bias as low as possible considering the noise in the voltage and current measurement. We can check that we are not significantly heating the thermistor by increasing the bias. If the measured resistance does not significant change, we are making a good measurement.

We operated P2 during one refrigerator run, Run 101; the bias resistor of S4 (1000 Å eutectic) was open, so we have no data for that sensor. We calibrated R(T) for the other three thermistors. Figure 3.10 shows the measurements and fits for S1, S2, and S3. Table 3.1 shows ρ_0 and Δ for the three sensors. One interesting result is that the parameters ρ_0 and Δ for all three of thermistors are in good agreement with each other. Because the two parameters are correlated in the fit, the spread in ρ_0 and Δ is a bit artificial (S1 has the lowest ρ_0 and highest Δ, while S3 has the lowest Δ and highest ρ_0.) If we fix the resistivity to the average value $\overline{\rho}_0 = 7.80 \cdot 10^{-3}$ Ω-cm, the quality of the fits are not affected, and we find $\Delta = 6.99$, 6.89, and 6.97 for S1, S2, and S3 respectively. The thermistors on P2 are remarkably similar in terms of ρ_0 and Δ. Previously, we have attributed difference in ρ_0 and Δ in different thermistors of the same type NTD Ge to stresses induced in mounting the thermistors to the crystal. This Au-Ge eutectic mounting method minimizes the stress induced in the thermistor, or at least makes the stress very reproducible. This is a desirable quality when we are designing a detector.

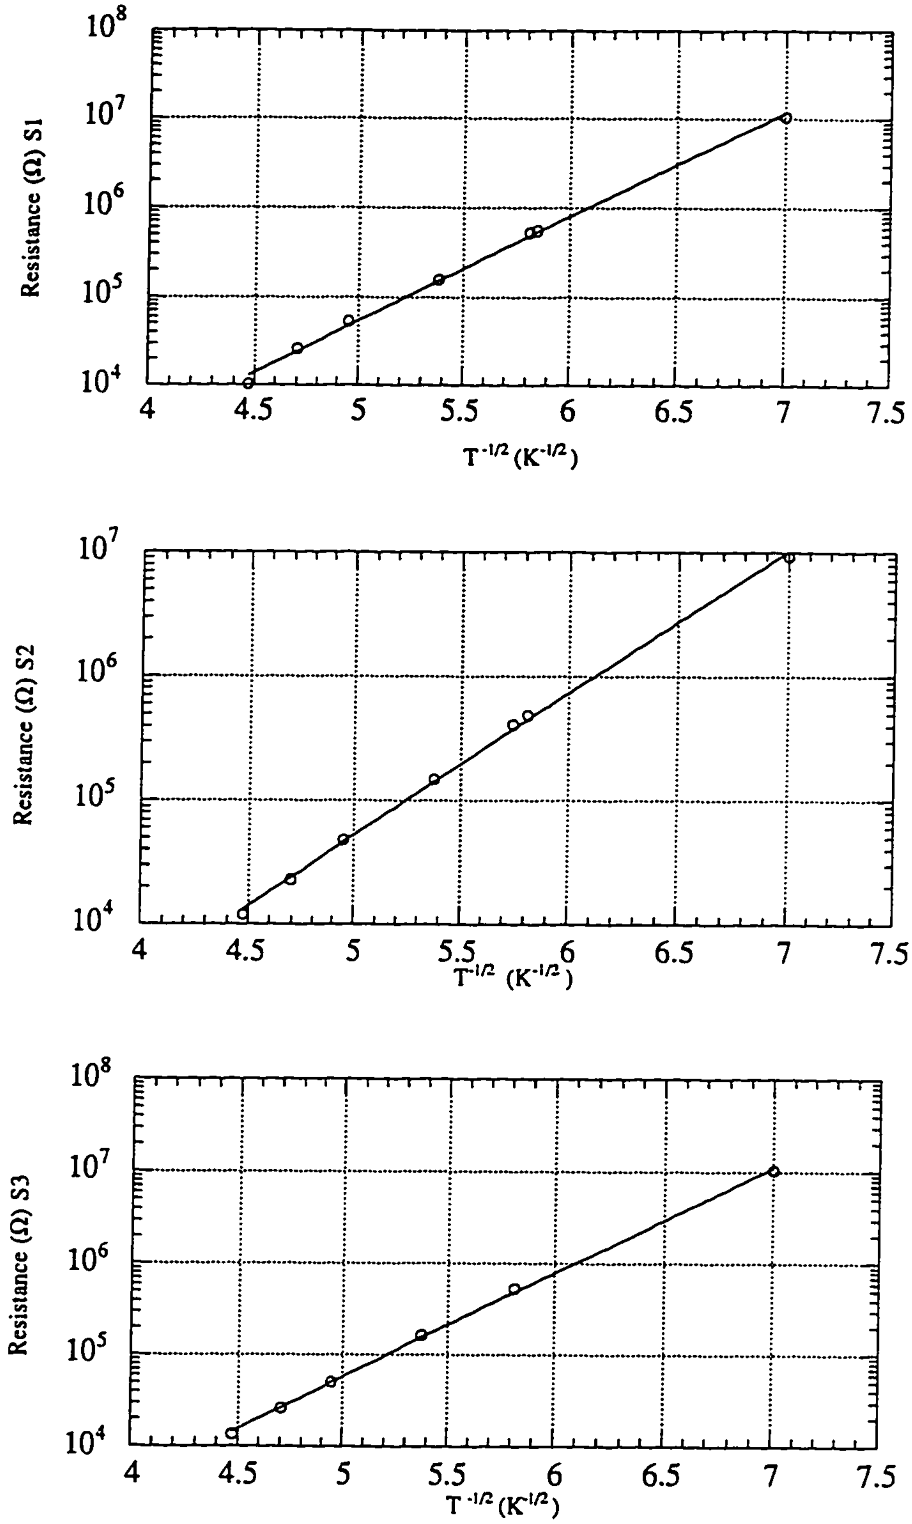

Figure 3.10 R vs. T Measurement for S1, S2 and S3 on P2
Shown is the zero-bias resistance versus refrigerator temperature for sensors S1, S2, and S3 on P2. The fit shown is $R = R_0 \exp(\Delta/T)^{1/2}$.

Table 3.1 Measured ρ0 and Δ for S1, S2 and S3

Sensor	ρ0 (Ω - cm)	Δ (K)
S1	$6.14 \cdot 10^{-3}$	7.25
S2	$7.92 \cdot 10^{-3}$	6.96
S3	$9.33 \cdot 10^{-3}$	6.89

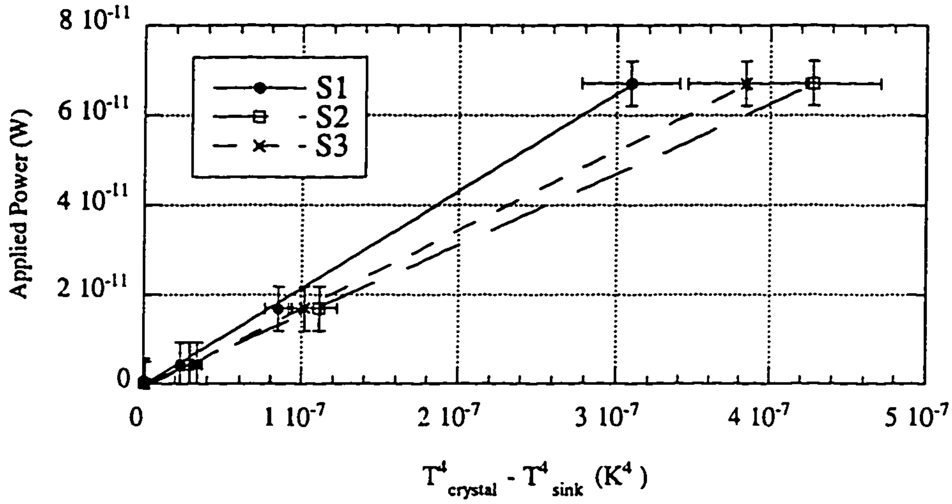

Figure 3.11 Measurement of the Phonon - Sink Coupling.
We heat the crystal through the heater pad and monitor the temperature of the crystal and the refrigerator. The refrigerator temperature did not change significantly during this measurement.

3.3.2 The Phonon - Sink Coupling: $g_{\phi s}$ and α.

Once we have the thermistor R(T) calibrated for each sensor, we can use the thermistors as thermometers on the detector crystal. We use the heater pad on the crystal to deposit a known amount of constant power into the crystal, and use the thermistors to monitor the temperature of the detector. We use this to measure the phonon-sink coupling, by fitting the applied power to the assumed coupling for heat sinking through the gold pads on the thermistors. Figure 3.11 shows this measurement made at the base refrigerator temperature of 18.8 mK. The temperature dependence of the heat sinking fits well the expected $\alpha = 4$. The three sensors find an average value of $g_{\phi s} = 1.8 \cdot 10^{-4}$ J K^{-4}, which is low compared with the predicted value of $4.7 \cdot 10^{-4}$ J K^{-4}.

The predicted value comes from scaling $g_{\phi s}$ measured on E2 by the ratio of heat sink areas on P2 to E2. This measurement could be improved by measuring to a greater value of $(T_{crystal}^4 - T_{sink}^4)$, using higher heater power. This is desirable considering that P2 was heat sunk differently than our other detectors.

3.3.3 The Electron - Phonon Coupling: $g_{e\phi}$ and β.

Once we have calibrated the sensor R(T), we can use the measured I-V curves to find the electron - phonon coupling parameters $g_{e\phi}$ and β. At each point on the I-V curve, the power flowing into the thermistor electron system is I·V. This power must flow into the phonon system and then to the sink. At each point we know R = V/I, and we can change this into T_e, the temperature of the electrons as they heat. The best way to make the measurement is to heat one thermistor while using another to monitor the temperature of the crystal. We did not do this, so we will estimate the temperature of the phonons using the known power into the thermistor and the measured phonon-sink coupling. Figure 3.12 shows the measurement of $g_{e\phi}$ made with S2 at four different base temperatures. The data fits well the expected $\beta = 6$. The measurement gives a value of $g_{e\phi} = 0.11$ J K^{-6} which is a factor of two higher than the measurement of $g_{e\phi}$ for NTD 29 made by Aubourg, *et al.*[29] This measurement with P2 is sensitive to the measurement of $g_{\phi s}$. The lower-than-predicted measured value of $g_{\phi s}$ makes the quantity $T_e^6 - T_\phi^6$ smaller at a given point, which leads to a larger measured $g_{e\phi}$. As stated above, the preferred method would be to monitor the phonon temperature by measuring the temperature of another thermistor at low bias at each point of the I-V curve.

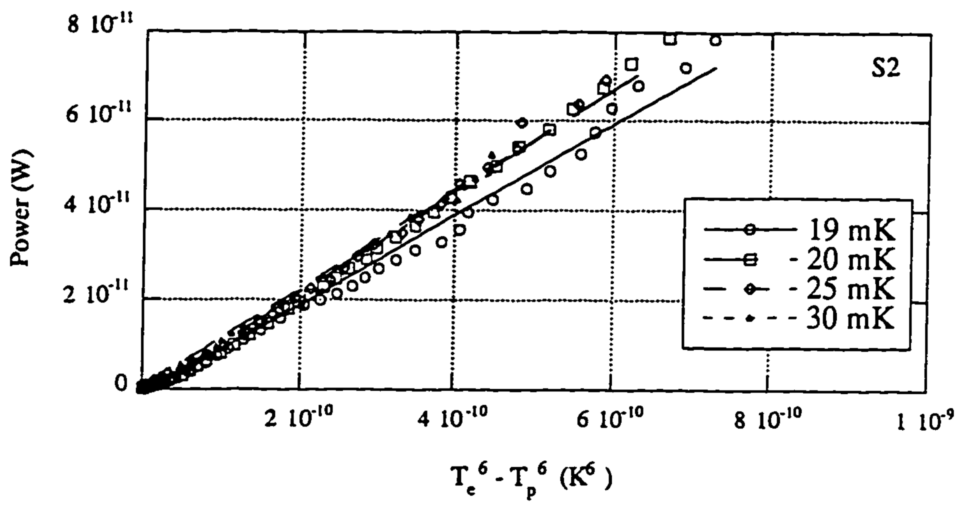

Figure 3.12 Measurement of the Electron-Phonon Coupling.
Shown is the measurement of the electron-phonon coupling of sensor S2 at four different base temperatures.

3.3.4 Heat Capacities: C_e and C_ϕ.

Calculating the dynamic behavior of the hot electron model requires the heat capacities of the electron and phonon systems, C_e and C_ϕ. The heat capacities cannot be measured statically -- we must measure them dynamically. We did not attempt a measurement of the heat capacities of P2. Where necessary in calculations of pulse height, we will assume the value of C_e published in Auborg, et al.[29], and a value for C_ϕ as given in Shutt.[30]

3.4 Tests of the Hot Electron Model with P2

Our phonon measurement relies on being able to measure voltage pulses on our sensors. To first order, we can improve the sensitivity of our phonon measurement by increasing the pulse height response of our sensors to a given energy deposition. Studying the pulse height response of our sensors under a variety of conditions (base temperature, biasing conditions, geometry, dopant level) is therefore very important to the success of our detectors. We will compare our measurements to the SHEM, and to our extensions to the SHEM where appropriate.

3.4.1 Pulse Height

We injected heat pulses of 106 keV into the detector crystal through the heater pads and measured the resulting event pulses in the thermistors. We applied a 100 mV, 10 μs pulse at the top of the heater voltage divider, giving us 106 keV deposited in the crystal for each pulse. We used a synchronous signal from the pulser to trigger the data acquisition hardware.

We measured the pulse height on S1, S2, and S3 for a series of bias currents, giving us the pulse height as a function of sensor resistance. The pulses were read out using our voltage amplifiers[31] through 40 kHz low pass filters to cut down high frequency noise and aliasing in the digitization. We averaged 20 events for each heater pulse height measurement. We naively expect the pulse height to increase with increasing bias current. At large bias, however, the electron system will be significantly hotter than the phonon system, and will become insensitive to changes in the phonon temperature.

Figure 3.13 shows the pulse heights measured at a series of bias powers and compares it to a calculation based on the SHEM using the parameters we measured on P2. We see the model does a fairly good job predicting the pulse height of heater events in sensor S1, but is low compared to the measured pulse height in S2 and S3. S1 has a lower pulse height than the other two sensors at all biases; this is presumably due to the fact that S1 was deliberately misaligned from the detector crystal axis. This could mean the phonon systems in the thermistor and crystal are not well coupled, which would decrease the pulse height. We can understand this as the phonons in the

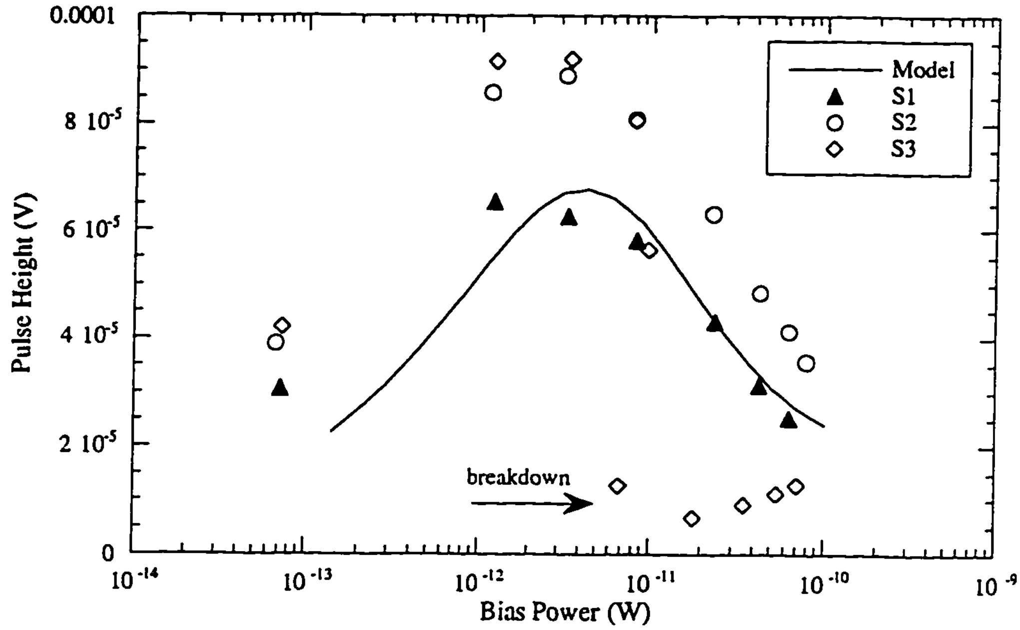

Figure 3.13 Comparison of Measured Pulse Heights to Calculation.
Shown is pulse height vs. bias power for S1, S2, and S3 on P2 and the pulse height calculated using the hot electron model. S1 is the sensor misaligned to the detector crystal axis. S3 suffered "breakdown" at high bias.

thermistors decoupling from the phonons in the crystal in much the same way as the electrons decouple from the phonons in the thermistor. Now a $\partial T_{crystal}$ must be "filtered" through a ∂T_{ϕ} to be sensed as a ∂T_e. This argument is not too compelling as the SHEM model fits the PH vs. bias data for S1 without a decoupling between the phonon systems, yet fails to fit S2 and S3 which should be well coupled. The lower pulse height in S1 vs. S2 and S3 could mean the additional pulse heights in S2 and S3 represent a significant athermal phonon collection compared to S1. Misaligning the S1's crystal axes from the detector axes has the effect of preventing epitaxial growth of the Ge at the interface between the thermistor and detector. Instead, the interface is a layer of polycrystalline Ge, ≈ 200 - 300 Å thick. This should greatly decrease the transmission of athermal phonons through the interface.

S3 suffered from "break down" at high bias. The thermistor can not be stably current biased because of negative dynamic resistance at high bias, and the thermistor changes to a state with a much lower resistance than we would otherwise expect for a given bias. We notice that before S3 breaks down, S2 and S3 look very similar. We should note that most of our detectors are constructed using the 1000 Å eutectic, which we were unable to test in this run of P2. It could be possible that the thicker eutectic would make a significant thermal impedance between the thermistor and detector or prevent transmission of athermal phonons into the thermistor, leading to a low pulse height, like S1.

We can compare the pulse heights measured in P2 to the pulse heights measured in P1b to begin to understand if a thick 1000 Å interface does indeed affect the pulse height. The germanium absorber used to construct P1b has the same dimensions and mass as the absorber on P2. In addition, all of the sensors on P1b and P2 were made from NTD 29. P1b:S3 (S3 on P1b) is the same dimensions as the sensors on P2. We would then expect the pulse height response to a given energy deposited in the crystal to be the same for P1b:S3 and P2:S2 and P2:S3 if the interface thickness has a negligible effect. Assuming the pulse height response of these sensors is linear in this energy range (we believe this to be true for our particle detectors E2 and E5), we can scale the pulse height measured by the energy of the heater events used to directly compare events between P1b and P2. This direct comparison might be complicated by the different heat sinking of the two detectors (P2 through the thermistors, P1b through the detector) and the presence of two large thermistors on P1b that make a significant contribution to P1b's heat capacity.

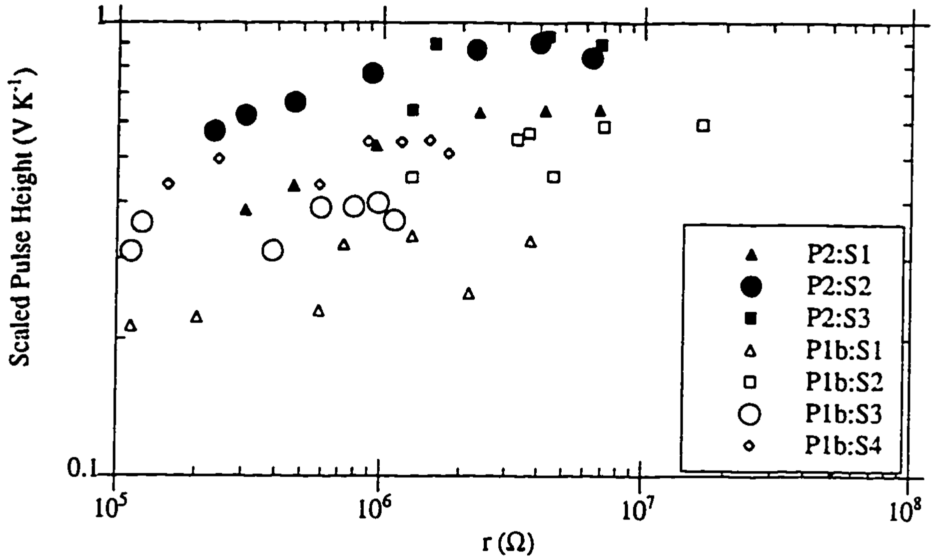

Figure 3.14 A Comparison of Scaled Pulse Height for Heater Events in P1b and P2. Shown is the pulse height measured scaled by the energy of the heater event and total heat capacity of the detector as a function of the resistance of the sensor at the bias point. The detector heat capacity includes the phonon heat capacity and the sensor electron heat capacity. Sensor P1b:S3 is the same size and NTD type as the P2 sensors.

Figure 3.14 is a comparison of the pulse heights measured scaled by the event energy and the total heat capacity for P1b and P2. The total heat capacity of each detector includes a contribution from the Ge phonon system and from the sensor electron systems. The total heat capacity of P1b is dominated by the electronic heat capacity of its sensors. We have not included the heat capacity of P1b:S2 in the pulse height scaling for P1b because it was glued to the detector crystal, and was too slow to affect the pulse height of the other sensors. The heat capacity of the sensor electrons and the crystal phonons are roughly equal in P2. We see in figure 3.14 that the scaled pulse height is a factor of ≈1.5 higher in P2:S2 than in P1b:S3. This is an indication that a thick 1000 Å interface may indeed be degrading the pulse height response of our detectors. This result is not firm; we used the nominal value of $c_e(18\text{ mK}) = 10^{-11}$ J K^{-1} mm^{-3}, where c_e is the electronic heat capacity per mm^3 for NTD 29 measured for

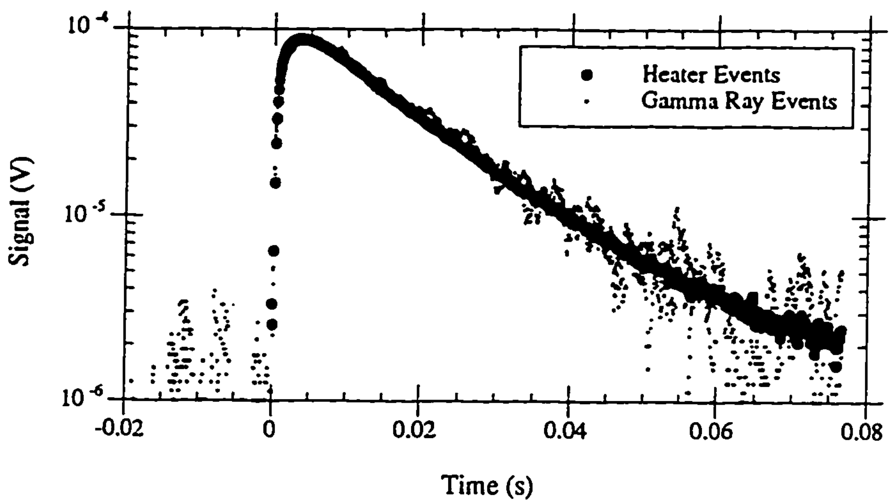

Figure 3.15 Comparison of Heater and Gamma Events.
Shown is a comparison of an average of twenty 104 keV heater events with three gamma ray events of comparable energy. The pulse shape is the same for the two classes of event. The events shown are from S2 at 20.4 mK.

sensors on P1a and P1b.[29] The scatter in values measured for the large sensors on P1b vary by about 20%. Because the heat capacity of P1b is dominated by the NTD Ge sensor electrons (particularly the large sensors), the scaled pulse height shown is only accurate to the same accuracy as the heat capacity.

3.4.2 Pulse Shape

Gamma ray events of comparable energy to the heater events were measured. We biased the charge collection contacts and triggered using the fast (μs) charge pulses to measure the slower (ms) phonon pulses. The gamma event pulses were filtered and measured the same way as the heater pulses. We cut on the pulse height of the event in S2 to find gamma events of comparable energy to the 104 keV heater events. We used the previous measurement of pulse height vs. bias for heater events, and defined a 2 μV window centered on the pulse height measured with the heater. We collected 20 clean (with no pileup) gamma events at each bias point. The gamma rays were presumably from Compton scatters of background sources in the room such as ^{40}K. The heater

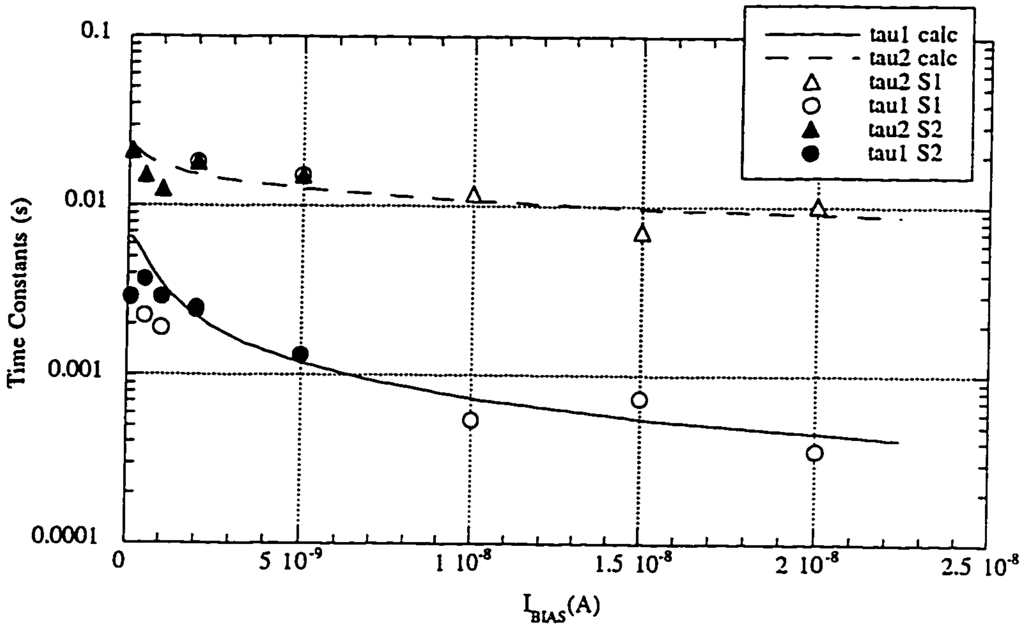

Figure 3.16 Comparison of Measured and Calculated Time Constants in P2.
Shown is a comparison of measured and calculated time constants for sensors S1 and S2 on P2 for 104 keV heater events at 20.4 mK as a function of bias current. Tau1 is the rise time, tau2 is the long fall time of the pulse. A third time constant is seen in these events, another fall time faster than tau2.

events and gamma ray events look identical. Figure 3.15 shows an average of 20 heater events with three gamma events of comparable energy overlaid in S2. The pulse heights have been normalized to 1 to directly compare the pulse shapes.

3.4.3 Comparison of Pulse Shapes to the Hot Electron Model

We can use the hot electron model to calculate the pulse shape of events. We use the measured conductances in P2 and measured properties of NTD-29 with equation 3.13 to calculate a pulse v(t). Figure 3.16 shows a comparison of a measured and calculated time constants in sensors S1 and S2 for heater events at 20.4 mK. Tau 1 is the rise time of the pulse, tau 2 is the long fall time of the pulse. There is a third time constant seen in these pulses, another fall time faster than tau 2. Figure 3.17 shows this fast fall time in S2 as a function of bias for two detector temperatures. In the calculation of the rise time and the long fall time, we have assumed a value of 10^{-11} J K^{-1} mm^{-3} for the heat

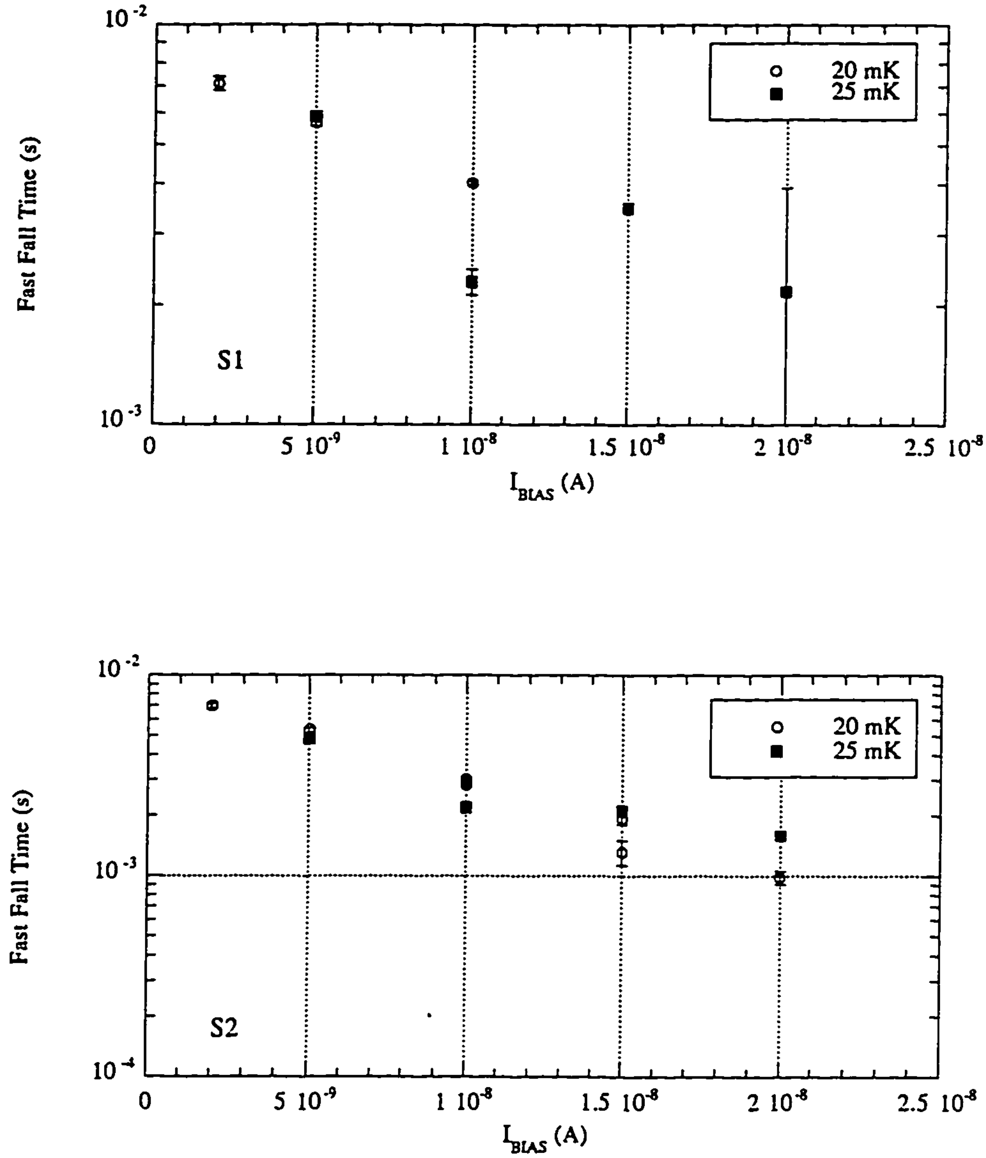

Figure 3.17 Fast Fall Time in S1 and S2.
Shown is the fast fall time in S1 (top) and S2 (bottom) for two temperatures. This fall time is not predicted by the simple HEM calculation.

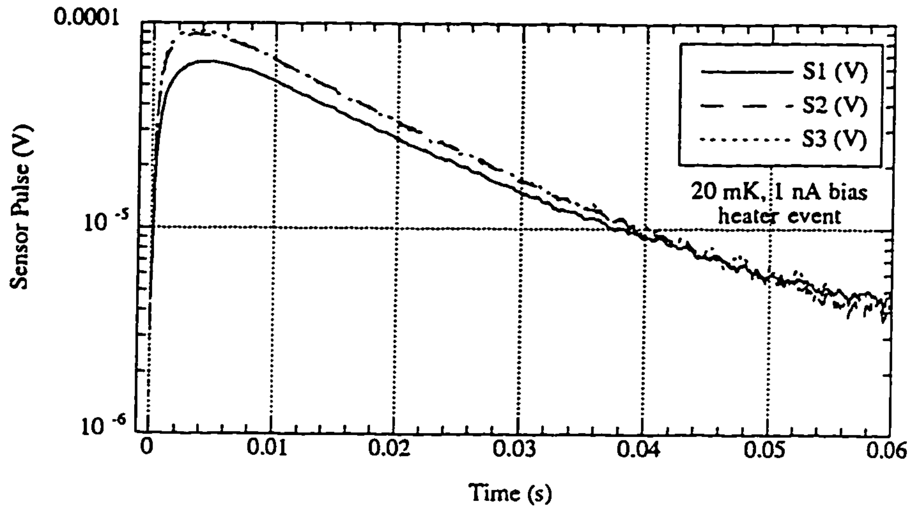

Figure 3.18 Comparison of Heater Pulses in S1, S2 and S3.
We show pulses measured simultaneously in S1, S2, and S3. S1 is markedly lower in pulse height than S2 and S3.

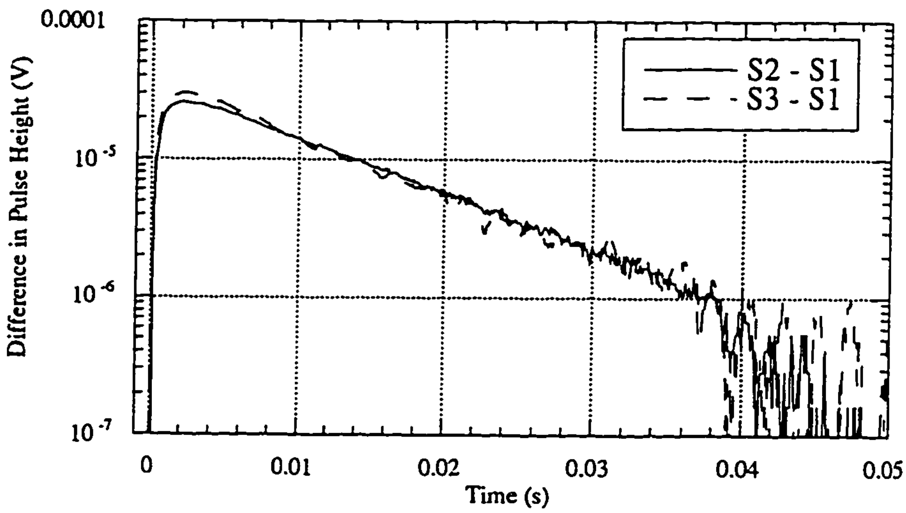

Figure 3.19 Difference in Pulse Height Between S2 and S3, and S1.
This shows the difference in pulse height between S1 and the two sensors S2 and S3. The fall time of the resulting difference is ≈ 10 ms for both.

capacity per mm^3 of the electrons in the thermistors at $T = 18$ mK. The total detector heat capacity is the heat capacity of the phonons in the germanium disk, the heat capacity of the lattice of four NTD Ge thermistors, and the heat capacity of the electrons in S4, the thermistor not being measured or biased. We assume the electrons in the unbiased thermistor are at the temperature of the phonons in the detector. We see the model fits both the rise time and long fall time of the pulse well, but of course fails to account for the presence of a third time constant.

The second fall time represents an enhancement to the pulse height of S2 and S3 compared to the pulse height measured by S1. We can see this directly by subtracting a pulse measured in S1 from a simultaneous pulse measured in S2 or S3. Figure 3.18 shows the average of 20 heater pulses measured simultaneously in S1, S2, and S3. S1 clearly has a lower pulse height than S2 and S3. Figure 3.19 shows the difference in pulse height, S2 - S1 and S3 -S2. The resulting difference has a fall time of ≈ 10 ms for both sensors. All three sensors have the fast fall time in their pulse; S1 differs from S2 and S3 by having a small amplitude component.

The fast fall time demonstrates an interesting dependence on the bias current through the sensor. Figure 3.17 shows the fast fall time measured in S1 and S2 for a series of bias currents and for two base temperatures, 20 mK and 25 mK. We might expect that a variation with bias current reflects a variation with the electron temperature in the sensors; increasing the bias current heats the electrons in the sensors and heats the detector as a whole as the increased power must be sunk through $g_{\phi s}$. This cannot be true, however, because the time constant is the same for a given bias at the both refrigerator temperatures, although the electron and detector temperatures are very different.

3.5 Interpretation

The performance of P2 raises some interesting issues with our understanding of how our phonon measurement works. The simplest hot electron model, while describing the basic phenomenology of the thermistors, fails to predict the detailed pulse shape of events in P2. We will review our three proposed extensions to the SHEM and compare their predictions to our data.

3.5.1 SHEM + Low Conductivity Interface

A poor interface could be a significant thermal impedance between the thermistor and detector phonon systems. We expect that S1 of all the sensors will have the largest thermal impedance to the detector. S2 and S3, because of their thin eutectic and epitaxial growth between the thermistor and detector,[24] should not have a significant thermal impedance between the thermistor and detector. In the limit of the heat capacity of the thermistor phonons being much less than the heat capacity of the detector phonons, the low conductivity interface acts to lower the pulse height response, but does not contribute an extra time constant to the pulses. In this case, we would expect that the performance of S2 and S3 would be well described by the SHEM and S1 would not. In particular, all sensors should have pulses with two time constants, and S1 should have lower pulse height. This is clearly untrue in P2; all sensors have three time constants, with S2 and S3 having a larger amplitude fast fall time than S1. In addition, the pulse height vs. bias prediction of the SHEM is a better fit to S1 than to S2 and S3.

If, however, our assumption that the interfaces between the detector crystal and sensors S2 and S3 are a negligible thermal impedance is wrong, the SHEM + low conductivity interface is still fails to describe the data. We would still expect only two time constants in the pulse, not the three observed. The low conductivity interface + SHEM extension alone is a poor model for P2.

3.5.2 SHEM + Hanging Heat Capacity

A significant heat capacity poorly coupled to the phonon system will have the effect of reducing the pulse height of all sensors, and could add another time constant to the pulse as the extra heat capacity heats to the phonon temperature. We would expect in this case that all detectors would have the same response, because the extra heat capacity is external to the sensors, and that the pulse height response would be lower than that predicted by the SHEM. This is in contradiction to the observed significant difference in response between S2 and S3, and S1. In addition, although the SHEM predicts the rise time and long fall time for S2 and S3, their actual pulse height is much greater than predicted. This model is not a good description of the behavior of P2.

3.5.3. SHEM + Collection of Athermal Phonons

If the thermistors are able to couple some of the initial recoil energy directly into the thermistor phonons, we should see a significant increase in pulse height in our sensors. The time constant with which the phonons couple to the sensors appears directly in the pulse shape as a time constant in addition to the SHEM time constants. We expect to see a dependence with the interface transmission, as presumably only the athermal phonons that propagate into the thermistor can heat the thermistor electrons directly. We expect then to see three time constants in the pulses, with S2 having the largest athermal component to the pulse, S3 somewhat less, and S1 very little. This extra time constant should be roughly independent of the detector temperature if these are high energy phonons coupling to the thermistor electrons. We also might expect to see a difference in sensor response to heater pulses and gamma ray events because their phonon spectrum is very different.

The data are somewhat supportive of this model, but again leave some questions. We do indeed see three time constants in our pulses, with S1 having the smallest amplitude addition of this extra time constant. If we think of this time constant as a reverberation time for phonons traveling ballistically through the detector crystal to leak into the thermistors, we can use equation 3.14 to estimate the transmissivity of the sensor interfaces, η. A time constant of $\approx$ 6 ms in S2 implies $\eta \approx 0.1$. This transmissivity is much lower than the transmission coefficient $\eta \approx 0.7$ that has been measured for 300 Å interfaces.[24] However, we are really convolving two effects, that of phonons propagating into the thermistor and the efficiency for the phonons to couple to the electrons. The effect shown in figure 3.17, that the fast fall time decreases with bias, could be interpreted as the efficiency of the phonon-electron coupling increasing with the number of electrons moving through the thermistor. A description of the exact mechanism and behavior of the coupling between the athermal phonons and the electrons in the thermistor should help decide if this effect we see is due to collection of athermal phonons.

3.6 Conclusion

This examination of the effect of the eutectic interface has answered some of our questions and raised some new ones. Our detectors have demonstrated the value of the Ge-Au-Ge eutectic bonding method of attaching NTD Ge thermistors to our Ge absorbers. The eutectic provides a reliable, consistent, stress-free bond. In addition, P2 has clearly shown that the interface is not the limiting factor in our failure to measure a significant athermal enhancement in the signal measured with our NTD Ge thermistors.

Because the bias wire on S4 was broken, we were not able to directly compare our standard 1000 Å eutectic with the thinner eutectics on S2 and S3. The comparison of P2 with sensor P1b:S3 is suggestive, however, that we could benefit by using the thinner eutectic on our detectors. We might expect the thinner eutectics to be better because these eutectics allow more epitaxial growth between the thermistor and detector. Another run of P2 with all four sensors working could answer this question.

The three time constants seen in pulses in P2 indicate that our simple hot electron model, while capturing the basic behavior of the thermistors, is not a complete description of the detector. A model that includes a direct coupling between the electrons and the phonons is consistent with the behavior seen in P2, although detailed calculations fail to accurately model the detector.

4. The Tower and FET Card

We have developed a novel system for for holding our detectors and cold electronics in the Icebox, the tower and FET card. I was part of the design team that developed the tower. I was responsible for most of the thermal conductivity measurements of the materials used in the tower and FET card. I designed and prototyped the FET card, and measured the cold FET noise characteristics on the card. I was initially responsible for integrating the tower and FET card with our detector E5 and our new electronics system. In addition, I developed our use of RuO_2 resistors as thermometers, that made possible the comprehensive monitoring of the tower and Icebox.

4.1. Requirements

Building a good detector is only part of the battle in our experiment. We must also hold the detector at 20 mK in a shielded, low radioactive background environment, and get low-noise signals to and from the detector. In addition, we would like to pack as much detector mass as possible into the icebox, so we want our detector packages to be as compact as possible. These ideas led us to develop our "tower" system of holding our detectors. The tower serves as the support structure that holds our detectors within the Icebox. As such, the tower is subject to a number of constraints arising from operating in a dilution refrigerator, operating in a low background environment, and operating with low noise signals. These constraints often push our design in opposing directions, and our final design will be a compromise between them. We will discuss these constraints below and how they affect our design.

4.1.1. Structural Integrity

We would like to use as much detector mass as possible in CDMS. Our current Berkeley Ge detectors are $\approx$ 160 g and we plan for each tower to hold six. To be safe, the tower should be able to support $\approx$ few kg. To complicate matters, we use the tower both in the Icebox, which loads from the top, and in Berkeley in a bottom loading refrigerator upside-down. We need to be able to flip the tower and detectors to either orientation. The tower must also carry the wires for each detector from the base temperature to the LHe can of the Icebox.

4.1.2 Thermal Conduction

The primary function of the tower is to hold the detectors at their nominal operating temperature of $\approx$20 mK inside the innermost can of the icebox. We cool the Icebox and detectors using a dilution refrigerator. The cooling power of a dilution refrigerator is very limited at base temperature. The Icebox can cool $\approx$ 1 W on the LHe can, $\approx$ 10 mW on the still, $\approx$ 10 μW on the cold plate, and $\approx$ few μW on the mixing chamber. The structure of the tower must carry the wiring from the LHe can to the detectors at the mixing chamber without compromising the thermal performance of the refrigerator. We typically minimize thermal conduction by using materials that are intrinsically poor conductors and by optimizing the geometry of the parts, i.e. minimizing A/L. In addition, careful heat sinking will minimize heat flow down the tower.

We must design the wiring carefully as well. Each tower will hold six detectors; each detector requires sixteen wires. We minimize heat flow down the wires in the same way we minimize heat flow down the tower. We use superconducting wires, which are poor thermal conductors, and heat sink the wires as they reach colder parts of the Icebox.

4.1.3 Low Background

We expect a very low event rate for CDMS (> 1 event kg^{-1} day^{-1}), and so must build a low background experiment. The tower must not make a significant contribution to the background in the experiment. We approach this problem from two directions. We try to insure we do not introduce any radioactive sources into the Icebox by using only pre-screened, radio-pure materials in constructing the tower. We have invested a considerable amount of money and time in measuring the level of radioactive contaminants in all of the material we use in building anything that will go inside the shield. In addition, we use the least amount of material necessary to construct the tower. This has two benefits. The total amount of radioactivity inside the shield from contamination in the materials will be minimized. Because we will initially be at a shallow site, muon induced neutron production is a problem. We can reduce the number of neutrons produced near the detectors by limiting the amount of mass near the detectors. If we are able to pack our detectors close together, we will also be able to use scattering in multiple detectors as an effective neutron tag.

4.1.4 Low Noise Signals

The operation of our detectors provides some electronics restraints on the tower. The Berkeley BLIPs use phonon sensors with resistances of ≈ few MΩ. To avoid low frequency noise, we typically use a lock-in measurement technique that requires us to bias the sensors at ≈ 1 kHz. This means we must keep the total capacitance of our front-end electronics to ≤ 100 pF to avoid RC roll off of the signal. At the same time, we need to have shielded signal wires to prevent cross-talk and pick-up problems. Microphonic noise can also be problem in our measurement. The tower should be designed to minimize susceptibility to microphonic pickup in the gate wires. This implies a few things; the gate wires should be as short as possible, and we need to eliminate insulators near the gate wires. Section 4.5.4 will discuss our experience with microphonics in more detail. Our desire to have short gate wires leads us to move our front-end FETs into the Icebox.

4.1.5 Infrared Shielding

Our detectors are both sensitive calorimeters and ionization detectors. Infrared radiation can easily sabotage both of these measurements. We are sensitive to pW type powers in the phonon measurement. IR radiation has the unwanted affect of producing a dark current in our ionization measurement. This can prevent us from applying an effective drift field across the detector. The IR photons also ionize impurity sites in the detector, creating regions of space charge that degrade our charge measurement. This is complicated by the presence of the JFETs in the Icebox which operate at 120 K, and therefore emit 120 K blackbody radiation.

Conventional methods of shielding IR radiation include coating the inside of each radiation shield with IR absorber. This is undesirable in the Icebox because of our low radioactivity background concerns, and because this absorber can make our wires microphonically susceptible.

4.2 Thermal Design Tools

Before we discuss in detail the design of the tower, we will review some of the theory of thermal conduction and practical applications that will help us in the thermal design of the tower.

4.2.1 Thermal Conduction

We first need to introduce a few useful equations relating thermal power, temperature, and thermal conductance. Thermal conduction is a diffusion process; one can write

$$\frac{P}{A} = -\kappa(T) \cdot \nabla T, \tag{4.1}$$

where P is the thermal power flowing across a surface of area A, $\kappa(T)$ is the thermal conductivity coefficient, T is the temperature at the surface, and ∇T the temperature gradient driving the power across that surface. We can write a simple expression for $\kappa(T)$ by considering that equation 4.1 describes the transport of heat. So we write

$$\kappa(T) = \frac{1}{3} c(T) \cdot v \cdot \lambda(T). \tag{4.2}$$

This can be understood as the product of the heat carried per unit volume (c(T)), the speed at which the carriers move (v), and how far the carriers move before being scattered ($\lambda(T)$). The factor of 1/3 comes from the fact that the carriers will be moving in three dimensions; we are only interested in movement along the temperature gradient. The carriers in equation 4.2 can be phonons or electrons. At very low temperatures scattering of both electrons and phonons is limited by impurities and defects and is thus temperature independent. For both phonons and electrons, the heat capacity per unit volume is proportional to a power of the temperature. This allows us to conveniently parameterize $\kappa(T) \equiv gT^{\alpha}$, where g has no dependence on temperature.

For our purposes it is often convenient to translate equation 4.1 into an integral relation between the total power flowing between two points at T_{hot} and T_{cold} and $\kappa(T)$. We can usually simplify the problem to one dimension, so that

$$P = \frac{A}{L} \int_{T_{cold}}^{T_{hot}} \kappa(T) \cdot dT, \tag{4.3}$$

where P is the thermal power flowing from T_{hot} to T_{cold} through an area A and across a length L. Substituting in our parameterization for $\kappa(T)$, we find

$$P = \frac{A}{L}\int_{T_{cold}}^{T_{hot}} gT^{\alpha} \cdot dT = \frac{A}{L} \cdot \frac{g}{\alpha+1}\left(T_{hot}^{\alpha+1} - T_{cold}^{\alpha+1}\right). \tag{4.4}$$

This will be a very useful equation to find the thermal power that will flow across a temperature difference. We can also invert equation 4.4 to find T_{hot},

$$T_{hot} = \sqrt[\alpha+1]{P \cdot \frac{L}{A} \cdot \frac{\alpha+1}{g} + T_{cold}} \, . \tag{4.5}$$

This will be useful in analyzing heat sinks.

4.2.2. The Wiedemann-Franz Law

It is often much easier in a metal to make an electrical conductivity measurement than a thermal conductivity measurement. Because the carriers in both instances are electrons, we can relate the charge carried by the electrons to the heat carried by the electrons. At low temperatures, the scattering length, λ_e, is dominated by impurity scattering and is independent of temperature. The electrical conductivity in this case will be independent of temperature too. C_e is proportional to temperature, so that the thermal conductivity will be proportional to temperature as well. It can be easily shown in kinetic theory that the ratio of the thermal conductivity to the electrical conductivity in a metal can be written as[32]

$$\frac{\kappa}{\sigma} = L_0 T \text{, where } L_0 = 24.5 \text{ nW} \cdot \Omega \cdot \text{K}^{-2}. \tag{4.6}$$

L_0, the Lorentz number, is the same for all conductors. Equation 4.6 is known as the Wiedemann-Franz Law, and is valid at low temperatures ($T \leq \theta_D/10$) and at high temperatures ($T \geq \theta_D$). At "medium" temperatures, the thermal conductivity is often lower than what the Wiedemann-Franz law predicts because of more complicated

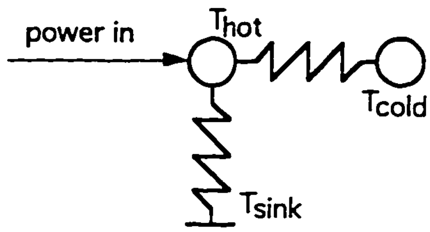

Figure 4.1 Thermal Divider
The thermal divider is composed of a high thermal impedance link between T_{hot} and T_{cold}, and a low thermal impedance link between T_{hot} and T_{sink}. Most of the power flowing onto the node at T_{hot} will flow onto the sink.

scattering processes.[33] Finally, let us restate the Wiedemann-Franz law in practical terms:

$$G(T) = 25.4\ \mathrm{nW} \cdot \frac{1}{R} \cdot T, \qquad [\mathrm{W/K}] \qquad (4.7)$$

where $G = A/L\ \kappa$, R is in ohms, and T is in Kelvins. With this version of the W-F law, there is no need to know the exact geometry -- it is taken care of in the measurement of R.

4.2.3. Thermal Dividers

We make use of thermal dividers throughout the tower. The idea is similar to an electrical divider with thermal impedances taking the place of electrical resistance and thermal power of electric current. The thermal analogy does not easily extend to voltage however; the relation between thermal power flowing through an impedance and the driving thermal gradient is not linear. Analyzing a thermal system like the tower can quickly become complicated because of the large number of elements at different temperatures. The analysis can be greatly simplified by a few realizations. Most of the tower thermal elements are designed to be thermal dividers -- power flows onto a node, and that power is divided between two conduction paths, one low impedance and one high impedance. This has two consequences for simplifying calculations. First, most (but not all!) of the thermal power will flow across the low impedance path in a well designed divider. Second, the power that flows through the high impedance path will be determined by the temperature on the hot side to first order. These two results allow us to break the complicated tower thermal system into easily analyzed subsystems consisting of individual thermal dividers.

We can design our divider in two steps. Figure 4.1 illustrates the thermal circuit. We will have a node at some temperature T_{hot}, a strap linking this node to a heat sink at T_{sink}, and a standoff holding the node above T_{cold}. We will have some thermal power flowing into the node, either from electrical power dissipated at the node, or as heat flowing from a hotter stage. Second, we have a thermal standoff that is connecting two different temperatures, T_{hot} and T_{cold}. It is usually the case that $T_{hot} >> T_{cold}$. The Icebox, for instance, is designed so that at each can, $T_{hot} \approx 10\ T_{cold}$. Most of the materials we might use for a standoff (e.g. Vespel, Graphite) have $\alpha \approx 2 - 3$.[34] Looking at equation 4.4, we see this means we can safely ignore T_{cold} when calculating the power flowing across the standoff.

$$P \approx \frac{A}{L} \cdot \frac{g}{\alpha + 1} T_{hot}^{\alpha+1}, \qquad T_{hot} >> T_{cold}. \tag{4.8}$$

This is helpful because we now can ignore small changes in temperature on colder layers. We need only calculate the temperature of the hotter layer, then the power that flows down the stand-off. Small changes in the lower temperature will not affect the upper temperature or the thermal power flow to first order.

4.3. Tower Design

The tower is hexagonal in plan. The hexagonal shape allows us to closely pack seven towers in an empty icebox. At the shallow Stanford Underground Facility, the icebox will have to be partially filled with polyethylene to moderate neutrons produced by muons in the Icebox itself; we will only pack three towers into the icebox at SUF.

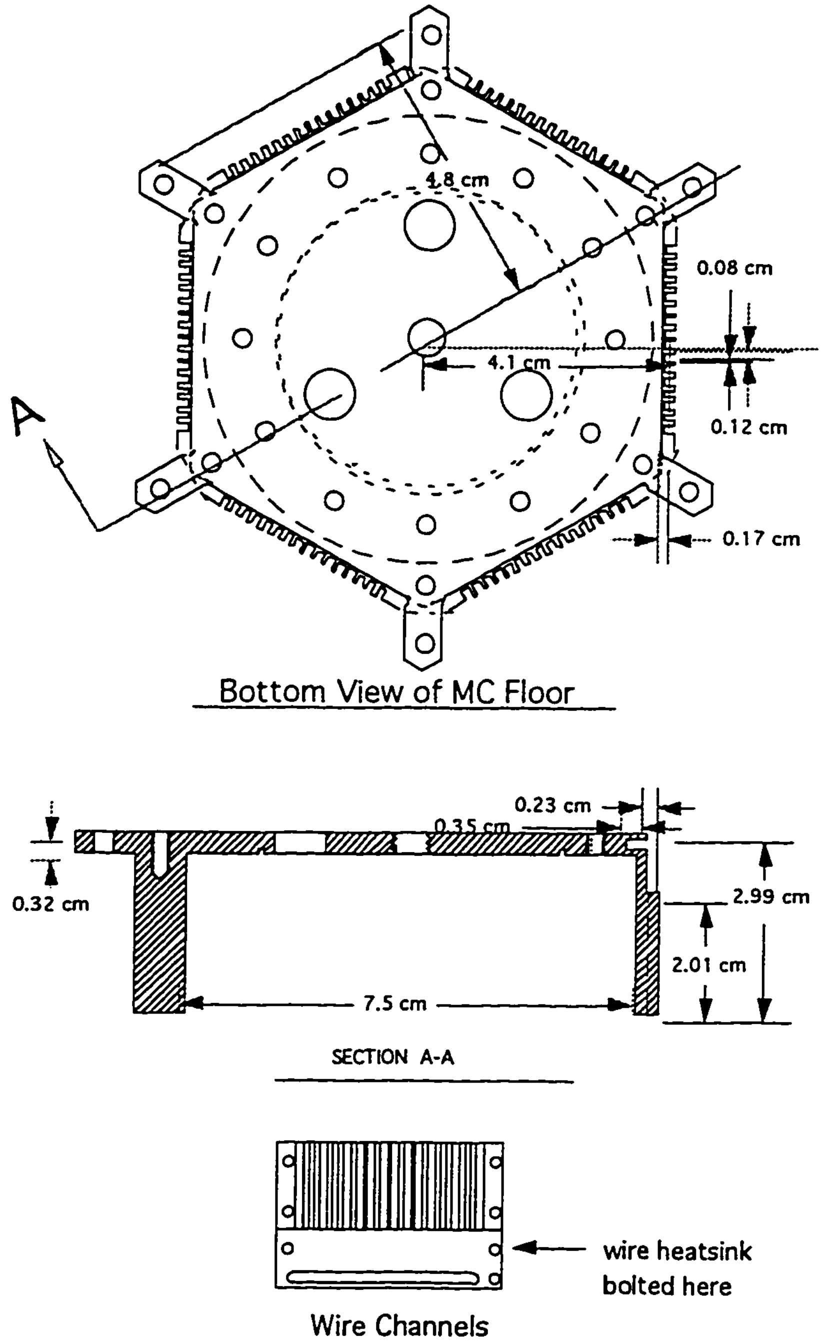

Figure 4.2 a. The MC Floor of the Tower.
This floor is bolted to the mixing chamber can. Shown are a view from the bottom, a section through the floor, and a view of one of the faces with the wire channels. Wire heat sinks are bolted to the six faces, and the stretched wires are soldered to the heat sinks. The graphite tube that act as the thermal standoff between the MC floor and the CP floor is glued into a groove on the inside of the floor.

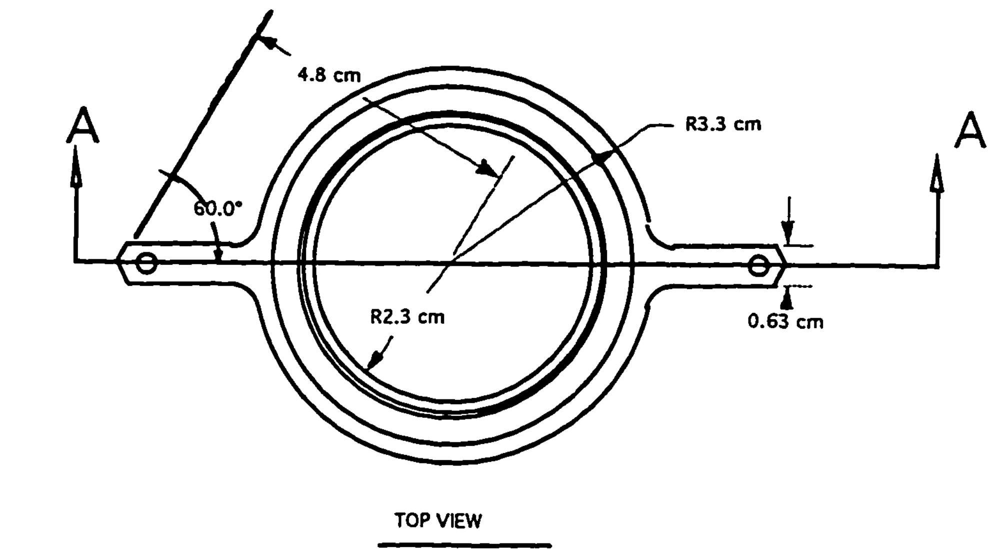

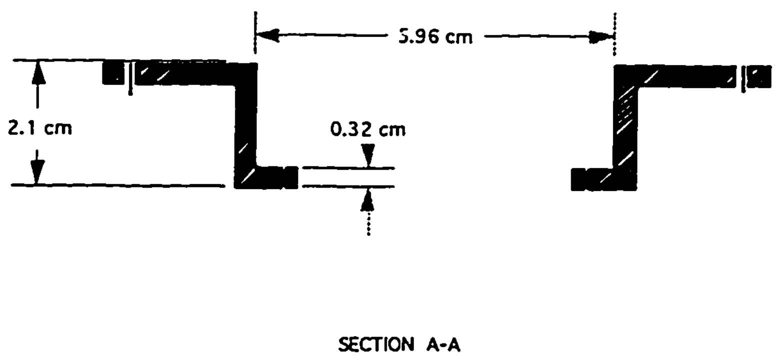

Figure 4.2 b. The CP Floor of the Tower.
The CP floor is thermally sunk to the cold plate can of the refrigerator. This floor serves to sink thermal power from the Still tower floor, so that the MC floor sees power from 100 mK, not 1 K. It was unnecessary to sink the wires on the cold plat, so this floor sits inside the MC floor and has no wire faces.

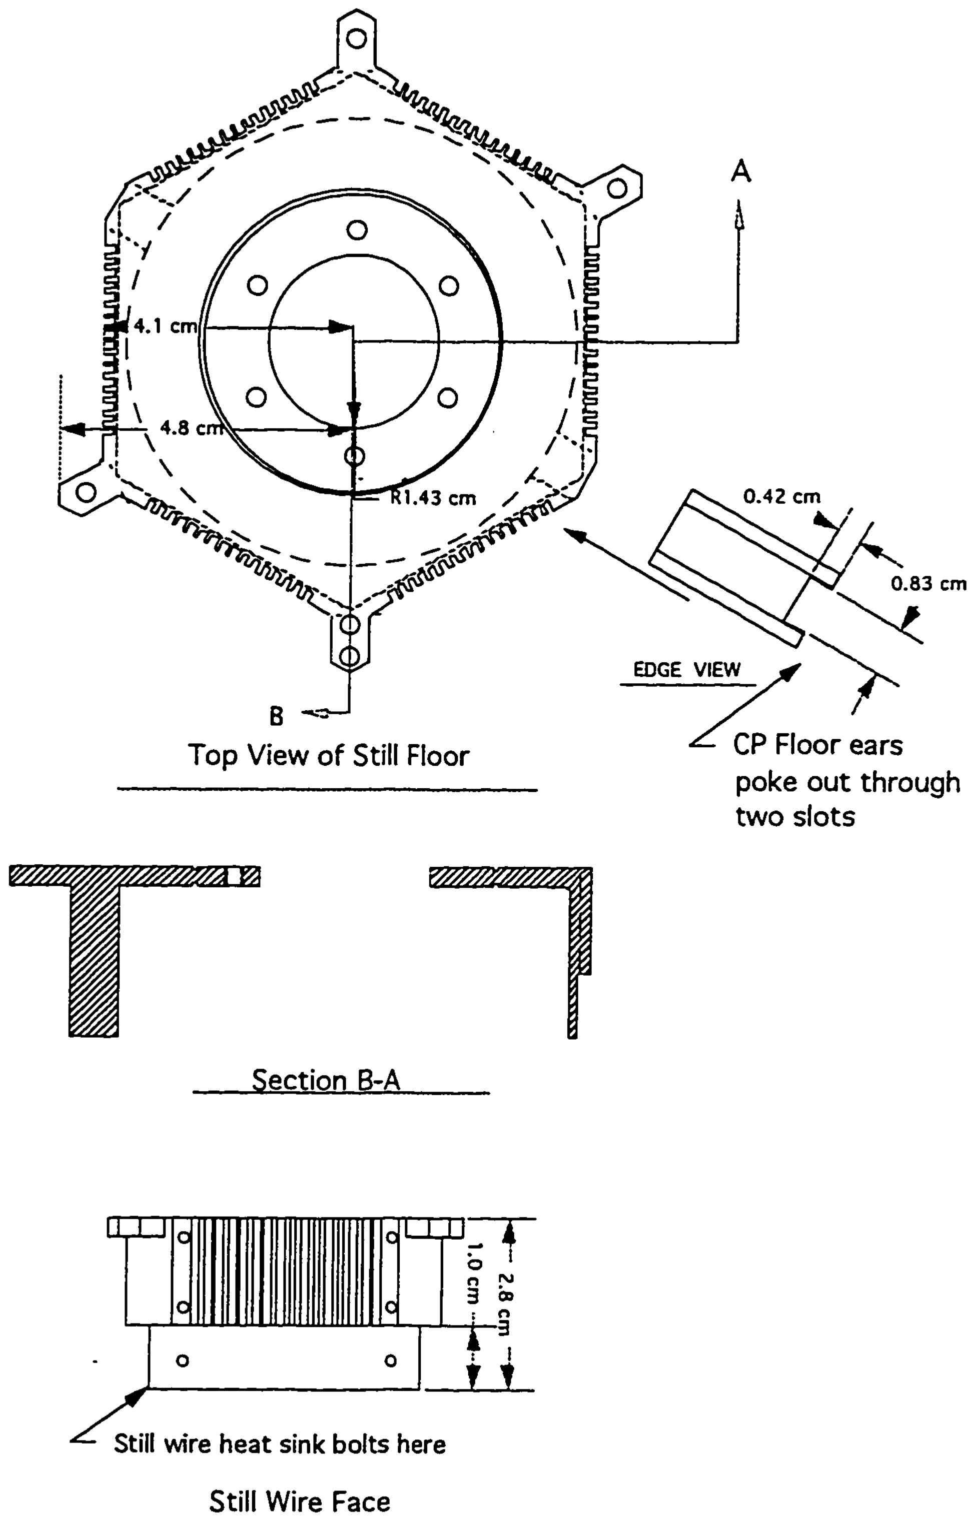

Figure 4.2 c. The Still Floor of the Tower.
The Still floor is thermally sunk to the still can of the refrigerator. Two of the ears are cut off above where the CP floor ears extend outside of the tower body. The Still wire heat sink is bolted to the bottom of each face. The wires are stretched over the heat sink and soldered to it.

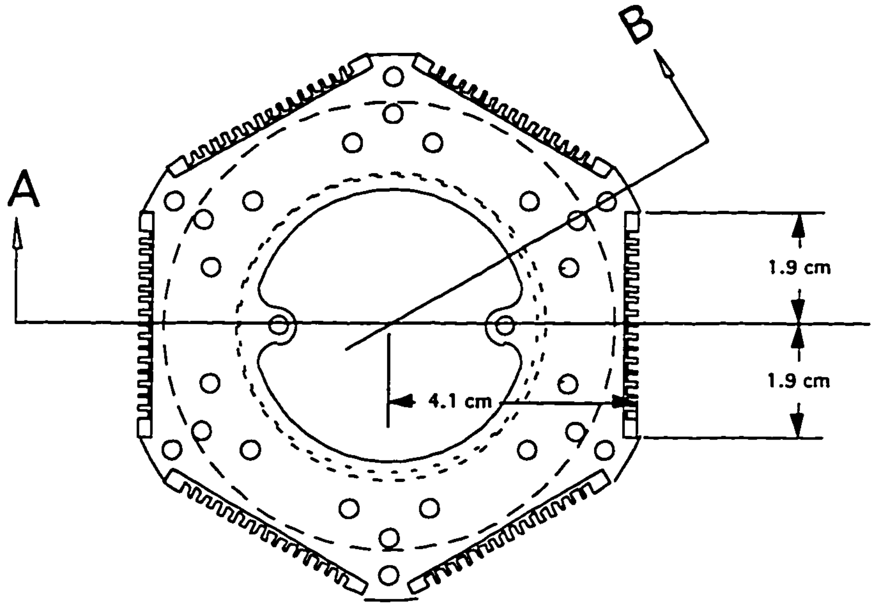

TOP VIEW OF 000LHe FLOOR

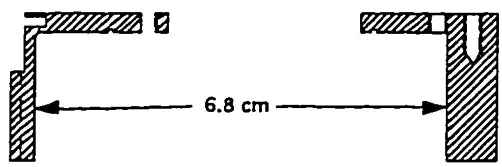

SECTION A-B

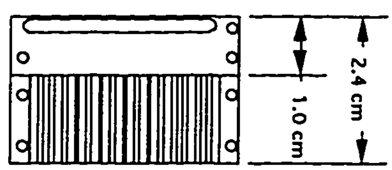

Wire channels

Figure 4.2 d. The LHe Floor of the Tower.
The LHe floor of the tower is strapped to the LHe can in the Icebox. The wires end in sockets at the top of the floor. Six FET cards plug in to the face tops and bolt onto the floor top.

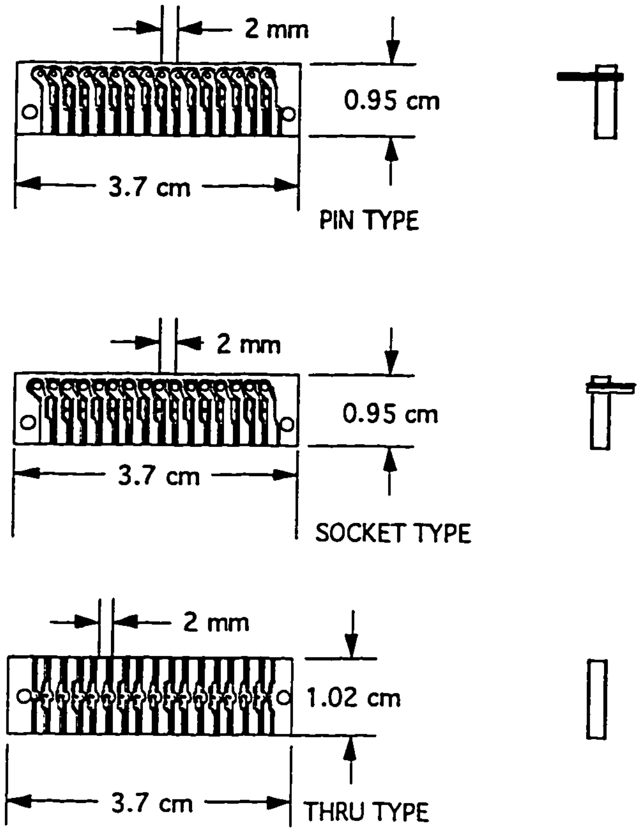

Figure 4.3 Wire Heat Sink Boards.
The tower system uses three types of wire heat sink boards. The LHe and MC tower floors use the socket-type boards; the Still floor uses the thru-type board. The side connector uses the pin-type board. The wires are soldered to the large pads on the board.

Wires are stretched down the six faces of the tower. These wires go from 4 K on the top LHe floor to 20 mK on the bottom MC floor in 8.3 cm. We fit sixteen wires down each face, allowing us to dedicate each face to the wiring for a separate detector. Thus each tower will naturally hold six detectors. The tower wiring terminates in female connectors on both the top and bottom of the tower. Figure 4.3 shows the Kapton circuit boards used as connectors and heat sinks on the tower. The tower consists of four copper floors separated by thermal standoffs. Figures 4.2a-d show the four floors of the tower, from bottom to top. The thermal stand-offs are thin-walled graphite tubes which act as the structural elements in the tower. Each floor is thermally coupled by a strap to one of the cans in the icebox, which are themselves thermally linked to the cooling elements of the dilution refrigerator.

The FET card plugs into the top of the tower. The FET card is bolted to the FET card gusset which is then bolted to the LHe floor of the tower. The FET card contains the front-end JFETs for our amplifiers, and acts as the physical interface between the tower wiring and the stripline wiring. Each stripline connects a detector's worth of wiring to the outside world through the e-stem of the icebox. The gusset acts both as a structural support and as an enclosure around the FETs.

The detectors are hung beneath the tower. The detector holders are independent of each other, so that individual detectors can be added or removed without dismantling the entire detector sub-assembly. Each detector holder includes circuit boards to hold the cold detector electronics for the detector. Electrical connections are made between the tower and detector holder by the side coaxial connector. The side coax consists of sixteen stretched vacuum coaxial wires, just like the wires stretched down the sides of the tower. The wires in the side coax terminate in male connectors to plug into the tower and the detector

Figure 4.4 shows schematically the tower floors and standoffs, the FET card and gusset, and the detector holder and side coax. The figure is oriented as installed in the Icebox, and all references to "top" and "bottom" refer to this orientation. With this convention, the tower installs "upside-down" in the modified Oxford 75 refrigerator at Berkeley.

4.3.1. Thermal Stand-offs

We need to minimize the power that flows down each stage of the tower. We accomplish this by placing large thermal impedances between the floors and small thermal impedances between the floors and cans. The large impedances between the

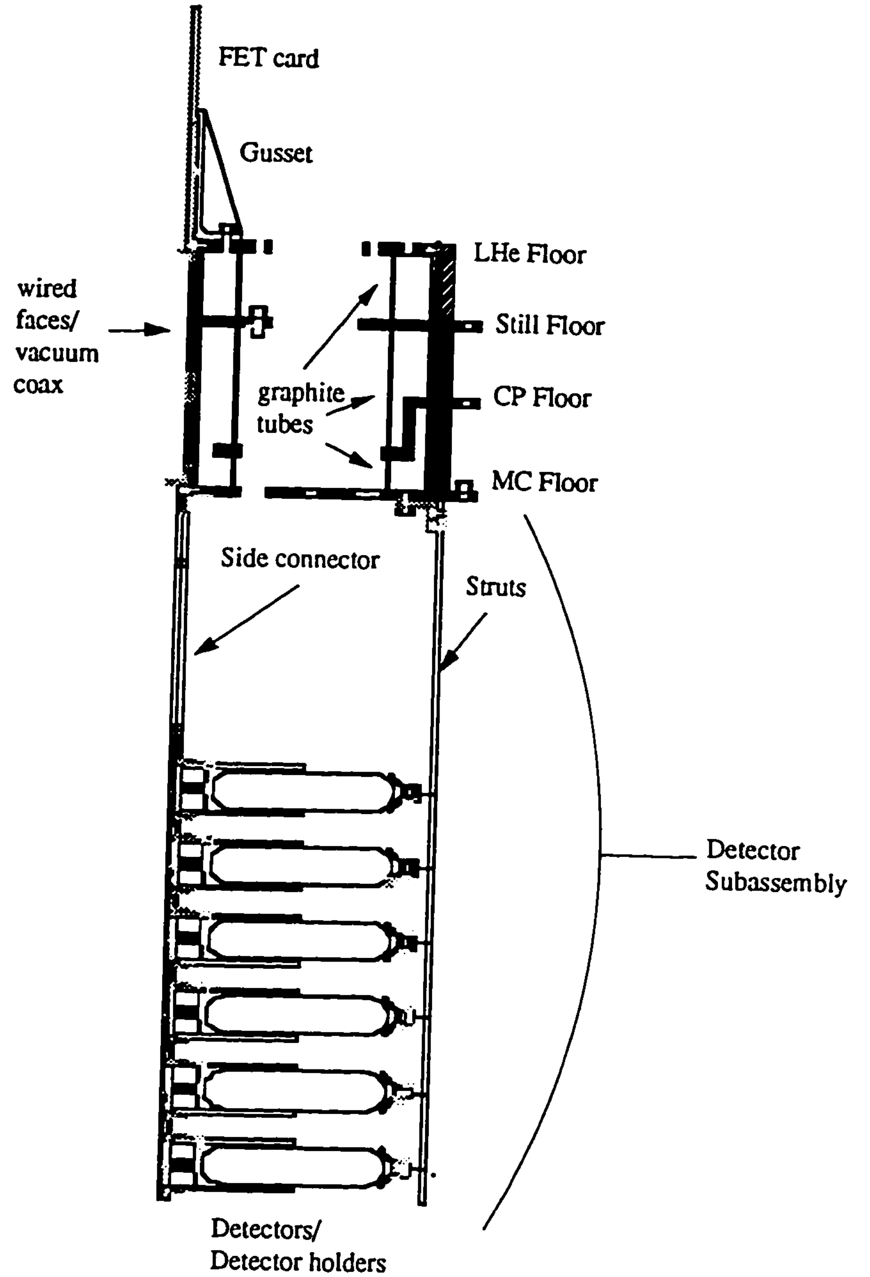

Figure 4.4 The Tower System
Shown are the major components of the tower system. This diagram is oriented as installed in the Icebox.

floors are the thermal stand-offs. Finding a material to make the thermal stand-offs is not an easy task. We use a material that is inherently a poor thermal conductor, but strong enough to act as the structural element holding the tower floors apart. Stainless steel and Vespel SP22 are common materials used in cryogenic engineering because of

Table 4.1 Thermal Conduction in Possible Thermal Standoffs.
Shown is the thermal power that would flow through standoffs between the tower floors for a geometry such that A/L = 1 cm. Thermal data is from Pobell.[33]

Material	LHe - Still	Still - CP	CP - MC
Stainless Steel	0.31	3.3E-3	3.3E-5
Vespel SP22	8.4E-4	8.4E-6	8.4E-8
Nuclear Graphite	1.6E-4	2.8E-6	4.9E-8

their poor thermal conductivity and strength. We are unable to use either of these because of radioactive contamination in the materials. For a given material, a structure with a smaller cross-section and a longer conduction path will conduct less power for a given temperature gradient. The spacing of the icebox cans roughly determines the spacing between the tower floors. We want to minimize A/L, where A is the cross-sectional area and L is the conduction path length. Since L is fixed by the can geometry, we are reduced to minimizing the cross-sectional area of the supports.

Obviously the cross-section of our support must be a function of the material we use -- a stronger material will allow us to use thinner supports. Table 4.1 compares the thermal power that would flow between the tower floors for standoffs made of three different materials. The material used to make the standoffs can make a huge difference in the conducted power. A standoff between the CP and MC floors made from stainless steel (which is radioactive anyway) would need an A/L $\approx 10^{-3}$ less than a standoff made from graphite.

Graphite has the lowest thermal conductivity at low temperatures of any common material.[34] Graphite is also fairly strong, and can be bought in radio-pure form. The tower uses thin-walled graphite tubes as thermal standoffs between each floor. The tubes are machined to a thickness of 28 mil and an OD. of 1.99". The length of each tube is roughly set by the vertical spacing between the icebox can lids, but we can adjust the lengths of the tubes with creative floor design. This is the reason the CP floor has the dropped shape; after calculating heat loads in the tower, we decided we needed a longer Still-CP tube than the simple length between the Still and CP cans. Using published values for thermal conduction in graphite, we can calculate the thermal conduction down the tubes. Thermal conduction in graphite can vary considerably from sample to sample. It would be best to measure the thermal conduction of the particular graphite we use in the tower; unfortunately, the graphite sample we measured proved to be too radioactive to use, so we were forced to use thermally unmeasured but radio pure graphite. Its properties proved to be close to the graphite shown in Pobell.

The graphite stand-offs are designed so that if the tower floors are near the nominal temperatures of the Icebox cans, the tower will not compromise the performance of the dilution refrigerator. The thin graphite tubes are surprisingly strong, and can easily support the entire weight of the tower and 1 kg of detectors if necessary. We have designed the tower such that only the MC floor supports the weight of the detectors when mounted in the Icebox. The supports are subjected to some stress when we bolt heat sink straps to the other floors.

4.3.2. Heat Sinking the Floors

A well designed tower should not affect the performance of the refrigerator. With no external power load to the tower, like the FETs, the power that flows down the graphite tubes will not significantly load the cooling power of any of the refrigerator stages. With the FETs turned on, we now have a significant extra power that must be sunk without disturbing the performance of the refrigerator. The tower can fail by letting more power flow onto a refrigerator stage than that stage can cool. This then either heats the mixing chamber directly, or degrades the performance of the refrigerator, so that the cooling power at the mixing chamber is reduced. Because the cooling power of the warmer stages are much greater than that of the cooler stages, our goal is to sink as

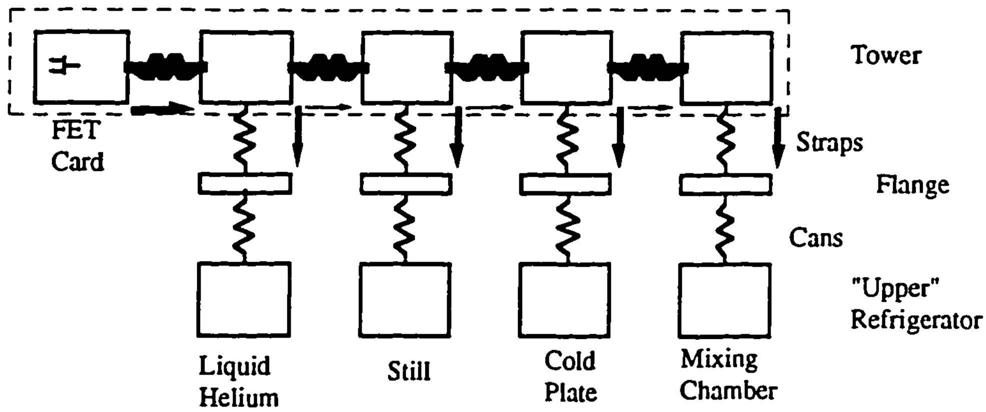

Figure 4.5 Tower and Refrigerator Thermal System.
This diagram shows the thermal power flow through the tower and refrigerator thermal system. The tower sections are separated by thermal standoffs, and each section is heat sunk to a refrigerator flange. Most of the thermal power will flow onto the flange, with some small amount flowing onto the next lower floor in the tower. The power flowing onto a given floor equals the power flowing off through the heat strap and the power flowing down to the next floor. We can use the impedance of the cans to measure the power flowing to each section of the "upper" refrigerator; this measures the power flow through the entire system.

much power as possible in the warm stages. We accomplish this by making the straps good thermal conductors. Figure 4.5 shows the entire thermal circuit for the tower and FET cards.

Our design of the thermal heat straps is driven by two main considerations. A heat strap must be a good thermal conductor and the strap must be flexible. Flexibility is important because the icebox cans are not fixed rigidly relative to each other. If the tower floors were rigidly fixed relative to the cans, movement of the cans could shear off floors from the tower. The thermal conduction of a wire is proportional to its cross-sectional area, $\propto r^2$. The stiffness of a wire increases like r^3, so if we increase the conductivity of our wire by making it thicker, we quickly make it very stiff. A solution to this is to use a braid for a heat strap. We use many small, flexible copper wires to make up the needed cross-sectional area. We have used two types of copper braid; Chemwick with the flux removed and a copper welding cable.

The heat strap still needs to be in good thermal contact with both the tower and the refrigerator. We weld copper lugs onto the ends of the cable or braid. We attach the

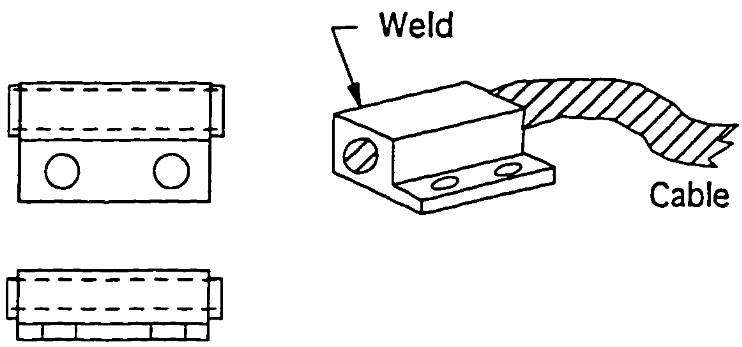

Figure 4.6 Heat Strap Lugs.
The heat strap lugs are basically small copper blocks drilled to allow a copper cable to be inserted through. The ends of the block have a lip around the hole. When the cable is welded to the block, the lip is melted into the cable.

straps to the tower by bolting the lug onto the tower "ears." We bolt the other end of the strap onto the flange (or can lid). This introduces four joints into the heat strap system -- two bolted lug/flange joints, and two welded cable/lug joints. The thermal impedance of these joints can easily dominate the impedance of the straps if we are not careful. Figure 4.6 shows the latest design for the heat strap lugs. The lugs are basically solid oxygen-free high-conductivity (OFHC) copper blocks with one hole drilled through the block to allow the insertion of the cable. The lug has two M3 clear holes drilled through to allow the lug to be bolted onto the tower or flange. The ends of the lug have a lip of copper left around the cable hole. The cables are attached to the lug by tungsten inert gas (TIG) welding. When the cables are welded onto the block, the lip on the end melts into the copper, giving a solid copper weld that penetrates about 1/4" into the block. We TIG weld to prevent the introduction of radioactive impurities into the weld.

The bolt joint can be a significant thermal impedance. The impedance of the bolt joint is affected by oxidation of the surfaces to be joined and the force pressing the two surfaces together.[35] The actual contact area is much less than the area of the lug because the surfaces are not perfectly smooth; thermal contact is made through many small "rough" spots. To combat oxidation, we use a plating pen to plate 24 K gold onto small spots on the tower ears and lugs where they are joined. Gold has the added

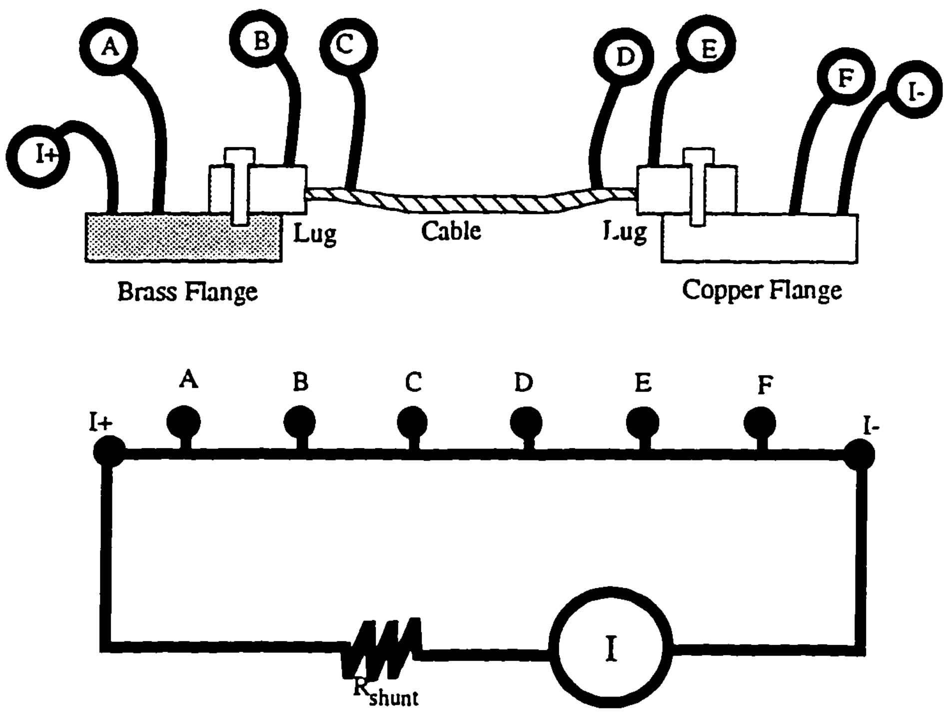

Figure 4.7 Wiedemann-Franz Test of Heat Straps.
We test the heat straps at 4 K by driving a large (3.5 A) current through the strap. The voltage across the bolted lug/flange joints, the lug/cable joints and the cable are measured. The actual current is monitored by measuring the voltage across a 10 mΩ shunt resistance.

advantage of being fairly soft, so it should deform and make a larger contact area than plain copper joints for a given pressure. We want to tighten the bolts as much as we can without stripping the threads on the tower or snapping the bolt heads off. We can use stainless steel M3 bolts in the modified Oxford 75 refrigerator, but we can only use counted brass M3 bolts in the icebox. With stainless steel SHC bolts, we can tighten to about 12 inch-pounds torque before we risk stripping threads; with brass SHC bolts, we can tighten to 8 inch-pounds torque before we risk breaking the bolt heads off. We have indirectly tested the thermal conductance of the straps and bolt joints. We do this by measuring the electrical impedance of the straps and bolt joints at 4 K, and inferring the thermal conductance using the Wiedemann-Franz law.

Figure 4.7 shows the experimental set up used to measure the electrical resistance of the heat straps at 4 K. We simulate both a copper flange and a brass

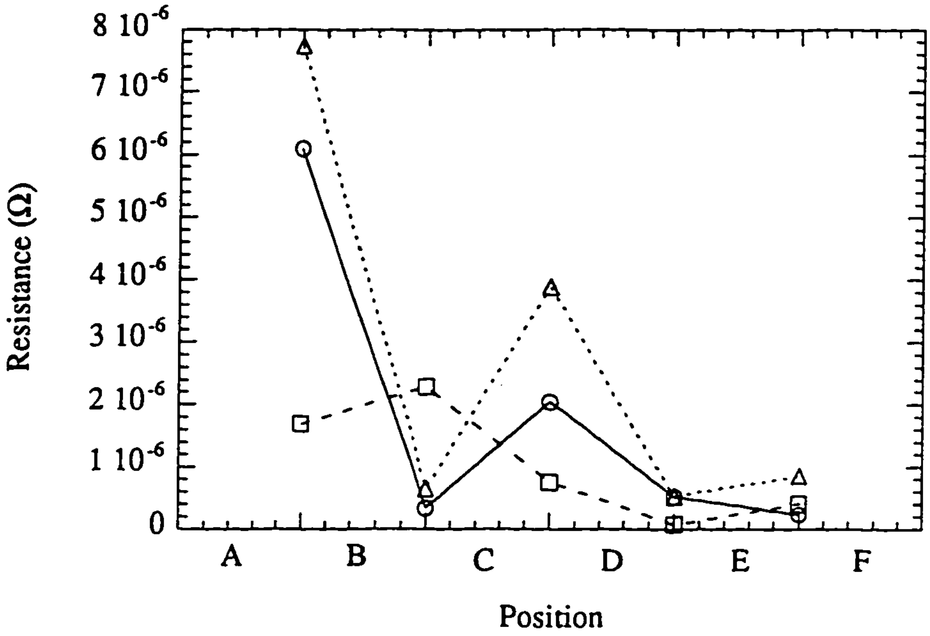

Figure 4.8 Resistance Measurements of Three Heat Straps at 4 K.
Shown is the measurement of electrical resistance across three heat straps. The position label is the same as figure 4.7. Note that the largest impedance in all cases is the lug/brass flange bolt joint. The triangle and circle data are from heat straps made with braids; the square data is from a heat strap made from oxygen-annealed OFHC copper wire.

flange; the modified Oxford 75 uses both brass and copper flanges, and the icebox uses only copper flanges and lids. The apparatus is built into a probe designed to be inserted into a LHe dewar. Electrical contact for the voltage monitoring points is made by bolting to the flange, and by alligator clips to the other points. The bolts for the shown data were stainless steel M3 bolts tightened to 8 inch-pounds. We mount the actual straps we use onto the apparatus and dunk the probe into the dewar. After the probe cools, we bias the strap with $\approx$ 3.5 A. The actual current is monitored by measuring the voltage across a warm 10 mΩ shunt resistance. We then measure the voltage across all the monitoring points. Figure 4.8 shows the resistance across the various joints for three straps used in mounting the tower in the modified Oxford 75. In this test, none of the bolt joints were gold plated. We note that in all cases, the lug/brass flange bolt joint is the highest resistance in the system. The data marked with triangles and circles are from straps made from copper braid; the data marked by squares is from a strap made from an oxygen-annealed OFHC copper wire. The total resistance of the braid straps is $\approx$ 10 - 14 $\mu\Omega$; the total resistance of the wire strap is $\approx$ 4 $\mu\Omega$. We can use this measurement as a benchmark to compare against the performance of the straps in the refrigerator.

The thermal links between the tower and the cans are crucial to the functioning of the tower. The tower design depends on efficient heat sinking of the tower floors to their cans. We use two types of "heat strap" to thermally link the floors to their cans in the modified Oxford 75 at Berkeley. One type, used on the liquid helium floor, consists of a thick OFHC copper wire that has been oxygen annealed to improve its thermal conductivity.[36] This process improved the RRR of the copper wire from 42 before treatment to $\approx$2200 after treatment. The second type of strap is made from copper braid. The braid is used on the still and cold plate floors. In all cases, copper lugs are welded to the ends of the strap. The lugs provide a solid connection that can be bolted to the tower and to the cans. The copper braid is used to ensure the strap is flexible; this is so that the strap cannot put significant shear forces onto the tower. In addition, all of the straps are bent so that they will flex, rather than transmit force directly along them. We can calculate the thermal conductance of the straps by treating them as wires of copper with bolt joints at each end. The actual design of the straps turns out to be not as important as the quality of the joints between the lugs and wire/braid and between the lug and the can or tower.

4.3.3. Wiring the Tower

Each tower holds six detectors and each detector uses sixteen wires. The tower must carry 96 wires from the FET card in the LHe can down to the detector at 20 mK.

The detector is much more sensitive to thermal power than the MC can. The Icebox can tolerate μW's of thermal power flowing onto the MC can, but the detector can tolerate only pW's. We need shielded wires to prevent cross-talk and pick-up on our gate wires and bias wires. We also want the gate wires as isolated from insulators as possible to reduce susceptibilty to microphonics. Our solution was to create "vacuum" coaxial cables using stretched wires inside copper channels on the tower face.

The tower sides are slotted to create sixteen channels per face. The face of each floor stops before it contacts the neighboring floors to prevent thermal contact. Heat sink circuit boards are bolted to the LHe, Still, and MC floors. The wires are stretched through the channels and soldered to the circuit boards to create the vacuum coaxial cables. The channels are covered with 10 mil copper covers to complete the wire electrical shield, except for the small gap between floors. We use copper clad NbTi superconducting wire on the tower. The copper is etched away from most of the wire, leaving copper only where we need it to solder the wire to its heat sinks. Superconducting wire is thermally advantageous in that when $T \leq T_C/10$, virtually all of the conduction electrons are in Cooper pairs and cannot conduct heat. Furthermore, the lattice of superconducting materials are generally poor thermal conductors because of the phonon-electron coupling needed to produce the Cooper pairing. T_C for NbTi wire is 10 K, so for temperatures less than 1 K, the NbTi wires will conduct heat like an insulator.

The heat sink design for the wires works in the same way as the floors. A highly thermal conductive path is provided from each wire to the floor (the copper-Kapton heat sink circuit board), and a poorly thermal conductive path (the wire) is provided to the next lower temperature. The wire is heat sunk on three of the tower floors (LHe, Still and MC) by soldering it to a trace on a copper-Kapton stack-copper circuit board. We use Kapton because G-10, the usual circuit board base, is radioactive.

Figure 4.3 shows the heat sink pads. We use the socket-type board on the LHe and MC floors, and the thru-type on the Still floor. The wire on the Still heat sink is soldered at each end and clipped in the middle, to ensure heat is not conducted straight through the wire. We use our own Pb-Sn solder manufactured from radio-pure materials. Each circuit board is bolted and epoxied with low radioactivity Stycast 1266 to its floor to heat sink it.

The wires are stretched by hanging a 30 g weight from one end as the wires are soldered. We need to stretch them so that as the tower cools and thermally contracts the wires will remain tight. Graphite thermally contracts a negligible amount, so the relevant change in length comes from the shrinking of the Still floor wire face. The Still floor wire heat sink will shrink towards the LHe floor and away from the MC floor heat sink. We need to stretch the wire enough so that the wire between the LHe and Still floors does not become slack, but not so tight that the wire between the Still and MC floors stretches past the elastic limit for the wire. [37]

The heat sink circuit boards are manufactured by laminating 5 mil Kapton in a stack with 1 oz. copper on the outside to create a total thickness for the circuit board of 50 mil. The circuit is photo etched onto this stack in the normal manner. The copper circuits on the boards are gold plated. The circuit boards on the LHe and MC floors are drilled to allow connector sockets to be soldered into the board. The connector holes are gold plated though, both to ensure the sockets are not microphonic susceptible and to allow solder to flow through the holes and make a better mechanical joint with the sockets. These boards act as heat sinks, anchors for the stretched wire and connectors at the ends of the tower.

The wires are carried down to the detector from the MC floor of the tower by the side coaxial connector. The side coax is made in the same way as the tower face. Slots are cut into a copper piece, and heat sink boards are bolted and epoxied to the ends.

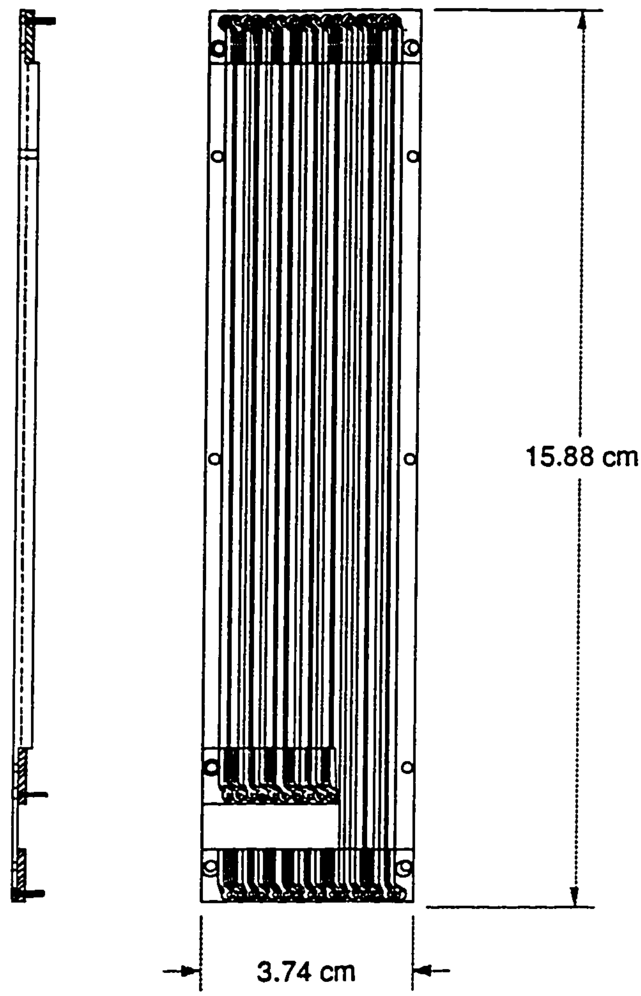

Figure 4.9 Side Coaxial Connector.
The side coax connects the wires for one detector from the MC floor of the tower to the detector holder. The connector on the bottom is split to plug into both top and bottom of the detector holder. The wires are stretched through slots like the wires on the side of the tower. Because detectors stack beneath the tower, we need a different length side coax for each detector slot.

These heat sink boards are the same as the tower boards, except they are drilled to accept connector male pins. Because the detectors are wired on both sides, the side coax must bring wires to both sides of the detector holder. Accordingly, the bottom heat sink board on the side coax is split. For the current Berkeley detectors (E5 and BLIPs), ten wires go to the phonon side of the detector and six wires go to the ionization side of the detector. The side coax has male connectors to plug into the detector holder and the bottom of the tower. The side coaxial connector is shown in figure 4.9.

4.3.4. The Detector Holder

The detector holder is perhaps the most sensitive part of the tower. Because of its proximity to the detector, we must be very careful in how the detector holder interacts with the detector. Background considerations (muon-induced neutron production, radio-purity, scattering between detectors) compel us to use the least amount of material near the detector as possible. In addition, we want to use the least number of different kinds of materials. The more different materials we use, the more materials we have to count for radio-purity and keep track of. This is expensive, and increases the chance we will inadvertently use contaminated material near the detector. Accordingly, we rely on the same materials we used in the rest of the tower: activity tested OFHC copper, Cu-Ka-Cu circuit boards, and activity tested Pb-Sn solder.

The skeleton of the detector holder is the detector holder ring. This hexagon of copper is the main structural element of the detector holder. The detector itself is held inside the ring by three pairs of copper "fingers." One of each pair of fingers is rigid and fixed to the ring; the other finger is made from thin copper sheet. The thin finger is bent to act as a spring, and is screwed onto the rigid finger. Small sapphire balls are glued to the ends of the fingers. The sapphire balls make the actual contact with the detector. Using sapphire balls allows us to hold the detector firmly, while making a very poor thermal connection to the detector. In the past we have used Delrin screws to hold the detector. These gave long time constants on the detector cool down. Sapphire is much harder than germanium. This led us to worry that the weight of the detector would lead to pressures above the yield point of germanium at the contact points with the sapphire balls. This becomes a real problem when considering that we will need to move the detector while it is in its holder. We developed a method of "softening" the sapphire balls to minimize this problem. We sputter a 10,000 Å gold film onto ≈1/3 the surface of the ball. The ball is installed in the fingers such that the part with the gold film will contact the detector. The gold is softer than the germanium, and deforms to increase the contact area. The gold does not extend to touch the fingers, so the detector remains thermally and electrically insulated from the detector holder. The gold will add to the heat capacity of the detector, however, so it is important to be careful with the amount of gold used on the balls.

Both sides of the detectors are instrumented. One side has the NTD phonon sensors bonded to it; the other side has the segmented ionization collection contacts. The detector holder has Cu-Ka-Cu circuit boards (called "Y"'s because of their shape) bolted to each side of the ring. Small copper spacers hold the Y's on either side of the ring and above and below the detector. The Y's contain the sockets to plug into the side coaxes and all cold electronics for each measurement. In addition, the detector itself is heat sunk to the Y. The detector heat sinking is accomplished through a patch of gold sputtered onto the surface of the detector. A gold wire is bonded between this gold patch and the phonon Y. The conductance between the detector and the tower in our design is dominated by the Kapitza resistance between the germanium detector and the gold pad. We control the Heat Sinking by varying the surface area of the gold heat sink pad on the detector. Figure 4.10 shows the detector holder and Y's.

4.3.5 Infrared Shielding

The conventional solution to infrared problems in cryogenics is twofold. First, the apparatus is made as light tight as possible between the FETs and the detector. FETs would be enclosed in a metal box with wiring fed through bulkhead connectors.

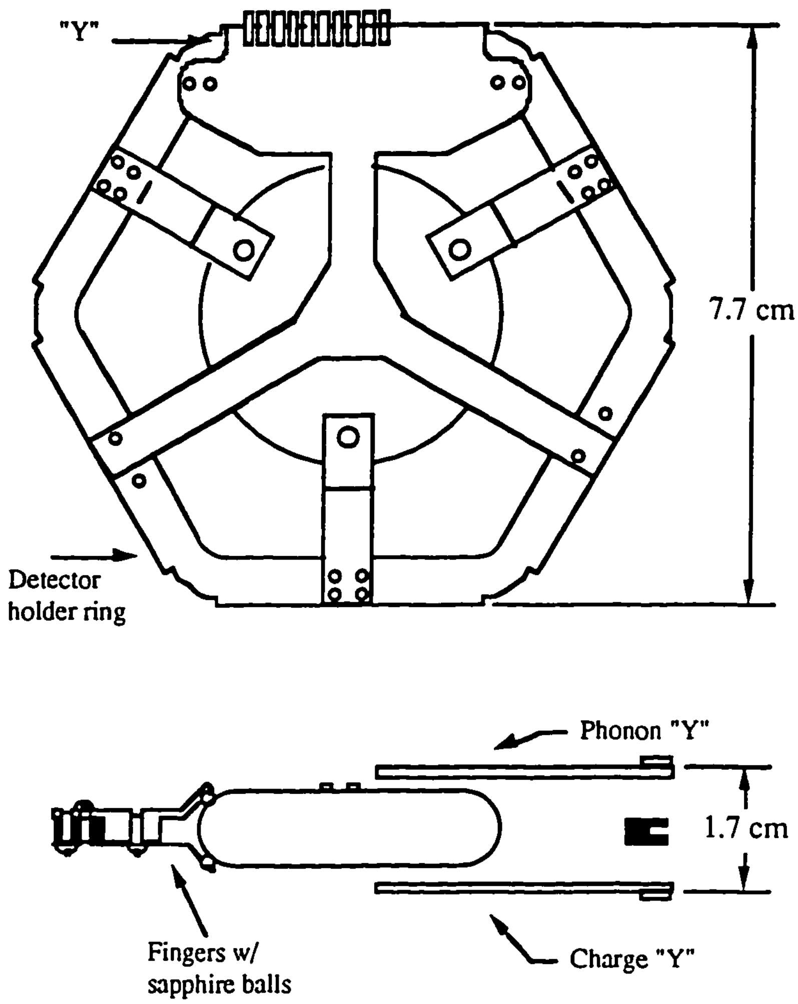

Figure 4.10 The Detector Holder.
The detector is held in the detector holder ring by being squeezed between copper fingers with sapphire balls on the tips. The sapphire balls are the only material actually in contact with the detector. The detector holder also carries 2 "Y"s which hold the cold electronics for each measurement. The side coax plugs into the two Y's. The detector holder ring is bolted to six struts which suspend the detector holders beneath the tower. The detector shown is drawn to scale for E5. The BLIPs will fill the ring.

The various temperature stage cans would be completely enclosed, with a solid lid on each can. Again, wiring would penetrate the cans using bulkhead feedthroughs. Second, the inner surface of each can is made to absorb infrared radiation by coating with a viscous black "goop" -- a mixture of carbon lamp black and epoxy. Thus any infrared that penetrates into a given can will absorbed on the can wall. Our tower

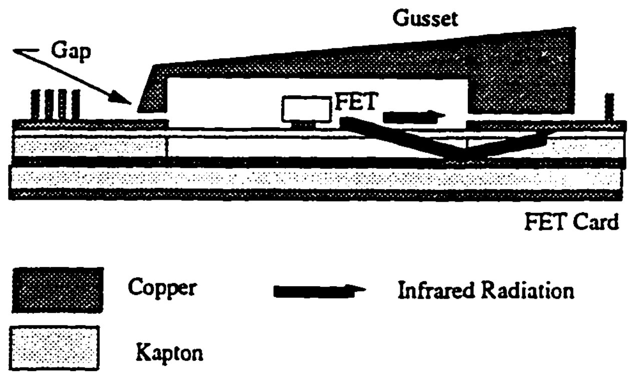

Figure 4.11 The FET Card as a "Leaky Box."
The FET card is a "leaky box" to IR radiation. The top surface of the FET card is patterned with electrical traces, so the FET card gusset cannot seal to the top face without shorting the traces. This leaves a ≈1" X 0.002" gap through which IR can leave the FET card directly. In addition, IR could leave the FET card through the Kapton if Kapton does not absorb it. Without some absorber in the FET card cavity, the IR radiation will bounce inside until it finds its way out.

design violates both of these principles, and so makes it hard to shield the detector from infrared.

The FET card and gusset assembly acts as a leaky box. The FET card bolts onto the gusset, with the FET region fitting into a hollow in the gusset. Figure 4.11 illustrates the light leaks in the FET card. Because the top surface of the FET card is patterned with traces, the gusset cannot seal against the FET card without shorting out these traces, except along the sides where there are no traces. This leaves an ≈ 1" X 0.002" gap on two sides of the FET enclosure. This gap is especially troublesome because of the way the FET card plugs into the tower; this gap looks down the tower vacuum coax. In addition, if the Kapton does not completely absorb IR radiation, then IR can leave the enclosure through the Kapton layer beneath the top surface. We have some evidence that Kapton absorbs IR, [38]but the absorption coefficient is a strong function of frequency, with definite peaks and valleys. Figure 4.12 shows the data on Kapton absorption in the infrared and compares it to the spectral density of a black body at 120 K. Unfortunately, the data does not extend to cover the high energy tail of the

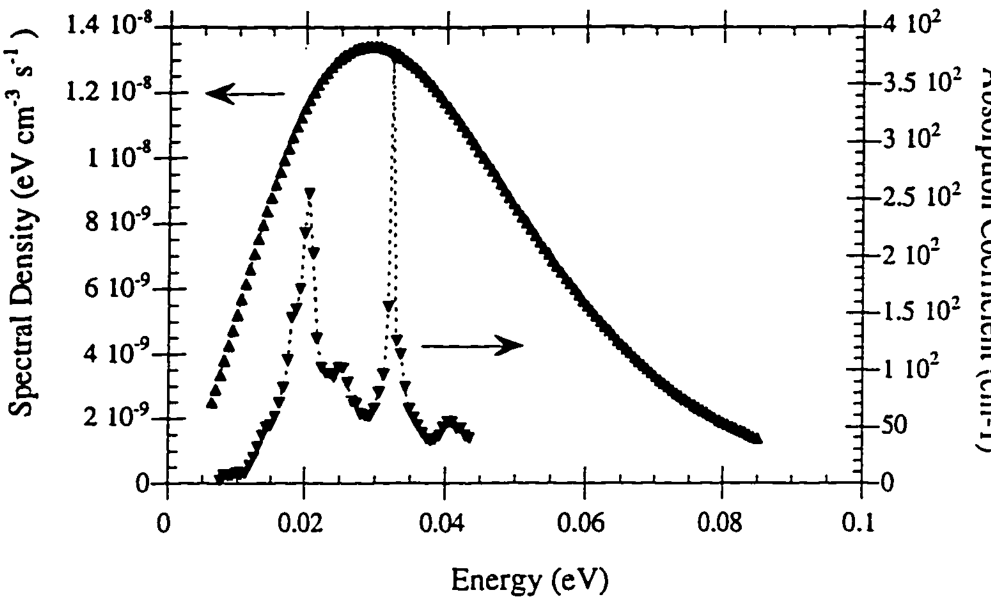

Figure 4.12 Comparison of Kapton Absorption with 120 K Black Body Emission.

This shows a comparison of the absorption coefficient measured by Smith, *et al*,[38] with the spectral density of a 120 K black body. We can see that although the Kapton has strong absorption peaks near the peak of the spectral density, the absorption over the high energy tail of the distribution is unmeasured.

black body distribution. Notice virtually all of the power in the black body spectrum is above the impurity gap energy, $\Delta E_i \approx 10$ meV, in our detector. The IR radiation can possibly escape the FET card side through ≈ 1/4" of Kapton. As long as the high energy absorption coefficient does not go to zero (it is not completely transparent in the visual range), the absorption through 1/4" of Kapton should be sufficient to completely absorb the IR. However, a thin (5 mil) Kapton sheet will not absorb infrared -- we need a good fraction of an inch to ensure the infrared is absorbed.

The tower penetrates all of the cans in the icebox and modified Oxford 75. This makes it very difficult, if not impossible, to make the cans light-tight. In particular, we are faces with two real problems in radiation shielding with the tower in place. The tower vacuum coaxes look straight through all the cans into the MC can. They are only half blocked by the copper-Kapton heat sinks on the LHe, Still, and MC floors. Second, we must always be careful we do not install rigid shields between the tower floors and can lids. We cannot allow our shields to put shear forces on the tower floors when he cans move. We make radiation shields from copper-Kapton sheet to complete the radiation shielding between the tower and can lids or can flanges. We rely on the copper-Kapton shield making a snug fit against the tower, but we cannot really ensure that the shielding is light-tight.

Another crucial difference between our shielding scheme in the icebox and conventional IR shielding technique is the use of absorber. Conventionally, we would coat the insides of all the cans with an IR absorbing goop. This would ensure that any high temperature IR radiation entering a can will be absorbed on the can wall and thermalized to the lower can temperature. Because we are designing a low radioactive background experiment, we cannot afford to have a large quantity of possibly radioactive IR absorber. The conventional recipe uses Stycast 2850 GT filled epoxy and lampblack. We know that the glass filler in Stycast 2850 GT is too radioactive to use in the icebox in large quantities. We use a mixture of Stycast 1266 and lampblack as our infrared absorber in these ratios: 100 parts resin: 28 parts catalyst: 10 parts lampblack by weight. This epoxy/carbon mixture is electrically conductive, which limits its utility on the wire heat sinks and FET card. When it must be electrically isolated from parts, we first use a coat of Stycast 1266 epoxy on the parts, then the Stycast 1266/carbon mixture. The infrared absorption of a similar mixture of Stycast 1266 and charcoal has been measured with the result that a 0.17 cm thick sample absorbed 80% of the infrared incident with wave number between 80 - 350 cm^{-1}.[39]

It is also important to note that using an infrared absorber near our signal wires is likely to cause increased microphonics. The epoxy/carbon mixture is a poor conductor, but good enough that we must use an insulating layer of epoxy near wires and electrical traces. This means that any absorber or the insulating epoxy layer could become electrically charged, and hold that charge, near our signals. This is another reason to use infrared absorber sparingly on the FET card, tower, and detector area.

We have developed a scheme for infrared shielding for our particular situation that takes into account our limitations. Basically, we divide our apparatus into radiation spaces. Some of these spaces are similar to a conventional system, i.e., the LHe can; some are unique to our system, i.e., the vacuum coaxial side connector. We make these spaces as light-tight as possible, but still regard them as leaky boxes. We focus on making the FET card and detector area especially light-tight. We place a small amount of infrared absorber in each radiation space, such that the area of absorber is much greater than the area of leaks and gaps. This is so that any radiation that enters the space will be more likely to be absorbed than to "leak" out of the space. We will discuss this strategy as it applies to our experience with the tower in the modified Oxford 75 at Berkeley.

Our first radiation space is at the FET card. We apply a thin layer of black IR absorber onto the FET card gusset above the FETs. We have combined this with one of two other attempts to seal the FET card. One method we call the envelope; the envelope was a piece of copper-Kapton sheet folded to fit around the FET card and taped shut with copper tape. The envelope was light-tight except for cutouts where the pins for the stripline connector and tower connector penetrated. The second method was to seal the FET card with the black absorber epoxy. The lampblack/epoxy mixture was smeared around the edge the card, around the connector pins, and over the gap between the gusset and the FET card. To the extent that the black absorber is actually opaque to infrared, this makes the FET card near light-tight; it leaves the gap between traces on the faces of the FET board uncovered by copper or absorber. This second method runs the risk of introducing microphonic susceptibility to the FET card by laying an insulating material over the traces.

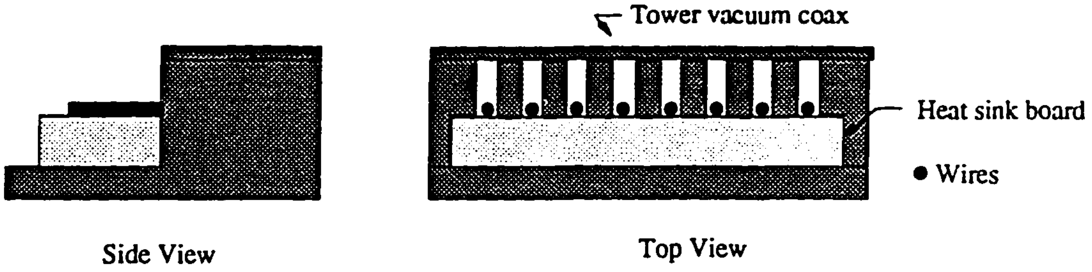

Figure 4.13 The Tower Vacuum Coax.
Each side of the tower is slotted, with wires stretched through the slots. The wires are held by soldering to heat sink boards on the LHe, Still and MC floors. The heat sink boards cover about 1/2 the slots. The top of the slots are covered by copper sheet. For clarity, only 8 slots are shown; each side of the tower has 16 slots and wires.

We use the cans as separate radiation spaces. We make small infrared absorber plates that we can bolt to the inside of the various cans. These absorber plates are pieces of copper sheet that are covered with a thick layer of roughened absorber. Each can is then fitted with a copper-Kapton sheet radiation shield that bolts to the flange and fits flush against the tower. These shields must be removed and installed with the tower much like the tower heat straps. The shields have the difficult task of making a light-tight seal against the flanges while the flanges contain other elements bolted to them: heaters, thermometers, and the heat straps. In addition, the monitoring wiring for the cans in the modified Oxford 75 are routed over the flanges. The wiring heat sinks are bolted to the flanges, and the wiring protected by a brass sheet shield. The wiring and heat sinks make a hole in any infrared shielding that is very difficult to plug. Our philosophy for shielding at the can flanges is very much the leaky box strategy. We know we are not making light-tight shields, so we bolt a few absorber plates to each flange. We are not hoping to block all of the infrared, but attenuate the total power leaking to the colder stages at each can. We install these copper-Kapton shields at each can flange: LHe, still, cold plate, and mixing chamber.

The wire channels penetrate all of the infrared shields, from the LHe can all the way to the MC can. Figure 4.13 is a diagram of a connector end of a wire channel on the tower or side coax. We treat the channels as a separate radiation space. The wires

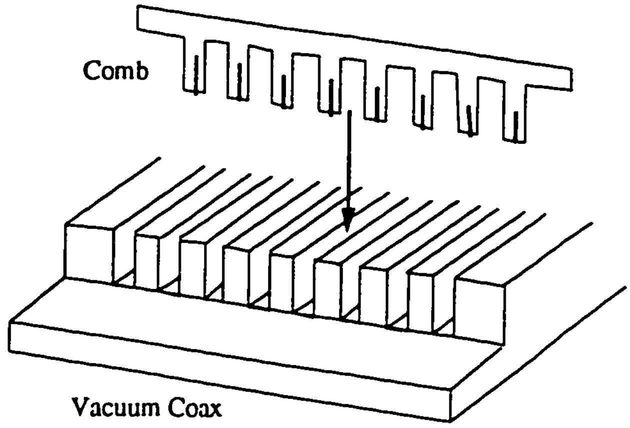

Figure 4.14 The Comb in the Vacuum Coax.
The comb is made from a copper-Kapton sheet. The comb fits into the wire channels in the vacuum coax on the tower or side connector. The comb has slits to fit over the wires, with the copper etched away from the slit to prevent the wires from shorting to the comb. The comb is an attempt to seal the vacuum coax to infrared radiation.

are stretched down the sides of the tower in the slots. The wires are held by soldering to heat sink boards at the top of the LHe floor, on the still floor, and at the bottom of the MC floor. The heat sink boards cover about 1/2 the cross section of the slots, leaving a large gap open for infrared radiation to penetrate directly into the detector area below the MC radiation shield. We cannot afford to put any infrared absorbing epoxy near the wires -- that would defeat the purpose of the vacuum coax, which is to minimize our susceptibility to microphonics noise. We have developed one strategy for blocking the slots which we call "combs". A comb is a copper-Kapton sheet cut to fit into the slots with a slit cut to fit over each wire. The copper will reflect infrared radiation entering the vacuum coax; the Kapton is too thin to absorb the infrared. Figure 4.14 is a diagram of the combs we have used in the wire channels. Combs have the undesirable property of placing a probably vibrating insulator next to the wires. Combs also require a large amount of time invested in their construction, as the wire slits must be very precise to make it worthwhile to use the combs. This requires measuring the position of the wires where they are soldered to the heat sinks to +/- 0.001". We have used combs at the top

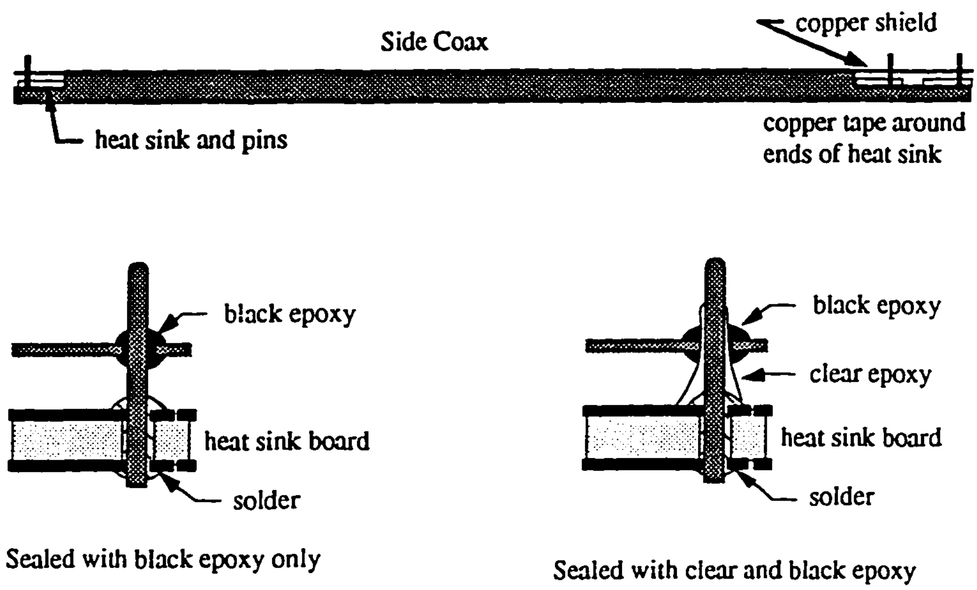

Figure 4.15 Side Coax Radiation Shields
We have used two variations of the same shielding scheme for the side coax. We seal the heat sink boards with a copper sheet and copper tape. The connector pins pierce the copper shield and the gap is sealed with black epoxy. Because the black epoxy is conductive, we modified this by using a coat of clear non-conductive epoxy on the pins before covering the gap with the black epoxy. This shielding scheme assumes the heat sink board Kapton is transparent to infrared.

of the tower, just below the LHe wire heat sink, and at the top of the side connector coax, just below the tower-side wire heat sink.

We also have installed radiation shields on the side coax. Because the side coax penetrates the MC can radiation shield, the side coax constitutes a potentially serious infrared leak into the MC can and detector area. We developed an elaborate shield scheme for the side coax before we realized the Kapton wire heat sinks were infrared absorbers, so future shielding strategies will not be as elaborate. We installed a comb by the tower-side wire heat sink. To further insure infrared was not entering the heat sink through the unmetalized gaps between traces, we wanted to install a copper shield completely around the wire heat sinks on both ends of the side coax. We made copper shields from 10 mil copper sheet to fit over the pins on each heat sink, and taped aroundthe heat sink and copper shield with counted copper tape. To keep the pins from

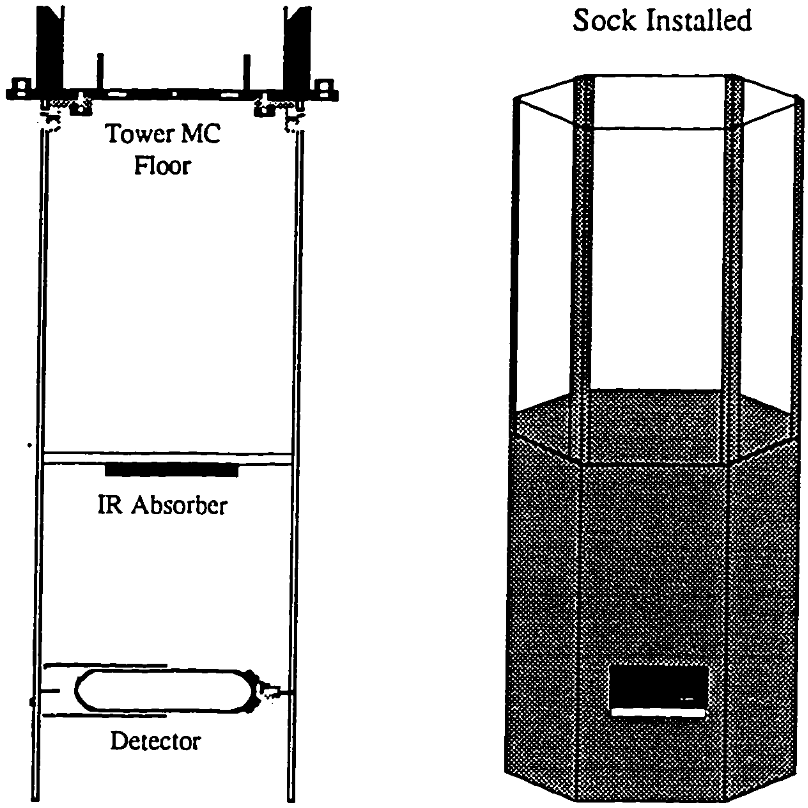

Figure 4.16 The Tower Sock.
The sock is the infrared radiation shield constructed around the detector. Pieces of copper sheet are taped onto the struts around the detector area. The top is sealed with a solid copper plate bolted to the struts; an IR absorber is bolted to the inside of this plate. The bottom is sealed with another taped copper piece. The side coaxial connector is not shown.

shorting to the shield and seal the holes, we ran a bead of epoxy around the pins. We first used our Stycast 1266 and carbon mixture to seal the pin holes, but the mixture shorted the pins to the shield, with a pin-pin resistance of $\approx 1\text{M}\Omega$ at 4 K. Our next shield used the same copper sheet shields, but we sealed the holes and coated the pins first with Stycast 1266 alone, then used the epoxy/carbon mixture to cover the pin-shield gap. The side coax shield is shown in figure 4.15.

Our final infrared radiation shield is called the "sock". The sock is built around the detectors, to create around the detectors a separate radiation space from the MC can. The detector holder ring is bolted to six struts that hold the detector holder below the tower MC floor. We bolt a solid copper plate with the same footprint as the detector ring inside the struts above the detector. This plate seals the space inside the struts. We cut thin copper sheet walls and tape them to the struts around the detector. We complete the sock with a thin copper hexagon taped to the bottom. The sock completely encloses the detector area except for a small area where the side coax plugs into the detector holder. We tape around the side coax to the sock after the side coax is installed to seal this connection. We bolt an epoxy/carbon infrared absorber to the solid copper plate at the top of the sock on the inside. Figure 4.16 shows the sock installed on around the detector holder assembly below the tower.

4.4 The FET Card

4.4.1. Design

The FET card was designed with similar constraints as the tower; i.e., radioactive purity, structural integrity, electrical properties and thermal properties. The FET card serves two purposes. First, it holds our front-end JFETs as physically close as practical to the detector at their operating temperature of ≈120 K. Having the FETs as close as possible is important to minimize input capacitance, microphonics, and cross talk and pickup problems. Second, the FET card serves as the physical interface between the tower wiring and the stripline wiring. We achieved these objectives by building a Copper-Kapton circuit board, with a window of Kapton to act as a thermal standoff to allow the FETs to be hotter than the 4K environment where the FET card is mounted.

We use Copper-Kapton circuits throughout the electronics system, from the stripline down to the Y's on the detector holder. Cu-Ka circuits are radioactively clean

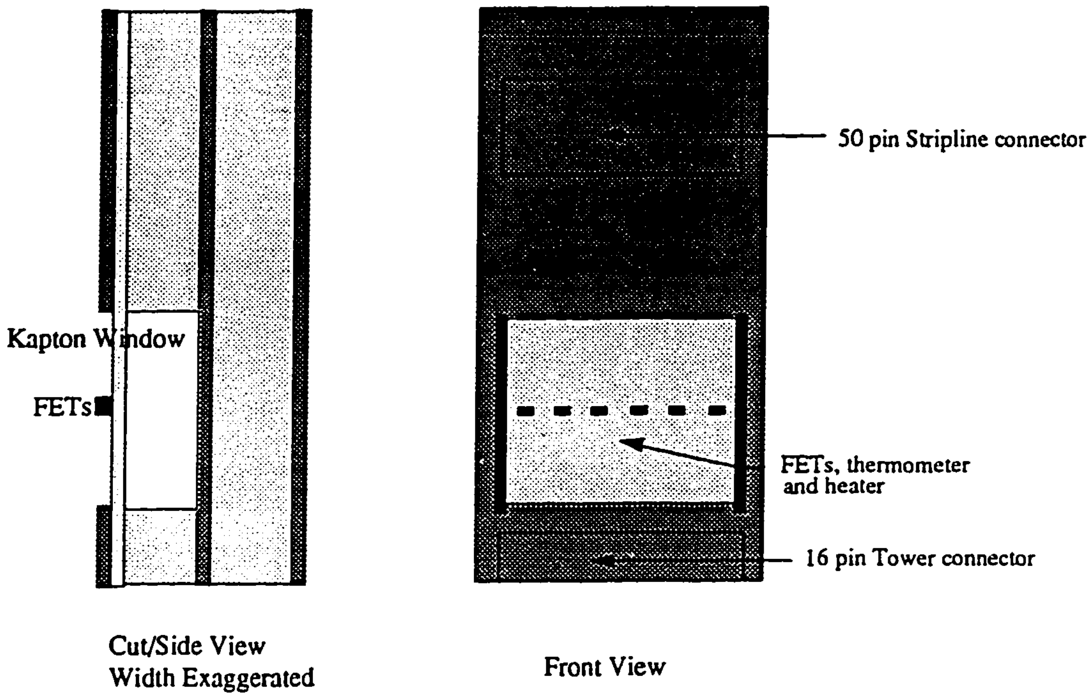

Figure 4.17 The FET Card.
The FET card is made by laminating a stack of Kapton and copper. The two outside faces are patterned with traces; the center copper acts as a ground plane. The FETs sit on a Kapton window over a hollow. The sides of the window are slit.

and not susceptible to microphonic noise like Cu-G10 circuits are. Although Kapton only comes in thin sheets (up to 5 mil thick), it can be layered with adhesive and bonded together to form sheets of arbitrary thickness. We used lay-ups of 5 mil Kapton to create a nominal 50 mil thick substrate to build the FET card.

Figure 4.17 is a drawing of the FET card. The "top" of the FET card has the 50 pin connector that mates to the stripline; the "bottom" has a row of 16 pins that mates to thetop of the tower. The FETs are mounted on the Kapton window in the center of the "front" of the FET card. The window is patterned to with contact pads on which to mount the FETs, heater, and thermometer. Electrical contact between the pads and the traces on the FET card are made by sputtered platinum traces on the window. A thin copper trace passes under the FETs. This trace acts both to make the center of the window an isotherm and to fix the potential under the FETs. The trace is connected to ground through two sputtered traces. The window also carries a chip resistor to allow us to heat the window, and a thermometer (the base-emitter junction of an npn transistor) to allow us to monitor the temperature of the window.

The FET card is manufactured from two 50 mil lay-ups of Cu-Ka-Cu, and one 5 mil Cu-Ka sheet. The first Cu-Ka-Cu stack has its copper on one side patterned and etched to be the back side of the FET card. The other side will act as a ground plane through the center of the card, and is patterned to etch copper away from where the connector pins penetrate the FET card. This is to prevent the pin holes from shorting to the center ground plane. A second stack of Kapton only is made and pressed, and a rectangular hole is cut through the stack where the window will be. A separate piece of 5 mil Cu-Ka is patterned and etched to be the front side. Now the three pieces are put together and pressed. Finally, the holes for pins and grounding the faces of the FET card are drilled and gold plated.

4.4.2. Thermal Design

The FET card is conceptually much simpler than the tower. In the case of the FET card, however, we are not worried about minimizing power flow to a cold stage. The FETs dissipate a given amount of power, and all their power will flow onto the LHe floor of the tower. Instead, we must arrange the conductance between the hot and cold areas of the FET card such that the thermal power flowing from the FETs is sufficient to hold the FETs at their operating temperature of ≈120 K. Our previous method of holding FETs in a dilution refrigerator (used in the Berkeley 75 μW refrigerator before we modified it) consisted of mounting each FET on a small fiberglass tube, and soldering discrete wires between the FET leads and the refrigerator wiring. We wished to design a system more reliable, more suited to mass production, less susceptible to microphonics and less radioactive.

4.4.2.1. Kapton Window

We have measured the thermal conductance of 5 mil Kapton over a wide range of temperatures. For the design of the FET card, we measured the thermal conductance of Kapton at temperatures from 4 K to 165 K. From this we find the value $\kappa = 9.2 \cdot 10^{-3} T^{1.02}$ [W/m/K]. We want to use a stretched Kapton membrane as both structural element to hold the FETs and thermal standoff. The power dissipated by the FETs should supply the heat necessary to hold the FETs in their operating range. Our current amplifiers dissipate about 8.9 mW per FET. The power dissipated is simply the product of the current pushed through the FET and the drain-source voltage. These are both set by the amplifier circuit, so the power dissipated by each FET can be changed easily. We have not done a complete study of the effect of the bias conditions of the FET and temperature on the noise of the FET. With four FET's per card, this gives us a total of 36 mW. In addition to the FETs, we would also like to have a diode thermometer and a resistor for a heater mounted with the FETs.

The width of the FET card is constrained by the size of the tower face. The width of our Kapton window must be wide enough to hold 4 FET's, a diode and a resistor. This means our window takes up most of the width of the FET card. We can then use our thermal data for Kapton, and this width, to decide on the height of the window.

We need to satisfy the equation where P is the power dissipated by the FETs, A, the area of thermal conduction, is the window width times the Kapton thickness, l is the length of the conduction path, g and α are the measured parameters for Kapton, T_{hot} is the temperature at which we wish the FETs to operate, and T_{cold} is the temperature of the rest of the FET card, assumed to be 4 K. T_{cold} will in fact be greater than this since the top of the tower will heat as the thermal power flows across its heat sinking straps to the LHe can, but this will not significantly affect our calculation. The FETs will be mounted in the center of the window; this means the window is effectively twice as wide for our calculation. We find l = 1.1", a convenient size.

4.2.2.2. Platinum Traces

Conventional circuit board imaging and etching limits one to trace widths greater than 5 mil. Copper traces this width on 1 oz copper would short-circuit our window as a thermal standoff. We need to have these traces fairly low in resistance to keep their Johnson noise from dominating the FET noise, but not so low that they become the dominant thermal conduction path between the FETs and the cold parts of the FET card. A noise analysis of the amplifier suggests that we want the resistance of the gate and source leads to be around 1 Ω. The drain trace can tolerate being a bit higher.

These resistances are too high to be easily realized with conventional copper etching techniques, so we decided to sputter a conductor onto the FET card window. The FET card is sputtered with traces of 30,000 Å thick Platinum on the window. The gate and source traces are nominally 40 mil wide; all others are 20 mil wide. The actual traces are a few mil wider due to spreading under the sputter mask. This gives a resistance of 1 Ω for the wide traces at room temperature, and 2 Ω for the narrow traces.

We can estimate the thermal conductance of the traces using the Wiedemann-Franz law. The Wiedemann-Franz law relates the thermal conductivity of conduction electrons to their electrical conductivity. We should remember that the W-F law is not very accurate for the temperature range in which we are interested. For temperatures in the range $\theta_D/10 < T < \theta_D$, the W-F law typically over-estimates the thermal conductivity of metals because of low angle scattering effects that can greatly affect thermal conduction, but doesn't affect electrical conduction. So we will take our estimate of the conductivity of the sputtered traces as an upper limit. We would like to find that the sputtered leads do not have a significant effect on the thermal design of the Kapton window.

The thermal power that flows is inversely proportional to the electrical resistance, and proportional to the temperature on the hot side squared. For a conservative estimate, we consider a trace of 1 Ω at 120 K. We find P = 0.18 mW. The FET card has eight 1 Ω traces and ten 2 Ω traces, all in parallel, so the power down the traces will be 13 times the 1 Ω result, 2.4 mW. This is only 8% of the nominal power that should flow through the Kapton window. At these temperatures, the Wiedemann-Franz law will probably overestimate the thermal conductivity of the platinum. In addition, the FETs have a very broad minimum in noise vs. temperature, so the actual operating temperature is not critical.

4.5. Tower Performance in the Modified Oxford 75

We have modified an Oxford 75 μW dilution refrigerator at Berkeley to use the tower system. The modification involved changing the refrigerator so that it mimics the icebox. This required the construction of cans connected to the various thermal stages of the refrigerator: liquid helium (the IVC), still, cold plate, and mixing chamber. The cans hang concentricly below the mixing chamber. Each can has a flange at the corresponding heights of the icebox lids. In addition to the cans, the fridge was fit with a tube and heat sink system (like the E-stem in the Icebox) to instrument the refrigerator, tower, and detectors using striplines instead of discrete wiring. The refrigerator currently has two striplines installed; one is used for monitoring refrigerator performance, and the other can be used to instrument the tower or to run a detector.

The refrigerator is instrumented using a system similar to the monitoring system on the icebox.[20] The monitoring system is controlled through a LabView 3.0 program connected to an electronics system containing a DAC, an ADC, a multiplexer, and the circuits necessary to bias the refrigerator thermometers and heaters. The refrigerator has thermometers on the can flanges and on the refrigerator itself. With the refrigerator monitoring system, we can measure the temperature at the MC, CP, and Still flanges; and at the refrigerator itself at the MC, CP, and still. The LHe can is immersed in the liquid helium bath, so we assume the can and its flange are at 4.2 K. The refrigerator also has metal film resistors mounted on the can flanges and refrigerator MC to act as heaters. The monitoring system also allows us to bias these heaters to put known amounts of thermal power on the different parts of the refrigerator.

Some nomenclature will be helpful in sorting through the following. We label thermometers with a two letter name: the first letter refers to the temperature stage where the thermometer is mounted (Mixing chamber, Cold plate, Still, Helium); the second letter refers to where the thermometer is mounted (Tower, Flange, Upper part of the refrigerator); because the mixing chamber temperature stage often has more than one thermometer mounted, these are distinguished by the addition of a number. Thus "MF2" refers to the #2 thermometer mounted on the mixing chamber flange. Heaters are named in the same way, except that they have the suffix "H" added. A heater mounted on the cold plate floor of the tower would be labeled "CTH."

4.5.1. Performance of the Modified Oxford 75

The modification of the refrigerator introduced many new joints into the thermal system between the can flanges, where heat would be applied, and the refrigerator, where this heat would be sunk. Before testing the thermal performance of the tower, we need to understand the thermal performance of our modified refrigerator.

The thermal impedance between the flanges and the refrigerator turned out to be non-negligible for the three inner cans (mixing chamber, cold plate, and still.) For the mixing chamber can, much of this impedance came from a poor design of the joint

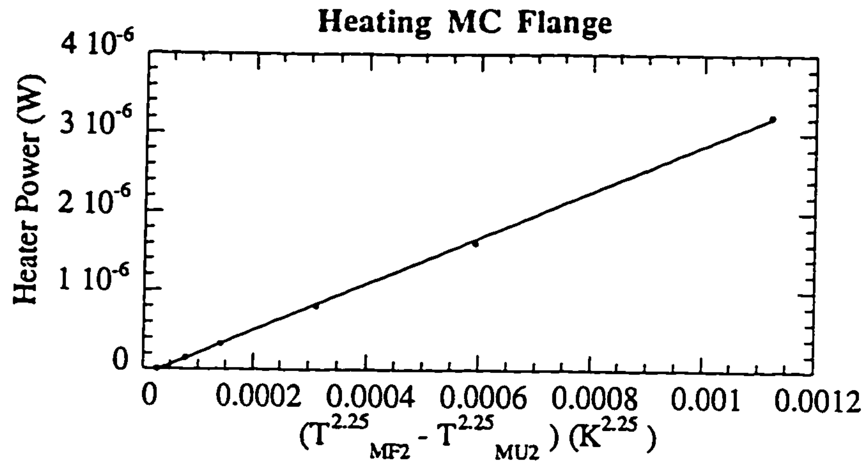

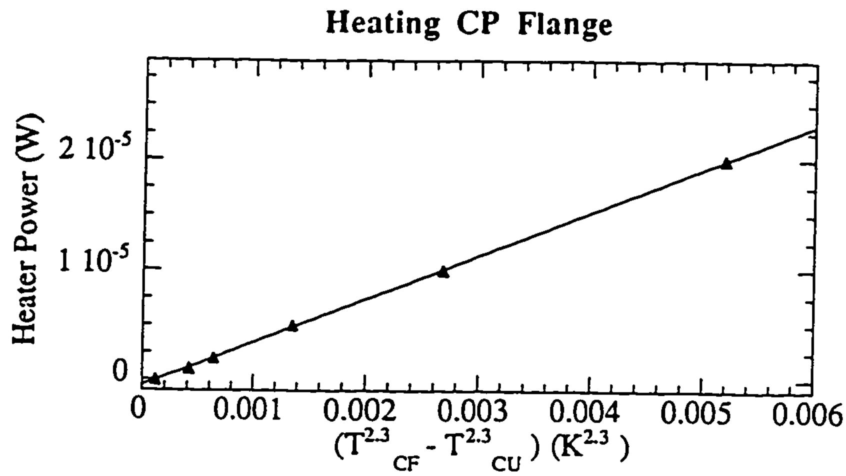

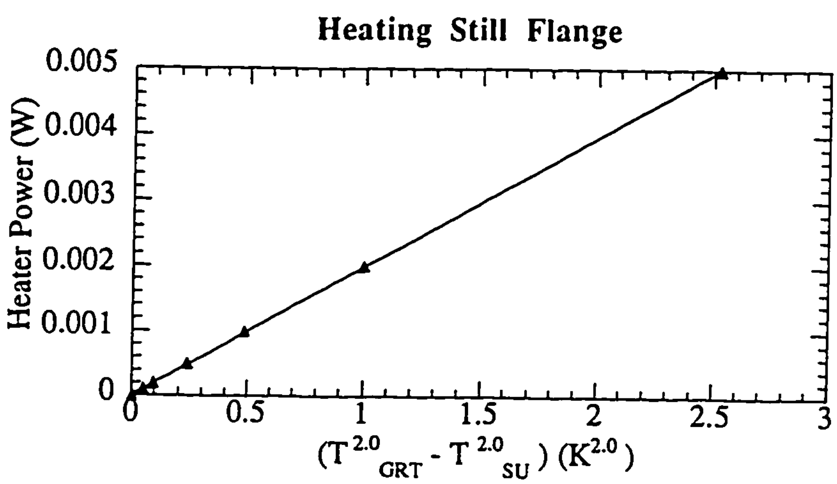

Figure 4.18 Measuring the Conductance of the Modified Refrigerator Cans
We heat the flange, measure the temperature on each side of the can, and fit the power.

Table 4.2 Thermal Conductance of the Modified Oxford 75 Cans

	G $(W/K^{\alpha+1})$	α
MC can	$2.90 \cdot 10^{-3}$	1.25
CP can	$3.93. \cdot 10^{-3}$	1.3
Still can	$1.97. \cdot 10^{-3}$	1.

where the can screwed into the mixing chamber. This was improved with a new design after run 135. The impedance in the other cans is probably dominated by the flanges themselves; they are made from brass. The large impedance of the cans is not what we desire for a refrigerator, since we want to use the cooling power of each stage of the refrigerator to sink thermal power from the tower. The impedance does, however, provide a good method of testing the tower. We can calibrate the impedance of the flange-to-refrigerator system, and use this to measure the amount of power the tower dumps at each flange.

We spent run 124 calibrating the thermal impedance of the cans. To do this, we apply a known power on the flange using the installed heaters, and measure the temperature at the flange and at the refrigerator. We then fit the measured temperatures and powers to the function

$$P_{app} + P_0 = G \cdot \left(T_F^{\alpha+1} - T_U^{\alpha+1}\right), \tag{4.9}$$

where P_{app} is the heater power we apply, T_F is the flange temperature, T_U is the upper refrigerator temperature, P_0 is some constant power on the flange, and G and α describe the thermal properties and geometries of the flange-upper can system. Figure 4.18 shows a measurement of the thermal impedance of the refrigerator cans.

P_0, G, and α are parameters to be fit. Table 4.2 lists the values we find for G and α for the three cans. P_0 is a power that apparently flows onto each flange with no applied power. This apparent power could be a real power input on the cans, as it seems to monotonically decrease with time as the run progresses. This is a consistent feature of our refrigerator, before and after the modifications. A good benchmark of this effect is that the mixing chamber base temperature cools ≈ few mK over ≈ week of running time. This could also arise through a miscalibration of our thermometers; if they do not read the same temperature when in fact they are at the same temperature, we will attribute the apparent temperature difference to an extra power. This is not likely to be a large effect on the MC and CP cans. The thermometers used were designed to be used at these temperatures and calibrated carefully in the range [18mK, ≈300 mK]. The Still thermometers are more problematic. We originally installed RuO_2 resistor thermometers on the Still flange and Still upper. These RuO_2 thermometers are not very sensitive at the temperatures of the still and still flange unless they are specifically calibrated over the temperature range [600 mK, 2 K]. For this measurement of the Still can conductance, we replaced the RuO_2 resistor at SF with our calibrated Lakeshore GRT. The Lakeshore GRT has been calibrated over the temperature range [50 mK, 4.5 K] and is our temperature standard over that range. We can also check the calibration of the RuO_2 thermometer at SU using the Lakeshore GRT mounted on the Still flange. We heat the still with a heater mounted on the still directly; the still flange temperature should equal the still temperature if no extra power is added to the still flange. After calibrating the SU thermometer, we measured the conductance of the three cans.

4.5.2. Thermal Performance of the Tower in the Modified Oxford 75

We wished to test the performance of the tower in the modified Oxford 75; we need to measure the thermal conductance of all the parts of the tower system, the heat straps and graphite tubes. The FET card was thermally tested separately first in an InfraLabs LHe dewar to ensure its design worked. To effectively measure the thermal performance of the tower, we need to have appropriate thermometers on each part of the tower and refrigerator system, and we need a way to introduce heat to each part of the system. This will allow us to measure temperature gradients across the heat straps and tubes in response to applied power.

4.5.2.1. Tower Thermometer and Heater Configuration

We mounted thermometers and heaters on each floor of the tower. The heaters were all 1 kΩ metal film resistors epoxied onto small pieces of copper sheet. The copper sheet is then bolted with an M3 screw onto the given tower floor. Metal film resistors were used because their resistance is roughly temperature independent and does not change with temperature cycling. The resistance of these resistors typically decreases less than 10% as they cool from room temperature down to 4 K or colder. The resistance of the wiring changes more than this, so we have to have independent measurements of the wiring in the system. Once we have measured the cold resistances of whatever wiring system we are using (stripline or discrete wires), we can measure the total resistance through the entire system (wires + heater) and calculate the heater resistance. We set and measure a voltage outside the refrigerator to bias the heater; we need to know the wire and heater resistance to know how much power our heater is dissipating in our system.

We used a variety of thermometers for monitoring the temperature of the tower floors. We used a Lakeshore Cryogenics DT-470-MT diode thermometer on the LHe floor of the tower. This thermometer operates over the range 1.4 K - 475 K. The diode is biased with a 10 μA current source and the voltage across the diode is measured. This diode thermometer is calibrated to a standard v(T) curve given by Lakeshore. We check the calibration of this thermometer at 77 K (LN) and 4.2 K (LHe). This thermometer was used as our calibration standard against which we calibrated the other diode thermometers we use for temperatures above 4 K.

We used a Lakeshore Cryogenics GR-200A-30 Germanium Resistance Thermometer (Lakeshore GRT) to measure the temperature of the still floor. This is the thermometer we previously used on the still flange to calibrate the still thermometers. This thermometer is calibrated in the temperature range 50 mK - 25 K. This thermometer has been our calibration standard over this temperature range. The Lakeshore GRT was bolted to the still floor through the access hole in the LHe floor.

The CP floor is difficult to instrument; we only easily have access to the two ears that extend outside of the tower. One ear is used to bolt to the CP heat strap. We constructed a small heater/thermometer package that could be bolted to the other CP floor ear, and still fit inside the modified refrigerator cans. The thermometer used was DF95A, a RuO_2 resistance thermometer. This thermometer was calibrated in refrigerator Run 95 against the Lakeshore GRT over the temperature range 20 mK - 400 mK.

We used GT1 as the thermometer on the MC floor. GT1 is a piece of NTD 23 germanium mounted in a copper casing. The NTD germanium looks much like one of our phonon sensors on our detectors. GT1 was calibrated against the Lakeshore GRT for temperatures above 50 mK, and was calibrated using ^{60}Co thermometry over the temperature range 18 mK - 25 mK. The specific R(T) for GT1 is such that is useless as a thermometer above about 100 mK -- its resistance drops quickly to $\approx 10\ \Omega$ around 100 mK.

All of the thermometers and heaters were mounted in the same way. The surface in contact with the tower was smeared with a small amount of vacuum grease and bolted with an M3 bolt to the tower. All thermometers were monitored using four-wire measurements. The bias point for each of the thermometers was adjusted individually before each new temperature to ensure a good measurement. At low temperatures, we must be certain we do not over bias the thermometer and heat it significantly above the temperature we are trying to measure. At higher temperatures, the resistance of both RuO_2 and NTD Ge thermometers decrease, so we must increase the bias to get a

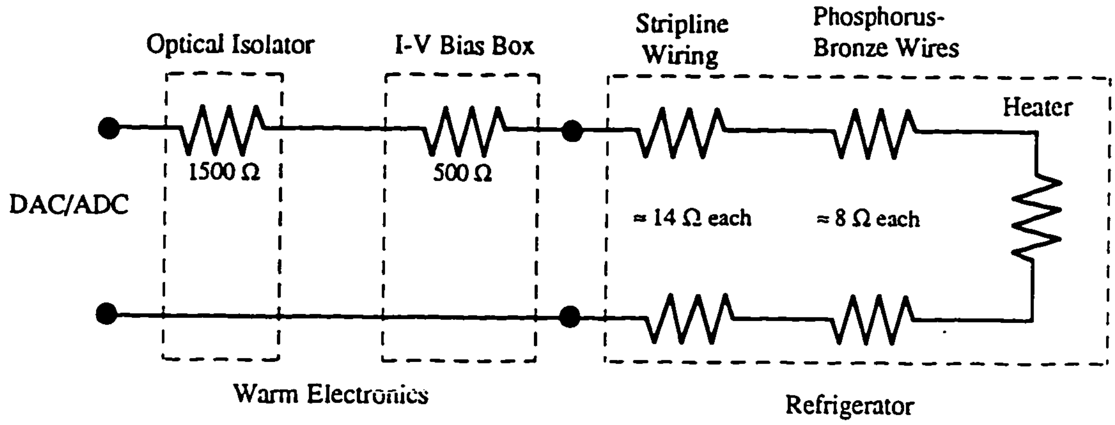

Figure 4.19 Heater Bias Circuit
We set and measure a voltage outside the refrigerator using a DAC/ADC controlled with LabView. The bias current goes through an optical isolator with 1500 Ω in series, a bias box with 500 Ω in series, the refrigerator wiring, and finally the heater.

measurable signal. Electrical contact was made using 5 mil phosphor-bronze wires between the heaters and thermometers and the refrigerator wiring harness at 4 K. Wires were heat sunk by wrapping and gluing them around a copper post that was then bolted to the same floor as the thermometer and heater. Heat conduction along these wires was negligible.

We will refer the heaters and thermometers using the nomenclature of section 4.5. Just to remind us, MT1 refers to the #1 thermometer mounted on the tower MC floor. CFH refers to a heater mounted on the CP flange.

4.5.2.2. Tower Thermal Measurements

It is easiest when examining the tower thermal performance to start with the MC floor and proceed sequentially to the warmer stages. This allows us to measure the thermal properties of the lower temperature stage and include its measured characteristics in the analysis of the higher temperature stage. This will also allow us to recheck our calibration of the thermal impedance of the cans. The thermal impedance of the MC can in particular varied between different refrigerator runs. This is because the main thermal impedance in this can seems to have been a loose screw joint, where the can

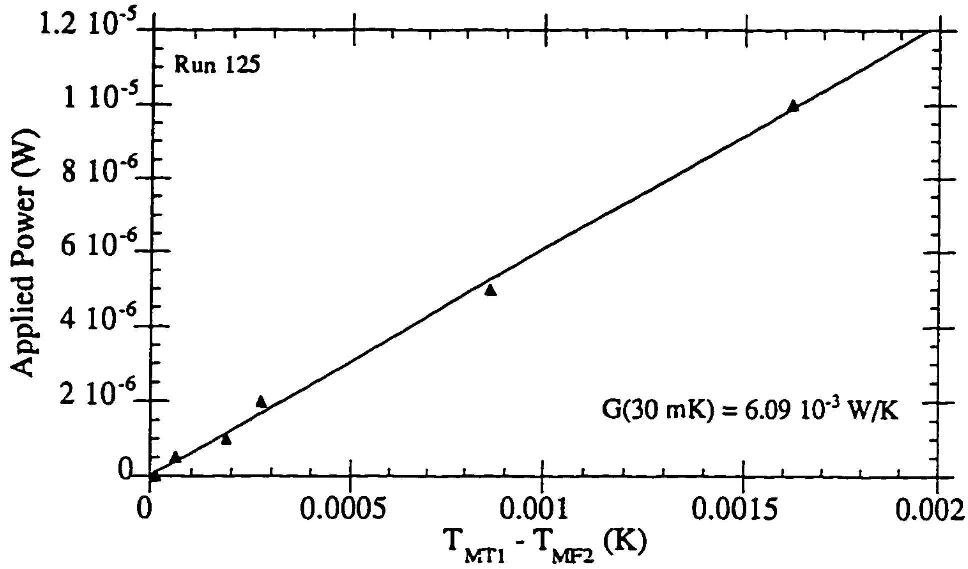

Figure 4.20 Tower MC Floor/MC Flange Joint Impedance
We apply heat to the MC floor using MTH and measure the temperature difference between MT1 and MF2. The temperature difference between the floor and flange is small enough that we linearize the conduction law to $P = G\,\Delta T$.

screws into the mixing chamber. This problem was fixed with an improved joint design after the measurements presented in this thesis. First we measure the thermal impedance of the tower/MC flange joint and the thermal impedance of the MC can by heating the MC floor of the tower using MTH. We monitor the temperature of the MC floor using MT1, the temperature of the MC flange using MF2, and the temperature of the MC refrigerator upper section using MU2. We apply heat to the MC floor by applying a voltage across MTH. We set and measure this voltage outside the refrigerator using a DAC/ADC controlled by a LabView routine. We need to understand the bias circuit to relate this voltage set and measured outside the refrigerator to the actual voltage across our heater. Figure 4.19 shows the actual bias circuit used for all of the heater measurements. We can easily see the power dissipated by the heater will be

$$P_{app} = \left(\frac{V_{app}^2}{R_{tot}^2}\right) \cdot R_H, \tag{4.10}$$

where P_{app} is the power dissipated by the heater, V_{app} is the voltage at the DAC/ADC, R_{tot} is the total resistance of the bias circuit, and R_H is the heater resistance. In practice, because the heaters are only connected with two wires, we measure the total resistance of the heater and the refrigerator wiring, and rely on separate measurements of the resistance of the stripline traces and other wiring. These measurements must be done cold because the resistance of all of the cold elements change as the refrigerator cools from room temperature to base. For the measurements we will present, $R_{tot} = 3030\ \Omega$ and $R_H = 983\ \Omega$, so $P_{app} = 1.07 \cdot 10^{-4}\ V_{app}^2$. Figure 4.20 shows a measurement of the tower MC floor/Flange joint. This joint consists of six ears bolted to a copper flange. From our 4 K resistance measurements of the heat straps, we expect a copper-copper bolt joint to be about $5 \cdot 10^{-7}\ \Omega$. We can use the Wiedemann-Franz law to estimate the thermal conductance of our six bolt joints; we find G(30 mK) = $9 \cdot 10^{-3}$ W/K, in fair agreement with the measured value of G(30 mK) = $6 \cdot 10^{-3}$ W/K shown in figure 4.20. This measurement was made before we gold plated any tower ears. We expect this to improve with gold plating.

We can use the same measurements to calibrate the MC can thermal impedance. We measure the temperature at MU2 and MF2 and correlate these with the applied power. Figure 4.21 shows the measurement of the MC can thermal impedance for Run 125. The conduction found from this data is $P = 1.7 \cdot 10^{-3}\ (T_{MF2}^2 - T_{MU2}^2)$ [W], in contrast to the previous relation $P = 2.9 \cdot 10^{-3}\ (T_{MF2}^{2.25} - T_{MU2}^{2.25})$ [W]. This change in the MC can thermal impedance can be attributed to a loosening of the screw joint that fixed the MC can to the mixing chamber. This joint was redesigned and fixed in May 1996 to address this problem.

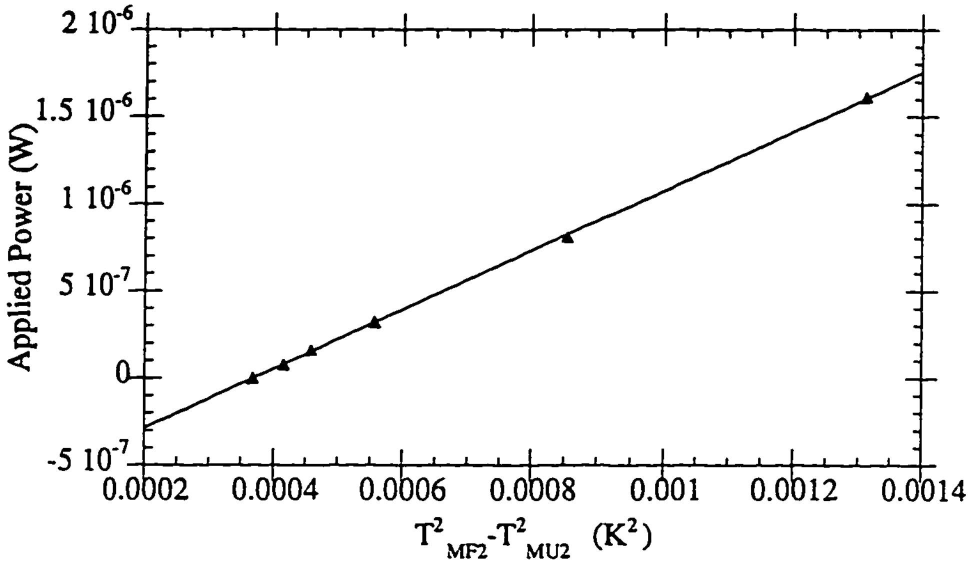

Figure 4.21 Run 125 MC Can Thermal Impedance
The thermal impedance of the MC can changed from our original measurement. This measurement finds P= 1.7 10^{-3} (T_{MF2}^2 - T_{MU2}^2). This change is due to a loosening of the screw joint that attaches the MC can to the mixing chamber.

We now heat the CP floor using CTH with a known power, P_{app}. We measure the temperature at MF2 and MU2 and use this to calculate the power flowing across the MC can, P_{MC}. P_{MC} must flow onto the MC floor through the graphite tube separating the MC and CP floors, so by measuring the temperature at MT1 and CT we can calculate the thermal conduction of the graphite tube. We measure the temperatures at CF and CU and compute the power flowing across the CP can, P_{CP}. PCP must flow onto the CP can through the CP heat strap, so by measuring the temperature at CT and CU we can calculate the thermal conduction of the CP heat strap. In a similar manner, we can measure the thermal conductance of the other graphite tubes and heat straps.
Measuring the conductance of the LHe floor heat strap is a little different. The LHe strap sinks the power on the LHe floor to the LHe can flange. This can is also the inner vacuum can (IVC) of the refrigerator, and sits in the LHe bath. Since the LHe flange is near the bottom of the IVC, it is immersed directly in the bath, even when the bath is

low. We have no way of measuring directly the power that flows onto this can, so we make the assumption that the flange is always at 4.2 K. We can sum the powers flowing onto the MC, CP and still cans; the difference between this sum and the power we apply to the LHe floor must be the power that flows through the LHe strap onto the LHe flange. In practice, the power flowing onto lower floors is much less than the applied power. We calculate the conductance of the LHe strap using this derived power and assuming the flange temperature is 4.2 K.

Figure 4.22 shows the measurement of the thermal conductance of the graphite tubes from Run 125. The graphite conductance is very close to measurements shown in Pobell, but about a factor of ten higher than another sample we had measured previously. This sample proved to radioactive to use in the tower, but we should keep in mind that we may be able to find clean samples of graphite with lower thermal conductivity than we use currently.

4.5.2.3 Tower Heat Straps

Each stage of the tower should act like an efficient thermal divider, diverting the thermal power flowing onto the floor through the strap onto the refrigerator. If the strap can sink the power without a significant temperature gradient across it, then the tower floor temperature will be unaffected, and the power down the tube will not change much from its quiescent value. This quality of a good heat strap makes it hard to measure the actual thermal conductance of the strap. The straps work well enough that we are limited by our thermometry in the Berkeley refrigerator. We will rely on the 4 K resistance measurements to check the performance of the heat straps.

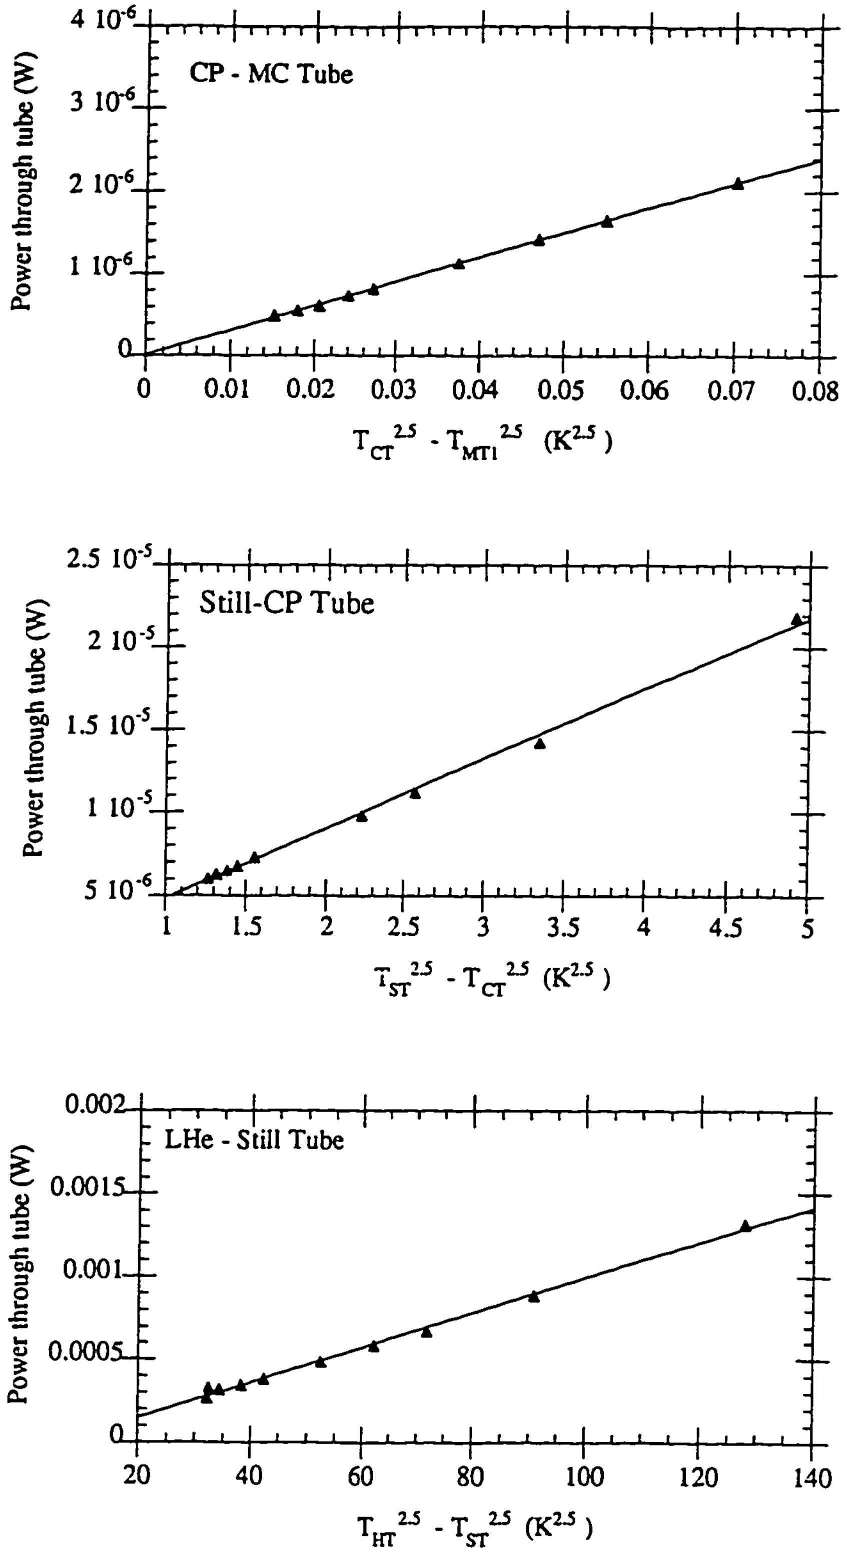

Figure 4.22 Thermal Conductance of the Tower Graphite Tubes
Shown are the thermal conductance measurements of the actual graphite stand-offs used in the tower.

4.5.3. Infrared Radiation in the Tower

One issue we did not adequately anticipate for in this first tower design is infrared radiation from the FETs mounted on the FET card. Let us estimate the effect of having 120 K FETs in close proximity to our detectors at 20 mK.

We have four FETs, one diode thermometer (actually an npn transistor packaged much like the FETs), and one chip resistor, all presumably at the same operating temperature of 120 K when the FETs are running. In addition, the FET card has a copper strip that under these components, keeping them all well coupled thermally. We use the law for a radiating black body to calculate J, the power per surface area radiated by these components:

$$J = \sigma_B T^4, \text{ where } \sigma_B = 5.67 \cdot 10^{-8} \frac{W}{m^2K^4}. \qquad (4.11)$$

The total exposed surface area of the components is $\approx 3.0 \cdot 10^{-4}$ m^2. We will ignore the copper strip, as its emissivity should be much less than that of the black components. We now calculate the total power radiated by the 120 K components on the FET card, P_{RAD}, using the above area, and equation 4.11. We find the remarkable result P_{RAD} = 3.5 mW! This is a significant fraction of the total power, 36 mW, dissipated by the FETs at their operating bias condition. We can also calculate the peak frequency emitted by a blackbody at this temperature:[32]

$$hf_{max} \approx 2.82 \cdot k_B T, \qquad (4.12)$$

where f_{max} refers to the frequency carrying the most energy per unit frequency range for a black body peak at temperature T. We note that this peak frequency corresponds to photons with an energy of 29 meV. This energy in uncomfortably close to the energy gap separating impurity states from the conduction band in our germanium detectors, $\Delta E_i \approx 10$ meV. We can conclude from these calculations that infrared radiation from the FETs constitutes two distinct problems. First, the large amount of power radiated in the infrared, O(mW), if it is absorbed in the cooler parts of the refrigerator, can easily overwhelm the cooling power of those stages. In particular, this is a factor 10^3 higher than what the mixing chamber can tolerate. Second, even if most of the total power of this infrared radiation is blocked from reaching the detector, individual IR photons will interact with the detector, and can disturb the charge measurement by reionizing impurity sites.

Table 4.3 Infrared Shield Configuration at Berkeley by Run

Run	Configuration				
	Can shields	FET card	Combs	Side Coax	Sock
127	crude	no shield	none	no shield	no
128, 129	final	"	yes	"	no
130	"	envelope	"	black seal	yes
131, 132	"	"	"	clear & black seal	"
133	"	black seal	"	"	"
134 - 136	"	"	none	"	"

4.5.3.1. Experience with Infrared in the Modified Oxford 75

We have run with the tower and E5 in the modified Oxford 75 with a variety of shields and shielding strategies. We have recognized the importance of infrared shielding from experience, and have seen infrared from the FETs affect both the temperature of the detector and the charge measurement of the detector. The most sensitive test of the shielding is measuring the amount of power falling on the detector.

Table 4.3 describes the infrared shielding situation broken into five categories for our test runs in the Berkeley refrigerator. "Can shields" refers to the copper-Kapton shields that cover the space between the can flanges and the tower; we use these on each can. The "crude" shields were not as carefully made as the "final" shields, and were probably much less light-tight. "FET card" refers to the shielding technique used on the FET card itself. Both the "envelope" and "black seal" include absorber on the inside of the gusset. "Combs" refers to the presence of combs on the LHe tower heat sink and side coax tower side coax. "Side coax" refers to any treatment of the side coax heat sinks. "Sock" refers to the presence of a sock infrared shield built around the detectors.

We saw in our first runs with E5 in the tower with minimal radiation shielding that the charge measurement was being ruined by infrared. The charge measurement should be relatively insensitive to a small power load on the detector; the charge measurement was being affected by large numbers of IR photons with an energy greater than 10 meV $\approx \Delta E_i$. We could see this as a very gross effect by measuring the I-V characteristic of the crystal with the FETs on and the FETs off. Figure 4.23 compares the leakage current across the detector as the bias voltage is varied when the FETs are off to when the FETs are on. With the FETs off, the detector shows a clear transition in leakage current at $\approx$ +0.5 V, -0.4 V. In contrast, with the FETs on, the detector has no transition, and conducts almost 100 times the current as when the FETs are off. This data is from run 127 with minimal radiation shielding.

Each run with E5 included an ^{241}Am source collimated and shielded to block the low energy γ's so that we should only see the 60 keV line from the source. In run 127, the detector was too hot to see a phonon spectrum, and we were unable to see the 60 keV line in the charge measurement. In run 129, the detector was again too hot to use the phonon measurement, but we could see the 60 keV line in the charge measurement. The width and stability of the 60 keV peak was much worse than previous measurements using E5 in the unmodified refrigerator. We could also see a large population of low energy events. This is consistent with a large number of >10 meV photons being absorbed in the detector and reionizing the impurity sites in E5. This would tend to create regions of space charge that would impede the full collection of charge from the 60 keV events

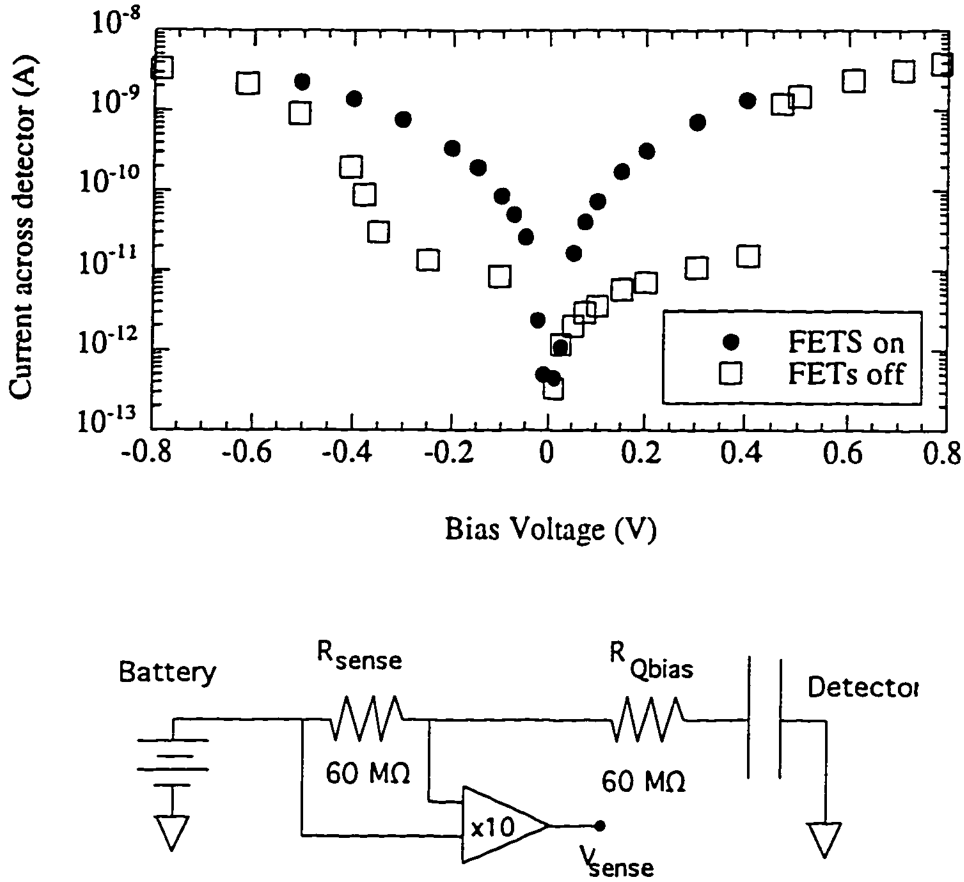

Figure 4.23 Measurement of the Leakage Current Across E5.
Shown is a measurement of the leakage current across E5 with the FETs off and the FETs on. With the FETs off, we can see a definite breakdown voltage for the crystal; with the FETs on, the leakage current is almost 2 orders of magnitude larger below the breakdown voltage when the FETs are on. This data is from run 127 with minimal radiation shielding.

by trapping the drifting charge and reducing the drift field in the detector. After we improved the infrared radiation shielding in run 130, we saw no degradation of the charge measurement.

A more sensitive test of the infrared being absorbed by the detector can be made using the thermistors on the detector. We had measured the conductance between the detector and the detector holder for E5 when it was mounted in our old detector holders for the unmodified refrigerator. We heat sink the crystal through a gold wire bonded to a gold pad sputtered onto the detector. We expect the impedance of the crystal heat sink to be dominated by the impedance of the gold-germanium interface,[30] so the detector-tower thermal impedance should be the same as the previous measurement. We find the values $G_{ds} = 4.7 \cdot 10^{-4}$ and $\alpha+1 = 3$, where $P = G_{ds}(T_{s1}{}^{\alpha+1} - T_{MT1}{}^{\alpha+1})$, and T_{s1} is the temperature measured by the sensor S1 on E5. We can also measure the thermal impedance of the MC flange to MC upper thermometers each run and use MF2 and MU2 to measure the power flowing onto the MC flange. We can measure the power onto the detector and onto the flange as a function of the FET temperature.

Unfortunately, we may not be measuring the actual temperature of the flange with MF2. MF2 is an NTD germanium thermometer we made, and is not packaged in a light-tight container. MF2 is mounted on the MC flange, and is outside the MC can radiation shield; the detector is protected by two more radiation spaces than MF2. It is reasonable to assume then that MF2 will see more infrared radiation than the detector, and this is in fact what we see -- the detector is colder than the flange. However, part of the apparent temperature difference could be due to MF2 directly absorbing infrared and heating above the flange. This effect has been seen in the icebox.[40] Apparently, RuO_2 thermometers are not as susceptible to infrared.

We can examine the effectiveness of the shielding by looking at the power flowing down the MC can (measure T_{MF2} and T_{MU2}) and the power flowing from detector to flange (measure T_{s1} and T_{MF2}). The power flowing onto the flange should be a measure of the effectiveness of the shielding above the MC can; the power onto the detector also includes the effectiveness of shielding inside the MC can, the sock and side coax shields. We can change the temperature of the FETs by adding power to the FET card heater; this should enable us to deconvolve the effect of infrared power from the FETs from other powers that affect the detector.

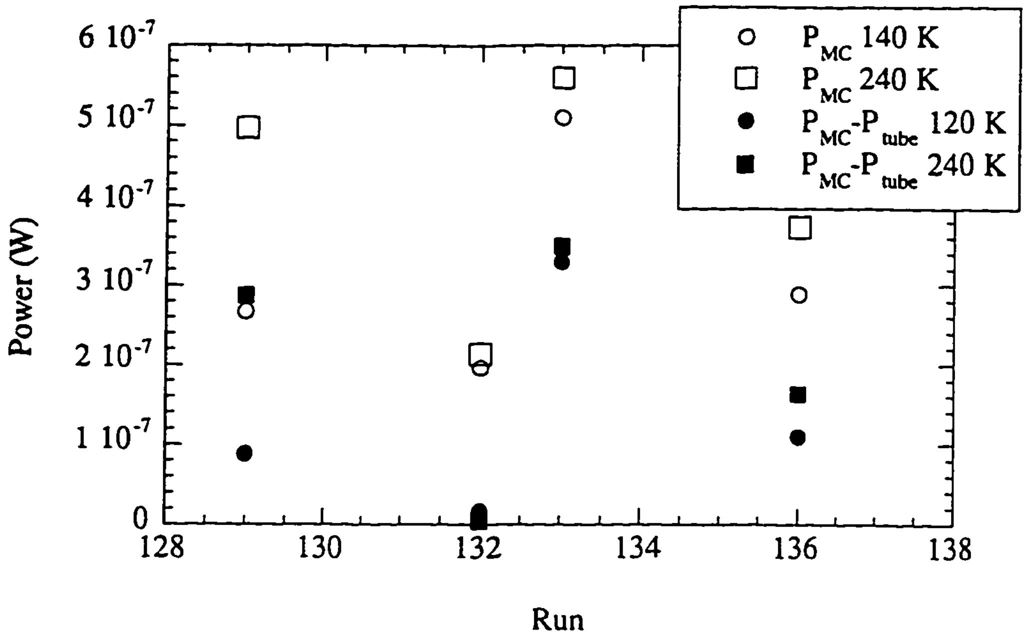

Figure 4.24 Power on MC flange from FETs.
This shows the power on the MC flange for four runs. P_{MC} is derived from a measurement of the MC can impedance each run. P_{tube} is estimated from run 129.

We start by looking at the power flowing down the MC can, P_{MC}, for the FETs at two temperatures, ≈140 K and ≈240 K. For reference, we note this is a factor of 8.6 in T_{FET}^4. We have measured values for MF2 and MU2, and thus P_{MC}, for runs 129, 132, 133 and 136. We also have a measurement of CU correlated with MF2 and MU2 for run 129; we can use this to estimate the temperature of the CP floor and the power flowing down the CP-MC graphite tube. We find $P_{tube} \approx 1.8 \cdot 10^{-7}$ W when the FETs are at 140 K, and $P_{tube} \approx 2.1 \cdot 10^{-7}$ W when the FETs are at 240 K. These estimates are probably good to about 20%, and are probably good for all the runs. The power estimate is dominated by the CP floor temperature, which should be dominated by thermal conduction down the graphite tubes and the efficiency of the heat sink straps. We believe our heat sink straps were done well by run 129. So we will assume these values for P_{tube} and subtract them from P_{MC} to find the effect of the FET radiation. Figure 4.24 shows the power on the relationship between P_{MC} and FET temperature

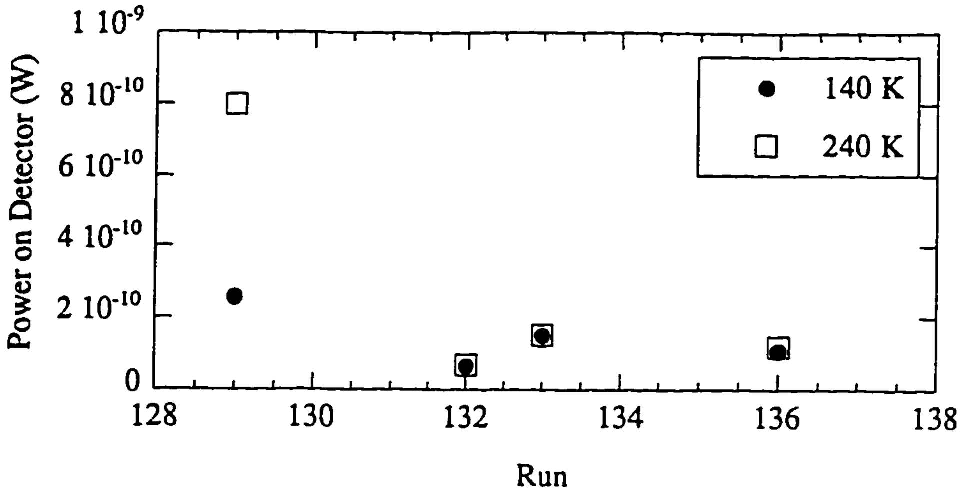

Figure 4.25 Power on E5 for Two FET Powers and Four Runs.
Shown is the power incident on E5 calculated from T_{s1} and T_{MF2} for two FET temperatures, 140 K and 240 K.

for four runs. If the quantity P_{MC}-P_{tube} really represented the power onto the MC flange from the FETs, we might expect the power on the flange to increase by $(240/140)^4 = 8.6$ when we increase the FET temperature. We can see that this is not true for any run, not even run 129 for which we know P_{MC} - P_{tube} the best. We notice there is a very small difference between the powers when the FETs are at 240 K and 120 K for runs 132, 133, and 136, while run 129 shows a large difference in the powers. If we don't see a difference in power when we heat the FETs, then the FET infrared radiation is probably not the power source heating the MC flange. This means the FET envelope installed before run 130 probably blocks most of the infrared. Note that before run 133 we switched from a FET envelope to sealing the FET card with black epoxy, and wired two more faces of the tower. The power on the MC flange was lower in run 136 than in run 133, though the only difference in shielding was that we removed the combs before run 136. This result, run 136 being better than run 133, indicates that the accuracy of the magnitude of the power flowing down the can is probably not the most trust worthy result of these measurements from run to run.

In the same manner as the above, we have measured the power flowing from the detector to the MC flange by measuring T_{S1} on E5 and T_{MF2}. We expect with no power incident on the detector that T_{S1} and T_{MF2} will be the same. We have found in all runs that this is not the case, the detector is hotter than the flange. We find that the power incident on the detector does not change much with increasing FET temperature compared to the total power. Figure 4.25 shows the power on the detector for two FET temperatures as a function of run. We see a similar result to the power on the MC flange. For run 129, with no treatment of the FET card or sock, we see a difference in power on the detector when we heat the FETs. For all later runs, with a sock and some treatment of the FET card, there is at most a 10% increase in power on the detector. This would seem to indicate that the sock and FET card seal works well in isolating the detector from the FET infrared radiation.

A large power incident on the detector not related to FET temperature (infrared) is hard to understand. However, for all of these runs we had an 0.5 μCi ^{214}Am source mounted on the struts close to the detector to use as a 60 keV line source. This source, in addition to the 60 keV gamma rays, emits 2 MeV alphas which are absorbed in a mylar sheet in the collimator. This source turns out to be a non-negligible power source, emitting about 15 nW. Since the source is mounted on the struts, this power must be sunk down the struts and through the 0-80 bolt joints that join the struts to the crown that bolts the detector sub-assembly to the tower MC floor. We expect the thermal impedance between the detector and the MC flange to be dominated by the gold-germanium interface of the detector heat sink. The thermal impedance between S1 on E5 and MF2 measured in run 129 was what we predict if the thermal impedance is solely due to the area of the gold heat sink pad on E5, $G_{ds}(20\text{ mK}) = 1.3\cdot 10^{-8}$ W/K. If the thermal impedance of the 0-80 bolt joints were significant, then the power from the

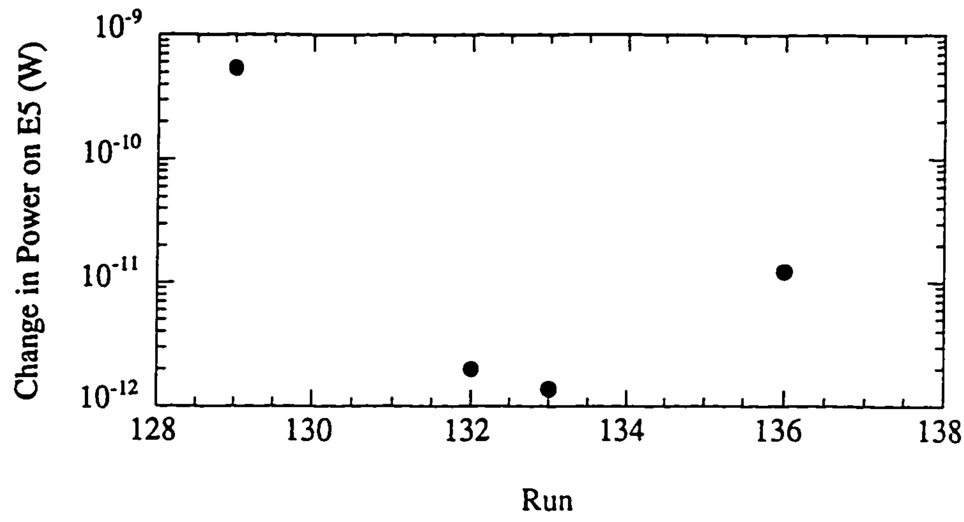

Figure 4.26 Difference in Power on E5 Between the FETs at 240 K and 140 K.
The vertical axis is the difference in power measured on E5 when the FETs are at 240K and 140 K. The difference should be $\approx(240/140)^4 = 8.6$ times the infrared power falling on the detector when the FETs are at 120K.

^{241}Am source could be heating the detector holder and E5 above the MC floor temperature. This temperature difference would then be interpreted as a power uncorrelated with the FET temperature flowing across the thermal impedance between the detector and floor. We would need $G_{bolt}(20\ \mathrm{mK}) \approx 10^{-5}$ W/K for the power from the ^{241}Am to account for $\approx$ few mK temperature difference between the detector and floor. This would imply the 0-80 bolt joints are $\approx 50\ \mu\Omega$, which is quite plausible when we consider that we can't tighten the 0-80 bolts very tight before they break and the few $\mu\Omega$ M3 bolt joints. This impedance would be unmeasurable in series with G_{ds}; we would need a heat source and thermometer on the detector holder to measure it. We would expect the temperature difference to disappear when the source is removed for installation in the icebox.

If we accept the large constant apparent power on E5 as a temperature gradient from the detector holder to the flange due to a power not related to the FET infrared (such as the ^{241}Am source), then the power difference we measure with the FETs hot and the FETs cool is meaningful. Figure 4.26 shows this power difference explicitly.

We should remember that the actual powers at a normal FET operating temperature of 120 K - 140 K will be a factor of 8.6 - 16 lower than the powers shown in figure 4.26. The precipitous drop in power between runs 129 and 132 must be due to a combination of the FET card envelope and the sock. Between runs 132 and 133 we changed from the FET card envelope to sealing the FET card with black stycast. This apparently had no adverse effect on the infrared shielding of E5. We see the power rise by a factor of 8.8 between runs 133 and 136. This must be due to the removal of the combs. This points to the importance of shielding the wire channels from infrared. This power falling on the detector in run 136 is about 12 pW/8.6 = 1.4 pW, which is comparable to the bias power of in the phonon measurement. This should be a warning that our infrared shielding is barely adequate, and we should be thinking of ways to improve shielding in the tower, particularly the wire channels.

4.5.4 Microphonics in the Tower

Microphonics is noise produced by vibrations in the experimental apparatus that is somehow coupled into the electronics chain. We have two ways of attacking a microphonics noise problem; we can reduce the source of the microphonics and we can reduce our system's susceptibility to microphonics. Microphonics sources in a dilution refrigerator include the circulation and pot pumps and vibrations in the needle valve on the pot. We believe microphonics noise arises through a capacitive coupling between our gate wire and a vibrating system at a different potential nearby. The difference in potential could be explicitly set by us, for instance if we held our gate wire at 1 V near a grounded shield, or could arise by the presence of static charge on nearby insulators. We can see this by examining the relationship between charge, potential, and capacitance:

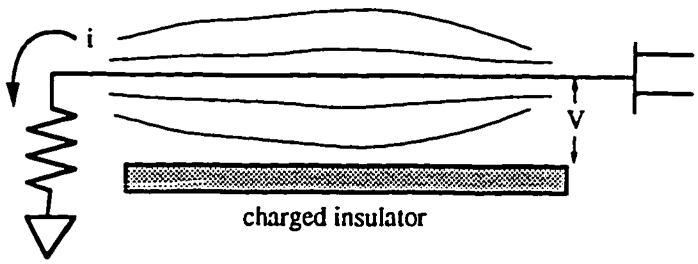

Figure 4.27 Microphonics Noise.
Microphonics noise arises as a capacitive coupling between our gate wire and a nearby object at a different potential. Vibrations drive a noise current that shows up as a voltage noise at the FET gate when the current flows across a resistance.

$$Q = C \cdot V, \tag{4.13a}$$

and its time-derivative,

$$i \equiv \dot{Q} = \dot{C} \cdot V + C \cdot \dot{V}, \tag{4.13b}$$

where Q is charge, C is capacitance, and V is potential. The first term in 4.13b is the microphonics noise; the second term is relevant to cross-talk created by pulses and our AC bias. If we ignore the cross-talk term, we see that we are looking at a current source, $\dot{Q}$, that is proportional to the changing capacitance and the potential between our wire and the capacitively-coupled system. Microphonics will induce a noise current on our gate wire that will flow across the thermistor in the phonon circuit or the feedback resistor in the ionization circuit and thus be seen as a voltage noise at the gate of the FET. Figure 4.27 illustrates this concept. Our strategy to fight microphonics noise will be to reduce $\dot{C}$ and V wherever possible.

We seek to minimize $\dot{C}$ systematically in the tower. We are interested in capacitive coupling to our gate wire that runs between the detector and the FET gate on the FET card. Where our gate wire is actually a trace on a copper-Kapton circuit board, we aim to fix the trace rigidly with respect to any nearby surfaces. In this way any vibrations through the tower will move the gate trace and nearby surfaces together. One area of worry might be the Kapton window on the FET card, but Kapton is very rigid when cold. The FET card, including the window, should be fairly stiff when cold.

Most of the length of our gate wire is in the vacuum coax along the face of the tower and in the side connector. We eliminate any nearby insulators by stretching the wire through copper channels. Stretching the wire keeps the wire under tension as the tower cools and contracts, but tensioning the wire can also have a beneficial effect on microphonics susceptibility. We can think of the microphonic sources as pumping power into our gate wire, making it vibrate. The amplitude of the vibration will act to change the capacitance between the wire and channel by moving the wire closer to one of the walls. The power in a vibrational mode on a string is related to the sound speed along the string, v_s, and amplitude of the vibration, A, by

$$P \propto v_s^3 \cdot A^2. \tag{4.14}$$

For a given amount of power into the system, A will decrease as v_s increases. The sound speed is proportional to $T^{1/2}$, where T is the wire tension. So if we increase the tension in the wire, we have the effect of decreasing the amplitude of the vibration of the wire for a constant power into the wire.

We minimize the potential difference between the gate wire and nearby objects through two main strategies. We maintain the gate at the same potential as its shield, and we minimize the presence of insulators near the gate. Our amplifiers are designed to work with the gate at ground, and have feedback circuits built to keep the gate at ground. We will always have a large potential difference between the gate and drain at the FET, ≈3 V. If vibrations in the FET change the gate-drain capacitance, the FET will be inherently susceptible to microphonics noise. Insulators can become charged as the refrigerator cools through the triboelectric effect.[41]. Charge induced on an insulator will remain, and will create a potential difference between the insulator and nearby objects. This is the reason why we have wherever possible removed insulators near the gate wire, and aim to have as short a gate wire as possible. The desire to remove insulators near the gate wire is often at odds with our infrared shielding desires. For instance, we coat the inside of the FET card gusset with black stycast to absorb the FET infrared. This is undesirable from a microphonics point of view, but perhaps necessary from an IR point of view.

4.5.4.1 Microphonics Tests at 4 K

We have tested the microphonic susceptibility of the FET card and tower in an InfraLabs LHe dewar at 4 K. This dewar has a large enough cold space to mount a tower and FET card without the detector holder subassembly on the 4 K cold plate. Constructing an objective measure of microphonic susceptibility is difficult. In the past, we typically have measured the system's response to a "knock" on the dewar with the knock being a standard bolt on a string swung from a standard height. This makes comparison of susceptibilities hard when the refrigerator is changed or when we test a system in a different refrigerator. Ultimately, we have relied on relative comparisons of noise while running in the refrigerator, with no way of making a quantitative comparison of the various configurations of detector/gate wire/FET mounting we have tried. So, we have tested the tower and FET card for microphonics in a qualitative, but not quantitative, manner.[26]

We will ground the gate wire through a resistor, so we will see the microphonically induced currents as a voltage across the resistor. We will see the Johnson noise of the resistor, and we will be looking for noise peaks above the Johnson noise. We benefit by cooling the resistor, as this decreases the Johnson noise by $T^{1/2}$. We benefit by using a large resistance as the microphonics noise increases directly with R, while the Johnson noise only increases with $R^{1/2}$.

We started by mounting the FET card and gusset directly onto the dewar cold plate. The gusset in these measurements had no IR absorbing black epoxy inside. We connect the FET gate to ground through a 2.2 MΩ metal film resistor. We amplify the noise using a voltage amplifier connected to the FET, and measure the noise spectrum

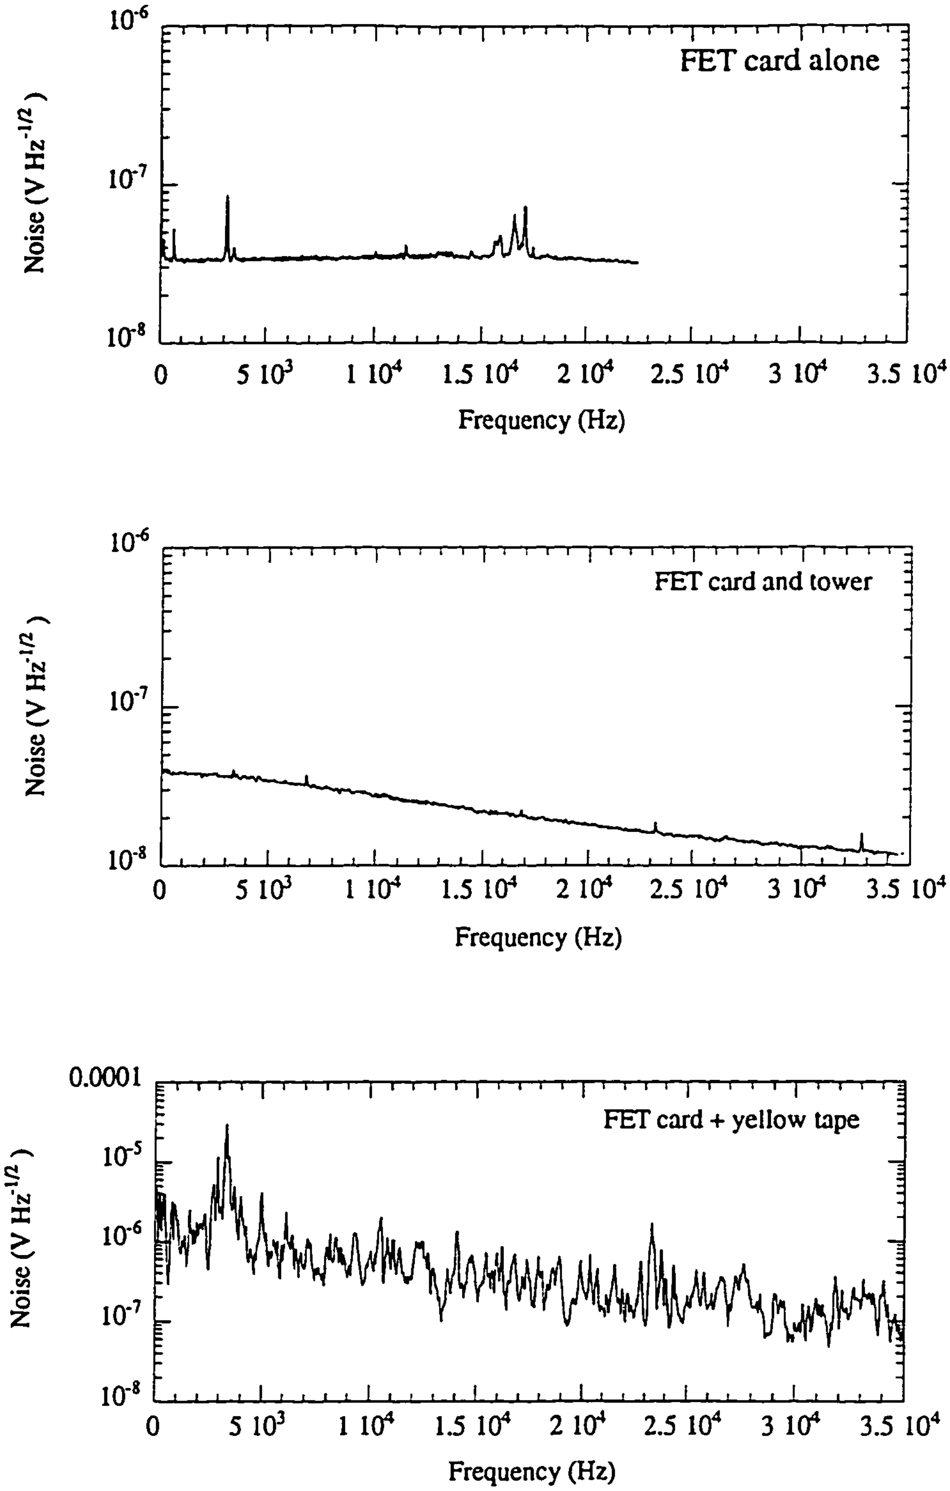

Figure 4.28 Microphonics Measurements with the FET card and Tower.
Shown are the microphonics response of the FET card alone, the FET card and tower, and FET card and yellow-taped gate wire. Note the change in noise scale of the last graph.

with an SR760 Spectrum Analyzer. Our standard "knock" in this case was made by tapping the side of the dewar with a screwdriver in a consistent manner. We then mounted the FET card on the tower and bolted the MC floor of the tower onto the dewar cold plate. The gate wire now includes the wire stretched down one face of the tower. We made the same test with our standard knock and the same 2.2 MΩ resistor. Finally, we measured the response of the FET card with a prototype gate wire used when the Oxford 75 was first used. This gate wire is a manganin wire taped between two layers of yellow tape. These tape covered wires were very microphonically noisy, and were part of our early recognition that microphonics are an important noise source.

Figure 4.28 shows the microphonics response of the three systems mentioned above, the FET card alone, the FET card and tower, and the FET card and yellow-taped gate wire. Note the change in vertical scale for the measurement of the FET card and yellow-taped gate wire. The falling noise in the FET card and tower measurement is due to RC roll off from the increased capacitance of the tower versus the FET card alone. We can immediately see that the FET card seems very good microphonically, and the FET card and tower even better. Hitting the dewar harder than the standard knock did not significantly increase any of the small noise spikes in the second measurement. In previous systems in the dilution refrigerator, we have typically seen a forest of microphonics lines. We are unable to see any real lines in this system of FET card and tower. We know that it is possible to see microphonics from the measurement with the taped gate wire. We should emphasize that these measurements were made before any infrared radiation shielding was added to the tower or FET card, and that the base noise level (the Johnson noise of the resistor) was at least a factor of 15 above what we achieve on a good day in the dilution refrigerator. It is possible that microphonics peaks are hidden under the Johnson noise in these tests. Apart from this consideration, we should expect much of the IR shielding measures, such as absorber in

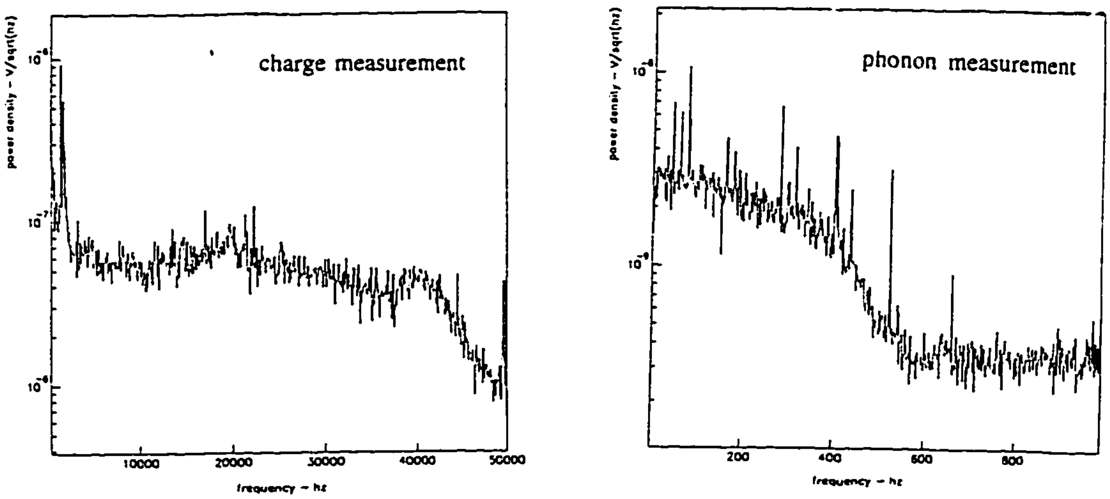

Figure 4.29 Noise Spectra from E5 in the Berkeley Refrigerator
Shown are noise spectra measured on E5 in the tower in the Berkeley refrigerator.

the gusset, black epoxy over traces, and the combs, to increase our susceptibility to microphonics.

4.5.4.2 Microphonics with E5 in the Modified Oxford 75

We have measured noise spectra of both the ionization and phonon measurements on E5 in the modified Oxford 75. We are in the process of redesigning our entire electronics chain, and have found many noise problems associated with purely warm electronics problems (i.e., ground loops, pickup and bad connectors.) Because of this we have no measurements of microphonics in the tower before we added much of our infrared shielding that could conceivably greatly affect the microphonic susceptibility of the tower. We also find that the microphonics environment varies greatly in the dilution refrigerator. The noise from the pot changes with the state of the pot and the bath. This has made it very difficult to gauge the severity of the microphonics problem. Usually, if we find the refrigerator and detector in a microphonically quiet state, we take data and hope the microphonics do not become a problem. Figure 4.29 shows two noise spectra -- one from the charge measurement and one from the phonon measurement. We are in the process of developing a quantitative method of measuring the microphonic susceptibility of the tower system in the refrigerator.

4.6 The Tower in the Ice Box

The tower must ultimately operate in the ice box. CDMS plans to run six detectors per tower, and 3 towers total, in the ice box at the Stanford Underground Facility. CDMS II plans to build another ice box in a deep site and run with six detectors per tower, seven towers total. We will calculate the effect of multiple towers running in the ice box.

4.6.1 Calculation Strategy

The thermal system in the ice box can be idealized in much the same way as the thermal system in the Modified Oxford 75 at Berkeley. Each tower floor is thermally connected to a can lid with a heat strap. The can lids are eventually thermally sunk to a cooling element in the Oxford 400 dilution refrigerator.

The FETs dissipate ≈9 mW each that must be sunk through the LHe floor and lid. We total the FET power for six detectors and assume all of this power flows across the heat strap. This power flows across the LHe can to the LHe bath. We calculate the thermal gradient across the can that this amount of power generates, and use this to find the temperature of the can lid. We will treat the cooling sources for the LHe, Still, and CP cans as "stiff"; i.e., we assume they can sink all the power incident on them without changing temperature significantly. Now we can calculate the thermal gradient across the heat strap, and find the temperature of the LHe floor. Once we know the LHe floor temperature, we can calculate the amount of power flowing down the graphite tube between the LHe and Still floors. The power down the tube will be small compared to the power down the strap for a well designed thermal divider. We now iterate: the actual power down the strap is the difference between the total FET power and the power that flows down the tube. We iterate until we have a consistent solution to the division of power between the strap and the tube. We then use the power down the tube as the incident power on the still floor, and do the process again: calculate the gradient across the can, the gradient across the strap, find the floor temperature. Use the floor temperature to calculate the power down the tube to the next colder floor. Iterate the power division.

This procedure works until we get to the MC floor. The MC is not a stiff cooling source, so we will need to describe the temperature of the MC as a function of the applied power. Once we know the MC temperature, the calculation proceeds as above, except we don't need to iterate because there is no power division.

The extension from one tower to multiple is simple. We simply multiply the power down each strap by the number of towers when we calculate the total power onto each can. We use this total power to calculate can thermal gradients and the rise in mixing chamber temperature.

4.6.2 The Ice Box Thermal Characteristics

The can lids are connected to their cooling source through a complicated system of joints and copper through the cold stem. Rather than analyze in detail the performance of the ice box, we will estimate the conductances of the cans using data

Table 4.4 Conductances and Temperatures in the Ice Box.
This table lists the values of the ice box can conductances and refrigerator temperatures used in the calculation of the tower performance in the ice box.
*We calculate the MC temperature with the applied power.
**This conductance is fit to a T^2 law.

Stage	Temperature (K)	Conductance (W/K)
LHe	4.2	0.813
Still	0.7	2.33×10^{-2}
CP	0.05	2.33×10^{-4}
MC	*	6.48×10^{-3} **

taken from ice box run 9.[20] This data usually only has one or two different applied powers per lid, so we are limited to a linear fit. This run was the first with the entire ice box installed, and an attempt was made to test the ice box by putting power on each can lid and measuring the resulting temperature changes at the lids and refrigerator. We will not attempt a detailed analysis of the ice box performance in this work, but will focus on the performance of the tower, using conservative assumptions where uncertainty in the ice box performance is important.

Table 4.3 lists the conductances and base temperatures used in the calculation. The thermometry on the still can was especially suspect; the still can lid was colder than the still itself. Also, most of the still thermometers gave unrealistically low temperatures (some as low as 435 mK.) So we will adopt the conservative estimate that the ice box still temperatures are at least as warm as the still on the Oxford 75 at Berkeley. The MC stage had better data, and we were able to fit the conductance to a more accurate $P=G(T_h^2-T_c^2)$. We start to run into trouble with multiple towers as the power down the cans can become large enough that our linear fit to the conductance is bad. The CP stage is critical for the functioning of the tower in the ice box, as it is the CP tower floor temperature that determines how much power flows down the CP - MC graphite tube. An improved measurement of the CP can conductance would greatly improve our ability to model the performance of multiple towers in the ice box.

4.6.3 The Tower in the Ice Box

We will assume the towers will all be constructed with the graphite used on the tower described in this work. We have measured a sample of graphite with a much lower thermal conductance, but it proved to be too radioactive to use in the icebox. In addition, we will assume the heat sinking straps have a total resistivity of 3 $\mu\Omega$. The straps used on the tower in the ice box are slightly different from the straps presented in this thesis, but are still constructed of copper lugs welded to a copper braid. The resistance of these braids have been measured at 4 K, and $\approx$3 $\mu\Omega$ is the value for a 4 1/2" braid. The actual braids are likely to be shorter. In the ice box, we have been using two straps on the LHe floor, and one strap each on the Still and CP floors. We assume six 0.3 $\mu\Omega$ bolt joints on the MC floor.

We will not attempt to calculate any effects of infrared radiation from the FETs. We will assume that we will be able to solve the shielding problem in the ice box. This assumption is more a statement of the difficulty of modeling the infrared, than a statement that the infrared will be negligible. In fact, we should think hard about our shielding in the ice box if we expect to run (4 x 6 x 3) = 72 FETs in the ice box given the difficulty we experienced in the Modified Oxford 75 with four FETs.

There will be some additional power load on the LHe tower floor from the striplines attached to the FET cards. We assume the thermal power from the striplines in negligible compared to the power from the FETs.

4.6.4 Results: Multiple Towers in the Ice Box

We present in table 4.4 the results of the calculation for one tower in the refrigerator. This calculation provides a set of numbers we can check against the actual performance of the tower as is it is currently in the ice box. Although the ice box itself is well monitored, there is no measurement of the crucial tower floor temperatures.

Table 4.5 Calculation for One Tower in the Ice Box.
Shown are the tower floor temperatures, the lid temperatures, and the power flowing onto each temperature stage for one tower with 24 operating FETs.
*This is the mixing chamber temperature, not the lid temperature.

Stage	Floor Temperature (K)	Lid Temperature (K)	Power onto Stage (W)
LHe	7.2	5.14	0.215
Still	0.96	0.76	1.45×10^{-3}
CP	0.073	0.066	3.81×10^{-6}
MC	0.013	0.0083*	4.4×10^{-8}

Table 4.6 Calculation for Three Towers in the Ice Box.
Shown are the tower floor temperatures, the lid temperatures, and the power flowing onto each temperature stage for three towers each with 24 operating FETs.
*This is the mixing chamber temperature, not the lid temperature.

Stage	Floor Temperature (K)	Lid Temperature (K)	Power onto Stage (W)
LHe	7.8	5.9	0.642
Still	1.13	0.93	5.3×10^{-3}
CP	0.127	0.120	17×10^{-6}
MC	0.014	0.0092*	2.3×10^{-6}

Table 4.7 Calculation for Seven Towers in the Ice Box.
Shown are the tower floor temperatures, the lid temperatures, and the power flowing onto each temperature stage for seven towers each with 24 operating FETs.
*This is the mixing chamber temperature, not the lid temperature.

Stage	Floor Temperature (K)	Lid Temperature (K)	Power onto Stage (W)
LHe	8.6	7.0	1.50
Still	1.56	1.37	16×10^{-3}
CP	0.355	0.355	70×10^{-6}
MC	0.32	0.032*	20×10^{-6}

This might become important on the LHe floor to ensure the floor temperature does not rise above the superconducting wire $T_C \approx 10$ K. The temperatures of the icebox lids in the calculation compare favorably with actual runs of the tower in the icebox. In particular, the MC temperature is roughly 8 mK with the tower installed. We need to apply ≈ 2 μW heat to the MC in order to bring the MC can lid up to $\approx$20 mK where we operate our detectors.

Table 4.6 presents the results of the calculation for three towers in the ice box. The calculation shows that three towers should run easily in the ice box. The MC floor temperature only rises to 14 mK in the calculation, so we will again need to heat the mixing chamber to operate our detectors. This is probably a bit low, but can depend on the exact operating conditions of the refrigerator. We will probably want to actively control the temperature of the MC can lid to prevent temperature drifts that could affect our detector resolution, so we would like to have a margin in MC temperature to use.

Table 4.7 shows the results of the seven tower calculation. The temperature of the mixing chamber, 32 mK, is unacceptable. This large temperature rise comes from the very hot cold plate floor at ≈360 mK. It is not clear that the conductances assumed are valid as the calculation would sink powers much larger than we have used to test the refrigerator. A more realistic T^2 conductance would sink more power across the cans at higher temperatures than our linear extrapolation, so it is likely that this calculation is unduly pessimistic. We should measure these conductances more accurately, however, to be sure our current ice box design will work with seven towers. The conductance of the CP can is too important to the success of running seven towers in the Icebox to leave it unmeasured. Measuring the conductance now will also allow us to make changes in the design of the Icebox, if necessary, before the construction of the CDMS II Icebox. We also have two simple ways of improving the performance of the tower. First, we can increase the number of heat sink straps at the LHe floor and Still floor to ensure the floor temperatures do not rise much above the lid temperature. Second, we know that graphite with a lower thermal conductance exists. We can look for a radio pure graphite with better thermal conduction than our current tubes.

5. Recent Results of CDMS

CDMS has taken a major step towards an actual dark matter search. We now have our first two cryogenic detectors, Berkeley's E5 and Stanford's FLIP, running in the icebox at the Stanford Underground Facility. Although E5 and FLIP are small in mass, we may be able to lower the existing limits on WIMP dark matter. We can also use these detectors to complete many important tests before installing other larger detectors. E5 will allow us to complete work on our electronic system, including many digitally controlled boards that will be required when running for long periods of time; in addition, we can now measure the actual photon and neutron backgrounds as seen by a cryogenic detector in the icebox and compare these with the predicted backgrounds from the shield design. We should be able as well to make a more accurate measurement of the background rejection capability of the detector, since the neutron background, with muon veto, should be considerably lower at SUF than at Berkeley. We will look at some of the results of measurements with E5 and FLIP at SUF.

5.1 Calibration Runs

The two detectors E5, Berkeley's 60 g Ge detector, and FLIP, Stanford's 100 g Si detector, have now been run simultaneously for more than 200 hours live time. The detectors triggers are logically ORed to allow a trigger in one detector to signal an event in the Icebox. This is important in searching for events which scatter in both detectors. Events which occur in more than one detector can safely be vetoed as background. Events which are identified as neutrons in both detectors provide a source of nuclear recoil events that can be used to monitor the performance of the detectors, especially at low energy where the response of the detectors to nuclear recoils is not well known.

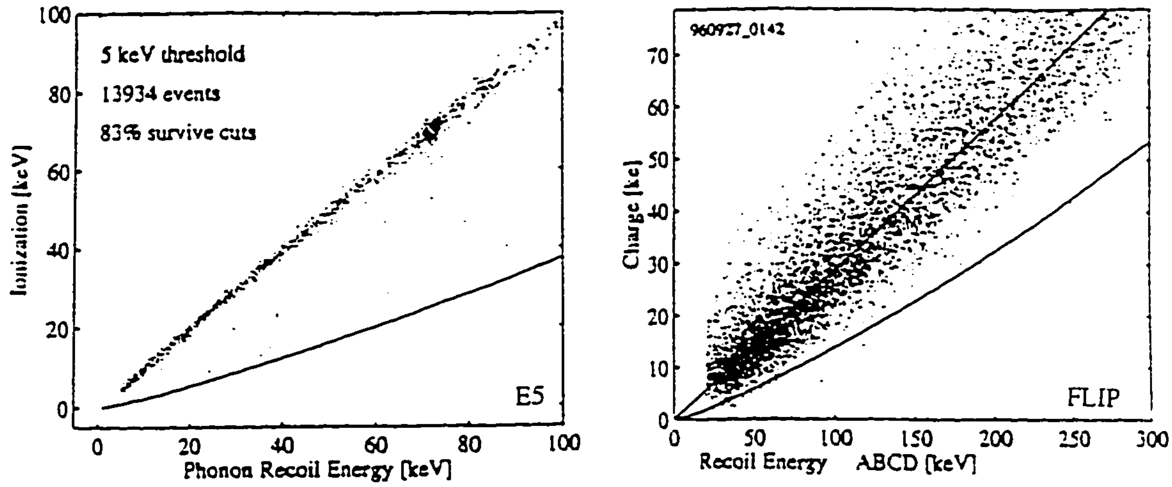

Figure 5.1 Electron Recoil Calibration Run.
Shown are the 2-D plots for E5 and FLIP from the ^{60}Co calibration run. The 2-D plot shows the distribution of measured energy in the charge measurement vs. the recoil energy of the events.

We need to characterize the detectors responses to both photon events and neutron events. By calibrating separately the two classes of events, we can measure the background rejection capability of our detectors as a function of energy. To calibrate the gamma response, we placed a strong ^{60}Co source above the Icebox lid, inside of the shield assembly and muon veto. This source produces two strong lines above 1 MeV (1176 keV and 1333 keV), as well as providing a bright continuum source at lower energy from Compton scattering in the detector. We recorded over 300,000 photon events in the detectors. Figure 5.1 shows the low energy spectra obtained by both E5 and FLIP for the ^{60}Co calibration; figure 5.2 shows a measure of the ionization yield distribution for these events.

To calibrate the response of the detectors to nuclear recoils we need a strong neutron source. Neutrons can endanger the low radioactivity environment of SUF and the Icebox because they create spallation products and activate nuclei, so we must limit our total neutron exposure. To avoid being dominated by our uncertainty in our

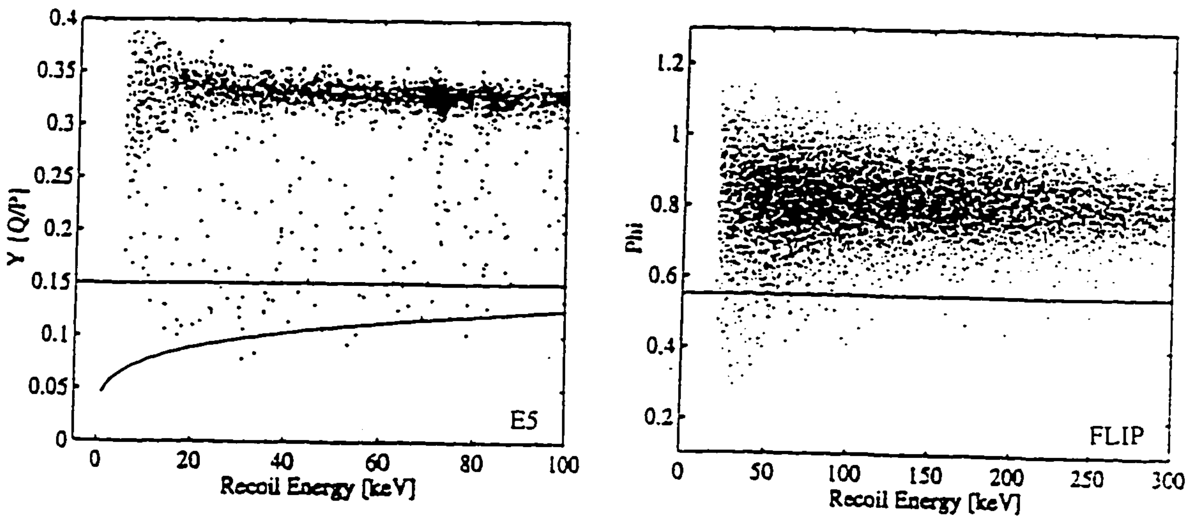

Figure 5.2 Ionization Yield Distribution for the Electron Recoil Calibration Run. Shown is a histogram of the ionization yield for the data from the ^{60}Co calibration run. Y is a measure of charge pairs produced per eV of recoil energy measured in E5. Phi is the arc tangent of the charge (keV) to recoil energy (keV) measured in FLIP.

knowledge of our background rejection, we need to take as many known nuclear recoil events as the total number of signal events we expect to see.[20] Luckily, this is a low rate experiment, so we do not need many neutron events, and thus do not need long neutron exposure times in the icebox. We correspondingly limit neutron calibration runs to ≈1 hour per refrigerator run. We remove the top of the polyethylene shield, and place the ^{252}Cf source on top of the shield lid, outside of the lead and the muon veto. The remaining polyethylene in the Icebox has the effect of softening the neutron spectrum, thus producing many neutrons that will deposit small amounts of energy (< 50 keV) in the detectors, much like the WIMP events we are searching for. Figure 5.3 shows the spectrum obtained for both E5 and FLIP for a neutron calibration run; figure 5.4 is the distribution of ionization yield for the same run.

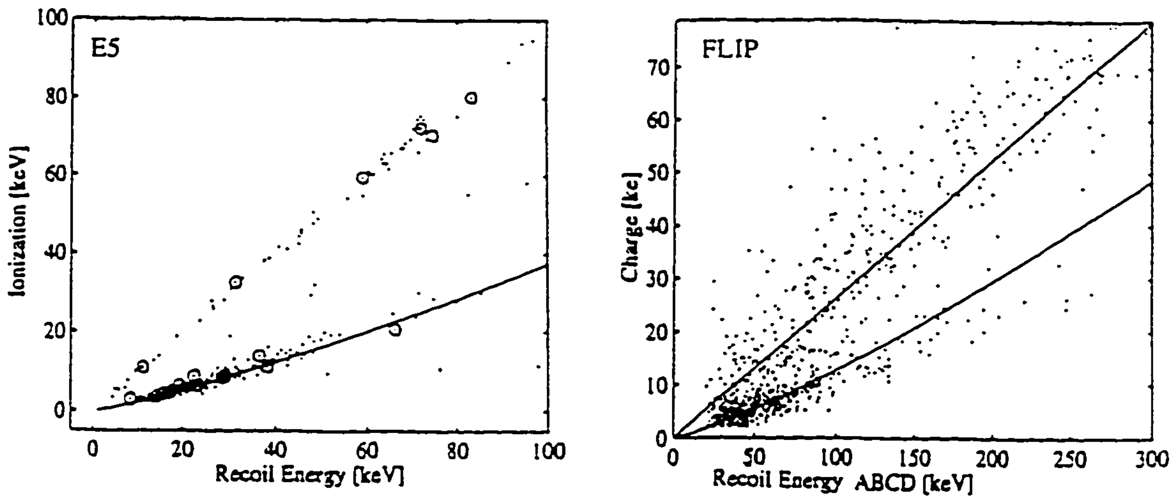

Figure 5.3 Nuclear Recoil Calibration Run.
Shown are the 2-D plots for E5 and FLIP from the ^{252}Cf calibration run. ^{252}Cf is a strong neutron source. The 2-D plot shows the distribution of measured energy in the charge measurement vs. the recoil energy of the events.

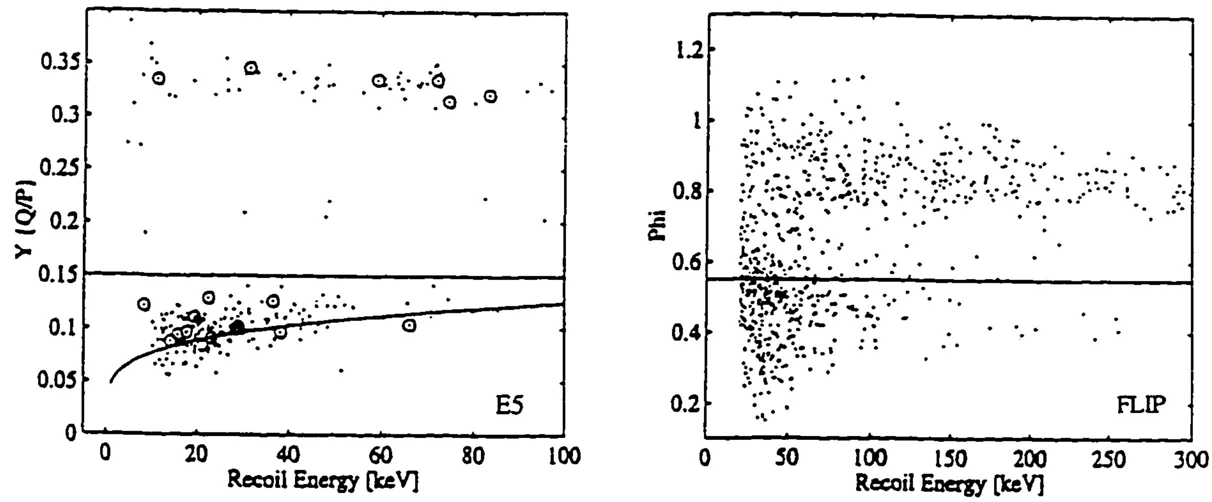

Figure 5.4 Ionization Yield Distribution for the Nuclear Recoil Calibration Run.
Shown is a histogram of the ionization yield for the data from the ^{252}Cf calibration run. Y is a measure of charge pairs produced per eV of recoil energy measured in E5. Phi is the arc tangent of the charge (keV) to recoil energy (keV) measured in FLIP.

5.2 Background Rejection Capability

5.2.1 Nuclear Recoil vs. Electron Recoil Discrimination

Both detectors demonstrate an ability to discriminate between electron recoil and nuclear recoil events. We can see with E5 a background rejection efficiency of 99% (β = 0.01) with a signal acceptance of 95% (α = 0.95) for events above 5 keV. Ideally, we would like to extend this rejection capability to lower energies, but a 5 keV threshold already provides us good sensitivity to WIMP masses > 20 GeV. Reducing the threshold to 1 keV would make us sensitive to WIMP masses > 10 GeV.[20] We are primarily limited by events that lie off the photon axis. We see a small fraction of events that appear to have poor charge collection, even though the resolution in the charge measurement is good. This may be due to charge trapping for events near the surface of the detector.

FLIP has demonstrated a background rejection factor of 99% (β = 0.01) for a signal acceptance of 73% (α = 0.73) for events above 20 keV. FLIP is primarily limited by its energy resolution, which is about ten times worse than E5. This seems to be due to quasiparticle loss in the superconducting films in the present design. FLIP does not seem to have the same problem with incomplete charge collection that E5 has.

5.2.2 The Muon Veto

We have also run the detectors without any external source to examine the muon veto efficiency. We use the scintillator in the shield to mark when muons go through the shield, and correlate this time with the events measured in our detectors. A fast muon through the Icebox will have a number of effects, producing bremsstrahlung γ's which can scatter or interact directly in the detector. More importantly, the muon can

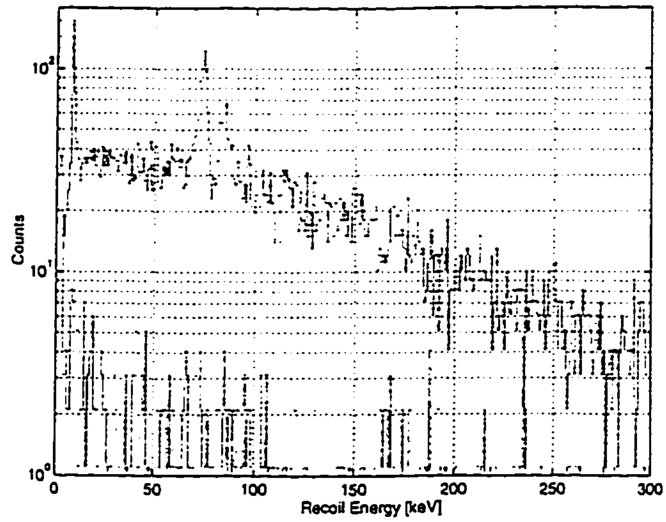

Figure 5.5 Background in the SUF with and without Muon Veto
We have ≈ 200 hours of data taken with E5 and FLIP in the Icebox at the SUF with no external source. The upper spectrum is the total spectrum measured, without any cuts based on the muon veto. The lower spectrum is the resulting background with a veto of any event measured within a 100 μs window after a hit in the shield scintillator.

produce neutrons inside the shield through photofission. These neutrons can then scatter in our detectors, mimicking a WIMP signal. The muon caused neutron production will be correlated with a muon event in the scintillator, but may be delayed by as much as 175 μs.[23] We therefore use the timing from the shield scintillator to cut events that fall within a time window after a scintillator hit. We do all of the actual vetoing of events off-line in our data analysis software. We have about 200 hours of live time with the detectors in a low background state. Figure 5.5 shows the spectrum obtained with E5 with and without the muon veto. The data shows the reduction in background gained by rejecting events measured within a window of 100 μs after a muon hit in the scintillator. The background is reduced by a factor of ≈ 100.

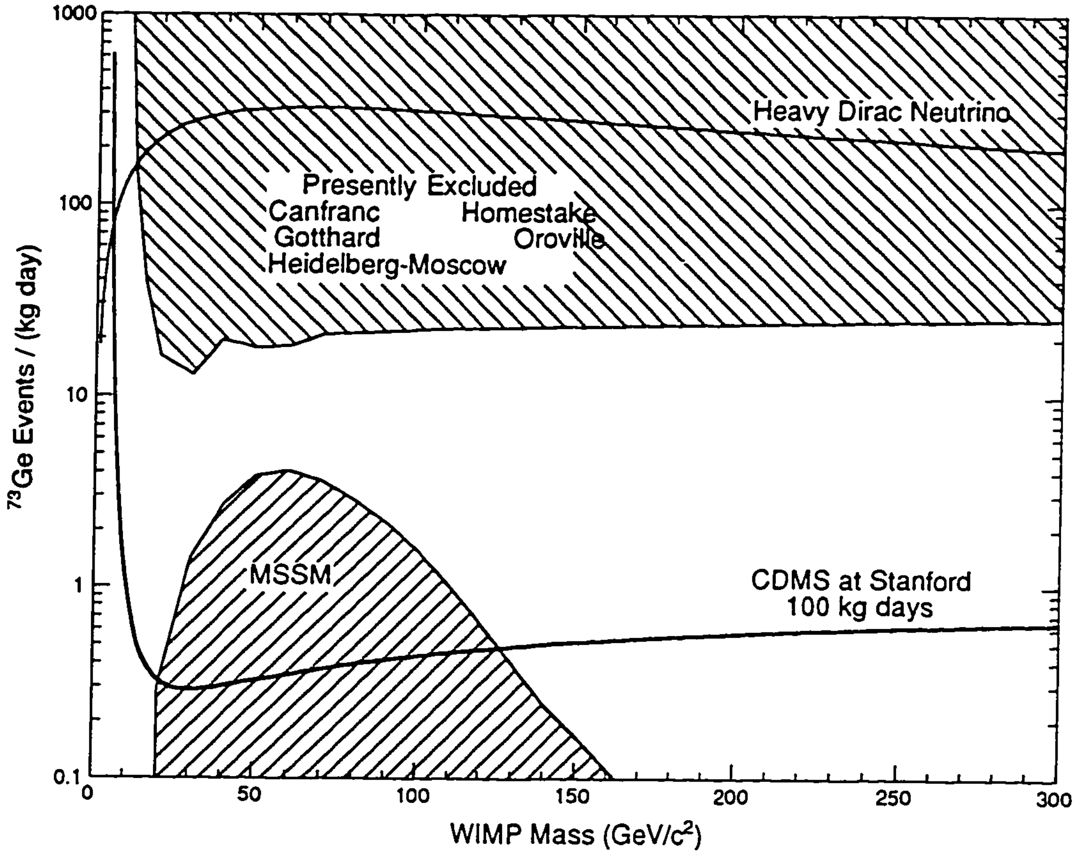

Figure 5.6 Current Limits on WIMP Dark Matter and CDMS.
Shown are the current exclusion limits on WIMP dark matter from Ge detectors, what CDMS should achieve at SUF with 100 kg days of counting, and expected event rates from plausible SUSY models (MSSM).

5.3 Prospects

We have shown that E5 has a background acceptance $\beta = 0.01$ and a signal acceptance $\alpha = 0.95$ at an energy of 5 keV. This gives us a quality factor (defined in section 2.2.1.1) Q(5 keV) $\approx\beta = 0.01$. We have attained 99% muon veto efficiency. We expect the BLIP detectors, which use the same technology as E5 but are $\approx$ 170 g, to attain the same level of resolution and background discrimination as E5. We should be able to achieve a 90% confidence upper limit on a WIMP signal of 0.3 events/kg/day for WIMPs for 100 kg days of counting at SUF. This is sensitive enough to begin exploring realistic models of SUSY particles as WIMP dark matter.[2] Figure 5.6

shows the current exclusion limits[42-46] on WIMP dark matter, and what we should be able to achieve at SUF.

References

[1] E. W. Kolb and M. S. Turner, *The Early Universe* (Addison-Wesley, Redwood City, CA, 1990).

[2] G. Jungman, M. Kamionkowski, and K. Griest, Phys. Rep. **267**, 5–6, 195–373 (1996).

[3] K. G. Begeman, A. H. Broeils, and R. H. Sanders, Monthly Notes of the Royal Astronomical Society **249**, 523 (1991).

[4] M. Fich and S. Tremaine, Ann. Rev. Astr. Astrophy. **29**, 409 (1991).

[5] V. Trimble, Ann. Rev. Astr. Astrophy. **25**, 245 (1987).

[6] S. D. M. White and e. al, Nature **366**, 429 (1993).

[7] A. Dekel, Ann. Rev. Astr. Astrophy. **32**, 371 (1994).

[8] E. Bertschinger and A. Dekel, Astrophysics Journal Letters **336**, L5 (1989).

[9] S. Dodelson, E. I. Gates, and M. S. Turner, Science **274**, 4 October 1996, 69-75 (1996).

[10] D. P. Bennett, *et al.*, in proceedings of 5th Annual Astrophysics Conference. Dark Matter, College Park, MD, USA, 1995, AIP Conference Proceedings #336 (American Institute of Physics, New York), 77-90.

[11] L. Moscoso, in proceedings of 16th International Conference on Neutrino Physics and Astrophysics, Eilat, Israel, 29 May - 3 June 1994, edited by A. Dar, G. Eilam and M. Gronau, Nuclear physics **B38**, (Proceedings Supplements), 387-93 (1995).

[12] A. Udalski, *et al.*, Acta Astronomica **45**, 1, 237-57 (1995).

[13] C. Alcock, *et al.*, Phys. Rev. Lett. **74**, 15, 2867-71 (1995).

[14] E. I. Gates, G. Gyuk, and M. S. Turner, Phys. Rev. Lett. **74**, 19, 3724-7 (1995).

[15] T. Shutt, *et al.*, in proceedings of Sixth International Workshop on Low Temperature Detectors, Beatenberg/Interlakken, Switzerland, 28 Aug.-1 Sept.

1995, Nuclear Instruments & Methods in Physics Research, Section A (Accelerators, Spectrometers, Detectors and Associated Equipment) **370**, 1, 165-7 (1996).

[16] N. Wang, *A Cryogenic Phonon Detector to Search for Dark Matter Particles*, Ph.D. thesis (University of California, Berkeley, 1991).

[17] A. C. Cummings, Ph.D. thesis (University of California, Berkeley, in preparation).

[18] P. N. Luke, J. Appl. Phys. **64**, 6858 (1988).

[19] T. Shutt, *et al.*, Nucl. Instr. Meth. Phys. Res. **A326**, 166-171 (1993).

[20] P. D. Barnes, Jr, *The Cryogenic Dark Matter Search (CDMS)*, Ph.D. thesis (U.C. Berkeley, 1996).

[21] B. Cabrera, *et al.*, in proceedings of the Fifth International Workshop on Low Temperature Detectors, Berkeley, CA, edited by S. Labov and B. A. Young, Journal of Low Temperature Physics **93**, 3/4, 365-375 (1993).

[22] K. D. Irwin, *et al.*, Rev. Sci. Instr. **66**, 11, 5322-6 (1995).

[23] A. Da Silva, *Development of a Low Background Environment for the Cryogenic Dark Matter Search*, Ph.D. thesis (University of British Columbia, 1996).

[24] W. B. Knowlton, *Structural and Phonon Transmission Study of Ge-Au-Ge Eutecticly Bonded Interfaces*, M.S. thesis (U. C. Berkeley, 1995).

[25] W. B. Knowlton, *et al.*, J. Low Temp. Phys. **93**, 3/4, 343-348 (1993).

[26] T. Shutt, *A Dark Matter Detector Based on the Simultaneous Measurement of Phonons and Ionization at 20 mK*, Ph.D. thesis (University of California, Berkeley, 1993).

[27] Y. B. Levinson, in *Nonequilibrium Phonons in Nonmetallic Crystals*, edited by W. Eisenmenger and Kaplyanskii (Elsevier Science Publishers, 1986), p. 91-142.

[28] G. P. Srivastava, *The Physics of Phonons* (IOP Publishing, LTD, Bristol, 1990), p. 178.

[29] É. Aubourg, *et al.*, in proceedings of the Fifth International Workshop on Low Temperature Detectors, Berkeley, CA, edited by S. Labov and B. A. Young, Journal of Low Temperature Physics **93**, 3/4, 289-294 (1993).

[30] T. Shutt, *E5 Design*, 6-23-94, Internal Report (CDMS Collaboration, 1994), p. 2.

[31] W. Stockwell, *et al.*, in proceedings of the Fifth International Workshop on Low Temperature Detectors, Berkeley, CA, edited by S. Labov and B. A. Young, Journal of Low Temperature Physics **93**, 3/4, 755-760 (1993).

[32] C. Kittel and H. Kroemer, *Thermal Physics* (W.H. Freeman, New York, 1980).

[33] C. Kittel, *Introduction to Solid State Physics* (John Wiley and Sons, New York, 1986, revised 1991).

[34] F. Pobell, *Matter and Methods at Low Temperatures* (Springer-Verlag, Berlin, 1992).

[35] O. V. Lounasmaa, *Experimental Principles and Methods Below 1 K* (Academic Press, London, 1971).

[36] D. Osheroff, personal communication, 1994.

[37] S. Margulies, *Measurement of Elastic Limit of NbTi Wires*, Internal Report (CfPA, 1993).

[38] D. R. Smith and E. V. Loewenstein, Journal of Applied Optics **14**, 6, 1335-1341 (1975).

[39] D. B. Tanner, G. G. Ihas, and K. A. Muething, Rev. Sci. Instr. **48**, 6, 610-612 (1977).

[40] S. White, personal communication, 1996.

[41] W. D. Greason, *Electrostatic Damage in Electronics: Devices and Systems* (John Wiley & Sons, New York, 1987).

[42] D. O. Caldwell, *et al.*, Phys. Rev. Lett. **61**, 61, 510-13 (1988).

[43] S. P. Ahlen, *et al.*, Phys. Lett. **B195**, 4, 603-8 (1987).

[44] M. Beck, *et al.*, Phys. Lett. **B336**, 2, 141-6 (1994).

[45] E. Garcia, *et al.*, Phys. Rev. **D51**, 4, 1458-64 (1995).

[46] D. Reusser, *et al.*, Phys. Lett. **B235**, 1, 143-5 (1991).