

Athermal Phonon Sensors in Searches for Light Dark Matter

By

Samuel L. Watkins

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:

Assistant Professor Matt Pyle, Chair
Professor of the Graduate School Bernard Sadoulet
Professor Martin White

Summer 2022

Athermal Phonon Sensors in Searches for Light Dark Matter

Copyright 2022
by
Samuel L. Watkins

Abstract

Athermal Phonon Sensors in Searches for Light Dark Matter

by

Samuel L. Watkins

Doctor of Philosophy in Physics

University of California, Berkeley

Assistant Professor Matt Pyle, Chair

In recent years, theoretical and experimental interest in dark matter (DM) candidates have shifted focus from primarily Weakly-Interacting Massive Particles (WIMPs) to an entire suite of candidates with masses from the zeV-scale to the PeV-scale to 30 solar masses. One particular recent development has been searches for light dark matter (LDM), which is typically defined as candidates with masses in the range of keV to GeV. In searches for LDM, eV-scale and below detector thresholds are needed to detect the small amount of kinetic energy that is imparted to nuclei in a recoil. One such detector technology that can be applied to LDM searches is that of Transition-Edge Sensors (TESs). Operated at cryogenic temperatures, these sensors can achieve the required thresholds, depending on the optimization of the design.

In this thesis, I will motivate the evidence for DM and the various DM candidates beyond the WIMP. I will then detail the basics of TES characterization, expand and apply the concepts to an athermal phonon sensor-based Cryogenic PhotoDetector (CPD), and use this detector to carry out a search for LDM at the surface. The resulting exclusion analysis provides the most stringent limits in DM-nucleon scattering cross section (comparing to contemporary searches) for a cryogenic detector for masses from 93 to 140 MeV, showing the promise of athermal phonon sensors in future LDM searches. Furthermore, unknown excess background signals are observed in this LDM search, for which I rule out various possible sources and motivate stress-related microfractures as an intriguing explanation. Finally, I will shortly discuss the outlook of future searches for LDM for various detection channels beyond nuclear recoils.

To My Family

Contents

1	The Motivation for Dark Matter	1
1.1	Cosmological Evidence for Dark Matter	1
1.1.1	Galaxy Cluster Dynamics	2
1.1.2	Galaxy Rotation Curves	3
1.1.3	Big Bang Nucleosynthesis	4
1.1.4	Colliding Galaxy Clusters	4
1.1.5	Cosmic Microwave Background and Structure Formation	6
1.1.6	Beyond Dark Matter: Dark Energy	8
1.2	Phenomenology of Dark Matter	10
1.2.1	The Classic WIMP	10
1.2.2	Light Dark Matter	13
1.2.3	Other Candidates	16
1.3	Direct Detection of Dark Matter	17
2	Athermal Phonon Sensors	23
2.1	Athermal Calorimeters	23
2.2	Transition-Edge Sensors	25
2.2.1	DC Characteristics	28
2.2.2	TES Response	30
2.3	TES Noise and Energy Sensitivity	42
2.3.1	Intrinsic Noise Sources	43
2.3.2	Noise Modeling	46
2.3.3	TES Baseline Energy Resolution	49
2.3.4	Parameterizing Excess Noise	51
2.4	QETs and Athermal Phonon Collection	51
3	Performance of a Cryogenic PhotoDetector	55
3.1	Motivation of a Large Area Photon Detector	55
3.1.1	Neutrinoless Double Beta Decay Experiments	55
3.1.2	Other Applications	56
3.2	QET-based Large Area Photon Detector	58
3.3	Experimental Setup	58

3.4	Characterization of the CPD	60
3.4.1	DC Characteristics and TES Response	61
3.4.2	Noise Modeling	63
3.4.3	Phonon Pulse Shape	65
3.5	Energy Calibration and Resolution	68
3.6	Discussion	73
4	A Light Dark Matter Search Operated Above Ground	74
4.1	Introduction	74
4.2	Experimental Setup and Data Acquisition	75
4.3	Data Reconstruction and Calibration	76
4.4	Data Selection and Efficiency	80
4.4.1	Baseline Cut	80
4.4.2	Low Frequency Chi-Square Cut	82
4.5	Signal Model and Results	86
4.5.1	Discussion of Flawed Reconstructed Signal Models	87
4.5.2	Reconstructed Signal Model for DM Search	89
4.5.3	Limit Setting and Results	94
4.6	Conclusion	101
5	Investigation of Excess Signals	103
5.1	Excess Signals in CPD	103
5.1.1	Expectation from Random Fluctuations	104
5.1.2	Other Possible Sources	107
5.2	The Case for Stress-Induced Microfractures	112
5.2.1	Glued TES Rectangles	113
5.3	Mitigation of Stress-Induced Events	119
5.3.1	Stress-Free Holding Schemes	119
5.3.2	Thermalization through Gold Pads	119
5.3.3	Vibrational Sensitivity	124
5.4	Summary	125
6	Conclusion and Outlook	126
	Bibliography	129
A	Thermal Conductance Measurements	146
A.1	Thermal Conductance Measurement Principles with TESs	146
A.1.1	Measuring Thermal Conductance between Absorber and Bath	146
A.1.2	Measuring Thermal Conductance between TES and Absorber	149
A.2	Thermal Conductance Measurements with G115	152
A.2.1	Cirlex Clamps	153

A.3	Thermal Conductance Measurements with QP.4	156
A.3.1	GE-7031 Varnish	157
A.3.2	Al Wire Bonds	160
A.3.3	Rubber Cement	164
A.3.4	Thermal Conductance with an IV Sweep	165
A.4	Thermal Conductance Measurements with a SPICE MELANGE Detector . .	167
A.4.1	Au Pad	168
A.5	Summary of Measurements	171
B	Reduction of Vibrational Noise	173
B.1	Derivation of Fundamental Frequencies	175
B.1.1	One Stage	175
B.1.2	Two Stages	181
B.1.3	Spring Mode Intuition	187
B.2	Extension Spring Design	189
B.2.1	Design Principles	189
B.2.2	Geometric Constraints and Expected Load	194
B.2.3	Proposed Extension Springs	195
B.3	Thermal Considerations	201
B.3.1	Thermal Conductance Between Stages	201
B.3.2	Stiffness of Copper Braids	203
B.3.3	Geometric Constraints	204
B.3.4	Proposed Thermal Links	206
B.3.5	Spring Cross-Checks	211
B.4	Intermission	215
B.5	Safety Structure	217
B.6	Final Remarks	221
C	Optimal Filtering	223
C.1	Optimal Filter Formalism	223
C.1.1	Extracting Pulse Heights	224
C.1.2	Extracting Pulse Heights and Pulse Times	225
C.1.3	An Optimal Filter Example	226
C.2	Two-Pulse Optimal Filters	227
C.2.1	Iterative Pileup Optimal Filter	228
C.2.2	Simultaneous Pileup Optimal Filter	229
C.3	Arbitrary Number of Pulses	231
C.4	Optimal Filter Code in QETPY	234

Acknowledgments

I'm fairly certain it is tradition to write your acknowledgments in one fell swoop the day before submitting a thesis, so for anybody that I miss in the below: thank you so much for your love and support!

I want to thank a few standout teachers and mentors that helped lead me down this path of life from before I started graduate school. Thank you Mr. Goldberg for being an incredibly passionate high school teacher and providing the first jumping off point to pursue physics. I would also like to thank Prof. Eric Hudson from my time as an undergrad at UCLA. Your courses, especially the quantum optics lab, showed me how amazing physics could be, as well as how awesome a physicist could be! Although I ended up doing dark matter detection, AMO will always have a soft spot in my heart. After my undergraduate education, I worked as an algorithm engineer at PNI Sensor Corporation in Santa Rosa, where my boss Andrew Taylor saw great potential in me and supported me instantly when I told him that I wanted to pursue a Physics PhD. Thank you for seeing something in me and giving me an extremely useful experience in industry.

I am so thankful that I started graduate school at UC Berkeley as a GSI for Physics 7A with an awesome group of people: Patrick, Isaac, Neha, Energy, Reed, Liz, the other Sam, and Best. I'm not sure how I would have gotten through the first year without you all—we absolutely must attend a 7A's game again one of these days! After that crazy first year, I joined the Pyle group in its infancy, joining the motley crew of Matt, Bernard, Bruno, Caleb, and Suhas. At that time, we didn't have much of a lab, and we frequently worked with our SuperCDMS collaborators at SLAC. There I was fortunate to work with Paul, Tsuguo, and Noah. Paul and Tsuguo, thank you eternally for running the lab that provided the majority of the data for this thesis! Noah, you filled the role of our group's senior grad student, and have always provided advice and mentorship at a moment's notice (even when you're in the midst of building your own group at SLAC, you still always found time for me). Here's looking forward to many more conferences where we wander the streets of wherever we are!

For the rest of SuperCDMS, I am immensely grateful to have been able to be a part of the collaboration. For the CPD DM Search in this thesis, I want to thank Wolfgang, Ray, Steve, Scott, and all the others that helped ensure this was an excellent result. Thank you Emanuele for helping shepherd it through the various steps to publication (as well as hosting Caleb and me in Italy for the hike of a lifetime in the Dolomites, though let's bring more water next time...). Thank you Belina for your incredible support from across the globe, working with you has been some of the best collaborating of my graduate school career, and your support in my postdoc search will always immensely appreciated.

Back to the Pyle group, I, as is standard in almost every SuperCDMS thesis, must thank Bruno! I'm not sure how you manage working what seems to be the job of ten (probably more) people, but you do it incredibly well. Working with you is always a great experience, and I always appreciate your patience when I ask a question that I probably should have known the answer to years before... Bernard, although you have mostly lived on my Zoom screen for the PhD for both Hawaii and COVID-related reasons, the time spent with you

has been amazing. Your depth of knowledge is awe-inspiring, as well as the perfectly timed quips here and there in meetings. Caleb (aka Dr. Best Man), you are definitely one of the main reasons I made it through graduate school. The friendship we’ve built over the years making it through every up and down has kept me sane. From the death metal concerts, after which neither of us can hold up our heads without pain, to the arm length burritos, CCW, and all the silly inside jokes we have, it’s awesome that we’re both going off to the same place! Long live Samleb.

Matt, I’m going to miss all of the random trips with the iPad to Yali’s, the business school cafe, or Strada! Your willingness to always be available to talk about literally any aspect of what I was doing (or trying to do) is something that is quite rare in a PI. Of the many lessons you’ve taught me, I’ll always remember the power of scaling laws and the power of the $\Delta\chi^2$. It’s amazing what these simple tools can teach us and help guide research. I am also incredibly thankful for the work-life balance of our group—it’s something that again is not a given. I’m looking forward to catching up every time we run into one another at conferences, or when I’m back in the bay area. Being part of the beginning of SPICE/HeRALD has been a fantastic experience, and I’m excited to see the awesome physics results that will come.

Liz, you have been the main source of support and encouragement in my life since we met in our first year here. You have always listened to me and helped whenever you could. I can’t imagine my life without you, and I’m so glad we both taught 7A together! This last year planning our wedding, finding academic jobs at the same place, and finishing off our theses has been a whirlwind and a testament to the strength of our relationship. I can’t wait for our next steps in NM!

I want to thank Jonah for being another amazing friend in graduate school. I’ll miss being able to just drop by your office and get some “tuny sammys” for lunch or some after-work beers. And of course, you were an awesome minister at Liz’s and my wedding—I’ll ask you for spiritual advice for years to come! To the Neaton group (and the Neaton group adjacent) in general, thank you for making me feel welcome and a part of the social aspect of the group when I started coming by to hang with Liz. And Ben, thanks for getting Caleb and I off to our legendary gym sessions, where we probably spent most of the time just chatting... But hey, those death metal shows probably were enough of a workout anyways.

Pursuing a PhD during COVID was quite the experience. I want to thank Sami, Kayee, and Andrew for being my online buddies and playing video games into the wee hours of the night (don’t worry Matt, it was only on weekends!). Having this outlet during the pandemic allowed me to let go and have fun no matter how much stress was going on in my life. Reconnecting was awesome, and I’m glad we’re keeping the gaming going.

I’ll finish these acknowledgments by thanking my family. Mom, Dad, Troy, Jen, you have all provided unconditional love and support throughout my academic career. I always knew I could count on one of you to listen and give advice when I needed it. And you’ve all supported my journey to the end of this PhD (and the beginning of what comes next). To my brother Robert, it was awesome to live close together, and I’m lucky to have an awesome bro like you. And to my Kissick, Richardson, and Watkins families, thank you! And no, I haven’t found dark matter... *yet...*

Chapter 1

The Motivation for Dark Matter

There are few unanswered questions about the universe that conjure up as much imagination as what makes up our universe beyond the baryonic matter: what is dark matter and what is dark energy? Through cosmological observations it has become clear that these two mysterious components make up the majority of the universe, as shown in Fig. 1.1. In this chapter, I will focus on the motivation, phenomenology, and direct detection methods for dark matter, keeping ourselves working within the matter of the universe and leaving a Ph.D. on dark energy for another life.

1.1 Cosmological Evidence for Dark Matter

From various independent measurements of our universe, there is strong cosmological evidence for the existence of dark matter. They span from galaxy rotation curves to big bang nucleosynthesis to precision measurements of the cosmic microwave background.




Figure 1.1: The percent of energy density of our universe's content for dark energy, dark matter, and baryonic matter. These data are reported in Ref. [1].

1.1.1 Galaxy Cluster Dynamics

Evidence for dark matter is apparent at the large scale of galaxy clusters. When observing galaxy clusters, the virial theorem can be used to estimate the mass of a cluster. The general form of the virial theorem is

$$\langle T \rangle = -\frac{1}{2} \sum_{k=1}^N \langle \mathbf{F}_k \cdot \mathbf{r}_k \rangle, \quad (1.1)$$

where T is the total kinetic energy N particles, $\mathbf{F}_k$ is the force on the k th particle, and $\mathbf{r}_k$ is the position of the k th particle. The brackets denote averaging of the enclosed quantities over time. For a potential that depends only on the distance between particles, i.e. $V(r) \propto r^\alpha$, this simplifies to

$$\langle T \rangle = \frac{\alpha}{2} \langle V \rangle. \quad (1.2)$$

For a galaxy cluster approximated as a sphere held together by gravity, we have that $\alpha = -1$, and this simplifies to

$$\langle T \rangle = -\frac{1}{2} \langle V \rangle, \quad (1.3)$$

where $T = \frac{1}{2} M v^2$, and $V = -\frac{3}{5} \frac{GM^2}{R}$ is the gravitational binding energy of the cluster. Plugging in these values and rearranging to solve for the mass of the cluster, we have that

$$M = \frac{5 \langle v^2 \rangle}{3G \langle 1/R \rangle}, \quad (1.4)$$

as shown and applied to galaxy clusters by Zwicky [2, 3]. When Zwicky applied the virial theorem to the Coma cluster, he was able to estimate its mass-to-light ratio, i.e. the ratio of the mass of a galaxy cluster to its luminosity. For a system completely composed of visible matter, the mass-to-light ratio in units of solar mass divided by solar luminosity would be expected to be about 1. Instead, Zwicky noted that the Coma cluster had a mass-to-light ratio of a few hundred, leading to the positing of dark (nonluminous) matter making up a significant majority of the cluster.

The amount of missing luminous mass in galaxy clusters can be more accurately determined through X-ray emission and gravitational lensing techniques to study the distributions of baryonic and dark matter. At these scales, the majority of baryonic matter remains captured within the gravitational potential well of the cluster, where it forms a gas of temperatures greater than 10^6 K and emits X-rays. From the distribution of this hot gas, and under the assumption that it is in hydrostatic equilibrium with the gravitational potential, the total mass of a cluster can be estimated [4]. For gravitational lensing, general relativity provides that the gravitational potential of the cluster's matter (or any massive object in general) will act as complex lens that distorts the images of more distant galaxies and other celestial objects, where the amount of lensing is directly related to the total mass of the cluster.




Figure 1.2: (Figure from Ref. [6]) The measured galaxy rotation curves for a variety of galaxies, showing the clear leveling off indicated that there must be more mass than is visibly apparent.

Measurements of galaxy cluster masses using each of these techniques have been shown to be statistically consistent [5], confirming that there appears to be significantly more mass than what can be accounted for by the visible matter.

1.1.2 Galaxy Rotation Curves

When one is asked about the existence of dark matter, they will likely start with the common explanation based on observations of the rotational dynamics of galaxies. Using our understanding of Doppler shifts of different spectral features, astronomers can calculate the galactic rotation speed as a function of radius [6, 7]. These spectral features include the visible H spectral lines, 21 cm H I line, and rotational transitions in CO.

If one were to assume that visible matter represents the entirety of the mass distributions of these galaxies, then the expectation would be that the orbital velocity of galactic matter would decrease with radius r as $\sim r^{-2}$ outside the visible disk (i.e. where the baryonic matter distribution no longer appreciably increases). However, as shown in Fig. 1.2, the observed velocity rotation curves have been observed to be constant at distances far from the center of the galaxy.

This constant velocity is consistent with a spherically symmetric distribution of mass whose enclosed mass increases linearly with radius. The suggestion becomes that there must be more mass than that which is visible, which can be explained by the existence of “dark matter halos” that extend much farther than the visible disk of matter. For non-

spiral galaxies, the rotation curves may not provide insight into the amount of missing mass (e.g. an elliptical galaxy does not appreciably rotate). However, from velocity dispersion measurements and applications of the virial theorem, dark matter has been shown to be needed to match observations of dwarf and elliptical galaxies [8, 9].

1.1.3 Big Bang Nucleosynthesis

Through dynamics on the scales of galaxies and galaxy clusters, we have seen that dark matter makes up the majority of the mass of these celestial objects. We can turn to observations of the early universe from standard cosmology to specify how much of the matter in our universe is dark. During the cooling of the universe after the Big Bang, light elements began to form as part of Big Bang Nucleosynthesis (BBN), most notably D (deuterium), ^3He , ^4He , and Li. Of these elements, D is a fragile nucleus that is destroyed within stars and no longer created in the modern universe, as opposed to the latter three which are created by stars. To understand the primordial abundances of ^3He , ^4He , and Li is therefore quite complex, and we are fortunate to be able to precisely measure the abundance of primordial D, which pins down the baryon density extremely accurately. These measurements are based on light from distant quasars being absorbed by intervening neutral hydrogen systems at redshifts of 3–4 (at which primordial abundances have yet to be altered). The absorption spectra provide key features for extracting the ratio of D to H, which are then related to the baryon–photon ratio η through standard BBN calculations [11]. These measurements give precision results on the baryon density, as shown in Fig. 1.3.

In this figure, the bands for each light element corresponds to predictions given various ratios of baryon–photon densities η , showing strong dependence. The vertical cyan band in this figure is the baryon density specified by primordial D measurements, corresponding to $\Omega_B h^2 = 0.019 \pm 0.0024$, or $\Omega_B \approx 0.04$ when dividing the square of the dimensionless Hubble constant h . When comparing the total matter density of $\Omega_M \approx 0.3$, it becomes clear that the majority of the matter content of our universe is nonbaryonic.

1.1.4 Colliding Galaxy Clusters

In the previous sections, it has been shown that the missing matter has quite different properties than baryonic matter: it is nonluminous, does not form compact structures (i.e. does not lose energy easily), does not emit X-rays, and is not made of baryons. From the baryonic density from Big Bang Nucleosynthesis and comparing to total mass estimates, we know that the amount of baryons in the universe is not enough to explain the mass of galaxy clusters. For a dramatic display of the different properties of baryonic matter and dark matter, one of the most well-known examples is that of the Bullet Cluster (1E 0657-56), which consists of two colliding clusters of galaxies. In Ref. [12], the distributions of hot gas and total mass have been reconstructed using the techniques of X-ray emission and gravitational lensing (see Fig. 1.4).




Figure 1.3: (Figure from Ref. [10]) The dependence on the primordial number of various nuclei on the amount of baryonic matter in the universe. The yellow regions show the observed abundances, with the pink vertical band being the constraint set by BBN. The constraints on the baryonic matter indicate that it cannot make up the total matter content of the universe.




Figure 1.4: (Figure from Ref. [12]) (Left) Image of the bullet cluster, where the green contours indicate measurements via gravitational lensing. (Right) Comparison of the mass distribution measurements via hot gas (baryonic matter) and via gravitational lensing (total matter). The decoupling of the two indicate the existence of nonbaryonic matter.

The results provide a clear decoupling between the baryonic matter and the dark matter in this cluster: the dark matter halos of each cluster have passed through one another with very little interactions (the green contours) as compared to the baryonic matter distributions that lag behind due to interactions between the hot gas of each cluster. The self-interaction cross section of dark matter can be quantitatively estimated from this cluster, giving that the cross section per mass is less than $1.25 \text{ cm}^2/\text{g}$ [13]. The majority of each cluster's mass is not the baryonic matter, but some other form of matter with very different properties (dark matter). This phenomenon has also been observed in the merging cluster MACS J0025.4-122 [14], providing more strong evidence of the differences between baryonic and dark matter. These colliding clusters furthermore give strong evidence against modified Newtonian dynamics theories (i.e. that we do not understand gravity at large scale) due to the large spatial separation between the centers of the total mass and the baryonic mass (8σ in the case of the Bullet Cluster).

1.1.5 Cosmic Microwave Background and Structure Formation

Through observations of the cosmic microwave background (CMB), we can find more evidence of the existence of dark matter. The CMB is the primordial blackbody radiation of the universe, which last scattered off electrons when the universe was about 400,000 years old. After this point, these photons have traveled freely through space and are observable today with a thermal blackbody spectrum at a temperature of 2.7 K. For decades, it was observed that the CMB was completely isotropic, until anisotropies were discovered in its temperature and polarization [15, 16]. These anisotropies follow a characteristic pattern that is predicted by inflation, as shown in Fig. 1.5




Figure 1.5: (Figure from Ref. [10]) Measurements of the temperature anisotropies in the CMB expanded in multipole moment. The location and magnitude of these peaks place strong limits on the energy contents of the universe.

These temperature fluctuations of the CMB are highly dependent on the baryon density, as they come from inhomogeneities in the photon-baryon fluid before the photons decoupled to become the CMB. Thus, the density of the baryonic matter will have a substantial effect on these fluctuations, and the results give a sensitive probe on the cosmological parameters in the theory of inflation. The results shown in the figure correspond to a baryonic matter density of $\Omega_B \approx 0.04$, and a majority of the matter in the universe must be nonbaryonic. Note that this value is consistent with the value from Big Bang Nucleosynthesis with a substantially different physics derivation.

These small anisotropies in the CMB have not had enough time to coalesce into the structures seen in today's universe. That is, if there were only baryons, these fluctuations are too small to have taken the early homogeneous universe and resulted in the observed structure of the universe today. Thus, there must be some matter that does not couple to photons, i.e. did not couple to the CMB. The fluctuations in this matter's density could then grow to make the structure we see without affecting the scale of the anisotropies of the CMB. Once radiation and baryonic matter decouple, then the baryonic matter will fall into the gravitational wells of this nonbaryonic matter, giving the structure of the universe. One

important characteristic of this nonbaryonic matter is that it must be nonrelativistic (“cold”), giving us the theory of cold dark matter. If dark matter had relativistic velocities, then this would result in structure formation that does not fit our cosmological observations, as the formation paradigm would be for large-scale structure to fragment into smaller structures. Instead, our observations are consistent of a hierarchical formation of large-scale structure originating from the coalescing of smaller structures [17, 18].

1.1.6 Beyond Dark Matter: Dark Energy

Although we have spent a few pages discussing the evidence for dark matter, the various types of matter in our universe cannot explain its full content (recall Fig. 1.1). Perhaps the most convincing first argument for this is that the universe is flat from CMB anisotropy measurements. If the universe were open, then the paths of photons from the surface of last scattering (i.e. the CMB photons) starting out parallel would slowly diverge, and the physical scale with maximal anisotropy (the first peak of the CMB anisotropies) would be shifted to smaller angular scales (higher l in the spherical harmonics, as the angular scale is $\theta \approx \pi/l$). A closed universe would similarly shift this first peak to larger angular scales. Precisely measuring the location of this first peak then gives a precise value of the curvature of the universe—this peak has been measured to occur at around $l = 200$ [20–22]. In models where the universe is nearly flat (where the curvature is described by the total density of the universe Ω_0 , and a flat universe has $\Omega_0 = 1$), the location of this peak is $l \approx 200/\Omega_0^{1/2}$. Thus, these measurements have concluded that $\Omega_0 = 1$ to high precision, giving convincing evidence that the universe is flat.

To demonstrate the importance of this, we turn to the Friedmann equations, which describe the expansion of space under the assumption of a homogeneous and isotropic universe. The first Friedmann equation can be written in terms of present day values as

$$\frac{H^2}{H_0^2} = \Omega_{\text{rad}} a^{-4} + \Omega_M a^{-3} + \Omega_k a^{-2} + \Omega_\Lambda a^{-3(1+w)}, \quad (1.5)$$

where a is the time-dependent scale factor of the universe, $H \equiv \dot{a}/a$ is the Hubble parameter, H_0 is the Hubble parameter today, Ω_{rad} is the radiation density today (about 10^{-4} and can be neglected), Ω_M is the total matter density today, $\Omega_k = 1 - \Omega_0$ is the “spatial curvature density” today, Ω_Λ is the cosmological constant today, and w is the equation of state parameter of dark energy (the ratio of the pressure to the energy density of a perfect fluid). For a universe with flat curvature, this would mean that $\Omega_k = 0$ and (neglecting Ω_{rad}) $\Omega_M + \Omega_\Lambda = 1$. Because $\Omega_M \approx 0.3$, the value of Ω_Λ must be nonzero—requiring the existence of some other phenomenon. Though this is not direct evidence of dark energy, it becomes quite difficult to reconcile the observed curvature of space without including a nonzero Ω_Λ .

The first direct evidence of dark energy was published in 1998 by Perlmutter *et al.* [19] and Riess *et al.* [23], where each analysis used high-redshift Type Ia supernovae observations to measure the expansion rate of the universe and show that this expansion is accelerating.




Figure 1.6: (Figure from Ref. [19]) Constraints set on Ω_Λ by The Supernova Cosmology Project in Ref. [19]. While the constraints on Ω_Λ have become more precise since this paper, these initial results provided direct evidence of a nonzero Ω_Λ and the existence of dark energy.

This acceleration is well explained by the existence of dark energy with an equation of state parameter of $w = -1$. In Fig. 1.6, we reproduce the results from Perlmutter *et al.*, showing that these measurements directly gave $\Omega_\Lambda > 0$.

Further evidence for dark energy comes from measurements of baryon acoustic oscillations [24] and structure formation measurements combined with the CMB anisotropies [25]. These observations have led to the Λ CDM model being the predominant description of the universe. Although the details of the Λ CDM model can be argued (e.g. perhaps $w \neq -1$, but is time varying, as proposed by Caldwell *et al.* as quintessence [26]), the overwhelming evidence for dark matter and dark energy means that we should figure out what they are. As mentioned in the first paragraph in this chapter, this thesis will focus on the dark matter side of these (as of 2022) unanswered questions.

1.2 Phenomenology of Dark Matter

With the strong evidence of dark matter from our cosmological observations, we have high confidence that dark matter must be nonbaryonic, less self-interacting than baryons, stable, and cold. At its core, these are remarks that dark matter interacts gravitationally, and any theory that fits these four criteria could be a plausible explanation of the nature of dark matter. In fact, there are theories of dark matter that cover the mass range of $\mathcal{O}(\text{zeV}/c^2)$ to $\mathcal{O}(\text{PeV}/c^2)$ up to 30 solar masses. Of these theories, a handful have received the most attention up until roughly the 2010s, with the most given to weakly interacting massive particles (WIMPs).

1.2.1 The Classic WIMP

Historically, in theories of dark matter, the WIMP seemed to be the most promising, due in part to the argument's relative simplicity. A WIMP χ is generally proposed as some massive nonbaryonic particle that was in thermal equilibrium with the universe at early times. At these times, when temperatures were much higher than the mass of the WIMP ($T \gg m_\chi$), the colliding particles of the thermal plasma had enough energy to efficiently create or annihilate WIMPs. As the universe expanded, the temperature of this plasma decreased until, at some point, the WIMP annihilation rate dropped below the expansion rate of the universe H , a process frequently described as “thermal freeze-out.” After freeze-out, the number of WIMPs in the universe remained approximately constant, creating the relic cosmological abundance that we see today.

To show this quantitatively, the derivation (e.g. see Ref. [27, 28]) starts with the Boltzmann equation for the time evolution of the WIMP number density $n_\chi(t)$

$$\frac{dn_\chi}{dt} + 3Hn_\chi = -\langle\sigma_A v\rangle \left[(n_\chi)^2 - (n_\chi^{\text{eq}})^2 \right], \quad (1.6)$$

where H is the Hubble expansion rate. The left hand side of Eq. (1.6) accounts for the expansion of the universe, while the right hand side accounts for the annihilation and creation




Figure 1.7: (Figure from Ref. [27]) The mechanism of the thermal freeze-out for WIMPs for achieving the DM relic density.

of WIMPs. If we approximate that $\langle\sigma_A v\rangle$ is energy-independent, then we can combine our understanding of the Hubble expansion rate during the radiation-dominated universe ($H(T) \propto T^2$) and the freeze-out condition ($\Gamma_A = n_\chi \langle\sigma_A v\rangle$) to approximately solve this equation. After using the present values of today's entropy density and critical density, one finds that the present DM mass density is

$$\Omega_\chi h^2 \simeq \frac{3 \times 10^{-27} \text{ cm}^3/\text{s}}{\langle\sigma_A v\rangle}. \quad (1.7)$$

The result in Eq. (1.7) is largely independent of the mass of the WIMP (there are small higher order corrections) and inversely proportional to its annihilation cross section. If we were to solve the Boltzmann equation numerically, then we similarly find that, as the annihilation cross section increases, the relic comoving number density decreases, as shown in Fig. 1.7.

To motivate the WIMP being a natural candidate, one can quickly estimate the expected annihilation cross section for some new particle with weak-scale interactions as

$$\begin{aligned} \langle\sigma_A v\rangle &\sim \alpha^2 (100 \text{ GeV})^{-2} \\ &\sim 10^{-25} \text{ cm}^3/\text{s}, \end{aligned} \quad (1.8)$$

where $\alpha \sim 10^{-2}$. This value is coincidentally very close to that which is needed to achieve the relic DM density, which provided a strong simple argument for DM to be a WIMP.




Figure 1.8: (Figure from Ref. [10]) The landscape of constraints on WIMPs as of 2021.

Furthermore, this is the same scale at which new physics was expected from supersymmetry, which was proposed in part as a solution to the hierarchy problem (that radiative corrections would drive the Higgs mass much higher unless there are new particles at the weak-scale) [29].

With the above argument in mind, experimentalists have searching for WIMP dark matter for decades to no avail. In Fig. 1.8, various DM searches over the expected WIMP mass scale have increasingly ruled out DM-nucleon cross sections down to $\sim 10^{-46} \text{ cm}^2$. In WIMP mass, Lee and Weinberg showed that WIMP DM should have a mass greater than $\sim 2 \text{ GeV}$ (the Lee-Weinberg bound), otherwise the DM would have frozen out much earlier in the universe's history, leading to an overabundance of DM and the overclosing of the universe [30]. As these WIMP searches have breached this bound, a huge amount of the WIMP parameter space has been excluded, and expectations from supersymmetric theories have been dashed. In Fig. 1.8, there remain theoretical supersymmetric models that predict unexplored DM-nucleon cross sections (many of which will be probed in coming experiments). However, the decades of null results on discovering supersymmetric particles at the Large Hadron Collider and WIMP dark matter particles in the expected mass range have brought recent interest to alternative dark matter theories [31–33].




Figure 1.9: (Figure from Ref. [31]) A diagram of the masses that are spanned by all dark sector candidates. There are many orders of magnitude in mass that are populated by non-WIMP candidates.

1.2.2 Light Dark Matter

Light dark matter (LDM) is based on the idea that these DM particles are instead neutral under standard model forces, but charged under undiscovered forces (frequently referred to as hidden or dark sectors). These hidden-sector interactions provide a possible origin for DM, but without the mass restrictions provided by the WIMP hypothesis (i.e. these non-WIMP theories circumvent the Lee-Weinberg bound).

In Fig. 1.9, these dark sector candidates cover a much wider range of DM masses, where LDM is usually thought of as the mass range between $\mathcal{O}(\text{keV})$ and $\mathcal{O}(\text{GeV})$. We note that the lower bound of $\mathcal{O}(\text{keV})$ also corresponds to the lower bound of warm dark matter mass [34, 35]. Although it is entirely possible that there are no interactions between this hidden sector and the Standard Model, general symmetry arguments allow for interactions between some generic hidden sector and the Standard model, where these “portal” interactions are generated by radiative corrections. These couplings also can explain the mechanism by which

the universe achieved its dark matter abundance via thermal contact.

When discussing hidden-sector DM, attention is generally given to simplified models of a new force that is mediated by a vector or scalar boson [31, 36]

$$\mathcal{L}_V \supset V_\mu \bar{f} (g_f^V \gamma^\mu + a_f^V \gamma^\mu \gamma^5) f \quad (1.9)$$

$$\mathcal{L}_S \supset \bar{f} (g_f^S + \gamma_5 a_f^S) f \phi, \quad (1.10)$$

where V_μ is a vector mediator, and ϕ is a scalar mediator. In these models, the structure of the g_f and a_f couplings will depend on how the mediator couples to ordinary matter. Two well-studied cases are the vector portal and the Higgs portal, for which there are unique renormalizable interactions of a Standard Model neutral boson that have all Standard Model symmetries [37–39]:

$$\mathcal{L} \supset -\frac{\epsilon}{2 \cos \theta_W} B_{\mu\nu} F'^{\mu\nu} \text{ (vector portal)} \quad (1.11)$$

$$\mathcal{L} \supset (\mu\phi + \lambda\phi^2) H^\dagger H \text{ (Higgs portal)}, \quad (1.12)$$

where $B_{\mu\nu}$ is the hypercharge $U(1)_D$ vector boson field strength, $F'^{\mu\nu}$ is the dark $U(1)_D$ vector boson field strength, and H is the Higgs doublet.

Focusing on the vector portal interactions, a popular model is that of the minimally mixed dark photon. The dark photon vector field A' has the Lagrangian

$$\mathcal{L}_{A'} = -\frac{1}{4} F'^{\mu\nu} F'_{\mu\nu} + \frac{\epsilon}{2 \cos \theta_W} B_{\mu\nu} F'^{\mu\nu} - \frac{1}{2} m_{A'}^2 A'^\mu A'_\mu, \quad (1.13)$$

where ϵ is some kinetic mixing parameter and $m_{A'}$. This model is one of the simplest dark sectors and could also represent the mediator of a larger dark sector. That is, this model can be extended to include a DM candidate that is not the mediator itself. The DM candidate could be a fermion or a scalar boson with coupling to the dark photon through dark-sector gauge interactions. The dark photon would be the mediator for interactions between DM and the Standard Model particles.

If the DM in this dark sector achieved its current abundance through the process of thermal freeze-out, then there are two distinct annihilation processes that it could follow, depending on the hierarchy of the DM and dark photon masses. If the DM candidate χ were heavier than the mediator ($m_\chi > m_{A'}$), then the DM would follow “secluded” annihilation, where it would annihilate into a pair of mediators, and the mediators would decay to Standard Model particles [40]. In this case, the annihilation rate scales as

$$\langle \sigma v \rangle \propto \frac{g_D^4}{m_\chi^2}, \quad (1.14)$$

where g_D is the coupling between the DM and the dark photon mediator. Note the lack of coupling between the mediator and the Standard Model $g_{\text{SM}} = \epsilon e$. This coupling could be vanishingly small and still be consistent with thermal freeze-out of DM. Alternatively, if the

DM candidate were lighter than the mediator ($m_\chi < m_{A'}$), then the DM would follow “direct” annihilation, where it would annihilate to Standard Model particles directly via exchange of a virtual mediator. In this case, the annihilation rate scales as

$$\langle\sigma v\rangle \propto \frac{g_D^2 g_{\text{SM}}^2 m_\chi^2}{m_{A'}^4}. \quad (1.15)$$

With the explicit dependence on the coupling between the mediator and the Standard Model, the direct annihilation channel leads to a predictive target for discovery or falsifiability. The dark coupling and the mass ratio $m_\chi/m_{A'}$ are at most $\mathcal{O}(1)$, giving a minimum value of the coupling between the mediator and the Standard Model. For the vector portal, this is at the level of $\epsilon \sim 10^{-7} m_{A'}^2/(m_\chi \text{MeV} \sqrt{\alpha_D})$, where $\alpha_D = g_D^2/4\pi$, providing a benchmark for mediator and DM searches.

It is also possible that the DM relic abundance was not set by a thermal freeze-out process. One such candidate is asymmetric dark matter (ADM), where the abundance is set by a primordial asymmetry, analogous to that of the baryon asymmetry leading to more matter than anti-matter. In a hidden-sector model of ADM where its asymmetry comes from the same mechanism that causes the baryon asymmetry, the abundances of the baryons and the DM are then naturally related by $\Omega_\chi \sim r \Omega_b m_\chi/m_p$, where r is an $\mathcal{O}(1)$ number that depends on the operator that maintains the chemical equilibrium between the dark matter sector and the Standard Model early in the universe [41, 42]. With the observed abundances, this motivates GeV-scale ADM. On the other hand, if the mechanism of asymmetry of DM is unrelated to the baryons, then r could have any value, and there would be no natural expectation of the ADM mass. Models with $r \ll 1$ would correspond to DM that is much lighter in mass than that of the proton, and those with $r \gg 1$ giving DM with very large masses.

Another possibility of DM relic abundance coming from a process that was not thermal freeze-out is for DM with masses near the QCD confinement scale of about 100 MeV, giving a hidden sector version of QCD [43, 44]. A number-changing process between three and two DM particles can deplete the DM abundance and achieve the correct relic density [45–47]. These Strongly Interacting Massive Particles (SIMPs) would require keeping the hidden sector and Standard Model particles in kinetic equilibrium via elastic scattering until the number-changing process freezes out. Instead of the number-changing process, the DM relic abundance could instead be determined by the cross-section of elastic scattering on Standard Model particles [48, 49], in this case the DM particles are Elastically Decoupling Relics (ELDERs). For either mechanism of determining the relic abundance, both ELDERs and SIMPs have predicted masses of few to a few hundred MeV, well within the LDM range.

The DM could also not be in thermal equilibrium, but still have a small interaction with ordinary matter, leading to a “freeze-in” mechanism where Standard Model particles slowly annihilate or decay into DM [50]. For hidden-sector DM with very weak mixing, this process could consist of the Standard Model particles freezing out into dark sector mediators, which then decay into DM. These couplings would be quite small with a light vector-portal

mediator, and these freeze-in models via the hidden sector predict DM candidates with masses in the LDM scale [51–53].

All of these LDM candidates represent the vibrant interest in this mass scale, giving experimentalists plenty of theoretical motivation to perform LDM searches. As the mass scales are much less than the WIMP mass scales, to search for these candidates requires the development of a different class of detectors optimized instead for LDM.

1.2.3 Other Candidates

In Fig. 1.9, there is also a whole regime of ultralight dark matter (ULDM) for DM that is bosonic and has a sub-keV mass. The DM for these masses cannot be fermionic, as the Fermi degeneracy pressure would prevent the galactic substructure formation at the scale of dwarf galaxies from clumped fermionic DM [31]. If the DM had a mass lower than 10^{-22} eV then the Compton wavelength would be larger than observed dwarf galaxies and would not be meaningfully bound to these galaxies. The most famous of these bosonic ULDM candidates is the QCD Axion, which was proposed as a solution to the strong CP problem [54–57]. These bosonic candidates could potentially be detected through absorption in some semiconducting material (such as the detectors that will be discussed in this thesis), where, e.g., the DM would excite an electron to the conduction band and the emitted photons or phonons could be detected after relaxation to the ground state.

At the very high mass end of DM candidates, there exist the possibility of primordial black holes (PBHs) as a DM candidate. These black holes could have formed in the early universe as some regions became so compressed that they experienced gravitational collapse to form black holes [58]. PBHs could have a very large range of masses, and the observation of merging black holes with masses of $30M_{\odot}$ by LIGO [59] has directed interest to this order of magnitude for PBHs as DM [60, 61]. To search for PBHs, a commonly used method is gravitational microlensing, where a lensing object crosses the line of sight of a background star and creates time-varying magnification the star [62–64]. Furthermore, this method has been used to look for Massive Compact Halo Objects (MACHOs) as a DM candidate, for which microlensing searches for MACHOs set an upper limit of 8% of the DM abundance at lower masses [65, 66]. At the mass scale of the LIGO black holes, studies of dwarf galaxies [67, 68] and the CMB [69] have ruled these out as DM candidates. However, there were complex assumptions and astrophysics involved in these constraints, whereas microlensing searches are largely independent of these, providing some motivation to continue searching for PBHs/MACHOs at these mass scales.

There is always the possibility that DM is solely gravitationally interacting. In this sometimes called “nightmare scenario,” there do exist novel proposals for directly detecting DM by probing its gravitational signatures, from pulsar-timing array probes [70] to an array of quantum-limited mechanical impulse sensors for Planck-scale DM [71] to accelerometer networks [72]. In many theoretical models of purely gravitationally-interacting DM, a common problem becomes that there is an overabundance of DM due to a lack of a mechanism for de-

pleting the primordial DM [33, 73], reiterating that having some coupling to Standard Model particles that is non-gravitational is well-motivated to achieve our current abundances.

1.3 Direct Detection of Dark Matter

When searching for DM, the different experimental detection methods are usually split up into three categories: direct detection, indirect detection, and DM production. While this section will focus on the principles of direct detection, the importance of indirect detection and DM production should not be dismissed. While direct detection experiments aim to detect the DM itself, indirect detection experiments attempt to detect Standard Model particles that are the products of the annihilations or decays of DM, such as cosmic rays, neutrinos, or photons. For DM production, the goal is use a high-energy particle collider, such as the LHC, to produce DM particles. These particles would escape the detector, which would leave a signature that there is an excess of events that have missing energy or momentum. These different methods, alongside astrophysical probes of the universe, work in tandem to search for DM, providing clear complementarity between them [74].

In direct detection, we have some arbitrary detector of volume V and mass density ρ_T , which we place in some environment to hopefully interact with DM. By using Fermi's Golden Rule, we can calculate the generic scattering rate for DM per unit target mass for this arbitrary detector [33, 75]

$$R_\chi = \frac{1}{\rho_T} \frac{\rho_\chi}{m_\chi} \int d^3\mathbf{v} f_\chi(\mathbf{v}) \frac{V d^3\mathbf{p}'_\chi}{(2\pi)^3} \sum_f |\langle f, \mathbf{p}'_\chi | \Delta H_{\chi T} | i, \mathbf{p}_\chi \rangle|^2 2\pi \delta(E_f - E_i + E'_\chi - E_\chi), \quad (1.16)$$

where ρ_χ is the DM energy density, m_χ is the DM mass, $f_\chi(\mathbf{v})$ is the DM velocity distribution at the detector, $|i\rangle$ is the initial detector state with energy E_i , $|f\rangle$ is the final detector state with energy E_f , and $\Delta H_{\chi T}$ is the Hamiltonian describing DM-target interactions. We will assume that this interaction Hamiltonian is nonrelativistic, which follows from the astrophysical observations that DM is cold.

In order to simplify Eq. (1.16), we follow the formalism in Ref. [33]. We make the following assumptions: the interaction Hamiltonian can be treated as a perturbation of the free-particle DM Hamiltonian (the unperturbed eigenstates are plane waves $|\mathbf{p}\rangle$), there is no entanglement between the target and the DM ($|a, \mathbf{p}_\chi\rangle \equiv |a\rangle \otimes |\mathbf{p}_\chi\rangle$, where a can be i or f), and that there is a single operator $\mathcal{O}$ that dominates the interaction Hamiltonian. With the assumption of no entanglement, we have that $\mathcal{O} = \mathcal{O}_\chi \otimes \mathcal{O}_T$, where $\mathcal{O}_\chi$ is the operator acting on the DM state and $\mathcal{O}_T$ is the operator acting on the target state. Thus, we can factorize

the matrix element into Fourier components $\mathbf{q}$, such that

$$\langle f, \mathbf{p}'_\chi | \Delta H_{\chi T} | i, \mathbf{p}_\chi \rangle \equiv \int \frac{d^3 \mathbf{q}}{(2\pi)^3} \langle \mathbf{p}'_\chi | \mathcal{O}_\chi(\mathbf{q}) | \mathbf{p}_\chi \rangle \times \langle f | \mathcal{O}_T(\mathbf{q}) | i \rangle \quad (1.17)$$

$$= \frac{1}{V} \sqrt{\frac{\pi \bar{\sigma}(q)}{\mu_\chi^2}} \langle f | \mathcal{O}_T(\mathbf{q}) | i \rangle, \quad (1.18)$$

where we have used our DM plane wave states and redefined the interaction potential in terms of a interaction potential strength $\bar{\sigma}(\mathbf{q})$ and a mass parameter μ_χ . Note that $\mathcal{O}_T$ could still depend on a DM model, e.g. the DM coupling strengths to Standard Model particles in the target system.

We can modify Eq. (1.16) further by using a dummy variable ω , such that

$$\delta(E_f - E_i + E'_\chi - E_\chi) = \int d\omega \delta(\omega + E'_\chi - E_\chi) \delta(E_f - E_i + \omega). \quad (1.19)$$

Thus, we apply our relations in Eqs. (1.18) and (1.19) to Eq. (1.16) to find a scattering rate of

$$R_\chi = \frac{1}{\rho_T} \frac{\rho_\chi}{m_\chi} \int d^3 \mathbf{v} f_\chi(\mathbf{v}) \int \frac{d^3 \mathbf{q}}{(2\pi)^3} d\omega \delta(\omega + E'_\chi - E_\chi) \frac{\pi \bar{\sigma}(q)}{\mu_\chi^2} \times S(\mathbf{q}, \omega), \quad (1.20)$$

where we define the dynamic structure factor

$$S(\mathbf{q}, \omega) \equiv \frac{2\pi}{V} \sum_f |\langle f | \mathcal{O}_T(\mathbf{q}) | i \rangle|^2 \delta(E_f - E_i - \omega). \quad (1.21)$$

The dynamic structure factor can be thought of as the response of the target material to the DM interaction, as it contains all of the dynamics of the target system (e.g. the electron or phonon excitations of the crystal).

Turning our attention to the kinematics of the DM scattering, we say there is incoming DM with momentum $\mathbf{p}_\chi = m_\chi \mathbf{v}$ that scatters off a detector target, exiting with momentum $\mathbf{p}'_\chi$. For nonrelativistic DM, the energy deposited in the target is

$$\begin{aligned} \omega_{\mathbf{q}} &= E_\chi - E'_\chi \\ &= \frac{1}{2} m_\chi v^2 - \frac{(m_\chi \mathbf{v} - \mathbf{q})^2}{2m_\chi} \\ &= \mathbf{q} \cdot \mathbf{v} - \frac{q^2}{2m_\chi}, \end{aligned} \quad (1.22)$$

where $E_\chi = p_\chi^2/2m_\chi$ and $E'_\chi = p'^2_\chi/2m_\chi$. Thus, Eq. (1.22) defines the kinematically-allowed region of DM scattering. For given ω and $\mathbf{q}$, the minimum DM initial velocity required for scattering is then

$$v_{\min}(q, \omega) = \frac{\omega_{\mathbf{q}}}{q} + \frac{q}{2m_\chi}. \quad (1.23)$$

Table 1.1: Example of some galactic constants used in DM searches.

Parameter	Value	Ref.
ρ_0 [GeV/cm ³]	0.3	[79]
v_0 [km/s]	220	[80]
v_E [km/s]	232	[76]
v_{esc} [km/s]	544	[81]

We can take the isotropic approximation of Eq. (1.20) in order to see this minimum velocity explicitly. Carrying out the ω integral, the isotropic rate is

$$R_{\chi}^{\text{iso}} = \frac{1}{\rho_T} \frac{\rho_{\chi}}{m_{\chi}} \int \frac{q dq}{(2\pi)^2} d\omega \eta(v_{\min}(q, \omega)) \times \frac{\pi \bar{\sigma}(q)}{\mu_{\chi}^2} \times S(q, \omega), \quad (1.24)$$

where $\eta(v_{\min}(q, \omega))$ is defined as

$$\eta(v_{\min}) \equiv \int_{v_{\min}}^{\infty} d^3\mathbf{v} \frac{f_{\chi}(\mathbf{v})}{v}. \quad (1.25)$$

In DM searches via annual modulation of the DM flux or using an anisotropic target system, the isotropic approximation will not hold, and proper care of these systems is required.

For the DM velocity distribution, it is common to use the Standard Halo Model (SHM), which describes the distribution as a truncated Maxwellian

$$f_{\text{R}}(\mathbf{v}) = \frac{1}{(2\pi\sigma_v^2)^{3/2} N_{\text{R,esc}}} \exp\left(-\frac{|\mathbf{v}|^2}{2\sigma_v^2}\right) \Theta(v_{\text{esc}} - |\mathbf{v}|), \quad (1.26)$$

where $N_{\text{R,esc}}$ normalizes this distribution to 1:

$$N_{\text{R,esc}} = \text{erf}\left(\frac{v_{\text{esc}}}{\sqrt{2}\sigma_v}\right) - \sqrt{\frac{2}{\pi}} \frac{v_{\text{esc}}}{\sigma_v} - \frac{v_{\text{esc}}^2}{2\sigma_v^2}. \quad (1.27)$$

In the SHM, the values used are the galactic escape velocity v_{esc} , the mean DM speed $v_0 = \sqrt{2}\sigma_v$, the local DM density ρ_0 . As discussed in Ref. [76], it is also common to include a small correction taking the Earth's velocity in the galaxy v_E into account, which will affect the tail end of the distribution. There are systematic errors associated with all of these values (most notably in ρ_0) [77, 78], such that the recommended values have changed over time. In Table 1.1, we include values that are historically used (and have been used in the DM search that will be discussed in this thesis).

For DM-nucleon scattering, we can assume that the DM mass is large enough that the kinetic energy is greater than any displacement energy of the target, and its momentum is greater than any zero-point lattice momenta of the target. This is not always true, and as

searches probe lower DM masses in the LDM range, much more care will need to be taken to ensure the expected differential rate is correct. With this assumption, we can treat the nuclear target as a free particle that is initially at rest. Using Eq. (1.22) and the energy deposited being $\omega = E_R = q^2/(2m_N)$ for a nucleus N , we find that

$$\mathbf{q} \cdot \mathbf{v} = \frac{q^2}{2\mu_{\chi N}}, \quad (1.28)$$

where $\mu_{\chi N}$ is the DM-nucleus reduced mass. From Eq. (1.28), we have that the maximum momentum transfer is $q_{\max} = 2\mu_{\chi N}v$, giving a maximum nuclear recoil energy of $E_{R,\max} = q_{\max}^2/(2m_N) = 2\mu_{\chi N}^2 v^2/m_N$. The best kinematic match occurs when $m_\chi \approx m_N$, and this energy transfer becomes quite inefficient when $m_\chi \ll m_N$. From Eq. (1.23), we can calculate the minimum DM velocity needed for a nuclear recoil of energy E_R

$$v_{\min} = \sqrt{\frac{m_N E_R}{2\mu_{\chi N}^2}}. \quad (1.29)$$

For elastic nuclear scattering, we can assume a contact potential between the DM and nucleon, such that the interaction Hamiltonian is

$$\Delta H_{\chi T} = \sqrt{\frac{\pi \bar{\sigma}_n}{\mu_{\chi n}^2}} \int \frac{d^3 \mathbf{q}}{(2\pi)^3} e^{i\mathbf{q} \cdot (\mathbf{r}_n - \mathbf{r}_\chi)}, \quad (1.30)$$

where $\bar{\sigma}_n$ is the DM-nucleon scattering cross section and $\mu_{\chi n}$ is the DM-nucleon reduced mass (as opposed to the entire nucleus), assuming the same coupling to protons and neutrons. Summing over the target nuclei N_{nuc} , the dynamic structure factor is

$$S(\mathbf{q}, \omega) = \frac{2\pi N_{\text{nuc}}}{V} A^2 |F_N(q)|^2 \delta\left(\omega - \frac{q^2}{2m_N}\right), \quad (1.31)$$

where A is the mass number and F_N is a nuclear form factor

$$F_N(q) = \langle N | \frac{1}{A} \sum_{\alpha=1}^A e^{i\mathbf{q} \cdot \mathbf{r}_\alpha} | N \rangle. \quad (1.32)$$

For these elastic nuclear recoils, this nuclear form factor is commonly taken to be the Helm form factor [82]. This form factor is relevant for WIMP DM, but is effectively unity for LDM. Using Eq. (1.24), we can calculate the isotropic rate and integrate over ω to arrive at the elastic nuclear recoil rate

$$R_\chi = N_T \frac{\rho_\chi}{m_\chi} \frac{A^2 \bar{\sigma}_n}{2\mu_{\chi n}^2} \int q dq \eta(v_{\min}(E_R)) F_N^2(q), \quad (1.33)$$




Figure 1.10: (Figure from Ref. [84]) The landscape of constraints on sub-GeV scale LDM through the DM-nucleon scattering channel.

where $N_T = N_{\text{nuc}}/(\rho_T V)$ is the number of target nuclei per detector mass. By making the substitution $q dq \rightarrow dE_R m_N$, we arrive at the differential rate for DM-nucleon scattering

$$\frac{\partial R_\chi}{\partial E_R} = m_T \frac{\rho_\chi}{m_\chi} \frac{A^2 \bar{\sigma}_n}{2\mu_{\chi n}^2} \eta(v_{\min}(E_R)) F_N^2(q), \quad (1.34)$$

where m_T is the mass of the target material. The $\eta(v_{\min})$ integral can be analytically solved, as done in Ref. [83]. We note that this function is exponential in nature (with some corrections from the cutoffs in the DM velocity distribution), with a characteristic form of

$$\eta(v_{\min}) \sim \exp\left(-\frac{m_N E_R}{2\mu_{\chi N}^2 \sigma_v^2}\right). \quad (1.35)$$

This differential rate spectrum has been used in many DM experiments to set limits on the DM-nucleon scattering cross section, with the “current” limits on LDM shown in Fig. 1.10.

For DM-electron scattering and DM-nucleon scattering with phonon-scale recoil energies, the dynamic structure factor is no longer that of simple elastic nuclear scattering, as seen above. Instead, one must take into account the crystal structure and its wavefunctions in order to correctly take into account the expected response of the target material. In Fig. 1.11, we show example electron and phonon band structures for silicon. These band structures can




Figure 1.11: (Figure from Ref. [92]) (Left) Electronic band structure for Si. (Right) Phonon band structure for Si.

be calculated from first principles via density functional theory (DFT). Common software to carry out these complex calculations are VASP [85], QUANTUM ESPRESSO [86], and PHONOPY [87]. The outputs from the DFT calculations must then be post-processed in order to calculate the dynamic structure factor for various DM models, which can be done using existing software such as QEDARK [51], QEDARK-EFT [88], EXCEED-DM [89], or DARKELF [90, 91].

Chapter 2

Athermal Phonon Sensors

In this chapter, I will start with a brief discussion on the need for athermal phonon sensors, derive the small signal model frequently used when characterizing Transition-Edge Sensors (TESs), discuss the intrinsic noise sources of the system, and expand these concepts to Quasiparticle-trap-assisted Electrothermal-feedback Transition-edge-sensors (QETs) and the collection of athermal phonons.

2.1 Athermal Calorimeters

When a particle interacts with some detector, the standard steps that follow are that the detector absorbs some energy from the particle, this absorption creates athermal excitations in the detector (e.g. prompt phonons, ionization, photons), inelastic scattering thermalizes these excitations, and the detector cools back to equilibrium. These steps are shown diagrammatically in Fig. 2.1, where the internal components of the detector are not shown (but would include some sensor and absorber).

In this process, the detector target (absorber) has some heat capacity C_{abs} which is connected to the thermal bath with a thermal conductance G_{ab} . When measuring thermal excitations, the sensor and the absorber should ideally have a strong thermal coupling, such that they act as a single thermal element. In Fig. 2.2, we show the pulse behavior that is experienced by the sensor (or thermometer). In this diagram, the temperature rises with some characteristic time constant and then falls with a different characteristic time constant. These time constants are determined by the thermalization time and C_{abs}/G_{ab} , where the rise corresponds to the faster of the two, and the decay to the slower. In this system, the ideal energy sensitivity (or baseline energy resolution) is limited by thermal fluctuations across the thermal link to the bath, giving $\sigma_E^2 \approx K_B T^2 C_{abs}$, implying that operating at low (cryogenic) temperatures will improve energy sensitivity.

The thermal conductance between the thermometer (sensor) and the absorber is driven in thin metal films by the electrons in the sensor being pushed far out of thermal equilibrium (the hot-electron effect [94]) and will scale as T^4 . On the other hand, the thermal conductance




Figure 2.1: (Figure from Ref. [93]) (a) A particle is absorbed in a detector. (b) The absorption creates athermal excitations. (c) These athermal excitations begin to thermalize. (d) The detector returns to equilibrium.




Figure 2.2: (Figure from Ref. [93]) Pulse behavior experienced by a sensor (or thermometer), where these a rise in temperature, a peak, and then a fall in temperature as it returns to equilibrium.

between the absorber and the bath scales as T^3 from the elastic constant mismatch between dissimilar semiconducting substrates (known as phonon mismatch or Kapitza coupling) [95]. Thus, we have the caveat that, as we decrease our temperature to improve energy sensitivity, it becomes difficult to thermally couple the thermometer and the absorber.

To get around this, one can use athermal calorimeters. That is, we put a sensitive sensor with fast response times that can measure the athermal excitations before they thermalize, such as a superconducting tunnel junction or a TES. In the case of TESs, this makes the limiting heat capacity that of the sensor and changes the pulse shape, where the time constants are determined by the athermal collection time (significantly faster than the thermalization time) and the sensor's thermal time constant (rather than being defined by the thermal time constant between the absorber and bath). TESs can successfully act as these sensors, especially when we expand the design to QETs, as will be discussed in the next sections.

2.2 Transition-Edge Sensors

At its core, the idea behind a TES is quite simple. A TES is some metal (commonly W in SuperCDMS) which has been heated by a suitable current to operate within its superconducting transition. This superconducting transition is very sharp with widths of the order of μK , with an example shown in Fig. 2.3. Thus, we have some temperature-dependent resistor (or thermistor), which is highly sensitive to temperature perturbations. Furthermore, this thermistor will be sensitive as well to current perturbations (current supplies Joule heating), and its dynamics will be described by two nonlinear coupled differential equations. When operating these sensors in the voltage-biased limit, they become self-regulating through the idea of negative electrothermal feedback. When heat (e.g. some event hits the TES) increases the resistance of the TES, the Joule heating ($P_J = V_{TES}^2/R_{TES}$) decreases—this decreases the temperature (or heat) of the TES, which then increases Joule heating, and we have that the TES self-regulates its temperature and returns to its steady-state. The alternative operating method of current-biasing leads instead to positive electrothermal feedback (where an increase in resistance instead increases the Joule heating, giving rise to thermal runaway), which made it historically difficult to keep TESs in this mode stable. First suggested by Irwin *et al.* [96], TESs have since been exclusively run in the voltage-biased mode.

Because the transition is sharp, this means that any small temperature change will result in a large change in resistance. Thus, by reading out the current through the TES via a circuit, one can measure a current change, which we will see can be related to the energy of some event. Generally, the actual circuit used is of the form shown in Fig. 2.4, where the readout circuit is based on a Superconducting Quantum Interference Device (SQUID). To define a few terms in Fig. 2.4, we have some applied bias current I_{Bias} , a shunt resistor R_{sh} , some parasitic resistance R_p , an inductor with inductance L , and our TES with a variable resistance R_{TES} . With this circuit, it is common to approximate it as the Thévenin equivalent voltage-biased circuit, as shown in Fig. 2.5, as we are generally operating in the voltage-biased mode to take advantage of negative electrothermal feedback.




Figure 2.3: (Figure from Ref. [97]) A characteristic plot of the transition of a TES from its normal state to its superconducting state with a change in operating temperature.




Figure 2.4: (Figure from Ref. [98]) A schematic of the usual circuit for reading out changes in current of the TES.




Figure 2.5: (Figure from Ref. [99]) Conversion of the circuit schematic (left) to the Thévenin equivalent circuit (right).

In the Thévenin equivalent circuit, we define the voltage bias $V_b = I_{Bias}R_{sh}$ and the load resistance $R_\ell \equiv R_p + R_{sh}$. The resulting equations are mathematically equivalent to those of the circuit with the bias current, but redefining the circuit in this way simplifies some of the algebra with these new definitions. From this circuit, we can use Kirckhoff's laws to come to a differential equation describing this circuit

$$L \frac{dI}{dt} = V_b - IR_\ell - IR(T, I) + \delta V, \quad (2.1)$$

where we have explicitly shown that the TES resistance is a function of both TES current and temperature, i.e. $R_{TES} \equiv R(T, I)$, and δV represents any change (assuming a first order expansion) in bias voltage from internal (e.g. noise fluctuations) or external sources (e.g. a change in TES resistance due to an upstream-related change in the voltage bias).

The other differential equation comes from thermal considerations of our system. Our TES will be in contact with some thermal bath (or reservoir), which we take to have a heat capacity of effectively infinity. In other words, the heat capacity is so large that its temperature does not change with a reasonable amount of heat (i.e. an amount that we would expect some particle to impart given some recoil). Given this, we can represent the thermal interactions between the TES and the bath through its heat capacity:

$$C \frac{dT}{dt} = P_J - P_{bath} + \delta P, \quad (2.2)$$

where $P_J = I^2 R(T, I)$ is the Joule heating, P_{bath} is the power flowing from the TES to the thermal bath, and δP represents any change (again assuming a first order expansion) in

power due to internal (e.g. noise fluctuations) or external sources (e.g. energy imparted into the TES). Generally, the P_{bath} term will follow a power law that is related to the temperature of the TES and the temperature of the bath. For metals with small volume and high power densities, the power flow P_{bath} from the electrons to the phonons will dominate with the form $P_{bath} = K(T_c^n - T_{bath}^n)$, where $K = \Sigma_{ep} V_{TES}$ and $n = 5$ [97]. The coefficient K is proportional to the volume of the TES by some material dependent constant which is usually on the order of $10^9 \text{ W} \cdot \text{m}^{-3} \cdot \text{K}^{-5}$. Note that this exponent of $n = 5$ is specific to our assumptions on TES volume and power density. In practice, this exponent should be verified through a G_{TA} measurement, as discussed in Ref. [98] and further detailed in Appendix A of this thesis. Thus, our thermal nonlinear differential equation becomes

$$C \frac{dT}{dt} = I^2 R(T, I) - K(T^n - T_{bath}^n) + \delta P. \quad (2.3)$$

At this point, we have arrived at the two nonlinear differential equations that govern the simplest TES systems. Before investigating the response times of these systems, it is worth gaining intuition on the zero-frequency components, or DC characteristics.

2.2.1 DC Characteristics

When discussing DC characteristics of TESs, we return to Kirckhoff's laws as they apply to the circuit in Fig. 2.5. Taking the zero-frequency component of Eq. (2.1) (i.e. setting the time derivatives and the δV terms to zero), the equation simplifies to

$$V_b = I(R_\ell + R(T, I)), \quad (2.4)$$

which returns us to our well-known Ohm's law of $V = IR$. There are two natural extremes to take this equation to: the limit of superconducting resistance and the limit of normal resistance. Both of these limits provide (while seemingly basic) important insight on the DC characteristics of the TES.

Starting with the superconducting resistance regime (i.e. when the temperature of the TES is well-below its superconducting transition temperature and has zero resistance), it is important to choose (as will be apparent later when we derive the baseline energy resolution equations) metals with low superconducting transition temperatures, and conventional BCS (Bardeen-Cooper-Schrieffer) [100] superconductors are commonly-used and preferred. In the superconducting regime of the TES circuit, the relation between the load resistance R_ℓ and the shunt resistance R_{sh} is

$$V_b = I_{Bias} R_{sh} = I R_\ell. \quad (2.5)$$

From an experimentalist's point of view, we should know I_{Bias} (our applied current bias to the circuit) and R_{sh} (the shunt resistor which we have installed on the circuit), and we can use this information to extract the parasitic resistance R_p with a strong degree of numerical certainty.

Continuing to the normal resistance regime (i.e. when the temperature of the TES is well-above its superconducting transition temperature), we have that the TES can be approximated to first order as having a constant resistance R_N . Generally, it is not true that normal resistance is constant for a metal (or other materials) with temperature from T_c to ∞ . However, it is a very good approximation at cryogenic temperatures for metals. When operating a TES above its transition at, e.g., 50 mK the change in resistance is quite negligible at 100 mK. For the normal regime, we have that Ohm's law gives

$$V_b = I(R_\ell + R_N), \quad (2.6)$$

where R_N is the normal resistance of the TES near its superconducting transition. This is again a linear relationship between the voltage bias and the current, both of which one will know in their experimental setup. Thus, the two limits give a strong reason to vary the voltage bias and observe the change in current, as this will provide direct insight on the normal resistance and the parasitic resistance.

For the approximation of TES dynamics, it is the transition region that is of upmost importance for science results because of its sensitivity to temperature fluctuations. Historically, one could voltage-bias a TES in transition and achieve a competitive result due to the novelty of the technology. Given the technology's maturity at the time of this thesis, one must further understand the mathematical model of the TES to begin to optimize its performance, and this begins with an understanding of the IV curve. In experimentally-specific terms, this is a measurement of how the variation of the bias current I_{Bias} relates to the measured current through the TES $I_{TES} \equiv I$ via, e.g., the SQUID array. Importantly, our understanding within the superconducting transition comes from both P_{bath} and the minuscule $\mathcal{O}(\mu\text{K})$ width of the superconducting transition. In other words, we can see that this small width allows us to approximately treat the equilibrium bias power ($P_0 = P_{bath}$) as a constant. Thus, within the transition with our constant bias power, we can use the voltage-current product law to give

$$V_b = P_0/I + IR_\ell, \quad (2.7)$$

which effectively follows a $1/x$ curve (in the usual case of $R_\ell \ll R_{TES}$).

Measuring TES DC Characteristics

Going through the above equations, we now have an expectation of how the current through the TES would change as we change the voltage bias: linearly in the normal region, approximately as $1/V_b$ within the transition to superconducting, and linearly in the superconducting region. In Fig. 2.6, we show this relation for a rectangular W TES of dimensions $200\ \mu\text{m} \times 50\ \mu\text{m} \times 40\ \text{nm}$. Each marker in the figure denotes a measurement of the TES current, while the connecting lines are simple linear interpolations between neighboring points. In practice, when one is measuring the current via a SQUID array, the measured current is only relative, not absolute (i.e. we can only measure changes in current). Because of this,




Figure 2.6: IV curve for a rectangular W TES of dimensions $200 \mu\text{m} \times 50 \mu\text{m} \times 40 \text{ nm}$.

one has to correct for the overall current offset by taking advantage of the linear normal and superconducting regions. In particular, one can linearly fit these regions, find the y -intercepts, and correct the IV curve such that both intercepts are zero. It is useful to do this linear fit for both of these regions, as this can correct for the possibility of some unknown offset in voltage bias (e.g. if there a voltage jitter sent down the TES bias line that has a small DC offset which adds to the current bias I_{Bias}).

Returning to Eq. (2.4), we can (given that we have already measured the load resistance from the superconducting region) rearrange the equation to give us both the TES resistance and the bias power as a function of voltage bias:

$$R_{TES} = \frac{V_b}{I} - R_\ell, \quad (2.8)$$

$$P_0 = IV_b - I^2 R_\ell. \quad (2.9)$$

Thus, it is simple to convert our IV curve in Fig. 2.6 to curves of resistance or bias power as a function of voltage bias, as shown in Figs. 2.7 and 2.8. We can then use our knowledge of our DC characteristics at various bias points as part of the goal of understanding the expected performance of the TES. However, to fully understand the expected performance, one must also understand the expected time-varying response of the TES.

2.2.2 TES Response

To understand the expected response of the TES to some Dirac delta deposit of energy, we need to solve our nonlinear coupled differential equations in Eqs. (2.1) and (2.3), which are




Figure 2.7: Resistance as a function of voltage bias for a rectangular W TES of dimensions $200 \mu\text{m} \times 50 \mu\text{m} \times 40 \text{ nm}$.




Figure 2.8: Bias power as a function of voltage bias for a rectangular W TES of dimensions $200 \mu\text{m} \times 50 \mu\text{m} \times 40 \text{ nm}$. In transition, the bias power is flat due to the sharpness of the superconducting transition in temperature and stabilized through electrothermal feedback.

not analytically solvable in their nonlinear form. However, we can Taylor expand about the equilibrium point assuming that the changes in current and temperature are small, i.e. we are making a small signal approximation. In making this approximation, we will need to make a series of parameter definitions, as done by Irwin and Hilton [97].

Starting with the power flowing to the bath, we have that $P_{bath} = K(T^n - T_{bath}^n)$. Expanding this about some equilibrium temperature T_0 , we have that

$$P_{bath} \approx P_{bath_0} + G\delta T, \quad (2.10)$$

where $G \equiv dP_{bath}/dT = nKT^{n-1}$ is the thermal conductance from TES to bath and $\delta T = T - T_0$. Next, we expand the TES resistance about its steady state values in resistance, temperature, and current

$$R(T, I) \approx R_0 + \left. \frac{\partial R}{\partial T} \right|_{I_0} \delta T + \left. \frac{\partial R}{\partial I} \right|_{T_0} \delta I. \quad (2.11)$$

The above equation is a somewhat messy to use in a derivation due to the partial derivatives, and it is convention to define the unitless temperature sensitivity

$$\alpha \equiv \left. \frac{\partial \log R}{\partial \log T} \right|_{I_0} = \frac{T_0}{R_0} \left. \frac{\partial R}{\partial T} \right|_{I_0}, \quad (2.12)$$

and the unitless current sensitivity

$$\beta \equiv \left. \frac{\partial \log R}{\partial \log I} \right|_{T_0} = \frac{I_0}{R_0} \left. \frac{\partial R}{\partial I} \right|_{T_0}, \quad (2.13)$$

which rearranges the TES resistance equation to

$$R(T, I) \approx R_0 + \alpha \frac{R_0}{T_0} \delta T + \beta \frac{R_0}{I_0} \delta I. \quad (2.14)$$

Next, we expand our Joule power to first order about the steady state values, such that

$$P_J = I^2 R(T, I) \approx P_{J_0} + 2I_0 R_0 \delta I + \alpha \frac{P_{J_0}}{T_0} \delta T + \beta \frac{P_{J_0}}{I_0} \delta I. \quad (2.15)$$

We also define the (unitless) loop gain

$$\mathcal{L} \equiv \frac{P_{J_0} \alpha}{GT_0}, \quad (2.16)$$

and the natural thermal time constant

$$\tau_0 \equiv \frac{C}{G}. \quad (2.17)$$

Thus, we can substitute all the above expansions and definitions into Eqs. (2.1) and (2.3), giving us the linearized versions

$$\frac{d\delta I}{dt} = -\frac{R_\ell + R_0(1 + \beta)}{L}\delta I - \frac{\mathcal{L}G}{I_0 L}\delta T + \frac{\delta V}{L}, \quad (2.18)$$

$$\frac{d\delta T}{dt} = \frac{I_0 R_0(2 + \beta)}{C}\delta I - \frac{1 - \mathcal{L}}{\tau_0}\delta T + \frac{\delta P}{C}. \quad (2.19)$$

One method to solve these two differential equations is to switch to frequency space through a Fourier transform. As convention is always up in the air, we define transform pair (for some function g) for the continuous case as

$$\tilde{g}(f) = \int_{-\infty}^{\infty} dt g(t) e^{-i\omega t}, \quad (2.20)$$

$$g(t) = \int_{-\infty}^{\infty} df \tilde{g}(f) e^{i\omega t}, \quad (2.21)$$

where $\omega \equiv 2\pi f$ for convenience, and for the discrete case as

$$\tilde{g}_n = \frac{1}{N} \sum_{k=-N/2}^{N/2-1} g_k e^{-i\omega_n t_k}, \quad (2.22)$$

$$g_k = \sum_{n=-N/2}^{N/2-1} \tilde{g}_n e^{-i\omega_n t_k}. \quad (2.23)$$

In the context of our differential equations (using the continuous case), we have that $d/dt \rightarrow i\omega$ (i.e. the Fourier transform of the derivative of some function f is equal to $i\omega$ times the Fourier transform of f). Thus, these equations become

$$i\omega \delta I = -\frac{R_\ell + R_0(1 + \beta)}{L}\delta I - \frac{\mathcal{L}G}{I_0 L}\delta T + \frac{\delta V}{L}, \quad (2.24)$$

$$i\omega \delta T = \frac{I_0 R_0(2 + \beta)}{C}\delta I - \frac{1 - \mathcal{L}}{\tau_0}\delta T + \frac{\delta P}{C}. \quad (2.25)$$

At this point, this system of equations can be put into matrix form, which we have done below after rearranging some of the terms

$$\begin{bmatrix} i\omega + \frac{R_\ell + R_0(1 + \beta)}{L} & \frac{\mathcal{L}G}{I_0 L} \\ -\frac{I_0 R_0(2 + \beta)}{C} & i\omega + \frac{1 - \mathcal{L}}{\tau_0} \end{bmatrix} \begin{bmatrix} \delta I \\ \delta T \end{bmatrix} = \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P}{C} \end{bmatrix}, \quad (2.26)$$

and we define the matrix at this point as M . Multiplying by the inverse of M , we can then come to our expected current and temperature responses to small power or voltage excitations

$$\begin{bmatrix} \delta I \\ \delta T \end{bmatrix} = M^{-1} \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P}{C} \end{bmatrix}, \quad (2.27)$$

where the inverse of the matrix is

$$M^{-1} = \frac{1}{|M|} \begin{bmatrix} i\omega + \frac{1-\mathcal{L}}{\tau_0} & -\frac{\mathcal{L}G}{I_0 L} \\ \frac{I_0 R_0(2+\beta)}{C} & i\omega + \frac{R_\ell + R_0(1+\beta)}{L} \end{bmatrix} \quad (2.28)$$

and

$$|M| = \left[i\omega + \frac{R_\ell + R_0(1+\beta)}{L} \right] \left[i\omega + \frac{(1-\mathcal{L})}{\tau_0} \right] + \frac{\mathcal{L}G R_0(2+\beta)}{L C}. \quad (2.29)$$

After going through the algebra, we have finally arrived at our desired form. From Eq. (2.27), we can extract two important relations for characterizing TESs: the complex admittance of the TES circuit $\partial I/\partial V$ and the power-to-current transfer function $\partial I/\partial P$. Starting with the complex admittance, we will assume that there is only a change in voltage δV (i.e. we are setting external signal power $\delta P = 0$). Thus, we can easily find the complex admittance, as it is related by a single matrix element

$$\frac{\partial I}{\partial V} \equiv \frac{\delta I}{\delta V} = \frac{(M^{-1})_{11}}{L}. \quad (2.30)$$

It can be shown that this form simplifies to

$$\frac{\partial I}{\partial V}(\omega) = \frac{1}{\underbrace{R_\ell + i\omega L + R_0(1+\beta) + \frac{R_0 \mathcal{L}}{1-\mathcal{L}} \frac{2+\beta}{1+i\omega\tau_0/(1-\mathcal{L})}}_{Z_{TES}(\omega)}} = \frac{1}{Z_{circ}(\omega)} \quad (2.31)$$

where we have defined the circuit complex impedance as

$$Z_{circ}(\omega) \equiv R_\ell + i\omega L + Z_{TES}(\omega), \quad (2.32)$$

and the TES complex impedance as

$$Z_{TES}(\omega) \equiv R_0(1+\beta) + \frac{R_0 \mathcal{L}}{1-\mathcal{L}} \frac{2+\beta}{1+i\omega\frac{\tau_0}{1-\mathcal{L}}}. \quad (2.33)$$

Note that the circuit complex impedance is what we work with, as in practice we are reading out all of the load resistor, inductor, and TES in series, rather than solely the TES.

For the power-to-current transfer function (also frequently referred to as the responsivity), the steps are nearly identical, but we set $\delta V = 0$ and keep δP nonzero. Thus, we find that the power-to-current transfer function becomes

$$\frac{\partial I}{\partial P} \equiv \frac{\delta I}{\delta P} = \frac{(M^{-1})_{12}}{C}. \quad (2.34)$$

Here we note that the denominator will be exactly the same as it was for the complex admittance, as it is set by the determinant of our original matrix M . In this way, it is clear

that there is a strong relationship between the complex admittance and the power-to-current transfer function. In fact, once we know the complex admittance, then we can fully define the power-to-current transfer function. A useful form of the power-to-current transfer function is

$$\frac{\partial I}{\partial P}(\omega) = -\frac{1}{Z_{circ}(\omega)I_0} \frac{Z_{TES}(\omega) - R_0(1 + \beta)}{R_0(2 + \beta)}, \quad (2.35)$$

from which we see the ease of which we can calculate the power-to-current transfer function from the complex impedance (the reciprocal of the complex admittance). This functional form applies to more complex TES models with multiple thermal bodies, as used by Maasilta in Ref. [101] and generally derived by Lindeman *et al.* in Ref [102] (in which it is shown that the $\partial I/\partial P$ is determined by the $\partial I/\partial V$ without an assumption on the complexity of the thermal modes).

With the knowledge that the complex admittance and power-to-current transfer function have the same frequency dependence (i.e. have the same poles), we can solve for these poles, which will give us the characteristic time constants of the system. One method to do this is to use that the determinant of a matrix is equal to the product of its eigenvalues, giving us a simple method to rearranging the denominator of these functions. Doing this, we have that

$$\frac{\partial I}{\partial V}(\omega) \propto \frac{1}{(i\omega + \lambda_+)(i\omega + \lambda_-)}, \quad (2.36)$$

where the eigenvalues are

$$\lambda_{\pm} \equiv \frac{1}{\tau_{\pm}} = \frac{1}{2} \left[\frac{1}{\tau_{el}} + \frac{1}{\tau_I} \pm \sqrt{\left(\frac{1}{\tau_{el}} - \frac{1}{\tau_I} \right)^2 - 4 \frac{\mathcal{L} R_0 (2 + \beta)}{L \tau_0}} \right], \quad (2.37)$$

and we have defined [97]

$$\tau_{el} \equiv \frac{L}{R_{\ell} + R_0(1 + \beta)}, \quad (2.38)$$

$$\tau_I \equiv \frac{\tau_0}{1 - \mathcal{L}}. \quad (2.39)$$

In these definitions, τ_{el} serves as the electronic time constant of the system, while τ_I is usually thought of as the constant-current time constant. When voltage-biasing, τ_I is not an intrinsic time constant as we are not operating with constant current, and we do not worry about negative values when $\mathcal{L} > 1$. If we were current-biased, then this would be an intrinsic time constant, and $\mathcal{L} > 1$ (leading to a negative τ_I) represents the thermal runaway experienced from positive electrothermal feedback, severely limiting the allowed values of $\mathcal{L}$ to less than 1. On the other hand, if a sensor's detection sensitivity is based on, e.g., the hopping conduction mechanism, such as neutron transmutation doped (NTD) Ge, then the temperature sensitivity α is negative, and current-biasing is the stable operating mode [103].

We have our relationship between the TES small-signal parameters and the physical time constants $\tau_{\pm}$, again noting that they are the same time constants for both the complex

admittance and the power-to-current transfer function. These time constants simplify greatly when we take the limit as the inductance L approaches zero (the low-inductance limit, i.e. $\tau_+ \ll \tau_-$), which allows us to have a little more intuition on their dependence on the various TES parameters than if we simply look at Eq. (2.37). Taking this limit for both time constants, we have the following definitions

$$\tau_{el} \equiv \lim_{L \rightarrow 0} \tau_+ = \frac{L}{R_\ell + R_0(1 + \beta)}, \quad (2.40)$$

$$\tau_{eff} \equiv \lim_{L \rightarrow 0} \tau_- = \tau_0 \frac{1 + \beta + R_\ell/R_0}{1 + \beta + R_\ell/R_0 + (1 - R_\ell/R_0)\mathcal{L}}, \quad (2.41)$$

where τ_{el} is the electrical time constant and τ_{eff} is the zero-inductance effective thermal time constant. The effective time constant can be taken to another limit

$$\lim_{\substack{\beta \rightarrow 0 \\ R_\ell \ll R_0}} \tau_{eff} = \frac{\tau_0}{1 + \mathcal{L}}, \quad (2.42)$$

which implies that to achieve fast sensor time constants, large loop gains $\mathcal{L}$ are needed (which equivalently means large temperature sensitivity α). Comparing to current-biased TESs with the constraint of $\mathcal{L} < 1$, their time constants are significantly slower than the voltage-biased counterparts, leading to a worse baseline energy resolution (this relationship is shown at the end of Section 2.3.2).

Measuring TES Response

With our expected TES response understood, the next step is to experimentally measure and fit the TES response for characterization. This can be done by injecting white noise and convolving the input and response [104]. However, most TES groups have since switched over to either sine wave sweeps or measuring square wave responses. The idea here is that we can bias the TES with some voltage bias and then add on top of the voltage bias some repeating signal (sine wave or square wave) and then deconvolve the measured response to extract the measured complex admittance. This is generally done in Fourier space, as the Fourier transform of a convolution becomes a simple multiplication of the two functions' Fourier transforms. Let's define some repeating signal $V_s(t)$ with Fourier transform $\tilde{V}_s(f)$ which is convolved with the complex admittance to return the TES response $I_{TES}(t)$ (or $\tilde{I}_{TES}(f)$ in frequency domain). We can relate each of these values through the convolution theorem

$$\tilde{V}_s(f) \frac{\partial I}{\partial V}(\omega) = \tilde{I}_{TES}(f), \quad (2.43)$$

remembering that $\omega = 2\pi f$. We divide by the Fourier transform of the signal, and we have that

$$\frac{\partial I}{\partial V}(\omega) = \frac{\tilde{I}_{TES}(f)}{\tilde{V}_s(f)}. \quad (2.44)$$

Thus, we can directly measure the complex admittance by measuring its time-domain response to some known repeating signal, which we can then fit to the small-signal model to extract all of our TES parameters. Practically, it is much faster to use a single square wave as far as data taking, as opposed to a single sine wave. This is because the Fourier transform of a sine wave returns Dirac delta functions at the sine wave frequency (and its negative frequency), whereas the Fourier transform of a square wave is defined across many discrete frequencies. For a single sine wave, we would have to take data over a sweep of frequencies, which can be quite slow when going from, e.g, 1 Hz to 10 kHz. To remedy this, one could use multiple overlapping sine waves, but calibrating the phase becomes difficult as compared to the well-defined start time from a single square wave.

To extract the measured complex admittance when using a square wave, we can start with the Fourier series of a square wave with frequency f_{sg} and peak-to-peak amplitude in voltage bias A_{sg}

$$V_s(t) = A_{sg} \frac{2}{\pi} \sum_{n=1,3,5,\dots}^{\infty} \frac{\sin(2\pi n f_{sg} t)}{n}. \quad (2.45)$$

Taking the (discrete) Fourier transform and using Eq. (2.44), we have that, for each frequency that is an odd integer multiple of the square wave frequency, the measured complex admittance is

$$\frac{\partial I}{\partial V}(\omega) = \begin{cases} \frac{\tilde{I}_{TES}(f)}{A_{sg} \frac{2}{n\pi}} & \omega = 2\pi n f_{sg} \text{ for } n = 1, 3, 5, \dots \\ 0 & \text{otherwise} \end{cases}, \quad (2.46)$$

and we can measure the complex admittance over a large range of frequencies in a single measurement. In practice, we will measure the response of the TES in time domain many times, select data that does not have external backgrounds (e.g. pulses, electronic glitches), Fourier transform and deconvolve the data, and then average the measure complex admittances (keeping track of the standard error of the mean for error propagation). This workflow has been written in `Python` as one of the many features of the `QETPY` package, co-created and co-maintained by C. W. Fink and myself [105].

In Fig. 2.9, we show an example fit to a TES response to a square wave, where the rise time of the exponential is $\tau_+ = 1.7 \mu\text{s}$ and the fall time is $\tau_- = 32.3 \mu\text{s}$, showing that this measurement truly gives us insight in these physical time constants. We can also show this fit in frequency domain, as done in Fig. 2.10. We can see that the roll-off from τ_- corresponds to the bandwidth $1/(2\pi\tau_-) \approx 5 \text{ kHz}$. We can also use the fitted parameters (R_0 and β) and Eq. (2.35) to give corresponding example $\partial I/\partial P$ curves, as shown in Fig. 2.11. The $\partial I/\partial P$ curve has the same poles as the complex admittance.

With these fits, there is an important difference in the number of TES parameters and the number of degrees of freedom in the complex admittance. Specifically, we have six TES parameters as defined in Eqs. (2.32) and (2.33): R_ℓ , L , R_0 , β , $\mathcal{L}$, and τ_0 . However, we only have four effective degrees of freedom. To explicitly show this, we reparameterize the complex impedance as

$$Z_{circ}(\omega) = A(1 + i\omega\tau_2) + \frac{B}{1 + i\omega\tau_1}, \quad (2.47)$$




Figure 2.9: Example complex admittance fit in time domain for a rectangular W TES of dimensions $200 \mu\text{m} \times 50 \mu\text{m} \times 40 \text{ nm}$.




Figure 2.10: (Left) Example absolute value of complex admittance fit in frequency domain for a rectangular W TES of dimensions $200 \mu\text{m} \times 50 \mu\text{m} \times 40 \text{ nm}$. (Right) Example phase of complex admittance fit in frequency domain for a rectangular W TES of dimensions $200 \mu\text{m} \times 50 \mu\text{m} \times 40 \text{ nm}$.




Figure 2.11: (Left) Example absolute value of $\partial I/\partial P$ using values from the complex admittance fit in frequency domain for a rectangular W TES of dimensions $200\text{ }\mu\text{m} \times 50\text{ }\mu\text{m} \times 40\text{ nm}$. (Right) Example phase of $\partial I/\partial P$ using values from the complex admittance fit in frequency domain for a rectangular W TES of dimensions $200\text{ }\mu\text{m} \times 50\text{ }\mu\text{m} \times 40\text{ nm}$.

where we define

$$A \equiv R_\ell + R_0(1 + \beta), \quad (2.48)$$

$$B \equiv \frac{R_0 \mathcal{L}(2 + \beta)}{1 - \mathcal{L}}, \quad (2.49)$$

$$\tau_1 \equiv \frac{\tau_0}{1 - \mathcal{L}}, \quad (2.50)$$

$$\tau_2 \equiv \frac{L}{R_\ell + R_0(1 + \beta)}. \quad (2.51)$$

Thus, we cannot actually fit all six TES parameters individually, as we have a degeneracy of parameters. To solve this problem, we go back to our IV curve analysis, where we solved for the TES resistance R_0 at different bias voltages, as well as the load resistance R_ℓ . We must measure these two quantities before finding the rest of TES parameters from a $\partial I/\partial V$ measurement in order to break these degeneracies. The optimal workflow for characterizing a TES is then carrying out an IV sweep as a function of bias voltage and taking $\partial I/\partial V$ data at each of these bias points, which allows us to fully characterize our TES. For completion's sake, we include the values of the our TES parameters and time constants in Table 2.1 for the TES used as our example.

As shown by I. Maasilta [101], our small signal model can become more complex if there are extra thermal bodies coupling to the TES and the thermal bath. However, the workflow proposed above will still be optimal, as this only means that we must work with a more complicated complex admittance from the same information gained from the IV and $\partial I/\partial V$ sweeps. It is also worth noting that these concepts can be applied to the normal and superconducting regions of the TES. That is, we can calculate the expected complex impedance

Table 2.1: Fitted and calculated TES parameters for the $200\ \mu\text{m} \times 50\ \mu\text{m} \times 40\ \text{nm}$ W TES at $R_0 = 113.5\ \text{m}\Omega$, noting that the shunt resistance R_{sh} is assumed to be $5\ \text{m}\Omega$. These measurements were carried out in the CryoConcept dilution refrigerator (with $T_{bath} \approx 8\ \text{mK}$) located at UC Berkeley in the Pyle lab. Values are grouped by the measurement method used to extract them.

Parameter	Value	Measurement Method
$R_\ell\ [\text{m}\Omega]$	9.35	<i>IV</i> curve
$R_0\ [\text{m}\Omega]$	113.5	
$I_0\ [\mu\text{A}]$	-0.901	
$P_0\ [\text{fW}]$	92.1	
$L\ [\text{nH}]$	261	Transition $\partial I/\partial V$ curve
β	0.470	
$\mathcal{L}$	137	
$\tau_0\ [\text{ms}]$	2.98	
$\tau_+\ [\mu\text{s}]$	1.76	
$\tau_-\ [\mu\text{s}]$	31.3	

for both these regions, which suffer from less degeneracies. The complex impedances for these two regions are

$$Z_{sc}(\omega) \equiv R_\ell + i\omega L \quad (2.52)$$

$$Z_N(\omega) \equiv R_\ell + i\omega L + R_N. \quad (2.53)$$

Thus, each of these equations have a single time constant: the electrical time constant. For each region, the time constant is L/R_ℓ when superconducting and $L/(R_\ell + R_N)$ when normal. After fitting these two regions' responses to a square wave, a natural cross-check is provided for the corresponding values obtained by the *IV* curve, as well as a quick way of measuring them before taking a full *IV* sweep. Though, the R_ℓ and R_N measurements generally have much less uncertainty in the *IV* sweep than from measurements to the complex admittance in these regions.

Infinite Loop Gain Limit and Energy Absorbed by a TES

When characterizing TESs that are operating with large loop gains (i.e. low in transition), the infinite loop gain limit can be useful both for quick estimation of parameters, as well as sanity checks of values. To show this, we will take the limit as $\mathcal{L} \rightarrow \infty$ for both the complex admittance and the power-to-current transfer function, and apply them to a few useful applications.

Starting with the complex admittance, the zero-frequency term in the infinite loop gain limit becomes

$$\lim_{\mathcal{L} \rightarrow \infty} \frac{\partial I}{\partial V}(0) = \frac{1}{R_\ell - R_0}, \quad (2.54)$$

where this is always negative for a voltage-biased TES (i.e. $R_0 \gg R_\ell$). We did not have to make any assumptions on the value of β , as it fortuitously cancels out in this limit. From this limit, we can measure the zero-frequency of our complex admittance and immediately have an estimate of R_0 (assuming we know R_ℓ), without needing an IV curve. This of course will just be an estimate, and the values from the IV curve will be much more accurate in cases of $\mathcal{L}$ is $\mathcal{O}(1-10)$.

For our example TES we have been studying, our fit in terms of the A , B , τ_1 , and τ_2 parameters gives $A = 0.176 \text{ m}\Omega$, $B = -0.282 \text{ m}\Omega$, $\tau_1 = -2.20 \mu\text{s}$, and $\tau_2 = 1.48 \mu\text{s}$. Thus, the zero frequency component in this form is

$$\frac{\partial I}{\partial V}(0) = \frac{1}{A + B}, \quad (2.55)$$

$$= -9.41 \Omega^{-1}. \quad (2.56)$$

Solving for R_0 in the infinite $\mathcal{L}$ limit (using $R_\ell = 9.35 \text{ m}\Omega$ from, e.g., the superconducting complex admittance), we find that $R_0 \approx 115.5 \text{ m}\Omega$. Comparing to the value in Table 2.1 of $R_0 = 113.5 \text{ m}\Omega$, this is only $2 \text{ m}\Omega$ (about 2% error), which is quite close due to $\mathcal{L}$ being $\mathcal{O}(100)$. We can go one step further and estimate the bias power

$$P_0 = \left(\frac{V_b}{R_0 + R_\ell} \right)^2 R_0 \quad (2.57)$$

using that $I_0 = V_b/(R_0 + R_\ell)$. As we should know the bias voltage ($V_b = -110.5 \text{ nV}$ for our example TES), we have estimates of each quantity, and we have that $P_0 \approx 90.5 \text{ fW}$, again only about 2% error off of the value in Table 2.1 of 92.1 fW . Thus, simply knowing the zero-frequency of the complex admittance is quite powerful in the estimates that we can make when in the case of infinite $\mathcal{L}$.

For a very quick back-of-the-envelope calculation of R_0 , one can estimate off of a time domain complex admittance plot. In Fig. 2.9, for example, the zero-frequency component of the complex admittance is the change in TES current divided by the change in voltage bias. In this data, we have a 2.5 nV peak-to-peak square jitter in bias voltage, and a 23.5 nA change in TES current (comparing the current change before and after the TES responds), such that we can estimate the zero-frequency value as $-9.4 \Omega^{-1}$ (recalling the minus sign due to $\partial I/\partial V(0) < 0$ when $\mathcal{L} > 1$) and follow the same steps to estimate various parameters. Thus, even if we were only looking at the TES response on, e.g., an oscilloscope, we could still quickly estimate R_0 and P_0 just from this change in TES current.

Next, we take the same infinite $\mathcal{L}$ limit for the power-to-current transfer function defined in Eq. (2.35) at zero-frequency, which gives

$$\lim_{\mathcal{L} \rightarrow \infty} \frac{\partial I}{\partial P}(0) = \frac{1}{I_0(R_\ell - R_0)}, \quad (2.58)$$

and can be positive or negative, depending on the current I_0 through the TES. In other words, this becomes the zero-frequency complex admittance divided by the TES current, and we can estimate the $\partial I/\partial P(0)$ quickly as well.

We can rearrange Eq. (2.58) such that

$$dP = I(R_\ell - R_0) dI, \quad (2.59)$$

$$= -(V_b - 2IR_\ell) dI, \quad (2.60)$$

where we have substituted $R_0 = V_b/I - R_\ell$. To calculate the energy absorbed by the TES, we note that the change in Joule heating power is equal and opposite in sign to the power absorbed by the TES, i.e. $dP_J = -dP_{abs}$, by energy conservation. Thus, we have that

$$dP_{abs} = (V_b - 2IR_\ell) dI, \quad (2.61)$$

and we can integrate both sides

$$\Delta P_{abs} = \int_{I_0}^{I_0 + \Delta I(t)} dI (V_b - 2IR_\ell), \quad (2.62)$$

$$= (V_b - 2I_0)\Delta I(t) - \Delta I(t)^2 R_\ell, \quad (2.63)$$

Finally, we integrate this change in power over time to calculate the energy absorbed by the TES, also known as the energy removed by electrothermal feedback in the large $\mathcal{L}$ limit

$$E_{\text{ETF}} = \int_0^\infty dt [(V_b - 2I_0)\Delta I(t) - \Delta I(t)^2 R_\ell]. \quad (2.64)$$

In practice, we are looking at a finite chunk of data, such that we cannot actually integrate to infinity. Thus, we must choose some time cutoff to the integral T_{trunc} , such that

$$E_{\text{ETF}} = \int_0^{T_{trunc}} dt [(V_b - 2I_0)\Delta I(t) - \Delta I(t)^2 R_\ell]. \quad (2.65)$$

where this cutoff is generally far enough from the change in current from injected energy to capture the entire change (e.g. if the change in the current has a characteristic fall time τ_- , then integrate up to $T_{trunc} = 7\tau_-$). Note that the direction of pulses have no effect on the integral, as a negative bias would give a negative V_b , a negative I_0 , and a negative $\Delta I(t)$. Whereas a positive bias would result in each of these quantities being positive. Thus, in either case, each term in the integral would end up with the same sign.

2.3 TES Noise and Energy Sensitivity

Understanding our TES parameters is only part of the battle. As we will be using TESs to set DM limits, we need to understand the baseline energy resolution of these devices in

order to know what energy thresholds are feasible. This will be done through the modeling and calculation of the power spectral densities (PSD) for a description of the power of the noise at various frequencies.

When working with these PSDs, it must be remembered that there are two common versions: one-sided and two-sided. The two-sided PSD is defined over positive and negative frequencies, while the one-sided PSD is only defined for positive frequencies. Because data being taken are real-valued, the two-sided PSD is an even function (i.e. the two-sided PSD has the same value for some frequency f and its negative counterpart). Thus, it is common practice to “fold over” the PSD, where we effectively dispose of the negative frequencies and multiply the positive frequencies by a factor of two, giving the one-sided PSD. In this way, the PSD becomes easier to plot on a log-log scale, and we have not lost the total power of the system. We note that all of the PSD definitions that will be made in Section 2.3.1 are using one-sided PSDs, and we have properly taken care of our factors of two. Furthermore, it is also common practice (as the reader will see in Section 2.3.2) to take the square root of the PSD when plotting, creating strange units such as $A/\sqrt{\text{Hz}}$, as the corresponding PSD would have units of A^2/Hz . We will see that there is convenience to doing this when we finally calculate the expected baseline energy resolution of a TES.

2.3.1 Intrinsic Noise Sources

With a TES, we have a handful of intrinsic noise sources that are directly related to the TES characteristics. We will begin with the Johnson noise, discovered by J. B. Johnson [106] and subsequently explained by H. Nyquist via a calculation of the electromotive force due to thermal agitation of electric charge carriers [107].

In our TES circuit, we have that any resistive element will have Johnson noise. For our purposes, it is easiest to split the Johnson noise into the load Johnson noise and the TES Johnson noise. The load Johnson noise will include the Johnson noise of the shunt resistor and any parasitic resistances. As Johnson voltage noise is flat in frequency space, we have that the load Johnson noise is represented by the one-sided PSD

$$S_{V_{load}} = 4k_B \left(\sum_i T_{sh_i} R_{sh_i} + \sum_j T_{p_j} R_{p_j} \right), \quad (2.66)$$

where we generalize that there could be multiple shunt resistances and parasitic resistances all at different temperatures. In practice, we do not have a way of splitting all of the degeneracies between these different resistors, and it is useful to reparameterize this in terms of an effective load temperature

$$T_\ell \equiv \frac{\sum_i T_{sh_i} R_{sh_i} + \sum_j T_{p_j} R_{p_j}}{R_\ell}, \quad (2.67)$$

which allows us to simplify the load Johnson noise in voltage to

$$S_{V_{load}} = 4k_B T_\ell R_\ell. \quad (2.68)$$

Since we know R_ℓ from our IV analysis, we only need to measure the effective load temperature T_ℓ to model this noise component, as will be discussed in Section 2.3.2. When actually measuring the noise of our system, we will be measuring the current noise (as we are reading out the current through the TES circuit). We can simply convert this voltage noise to current noise by using our complex admittance (i.e. our voltage-to-current transfer function)

$$S_{I_{load}}(\omega) = 4k_B T_\ell R_\ell \left| \frac{\partial I}{\partial V}(\omega) \right|^2, \quad (2.69)$$

showing the frequency dependence of the load Johnson noise in current space.

The TES Johnson noise is more complex due to the thermal and current dependencies of the TES creating nonlinear effects. Using nonlinear Markov fluctuation-dissipation techniques for nonlinear resistors developed by R. L. Stratonovich [108, 109], Irwin and Hilton show that the TES Johnson noise in voltage can be approximated as [97]

$$S_{V_{TES}} = 4k_B T_0 R_0 (1 + \beta)^2, \quad (2.70)$$

showing that the voltage noise is elevated by our current sensitivity β . Converting this to current noise is more complex than the using the complex admittance due to the current and temperature dependence of the TES. In Ref. [99], M. Pyle shows that the conversion to current space follows the compact form

$$S_{I_{TES}}(\omega) = 4k_B T_0 R_0 (1 + \beta)^2 \left| \frac{\partial I}{\partial V}(\omega) - I_0 \frac{\partial I}{\partial P}(\omega) \right|^2. \quad (2.71)$$

We now move to the next intrinsic noise source: thermal fluctuation noise (TFN), also referred to as phonon noise. Because our TES is by design sensitive to temperature changes in transition, this brings a sensitivity to thermal fluctuations between the TES and the thermal bath (each having different temperatures). As shown by R. F. Voss and J. Clarke [110], and applied to TESs in, e.g., M. Pyle's thesis [99], these fluctuations across the thermal link will give a noise in power space that is flat across all frequencies:

$$S_{P_{TFN}} = 4k_B T_0^2 G F_{TFN}(T_0, T_{bath}), \quad (2.72)$$

where the unitless F_{TFN} factor [111–113]

$$F_{TFN}(T_0, T_{bath}) = \begin{cases} \frac{(T_{bath}/T_0)^{n+1} + 1}{2} & \text{ballistic phonons} \\ \frac{n}{2n+1} \frac{(T_{bath}/T_0)^{2n+1} - 1}{(T_{bath}/T_0)^n - 1} & \text{diffuse phonons} \end{cases}, \quad (2.73)$$

takes a value from 0.5 to 1. Usually, we are in the case of ballistic phonons. As we want the TFN contribution in current space, we easily do this via our power-to-current transfer function

$$S_{I_{TFN}}(\omega) = 4k_B T_0^2 G F_{TFN}(T_0, T_{bath}) \left| \frac{\partial I}{\partial P}(\omega) \right|^2. \quad (2.74)$$

Lastly, as we are measuring the current through the TES via a SQUID amplifier, this will give us a SQUID and downstream electronics noise component. The treatment of these noise sources can be quite complex, but we generally use a simplified model of $1/f$ noise [114]

$$S_{ISQUID}(\omega) = \left\{ DC_{SQUID} \left[1 + \left(\frac{\omega_{SQUID}}{\omega} \right)^{n_{SQUID}} \right] \right\}^2, \quad (2.75)$$

where DC_{SQUID} is some constant value (usually on the order of a few pA/ $\sqrt{\text{Hz}}$), ω_{SQUID} determines the characteristic frequency of the $1/f$ behavior, and n_{SQUID} is some exponent for generalization of the $1/f$ behavior. In practice, each of these values will be fit from noise measurements in the normal regime of the TES, as explained in Section 2.3.2.

We can combine all of our noise sources to find the total expected noise spectrum in current for our TES in transition

$$S_{I_{tot}}(\omega) = S_{I_{load}}(\omega) + S_{I_{TES}}(\omega) + S_{I_{TFN}}(\omega) + S_{ISQUID}(\omega). \quad (2.76)$$

For the normal and superconducting regimes, the derivations are much simpler, as the resistors act as normal resistors (i.e. there is no TFN). For the normal regime, the load Johnson noise has the same voltage form as Eq. (2.68), but we use our normal resistance complex impedance from Eq. (2.53) to convert to current noise

$$S_{I_{load}}^N(\omega) = \frac{4k_B T_\ell R_\ell}{|Z_N(\omega)|^2}. \quad (2.77)$$

The TES Johnson noise simplifies greatly, where we assume that the temperature of the TES does not change appreciably from its equilibrium (transition) temperature

$$S_{V_{TES}}^N(\omega) = 4k_B T_0 R_N, \quad (2.78)$$

with the corresponding current noise being

$$S_{I_{TES}}^N(\omega) = \frac{4k_B T_0 R_N}{|Z_N(\omega)|^2}. \quad (2.79)$$

Thus, the total expected normal state current noise becomes

$$S_{I_{tot}}^N(\omega) = S_{I_{load}}^N(\omega) + S_{I_{TES}}^N(\omega) + S_{ISQUID}(\omega), \quad (2.80)$$

where the SQUID current noise term is unchanged, as it has no dependence on the status of the TES. In a well-designed system, the SQUID noise term will dominate when a TES is normal and can be easily measured.

When superconducting, the TES Johnson noise will go to zero (the TES resistance is zero), and we have only the load Johnson noise and the SQUID noise. The load Johnson current noise when superconducting is determined by Eq. (2.52), such that

$$S_{I_{load}}^{sc}(\omega) = \frac{4k_B T_\ell R_\ell}{|Z_{sc}(\omega)|^2}, \quad (2.81)$$

and the total superconducting current noise is

$$S_{I_{tot}}^{sc}(\omega) = S_{I_{load}}^{sc}(\omega) + S_{ISQUID}(\omega). \quad (2.82)$$

With the expected current noise and our transfer functions known for the normal, superconducting, and in transition regimes of the TES, this allows us to model the expected noise performance of our system, given that we have knowledge of all of our TES parameters.

2.3.2 Noise Modeling

As we have been using an example W TES of dimensions $200\ \mu\text{m} \times 50\ \mu\text{m} \times 40\ \text{nm}$ throughout this chapter, we will use what we learned about the intrinsic TES noise to allow us to model it, via the small-signal approximation. In order to measure our TES current noise, we generally record many random time domain traces of data at a single bias point. We then select traces that should describe the intrinsic noise well (i.e. remove traces that have noise from external sources such as pulses, vibrations, electronic glitches, etc.), Fourier transform these traces, take the modulus squared of the Fourier transforms, properly normalize (divide by the digitization rate and the number of time bins) to calculate the power spectral density (PSD), and take the mean at each frequency. This is programmatically done in the `calc_psd` function in QETPY [105].

At this point, we have all of our TES parameters, but we will need to use our noise measurements to help us understand our SQUID noise and our load Johnson noise when modeling all of our noise components. To do this, we first start with our normal noise. In this regime, the SQUID noise is dominant, with a small correction from the Johnson noise of the normal TES. Thus, we can model our SQUID noise by fitting the data to our SQUID noise model (including the correction from the TES Johnson noise). From this fit, we have the parameters to model our SQUID noise in the other regimes, where it may not be (and usually is not) the dominant noise source. We show an example of this in Fig. 2.12. Note that, while we have defined a load Johnson noise component when normal in our equations, it is so small (less than $1\ \text{pA}/\sqrt{\text{Hz}}$ over all frequencies) in the normal regime that we neglect it. We have give our fit values in Table 2.2, which also contains the superconducting fit values and will be discussed in the next paragraph.

After modeling our SQUID noise, we return to the effective load temperature T_ℓ . When defining T_ℓ in Section 2.3.1, we noted that it must be measured for proper noise modeling. Fortunately, it can be measured from the superconducting noise. As we know the rest of the parameters in the load Johnson noise component, we can use our superconducting PSD data to estimate (i.e. fit for) T_ℓ , where we use our total noise model and our previous SQUID noise fit. In Fig. 2.13, we show both of these components and the total noise, where we have fit the superconducting noise model to find our effective load temperature, which was found to be $T_\ell = 68.5\ \text{mK}$, as reported in Table 2.2.

With a TES, the most interesting region is the transition region. Our work on the normal and superconducting regions was simply to allow us to fully model the noise of a TES in




Figure 2.12: Modeling of normal current noise for the $200\ \mu\text{m} \times 50\ \mu\text{m} \times 40\ \text{nm}$ W TES. The $1/f$ noise is attributed to the SQUID, with a small discrepancy between the model and the data at $\sim 10\ \text{Hz}$ due to the fit range being restricted to $0.1\text{--}1\ \text{kHz}$ to be above the $60\ \text{Hz}$ spike. The many peaks at $10\ \text{kHz}$ and above can be attributed to RF pickup in the electronics.

transition. There is one more step that we have neglected so far, which is calculating the thermal conductance G . If we are in the limit that $T_0 \gg T_{\text{bath}}$ and small parasitic power (which is usually the case), then we can estimate the thermal conductance as

$$G \approx \frac{nP_0}{T_0}, \quad (2.83)$$

giving $G \approx 7.75\ \text{pW/K}$. In Appendix A, we explain methods to directly measure the thermal conductances of the system, but we will use this approximation for our noise modeling purposes. With all of our parameters now known and in Table 2.2, we can use our derived equations from Section 2.3.1 to calculate our expected total current noise model and compare to the measured current noise, as compared in Fig. 2.14.

In our comparison, we see that the model does not perfectly overlap with what we measured. This is common in the TES field. There can be external noise sources that increase our noise, such as infrared photons or electromagnetic interference. In general, excess noise is frequently experienced, and there have been many explanations put forward [98, 101, 116], but it is always a difficult problem to understand due to the nature of TES sensitivity.

At this point, we have only discussed current noise, due to the easy comparison between the different regions from normal to transition to superconducting. However, we can calculate




Figure 2.13: Modeling of superconducting (S.C.) current noise for the $200\text{ }\mu\text{m} \times 50\text{ }\mu\text{m} \times 40\text{ nm}$ W TES.

Table 2.2: TES noise parameters for the $200\text{ }\mu\text{m} \times 50\text{ }\mu\text{m} \times 40\text{ nm}$ W TES.

Parameter	Value
DC_{SQUID} [pA/ $\sqrt{\text{Hz}}$]	4.52
ω_{SQUID} [s^{-1}]	594
n_{SQUID}	0.812
T_ℓ [mK]	68.5
G [pW/K]	7.75
σ_E [meV]	45.9




Figure 2.14: Modeling of transition current noise for the $200\ \mu\text{m} \times 50\ \mu\text{m} \times 40\ \text{nm}$ W TES. There is a small amount of excess noise above the total modeled noise, which is of unknown origin. It is seldom seen that the noise model from the small signal approximation perfectly matches the measured noise, and most groups see similar excesses [98, 115].

the expected energy sensitivity of a TES via an understanding of its power noise. Since we already know the power-to-current transfer function, we can convert from TES current noise to TES power noise in one step

$$S_{P_{tot}}(\omega) = \frac{S_{I_{tot}}(\omega)}{\left| \frac{\partial I}{\partial P}(\omega) \right|^2}. \quad (2.84)$$

As discussed earlier, we will work instead with the square root of the power noise, which we will define as the Noise-Equivalent Power (NEP)

$$\text{NEP}(\omega) \equiv \sqrt{S_{P_{tot}}(\omega)}. \quad (2.85)$$

In Fig. 2.15, we show this conversion, where each component of the total modeled power noise has also been converted from current noise to power noise by taking advantage of our (at this point) well-known power-to-current transfer function.

2.3.3 TES Baseline Energy Resolution

The importance of calculating the NEP is best shown when using all of our knowledge gained to calculate the expected baseline energy resolution (the expected standard deviation of zero-




Figure 2.15: Modeling of transition power noise, or Noise-Equivalent Power (NEP), for the $200 \mu\text{m} \times 50 \mu\text{m} \times 40 \text{ nm}$ W TES.

energy fluctuations) of a TES to Dirac delta deposits. This is usually done via the optimal filter (OF) formalism, as discussed in detail in, e.g., Refs. [97–99, 112, 117]. The common derived result is that this baseline energy resolution can be expressed as (in our Dirac delta deposit approximation)

$$\sigma_E = \left[\int_0^\infty \frac{d\omega}{2\pi} \frac{4}{S_P(\omega)} \right]^{-1/2}. \quad (2.86)$$

From our measured (one-sided) $S_P(\omega)$, we can numerically integrate this baseline energy resolution, returning the reported value in Table 2.2 of $\sigma_E = 45.9 \text{ meV}$.

In practice, we can find an approximate analytic solution to this integral if we have an estimate of the NEP. The idea is that $S_P(\omega)$ can be approximated as

$$S_P(\omega) \approx NEP(0)^2 (1 + i\omega\tau_-). \quad (2.87)$$

Using this form, we can analytically solve Eq. (2.86), which gives us an expected baseline energy resolution to Dirac delta deposits of

$$\sigma_E \approx NEP(0) \sqrt{\tau_-}, \quad (2.88)$$

and an estimated baseline energy resolution of $\sigma_E \approx 47.8 \text{ meV}$. Comparing to the value obtained in Table 2.2, this is quite close, which allows us to achieve a simple way of estimating the expected baseline energy resolution without needing to carry out numerical integration.

2.3.4 Parameterizing Excess Noise

In Fig. 2.14, we noted the existence of excess noise. It can be useful to parameterize this noise into two different categories: excess TFN-like noise and excess Johnson-like noise. To describe excess TFN-like noise, we define the parameter ξ as related to the previously derived S_{PTFN} in Eq. (2.72) by

$$S_{PTFN}^* = (1 + \xi)S_{PTFN}, \quad (2.89)$$

where we have added an asterisk to denote the inclusion of an excess noise factor. We assume ξ is frequency-independent, such that this same relation holds in current space. Similarly, the excess Johnson-like noise can be described using the M factor [118], defined in relation to the modeled Johnson voltage noise by

$$S_{V_{TES}}^* = S_{V_{TES}}(1 + M^2), \quad (2.90)$$

where again have added an asterisk. Assuming that M is frequency-independent, this same relation holds for the TES Johnson noise in current space. In order to actually determine ξ and M , the predicted TES noise must be compared to the measured noise. One method is to carry out a least squares fit to the data using the total modeled noise including the excess, defined in current space as

$$S_{I_{tot}}^*(\omega) = S_{I_{load}}(\omega) + S_{I_{TES}}(\omega)(1 + M^2) + S_{I_{TFN}}(\omega)(1 + \xi) + S_{I_{SQUID}}(\omega). \quad (2.91)$$

For the TES we have been studying in this section, we can use the current noise in Fig. 2.14 and parameterize the excess noise with ξ and M . Fitting Eq. (2.91) with ξ and M varying to the data, we find that $\xi = 1.15$ and $M = 0.71$ provide a good parameterization of the excess noise, the result of which is shown in Fig. 2.16. There does still remain a small amount of excess current noise at around 10 kHz that this parameterization was not able to fully explain, suggesting that there may be another source of excess noise which does not behave as TFN or Johnson noise.

2.4 QETs and Athermal Phonon Collection

Now that we understand the basic concepts of a TES, we can further extend this technology to the QET (Quasiparticle-trap-assisted Electrothermal-feedback Transition-edge-sensor) [96]. A QET is a TES that has been fabricated with a small region that overlaps with superconducting fin structures (typically Al). These fins will increase the total sensor surface area without increasing the sensor heat capacity.

With these fin structures, athermal excitations will interact with the superconducting Al fins and break Cooper pairs, creating quasiparticles. These quasiparticles will randomly travel throughout the Al fin, eventually reaching and becoming trapped in the overlap region, which will have a suppressed T_c due to the W. These quasiparticles then travel into the W




Figure 2.16: Parameterized excess noise using the ξ and M parameters for the $200\ \mu\text{m} \times 50\ \mu\text{m} \times 40\ \text{nm}$ W TES in current space.

TES due to the spatial T_c gradient, as depicted in Fig. 2.17 with a cross section of a typical QET shown. Because of the sharp T_c variation, the majority of the potential energy of the Al quasiparticles will convert into kinetic energy and potential energy of the W quasiparticles, which then thermalize in the TES. The temperature of the TES will increase and then respond via the equations we derived in Section 2.3.1.

These athermal excitations come in the form of athermal phonons that are the product of some particle interaction in our substrate. The particle interaction will create electron-hole pairs and optical phonons, with the electron-hole pairs recombining into optical phonons. These optical phonons very quickly (on the order of fs to ps) downconvert to acoustic phonons, which then travel ballistically throughout the substrate that our QETs have been deposited on. As they travel throughout the substrate, they will eventually interact with a superconducting Al fin, breaking Cooper pairs, and starting the process discussed in the previous paragraph. In this process, there are efficiency losses at each interface (e.g. when phonons reflect off surfaces, when they enter into the Al fins, when quasiparticles travel between different regions of the QET), as well as efficiency losses due to downconversion of the phonons to energies below the Al superconducting bandgap, but the increased sensor coverage is worth these losses (which in total are on the order of $\varepsilon_{ph} \approx 10\text{--}20\%$). There is also a characteristic time constant τ_{ph} for collecting these athermal phonons, and we are no longer in the Dirac delta deposit regime, as discussed in Section 2.3.2.

Thus, we must change our expected baseline energy resolution for collecting athermal




Figure 2.17: (Top) The spatial T_c gradient of the superconducting bandgap in a QET. (Bottom) A cross section of a typical QET, showing quasiparticles being collected in the the W TES.

phonons. Instead of the Dirac delta deposit, we will have a single-exponential (assuming one phonon collection time), which in Fourier space has the form

$$p(\omega) = \frac{\varepsilon_{ph}}{1 + i\omega\tau_{ph}}. \quad (2.92)$$

The expected baseline energy resolution equation from Eq. (2.86) becomes

$$\sigma_E = \left[\int_0^\infty \frac{d\omega}{2\pi} \frac{4|p(\omega)|^2}{S_P(\omega)} \right]^{-1/2}. \quad (2.93)$$

Taking the same approximation that $S_p(\omega) \approx NEP(0)(1 + i\omega\tau_-)$, this equation can be analytically solved, giving

$$\sigma_E \approx \frac{1}{\varepsilon_{ph}} NEP(0) \sqrt{\tau_{ph} + \tau_-}, \quad (2.94)$$

which is quite similar to Eq. (2.88), but with τ_{ph} in the square root. In the case of the QET deposited on a substrate, this tells us that we must understand the phonon collection time τ_{ph} and the collection efficiency ε_{ph} when calculating baseline energy resolutions.

In the next chapter, we will apply all of these concepts to a QET-based detector in order to fully characterize it.

Chapter 3

Performance of a Cryogenic PhotoDetector

In this chapter, I will present the motivation, design, and characterization of a large-area Cryogenic PhotoDetector (CPD). The analysis herein builds upon the TES characterization concepts discussed in Chapter 2, applying them to a fully-realized detector concept for active particle identification in rare event searches. An abbreviated version of this chapter was originally published in *Applied Physics Letters* as Ref. [115].

3.1 Motivation of a Large Area Photon Detector

In rare event searches, experimental sensitivity is often limited by background signals [119–128]. Developing precision detectors to veto background and noise signals has been a high priority in these fields.

3.1.1 Neutrinoless Double Beta Decay Experiments

Much interest in low temperature cryogenic detector technology has been shown by groups carrying out searches for neutrinoless double beta decay [129] ($0\nu\beta\beta$), such as the CUORE [119, 120], CUPID [130], and AMoRE [131] experiments. In these low-temperature calorimeters, the dominant source of background events consists of α decays from surface contaminants on the crystal and the surrounding environment [119, 120, 132]. In Fig. 3.1, we show a recent result from CUORE which is indeed limited by these background signals [120].

It has been shown that Cherenkov emission or scintillation light can be used to positively identify the signal β s, allowing for background discrimination [133]. In order for these experiments to achieve a high level of rejection for these α backgrounds, photon detectors with large surface areas and baseline energy resolutions below 20 eV (RMS) for Cherenkov signals [134], or of $\mathcal{O}(100)$ eV for scintillation signals [130], are required.




Figure 3.1: (Figure from Ref. [120]) Limits from 2020 on $0\nu\beta\beta$ decay set by CUORE, showing the spectrum used for setting the limit. The dominant background is from α decays of the surrounding experimental setup.

To reject the pileup background from multiple ordinary (two neutrino) double beta decay ($2\nu\beta\beta$) events, experiments need timing resolutions down to $10\ \mu\text{s}$ (for the ^{100}Mo isotope) [130].

3.1.2 Other Applications

There is also the theoretical and experimental motivation to search for dark matter in the mass range of keV/c^2 to GeV/c^2 [31, 32, 53, 135]. However, current experiments have been limited by unknown background signals in the energy range of $\mathcal{O}(1\text{--}100)\ \text{eV}$ [122–128, 136]. If the source of such backgrounds are high energy photons that deposit only an extremely small fraction of their energy in the target [137], then a nearly 4π active shield composed of high- Z scintillating crystals read out by these large area photon detectors could be highly efficient at suppressing these backgrounds. Additionally, a sensitive large area cryogenic detector could be useful for discriminating small energy depositions due to radiogenic surface backgrounds.

Other potential DM applications for this detector technology include searches for inelastic electronic recoils off scintillating crystals [138, 139] and searches for interactions with superfluid ^4He [140], shown schematically in Figs. 3.2 and 3.3, respectively.




Figure 3.2: (Figure from Ref. [138]) A diagram of an DM-electron recoil with a scintillating crystal. The photodetectors would surround the scintillation material and read out the energy of emitted photon from the recoil.




Figure 3.3: (Figure from Ref. [140]): A diagram of the various byproducts of a recoil with a superfluid ^4He (light blue). The photodetectors in this concept at the red rectangles, which could read out the various signals associated with some DM recoil.




Figure 3.4: A picture of the CPD installed in a copper housing. The instrumented side is shown facing up.

3.2 QET-based Large Area Photon Detector

The (100)-oriented substrate of the CPD is a 10.6 g Si wafer of thickness 1 mm and a surface area of 45.6 cm². A parallel network of 1031 QETs [96, 97] with $T_c = 41.5$ mK was deposited on one side of the wafer. The QETs are uniformly distributed over the wafer's surface and connected to a single readout channel. The uniform and distributed nature of the channel allows for the fast collection of athermal phonons with minimal positional dependence, reducing efficiency penalties from effects such as athermal phonon down-conversion [141, 142]. The opposite side of the Si wafer is unpolished and noninstrumented. The detector and QET mask design can be seen in Figs. 3.4 and 3.5. In Table 3.1, the QET design specifications for the CPD are listed.

3.3 Experimental Setup

The detector was studied mostly by our Berkeley group at the SLAC National Accelerator Laboratory in a cryogen-free dilution refrigerator at a bath temperature (T_B) of 8 mK. The detector was placed in a copper housing and was held mechanically with the use of six cirlex clamps. The cirlex clamps also provided the thermal link between the detector and the copper housing. The QET arrays were voltage biased and the current through the TES was




Figure 3.5: Left: Mask design of the CPD. Right: The design of the QETs used for the detector. (Blue: Al fins, Purple: W TES.)

Table 3.1: QET design specifications for the CPD describing the W TESs and the Al fins that each QET consists of. The active surface area refers to the amount of substrate that is covered by the Al fins of the QETs, while the passive surface area is that which is not covered by the Al fins, but by the Al bias rails, bonding pads, and other structures that absorb athermal phonons, but do not add to the signal.

Specification	Value
TES Length [μm]	140
TES Thickness [nm]	40
TES Width [μm]	3.5
Number of Al Fins	6
Al Fin Length [μm]	200
Al Fin Thickness [nm]	600
Al-W Overlap [μm]	10
Number of QETs	1031
Active Surface Area [%]	1.9
Passive Surface Area [%]	0.2




Figure 3.6: Design of collimator used in calibrating the CPD. Note that the diagram is not to scale. We have also included the various dimensions of the design, where horizontal measurements refer to radii.

measured with a DC superconducting quantum interference device (SQUID) array with a measured noise floor of $\sim 4 \text{ pA}/\sqrt{\text{Hz}}$, similar in design to the one described in Ref. [143].

A collimated ^{55}Fe X-ray source was placed inside the cryostat and was incident upon the noninstrumented side of the CPD in the center of the detector. A layer of Al foil was placed inside the collimator to provide a calibration line from fluorescence at 1.5 keV [144, 145]. The collimator was tuned such that there was $\sim 5 \text{ Hz}$ of the ^{55}Fe K_α and K_β decays incident on the detector. A diagram of the collimator is shown in Fig. 3.6. The detector was held at a bath temperature $T_B \ll T_c$ for approximately two weeks to allow any parasitic heat added by the cirlex clamps to dissipate. During this time, we attempted to neutralize potential charged impurities within the Si wafer as much as possible with ionization produced by a $9.13 \mu\text{Ci}$ ^{137}Cs source placed outside of the cryostat.

3.4 Characterization of the CPD

Following the characterization steps that have been outlined in Chapter 2, we can understand the DC characteristics, TES response, and noise of this device. Furthermore, we will discuss the phonon pulse shape for this athermal phonon sensor.




Figure 3.7: TES current as a function of voltage bias for the CPD at base temperature.

3.4.1 DC Characteristics and TES Response

To characterize the QETs, IV sweeps were taken at various bath temperatures by measuring TES quiescent current as a function of bias current, with superimposed small square pulses providing complex admittance [97] at each point in the IV curve [98, 99, 114]. Since all the QETs are connected in parallel in a single channel, the channel was treated as if it were a single QET, describing the average characteristics of the total array. The IV data allowed for the estimation of the parasitic resistance in the TES line (R_p), the normal state resistance (R_N), and the nominal bias power (P_0). In Figs. 3.7–3.9, we show the TES current, resistance, and bias power as a function of voltage bias ($V_b = I_b R_{sh}$) at base temperature. The effective thermal conductance between the QETs to the Si wafer (G_{TA}) and T_c were measured by fitting a power law to the measured bias power as a function of bath temperature [98], also discussed further in Appendix A. This measurement is a lower bound of both these values, as it assumes no parasitic bias power in the system. If there were parasitic power, then this would cause the measured P_0 curve as a function of temperature to be systematically shifted down as compared to the intrinsic curve, such that the measured values for G_{TA} and T_c are lower (see Appendix A for examples of these curves). We summarize the DC characteristics of the detector in Table 3.2.

The complex admittance data allows us to estimate the dynamic properties of the sensors. Throughout the superconducting transition, primary and secondary thermal fall times were observed, e.g. $58 \mu\text{s}$ and $370 \mu\text{s}$, respectively, at $R_0 \approx 35\% R_N$. The origin of this addi-




Figure 3.8: TES resistance as a function of voltage bias for the CPD at base temperature.




Figure 3.9: TES bias power as a function of voltage bias for the CPD at base temperature.

Table 3.2: DC characteristics of the TES from IV curves, where a specific P_0 is reported at the bias point given by R_0 . The systematic errors on G_{TA} and T_c represent the upper bounds on these values, using the hypothesis that the observed excess noise in the sensor bandwidth is entirely due to parasitic bias power.

Parameter	Value
R_{sh} [m Ω]	5 ± 0.5
R_p [m Ω]	8.7 ± 0.8
R_N [m Ω]	88 ± 10
G_{TA} [nW/K]	0.48 ± 0.04 (stat.) $^{+0.49}_{-0.00}$ (syst.)
T_c [mK]	41.5 ± 1.0 (stat.) $^{+10}_{-0}$ (syst.)
R_0 [m Ω]	31 ± 3
P_0 [pW]	3.85 ± 0.45

Table 3.3: TES response parameters from the complex impedance data for the bias point of $R_0 \approx 35\% R_N$.

Parameter	Value
τ_0 [μ s]	1700 ± 200
L [nH]	190 ± 10
β	1.1 ± 0.1
$\mathcal{L}$	80 ± 15

tional time constant could be due to a more complex thermal or electrical system, e.g. phase separation [99, 146] or an extra heat capacity connected to the TES heat capacity [101]. The electronic rise time of the TES at $R_0 \approx 35\% R_N$ was also measured to be 3μ s. Characteristic plots of the complex impedance of the TES circuit in time domain and frequency domain can be seen in Figs. 3.10 and 3.11, respectively. Note the excellent agreement between the fits and the model, showing that we are modeling the TES response well. We have summarized the TES parameters at this bias point in Table 3.3, which will be used for noise modeling.

3.4.2 Noise Modeling

Knowledge of the TES parameters, given in Tables 3.2 and 3.3, allowed for the calculation of the power-to-current responsivity, which was used to convert the measured current-referred power spectral density (PSD) to the noise equivalent power (NEP). These parameters were used to predict the expected noise spectrum using the single-heat-capacity thermal model. A comparison of the NEP to the model at $R_0 \approx 35\% R_N$ can be seen in Fig 3.12. The excess




Figure 3.10: The TES response of the CPD to a square wave jitter in time domain (black). The model for the fit (red) includes the extra time constant beyond the usual TES thermal time constant.




Figure 3.11: The magnitude and phase of the measured complex impedance are shown as the black and blue markers, respectively. The modeled complex impedance is shown as the cyan solid line. The black dotted line denotes the corresponding bandwidth of 2.7 kHz for the thermal time constant $\tau_+ = 58 \mu\text{s}$.

noise spikes above approximately 500 Hz have been experimentally confirmed to be largely caused by vibrations from the operation of the pulse tube cryocooler. The observed noise is also elevated above our model at frequencies in the effective sensor bandwidth interval (approximately the inverse of the thermal time constant τ_- [97]) by a factor of ~ 2 , as compared to the prediction. This “in-band” excess noise is consistent with two different hypotheses: a white power noise spectrum incident on the detector of $8 \times 10^{-18} \text{ W}/\sqrt{\text{Hz}}$ (e.g. a light leak) or a parasitic DC power in the bias circuit of approximately 6 pW. If we assume the latter is the source, this allows us to calculate the upper bounds on our estimates of G_{TA} and T_c , as reported in Table 3.2. There remains bias-dependent excess noise above the sensor bandwidth. We parameterize the excess TES Johnson-like noise with the M factor, as discussed in Section 2.3.4. Using values of M up to 1.8, depending on bias point, can account for the discrepancy between observation and prediction at these frequencies. We note that this excess noise could possibly also be explained with a more complex thermal model.

The lowest integrated NEP was achieved at an optimum bias point of $R_0 = 31 \text{ m}\Omega \approx 35\% R_N$. In addition to the characterization data, approximately 500,000 threshold triggered events and 80,000 randomly triggered events were recorded at this bias. These data will be used to understand the phonon pulse shape of events, energy calibration, and baseline energy resolution.

3.4.3 Phonon Pulse Shape

For the measured phonon-pulse shape, there are multiple characteristic time constants. The pulse rise time was measured as $\tau_{ph} = 20 \mu\text{s}$, which is the expected characteristic time scale for athermal phonons being absorbed by the Al collection fins of the QETs for this design. The dominant pulse fall time is consistent with the expectation from the complex impedance as we approach zero-energy, where we confirmed the expected thermal time constant $\tau_- = 58 \mu\text{s}$ via a fit of the rise and fall times of the pulses. The secondary time constant from the complex impedance of $370 \mu\text{s}$ was also seen in these low-energy pulses. The secondary time constant from the complex impedance of $370 \mu\text{s}$ was also seen in these low-energy pulses, with an amplitude ratio of less than 2% to the dominant decay exponential.

We observed an additional long-lived behavior in the pulses, which can be estimated as a low-amplitude $\sim 3 \text{ ms}$ exponential tail whose magnitude scales linearly with the event energy. This time constant is not seen in the complex impedance data, which implies that it is not due to intrinsic TES response. Due to phonon down-conversion, there will be a phonons with energies less than the superconducting Al bandgap. In this case, these phonons still have a chance of being directly absorbed by the W TES (as opposed to being caught by the Al fins). Because the W TES coverage is significantly smaller than the Al coverage, we would expect the corresponding time constant to be much slower. The ratio of Al coverage to W coverage is about a factor of 340, while the ratio of Al coverage to W coverage, including the Al/W overlap, is about a factor of 33. Estimating the subgap phonon collection time as this factor times the collection time of the above-gap phonons $\tau_{ph} = 20 \mu\text{s}$ (the phonon




Figure 3.12: Modeled noise components: TES Johnson noise (blue solid), load resistor Johnson noise (blue dots), electronics noise (purple dashed), thermal fluctuation noise (TFN) between the TES and the bath (yellow solid), and total modeled noise (green solid) compared with the measured NEP (black solid) for $R_0 \approx 35\%R_N$. We additionally show the total noise model (green alternating dashes and dots), which includes a hypothetical environmental noise source of $8 \times 10^{-18} \text{ W}/\sqrt{\text{Hz}}$ and excess TES Johnson noise with $M = 1.8$. The light-purple line in upper portion of the figure denotes the power-pulse shape (arbitrarily scaled), which consists of a single pole at the observed rise time of $1/(2\pi\tau_{\text{ph}}) = 8 \text{ kHz}$.

collection time is a geometric quantity that is inversely proportional to sensor coverage, as shown, e.g., in C. W. Fink’s thesis [147]), we find that the subgap phonon collection time could be in the range of 0.9–6.9 ms, depending on how much of the Al/W overlap is used, where 3 ms sits firmly in this range. Thus, this long-lived behavior might be due to direct absorption of phonons with energy smaller than the Al superconducting band gap into the TES, as previously suggested by K. D. Irwin *et al.* in Ref. [96].

For energies above 300 eV, we observed a local saturation effect that manifests as the dominant fall time lengthening with increased energy. In Fig. 3.13, we show averaged pulses for various event amplitudes, showing the dependence of the pulse fall time on energy. We associate this effect with high-energy, single-particle events pushing nearby QETs into the normal resistance regime, slowing down the response of the total single-channel device. We also note that there is a position-dependent effect for a subset of high-energy events, notable by a varying fall time for events with the same amplitude. Our hypothesis for this phe-




Figure 3.13: We show averaged pulse shapes (green solid) normalized by the fitted OF amplitude (which explains the small amount of dispersion of the maximum values at the peak current), for which the shade of green lightens with increased OF amplitude. For reference, $0.1 \mu\text{A}$ corresponds to about 0.1 keV , $1.1 \mu\text{A}$ corresponds to about 1.5 keV , and $2.0 \mu\text{A}$ corresponds to about 3.4 keV . Each averaged pulse consists of about 100 events averaged in $0.04 \mu\text{A}$ bin-widths. The lengthened fall time of the averaged pulse with increased OF amplitude (an energy estimator) is evident. The phonon-pulse template used in this analysis (black dashed) shows good agreement with the low energy (dark green) pulses. We also show an analytic phonon-pulse with only the first sensor fall time (gray dashed). Comparing to the phonon-pulse template, we see that the second sensor fall time has a small effect in this limited time interval.

nomenon is that events close to the edge of the detector have less solid angle to deposit the energy, which leads to longer recovery times as opposed to events in the center of the detector (e.g. the calibration events). These effects are specific to the single-particle nature of the measured events. For scintillation events, the isotropic nature of the photons would spread out the event energy across the entire detector channel, avoiding these local saturation and position-dependent effects.

3.5 Energy Calibration and Resolution

To reconstruct event energies, two energy estimators were used in this analysis: the optimal filter (OF) amplitude (see Appendix C and Refs. [117, 148, 149]) and the energy removed by electrothermal feedback (E_{ETF}) [97]. For the OF, we used an offline algorithm to reconstruct energies. A single noise spectrum was used, which was computed from the randomly triggered events. The phonon-pulse template used was an analytic template that matches the measured low-energy pulse shape, neglecting the 3 ms low-amplitude tail. Because we could not directly measure the low-energy phonon-pulse shape with high statistics, we used a template without the long-lived behavior.

The integral estimator E_{ETF} was calculated for each triggered event by measuring the decrease in Joule heating via

$$E_{\text{ETF}} = \int_0^{T_{\text{trunc}}} [(V_b - 2I_0 R_\ell) \Delta I(t) - \Delta I(t)^2 R_\ell] dt, \quad (3.1)$$

where T_{trunc} is the time at which the integral is truncated, $\Delta I(t)$ is the baseline-subtracted pulse in current, I_0 is the quiescent current through the TES, R_ℓ is the load resistance, and V_b is the voltage bias of the TES circuit, as discussed at the end of Section 2.2.2.

In comparison to the OF amplitude, this integral estimator was less sensitive to saturation effects, but had a worse baseline energy resolution. When characterizing this device, we used the integral truncation of $T_{\text{trunc}} \approx 7\tau_-$ for E_{ETF} . This was done to preserve good baseline energy sensitivity in this integral estimator when calibrating the OF amplitude energy estimator at low energies.

For pulse-shape saturation at high energies, we use the following empirical model:

$$E_{\text{ETF}} = a \left(1 - \exp \left(-\frac{E_{\text{true}}}{b} \right) \right). \quad (3.2)$$

This functional form has the expected behavior: it intercepts zero, approaches an asymptotic value at high energies, and becomes linear for small values of E_{true} . In Fig. 3.14, the fitted saturation model, as well as the calibrated and uncalibrated E_{ETF} spectra, are shown, as compared to the energies of various spectral peaks in both energy scales. For the event spectra, we observed an unknown background at low energies, which we will discuss in detail in Chapter 5.

The absolute phonon collection efficiency (ε_{ph}) of the detector was estimated by measuring E_{ETF} at the lowest energy calibration line (Al fluorescence) and dividing by the known energy of that line. Because of the long-lived behavior in the phonon-pulse shapes, the measured collection efficiency of this detector depends on the integration truncation time T_{trunc} . If it is chosen to only include energy collected over a few times the first sensor fall time τ_- (e.g. $T_{\text{trunc}} \approx 7\tau_-$), then we find that $\varepsilon_{ph} = 13 \pm 1\%$. Alternatively, if we integrate to effectively infinity, this includes the low-amplitude long-lived behavior of the phonon pulses. In this case, the collection efficiency increases to $\varepsilon_{ph}^\infty = 17 \pm 1\%$, which implies that about 30% of




Figure 3.14: Upper: The calibrated E_{ETF} (which estimates E_{true}) spectrum for the CPD (solid black) grouped in bins of width 70 eV. Right: The energy spectrum in E_{ETF} (solid black) grouped in bins of width 7.4 eV. Lower left: The fitted saturation model using Eq. (3.2) (solid black). In each of these panels, we have shown, for both the calibrated and uncalibrated E_{ETF} energy scales, the location of the K_{α} , K_{β} , and Al fluorescence calibration peaks (pink dashed, blue dotted, and cyan alternating dashes and dots, respectively). In the lower left panel, the intersections of the lines corresponding to each spectral peak represent the points used for calibration of E_{ETF} via Eq. (3.2). The unmarked peaks at 4.2 keV and 4.8 keV in calibrated E_{ETF} are the Si escape peaks [150].




Figure 3.15: How the collection efficiency of the Al line changes with amount of integration of E_{ETF} .

the collected energy for a given event is associated with the low-amplitude tail of the phonon-pulse shape (about 8% and 22% from the $370\,\mu\text{s}$ and $3\,\text{ms}$ components, respectively). This is shown in Fig. 3.15, where we are plotting the collection efficiency calculated at the Al fluorescence line as a function of number of fall times τ_- integrated. Without the long-lived behavior of the phonon pulses, we would have expected the collection efficiency to asymptote at 13%. Because we see the collection efficiency continue to increase long after the first sensor fall time, this shows the importance of the long-lived behavior in determining the amount of energy collected by the TES for an event.

To calibrate the OF amplitude to units of energy, we fit the relationship between the calibrated E_{ETF} and the OF amplitude to a linear slope at low energies (below approximately 300 eV), as shown in Fig. 3.16. The value of the slope of the red line in the plot gives a calibration constant of $1.045 \pm 0.009\,\text{keV}/\mu\text{A}$. This method does not provide a calibration of the OF amplitude at high energies, but allows for the calculation of the baseline energy resolution.

For the calibration method used, the main source of systematic error is the saturation model in Eq. (3.2). Since it is empirical, its use introduces uncertainty in its applicability. We can estimate a rough upper bound of the effect of this systematic on the baseline energy resolution as the value that would be reached if we instead calibrated E_{ETF} linearly using




Figure 3.16: Correction of OF amplitude to energy, using linear region. Note that the variation of the black dashed lines that define the fit region result in negligible change in the calibration.

the Al fluorescence line. That is, the formula for calibration in this case would be

$$E_{\text{ETF}} = a \cdot E_{\text{true}}|_{\text{Al line}}. \quad (3.3)$$

In this case, this worsens the baseline energy resolution, as we are not taking into account the expected response that is seen in Fig. 3.14, and we know that there is saturation from our observations (however, we will use this as a conservative upper bound nonetheless). The full calibrated E_{ETF} spectrum would then become what is shown in Fig. 3.17, where the true locations of the ^{55}Fe K_α and K_β lines are clearly no longer reconstructed well. At low energies, this also has the effect of increasing the conversion factor from OF amplitude to units of energy to 1.098 ± 0.009 keV/ μA . Because the baseline energy resolution is simply the OF amplitude baseline energy resolution in units of Amperes multiplied by this conversion factor, this shows that this calibration method returns a worsened baseline energy resolution when not treating saturation effects with our model in Eq. (3.2). This worsened baseline energy resolution will henceforth be reported as the upper bound of the systematic error on this value.

The baseline energy resolution was calculated as the RMS of 46,000 randomly triggered events, after removing data contaminated by pileup events, electronic glitches, or thermal tails. This gave a resolution of $\sigma_E = 3.86 \pm 0.04$ (stat.) $^{+0.19}_{-0.00}$ (syst.) eV (RMS) for the OF energy estimator, where these data are consistent with a normal distribution. This is in




Figure 3.17: Calibrated E_{ETF} energy spectrum if calibrating directly to the Al line via Eq. (3.3). This alternative calibration determines the systematic upper bound on the baseline energy resolution reported for the CPD.

agreement with our estimation from the observed NEP and the power-referred phonon-pulse shape (a single-exponential with fall time τ_{ph} and collection efficiency ε_{ph}), which gave an expected baseline energy resolution of $\sigma_E^{th} = 3.9 \pm 0.4$ eV (RMS), as discussed in Section 2.4 and was similarly done in Ref. [98].

Using the OF formalism, we can also calculate the expected timing resolution σ_t (see Section C.1.2) of the CPD, which provides an estimate of the minimum resolving time for two pileup events (i.e. we would expect to be able to determine that two pulses spaced a few σ_t apart to be detectable as two pulses). For a $5\sigma_E$ event, the corresponding timing resolution of this detector is $2.3 \mu\text{s}$, and this timing resolution improves for higher energies, as it is inversely proportional to energy. For many $0\nu\beta\beta$ experiments, the minimum resolving time requirement to make pileup of multiple $2\nu\beta\beta$ events a negligible background is on the order of 1 ms [151–154]. For the CUPID and CUPID-1T experiments, this requirement is about $300 \mu\text{s}$ and $10 \mu\text{s}$, respectively [130]. In the future, one should study the minimum resolving time using an OF-based pileup detection algorithm with a simulation based on the expected $2\nu\beta\beta$ spectrum for ^{100}Mo , where one could compare the χ^2 goodness-of-fit metric from the OF fits for one pulse and two pulses (i.e. calculate the $\Delta\chi^2$) and find the minimum separation time of two pulses for which this comparison breaks down.

Table 3.4: Comparison of this work to various state-of-the-art devices for degraded α rejection in $0\nu\beta\beta$ experiments. The table is sorted by decreasing $\frac{\sigma_E}{\sqrt{\text{Area}}}$, a common figure-of-merit of devices for this application, for which smaller is better. The column labeled “NTL?” denotes whether or not each detector relies on NTL amplification to achieve the corresponding result.

Device	Area [cm ²]	σ_E [eV]	$\frac{\sigma_E}{\sqrt{\text{Area}}}$ [$\frac{\text{eV}}{\text{cm}}$]	NTL?
MKID [157]	4.0	26	13	No
W-TES [158]	12.6	23	6.5	No
Ge-NTD [159]	15.6	20	5.1	No
Ge-NTD [160]	19.6	19	4.3	Yes
IrAu-TES [161]	4.0	7.8	3.9	Yes
Ge-NTD [162]	4.9	7.6	3.5	Yes
Ge-NTD [163]	15.2	10	2.6	Yes
Ge-NTD [164]	15.2	8	2.1	Yes
W-TES [165]	12.6	4.1	1.2	No
W-TES (this)	45.6	3.9	0.6	No

3.6 Discussion

When comparing the baseline energy resolution of the CPD to the requirements of the CUPID experiment, the value surpasses the requirement of less than 20 eV (RMS) by a factor of five. While the CPD is a TES-based detector, it has been shown that Microwave Kinetic Inductance Detectors (MKIDs) and Neutron-Transmutation-Doped (NTD) Ge detectors are also promising avenues for achieving the sub-20 eV baseline goal. In Table 3.4, we report this result alongside those of other detectors for this application. In comparison to the devices that have met or exceeded the requirement, the CPD does not require Neganov-Trofimov-Luke (NTL) amplification [155, 156] (which often results in excess dark counts due to current injection at large electric field) and has the best baseline energy sensitivity for its size.

The measured baseline energy resolution of 3.86 ± 0.04 (stat.) $^{+0.19}_{-0.00}$ (syst.) eV and the expected timing resolution of $2.3 \mu\text{s}$ (at $5\sigma_E$), combined with its large surface area, makes this detector an excellent candidate for background rejection in both $0\nu\beta\beta$ and DM experiments. Because of the energy sensitivity, this device can be used as a dark matter detector itself, as we have done in collaboration with SuperCDMS to set limits on spin-independent dark matter-nucleon interactions for sub-GeV/ c^2 dark matter particle masses [84]. Similarly, this 10-gram-scale device could be applied to coherent elastic neutrino-nucleus scattering experiments [166]. The performance of the CPD can be further optimized through adjustment of characteristics such as the Al-W overlap and overall Al coverage. From these considerations, we anticipate up to a factor of two improvement in baseline energy resolution for a future iteration of the CPD, which has been designed by S. Zuber and C. W. Fink, and is currently being tested (as of 2022).

Chapter 4

A Light Dark Matter Search Operated Above Ground

In this chapter, we use the CPD to carry out an exclusion analysis that set the most stringent (comparing to contemporary searches) dark matter-nucleon scattering cross-section limits achieved by a cryogenic detector for dark matter particle masses from 93 to 140 MeV/ c^2 , with a raw exposure of 9.9 g d acquired at an above-ground facility. This work illustrates the scientific potential of detectors with athermal phonon sensors with eV-scale energy resolution for future dark matter searches. An abbreviated version of this chapter was originally published in *Physical Review Letters* as Ref. [84].

4.1 Introduction

As detailed in Chapter 1, numerous observations have shown that the majority of the Universe is composed of dark matter [12, 167, 168], for which the WIMP [169] has long been a favored candidate for this dark matter. However, direct detection experiments have ruled out a significant portion of the most compelling WIMP parameter space [170–172], which has provided the motivation for both theoretical and experimental exploration of alternative DM models [31]. In this chapter, we will detail a search for light dark matter (LDM) with a mass in the keV/ c^2 to GeV/ c^2 range and coupling to Standard Model particles via a new force mediator, which provides a well-motivated alternative to the WIMP hypothesis [32, 53, 135]. While many LDM searches focus on DM-electron interactions [125–127, 173, 174], detectors with eV-scale energy thresholds can also be used to study LDM via DM-nucleon interactions.

Though the CPD was designed for active particle identification in rare event searches, we will use its excellent baseline energy resolution to probe LDM parameter space. As a combined effort of the SuperCDMS and CPD collaborations, a DM search was carried out with 9.9 g d of raw exposure from Sept. 9th to 10th 2018. The data were acquired at the SLAC National Accelerator Laboratory in a surface facility of ~ 100 m in elevation. In the

following sections, we will discuss data acquisition techniques, device performance, and the results of an exclusion analysis for spin-independent DM-nucleon interactions.

4.2 Experimental Setup and Data Acquisition

The experimental setup for the DM search is unchanged as compared to that described in Chapter 3, as the DM search dataset was taken as part of the same analysis run as the characterization of the detector. Because of project time constraints and large cosmogenic backgrounds, the DM search was limited to 22 hrs. Data were acquired over this period using a field-programmable gate array (FPGA) triggering algorithm based on the optimal filter (OF) formalism (see Appendix C and Refs. [117, 148, 149]). Throughout the exposure, randomly triggered samples of the baseline noise were acquired (“in-run random triggers”), which allowed us to observe any changes in the noise over the course of the search and to calculate and monitor the baseline energy resolution.

Using the OF formalism (see Appendix C), the trigger threshold was set at $4.2\sigma_E$ above the normally-distributed baseline noise level (i.e. the baseline energy resolution of randomly triggered data using the OF with no delay), corresponding to 16.3 eV after calibration. Any pulse with an OF amplitude above this threshold was tagged as an event and saved. The phonon-pulse template used for the FPGA triggering algorithm was a double-exponential pulse with a rise time of $\tau_r = 20\mu\text{s}$ and a fall time of $\tau_f = 58\mu\text{s}$. The rise time was taken from the expected collection time of athermal phonons, and the fall time was taken from the thermal response time of the QET estimated from a measurement of the complex admittance [97]. Each of these time constants was confirmed by a nonlinear least squares fit to nonsaturated pulses. Although Chapter 3 discusses the existence of extra fall times, their effect on the OF amplitude measured for each event is negligible. Before starting the DM search, a separate, small subset of random triggers was collected. After removing data contaminated by effects such as elevated baselines and phonon pulses, the noise spectrum used by the FPGA algorithm was generated from these random triggers.

For overlapping triggered pulse traces, the triggering algorithm was set to save a trace centered on the pulse with the largest OF amplitude. We note that the FPGA triggering algorithm acted on a trace that was downsampled by a factor of 16, from the digitization rate of 625 kHz to 39 kHz. Additionally, the FPGA triggering algorithm considered only 26.2 ms of the total 52 ms long time trace saved for each triggered pulse trace (“event”). Because of these factors, the energy resolution of the FPGA triggering algorithm is not as good as can be achieved by reconstructing event energies using an offline OF, as described in the following sections.




Figure 4.1: A zoomed-in portion of an example event within the analysis ROI. The raw pulse (gray) is compared to the offline optimal filter result (blue), the pulse template scaled by the fit result (black dashed), the FPGA filter result (red with dots), and the FPGA trigger threshold (black dotted). The noticeable spacing between the FPGA optimal filter data points is due to its downsampling by a factor of 16 as compared to the offline optimal filter. The offline and FPGA optimal filters are highly correlated, but not exactly the same, with corresponding energy estimates for this event of 187 eV and 179 eV, respectively. The offset between the optimal filters and the raw pulse is by design of the filters, as they were set up to determine the time of the beginning of the pulse, as opposed to the maximum of the pulse.

4.3 Data Reconstruction and Calibration

While the FPGA-based OF was used to trigger the experiment in real time, we ultimately used our offline algorithm to reconstruct event energies also based on the OF formalism. For this offline OF, we were able to use a single noise spectrum computed from the in-run random triggers to represent the entire data set because there was negligible time variation of the noise over the course of the full exposure. Pulse amplitudes and start times were reconstructed using the same phonon-pulse template as in the FPGA triggering algorithm. Thus, there are two different pulse amplitudes for each event—one from the FPGA triggering algorithm and one from the offline OF. In Fig. 4.1, we compare the different energy estimators for a representative event.

This detector was optimized for maximum energy sensitivity at low energies and does




Figure 4.2: Measured energy spectrum in the DM-search ROI for the full exposure after application of the quality cuts. The data have been normalized to events per gram per day per eV and have been corrected for the event-selection efficiency, but not the trigger efficiency. The inset shows the calibrated E_{ETF} spectrum up to 7 keV, noting the locations of the different spectral peaks. The known values of the dashed lines are 1.5, 5.9, and 6.5 keV for the Al fluorescence (pink), ^{55}Fe K_α (blue), and ^{55}Fe K_β (cyan) lines, respectively. The two dotted gray lines between 4 and 5 keV in calibrated E_{ETF} are the Si escape peaks [150].

not have a large enough dynamic range to observe the calibration lines without nonlinear effects from saturation of the QETs. The nonlinearity is minimal within our region-of-interest (ROI), which is below 240 eV. Above the ROI, the fall time of the pulses increases monotonically with energy, which can be explained by effects of local saturation. Localized events can saturate nearby QETs to above the superconducting transition, while QETs far from the event stay within the superconducting transition. Because this is a single-channel device, the saturated and unsaturated QETs are read out in parallel and thus effectively combine into a single phonon pulse with an increased fall time. The calibration method as outlined in Section 3.5 is used to correct out these effects. In Fig. 4.2, we show the differential rate spectrum of the calibrated offline OF amplitude within the ROI, with the inset showing the differential rate spectrum for the calibrated E_{ETF} , which matches the spectrum shown in Fig. 3.14.

For the calibration of the FPGA trigger energy estimator, we know the conversion of the arbitrary FPGA units to Amperes. In this analysis, we had a 16-bit FPGA trigger, which meant that the largest FPGA amplitude that could be saved was 32767 in arbitrary FPGA




Figure 4.3: Calibration of the trigger amplitude to energy using the reconstructed energy. Data to the left of the dashed black line were not included in the calibration.

units. This corresponds to $0.47 \mu\text{A}$ after converting (about 500 eV), which is far below the Al fluorescence line. Any events that had an amplitude greater than $0.47 \mu\text{A}$ effectively have no FPGA trigger energy estimator associated with it. This means that we are unable to carry out an independent calibration of the FPGA trigger energy estimator, where we would apply the same method as the offline energy estimator. One possible method would be to apply the same calibration of the offline energy estimator E' to the FPGA trigger energy estimator E_T . However, there could be an error in the factors we are using to convert from the arbitrary units of the FPGA to Amperes or some other effect such that it is not guaranteed that we can use the same calibration as we did for the offline reconstructed energy.

Rather than using the same calibration as the reconstructed energy, an alternative calibration will be done where we calibrate the trigger amplitude to the reconstructed energy using a relative normalization, defined by $E_T = m \cdot A_T$. This normalization ensures that zero energy events have zero energy. As we do not know the true energy of the events in our spectrum, we will use the reconstructed energy instead as an estimate. This should be valid, as we expect that the offline reconstructed energy should have an average that is equal to the true energy. We will also not include the low energy events, as these events are at low enough energies where the energy estimators become significantly biased (where the nature of the bias may be slightly different in the two estimators) due to an OF being biased towards large noise fluctuations. For this fit, as shown in Fig. 4.3, we did not include any data points with a trigger amplitude below 50 nA . Note that that figure shows values below




Figure 4.4: FPGA trigger energy estimator spectrum in units keV. The large population of events at about 0.5 keV is due to saturation of the 16 bit FPGA trigger amplitude.

this cutoff, but they were not included in the calibration.

The resulting spectrum for the FPGA trigger amplitude is shown in Fig. 4.4. We can see that, in the figure, we only have FPGA trigger energy estimator information for event up to about 0.5 keV, which is above our ROI. At this point, we see a large population of data, which is all events with an FPGA trigger amplitude of 32767 (i.e. all of these events do not have a meaningful value for the FPGA trigger amplitude). This does not change anything for the DM search, as we only need FPGA trigger amplitudes for events within our ROI.

As previously noted, we did not include any data points with an FPGA trigger amplitude below 50 nA in its calibration. To show the effect of different lower energy cutoffs, we can plot the fitted slope as a function of cutoff, as done in Fig. 4.5. We see that our fitted value is within the statistical deviation of the slope for fits with cutoff above 25 nA in trigger amplitude. Below this point, the low energy fluctuations bias the fit, showing that we should not include them. The chosen cutoff of 50 nA (rather than 25 nA) ensures that we are not too close to where the bias begins. Because the error bars are consistent with the chosen value throughout, the effect of a different cutoff would be negligible. We also note that, since the plotted fit is within the statistical deviation of the fitted slope for all other cutoffs, this implies that the cutoff itself does not inject systematic error into our analysis. The statistical error in the fit ($\sim 0.1\%$ error) is very small, which tells us that this statistical deviation is negligible.




Figure 4.5: Dependence of the relative normalization of the trigger energy on the trigger amplitude cutoff. The red, dotted line is the fitted value with cutoff of 50 nA, and the black, dashed line shows the 25 nA cutoff. Note that the slope is very consistent above 25 nA, and then changes drastically below this point. This is due to the effect of the low energy events with a biased amplitude towards positive fluctuations.

4.4 Data Selection and Efficiency

We make our final event selection with a minimal number of selection criteria (“cuts”) to remove poorly reconstructed events without introducing energy dependence into the selection efficiency. This approach helps to reduce the complexity of the analysis and thus avoid introduction of systematic uncertainties. We apply two data-quality cuts: a prepulse baseline cut and a chi-square cut.

4.4.1 Baseline Cut

We define the event baseline as the average output in the prepulse section of each event, which is the first 25.6 ms of each trace. Large energy depositions have a long recovery time, which may manifest itself as a sloped baseline for subsequent events. Our trigger has reduced efficiency for any low-energy events occurring on such a baseline. We expected roughly 10% of the events to sit on the tail of a high energy event in part because of the high muon flux at the surface of $\sim 1 \text{ muon/cm}^2/\text{min}$ [175]. The baseline cut is performed by binning the data across the search in 400 s long bins and removing from each bin the 10% of events that




Figure 4.6: The data passing the baseline cut over time, to show the relatively constant baseline. The greyed out data are those that were deemed to have high baselines and were removed from the DM search dataset.

have the highest baseline, as shown in Fig. 4.6.

We see that there is a small amount of variation in the baseline over time, meaning that we should check if the TES bias point is changing appreciably. From Fig. 4.6, we can see that the largest change in current is approximately 30 nA over the course of the experiment. To estimate the corresponding change in resistance, we can estimate the bias voltage as constant with some change in current ΔI and related change in TES resistance ΔR

$$V_b = I_0(R_0 + R_\ell) = (I_0 + \Delta I)(R_0 + \Delta R + R_\ell). \quad (4.1)$$

Thus, rearranging, we have the the change in resistance given some change in current (i.e. baseline) is

$$\Delta R = \frac{-\Delta I}{I_0 + \Delta I}(R_0 + R_\ell). \quad (4.2)$$

Using the values in Table 3.2, a change in current at this bias point of $\Delta I = 30$ nA corresponds to a change in resistance of $\Delta R = 0.12$ m Ω . Since the operating resistance of the TES at




Figure 4.7: Using the in-run randomly triggered events, we have calculated the baseline energy resolution (RMS) of the reconstructed energy (i.e. calibrated OF amplitude) for two-hour sections of the data, which are shown in the plot as black error-bar markers (where the error is of the 1σ level). We have also plotted the energy resolution for the entire run as the black dashed line with 68% confidence level bands. Their agreement over time represents the stability of the system.

this bias point is $R_0 = 31$ mOhms, this is a less than 0.4% effect, and we can consider our baseline to be stable. We can further measure the baseline energy resolution throughout the DM search exposure, as we have done in Fig. 4.7. We see from above that the measured energy resolution for any two hour section is consistent with the value for the entire run, as they are all within 3σ of the value for the entire run. This is expected from our discussion of the stable baseline, where we argued that we should have a stable energy resolution. We are confident that our energy resolution is stable and represents the system well.

4.4.2 Low Frequency Chi-Square Cut

The chi-square cut is a general cut on our goodness-of-fit metric, for which we use the low-frequency chi-square χ_{LF}^2 calculated from the offline OF fit [117]. This metric is similar to the χ^2 from the offline OF fit, but we exclude frequencies over f_{cutoff} from the sum. That is, in the OF formalism, we have the following quantities: the Fourier transform of our signal $\tilde{v}(f)$, the Fourier transform of our pulse template $\tilde{s}(f)$, the fitted OF amplitude $\hat{A}$, and

the two-sided noise PSD $J(f)$. To calculate the chi-square for a given fit, as discussed in Appendix C, it is given by

$$\chi^2 = \int_{-\infty}^{\infty} df \frac{|\tilde{v}(f) - \hat{A}\tilde{s}(f)|^2}{J(f)}, \quad (4.3)$$

where we use the continuous Fourier transform for convenience (we convert to discrete when calculating these values on data). However, this equation includes high frequency values outside of our signal band. Because the mean of a chi-square distribution is n_{dof} and its variance is $2n_{dof}$, where n_{dof} is the number of degrees of freedom is the number of frequencies that we are integrating over, the inclusion of the frequencies above our signal band serve to reduce the sensitivity of the chi-square metric. Thus, to improve this sensitivity, we truncate the integral at some specified f_{cutoff} , giving

$$\chi_{\text{LF}}^2 = \int_{-f_{\text{cutoff}}}^{f_{\text{cutoff}}} df \frac{|\tilde{v}(f) - \hat{A}\tilde{s}(f)|^2}{J(f)}. \quad (4.4)$$

This truncation allows us to remove sensitivity to superfluous degrees of freedom outside of our signal band from the chi-square, thereby reducing both the expected mean and the expected variance of the chi-square distribution. In this analysis, we used $f_{\text{cutoff}} = 16 \text{ kHz}$ because the rise and fall times of our expected pulse shape correspond to frequencies of 8.0 kHz and 2.7 kHz respectively. The inclusion of this cutoff reduces n_{dof} from 32500 to 1663, improving the chi-square variance substantially (changing its standard deviation from 255 to 58). The pulse-shape variation within the DM-search ROI is minimal; this leads to a chi-square distribution that is largely independent of reconstructed event energy within this range. This in turn allows us to set an energy-independent cutoff value for this metric at $\chi_{\text{LF}}^2 = 2000$, as applied to the data in Fig. 4.8. We note that this cutoff value is nearly 6 standard deviations away from the mean, suggesting that this cut should have a high efficiency for accepting events.

Our measured events cannot be used directly to measure the signal efficiency of the chi-square cut because they include some that are not representative of the expected DM signal, e.g. vibrationally-induced events, electronic glitches, pileup events, etc. Therefore, we created a pulse simulation by adding noise from the in-run random triggers to the pulse template, systematically scaling the latter over the range of energies corresponding to the DM-search ROI. We then process and analyze the simulated data in the same way as the DM-search data. This simulation workflow has been included in the now-archived GitHub repository RQPY, co-created by C. W. Fink and myself [176]. For the chi-square cut, we use randomly triggered data that have passed the baseline cut, and have no (real) pulses in them that have an energy outside of the ROI. The passage fraction of this chi-square cut, which has an energy-independent value of $98.53 \pm 0.01\%$, represents the cut's efficiency, as shown in Fig. 4.9. While this value does not match the expectation of 6 standard deviations (which would be 99.99...%), this is not unexpected: the PSD we use to calculate the




Figure 4.8: The data passing the low frequency chi-square cut at $\chi_{\text{LF}}^2 = 2000$. The grey region represents the DM search ROI.

chi-square assumes that the noise is stationary (we know there is some time-variation in the noise, but it is small from Fig. 4.7) and uncorrelated (of which we cannot make a claim due to the single-channel nature of the device). Furthermore, the randomly triggered events can include real pulses, such that we could have both a simulated and a real pulse in the same event. This scenario would be excluded and is not included in the expected efficiency when assuming uncorrelated and stationary noise. Thus, it is reasonable to expect the observed deviation.

We can also compare the value of the simulated passage fraction due to the chi-square cut to our passage on in-run randoms, which are equivalent to zero-energy dark matter events. We would expect to see that the two are approximately equal, as our chi-square cut is seen to be energy independent in our region of interest. The passage fraction for our in-run randoms is $98.6 \pm 0.6\%$, showing that our result matches expectations.

As previously mentioned, the pulse-shape variation within the DM-search is minimal. To show that this possible source of systematic is negligible, we will show that inclusion of pulse-shape variation does not change the chi-square cut's efficiency. Within the ROI, there are two




Figure 4.9: The simulated passage fraction as a function of energy for the low frequency chi-square cut.

Table 4.1: The distribution of fall times and the estimated standard deviation due to (presumably) positional dependence at different pulse energies within the ROI.

Pulse Energy [eV]	τ_f [μ S]	$\sigma_{\text{positional}}$ [μ S]
16	58	1
80	62	1.5
240	63	5.8

possible sources of pulse-shape variation: local saturation and positional dependence. The main observable of local saturation is an increasing fall time of the pulses as energy increases. From Fig. 4.8, we see that chi-square values are roughly flat within the ROI, suggesting that we are likely unaffected by local saturation within this range. The main observable feature of positional dependence is, even with no local saturation, there being a variance of measured fall times that is greater than the expected statistical variance, regardless of energy. To estimate the statistical variance in a given population of the data, we average the variances for each data point taken from a nonlinear fit to the event fall times. Then, we estimate the total variance of the measured fall times from its distribution, which should be a value that is greater than or equal to the statistical variance. Given the single-channel nature of the CPD, we cannot prove its existence, but it is a reasonable explanation for variance of pulse shapes large than that expected from statistical variance (within CDMS, positional dependence has been frequently observed in similarly designed multi-channel devices). In order to extract the estimated variance due to (presumably) positional dependence, we subtract the statistical variance from the total variance

$$\sigma_{\text{positional}}^2 \approx \sigma_{\text{total}}^2 - \sigma_{\text{statistical}}^2. \quad (4.5)$$

In order to study the effect of the chi-square cut value, we have measured the distribution of the fall times for pulses of energies 16 eV, 80 eV, and 240 eV, where the various distributions of the fall times τ_f due to this assumption of positional dependence are reported in Table 4.1.

To study the effect of the pulse shape variation on the chi-square, we can use the pulse shapes at 240 eV for the three cases of pulse shapes, as this is where the effect should be most apparent. The three cases are: pulses have 58 μs fall times with no variation, pulses have 64 μs fall times with no variation, and pulses have 64 μs fall times with an estimated 5.8 μs standard deviation in fall time variation. The assumption that we used when measuring the chi-square passage fraction is the first case. With these three cases, we will see what the maximum effect of the different pulse shapes is. The expectation would be that the last case would have the largest effect, due to there being pulses with fall times far from the template's fall time of 58 μs . To test this idea, we have simulated data from recoil energies 0 to 240 eV for the three cases above, run the data through the same processing routines as the DM search, and calculated the passage fraction as a function of energy for the chi-square cut in each case. In Fig. 4.10, we show the chi-square cut efficiency for these three cases, showing that the pulse shape variation within the ROI is indeed minimal and can be neglected.

We do not apply any other cuts to the DM-search data. The total signal efficiency is thus 88.7% and is independent of energy. A variation of the cut values within reasonable bounds was found to have no significant impact on the experimental sensitivity.

4.5 Signal Model and Results

In our DM signal model for spin-independent nuclear-recoil interactions [83], we use the standard astrophysical parameters for the dark matter velocity distribution [80, 81, 177]:




Figure 4.10: The calculated passage fraction of the chi-square cut for different simulated datasets corresponding to different pulse shapes used, showing that the pulse shape variation does not have an effect on this efficiency.

a velocity of the Sun about the galactic center of $v_0 = 220$ km/s, a mean orbital velocity of the Earth of $v_E = 232$ km/s, a galactic escape velocity of $v_{\text{esc}} = 544$ km/s, and a local DM density of $\rho_0 = 0.3$ GeV/cm³. To relate this to the reconstructed energy, there are a multitude of possible techniques one might try. In the following paragraphs, we discuss the different possibilities and motivate the one used in this analysis.

4.5.1 Discussion of Flawed Reconstructed Signal Models

The first reconstructed signal model option would be a method where we simply fit the trigger efficiency in reconstructed energy and multiplied that curve by the DM differential rate model

$$\frac{\partial R}{\partial E'}(E') = \varepsilon(E') \frac{\partial R}{\partial E_0}(E'). \quad (4.6)$$

In this method, we are modelling the trigger efficiency as an error function, where this is the efficiency curve of the data after applying a cut on the trigger energy E_T . Note that we have written the differential rate as a function of the reconstructed energy. This is equivalent to assuming that the reconstructed energy has zero baseline energy resolution (RMS), which is a false statement. This immediately tells us that this is a flawed model, and that the nonzero baseline energy resolution must be correctly taken into account.

One possible method of taking into account the nonzero baseline energy resolution of the reconstructed energy E' is to convolve the signal model in true energy with a Gaussian that describes the resolution of our reconstructed energy

$$\frac{\partial R}{\partial E'}(E') = \varepsilon(E') \int_0^\infty dE_0 \mathcal{N}(E'|E_0, \sigma_{E'}) \frac{\partial R}{\partial E_0}(E_0). \quad (4.7)$$

However, this model is also flawed. This would only be applicable if we were setting our trigger threshold on E' (in which case $\varepsilon(E')$ would be a step function centered on the threshold). For the CPD dataset, we have set our threshold instead on the FPGA amplitude E_T , while using a separate energy estimator E' for our reconstructed energy. We must take into account the relation between these two (highly correlated) energy estimators when modeling the signal in reconstructed energy. We do note that if this model were used, then the Gaussian used would have to be truncated at some value of E' , as we have finite statistics in our sample. There is some range where we would not have enough statistics to be able to confidently say that we have a Gaussian to some level. In this case, we would have to decide on the extent of the smearing based on our confidence in the Gaussianity of the distribution far from the mean value.

The next model we consider takes into account the relationship of the two highly correlated energy estimators E' and E_T by modelling their relationship as a bivariate normal distribution (with the baseline energy resolutions of each estimator being related through some covariance matrix Σ) and convolving the signal with this distribution

$$\frac{\partial R}{\partial E'}(E') = \int_0^\infty dE_T \int_0^\infty dE_0 \Theta(E_T - \delta) \mathcal{N}(E', E_T|E_0, \Sigma) \frac{\partial R}{\partial E_0}(E_0). \quad (4.8)$$

In this signal model, we note the following assumptions that have been made: the trigger efficiency function the FPGA energy estimator E_T is a step function (as shown by the Heaviside step function, where the step occurs at the trigger threshold δ), the relationship between the energy estimators E' and E_T is described by a multivariate normal distribution, and the covariance matrix describing the multivariate normal distribution is independent of the true energy E_0 . However, the bivariate normal distribution does not describe the relationship between the two energy estimators near threshold. Because of the time-shifting aspects of both the FPGA trigger and the offline OF, the algorithms will latch onto the largest noise fluctuation in their search windows (as discussed in Appendix C, the largest OF amplitude corresponds to the smallest chi-square), which are 26.2144 ms and 160 μ s respectively. Due to latching onto the largest fluctuation, the distribution that relates the two energy estimators is non-normal below approximately 30 eV.

4.5.2 Reconstructed Signal Model for DM Search

To correctly take into account the trigger efficiency, we convolve the differential rate with the joint probability density function relating our two energy estimators, including the effects of the applied cuts. The signal model, which includes the estimated trigger efficiency, is

$$\frac{\partial R}{\partial E'}(E') = \int_0^\infty dE_T \int_0^\infty dE_0 \Theta(E_T - \delta) \varepsilon(E', E_T, E_0) P(E', E_T | E_0) \frac{\partial R}{\partial E_0}(E_0), \quad (4.9)$$

where E_0 is the true recoil energy, E' is the recoil energy measured by the offline OF, E_T is the recoil energy measured by the FPGA triggering algorithm, δ is the trigger threshold set on the FPGA triggering algorithm, ε is the efficiency of the two quality cuts and two cuts that are applied to simulated data (as described in the following paragraphs), Θ represents the trigger threshold cut (a Heaviside function), and $P(E', E_T | E_0)$ is the probability to extract E' and E_T using the two energy reconstruction algorithms given the true recoil energy E_0 . For the efficiency ε , we have generalized its form to be a function of energy, knowing that the baseline and chi-square cuts themselves are energy independent. The yield factor of heat signals (i.e. the ratio of heat signals produced by nuclear and electron recoils of the same energy that accounts for effects such as displacement damage) has been assumed to be unity for this work. Though measurements of the heat yield factor have not been made for Si, similar work has been undertaken for Ge, where the heat quenching factor was shown to be very close to unity [178, 179].

The model in Eq. (4.9) was evaluated numerically, taking advantage of our pulse simulation. The pulse simulation includes a software simulation of the FPGA triggering algorithm, which had the same output as the hardware version when run on the DM-search data. With this simulation of the FPGA triggering algorithm, we can use the pulse simulation to determine $P(E', E_T | E_0)$ directly. Low pulse height events may have their OF energy estimate affected by a shift of the start time estimate, but the simulation automatically takes this effect into account.

We also added two cuts to the simulated data only: a confidence ellipse cut and a trigger time cut. The confidence ellipse cut removed any events with an energy estimator value outside of the 99.7% confidence ellipse, which is defined by the covariance matrix of our two energy estimators for zero-energy events. This cut was implemented to exclude the possible scenario of calculating a finite sensitivity to zero-energy DM, which would be a nonphysical result. The trigger time cut removed events that were not within $29 \mu s$ —half of a fall time of a pulse—of the true event time, as determined by the energy-scaled pulse template. This cut ensured that the triggering algorithm was able to detect the signal added, as opposed to a large noise fluctuation elsewhere in the trace. These two cuts required knowledge of the true energy of the pulse—they cannot be applied to the data, but can be applied to the signal model—and helped to ensure that our signal modeling was conservative. In adding each of these cuts, we reduced our signal efficiency estimate, which necessarily biased the results in the conservative direction.

Carrying out the pulse simulation by injecting events to the in-run randoms with energies ranging from 0 to 40 eV in 0.4 eV steps, we can take the energy estimators, relate them to the known injected energies, and extract the probability density function. With our pulse simulation and cuts, we have shown three characteristic estimations of the probability density function in Eq. (4.9) in Figs. 4.11–4.13 with three cuts being applied to the simulated dataset (purple): trigger time cut (blue), 99.7% confidence ellipse cut (turquoise), and the trigger threshold cut (lime green). Each figure shows the probability density function projected to the reconstructed energy axis and to the FPGA trigger energy axis (for easier demonstration of the effects of the various cuts). In Fig. 4.11, we see that no events pass all of the cuts, as desired from the inclusion of the 99.7% confidence ellipse cut (i.e. that we do not have sensitivity to zero-energy events). In Fig. 4.12, we see that there is a nonzero population of events that have been reconstructed with energies above our trigger threshold. As these events all have a true energy of 10.4 eV, this demonstrates the concept of being sensitive to noise-boosted events, as well as the distribution being non-normal. In Fig. 4.13, the injected energy is now far above threshold, and the distributions now appear to be normal. In fact, the width of the distribution matches the baseline energy resolution in both energy estimators, as we would expect.

We can have further conceptual understanding of our trigger model by looking at the probability that an event with true energy E_0 is triggered (this can be thought of as the trigger efficiency in true energy), which we define as

$$P_{\text{trig}} = \int_0^\infty dE_T \int_0^\infty dE' \Theta(E_T - \delta) \varepsilon(E', E_T, E_0) P(E', E_T | E_0), \quad (4.10)$$

This equation is quite similar to Eq. (4.9), with the difference being that the true differential rate $\partial R / \partial E_0$ is replaced by 1, and we integrate over both of the energy estimators. The expectation would be that the shape is similar to an error function, but not exactly that same due to the non-Gaussianity of our joint probability density function at and below our trigger threshold.

At this point it is worth noting that plots of our probability density function were shown for just one simulation. In order to have an understanding of the variation of the trigger efficiency and our eventual limits, we have redone the simulations an additional seven times. Each simulation ensures that the random seed is uniquely set and the pulses are allowed to be injected over a 10 ms window in the center of each random trace (to ensure edge effects are negligible in reconstructing the energies). Thus, we carry out the same workflow as far as measuring the probability density function, as well as calculating the trigger efficiency. There are two ways to calculate the trigger efficiency for the simulated data: a direct measure of the passage fraction as a function of true energy and calculating it from our interpolated probability density. The expectation is that these values should agree well for all energies. For the direct measure of the passage fraction as a function of true energy (“measured”), this is the passage fraction of the 99.7% confidence ellipse cut, FPGA trigger time cut, and the FPGA trigger threshold cut. To calculate the uncertainty, we use the formalism described by Paterno in Ref. [180], such that we avoid nonphysical values for efficiencies near 0% or




Figure 4.11: Dependence of the reconstructed energy (a) and the FPGA trigger energy (b) on the true energy inputted, including cuts on simulated data. The data consist of simulated events all with injected energies of 0.4 eV, and no events pass all cuts (as shown by the lack of the lime-green line). We do not have sensitivity to events with 0.4 eV energies or less.




Figure 4.12: Dependence of the reconstructed energy (a) and the FPGA trigger energy (b) on the true energy inputted (10.4 eV), including cuts. The distributions are non-Gaussian, and the trigger threshold cut is a step function in trigger energy.




Figure 4.13: Dependence of the reconstructed energy (a) and the FPGA trigger energy (b) on the true energy inputted, including cuts on simulated data. The data consist of simulated events all with injected energies of 35.5 eV, and the distributions are now Gaussian, as we are far from threshold.

100%. For the calculation for our interpolated PDF (“calculated”), we again use Eq. (4.10). For each simulation and each calibration, we have estimated the trigger efficiency using both methods, as shown in Fig. 4.14.


We see that the measured and calculated trigger efficiency for each curve agree well with each other. More importantly, we see that there is some variation between the curves (both calculated and measured) for values of true energy below 5 eV. We also see that the calculated curve finds nonzero values of the trigger efficiency below the lowest nonzero value of the measured trigger efficiency. This is due to the binning between energies when interpolating between energies in the determination of the PDF. Thus, we need to add a cut on the true energy to ensure we are not allowing nonzero trigger efficiencies where it should be zero. In this case, we have added a cut on the values where, below this number (2.7 eV), the measured trigger efficiency is zero for at least one simulation. To get a more intuitive idea of the uncertainties at these low energies, we have shown the relative error in the measured trigger efficiency in Fig. 4.15.

We see that the relative error decreases with increasing energy, as expected (we have more passing events, which implies better statistics). We also see that below about 5 eV, we have that the relative error increases above 10%. This tells us to expect a similar behavior for the limits set at low masses, as many of these are being set by our sensitivity to events with recoil energies this low.


4.5.3 Limit Setting and Results

The objective of this DM search was to set conservative limits on the spin-independent interaction of dark matter particles with masses below $1.5 \text{ GeV}/c^2$. For the lower edge of the limit contour, we use the optimum interval (OI) method with unknown background, which is a method for selecting an interval within the ROI that contains few events as compared to what would have been expected from (in this case) a DM signal.[181, 182]. Practically, this is done by scaling the DM-nucleon scattering cross section of the expected signal for a given DM mass, and finding the cross section above which (at some specified confidence level) we would have expected to see a DM signal in some range of values above the measured spectrum. For the upper edge of the limit contour, we use a modified version of the publicly available VERNE code [183], which uses a Poisson counting method to calculate the effects of overburden [184–186] on the DM signal, as similarly used in Refs. [122, 187, 188]. Both limit-setting methods assume that the full measured event rate could be due to a DM signal and set the limits at the 90% confidence level (C.L.).

Because the velocity distribution for DM with large interacting cross sections has a dependence on the interaction cross section itself, this means that the calculated DM spectra for the overburden are not simply proportional to the scattering cross section (which prevents use of the OI method). To set this Poisson limit for the overburden, a grid search is performed in both DM mass and scattering cross section. For the ROI in the DM search, the code compares the number of events observed with the number of expected events for each DM mass and interaction cross section when including the overburden effects. The



(a)



(b)

Figure 4.14: (a) The measured and calculated trigger efficiency with the FPGA trigger time and 99.7% confidence ellipse cuts on a linear scale. (b) Zoomed in to low energies and plotted on a log-scale.




Figure 4.15: Relative error in the measured trigger efficiency.

90% Poisson limit is calculated by interpolating the expected event numbers at each DM mass as a function of cross section, finding the cross section that corresponds to the 90% Poisson limit on the total observed number of events. This code does make the assumption that DM particles scatter continuously and travel along straight line trajectories, which is a better approximation for heavy DM. However, it has been shown that, for $\mathcal{O}(1)$ GeV masses and below, the VERNE code and Monte Carlo techniques without this assumption (e.g. DAMASCUS-CRUST) achieve similar results [186]. For the overburden assumption, we include the 5 cm thick, five-walled Cu “box” surrounding the section of the dilution refrigerator that contains the detector (which dominates over the shielding of the building), the shielding from the atmosphere, and the shielding from the Earth.

To estimate the statistical error in the limit contour, we compared the results obtained by calculating the signal model using the eight different sets of pulse simulations. Showing this in Fig. 4.16, there is visual noticeable variation in the lower edge of the limits for DM masses below $200 \text{ MeV}/c^2$. The corresponding plot of the relative standard deviation (Fig. 4.17) of the cross section for each DM mass shows that, above $200 \text{ MeV}/c^2$, the error is less than 1%. However, below this mass, the error increases to $\mathcal{O}(10\%)$. This matches our expectations looking at where the OI limits appear to “split” in Fig. 4.16, as well as confirms the idea that the lowest masses in the limit will have more variation due to the higher uncertainty in the trigger efficiency. For the overburden limits, we see that the errors on the overburden cross section (Fig. 4.18) are also of $\mathcal{O}(10\%)$ as compared to the errors on the DM OI limit.




Figure 4.16: The OI and overburden limit for each simulation. The overburden calculation rules out sensitivity to DM masses below $93 \text{ MeV}/c^2$. Though the two curves do not connect, this is because the overburden drops effectively vertically in cross section below this, and a very fine grid search would be required to capture the behavior. Because this would have been computationally expensive, and initial studies showed that the reach would have changed by less than $1 \text{ MeV}/c^2$, we will simply cut the exclusion limit off at a DM mass of $93 \text{ MeV}/c^2$.

The error on the overburden limit does increase as we go lower in DM mass, but it does not increase much more than 10%, and it only does so at the very lowest masses below $100 \text{ MeV}/c^2$.

With variation in both the OI limit and the overburden, we should understand how this variation propagates to the final limit. To do this, we take the central value (median) of the limits for each DM mass, and estimate its error as the standard deviation of the limits—as opposed to taking the mean value and the standard error of the mean, which would reduce the error by a factor of $\sqrt{8}$. Compared to the latter, the central value method is more conservative in its error estimation (i.e. it will give larger errors), chosen to ensure that we are at worst overestimating the uncertainty. We show the effect of this central value method for the calculation of the 1σ band in Fig. 4.19.

We see that the 1σ band for the OI limit is quite small, which is as expected when plotting 10% errors on a log-scale over many orders of magnitude. For the overburden limit, the 1σ band is effectively not visible on this scale, which is as expected due to its smaller error. We also see that the overburden calculation removed our ability to set limits on some of the lowest




Figure 4.17: The relative error of the OI limit cross sections at each DM mass




Figure 4.18: The relative error of the cross sections in the overburden limit at each DM mass.




Figure 4.19: The DM exclusion contour set by the upper limit and the overburden calculation, including the 1σ error band in the limit set from the eight pulse simulations.

masses in the upper limit calculation, as shown by the vertical line connecting the two limit techniques. One might wonder if there should be a smoother connection between the two, as the overburden limit should continue to decrease in cross section for lower DM masses. In practice, this is true, but it would require a very fine grid-search to find this exact path due to how quickly the limit changes after this lowest mass. Due to computational limitations and the negligible improvement in the limit, this was not pursued.

The results of the dark matter search are shown in Fig. 4.20, where we have taken the median of the limits calculated for the eight simulations at each DM mass, and compared to other pertinent DM searches in the same parameter space [122, 123, 128, 189–192]. The 10% variation is not plotted, as it would not be visible in the figure. For DM masses between 93 and 140 MeV/c^2 , these results provide the most stringent limits for nuclear-recoil DM signals using a cryogenic detector. For DM masses between 220 MeV/c^2 and 1.35 GeV/c^2 , they are the most stringent limits achieved in an above-ground facility. For these low DM masses, the large cross sections approach the level at which the Born approximation used in the standard DM signal model begins to fail [193]. However, in the absence of a generally accepted alternative model and to be comparable to other experiments (all of which also use the Born approximation in this regime), we decided to keep it in our signal model as well.

In Fig. 4.21, we show the data spectrum for reconstructed energies below 40 eV and DM signal curves for various DM masses for a single pulse simulation, where the cross sections




Figure 4.20: The 90% C.L. limits on the spin-independent DM-nucleon cross section as a function of DM mass for this work (solid red line), compared to results from other experiments [122, 123, 128, 189–192]. For above-ground experiments with overburden calculations, the previously ruled out parameter space is shown as the gray shaded region, and the new parameter space ruled out from this search is shown as the red shaded region. For the Collar 2018 surface result, which uses a liquid scintillator cell operated at 1°C , an overburden calculation would be useful for comparison to the upper edges of the various contours for the surface searches. We note that the systematic error in the baseline energy resolution changes the result within the error of the limit’s line width, thus we only include the result from the 3.86 eV calibration.

from the OI limit are used. The approximate location of the optimum interval is apparent for each dark matter mass. We see that the sensitivity to DM masses below $400\text{ MeV}/c^2$ corresponds to recoil energies below 50 eV, with the lowest masses requiring energy sensitivity down to 15 eV. In future experiments, we will need to continue to improve our baseline energy resolution in order to be able to set thresholds at lower values.

In this search, we see an excess of events for recoil energies below about 100 eV, emerging above the roughly flat rate from Compton scattering of the gamma-ray background. This excess of events could be from an unknown external background or due to detector effects such as stress-induced microfractures [194]. As other experiments have observed excess events in searches for low-mass nuclear-recoiling DM [122, 123, 128, 136], understanding this background is of pivotal importance. In Chapter 5, we will discuss further the possible




Figure 4.21: The event spectrum for the DM search data below 40 eV in reconstructed energy. The data have been normalized to events per gram per day per eV and have been corrected for the signal efficiency of the data-quality cuts, but not the confidence ellipse and trigger time cuts. The colored dashed lines represent the calculated event rates for selected DM cross sections and masses from the 90% C.L. OI limit for a single pulse simulation, where the optimum intervals in recoil energy are below 40 eV. Sensitivity to DM masses below 400 MeV/ c^2 corresponds to recoil energies below 40 eV, with the lowest masses requiring energy sensitivity down to about 15 eV.

various backgrounds, as well as studies that have been done to rule some out for the CPD.

4.6 Conclusion

Using a detector with $\sigma_E = 3.86 \pm 0.04$ (stat.) $^{+0.19}_{-0.00}$ (syst.) eV baseline energy resolution operated in an above-ground facility with an exposure of 9.9 g d, we probe parameter space for spin-independent interactions of DM with nucleons for dark matter particles with masses above 93 MeV/ c^2 . The range from 93 to 140 MeV/ c^2 was previously not accessible to cryogenic detectors. These results also set the most stringent limits for above-ground nuclear-recoil signals from dark matter for masses between 220 MeV/ c^2 and 1.35 GeV/ c^2 . This was achieved using a single readout channel composed of QETs distributed on a Si substrate, with a recoil energy threshold set at 16.3 eV.

The results of this work were accomplished despite the high background rates in our

surface facility because of the excellent baseline energy resolution of the detector. There are plans to operate this detector in an underground laboratory, where we expect to have a significantly lower Compton scattering background rate. This will allow further study of the excess events observed in the ROI, hopefully providing insight into the origin of the event rate that is limiting the results reported here, as discussed in Chapter 5.

These results also demonstrate the potential of athermal phonon sensors with eV-scale baseline energy resolution for future dark matter searches via DM-nucleon interactions. Because this detector has a large surface area relative to its small volume, it is not optimal for a DM search. The baseline energy resolution of such devices scales with the square root of the number of QETs, which itself is proportional to the instrumented area (assuming the same QET design used by the CPD) [139, 195]. Thus, a decrease in the instrumented area, with an increase in volume, should lead to improvements in baseline energy resolution. Future work is planned to design detectors of volume $\sim 1 \text{ cm}^3$ (detailed in C. W. Fink’s thesis [147]), for which it is reasonable to expect roughly an order of magnitude improvement in baseline energy resolution through these geometric considerations alone. With improved baseline energy resolution comes a lower energy threshold, allowing a search for spin-independent DM-nucleon interactions for even lower DM masses and a clear path to surpassing the existing noncryogenic detector constraints on sub-100 MeV/ c^2 DM interacting with nucleons.

Chapter 5

Investigation of Excess Signals

Many LDM experiments have observed various excess signals above the expected background, as detailed in Ref. [196]. As noted in the previous chapter, the SuperCDMS-CPD DM search too has its own excess signals. In this chapter, we will motivate our hypothesis of stress-induced microfracture backgrounds via the ruling out of other backgrounds, as well as showing its likely existence in other devices run by our research group. We also propose a solution to removing this background from our detectors via stress-free holding schemes and passive vibration isolation.

5.1 Excess Signals in CPD

There are a multitude of possible sources of the excess signals observed in the CPD analysis that rise over the roughly flat expected differential rate of Compton scattering of the gamma ray background. Before going through them, we first return to the spectrum from the DM search. For convenience in studying the excess, we will integrate the observed reconstructed energy spectrum (i.e. the main plot of Fig. 4.2) in terms of the reconstructed energy via

$$R(E') = \int_x^{E_{ROI}} dE' \frac{\partial R}{\partial E'}(E'), \quad (5.1)$$

where $E_{ROI} = 240$ eV is the maximum energy within the DM search ROI and $\partial R/\partial E'$ is defined by

$$\frac{\partial R}{\partial E'}(E') = \int_0^\infty dE_T \int_0^\infty dE_0 \Theta(E_T - \delta) \varepsilon(E', E_T, E_0) P(E', E_T | E_0) \frac{\partial R}{\partial E_0}(E_0), \quad (5.2)$$

which has been reproduced from Eq. (4.9). We can count up the total number of events within the ROI and divide by the 22 hours of live time, which returns the total rate within the ROI of 1.9 Hz. If we normalize this to the CPD mass of 10.6 g, this is a rate of 0.18 Hz/g (i.e. the rate of events that were tagged with FPGA trigger energies above threshold). Calculating $R(E')$ over the entire ROI, we have the mass-normalized rate above some offline reconstructed




Figure 5.1: The mass-normalized integrated rate of the SuperCDMS-CPD DM search spectrum within the ROI, with the location of the FPGA trigger threshold noted. Because we are plotting as a function of the offline reconstructed energy, the smearing of the energies shows events below threshold—these events are above threshold in FPGA trigger energy.

energy E' , as plotted in Fig. 5.1. In this figure, we also compare to the rate expected from the flat background at 2×10^5 events/(kg keV day), noting that the functional form is simply a constant multiplied by $(E_{ROI} - E')$. In this case, that constant is 2.5 mHz/(g keV), taking care to convert units correctly. In other words, the y -intercept of the flat background is 0.6 mHz/g, as seen in the figure.

We see that there are two characteristic exponential excess backgrounds from that expected from the flat background. There is one that is dominant from the energy threshold to about 25 eV and another that is dominant from 25 eV to about 100 eV. To make an educated guess at the origin of these backgrounds, we must go through all of the possibilities, which are: random fluctuations, Cherenkov interactions, transition radiation, luminescence, other low energy interactions with high energy particles, neutrons, radon, electromagnetic interference (EMI) signals, and stress-induced microfractures.

5.1.1 Expectation from Random Fluctuations

To understand what is expected from random fluctuations above threshold, we can generate events from our understanding of the noise of the CPD. To do this, we take our noise PSD, as depicted in Fig. 5.2, and generate simulated randoms from it via the function `gen_noise` in




Figure 5.2: Comparison of the PSD, template, and optimal filter (see Appendix C for the definition) for offline processing of both the DM search data and any simulated data.

QETPY [105]. The workflow of this function is to sample from a standard normal distribution to create white noise, take the Fourier transform of the sampled data, multiply by the square root of the PSD, and then inverse Fourier transform back to time domain. This is equivalent to applying a transfer function to white noise to have the desired frequency response (our noise spectrum). Note that this randoms generation assumes that the noise is uncorrelated and stationary, which is likely not a true representation of our noise environment, but is the best that we can do with our data from the single-channel CPD.

For this study, we created 50,000 simulated randoms using the above noise PSD and ran an unconstrained time degree-of-freedom OF to roughly simulate the FPGA triggering algorithm, as the FPGA triggering algorithm is effectively an unconstrained OF due to its convolution over traces with a 25 ms long window and choosing the maximum amplitude. For the generated randoms, this means that we do not put the data fully through the same analysis pipeline as the DM search dataset, but we expect that this should give a good rough estimate.

Taking the OF amplitude spectrum from these generated randoms, we can calculate the




Figure 5.3: Comparison of integrated spectrum of DM Search to the randoms spectrum calculated from the noise generation routine, where these data have been generated from the experimentally measured noise PSD. The black dashed line cuts off at the FPGA trigger threshold with a rate of 0.2 Hz/g, as we set the energy threshold on the generated randoms at the FPGA threshold.

integrated rate via Eq. (5.1) and compare it to the DM search spectrum, as well as the expectation from the flat background. The result of this is shown in Fig. 5.3.

We see that the randoms spectrum is roughly within a factor of two of the sub-30 eV region of the DM search spectrum. Because the simulated data from the experimental PSD assumes a normal distribution, nonstationary noise, and no correlations between frequencies, the factor of two difference is not unexpected, as each of these phenomena are likely existent in the DM search data to some degree. With this small difference, it is reasonable that the sub-30 eV background could indeed be due to randoms and does not require explanation via some physical background. With this in mind, pushing towards lower threshold DM searches will also mean pushing of this randoms background to lower energies, and the main background to be worried about is the apparent sub-100 eV excess background.

To understand the remaining component of the excess background, we carry out a least squares fit assuming an exponential background to the residual of its shape (after subtracting the expectation from the flat background). Doing this for the 30 to 100 eV region, we find that the exponential background has a characteristic energy scale of $E_{char.} = 16$ eV (that is, it goes as $\exp(-E'/E_{char.})$) and a y -intercept of 4.4 mHz/g. It remains to be seen if this background is well-modeled as a single exponential below threshold, which can be confirmed




Figure 5.4: Comparison of the summed expected spectra from randoms, the flat background, and the fitted exponential excess to the observed DM search spectrum.

in future low-threshold searches.

In Fig. 5.4, we compare the randoms spectrum, the DM search spectrum, the flat background, and the exponential background. Together, the flat and exponential backgrounds provide excellent agreement with the observed spectrum above 30 eV, with a small amount of discrepancy below this point likely due to the generated randoms not being truly representative of random events in the DM search. Thus, we have a reasonable explanation of the sub-30 eV background, while we still lack an explanation of the remaining exponential background in the 30 to 100 eV region. This requires us to continue ruling out other possible sources of such a background.

5.1.2 Other Possible Sources

To recap, the remaining other sources of this exponentially increasing excess background could be one of the following: Cherenkov interactions, transition radiation, luminescence, other low energy interactions with high energy particles, neutrons, radon, EMI signals, or stress-induced microfractures. In the next few paragraphs, we will discuss whether or not any of these are likely explanations.




Figure 5.5: (Figure from Ref. [197]) (Left) An arbitrary detector being held in place and its signals read out by some type of cabling system. Generally, these detectors are held inside a Cu housing. The purple arrows represent high-energy charged particles (electrons or muons) interacting with the various components of the system. The red lines denote photons emitted from Cherenkov interactions, and the blue lines denote those emitted from transition radiation. (Right) For the same arbitrary setup, a diagram of possible interactions of a high energy particle exciting electrons and holes, which then relax and recombine, releasing energy in the form of photons (luminescence) or phonons (the green lines).

Cherenkov Interactions, Transition Radiation, or Luminescence

Cherenkov interactions, transition radiation, and luminescence have the potential to create low-energy events that could be detected by in low-threshold dark matter searches, such as in the CPD DM search. Cherenkov interactions are related to when charged particles pass through dielectric materials with a velocity greater than the in-medium speed of light. When this occurs, photons with eV-scale energies can be emitted via these high-energy particle interactions. Similarly, transition radiation occurs at the interface between two regions with different dielectric properties, where the passage of high-energy charged particles can too emit photons of with eV-scale energies. Luminescence (or scintillation) occurs when some particle (e.g. a photon, electron, etc.) scatters off a material, excites an electron into the conduction band, which then releases energy after returning to the valence band in the form of photons (or possibly phonons if an indirect bandgap or there are in-gap states from impurities). These processes are diagrammatically shown in Fig. 5.5 for an arbitrary detector setup.

In the case of the CPD DM search, the CPD is a Si substrate held by cirlex clamps and read out by our PCB electronics that are within the hexagonal Cu housing. Thus, there are multiple nonconductive materials with various dielectric properties which all can emit low-energy photons via Cherenkov interactions, transition radiation, or luminescence. In Ref. [197], P. Du *et al.* have theoretically calculated the expected rates of each of these backgrounds and compared to the experimentally observed excess background in the SuperCDMS-CPD DM search. They find that these backgrounds can only account for up to 10% of the

observed excess exponential background, and thus cannot be the dominant source.

Other Low Energy Interactions with High Energy Particles, Neutrons

There does remain backgrounds from other low-energy interactions with high-energy particles, as well as interactions with neutrons. From only the CPD DM dataset at SLAC, we cannot claim anything about these backgrounds. Fortunately, we were able to have our SuperCDMS collaborators run the CPD at their underground facility CUTE (Cryogenic Underground TEST facility), located at SNOLAB in Sudbury, Ontario [198, 199]. In particular, we were able to measure event spectra in their low-background cryostat, though without a calibration source. However, we can achieve a rough energy calibration through our estimate of the energy resolution via the NEP, as discussed in Chapter 2. In this scenario, we can take the measured OF amplitudes in current, and convert to energy scales via the ratio of the measured OF amplitude resolution and the expected baseline energy resolution of the CPD. Note that the expected baseline energy resolution has a dependence on the collection efficiency of the detector, which was measured to be $\varepsilon_{ph} = 13\%$ (recall Section 3.4.3). In the case of running this detector at CUTE, we would expect this efficiency to be roughly the same, where changes could be due to an increase in charged impurities as compared to the state of the detector when ran at SLAC, or perhaps oxygen and other gases remaining on the surface of the detector due to not pumping to vacuum efficiently. Thus, we can only use this as a “rough calibration” (that becomes nonlinear at a few hundred eV), but it should be roughly correct at low energies.

The CPD was operated at CUTE in a few runs, and we will compare the spectra from the CPD DM search at SLAC to those from Runs 11 and 14 at CUTE. In Run 11, there was a single 24-hour dataset that we will compare to from mid-January, 2020. In Run 14, there were two datasets that we will compare to: 23200301_140601 (22 hours starting on March 1, 2020) and 232200303_231114 (8 hours starting on March 3, 2020). These numeric sequences denote the date and time of the beginning of data taking—both in early March 2020, but two days apart. Extracting offline OF amplitudes in the same way as in the DM search at SLAC, we can compare the spectra of all of these datasets, as done in Fig. 5.6.

Between the CUTE runs and SLAC, there are many similarities. Comparing SLAC and CUTE Run 11, we see that the spectra are actually quite similar above about 150 eV. However, due to an unknown change between CUTE Runs 11 and 14, this background disappears for the Run 14 datasets. However, between all four datasets, we see a very similar exponential background with nearly the same characteristic energy scale. The continued existence of this background and similar order of magnitude suggests that going underground did not reduce the background. Because of the excellent shielding at CUTE from both the Earth and the facility’s water shield, the lack of change means that low-energy interactions of high-particles (beyond Cherenkov interactions, transition radiation or luminescence), as well as neutron backgrounds, are likely not the source of this background. The figure also shows that the rate has decreased after two days between the two Run 14 datasets, suggesting a time-dependence that we would not expect for these types of backgrounds.




Figure 5.6: Comparison of the various datasets with the rough calibration. The SLAC R44 line corresponds to a rough calibration of the CPD DM search at SLAC spectrum. The y values are the differential rate spectra for each dataset in units of DRU (Differential Rate Unit), a common abbreviation used to denote events/(kg keV day). Comparing the two CUTE R14 datasets, we see that there is a decrease in the rate after only two days, providing another reason that this background is likely not from low-energy interactions with high-energy particles (we would not expect a decreasing time-dependence from this type of background). This figure was also reproduced in Matt Pyle’s talk at the EXCESS2022 Workshop [200].

Radon Contamination

Furthermore, at keV-scale energies in the CUTE Run 14 datasets, the differential rate is on the order of 1000 events/(kg keV day). One source that could be creating these events is ^{210}Pb surface contamination of the detector or its housing (resulting from radon exposure), as the detector had not been subject to radon control for up to two years of unmitigated exposure. In order to estimate the expected rate of events due to from ^{210}Pb surface contamination, we start with the national average of radon concentration indoors being about 50 Bq/m³. This concentration will affect our detector through radon progeny depositing onto the material surfaces, a process called “plate-out”. For a surface with some area, the amount of radon progeny deposition on that surface is related to an effective plate-out height, such that all the progeny of the radon above that area up to this height will be deposited. For a typical cleanroom, this value is about $h \approx 15\text{ cm}$ [201]. This detector was enclosed in a housing




Figure 5.7: Energy spectra for ^{210}Pb daughter decay (plot by Peter Redl), simulated as part of the SuperCDMS SNOLAB projected sensitivity [202].

while in a cleanroom, and thus is more likely that the effective plate-out height is about 1 cm. In either case, we will estimate the expected rate due to this contamination for both scenarios for a duration two years.

To estimate the ^{210}Pb activity, we multiply the concentration of radon indoors by the plate-out height, the surface area of the CPD (93.6 cm^2), and the ratio of contamination time to the mean lifetime of ^{210}Pb , which gives a total ^{210}Pb activity of 4.5 mBq and $30\text{ }\mu\text{Bq}$ for plate-out heights of 15 cm and 1 cm, respectively. To convert these values to an expected differential rate on the CPD, we refer to the energy recoil distributions for the various progeny, as shown in Fig. 5.7, focusing on the dominant ^{210}Pb spectrum. We can estimate the expected DRU within the energies that have been measured in the CPD (up to 6 keV) by taking the fractional integral of the ^{210}Pb spectrum up to 6 keV (about 0.27), multiplying it by the ^{210}Pb activity, and dividing by the CPD mass, which gives an expected differential rate from ^{210}Pb of 1760 events/(kg keV day) and 117 events/(kg keV day) for plate-out heights of 15 cm and 1 cm, respectively, where the maximal value is on the order of the observed 1000 events/(kg keV day) rate. However, the effective plate-out height for this scenario is quite conservative because the detector was stored in a covered housing, as opposed to being left out uncovered in a cleanroom. Thus, we do not expect that nearly as much radon had

plated out onto the substrate, and the rate corresponding to effective plate-out height of 1 cm is a much better estimate (an order of magnitude below what was observed). Though there are other progeny (^{210}Bi , ^{210}Po) which could increase these rates, their differential rates are expected to be significantly lower than that from the ^{210}Pb , based on takeaways from the SuperCDMS SNOLAB sensitivity projections [202].

Surface contamination from ^{210}Pb could, in principle, have been the origin in the absolute worst case scenario, but it is highly unlikely given that the expectation for a CPD stored in a housing can only explain up to roughly 10–15% of these excess signals. If one were to consider ^{210}Pb surface and bulk contamination in the Cu housing, this may increase the contribution to as much 50% of the excess signals, but still too low to fully explain our observations.

EMI Signals

Signals originating from electromagnetic interference (EMI) are common sources of excess backgrounds in rare event searches. For the CPD, these signals would couple into the TES electronically, rather than coupling into the substrate. Because EMI signals are extremely fast as compared to the time constants of a TES, they can be treated as a dirac Delta pulse. In this scenario, a TES would respond to an EMI pulse with an electronic rise time and the TES thermal fall time, similarly to what is seen in response to a square wave jitter down the bias line (e.g. Fig. 2.9 in Chapter 2). For the CPD, the electronic rise time was measured to be $3\mu\text{s}$ and the primary TES fall time was measured to be $58\mu\text{s}$. However, as noted in Chapter 3, the characteristic pulse shape at low energies was seen to have a $20\mu\text{s}$ rise time, which is attributed to the collection of athermal phonons. As the excess exponential background becomes dominant at low energies, then we would expect a large population of events $3\mu\text{s}$ rise time, of which none were observed in the CPD datasets, ruling out EMI signals as the source of the excess exponential background.

Of our proposed sources of these excess signals, we have yet to rule out stress-induced microfractures. In the following sections, we will further motivate this background through other tests and discuss methods to mitigate it.

5.2 The Case for Stress-Induced Microfractures

When the CPD is installed in its housing with the cirlex clamps, the force between the clamps and the wafer provide the thermal connection for heat to flow from the Si substrate to the bath. As the system is cooled from 300 K to 8 mK, both the clamps and the Si will undergo thermal contraction each at a rate defined by their coefficient of thermal expansion (CTE). If the CTEs are different enough, then this can create a stress at the interface after cooling down due to the difference in how much the two materials have contracted. In this metastable state, stress can be released through energy fluctuations or athermal events (e.g. recoils), dropping the interface to lower energy state while releasing energy over time into the



Figure 5.8: (Left) The QP.4 chip with the $25\,\mu\text{m} \times 100\,\mu\text{m} \times 40\,\text{nm}$ and $50\,\mu\text{m} \times 200\,\mu\text{m} \times 40\,\text{nm}$ rectangles being read out, which is simply resting on the Cu housing. (Right) The QP.4 chip with the $25\,\mu\text{m} \times 100\,\mu\text{m} \times 40\,\text{nm}$ and $50\,\mu\text{m} \times 200\,\mu\text{m} \times 40\,\text{nm}$ rectangles being read out, which is glued to the Cu housing with GE-7031 varnish. Photo by William Page.

Si. These events are what can be thought of as microfractures, i.e. some micro-dislocation of the materials at their interface.

A common alternative holding scheme to clamps for TESs on $1\,\text{cm}^2$ chips is to glue them with some type of adhesive, such as GE-7031 varnish, which should be subject to the same concept of stress induced by CTE mismatch. While cirlex has been reported to have CTEs of 30×10^{-6} and $118 \times 10^{-6}\,\text{K}^{-1}$ depending on orientation [203], GE-7031 varnish has been estimated based on other polymers to have a CTE of about $50 \times 10^{-6}\,\text{K}^{-1}$ [204]. When comparing to crystalline Si, the CTE is on the order of $1 \times 10^{-6}\,\text{K}^{-1}$ [205], showing that a CTE mismatch is reasonable to expect in these detector holding schemes. Given this, we would expect to see an excess amount of events on a TES-based sensor whose substrate has been glued with GE-7031 varnish as opposed to an unglued (or resting) version, if stress-induced microfractures are a significant background. In the next section, we will do this comparison with simple TES rectangles.

5.2.1 Glued TES Rectangles

In our cryogen-free dilution refrigerator at UC Berkeley, we installed two $1\,\text{cm}^2$ Si substrates, as pictured in Fig. 5.8, for which each had the mask design shown in Fig. 5.9. These




Figure 5.9: Full mask design of the QP.4 chip with TES rectangles of various areas, with and without antenna structures. Red is W, and blue is Al. Note that the surface of this chip is almost entirely covered with passive W (i.e. W that is not read out by our electronics). The other side of each of these 1 cm^2 chips is bare Si (i.e. there is no W or Al added to the Si on that side).




Figure 5.10: Zoomed in on the TES rectangles without antenna structures, these TESs scale by factors of 4 in area, while keeping the aspect ratio the same. Thus, it is expected that the normal resistances stay the same, while the in-transition bias powers will scale with area.

Table 5.1: Various parameters for the different TES rectangles on the QP.4 chips. Each dataset was taken with the cryocooler pulse-tube on, as phonons from vibrationally induced events are collected in the large amount of passive W, rather than the small TES rectangles.

Device	P_0 [fW]	τ_+ [μ s]	τ_- [μ s]	σ_A [nA]	$\frac{\partial I}{\partial P}(0)$ [$\frac{\text{A}}{\mu\text{W}}$]	σ_E [meV]
TES25x100 Glued	17.9	1.4	94.6	1.07	21.6	31.0
TES50x200 Glued	73.2	1.6	38.9	1.91	10.1	52.4
TES25x100 Resting	18.0	1.6	60.3	1.34	22.1	24.8
TES50x200 Resting	82.6	1.6	43.5	1.50	9.7	49.6

$T_c \approx 55$ mK chips are from the QP.4 detector mask, designed by Suhas Ganjam in 2017. The original purpose of these simple rectangular TESs was to test power-scaling laws via the rectangles changing area with the same aspect ratios, where we have zoomed in to show more detail in Fig. 5.10. The same rectangles were also included on the chip with antenna-like structures for EMI tests (coupling with RF), but these TESs were not read out in this study.

For each chip in Fig. 5.8, we are reading out the TES rectangles with dimensions $25 \mu\text{m} \times 100 \mu\text{m} \times 40 \text{ nm}$ and $50 \mu\text{m} \times 200 \mu\text{m} \times 40 \text{ nm}$, where the left chip is simply resting on the Cu housing and the right chip has been glued to the housing with GE-7031 varnish. In this configuration, resting chip thermalizes through the superconducting Al wire bonds that are used for reading out the current change in the TES, while the glued chip thermalizes mainly through the GE-7031 varnish connection to the Cu plate.

With this setup, we can follow the TES characterization steps as outlined in Chapter 2, where we analyze IV and $\partial I/\partial V$ sweep data for understanding our TESs. In this case, we have reproduced the pertinent parameters for studying stress-induced microfractures in Table 5.1, where P_0 is the TES bias power, τ_+ is the electronic TES time constant, τ_- is the thermal TES time constant, σ_A is the OF amplitude baseline resolution (RMS), $\frac{\partial I}{\partial P}(0)$ is




Figure 5.11: (Left) Current-referenced PSDs for each of the TESs on the two QP.4 chips. (Right) NEPs for each of the TESs on the two QP.4 chips. The smaller two TESs on each chip have slightly different high frequency responses in power space, which could be due to slightly different properties of the TESs originating from deviations in the fabrication process. The low frequency (flat) NEP is about the same for TESs of the same dimensions.

the zero-frequency value of the power-to-current transfer function, and σ_E is the OF energy baseline energy resolution (RMS). In order to calculate σ_E , we convert OF amplitudes from current to energy for an event via

$$E_{OF} = A_{OF} \frac{\int dt s(t)}{\frac{\partial I}{\partial P}(0)}, \quad (5.3)$$

where we further define E_{OF} as the OF energy of this event, A_{OF} as the OF amplitude (in current) of this event, and $s(t)$ as the amplitude-normalized pulse template. For $s(t)$, we use a template with a rise time of τ_+ and a fall time of τ_- for each device, as this is what would be expected for an event that is directly incident on the TES (as opposed to phonon collection via superconducting Al fins in QETs). The energy estimate provided by E_{OF} only provides a rough energy estimate for single events on a TES (as it neglects phenomena such as collection efficiency or pulse shape saturation). Because substrate events would have widely varying phonon collection times due to the large amount of passive W in this QP.4 chip (see Fig. 5.9), any events that are coincident on multiple TESs on the same chip would not have an accurate energy estimate.

To study these four TES rectangles, we took continuous data for four hours. We used the triggering algorithm in PYTESDAQ [206] for extraction of OF amplitudes from the continuous data and saving events with amplitudes that were greater than a threshold of $6\sigma_A$. In PYTESDAQ, there is support for flagging when events occur at the same time (coincident events) or if an event was only on one TES (singles), which we used when running the event triggering algorithm. For extracting OF amplitudes, we have our previously defined pulse template $s(t)$, and we use the measured current-referenced PSDs as shown in Fig. 5.11.




Figure 5.12: (Left) Comparison of the spectra of the TES25x100 rectangle on the glued and resting chips, where each of the events were marked as coincident with the TES50x200 rectangle on the same chip. (Right) Comparison of the spectra of the TES50x200 rectangle on the glued and resting chips, where each of the events were marked as coincident with the TES25x100 rectangle on the same chip. In each plot, we keep the OF amplitudes in units of current because the energies of these shared events are not well defined, as a significant (and variable) fraction of energy is collected by the passive W on the chip.

Because we have fully characterized our TESs, we can convert these current-referenced PSDs to NEPs, also shown in Fig. 5.11, where we see that the NEPs scale with a factor of two between the two TES geometries, which is expected for TESs dominated by phonon noise.

Because the triggering algorithm flags events as coincidences or singles, we can select data that are coincidences or singles, respectively defined as when two TESs on the same chip both have reconstructed OF amplitudes greater than $6\sigma_A$, or only a single TES does. Furthermore, we selected the region of data without large baseline fluctuations and added a loose (5σ) chi-square goodness-of-fit cut to all four TESs individually, each of which had a passage fraction of $\sim 90\%$. With our energies estimated and our data quality cuts applied, we can measure the event spectra seen on each of the TESs. We split this up into comparing the spectra of the coincidences and the spectra of the singles.

We first select events that are coincident between TES25x100 and TES50x200 on each of the chips. To understand if there is a change in the coincident event spectrum with and without the chip being glued, we will compare the TES25x100 glued spectrum to the TES25x100 resting spectrum and the TES50x200 glued spectrum to the TES50x200 resting spectrum, as shown in Fig. 5.12. Note that these spectra are shown in units of current due to the large amount of passive W in the mask design, as previously discussed. The difference between the glued and resting spectra is striking: there is a large increase in the near-threshold event rate for the TESs on the glued chip, which is completely absent in the spectra of TESs on the resting chip. Because the only significant difference between the two chips are that one is glued and one is not, this provides strong evidence that the presence




Figure 5.13: (Left) Comparison of the spectra of the TES25x100 rectangle on the glued and resting chips, where each of the events were marked as occurring independently of the TES25x100 rectangle on the same chip (i.e. was a single). (Right) Comparison of the spectra of the TES50x200 rectangle on the glued and resting chips, where each of the events were marked as occurring independently of the TES50x200 rectangle on the same chip (i.e. was a single). As opposed to Fig. 5.12, events collected into a single TES are better defined in energy, and we can roughly estimate the collected energy using Eq. (5.3), remembering that this is only a rough estimate and that the main takeaway from this figure is that the single rates with and without glue are about the same except very near the energy threshold.

of the glue has increased the low-energy event rate on the TESs. Furthermore, the higher amplitude event rates are the same, meaning that the glue did not have an effect at these higher energies.

We next can look at the singles event rate on each of the TESs and make the same comparison. That is, we would like to know if there a change in the singles spectra with or without the presence of glue. In Fig. 5.13, we compare these spectra and again see that there is an excess of events for TESs on the glued chip near threshold, which is not seen for the TESs on the resting chip. Though this is seen in the singles spectrum, it is natural to explain this excess as substrate events that were not appreciably detected by the other TES on the same chip. Because of the large amount of passive W, it is not unlikely that the phonons produced by some substrate event would only be seen by one of the TESs depending on the location of the event (i.e. the energy was collected by the passive W or another TES not being read out).

The significant difference in the low-energy coincident spectra between TESs on glued and resting substrates is consistent with the idea of stress-induced microfractures. The GE-7031 varnish that provides the thermal link for the glued chip can have different thermal contraction properties than the Si substrate, such that the interfaces contracting at different scales can cause deviations of nuclei at the surface of the Si, releasing phonons into the substrate. In this study of the TES rectangles, we have been limited in understanding the

energy scale of these possible microfracture events due to the passive W. In order to estimate the energy scale, we would need a device with very low amounts of passive structures, such that we have a high and position-independent phonon collection efficiency, as well as an optical calibration source. This study will be done in future work in the Pyle lab.

5.3 Mitigation of Stress-Induced Events

With stress-induced events motivated as an origin to these excess events, we must consider how to change our experimental setup for the CPD such that these types of events are mitigated. To do this, we will propose a stress-free holding scheme, thermalization through a gold bonding pad, and reduction of vibrational sensitivity via a passive vibration isolation system.

5.3.1 Stress-Free Holding Schemes

To hold a CPD with minimal stress on the substrate, we know that we need to abandon the circlx clamps, and we also should not use some type of glue given our the studies of the Si chips from the previous section. The natural solution is to simply have the CPD resting directly on Cu, such as with the resting Si chip as well as other cryogenic light detectors [159]. However, we note that the surface area of contact will be much higher with the 3-inch diameter CPD than for a 1 cm^2 chip. Because of this large surface area, the resting CPD will have many points of contact for frictional rubbing. To reduce these points of contact, we can adopt the design pictured in Fig. 5.14.

In this design, the points of contact between the Cu holder and a CPD would be only near the perimeter of the wafer. We also note that this region of contact is inset from the connection of the holder to the housing, in order to restrict movement of the CPD when installing the detector. This design also includes three symmetrically placed tapped holes, which allow for the installation of screws to raise the CPD from this region of contact. In this case, the three screws would be the only point of contact, which may further reduce possible frictional rubbing events. Note that, because a standard 3-inch wafer generally has primary and secondary flats for alignment, we have also included these in the cutout geometry of the holder. In Table 5.2, we have placed the relevant SEMI dimensions for a standard 3-inch wafer, which are from the SEMI M1-0302 specifications.

5.3.2 Thermalization through Gold Pads

With this holder, however, we would not be able to thermalize a CPD without a thermal connection beyond that given by the superconducting Al wire bonds. In order to provide an excellent thermal connection, we can add a gold wire-bonding pad to a CPD, from which we can use gold wire to connect to the stress-free holder. In practice, the Cu holder should be gold-plated, as our group has found it difficult to achieve a good bonding connection with




Figure 5.14: Stress free holder for the CPD for removal of stress-induced events, designed by Samuel Watkins. This holder will be used in a future run, and has not been run yet as of June 2022.

Table 5.2: The SEMI specifications for a standard 3-inch wafer.

Dimension	Length [μm]
Diameter	76200 ± 630
Primary Flat Length	22220 ± 3170
Secondary Flat Length	11180 ± 1520

Table 5.3: The thermal conductive of thin gold wire for different levels of impurity, values from Ref. [207].

Purity [%]	$\lambda(T)$ $\left[\frac{\text{W}}{\text{K cm}}\right]$
99.7	$0.115T + 7 \times 10^{-4}T^2$
99.9	$0.325T$
99.999	$4.35T$

gold wire onto bare Cu. To ensure that the proposed gold wire thermalization scheme can handle the 3.86 pW bias power, we must ensure that both the wire itself has a good thermal conductance, as well as the gold pad.

For the conductance through the gold wire, the process that limits the conductivity is scattering of electrons by impurities. In Table 5.3, we give the thermal conductivity values that have been reported in Ref. [207], which we can use to check how well 3.86 pW of power flows through a single gold wire bond of varying purity. We can relate the thermal conductivity $\lambda(T)$ of a wire to the power flow through it via

$$P_{\text{wire}} = \frac{A_{\text{wire}}}{\ell_{\text{wire}}} \int_{T_b}^{T_{\text{Au pad}}} dT \lambda(T), \quad (5.4)$$

where A_{wire} is the cross sectional area of the wire, ℓ_{wire} is the length of the wire, $T_{\text{Au pad}}$ is the temperature of the Au pad, and T_b is the bath temperature. For our purposes, we will assume that we have bonded a single gold wire with a diameter of $25.4 \mu\text{m}$ and a length of 1 cm, and that the bath temperature is $T_b = 8 \text{ mK}$. With these values, we can plot P_{wire} as a function of $T_{\text{Au pad}}$ and compare to the bias power of the CPD, as done in Fig. 5.15.

For the lowest purity gold wire, the power load of the CPD does not increase the temperature of the Au pad to more than 9 mK, which is far from the T_c of the device. Because commercial-grade gold wire can be found with purities of 99.99% (or 4N in the industry nomenclature), this means that one gold wire should easily be able to hand the CPD heat load. In practice, it would be advisable to add more wires in the event of failure, which would also serve only to improve the thermal conductance.

With the gold wire known to be able to handle the CPD bias power, the next step is to ensure that the thermal conductance between the Si absorber and the gold pad is large enough, such that the Si absorber of the CPD does not equilibrate at a high temperature (i.e. significantly close to T_c). As measured in Appendix A, the heat flow between a Si absorber and a gold pad is

$$P_{\text{pad}} = \Sigma_{ep,Au} V_{\text{pad}} (T_{\text{abs}}^n - T_{\text{Au pad}}^n), \quad (5.5)$$

where $\Sigma_{ep,Au}$ is the gold electron-phonon coupling constant, V_{pad} is the gold bonding pad volume, T_{abs} is the temperature of the Si absorber, and $n = 5$ is the power-law exponent. In




Figure 5.15: The heat flow through a single gold wire of diameter of $25.4\,\mu\text{m}$ and a length of $1\,\text{cm}$ as a function $T_{\text{Au pad}}$, comparing to the bias power of the CPD ($P_0 = 3.86\,\text{pW}$) and its $T_c = 41.5\,\text{mK}$.

Appendix A, we noted that our fabrication facilities at Texas A&M University are currently creating gold pads with $750\,\text{nm}$ thickness. Thus, we will assume this thickness for a prospective gold pad on the CPD and calculate what area of pad will be needed for its $3.86\,\text{pW}$ bias power.

In Fig. 5.16, we show the heat flow from the absorber to the gold pad as a function of absorber temperature for different size gold pads, where we assume that the gold wire itself is pure enough such that the pad has a temperature of $T_{\text{Au pad}} = 8\,\text{mK}$. At this interface, we clearly must be careful about the size of the pad, as a volume of $V_{\text{pad}} = 1\,\text{mm}^2 \times 750\,\text{nm}$ would allow the absorber to cool down to about $T_{\text{abs}} = 25\,\text{mK}$. The thermal flow between the TES and the absorber similarly goes as

$$P_{\text{TA}} = \Sigma_{ep,W} V_{\text{TES}} (T_c^n - T_{\text{abs}}^n), \quad (5.6)$$

where $\Sigma_{ep,W}$ is the W electron-phonon coupling constant, V_{TES} is the total TES volume, and T_c is the TES superconducting critical temperature. Thus, the fractional change in bias power for a device with an absorber at a temperature at T_{abs} rather than effectively at the bath temperature T_b is given by

$$\varepsilon_{\text{abs}} = \frac{T_c^n - T_{\text{abs}}^n}{T_c^n - T_b^n}. \quad (5.7)$$




Figure 5.16: The heat flow from absorber to a gold pad of thickness 750 nm and varying area as a function T_{abs} , comparing to the bias power of the CPD ($P_0 = 3.86$ pW) and its $T_c = 41.5$ mK.

Table 5.4: The expected Si absorber temperature and fractional change in bias power ε_{abs} of the CPD for different pad areas, where ε_{abs} is defined in Eq. (5.7). Each of these values assume a bath temperature of $T_b = 8$ mK.

Pad Area	T_{abs} [mK]	ε_{abs}
1 mm ²	25	0.921
4 mm ²	18	0.985
1 cm ²	10	0.999

For $T_{\text{abs}} = 25$ mK and $T_b = 8$ mK, this would give a fractional change in bias power of $\varepsilon_{\text{abs}} = 0.92$, meaning that the expected measured bias power of the CPD would be lowered by 8% to 3.55 pW due to the introduction of this parasitic power, significantly degrading the detector's performance. For the 4 mm² and 1 cm² gold pad areas, we can follow the same calculation, and we have summarized the results in Table 5.4.

The effect on the CPD becomes smaller as we increase the gold pad size. However, we must remember that we do not want the pad to be so large that its passive area begins to appreciably collect the signal. From Table 3.1 when we characterized the CPD, we have that it has an active surface area of 1.9% and passive surface area of 0.2%, where the total area




Figure 5.17: (Left) Comparison of characteristic time-domain traces for the 0.25% coverage SPICE Melange detector with PT on and off. (Right) Comparison of current-referenced PSD for the 0.25% coverage SPICE Melange detector with PT on and off.

of the CPD wafer is about 93.6 cm^2 . Thus, the 1 mm^2 , 4 mm^2 , and 1 cm^2 gold pads would increase the passive surface area to 0.21%, 0.24%, and 1.27%, respectively. Although the 1 cm^2 pad would have excellent heat flow, the passive area would be on the order of the active area, and this pad would likely act as a significant area for collection of phonons that we would not read out, thus degrading the phonon collection efficiency and the detector resolution. Instead, a pad on the order of 4 mm^2 would likely have a small effect on the phonon signal, while ensuring that the Si absorber can cool down to a low enough temperature to have a negligible effect on the bias power.

With care taken to ensure that the thermal connections of the gold pad are designed well, we should be able to thermalize a stress-free CPD, where we have added a gold pad of area 4 mm^2 . The next-generation CPD will have a mask design such that the equilibrium bias power will be lower due the optimization of the QETs for lower surface coverage, which should allow thermalization with a smaller gold pad than would be needed for this CPD with a 3.86 pW bias power.

5.3.3 Vibrational Sensitivity

Although we have proposed a scheme for holding a CPD without creating stress-induced events and its thermalization, we must also take into account that these sensitive devices will be run in a cryogen-free dilution refrigerator. In these systems, a common source of environmental noise can be attributed to the pulse-tube (PT) cryocooler, which is critical for achieving the sub-10 mK base temperatures. The oscillation of the PT at 1.4 Hz will create a source of vibrational noise within the dilution refrigerator, which can then be seen by sensitive detectors. In the case of resting detectors, the effects of these vibrations can be stark, as shown for one of our devices in Fig. 5.17. The device shown is the 0.25%

coverage SPICE Melange detector, designed by C. W. Fink [147]. Here, we use this detector to show that the effect of PT vibrations is significant for a resting device, in that it is visibly apparent in time domain, and is seen as excess low frequency noise in the PSD. Thus, we would expect these events to degrade our baseline energy resolution as well as create excess vibrationally-induced events.

The decoupling of sensitive detectors or instruments from environmental vibrations has been pursued by many cryogenic experiments, usually using a three-dimensional elastic pendulum [208] for passive vibration isolation [209–215]. These systems have all had success in reducing vibrationally-induced backgrounds, and such a system will be necessary to run a detector in a stress-free holding scheme. In Appendix B, we propose such a system for the Pyle group dilution refrigerator, which has been designed to achieve excellent attenuation of vibrations with frequencies above about 10 Hz.

5.4 Summary

In this chapter, we have motivated the hypothesis of stress-induced events as a likely source for the excess signals observed in the DM search dataset through ruling out many other backgrounds and a comparison of TEs on a glued and resting substrate. With the stress-free strategies put forward in the preceding sections, it should be possible to make such a background negligible. In this scenario, we would expect a CPD with improved baseline energy resolution to significantly improve upon the low-mass reach of the SuperCDMS-CPD DM search.

Chapter 6

Conclusion and Outlook

In this thesis, we started with basic TES concepts, applied them to a QET-based detector, carried out an LDM search that achieved world-leading limits for a cryogenic device, investigated stress-induced events as a source of excess signals in that search, and proposed their mitigation through stress-free holding schemes. In the near future, our lab at UC Berkeley will have an installed passive vibration isolation system, and then stress-free CPDs can be run to improve the reach of the SuperCDMS-CPD DM search. A next generation CPD energy sensitivity would be improved upon through a lowering of T_c and a lowering of the surface coverage of the sensors, and such a device has been designed by S. Zuber and C. W. Fink with an expectation of 1 eV baseline energy resolution. This CPD would improve upon the results in this thesis, especially if ran in a low-background cryostat. However, orders of magnitude improve to lower DM masses through the nuclear recoil channel will require orders of magnitude of improvement in baseline energy resolution.

Ultimately, these detectors will be limited by kinematic mismatching of sub-GeV/ c^2 LDM with masses less than the detector nucleus when searching for LDM through elastic nuclear recoils. This concept is demonstrated in Fig. 6.1, which shows the maximum recoil energies of the expected spectrum of DM-nucleon interactions for various detector materials. Significantly improving the current LDM mass reach (from 100 MeV/ c^2 to 1 MeV/ c^2) through this interaction channel with, e.g., Si-based detectors will require athermal sensors with meV-scale energy thresholds. Given that state-of-the-art TESs have been shown to be capable of energy thresholds on the order of hundreds of meV [98], TES and QET detector R&D will need to produce detectors with improvements of energy sensitivities by a few orders of magnitude to vastly improve the DM search results through nuclear recoils.

While R&D continues on lowering energy thresholds, DM detection schemes via inelastic recoils have been proposed to circumvent the limitations brought forward by kinematic matching. For DM-nucleon interactions, limits can be set using DM signal models based on Bremsstrahlung [216] or the Migdal effect [217, 218]. Each of these processes are based on secondary radiation that can accompany a nuclear recoil. For Bremsstrahlung, there is some probability of photon emission when a dark matter particle recoils with a nucleus, where detection of this photon energy with a calorimeter can allow probes of much lower LDM




Figure 6.1: Maximum recoil energy of the differential rate spectra for various target materials. As DM mass drops below the mass of the target nucleus, the imparted energy to the target becomes vanishingly small.

masses. In the case of the Migdal effect, the process is quite analogous, where an electron is ionized from a nucleus rather than a photon. In, e.g., limits set by XENON1T [188], the Migdal effect outperforms the low-mass reach of Bremsstrahlung. In fact, it has been shown that the Migdal effect dominates over Bremsstrahlung, essentially making Bremsstrahlung irrelevant for LDM direct detection [219]. However, the Migdal effect has not been confirmed experimentally, thus the limits set by them are mostly an expectation rather than an exclusion.

For DM-electron interactions, interactions with polar materials [139], superfluid ^4He [140], scintillating crystals [138], and excitation and readout of electron-hole pairs with narrow bandgap materials via cryogenic charge amplifiers [220–223] have all been pursued as promising experimental avenues for probing of lower mass LDM. For scintillating crystals such as GaAs, a CPD with sub-eV resolution would prove to be an excellent match for the 1.33 eV scintillation photons. The same holds true for the scintillation signals expected from superfluid ^4He , which are of order 1 eV, as well as potentially the triplet excimer or He adsorption signals in such an experiment. All of these experiments are quite novel and require R&D for background reduction and low-energy calibration, but all offer intriguing and promising directions for future LDM searches.

The recent experimental push towards LDM searches has created an exciting new direction within the world of DM detection, where small-scale labs have the potential to make large-scale discoveries. With many different experimental directions and methods that can

be pursued, the future is bright for pushing to lower DM masses and potentially making a discovery—answering the fundamental question I started with in this thesis: what makes up our universe?

Bibliography

- [1] N. Aghanim *et al.* (Planck Collaboration), Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.* **641**, A6 (2020).
- [2] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helv. Phys. Acta* **6**, 110 (1933).
- [3] F. Zwicky. On the Masses of Nebulae and of Clusters of Nebulae. *Astrophys. J.* **86**, 217 (1937).
- [4] S. Ettori, A. Donnarumma, E. Pointecouteau, T. H. Reiprich, S. Giodini, L. Lovisari, and R. W. Schmidt. Mass profiles of Galaxy Clusters from X-ray analysis. *Space Sci. Rev.* **177**, 119 (2013).
- [5] X.-P. Wu and L.-Z. Fang. A Statistical Comparison of Cluster Mass Estimates from Optical/X-Ray Observations and Gravitational Lensing. *Astrophys. J.* **483**, 62 (1997).
- [6] Y. Sofue, Y. Tutui, M. Honma, A. Tomita, T. Takamiya, J. Koda, and Y. Takeda. Central Rotation Curves of Spiral Galaxies. *Astrophys. J.* **523**, 136 (1999).
- [7] Y. Sofue and V. Rubin. Rotation Curves of Spiral Galaxies. *Annu. Rev. Astron. Astrophys.* **39**, 137 (2001).
- [8] J. D. Simon and M. Geha. The Kinematics of the Ultra-faint Milky Way Satellites: Solving the Missing Satellite Problem. *Astrophys. J.* **670**, 313 (2007).
- [9] M. Loewenstein and R. E. White III. Prevalence and Properties of Dark Matter in Elliptical Galaxies. *Astrophys. J.* **518**, 50 (1999).
- [10] P. Zyla *et al.* (Particle Data Group), Review of Particle Physics. *Prog. Theor. Exp. Phys.* **2020**, 083C01 (2020).
- [11] S. Burles and D. Tytler. The Deuterium abundance towards Q1937–1009. *Astrophys. J.* **499**, 699 (1998).
- [12] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky. A Direct Empirical Proof of the Existence of Dark Matter. *Astrophys. J.* **648**, L109 (2006).
- [13] S. W. Randall, M. Markevitch, D. Clowe, A. H. Gonzalez, and M. Bradač. Constraints on the Self-Interaction Cross Section of Dark Matter from Numerical Simulations of the Merging Galaxy Cluster 1E 0657-56. *Astrophys. J.* **679**, 1173 (2008).

-
- [14] M. Bradač, S. W. Allen, T. Treu, H. Ebeling, R. Massey, R. G. Morris, A. von der Linden, and D. Applegate. Revealing the Properties of Dark Matter in the Merging Cluster MACS J0025.4-1222. *Astrophys. J.* **687**, 959 (2008).
 - [15] W. Hu and S. Dodelson. Cosmic Microwave Background Anisotropies. *Annu. Rev. Astron. Astrophys.* **40**, 171 (2002).
 - [16] W. Hu and M. White. A CMB polarization primer. *New Astron.* **2**, 323 (1997).
 - [17] E. W. Kolb and M. S. Turner. *The Early Universe*. Ed. by D. Pines. Vol. 69. Boca Raton, Florida: CRC Press, (1990). ISBN: 978-0-201-62674-2.
 - [18] V. Springel *et al.* Simulations of the formation, evolution and clustering of galaxies and quasars. *Nature* **435**, 629 (2005).
 - [19] S. Perlmutter *et al.* (The Supernova Cosmology Project), Measurements of Ω and Λ from 42 High-Redshift Supernovae. *Astrophys. J.* **517**, 565 (1999).
 - [20] P. de Bernardis *et al.* (BOOMERanG Collaboration), A flat universe from high resolution maps of the cosmic microwave background radiation. *Nature* **404**, 955 (2000).
 - [21] A. Balbi *et al.* Constraints on cosmological parameters from MAXIMA-1. *Astrophys. J.* **545**, L1 (2000).
 - [22] A. D. Miller, R. Caldwell, M. J. Devlin, W. B. Dorwart, T. Herbig, M. R. Nolta, L. A. Page, J. Puchalla, E. Torbet, and H. T. Tran. A Measurement of the Angular Power Spectrum of the Cosmic Microwave Background from $l = 100$ to 400. *Astrophys. J.* **524**, L1 (1999).
 - [23] A. G. Riess *et al.* Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *Astron. J.* **116**, 1009 (1998).
 - [24] É. Aubourg *et al.* (BOSS Collaboration), Cosmological implications of baryon acoustic oscillation measurements. *Phys. Rev. D* **92**, 123516 (2015).
 - [25] G. Efstathiou *et al.* Evidence for a non-zero Λ and a low matter density from a combined analysis of the 2dF Galaxy Redshift Survey and cosmic microwave background anisotropies. *Mon. Not. R. Astron. Soc.* **330**, L29 (2002).
 - [26] R. R. Caldwell, R. Dave, and P. J. Steinhardt. Cosmological Imprint of an Energy Component with General Equation of State. *Phys. Rev. Lett.* **80**, 1582 (1998).
 - [27] J. Silk *et al.* *Particle Dark Matter: Observations, Models and Searches*. Ed. by G. Bertone. Cambridge: Cambridge Univ. Press, (2010). ISBN: 978-1-107-65392-4.
 - [28] G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric dark matter. *Phys. Rep.* **267**, 195 (1996).
 - [29] H. Haber and G. Kane. The search for supersymmetry: Probing physics beyond the standard model. *Phys. Rep.* **117**, 75 (1985).

-
- [30] B. W. Lee and S. Weinberg. Cosmological Lower Bound on Heavy-Neutrino Masses. *Phys. Rev. Lett.* **39**, 165 (1977).
 - [31] M. Battaglieri *et al.* US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report (2017). arXiv: 1707.04591 [hep-ph].
 - [32] J. Alexander *et al.* Dark Sectors 2016 Workshop: Community Report (2016). arXiv: 1608.08632 [hep-ph].
 - [33] Y. Kahn and T. Lin. Searches for light dark matter using condensed matter systems. *Rep. Prog. Phys.* (2022).
 - [34] M. Viel, G. D. Becker, J. S. Bolton, and M. G. Haehnelt. Warm dark matter as a solution to the small scale crisis: New constraints from high redshift Lyman- α forest data. *Phys. Rev. D* **88**, 043502 (2013).
 - [35] L. Lopez-Honorez, O. Mena, S. Palomares-Ruiz, and P. Villanueva-Domingo. Warm dark matter and the ionization history of the Universe. *Phys. Rev. D* **96**, 103539 (2017).
 - [36] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro. Analyzing the Discovery Potential for Light Dark Matter. *Phys. Rev. Lett.* **115**, 251301 (2015).
 - [37] P. Galison and A. Manohar. Two Z's or not two Z's? *Phys. Lett. B* **136**, 279 (1984).
 - [38] B. Holdom. Two U(1)'s and ϵ charge shifts. *Phys. Lett. B* **166**, 196 (1986).
 - [39] B. Patt and F. Wilczek. Higgs-field portal into hidden sectors (2006). arXiv: hep-ph/0605188.
 - [40] M. Pospelov, A. Ritz, and M. Voloshin. Secluded WIMP dark matter. *Phys. Lett. B* **662**, 53 (2008).
 - [41] D. E. Kaplan, M. A. Luty, and K. M. Zurek. Asymmetric dark matter. *Phys. Rev. D* **79**, 115016 (2009).
 - [42] K. M. Zurek. Asymmetric Dark Matter: Theories, signatures, and constraints. *Phys. Rep.* **537**, 91 (2014).
 - [43] M. J. Strassler and K. M. Zurek. Echoes of a hidden valley at hadron colliders. *Phys. Lett. B* **651**, 374 (2007).
 - [44] R. N. Mohapatra and V. L. Teplitz. Mirror dark matter and galaxy core densities. *Phys. Rev. D* **62**, 063506 (2000).
 - [45] Y. Hochberg, E. Kuflik, T. Volansky, and J. G. Wacker. Mechanism for Thermal Relic Dark Matter of Strongly Interacting Massive Particles. *Phys. Rev. Lett.* **113**, 171301 (2014).
 - [46] Y. Hochberg, E. Kuflik, H. Murayama, T. Volansky, and J. G. Wacker. Model for Thermal Relic Dark Matter of Strongly Interacting Massive Particles. *Phys. Rev. Lett.* **115**, 021301 (2015).

-
- [47] Y. Hochberg, E. Kuflik, and H. Murayama. SIMP Spectroscopy. *J. High Energy Phys.* **2016**, 90 (2016).
 - [48] E. Kuflik, M. Perelstein, N. R.-L. Lorier, and Y.-D. Tsai. Elastically Decoupling Dark Matter. *Phys. Rev. Lett.* **116**, 221302 (2016).
 - [49] E. Kuflik, M. Perelstein, N. R.-L. Lorier, and Y.-D. Tsai. Phenomenology of ELDER Dark Matter. *J. High Energy Phys.* **2017**, 78 (2017).
 - [50] L. J. Hall, K. Jedamzik, J. March-Russell, and S. M. West. Freeze-In Production of FIMP Dark Matter. *J. High Energy Phys.* **2010**, 80 (2010).
 - [51] R. Essig, M. Fernandez-Serra, J. Mardon, A. Soto, T. Volansky, and T.-T. Yu. Direct Detection of sub-GeV Dark Matter with Semiconductor Targets. *J. High Energy Phys.* **2016**, 46 (2016).
 - [52] X. Chu, T. Hambye, and M. H. G. Tytgat. The Four Basic Ways of Creating Dark Matter Through a Portal. *J. Cosmol. Astropart. Phys.* **05**, 034 (2012).
 - [53] R. Essig, J. Mardon, and T. Volansky. Direct detection of sub-GeV dark matter. *Phys. Rev. D* **85**, 076007 (2012).
 - [54] R. D. Peccei and H. R. Quinn. CP Conservation in the Presence of Pseudoparticles. *Phys. Rev. Lett.* **38**, 1440 (1977).
 - [55] R. D. Peccei and H. R. Quinn. Constraints imposed by CP conservation in the presence of pseudoparticles. *Phys. Rev. D* **16**, 1791 (1977).
 - [56] S. Weinberg. A New Light Boson? *Phys. Rev. Lett.* **40**, 223 (1978).
 - [57] F. Wilczek. Problem of Strong P and T Invariance in the Presence of Instantons. *Phys. Rev. Lett.* **40**, 279 (1978).
 - [58] B. J. Carr and S. W. Hawking. Black holes in the early Universe. *Mon. Not. R. Astron. Soc.* **168**, 399 (1974).
 - [59] B. P. Abbott *et al.* (LIGO Scientific Collaboration and Virgo Collaboration), Observation of Gravitational Waves from a Binary Black Hole Merger. *Phys. Rev. Lett.* **116**, 061102 (2016).
 - [60] B. Carr, F. Kühnel, and M. Sandstad. Primordial black holes as dark matter. *Phys. Rev. D* **94**, 083504 (2016).
 - [61] S. Bird, I. Cholis, J. B. Muñoz, Y. Ali-Haïmoud, M. Kamionkowski, E. D. Kovetz, A. Raccanelli, and A. G. Riess. Did LIGO Detect Dark Matter? *Phys. Rev. Lett.* **116**, 201301 (2016).
 - [62] B. Paczynski. Gravitational Microlensing by the Galactic Halo. *Astrophys. J.* **304**, 1 (1986).
 - [63] K. Griest *et al.* Gravitational Microlensing as a Method of Detecting Disk Dark Matter and Faint Disk Stars. *Astrophys. J.* **372**, L79 (1991).

-
- [64] H. Niikura *et al.* Microlensing constraints on primordial black holes with Subaru/HSC Andromeda observations. *Nature Astron.* **3**, 524 (2019).
- [65] T. Lasserre *et al.* (EROS Collaboration), Not enough stellar mass machos in the galactic halo. *Astron. Astrophys.* **355**, L39 (2000).
- [66] P. Tisserand *et al.* (EROS-2 Collaboration), Limits on the Macho Content of the Galactic Halo from the EROS-2 Survey of the Magellanic Clouds. *Astron. Astrophys.* **469**, 387 (2007).
- [67] T. D. Brandt. Constraints on MACHO Dark Matter from Compact Stellar Systems in Ultra-Faint Dwarf Galaxies. *Astrophys. J. Lett.* **824**, L31 (2016).
- [68] T. S. Li *et al.* (DES Collaboration), Farthest Neighbor: The Distant Milky Way Satellite Eridanus II. *Astrophys. J.* **838**, 8 (2017).
- [69] Y. Ali-Haïmoud and M. Kamionkowski. Cosmic microwave background limits on accreting primordial black holes. *Phys. Rev. D* **95**, 043534 (2017).
- [70] N. K. Porayko *et al.* Parkes Pulsar Timing Array constraints on ultralight scalar-field dark matter. *Phys. Rev. D* **98**, 102002 (2018).
- [71] D. Carney, S. Ghosh, G. Krnjaic, and J. M. Taylor. Proposal for gravitational direct detection of dark matter. *Phys. Rev. D* **102**, 072003 (2020).
- [72] N. L. Figueroa, D. Budker, and E. M. Rasel. Dark matter searches using accelerometer-based networks. *Quantum Sci. Technol.* **6**, 034004 (2021).
- [73] T. Moroi, H. Murayama, and M. Yamaguchi. Cosmological constraints on the light stable gravitino. *Phys. Lett. B* **303**, 289 (1993).
- [74] D. Bauer *et al.* Dark matter in the coming decade: Complementary paths to discovery and beyond. *Phys. Dark Universe* **7-8**, 16 (2015).
- [75] T. Trickle, Z. Zhang, K. M. Zurek, K. Inzani, and S. M. Griffin. Multi-channel direct detection of light dark matter: theoretical framework. *J. High Energy Phys.* **2020**, 36 (2020).
- [76] R. W. Schnee. Introduction to dark matter experiments. *Theoretical Advanced Study Institute in Elementary Particle Physics: Physics of the Large and the Small*. (2011).
- [77] N. W. Evans, C. A. J. O’Hare, and C. McCabe. Refinement of the standard halo model for dark matter searches in light of the Gaia Sausage. *Phys. Rev. D* **99**, 023012 (2019).
- [78] D. Baxter *et al.* Recommended conventions for reporting results from direct dark matter searches. *Eur. Phys. J. C* **81**, 907 (2021).
- [79] E. I. Gates, G. Gyuk, and M. S. Turner. The Local Halo Density. *Astrophys. J.* **449**, L123 (1995).
- [80] F. J. Kerr and D. Lynden-Bell. Review of galactic constants. *Mon. Not. R. Astron. Soc.* **221**, 1023 (1986).

-
- [81] M. C. Smith *et al.* The RAVE survey: constraining the local Galactic escape speed. *Mon. Not. R. Astron. Soc.* **379**, 755 (2007).
 - [82] R. H. Helm. Inelastic and Elastic Scattering of 187-Mev Electrons from Selected Even-Even Nuclei. *Phys. Rev.* **104**, 1466 (1956).
 - [83] J. D. Lewin and P. F. Smith. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. *Astropart. Phys.* **6**, 87 (1996).
 - [84] I. Alkhatib *et al.* (SuperCDMS Collaboration), Light Dark Matter Search with a High-Resolution Athermal Phonon Detector Operated above Ground. *Phys. Rev. Lett.* **127**, 061801 (2021).
 - [85] J. Hafner. Ab-initio simulations of materials using VASP: Density-functional theory and beyond. *J. Comput. Chem.* **29**, 2044 (2008).
 - [86] P. Giannozzi *et al.* QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials. *J. Phys. Condens. Matter* **21**, 395502 (2009).
 - [87] A. Togo and I. Tanaka. First principles phonon calculations in materials science. *Scr. Mater.* **108**, 1 (2015).
 - [88] R. Catena, T. Emken, M. Matas, N. A. Spaldin, and E. Urdshals. Crystal responses to general dark matter-electron interactions. *Phys. Rev. Research* **3**, 033149 (2021).
 - [89] S. M. Griffin, K. Inzani, T. Trickle, Z. Zhang, and K. M. Zurek. Extended calculation of dark matter-electron scattering in crystal targets. *Phys. Rev. D* **104**, 095015 (2021).
 - [90] S. Knapen, J. Kozaczuk, and T. Lin. PYTHON package for dark matter scattering in dielectric targets. *Phys. Rev. D* **105**, 015014 (2022).
 - [91] B. Campbell-Deem, S. Knapen, T. Lin, and E. Villarama. Dark matter direct detection from the single phonon to the nuclear recoil regime (2022). arXiv: 2205.02250 [hep-ph].
 - [92] S. M. Griffin, K. Inzani, T. Trickle, Z. Zhang, and K. M. Zurek. Multichannel direct detection of light dark matter: Target comparison. *Phys. Rev. D* **101**, 055004 (2020).
 - [93] M. A. Lindeman. Microcalorimetry and the transition-edge sensor. PhD thesis. University of California, Davis, (2000).
 - [94] F. C. Wellstood, C. Urbina, and J. Clarke. Hot-electron effects in metals. *Phys. Rev. B* **49**, 5942 (1994).
 - [95] W. A. Little. The Transport of Heat Between Dissimilar Solids at Low Temperatures. *Can. J. Phys.* **37**, 334 (1959).
 - [96] K. D. Irwin, S. W. Nam, B. Cabrera, B. Chugg, and B. A. Young. A quasiparticle-trap-assisted transition-edge sensor for phonon-mediated particle detection. *Rev. Sci. Instrum.* **66**, 5322 (1995).

-
- [97] K. D. Irwin and G. C. Hilton. Transition-Edge Sensors. *Cryogenic Particle Detection*. Ed. by C. Enss. Berlin, Heidelberg: Springer Berlin Heidelberg, (2005), pp. 63–150. ISBN: 978-3-540-31478-3.
- [98] C. W. Fink, S. L. Watkins, T. Aramaki, P. L. Brink, S. Ganjam, B. A. Hines, M. E. Huber, N. A. Kurinsky, R. Mahapatra, N. Mirabolfathi, *et al.* Characterizing TES power noise for future single optical-phonon and infrared-photon detectors. *AIP Adv.* **10**, 085221 (2020).
- [99] M. C. Pyle. Optimizing the design and analysis of cryogenic semiconductor dark matter detectors for maximum sensitivity. PhD thesis. Stanford University, (2012).
- [100] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of Superconductivity. *Phys. Rev.* **108**, 1175 (1957).
- [101] I. J. Maasilta. Complex impedance, responsivity and noise of transition-edge sensors: Analytical solutions for two- and three-block thermal models. *AIP Adv.* **2**, 042110 (2012).
- [102] M. A. Lindeman, B. Dirks, J. van der Kuur, P. A. J. de Korte, R. H. den Hartog, L. Gottardi, R. A. Hijmering, H. F. C. Hoevers, and P. Khosropanah. Relationships Between Complex Impedance, Thermal Response, and Noise in TES Calorimeters and Bolometers. *IEEE Trans. Appl. Supercond.* **21**, 254 (2011).
- [103] N. Wang, F. C. Wellstood, B. Sadoulet, E. E. Haller, and J. Beeman. Electrical and thermal properties of neutron-transmutation-doped Ge at 20 mK. *Phys. Rev. B* **41**, 3761 (1990).
- [104] M. A. Lindeman, S. Bandler, R. P. Brekosky, J. A. Chervenak, E. Figueroa-Feliciano, F. M. Finkbeiner, M. J. Li, and C. A. Kilbourne. Impedance measurements and modeling of a transition-edge-sensor calorimeter. *Rev. Sci. Instrum.* **75**, 1283 (2004).
- [105] C. W. Fink and S. L. Watkins. *QETpy*. Version 1.3.1. (2021).
- [106] J. B. Johnson. Thermal Agitation of Electricity in Conductors. *Phys. Rev.* **32**, 97 (1928).
- [107] H. Nyquist. Thermal Agitation of Electric Charge in Conductors. *Phys. Rev.* **32**, 110 (1928).
- [108] R. L. Stratonovich. Thermal noise in nonlinear resistors. *Radiophys. Quantum Electron.* **13**, 1164 (1970).
- [109] R. L. Stratonovich. *Nonlinear nonequilibrium thermodynamics I: linear and nonlinear fluctuation-dissipation theorems*. Vol. 57. Berlin, Heidelberg: Springer Science & Business Media, (1992). ISBN: 978-3-642-77345-7.
- [110] R. F. Voss and J. Clarke. Flicker ($\frac{1}{f}$) noise: Equilibrium temperature and resistance fluctuations. *Phys. Rev. B* **13**, 556 (1976).

-
- [111] W. S. Boyle and K. F. Rodgers. Performance Characteristics of a New Low-Temperature Bolometer. *J. Opt. Soc. Am.* **49**, 66 (1959).
- [112] D. McCammon. *Thermal Equilibrium Calorimeters – An Introduction*. Ed. by C. Enss. Berlin, Heidelberg: Springer Berlin Heidelberg, (2005), pp. 1–34. ISBN: 978-3-540-31478-3.
- [113] J. C. Mather. Bolometer noise: nonequilibrium theory. *Appl. Opt.* **21**, 1125 (1982).
- [114] N. A. Kurinsky. The Low-Mass Limit: Dark Matter Detectors with eV-Scale Energy Resolution. PhD thesis. Stanford University, (2018).
- [115] C. W. Fink *et al.* (CPD Collaboration), Performance of a large area photon detector for rare event search applications. *Appl. Phys. Lett.* **118**, 022601 (2021).
- [116] N. A. Wakeham *et al.* Thermal fluctuation noise in Mo/Au superconducting transition-edge sensor microcalorimeters. *J. Appl. Phys.* **125**, 164503 (2019).
- [117] S. R. Golwala. Exclusion limits on the WIMP nucleon elastic scattering cross-section from the Cryogenic Dark Matter Search. PhD thesis. University of California, Berkeley, (2000).
- [118] N. Jethava, J. N. Ullom, K. D. Irwin, W. B. Doriese, J. A. Beall, G. C. Hilton, L. R. Vale, and B. Zink. Dependence of Excess Noise on the Partial Derivatives of Resistance in Superconducting Transition Edge Sensors. *AIP Conf. Proc.* **1185**, 31 (2009).
- [119] C. Alduino *et al.* (CUORE Collaboration), First Results from CUORE: A Search for Lepton Number Violation via $0\nu\beta\beta$ Decay of ^{130}Te . *Phys. Rev. Lett.* **120**, 132501 (2018).
- [120] D. Q. Adams *et al.* (CUORE Collaboration), Improved Limit on Neutrinoless Double-Beta Decay in ^{130}Te with CUORE. *Phys. Rev. Lett.* **124**, 122501 (2020).
- [121] E. Andreotti *et al.* ^{130}Te Neutrinoless Double-Beta Decay with CUORICINO. *Astropart. Phys.* **34**, 822 (2011).
- [122] E. Armengaud *et al.* (EDELWEISS Collaboration), Searching for low-mass dark matter particles with a massive Ge bolometer operated above ground. *Phys. Rev. D* **99**, 082003 (2019).
- [123] G. Angloher *et al.* (CRESST Collaboration), Results on MeV-scale dark matter from a gram-scale cryogenic calorimeter operated above ground. *Eur. Phys. J. C* **77**, 637 (2017).
- [124] Q. Arnaud *et al.* (EDELWEISS Collaboration), First Germanium-Based Constraints on Sub-MeV Dark Matter with the EDELWEISS Experiment. *Phys. Rev. Lett.* **125**, 141301 (2020).
- [125] R. Agnese *et al.* (SuperCDMS Collaboration), First Dark Matter Constraints from a SuperCDMS Single-Charge Sensitive Detector. *Phys. Rev. Lett.* **121**, 051301 (2018).

-
- [126] O. Abramoff *et al.* (SENSEI Collaboration), SENSEI: Direct-Detection Constraints on Sub-GeV Dark Matter from a Shallow Underground Run Using a Prototype Skipper CCD. *Phys. Rev. Lett.* **122**, 161801 (2019).
- [127] A. Aguilar-Arevalo *et al.* (DAMIC Collaboration), Constraints on Light Dark Matter Particles Interacting with Electrons from DAMIC at SNOLAB. *Phys. Rev. Lett.* **123**, 181802 (2019).
- [128] A. H. Abdelhameed *et al.* (CRESST Collaboration), First results from the CRESST-III low-mass dark matter program. *Phys. Rev. D* **100**, 102002 (2019).
- [129] F. T. Avignone III, S. R. Elliott, and J. Engel. Double beta decay, Majorana neutrinos, and neutrino mass. *Rev. Mod. Phys.* **80**, 481 (2008).
- [130] W. R. Armstrong *et al.* (CUPID Collaboration), CUPID pre-CDR (2019). arXiv: 1907.09376 [physics.ins-det].
- [131] V. Alenkov *et al.* First results from the AMoRE-Pilot neutrinoless double beta decay experiment. *Eur. Phys. J. C* **79**, 791 (2019).
- [132] O. Azzolini, J. W. Beeman, F. Bellini, M. Beretta, M. Biassoni, C. Brofferio, C. Bucci, S. Capelli, L. Cardani, P. Carniti, *et al.* Background Model of the CUPID-0 Experiment. *Eur. Phys. J. C* **79**, 583 (2019).
- [133] T. Tabarelli de Fatis. Cerenkov emission as a positive tag of double beta decays in bolometric experiments. *Eur. Phys. J. C* **65**, 359 (2010).
- [134] N. Casali, M. Vignati, J. W. Beeman, F. Bellini, L. Cardani, I. Dafinei, S. Di Domizio, F. Ferroni, L. Gironi, S. Nagorny, *et al.* TeO₂ bolometers with Cherenkov signal tagging: towards next-generation neutrinoless double-beta decay experiments. *Eur. Phys. J. C* **75**, 12 (2015).
- [135] R. Essig *et al.* Dark Sectors and New, Light, Weakly-Coupled Particles (2013). arXiv: 1311.0029 [hep-ph].
- [136] N. Kurinsky, D. Baxter, Y. Kahn, and G. Krnjaic. Dark matter interpretation of excesses in multiple direct detection experiments. *Phys. Rev. D* **102**, 015017 (2020).
- [137] A. E. Robinson. Coherent photon scattering background in sub-GeV/ c^2 direct dark matter searches. *Phys. Rev. D* **95**, 021301 (2017).
- [138] S. Derenzo, R. Essig, A. Massari, A. Soto, and T.-T. Yu. Direct detection of sub-GeV dark matter with scintillating targets. *Phys. Rev. D* **96**, 016026 (2017).
- [139] S. Knapen, T. Lin, M. Pyle, and K. M. Zurek. Detection of light dark matter with optical phonons in polar materials. *Phys. Lett. B* **785**, 386 (2018).
- [140] S. A. Hertel, A. Biekert, J. Lin, V. Velan, and D. N. McKinsey. Direct detection of sub-GeV dark matter using a superfluid ⁴He target. *Phys. Rev. D* **100**, 092007 (2019).

-
- [141] W. Knaak, T. Hauß, M. Kummrow, and M. Meißner. Thermalization of Ballistic Phonon Pulses in Dielectric Crystals Below 1K Using Time-Resolved Thermometry. *Phonon Scattering in Condensed Matter V*. Ed. by A. C. Anderson and J. P. Wolfe. Berlin, Heidelberg: Springer Berlin Heidelberg, (1986), pp. 174–176. ISBN: 978-3-642-82912-3.
- [142] A. G. Kozorezov, A. F. Volkov, J. K. Wigmore, A. Peacock, A. Poelaert, and R. den Hartog. Quasiparticle-phonon downconversion in nonequilibrium superconductors. *Phys. Rev. B* **61**, 11807 (2000).
- [143] S. Hansen, F. DeJongh, J. Hall, B. A. Hines, M. E. Huber, T. Kiper, V. Mandic, W. Rau, T. Saab, D. Seitz, *et al.* The Cryogenic Dark Matter Search test stand warm electronics card. *IEEE Nuclear Science Symposium Medical Imaging Conference*. Knoxville, TN, (2010), pp. 1392–1395.
- [144] J. Schweppe, R. D. Deslattes, T. Mooney, and C. J. Powell. Accurate measurement of Mg and Al $K\alpha_{1,2}$ X-ray energy profiles. *J. Electron Spectrosc.* **67**, 463 (1994).
- [145] G. Hölzer, M. Fritsch, M. Deutsch, J. Härtwig, and E. Förster. $K\alpha_{1,2}$ and $K\beta_{1,3}$ x-ray emission lines of the 3d transition metals. *Phys. Rev. A* **56**, 4554 (1997).
- [146] B. Cabrera. Introduction to TES Physics. *J. Low Temp. Phys.* **151**, 82 (2008).
- [147] C. W. Fink. A Gram-Scale Low- T_c Low-Surface-Coverage Athermal-Phonon Sensitive Dark Matter Detector. PhD thesis. University of California, Berkeley, (2022).
- [148] L. A. Zadeh and J. R. Ragazzini. Optimum Filters for the Detection of Signals in Noise. *Proc. IRE* **40**, 1223 (1952).
- [149] E. Gatti and P. F. Manfredi. Processing the Signals From Solid State Detectors in Elementary Particle Physics. *Riv. Nuovo Cim.* **9**, 1 (1986).
- [150] S. J. B. Reed and N. G. Ware. Escape peaks and internal fluorescence in X-ray spectra recorded with lithium drifted silicon detectors. *J. Phys. E Sci. Instrum.* **5**, 582 (1972).
- [151] D. M. Chernyak, F. A. Danevich, A. Giuliani, E. Olivieri, M. Tenconi, and V. I. Tretyak. Random coincidence of $2\nu 2\beta$ decay events as a background source in bolometric $0\nu 2\beta$ decay experiments. *Eur. Phys. J. C* **72**, 1989 (2012).
- [152] D. R. Artusa *et al.* Exploring the neutrinoless double beta decay in the inverted neutrino hierarchy with bolometric detectors. *Eur. Phys. J. C* **74**, 3096 (2014).
- [153] M. Pyle, E. Figueroa-Feliciano, and B. Sadoulet. Optimized Designs for Very Low Temperature Massive Calorimeters (2015). arXiv: 1503.01200 [astro-ph.IM].
- [154] N. Casali, L. Cardani, I. Colantoni, A. Cruciani, S. Di Domizio, M. Martinez, G. Pettinari, and M. Vignati. Phonon and light read out of a Li_2MoO_4 crystal with multiplexed kinetic inductance detectors. *Eur. Phys. J. C* **79**, 724 (2019).
- [155] B. S. Neganov and V. N. Trofimov. Calorimetric method measuring ionizing radiation. *Otkryt. Izobret.* **146**, 215 (1985).

-
- [156] P. N. Luke. Voltage-assisted calorimetric ionization detector. *J. Appl. Phys.* **64**, 6858 (1988).
- [157] L. Cardani, N. Casali, A. Cruciani, H. le Sueur, M. Martinez, F. Bellini, M. Calvo, M. G. Castellano, I. Colantoni, C. Cosmelli, *et al.* Al/Ti/Al phonon-mediated KIDs for UV–vis light detection over large areas. *Supercond. Sci. Technol.* **31**, 075002 (2018).
- [158] K. Schäffner, G. Angloher, F. Bellini, N. Casali, F. Ferroni, D. Hauff, S. S. Nagorny, L. Pattavina, F. Petricca, S. Pirro, *et al.* Particle discrimination in TeO₂ bolometers using light detectors read out by transition edge sensors. *Astropart. Phys.* **69**, 30 (2015).
- [159] M. Barucci, J. W. Beeman, V. Caracciolo, L. Pagnanini, L. Pattavina, G. Pessina, S. Pirro, C. Rusconi, and K. Schäffner. Cryogenic light detectors with enhanced performance for rare event physics. *Nucl. Instrum. Methods Phys. Res. A* **935**, 150 (2019).
- [160] L. Pattavina, N. Casali, L. Dumoulin, A. Giuliani, M. Mancuso, P. de Marcillac, S. Marnieros, S. S. Nagorny, C. Nones, E. Olivieri, *et al.* Background Suppression in Massive TeO₂ Bolometers with Neganov–Luke Amplified Light Detectors. *J. Low Temp. Phys.* **184**, 286 (2015).
- [161] M. Willers, F. v. Feilitzsch, A. Gütlein, A. Münster, J.-C. Lanfranchi, L. Oberauer, W. Potzel, S. Roth, S. Schönert, M. v. Sivers, *et al.* Neganov-Luke amplified cryogenic light detectors for the background discrimination in neutrinoless double beta decay search with TeO₂ bolometers. *J. Instrum.* **10**, P03003 (2015).
- [162] E. Armengaud *et al.* Development of ¹⁰⁰Mo-containing scintillating bolometers for a high-sensitivity neutrinoless double-beta decay search. *Eur. Phys. J. C* **77**, 785 (2017).
- [163] L. Bergé *et al.* Complete event-by-event $\alpha/\gamma(\beta)$ separation in a full-size TeO₂ CUORE bolometer by Neganov-Luke-magnified light detection. *Phys. Rev. C* **97**, 032501 (2018).
- [164] V. Novati, L. Bergé, L. Dumoulin, A. Giuliani, M. Mancuso, P. de Marcillac, S. Marnieros, E. Olivieri, D. V. Poda, M. Tenconi, *et al.* Charge-to-heat transducers exploiting the Neganov-Trofimov-Luke effect for light detection in rare-event searches. *Nucl. Instrum. Methods Phys. Res. A* **940**, 320 (2019).
- [165] J. Rothe, G. Angloher, P. Bauer, A. Bento, C. Bucci, L. Canonica, A. D’Addabbo, X. Defay, A. Erb, F. v. Feilitzsch, *et al.* TES-Based Light Detectors for the CRESST Direct Dark Matter Search. *J. Low Temp. Phys.* **193**, 1160 (2018).
- [166] D. K. Papoulias, T. S. Kosmas, and Y. Kuno. Recent Probes of Standard and Non-standard Neutrino Physics With Nuclei. *Front. Phys.* **7**, 191 (2019).
- [167] M. Tanabashi *et al.* (Particle Data Group), Review of Particle Physics. *Phys. Rev. D* **98**, 030001 (2018).

-
- [168] P. A. R. Ade *et al.* (Planck Collaboration), Planck 2015 results - XIII. Cosmological parameters. *Astron. Astrophys.* **594**, A13 (2016).
- [169] G. Steigman and M. S. Turner. Cosmological constraints on the properties of weakly interacting massive particles. *Nucl. Phys. B* **253**, 375 (1985).
- [170] E. Aprile *et al.* (XENON Collaboration), Dark Matter Search Results from a One Ton-Year Exposure of XENON1T. *Phys. Rev. Lett.* **121**, 111302 (2018).
- [171] D. S. Akerib *et al.* (LUX Collaboration), Results from a Search for Dark Matter in the Complete LUX Exposure. *Phys. Rev. Lett.* **118**, 021303 (2017).
- [172] X. Cui *et al.* (PandaX-II Collaboration), Dark Matter Results from 54-Ton-Day Exposure of PandaX-II Experiment. *Phys. Rev. Lett.* **119**, 181302 (2017).
- [173] R. Essig, A. Manalaysay, J. Mardon, P. Sorensen, and T. Volansky. First Direct Detection Limits on Sub-GeV Dark Matter from XENON10. *Phys. Rev. Lett.* **109**, 021301 (2012).
- [174] E. Aprile *et al.* (XENON Collaboration), Light Dark Matter Search with Ionization Signals in XENON1T. *Phys. Rev. Lett.* **123**, 251801 (2019).
- [175] M. P. De Pascale, A. Morselli, P. Picozza, R. L. Golden, C. Grimani, B. L. Kimbell, S. A. Stephens, S. J. Stochaj, W. R. Webber, G. Basini, *et al.* Absolute spectrum and charge ratio of cosmic ray muons in the energy region from 0.2 GeV to 100 GeV at 600 m above sea level. *J. Geophys. Res.* **98**, 3501 (1993).
- [176] C. W. Fink and S. L. Watkins. *RQpy*. <https://github.com/spice-herald/rqpy>. (2022).
- [177] R. Schönrich, J. Binney, and W. Dehnen. Local kinematics and the local standard of rest. *Mon. Not. R. Astron. Soc.* **403**, 1829 (2010).
- [178] A. Benoit *et al.* (EDELWEISS Collaboration), Measurement of the response of heat-and-ionization germanium detectors to nuclear recoils. *Nucl. Instrum. Methods Phys. Res. A* **577**, 558 (2007).
- [179] R. Agnese *et al.* (SuperCDMS Collaboration), Energy loss due to defect formation from ^{206}Pb recoils in SuperCDMS germanium detectors. *Appl. Phys. Lett.* **113**, 092101 (2018).
- [180] M. Paterno. Calculating efficiencies and their uncertainties (2004). Report Number: FERMILAB-TM-2286-CD.
- [181] S. Yellin. Finding an upper limit in the presence of an unknown background. *Phys. Rev. D* **66**, 032005 (2002).
- [182] S. Yellin. Extending the optimum interval method (2007). arXiv: 0709.2701 [physics.data-an].
- [183] B. J. Kavanagh. *bradkav/verne: Release*. Version v1.0. (2017).

-
- [184] G. D. Starkman, A. Gould, R. Esmailzadeh, and S. Dimopoulos. Opening the window on strongly interacting dark matter. *Phys. Rev. D* **41**, 3594 (1990).
- [185] G. Zaharijas and G. R. Farrar. Window in the dark matter exclusion limits. *Phys. Rev. D* **72**, 083502 (2005).
- [186] T. Emken and C. Kouvaris. How blind are underground and surface detectors to strongly interacting dark matter? *Phys. Rev. D* **97**, 115047 (2018).
- [187] B. J. Kavanagh. Earth scattering of superheavy dark matter: Updated constraints from detectors old and new. *Phys. Rev. D* **97**, 123013 (2018).
- [188] E. Aprile *et al.* (XENON Collaboration), Search for Light Dark Matter Interactions Enhanced by the Migdal Effect or Bremsstrahlung in XENON1T. *Phys. Rev. Lett.* **123**, 241803 (2019).
- [189] A. Aguilar-Arevalo *et al.* (DAMIC Collaboration), Results on Low-Mass Weakly Interacting Massive Particles from an 11 kg d Target Exposure of DAMIC at SNOLAB. *Phys. Rev. Lett.* **125**, 241803 (2020).
- [190] G. Angloher *et al.* (CRESST Collaboration), Results on light dark matter particles with a low-threshold CRESST-II detector. *Eur. Phys. J. C* **76**, 25 (2016).
- [191] Q. Arnaud *et al.* (NEWS-G Collaboration), First results from the NEWS-G direct dark matter search experiment at the LSM. *Astropart. Phys.* **97**, 54 (2018).
- [192] J. I. Collar. Search for a nonrelativistic component in the spectrum of cosmic rays at Earth. *Phys. Rev. D* **98**, 023005 (2018).
- [193] M. C. Digman, C. V. Cappiello, J. F. Beacom, C. M. Hirata, and A. H. G. Peter. Not as big as a barn: Upper bounds on dark matter-nucleus cross sections. *Phys. Rev. D* **100**, 063013 (2019).
- [194] J. Åström, P. C. F. Di Stefano, F. Pröbst, L. Stodolsky, J. Timonen, C. Bucci, S. Cooper, C. Cozzini, F. v. Feilitzsch, H. Kraus, *et al.* Fracture processes observed with a cryogenic detector. *Phys. Lett. A* **356**, 262 (2006).
- [195] Y. Hochberg, M. Pyle, Y. Zhao, and K. M. Zurek. Detecting superlight dark matter with Fermi-degenerate materials. *J. High Energ. Phys.* **2016**, 57 (2016).
- [196] P. Adari *et al.* *EXCESS workshop: Descriptions of rising low-energy spectra*. Ed. by A. Fuss, M. Kaznatcheeva, F. Reindl, and F. Wagner. (2022). arXiv: 2202.05097 [astro-ph.IM].
- [197] P. Du, D. Egana-Ugrinovic, R. Essig, and M. Sholapurkar. Sources of Low-Energy Events in Low-Threshold Dark-Matter and Neutrino Detectors. *Phys. Rev. X* **12**, 011009 (2022).
- [198] P. Camus, A. Cazes, A. Dastgheibi-Fard, K. Dering, G. Gerbier, W. Rau, S. Scorza, and X. Zhang. CUTE: A Low Background Facility for Testing Cryogenic Dark Matter Detectors. *J. Low Temp. Phys.* **193**, 813 (2018).

-
- [199] W. Rau, G. Gerbier, P. Camus, K. Dering, A. Cazes, S. Scorza, X. Zhang, and A. Dastgheibi-Fard. CUTE - A Cryogenic Underground Test Facility at SNOLAB. *J. Phys. Conf. Ser.* **1342**. Ed. by K. Clark, C. Jillings, C. Kraus, J. Saffin, and S. Scorza, 012128 (2020).
- [200] *EXCESS2022 Workshop*. <https://indico.scc.kit.edu/event/2575/>. (2022).
- [201] E. S. Morrison, T. Frels, E. H. Miller, R. W. Schnee, and J. Street. Radon daughter plate-out onto Teflon. *AIP Conf. Proc.* **1921**, 090002 (2018).
- [202] R. Agnese *et al.* (SuperCDMS Collaboration), Projected sensitivity of the SuperCDMS SNOLAB experiment. *Phys. Rev. D* **95**, 082002 (2017).
- [203] M. Daal, N. Zobrist, N. Kellaris, B. Sadoulet, and M. Robertson. Properties of selected structural and flat flexible cabling materials for low temperature applications. *Cryogenics* **98** (2019).
- [204] C. R. Morelock, M. R. Suchomel, and A. P. Wilkinson. A cautionary tale on the use of GE-7031 varnish: low-temperature thermal expansion studies of ScF_3 . *J. Appl. Crystallogr.* **46**, 823 (2013).
- [205] R. Hull. *Properties of Crystalline Silicon*. London, UK: INSPEC, The Institution of Electrical Engineers, (1999). ISBN: 978-0-85296-933-5.
- [206] SPICE/HeRALD Collaboration. *pytesdaq*. <https://github.com/spice-herald/pytesdaq>. (2022).
- [207] V. Novotny. An electronic vapour pressure controller. *Cryogenics* **17**, 451 (1977).
- [208] P. Lynch. Resonant motions of the three-dimensional elastic pendulum. *Int. J. of Nonlin. Mech.* **37**, 345 (2002).
- [209] R. Maisonobe, J. Billard, M. De Jesus, A. Juillard, D. Misiak, E. Olivieri, S. Sayah, and L. Vagneron. Vibration decoupling system for massive bolometers in dry cryostats. *J. Instrum.* **13**, T08009 (2018).
- [210] L. Gottardi, H. van Weers, J. Dercksen, H. Akamatsu, M. P. Bruijn, J. R. Gao, B. Jackson, P. Khosropanah, J. van der Kuur, K. Ravensberg, *et al.* A six-degree-of-freedom micro-vibration acoustic isolator for low-temperature radiation detectors based on superconducting transition-edge sensors. *Rev. Sci. Instrum.* **90**, 055107 (2019).
- [211] H. W. Chan, J. C. Long, and J. C. Price. Taber vibration isolator for vacuum and cryogenic applications. *Rev. Sci. Instrum.* **70**, 2742 (1999).
- [212] S. Caparrelli, E. Majorana, V. Moscatelli, E. Pascucci, M. Perciballi, P. Puppo, P. Rapagnani, and F. Ricci. Vibration-free cryostat for low-noise applications of a pulse tube cryocooler. *Rev. Sci. Instrum.* **77**, 095102 (2006).
- [213] S. Pirro, A. Alessandrello, C. Brofferio, C. Bucci, O. Cremonesi, E. Coccia, E. Fiorini, V. Fafone, A. Giuliani, A. Nucciotti, *et al.* Vibrational and thermal noise reduction for cryogenic detectors. *Nucl. Instrum. Methods Phys. Res. A* **444**, 331 (2000).

-
- [214] S. Pirro. Further developments in mechanical decoupling of large thermal detectors. *Nucl. Instrum. Methods Phys. Res. A* **559**, 672 (2006).
- [215] M. Pelliccione, A. Sciambi, J. Bartel, A. J. Keller, and D. Goldhaber-Gordon. Design of a scanning gate microscope for mesoscopic electron systems in a cryogen-free dilution refrigerator. *Rev. Sci. Instrum.* **84**, 033703 (2013).
- [216] C. Kouvaris and J. Pradler. Probing Sub-GeV Dark Matter with Conventional Detectors. *Phys. Rev. Lett.* **118**, 031803 (2017).
- [217] M. Ibe, W. Nakano, Y. Shoji, and K. Suzuki. Migdal Effect in Dark Matter Direct Detection Experiments. *J. High Energy Phys.* **03**, 194 (2018).
- [218] S. Knapen, J. Kozaczuk, and T. Lin. Migdal Effect in Semiconductors. *Phys. Rev. Lett.* **127**, 081805 (2021).
- [219] N. F. Bell, J. B. Dent, J. L. Newstead, S. Sabharwal, and T. J. Weiler. Migdal effect and photon bremsstrahlung in effective field theories of dark matter direct detection and coherent elastic neutrino-nucleus scattering. *Phys. Rev. D* **101**, 015012 (2020).
- [220] P. Rosa *et al.* Colossal magnetoresistance in a nonsymmorphic antiferromagnetic insulator. *npj Quantum Mater.* **5** (2020).
- [221] A. Juillard *et al.* Low-Noise HEMTs for Coherent Elastic Neutrino Scattering and Low-Mass Dark Matter Cryogenic Semiconductor Detectors. *J. Low Temp. Phys.* **199**, 798 (2019).
- [222] A. Phipps, A. Juillard, B. Sadoulet, B. Serfass, and Y. Jin. A HEMT-based cryogenic charge amplifier with sub-100 eVee ionization resolution for massive semiconductor dark matter detectors. *Nucl. Instrum. Methods Phys. Res. A* **940**, 181 (2019).
- [223] R. Essig, G. K. Giovanetti, N. Kurinsky, D. McKinsey, K. Ramanathan, K. Stifter, and T.-T. Yu. Snowmass2021 Cosmic Frontier: The landscape of low-threshold dark matter direct detection in the next decade. *2022 Snowmass Summer Study*. (2022). arXiv: 2203.08297 [hep-ph].
- [224] P. Virtanen *et al.* SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nat. Methods* **17**, 261 (2020).
- [225] J. T. Karvonen, L. J. Taskinen, and I. J. Maasilta. Electron-phonon interaction in thin copper and gold films. *Phys. Status Solidi C* **1**, 2799 (2004).
- [226] J. T. Karvonen, L. J. Taskinen, and I. J. Maasilta. Observation of disorder-induced weakening of electron-phonon interaction in thin noble-metal films. *Phys. Rev. B* **72**, 012302 (2005).
- [227] V. F. Gantmakher. The experimental study of electron-phonon scattering in metals. *Rep. Prog. Phys.* **37**, 317 (1974).
- [228] M. Kanskär and M. N. Wybourne. Crossover between Dissipative and Nondissipative Electron Transport in Metal Wires. *Phys. Rev. Lett.* **73**, 2123 (1994).

-
- [229] A. Schmid. Electron-phonon interaction in an impure metal. *Z. Phys.* **259**, 421 (1973).
- [230] A. Sergeev and V. Mitin. Electron-phonon interaction in disordered conductors: Static and vibrating scattering potentials. *Phys. Rev. B* **61**, 6041 (2000).
- [231] F. Giazotto, T. T. Heikkilä, A. Luukanen, A. M. Savin, and J. P. Pekola. Opportunities for mesoscopics in thermometry and refrigeration: Physics and applications. *Rev. Mod. Phys.* **78**, 217 (2006).
- [232] D. Rosenberg, A. Lita, A. Miller, S. Nam, and R. Schwall. Performance of photon-number resolving transition-edge sensors with integrated 1550 nm resonant cavities. *IEEE Trans. Appl. Supercond.* **15**, 575 (2005).
- [233] R. G. Budynas and J. K. Nisbett. *Shigley's Mechanical Engineering Design, 9th Edition*. New York, NY: McGraw-Hill, (2011). ISBN: 978-0-07-352928-8.
- [234] R. C. Dhuley, M. Ruschman, J. T. Link, and J. Eyre. Thermal conductance characterization of a pressed copper rope strap between 0.13 K and 10 K. *Cryogenics* **86**, 17 (2017).
- [235] J. Tuttle, E. Canavan, and M. DiPirro. THERMAL AND ELECTRICAL CONDUCTIVITY MEASUREMENTS OF CDA 510 PHOSPHOR BRONZE. *AIP Conf. Proc.* **1219**, 55 (2010).
- [236] J. W. Ekin. *Experimental Techniques for Low-Temperature Measurements: Cryostat Design, Material Properties and Superconductor Critical-Current Testing*. New York, NY: Oxford University Press, Inc., (2006). ISBN: 978-0-19-857054-7.
- [237] R. Berman and C. F. Mate. Thermal Contact at Low Temperatures. *Nature* **182**, 1661 (1958).
- [238] R. Berman. Some Experiments on Thermal Contact at Low Temperatures. *J. Appl. Phys.* **27**, 318 (1956).
- [239] R. C. Dhuley. Pressed copper and gold-plated copper contacts at low temperatures – A review of thermal contact resistance. *Cryogenics* **101**, 111 (2019).
- [240] A. Berryhill, M. Coffey, C. Danaher, L. Mauritsen, B. Pollard, C. Wang, and Z. Zhao. *Cryogenic Engineering and Technologies: Principles and Applications of Cryogen-Free Systems*. Ed. by Z. Zhao and W. Chao. Boca Raton, FL: CRC Press, (2019). ISBN: 978-0-42-919472-6.
- [241] G. A. Costello. *Theory of Wire Rope*. New York, NY: Springer New York, NY, (1997). ISBN: 978-1-4612-1970-5.
- [242] F. Foti and L. Martinelli. Mechanical modeling of metallic strands subjected to tension, torsion and bending. *Int. J. Solids Struct.* **91**, 1 (2016).
- [243] *Technology Applications, Inc.* <https://www.techapps.com/>.

-
- [244] R. Schmitt, G. Tatkowski, M. Ruschman, S. Golwala, N. Kellaris, M. Daal, J. Hall, and E. Hoppe. Thermal conductance measurements of bolted copper joints for SuperCDMS. *Cryogenics* **70**, 41 (2015).
- [245] J. W. Arblaster. Thermodynamic Properties of Copper. *J. Phase Equilib. Diffus.* **36**, 422 (2015).
- [246] H. M. Ledbetter. Temperature behaviour of Young’s moduli of forty engineering alloys. *Cryogenics* **22**, 653 (1982).
- [247] H. M. Ledbetter. Ratio of the shear and Young’s Moduli for polycrystalline metallic elements. *Mater. Sci. Eng.* **27**, 133 (1977).
- [248] N. J. Simon, E. S. Drexler, and R. P. Reed. Properties of copper and copper alloys at cryogenic temperatures. Final report (1992). Report Number: PB-92-172766/XAB; NIST/MONO-177.
- [249] *McMaster-Carr Supply Company*. <https://www.mcmaster.com/>.
- [250] T. A. Schutt. A dark matter detector based on the simultaneous measurement of phonons and ionization at 20 mK. PhD thesis. University of California, Berkeley, (1993).
- [251] A. T. J. Phipps. Ionization Collection in Detectors of the Cryogenic Dark Matter Search. PhD thesis. University of California, Berkeley, (2016).
- [252] N. G. Mast. Cryogenic Detector Development for SuperCDMS. PhD thesis. University of Minnesota, Twin Cities, (2020).
- [253] P. Sprent. Inversion of Nearly Diagonal Matrices. *Math. Gaz.* **49**, 184 (1965).
- [254] R. Storn and K. Price. Differential Evolution - A Simple and Efficient Heuristic for Global Optimization over Continuous Spaces. *J. Glob. Optim.* **11**, 341 (1997).
- [255] J. Lampinen. A constraint handling approach for the differential evolution algorithm. *Proceedings of the 2002 Congress on Evolutionary Computation. CEC’02 (Cat. No.02TH8600)*. Vol. 2. (2002), pp. 1468–1473.

Appendix A

Thermal Conductance Measurements

In this appendix, we will discuss the principles of how we can measure thermal conductances with TESs and measure the conductances for various systems.

A.1 Thermal Conductance Measurement Principles with TESs

The general setup for an arbitrary TES in a dilution refrigerator is pictured in Fig. A.1. There is a TES on an absorber, which itself is connected to the thermal bath (e.g. through detector holders or through a thermalizing Au wire bond). There is also some type of heater that can change the temperature of the thermal bath. Since the TES is in thermal contact with an absorber, which is in thermal contact with the bath, we must measure two thermal conductances, G_{TA} and G_{AB} , to understand the thermal properties of our system. In the following sections, we explain the principles behind the measurements of these two values.

A.1.1 Measuring Thermal Conductance between Absorber and Bath

We begin with measuring G_{AB} , which requires two TESs on the same absorber, one which acts as an absorber thermometer and the other as an absorber heater. The underlying equation that we want to fit is

$$P_{AB} = K_{AB}(T_a^n - T_b^n), \quad (\text{A.1})$$

where P_{AB} is the power flowing from absorber to bath, T_a is the temperature of the absorber, T_b is the temperature of the bath, n is a power-law exponent, and K_{AB} is related to G_{AB} via

$$G_{AB} = \frac{dP_{AB}}{dT_a} = nK_{AB}T_a^{n-1}. \quad (\text{A.2})$$




Figure A.1: The setup of an arbitrary TES. The TES has been fabricated on some type of absorber, which itself is thermally connected to the thermal bath (i.e. the Cu housing in this diagram). The squiggly black line represents some type of thermal connection, which will depend on the detector holding scheme.

To extract G_{AB} , we must fit Eq. (A.1), which requires measuring P_{AB} as a function of T_b , while keeping T_a constant. Because these measurements will be made when the system is in thermal equilibrium, the power flowing from the TESs to the absorber will be equal to the power flowing from the absorber to the bath. Thus, we can bias a TES in the normal regime with a Joule heating of

$$P_{TES} = I_{TES}^2 R_N, \quad (\text{A.3})$$

where I_{TES} is the current flowing through the TES and R_N is the TES normal resistance, both of which can be known through the characterization steps outlined in Chapter 2. The TES current can be related back to the bias current I_{bias} and various resistances of the system via

$$I_{TES} = I_{bias} \frac{R_{sh}}{R_\ell + R_N}, \quad (\text{A.4})$$

where R_{sh} is the shunt resistance, R_ℓ is the load resistance, and R_N is the normal resistance.

To measure P_{AB} as a function of T_b , we can use a heater placed in the bath (usually the mixing chamber in a dilution refrigerator) and measure the change in power as we change the heat load on the bath. To keep T_a constant, we use our two TESs that have been fabricated on the same absorber. The thermometer TES is biased at some point in its superconducting transition, such that any change in temperature in the absorber will change the TES bias




Figure A.2: Diagram of the G_{AB} measurement setup, where TES₁ is the thermometer, and TES₂ is the heater. For TES₁, the inset plot shows a characteristic in-transition TES response in time domain to a square wave. For TES₂, the inset plot shows the bias power as a function of bias current, where the green circle shows the approximate starting area in the normal regime for the G_{AB} measurement.

point. The heater TES is biased such that it is normal, and its Joule heating will heat up the absorber. This setup is diagrammatically shown in Fig. A.2.

The measurement carries out as follows: we bias TES₁ low in transition and heat the bath up to a few mK below the T_c of TES₁ (e.g. start at 42 mK if $T_c = 45$ mK), such that biasing TES₂ normal will keep TES₁ in transition. TES₂ is then biased normal with a power that should be much greater than the bias power of TES₁, such that the latter can be neglected (this should be simple if TES₁ is the more sensitive device). This first point gives P_{AB} at the starting T_b , as well as our starting bias point for TES₁. We then decrease the heat load on the bath, which decreases T_a and equivalently changes the bias point of TES₁, as the power flowing in equilibrium has not yet changed (it is still equal to the Joule heating of TES₂). To remedy this, we increase the bias current on TES₂ until TES₁ has returned to its initial bias point and thus so has T_a . This gives us our second data point for P_{AB} at a new T_b . This procedure is then continued until either TES₂ can no longer heat the absorber up to the same T_a (due to some limitation with the amount of bias current that can be supplied) or we have reached the base temperature of the dilution refrigerator.

In Fig. A.2, the bias point of TES₁ is represented by its response to a square wave. In this case, we would know that, when changing the Joule heating of TES₂, TES₁ has returned to its initial bias point when zero frequency change of its response (i.e. $\partial I/\partial V(0)$) is the same. Alternatively, an equivalent method is to not send a square wave jitter down the bias line, but instead follow the DC value of the TES current. After the DC value changes as the bath is heated, then TES₁ will have returned to its initial point by changing the Joule heating of TES₂ until TES₁ has the same TES current. It is up to the person doing the measurement which they prefer, the former gives more TES information, while the latter may be faster in terms of data taking time (no need to fit the TES response, but simply track the DC value of the current).

At the conclusion of the measurement, we will have P_{AB} as function of T_b at hopefully many bias points, and we can simply fit Eq. (A.1) via a least squares routine to extract G_{AB} , T_a , and n . It may also be useful to calculate the natural time constant $\tau = C/G_{AB}$ where C is the heat capacity of the absorber, which is the characteristic time for the absorber to cool back down to base temperature after being warmed up to about T_c (e.g. a muon hits the detector and pushes the TESs normal). If the absorber is a semiconductor or insulator (usually the case), then we can estimate its heat capacity in the low-temperature Debye limit

$$C^{th} = \frac{12\pi^4 k_b}{5A} \rho V N_A \left(\frac{T_a}{\Theta_D} \right)^3, \quad (\text{A.5})$$

where $k_b = 1.38 \times 10^{-23}$ J/K is the Boltzmann constant, ρ is the density of the absorber, V is the volume of the absorber, A is the atomic mass of the absorber, $N_A = 6.022 \times 10^{23} \text{ mol}^{-1}$ is the Avogadro constant, and Θ_D is the Debye temperature of for the absorber.

A.1.2 Measuring Thermal Conductance between TES and Absorber

To measure G_{TA} , we only need a single TES, but it is dependent on already having done the G_{AB} . The underlying equation that we want to fit is

$$P_{TA} = K_{TA}(T_c^n - T_a^n), \quad (\text{A.6})$$

where T_c is the superconducting critical temperature of the TES, n is the power law exponent, and K_{TA} is related to G_{TA} via

$$G_{TA} = \frac{dP_{TA}}{dT_c} = nK_{TA}T_c^{n-1}. \quad (\text{A.7})$$

As discussed in Section 3.5 of Irwin and Hilton [97], the thermal impedance between the electrons and phonons in a metal with small volume and high power density (e.g. a W TES) dominates. In this case, the power flow is from the electrons to the phonons via $P = \Sigma_{ep} V_{TES} (T_{el}^5 - T_{ph}^5)$, where V_{TES} is the TES volume and Σ_{ep} is a material-dependent

electron-phonon coupling constant (usually on the order of $10^9 \text{ W m}^{-3} \text{ K}^{-5}$). Thus, it is expected that $n = 5$, $K_{TA} = \Sigma_{ep} V_{TES}$, and $G_{TA} = 5 \Sigma_{ep} V_{TES} T_c^4$ for a W TES.

Note that Eq. (A.6) is dependent on T_a , but the experimental setup generally only has the ability to measure T_b , showing why it is important to measure G_{AB} first in order to accurately calculate T_a . Although, in practice, a good thermal connection between absorber to bath will mean that $G_{AB} \gg G_{TA}$ and $T_a \approx T_b$. Thus, in single-channel devices (i.e. only one TES channel on an absorber), it is common to make the assumption that $T_a \approx T_b$ if the thermalization scheme is known to have a large G_{AB} from historical measurements and continue to do a G_{TA} measurement.

For the G_{TA} measurement, the setup is simple: we have a single TES on a substrate, which is thermally connected to the absorber, and the absorber is thermally connected to the bath. For this measurement, we keep the TES at the same transition point while changing the bath temperature via the heater, and we calculate the TES bias power as a function of bath temperature. Since we wait for the system to reach equilibrium, the bias power of the TES P_0 will be equal to the power from the TES to the absorber P_{TA} , which in turn is equal to the power from the absorber to the bath P_{AB} . P_0 is easily defined via

$$P_0 = I_0^2 R_0, \quad (\text{A.8})$$

where I_0 is the current through the TES, and R_0 is TES resistance. There are two common ways to calculate P_0 as a function of T_b for the G_{TA} measurement: the “quick” way and the “clean” way.

The Quick G_{TA} Measurement

The quick measurement is to bias the TES in transition and measure the TES response. If the TES is low enough in transition, then we can use that, in the limit of large loop gain $\mathcal{L}$, the TES complex impedance (Eq. (2.32)) becomes

$$\lim_{\mathcal{L} \rightarrow \infty} Z_{\text{circ}}(\omega) = R_\ell + i\omega L - R_0, \quad (\text{A.9})$$

and the zero-frequency value is $Z_{\text{circ}}(0) = R_\ell - R_0$. Thus, we can solve for $R_0 = |Z_{\text{circ}}(0)| + R_\ell$ and $I_0 = I_{\text{bias}} R_{sh} / (R_0 + R_\ell)$, and we have an estimate of P_0 .

The quick G_{TA} measurement procedure is then as follows (and pictured in Fig. A.3): we bias the TES low in transition and heat the bath up to a few mK below the T_c of TES (e.g. start at 42 mK if $T_c = 45$ mK). We measure the zero-frequency value of the circuit complex impedance, and we immediately have an estimate of P_0 and, equivalently, P_{TA} for at this T_b . In order to convert this T_b to T_a (if we have a G_{AB} measurement), we use Eq. (A.2), giving the first data point for P_{TA} as a function of T_a . We then, as we did with G_{AB} , decrease the heat load on the bath, which decreases T_a and equivalently changes the bias point of the TES, as the power flowing in equilibrium has not yet changed (it is equal to the Joule heating of the TES). Note that the TES in transition is also supplying the Joule heating, in contrast to the G_{AB} measurement system (in which the in-transition TES bias power was




Figure A.3: Diagram of the G_{TA} measurement setup, where TES_1 is kept at the same point in transition as the bath temperature changes. For TES_1 , the inset plot shows a characteristic in-transition TES response in time domain to a square wave.

negligible). When we increase the bias current to the TES so that it returns to the same bias point (i.e. the same R_0) the bias power will increase, and T_a will decrease. Doing this, we have a second data point for the bias power of the TES as a function of T_a . This procedure continues until we have reached the base temperature of the dilution refrigerator (that is, the heater on the mixing chamber has zero heat load).

This measurement can go quickly, but the approximation of $\mathcal{L} \rightarrow \infty$ may introduce a systematic on R_0 , and we may actually have a different true bias power. To avoid this possible systematic, the G_{TA} measurement can be done using clean way, at the expense of longer data taking.

The Clean G_{TA} Measurement

In essence, the clean G_{TA} measurement is done by taking IV sweeps at many different bath temperatures. These IV sweeps can then be interpolated such that we can choose the same TES resistance for each bath temperature from the sweep, immediately extracting the measured I_0 and R_0 and avoiding possible systematic error. The drawback is that these IV sweeps take a fair amount of time (perhaps an hour per sweep), depending on how many points are desired. When doing these for many bath temperatures, the data taking time can quickly become nearly the entire day. Because of this, usually the quick way to do the measurement is done first, with the clean way done in the case of analysis for publication.




Figure A.4: (Figure from Ref. [114]) Left: Full image of the G115 mask design, showing the Al (blue), amorphous Si (green), and W (black). There are four channels, where three are inner (note the three line denoting the split) and one is on the outer ring. Right: Zoom in of the center of the detector, showing the parquet pattern of the QETs.

We show an example of the clean way in Section A.3.4.

With our methods of measuring G_{AB} and G_{TA} explained, we will next show a few example measurements, where all nonlinear least squares fits are done using `optimize.curve_fit` in `SCIPY` [224].

A.2 Thermal Conductance Measurements with G115

Developed by Matt Pyle and Suhas Ganjam, G115 is a prototype detector, consisting of a 3 inch diameter and 1 mm thick Ge wafer, with a four-channel QET mask deposited on one side (mask design shown in Fig. A.4). This detector was made to study QET performance, as well as for understanding the effect high energy events (muons and photons) on high voltage noise, as discussed in detail in Noah Kurinsky's thesis [114]. In this section, we will bypass that discussion and instead summarize the thermal conductance measurements made on this device in a cryogen-free dilution refrigerator as SLAC, in which the detector was thermally connected to bath via cirlex clamps.

For these measurements, we will use two of the inner channels of G115, henceforth known as channels A and B in this section. The measured normal resistances of the channels A

Table A.1: Values measured and calculated for channel B, calculating the power flowing from absorber to bath.

T_b [mK]	I_{bias} [μA]	P_{AB} [pW]
42.15	500.00	32.5
41.75	585.00	44.5
41.45	663.00	57.2
41.05	733.00	69.9
40.75	797.00	82.6
40.40	853.00	94.7
40.09	909.75	107.7
39.05	1049.00	143.2
37.95	1161.00	175.4
37.16	1258.00	205.9
35.87	1399.00	254.6

and B are $200\text{ m}\Omega$ and $182\text{ m}\Omega$, respectively. For each channel, the parasitic resistance is $R_p = 12\text{ m}\Omega$.

A.2.1 Cirlex Clamps

The cirlex clamps holding scheme will have a interfacial thermal resistance due to the elastic constant mismatch between the dissimilar semiconducting substrate and the clamps themselves [95], and the thermal conductance is expected to go as $G \propto T^3$.

G_{AB} Measurement

To measure G_{AB} , we follow the method as outlined in Section A.1.1, where channel A acted as a thermometer, and channel B acted as a heater. The data from this measurement are reported in Table A.1, where the error in P_{AB} is set to be 10% of the measured power, based on an estimated 10% uncertainty in R_{sh} . The dominant systematic when taking these measurements was that the fridge had a stabilization time of $\mathcal{O}(20\text{ min})$, meaning that it was difficult to tell if we had fully stabilized at each point. Furthermore, the size of the detector meant that it had a substantial rate of high energy events (muons and gammas) at the surface, which would heat the substrate. Thus, there may be some bias in the measurements, which is difficult to quantify.

We next fit Eq. (A.2) with n allowed to be any value, and $n = 4$ fixed as expected from the phonon mismatch between the absorber and the clamps. In Fig. A.5, we show the results of these two fits and compare to the measured data, with the various fit parameters reported in Table A.2. It is clear from the fits that we did not take enough data (i.e. go to a low enough bath temperature) to be able to describe the power law behavior with any




Figure A.5: The power flowing from absorber to bath through the cirlex clamps at the various bath temperatures taken for G115. Also shown are the least squares fits with n varying and $n = 4$ fixed.

Table A.2: Values from the least-squares fits to the power flowing from absorber to bath the cirlex clamps. Included in the table are the estimated heat capacity via the low-temperature Debye model using Eq. (A.5), as well as the related thermal time constant τ .

Fit	G_{AB} [$\frac{\text{nW}}{\text{K}}$]	T_a [mK]	n	C^{th} [$\frac{\text{nJ}}{\text{K}}$]	τ [ms]
Any n	29.2 ± 5.0	43.04 ± 0.19	1.6 ± 2.3	3.96 ± 0.05	136 ± 25
$n = 4$	34.8 ± 1.8	42.88 ± 0.09	4	3.92 ± 0.02	112 ± 6

Table A.3: Values measured and calculated for channel A of G115, showing that we are approximately at the same operating resistance as the bath temperature changes. Note that T_a is calculated using the G_{AB} fit.

T_b [mK]	T_a [mK]	I_{bias} [μA]	R_0 [$\text{m}\Omega$]	P_{TA} [pW]
10.90	11.02	170.0	134.79	4.23
13.40	13.52	168.0	132.20	4.19
17.80	17.91	164.5	132.43	4.01
26.30	26.40	151.0	132.43	3.38
31.80	31.87	129.5	132.20	2.49
36.25	36.29	102.0	131.51	1.55
40.80	40.81	50.0	132.42	0.37

Table A.4: The corresponding fitted values from the G_{TA} measurement for channel A of G115, where the errors come from propagating the 10% error in the power measurement. Included is estimated electron-phonon coupling constant Σ_{ep} for this array of W TESs.

Fit	G_{TA} [$\frac{\text{pW}}{\text{K}}$]	T_c [mK]	n	Σ_{ep} [$\frac{\text{GW}}{\text{K}^n \text{m}^3}$]
Any n	318 ± 41	42.00 ± 0.21	3.1 ± 0.5	—
$n = 5$	447 ± 20	41.64 ± 0.11	5	1.24 ± 0.06

confidence. However, the thermal conductance itself is not very sensitive to the power law exponent, nor is the absorber temperature, so we were still able to estimate the absorber heat capacity and the thermal time constant to a reasonable degree.

G_{TA} Measurement

To measure G_{TA} , we used channel A and the quick G_{TA} measurement method (i.e. estimating the TES resistance from the zero-frequency value of the circuit impedance). The data from this measurement are reported in Table A.3, and the error in P_{TA} is estimated to be 10% of the measured power. The sources of systematic error are the same as the discussion for the G_{AB} measurement, such that there may be a bias in these measurements due to the long thermalization times of the fridge and the high rate of high-energy events.

Using Eq. (A.7), we fit the power law with n varying and $n = 5$ fixed, the latter as expected for a W TES. In Fig. A.6, we show the results of the two fits as compared to the measured data, and report the various fit parameters in Table A.4. With this data, we took enough data points, but we instead have an anomalous power exponent of $n = 3.1 \pm 0.5$ when letting n vary. As this is far from expectation, and we know that we were having trouble with the long fridge thermalization times, we can explain this value as due to the




Figure A.6: The power flowing from TES to absorber at various absorber temperatures taken for G115. Also shown are the least squares fits with n varying and $n = 5$ fixed.




Figure A.7: The TES rectangles to be used for thermal conductance measurements, which are on a $1 \text{ cm}^2 \times 1 \text{ mm}$ Si substrate. These TESs scale by factors of 4 in area, while keeping the aspect ratio the same.

systematics from the measurement.

A.3 Thermal Conductance Measurements with QP.4

The entire QP.4 chip we will use for these measurements is shown in Fig. 5.9, where we use the TES rectangles, as pictured in Fig. A.7 for convenience. We will measure thermal

Table A.5: Values used to calculate the change in power flowing from the absorber to bath via the GE-7031 varnish at different bath temperatures, where TES200x50 was the absorber heater.

T_b [mK]	I_{bias} [μA]	P_{AB} [pW]
56	44.8	0.12
55.5	161.05	1.52
55	223.875	2.93
54.5	270.5	4.27
53	368	7.91
52	417.5	10.18
51	457.75	12.24
50	492.25	14.16

conductances of QP.4 in the following configurations: glued with GE-7031 varnish, resting on Cu and thermalized through the electronic wire bonds, and glued with rubber cement. We will also show the clean method of measuring G_{TA} via IV curves at different bath temperatures. These TES rectangles have dimensions $800\text{ }\mu\text{m} \times 200\text{ }\mu\text{m} \times 40\text{ nm}$, $400\text{ }\mu\text{m} \times 100\text{ }\mu\text{m} \times 40\text{ nm}$, $200\text{ }\mu\text{m} \times 50\text{ }\mu\text{m} \times 40\text{ nm}$ and $100\text{ }\mu\text{m} \times 25\text{ }\mu\text{m} \times 40\text{ nm}$, and will henceforth be referred to as TES800x200, TES400x100, TES200x50, and TES100x25, respectively. Each of the following measurements were done on a QP.4 chip from the same wafer, but each are different chips, thus giving variation on the various values of G_{TA} that we will measure for different TESs with the same rectangular dimensions.

A.3.1 GE-7031 Varnish

GE-7031 varnish is a common polymer-based adhesive for cryogenic applications. In our lab it has been frequently used to glue 1 cm^2 substrates to our Cu housing. Because the GE-7031 varnish is dissimilar to the semiconductor substrate, just as the case with the cirlex clamps, the thermal conductance at the interface will be due to the mismatch of the elastic constants [95], giving an expected thermal conductance temperature dependence of $G \propto T^3$ (i.e. the power law exponent of the heat flow is expected to be $n = 4$). Practically, we have had some trouble in the past with the GE-7031 varnish thermal contraction being large enough that our 1 cm^2 chips would frequently pop off of the Cu housing, and, depending on the device, the substrate may no longer have a good thermal connection to bath.

G_{AB} Measurement

To measure G_{AB} , we use TES200x50 as the heater and TES100x25 as the thermometer. In addition to the $5\text{ m}\Omega$ shunt resistor, the normal resistance of TES200x50 in this measurement is $413\text{ m}\Omega$, with a parasitic resistance of $2.4\text{ m}\Omega$. The data from this measurement are




Figure A.8: The power flowing from absorber to bath through the GE-7031 varnish at various bath temperatures taken for TES200x50. Also shown are the least squares fits with n varying and $n = 4$ fixed.

Table A.6: Values from the least-squares fits to the power flowing from absorber to bath through GE-7031 varnish. Included in the table are the estimated heat capacity via the low-temperature Debye model using Eq. (A.5), as well as the related thermal time constant τ .

Fit	$G_{AB} \left[\frac{\text{nW}}{\text{K}} \right]$	$T_a \text{ [mK]}$	n	$C^{th} \left[\frac{\text{pJ}}{\text{K}} \right]$	$\tau \text{ [ms]}$
Any n	2.89 ± 0.22	56.04 ± 0.01	4.9 ± 2.3	10.572 ± 0.003	3.65 ± 0.27
$n = 4$	2.83 ± 0.11	56.04 ± 0.01	4	10.572 ± 0.003	3.74 ± 0.15




Figure A.9: For the TESs on the chip thermalized through GE-7031 varnish: (Left) The power flowing from TES to absorber for TES100x25 as a function of absorber temperature. (Right) The power flowing from TES to absorber for TES200x50 as a function of absorber temperature. Both plots show the least squares fits with $n = 5$ fixed and n allowed to vary.

Table A.7: Values measured and calculated for TES100x25 on the GE-7031 varnished chip, showing that we are approximately at the same operating resistance as the bath temperature changes during the G_{TA} measurement.

T_b [mK]	T_a [mK]	I_{bias} [μA]	R_0 [$\text{m}\Omega$]	P_{TA} [fW]
30	30	12.45	264	13.88
40	40	11.725	269	12.09
45	45	10.975	262	10.86
50	50	9.275	260	7.81
52	52	7.8	252	5.69
54	54	6	255	3.33
55	55	4.875	261	2.15
56	56	3.175	261	0.91

reported in Table A.5. The error in P_{AB} is set to be 10% of the measured power, based on an estimated 10% uncertainty in R_{sh} . We fit Eq. (A.2) with n allowed to be any value, and $n = 4$ fixed as expected from the phonon mismatch between the absorber and the GE-7031 varnish. In Fig. A.8, we show the results of these two fits and compare to the measured data, with the various fit parameters reported in Table A.6.

Table A.8: Values measured and calculated for TES200x50 on the GE-7031 varnished chip, showing that we are approximately at the same operating resistance as the bath temperature changes during the G_{TA} measurement.

T_b [mK]	T_a [mK]	I_{bias} [μA]	R_0 [$\text{m}\Omega$]	P_{TA} [fW]
20	30	27.35	293	60.68
30	40	27.35	293	60.68
40	45	25.4	293	52.33
50	50	20.75	293	34.93
52	52	18.075	292	26.59
54	54	15.125	299	18.20
55	55	12.45	285	12.91
56	56	9.75	290	7.79

Table A.9: The corresponding fitted values from the G_{TA} measurement for TES100x25 and TES200x50 on the GE-7031 varnished chip. Included are estimated electron-phonon coupling constants Σ_{ep} for these W TESs.

	Fit	G_{TA} [$\frac{\text{pW}}{\text{K}}$]	T_c [mK]	n	Σ_{ep} [$\frac{\text{GW}}{\text{K}^n \text{m}^3}$]
TES100x25	Any n	1.47 ± 0.14	56.63 ± 0.10	5.8 ± 0.9	—
	$n = 5$	1.36 ± 0.06	56.69 ± 0.08	5	0.264 ± 0.013
TES200x50	Any n	6.37 ± 0.77	57.29 ± 0.24	6.0 ± 1.0	—
	$n = 5$	5.61 ± 0.28	$57.51 \pm 0.0.17$	5	0.256 ± 0.015

G_{TA} Measurement

For the G_{TA} measurements with GE-7031 varnish, we used the quick measurement method with TES100x25 and TES200x50. The data from this measurement are reported in Tables A.7 and A.8, where the error in P_{TA} is estimated to be 10% of the measured power. Using Eq. (A.7), we can fit the power law with n varying and $n = 5$ fixed, the latter as expected for a W TES. In Fig. A.9, we show the results of the fits as compared to the measured data, and report the various fit parameters in Table A.9 for the two TESs.

A.3.2 Al Wire Bonds

To study a configuration without glue, we ran a set of rectangles where the thermalization of the chip was achieved through the superconducting Al wire bonds that are used for electrical readout. In this scenario, we do not expect the thermal conductance to be from an interface, but instead dominated by the lattice wave (phonon) contribution. Because Al has a critical temperature $T_c = 1.2$ K, and we are operating these TESs far below this value (tens

Table A.10: Values used to calculate the change in power flowing from the absorber to bath via the superconducting Al wire bonds at different bath temperatures, where TES200x50 was the absorber heater.

T_b [mK]	I_{bias} [μA]	P_{AB} [fW]
50	44.8	113.7
49	53.7	163.4
48	60.05	204.4
45	75.25	320.9
42	85	409.5
38	94	500.8
32	103	601.2
26	107.5	654.9
20	110.5	692.0
15	114	736.5

Table A.11: Values from the least-squares fits to the power flowing from absorber to bath through the superconducting Al wire bonds. Included in the table are the estimated heat capacity via the low-temperature Debye model using Eq. (A.5), as well as the related thermal time constant τ .

Fit	G_{AB} [$\frac{\text{pW}}{\text{K}}$]	T_a [mK]	n	C^{th} [$\frac{\text{pJ}}{\text{K}}$]	τ [ms]
Any n	53.0 ± 7.4	52.32 ± 0.45	3.9 ± 0.7	8.60 ± 0.22	162 ± 26
$n = 4$	54.2 ± 2.4	52.26 ± 0.25	4	8.57 ± 0.12	156 ± 9

of mK), the electrons have no contribution to the thermal transport due to being bound in Cooper pairs, and the thermal conductance is expected to be $G \propto T^3$ from the phonon contribution to the thermal transport (i.e. we expect $n = 4$ for the P_{AB} power law).

G_{AB} Measurement

To measure G_{AB} , we use TES200x50 as the heater and TES100x25 as the thermometer. In addition to the $5\text{ m}\Omega$ shunt resistor, the normal resistance of TES200x50 in this measurement is $426\text{ m}\Omega$, with a parasitic resistance of $2.5\text{ m}\Omega$. The data from this measurement are reported in Table A.10, where error in P_{AB} is set to be 10% of the measured power, based on an estimated 10% uncertainty in R_{sh} . We then fit Eq. (A.2) with n allowed to be any value, and $n = 4$ fixed as expected from the phonon contribution to the thermal conductance. In Fig. A.10, we show the results of these two fits and compare to the measured data, with the various fit parameters reported in Table A.11.




Figure A.10: The power flowing from absorber to bath through the superconducting Al wire bonds at various bath temperatures taken for TES200x50. Also shown are the least squares fits with n varying and $n = 4$ fixed.

Table A.12: Values measured and calculated for TES100x25, showing that we are approximately at the same operating resistance as the bath temperature changes.

T_b [mK]	T_a [mK]	I_{bias} [μA]	R_0 [$\text{m}\Omega$]	P_{TA} [fW]
30	30.2	11.975	269	12.61
40	40.2	10.75	275	9.96
45	45.1	9.625	270	8.12
50	50.0	6.6	264	3.90
51	52.0	5.375	270	2.53
52	56.0	4.475	240	1.96

G_{TA} Measurement

For the G_{TA} measurement, we use the quick measurement method for TES100x25 and TES200x50. The data from these measurements are reported in Tables A.12 and A.13, and the error in P_{TA} is estimated to be 10% of the measured power. Using Eq. (A.7), we fit the power law with n varying and $n = 5$ fixed, the latter as expected for a W TES. In Fig. A.11, we show the results of the two fits as compared to the measured data, and report




Figure A.11: For the TESs on the chip thermalized through the superconducting Al wire bonds: (Left) The power flowing from TES to absorber for TES100x25 as a function of absorber temperature. (Right) The power flowing from TES to absorber for TES200x50 as a function of absorber temperature. Both plots show the least squares fits with $n = 5$ fixed and n allowed to vary.

Table A.13: Values measured and calculated for TES100x25, showing that we are approximately at the same operating resistance as the bath temperature changes.

T_b [mK]	T_a [mK]	I_{bias} [μA]	R_0 [$\text{m}\Omega$]	P_{TA} [fW]
30	30.2	27.35	252	12.61
40	40.2	25.15	249	9.96
45	45.1	22	252	8.12
50	50.0	17.825	256	3.90
52	52.0	14.4	261	2.53
54	56.0	8.075	247	1.96

Table A.14: The corresponding fitted values from the G_{TA} measurement for TES100x25 and TES200x50 on the chip thermalizing through the Al wire bonds. Included are estimated electron-phonon coupling constants Σ_{ep} for these W TESs.

	Fit	G_{TA} [$\frac{\text{pW}}{\text{K}}$]	T_c [mK]	n	Σ_{ep} [$\frac{\text{GW}}{\text{K}^n \text{m}^3}$]
TES100x25	Any n	1.16 ± 0.21	53.71 ± 0.40	4.5 ± 1.3	—
	$n = 5$	1.23 ± 0.08	53.59 ± 0.21	5	0.299 ± 0.022
TES200x50	Any n	8.03 ± 1.00	54.92 ± 0.14	6.2 ± 1.3	—
	$n = 5$	7.11 ± 0.37	55.03 ± 0.11	5	0.387 ± 0.022

Table A.15: Values used to calculate the change in power flowing from the absorber to bath via rubber cement at different bath temperatures, where TES200x50 was the absorber heater.

T_b [mK]	I_{bias} [μA]	P_{AB} [pW]
52	501.2	14.2
51.5	734	30.5
51	894	45.2
50.5	1030	60.0
50	1140	73.5
49	1310	97.0
48	1440	117.2
46	1640	152.1
43	1830	189.4
40	1960	217.2

Table A.16: Values from the least-squares fits to the power flowing from absorber to bath through the rubber cement. Included in the table are the estimated heat capacity via the low-temperature Debye model using Eq. (A.5), as well as the related thermal time constant τ .

Fit	G_{AB} [$\frac{\text{nW}}{\text{K}}$]	T_a [mK]	n	C^{th} [$\frac{\text{pJ}}{\text{K}}$]	τ [ms]
Any n	34.8 ± 52.4	52.42 ± 0.06	7.3 ± 1.5	8.65 ± 0.03	0.248 ± 0.022
$n = 4$	28.4 ± 1.2	52.55 ± 0.06	4	8.72 ± 0.03	0.306 ± 0.014

the various fit parameters in Table A.14.

A.3.3 Rubber Cement

As the GE-7031 varnish had problems with glued chips popping off after cooling, it was recommended by Aritoki Suzuki that we try rubber cement (Elmer's) as an alternative adhesive, as the Suzuki group had never experienced this in their lab. As with GE-7031 varnish and the cirlex clamps, rubber cement is dissimilar to the semiconductor substrate, and the thermal conductance at the interface will be due to the mismatch of the elastic constants [95], giving an expected thermal conductance temperature dependence of $G \propto T^3$.

G_{AB} Measurement

To measure G_{AB} , we use TES200x50 as the heater and TES100x25 as the thermometer. In addition to the $5 \text{ m}\Omega$ shunt resistor, the normal resistance of TES200x50 in this measurement




Figure A.12: The power flowing from absorber to bath through the rubber cement at various bath temperatures taken for TES200x50. Also shown are the least squares fits with n varying and $n = 4$ fixed.

is $423\text{ m}\Omega$, with a parasitic resistance of $4.3\text{ m}\Omega$. The data from this measurement are reported in Table A.15. The error in P_{AB} is set to be 10% of the measured power, based on an estimated 10% uncertainty in R_{sh} . We then fit Eq. (A.2) with n allowed to be any value, and $n = 4$ fixed as expected from the phonon mismatch between the absorber and the rubber cement. In Fig. A.12, we show the results of these two fits and compare to the measured data, with the various fit parameters reported in Table A.16.

A.3.4 Thermal Conductance with an IV Sweep

For each of the previous sections, we showed examples of the quick G_{TA} measurement method, which is fast, but could be subject to systematic errors as noted in the explanation of the method. In this section, we show examples of clean G_{TA} measurements via IV sweeps at various bath temperatures for TES800x200, TES400x100, and TES200x50.

G_{TA} Measurement

In Figs. A.13–A.15, we show the IV , RV , and PV curves for each of the TESs, as well as the G_{TA} fit. Note that we chose an operating resistance of $R_0 = 150\text{ m}\Omega$ and interpolated between data points to ensure estimate the bias power of the TES at this value. We fit Eq. (A.7)




Figure A.13: Example of the clean G_{TA} measurement with TES200x50. (Top left) TES current as a function of bias current. (Top right) TES resistance as a function of bias current. (Bottom left) TES bias power as a function of bias current. For these three plots, the superconducting regions of the curves were not plotted for cosmetic reasons, and the lines connecting data points are simple linear interpolations. (Bottom right) Fit to G_{TA} with n varying and $n = 5$ fixed.

Table A.17: The corresponding fitted values from the G_{TA} measurement for TES200x50, TES400x100, and TES800x200 on the GE-7031 varnished chip used as an example of a clean G_{TA} measurement. Included are estimated electron-phonon coupling constants Σ_{ep} for these W TESs.

	Fit	G_{TA} [$\frac{\text{pW}}{\text{K}}$]	T_c [mK]	n	Σ_{ep} [$\frac{\text{GW}}{\text{K}^n \text{m}^3}$]
TES200x50	Any n	12.8 ± 3.7	60.3 ± 1.2	6.2 ± 1.9	—
	$n = 5$	10.5 ± 0.6	61.1 ± 0.7	5	0.379 ± 0.034
TES400x100	Any n	45.1 ± 13.1	60.7 ± 1.2	6.0 ± 1.9	—
	$n = 5$	37.9 ± 1.9	61.5 ± 0.7	5	0.331 ± 0.028
TES800x200	Any n	145 ± 42	60.8 ± 1.3	5.9 ± 1.8	—
	$n = 5$	123.8 ± 6.2	61.5 ± 0.7	5	0.270 ± 0.023




Figure A.14: Example of the clean G_{TA} measurement with TES400x100. (Top left) TES current as a function of bias current. (Top right) TES resistance as a function of bias current. (Bottom left) TES bias power as a function of bias current. For these three plots, the superconducting regions of the curves were not plotted for cosmetic reasons, and the lines connecting data points are simple linear interpolations. (Bottom right) Fit to G_{TA} with n varying and $n = 5$ fixed.

with n varying and $n = 5$ fixed, and report the various fit parameters in Table A.17. The power law curves in these measurements follow $n = 5$ much more convincingly in these fits, due to the lower amount of systematics as opposed to the quick measurement.

A.4 Thermal Conductance Measurements with a SPICE MELANGE Detector

In Fig. A.16, we show the detector setup from one of our runs in the dilution refrigerator at UC Berkeley, where we will use the 4% detector with disconnected fins to measure the thermal conductance from absorber to bath of the gold pad. Both of these chips are resting on the gold-plated Cu plate, and each has a TES rectangle of dimensions $295 \mu\text{m} \times 40 \mu\text{m} \times 40 \text{nm}$, henceforth referred to as TES295x40. Because the TES rectangle on the connected fin device




Figure A.15: Example of the clean G_{TA} measurement with TES800x200. (Top left) TES current as a function of bias current. (Top right) TES resistance as a function of bias current. (Bottom left) TES bias power as a function of bias current. For these three plots, the superconducting regions of the curves were not plotted for cosmetic reasons, and the lines connecting data points are simple linear interpolations. (Bottom right) Fit to G_{TA} with n varying and $n = 5$ fixed.

had a short, we used the disconnected fin device instead.

A.4.1 Au Pad

On the Si substrate of the detector, a gold pad was deposited with dimensions $300\text{ }\mu\text{m} \times 300\text{ }\mu\text{m} \times 750\text{ nm}$ for the purpose of thermalization through a gold wire. The sizing of the pad was decided by C. W. Fink, and we will use it as a data point for the thermal conductance of a thin gold film. As noted in Refs. [225, 226], the power law exponent for a thin gold film can be expected to be either $n = 5$ for an ordered metal [94, 227, 228], or $n = 6$ for a disordered metal [229, 230].




Figure A.16: (Figure from Ref. [147]) The right $1\text{ cm}^2 \times 1\text{ mm}$ chip was used to measure the thermal conductance from absorber to bath for a Au wire bonding pad, which is a 4% coverage design with “disconnected” Al fins (cyan inset). The TES295x40 on the chip is circled in yellow. These QET-based detectors were designed by C. W. Fink and are studied in detail in his thesis [147].

Table A.18: Values used to calculate the change in power flowing from the absorber to bath via the Au wire bond pad at different bath temperatures, where TES295x40 was the absorber heater.

T_b [mK]	I_{bias} [μA]	P_{AB} [pW]
47	133	1.23
46	179	2.23
44	232	3.74
42	266	4.92
39	300	6.26
36	321	7.17
30	347	8.38
22	362	9.12
15	368	9.42
10	370	9.52




Figure A.17: The power flowing from absorber to bath through the Au wire bond pad at various bath temperatures taken for TES295x40. Also shown are the least squares fits with n varying, $n = 5$ fixed, $n = 6$ fixed.

Table A.19: Values from the least-squares fits to the power flowing from absorber to bath through the Au wire bond pad. Included in the table are the estimated heat capacity via the low-temperature Debye model using Eq. (A.5), the related thermal time constant τ , and the electron-phonon coupling constant measured for the Au pad.

Fit	G_{AB} [$\frac{nW}{K}$]	T_a [mK]	n	C^{th} [$\frac{pJ}{K}$]	τ [ms]	Σ_{ep} [$\frac{GW}{K^n m^3}$]
Any n	1.04 ± 0.12	48.28 ± 0.22	5.4 ± 0.8	6.76 ± 0.09	6.45 ± 0.84	—
$n = 5$	0.98 ± 0.04	48.38 ± 0.15	5	6.80 ± 0.06	6.93 ± 0.32	0.531 ± 0.026
$n = 6$	1.12 ± 0.04	48.7 ± 0.13	6	6.71 ± 0.0	5.97 ± 0.26	10.7 ± 0.5

G_{AB} Measurement

To measure G_{AB} , we use TES295x40 as the heater and array of disconnected fin QETs as the thermometer. In addition to the 5 m Ω shunt resistor, the normal resistance of TES295x40 in this measurement is 344 m Ω , with a parasitic resistance of 2.6 m Ω . The data from this measurement are reported in Table A.18, where the error in P_{AB} is set to be 10% of the measured power, based on an estimated 10% uncertainty in R_{sh} . We then fit Eq. (A.2) with n allowed to be any value, and $n = 5$ and $n = 6$ fixed as expected from ordered and disordered metals, respectively. In Fig. A.17, we show the results of these fits and compare to the measured data, with the various fit parameters reported in Table A.19. Both $n = 5$ and $n = 6$ seem to reasonably describe the power law, so we do not make a claim as to which is correct, but instead report the values for both.

A.5 Summary of Measurements

In this appendix, we discussed methods for measuring the thermal conductances from TES to absorber and absorber to bath and used the methods for a variety of different measurements, where the substrate holding scheme varied significantly. The various values from that we fit are summarized in Table A.20 for the G_{AB} fits and Table A.21 for the G_{TA} fits.

For the Σ_{ep} values, previous measurements have seen $\Sigma_{ep, Au} = 2.4 \text{ GW}/(\text{K}^5 \text{ m}^3)$ [231] and $\Sigma_{ep, W} = 0.3 \text{ GW}/(\text{K}^5 \text{ m}^3)$ [232] for Au and W, respectively. Our measurements for W are completely consistent with that value, while the $n = 5$ value for Au is a factor of 5 smaller than previously measured. It is unknown why this is the case. For the $n = 6$ power law for Au (i.e. a disordered film), it has been measured as $\Sigma'_{ep, Au} = 60 \text{ GW}/(\text{K}^6 \text{ m}^3)$ [225, 226], a value which itself depends on the level of disorder. If this were the case, then the $n = 6$ value we measured is a factor of 6 lower, which would imply that the level of disorder is different for our film (i.e. less disordered). This could be the case, as our Au film is about a factor of 10 thicker (750 nm vs. 57 nm), and the boundary scattering effects would be smaller. Another possibility could be that the Au pad itself is not well-connected to the substrate, such that there is an extra interfacial thermal resistance that we cannot measure. More

Table A.20: Summary table of the G_{AB} measurements in this appendix. We have added an asterisk on the cirlex clamps measurement due to the poor quality of that dataset.

Thermal Bath Connection	Fixed n	$G_{AB} \left[\frac{\text{nW}}{\text{K}} \right]$
Cirlex Clamps*	$n = 4$	34.8 ± 1.8
GE-7031 Varnish	$n = 4$	2.83 ± 0.11
Al Wire	$n = 4$	0.054 ± 0.002
Rubber Cement	$n = 4$	28.4 ± 1.2
Au Pad	$n = 5$	0.98 ± 0.04

Table A.21: Summary table of the Σ_{ep} measurements in this appendix. We have added an asterisk on the G115 measurement due to the poor quality of that dataset.

Section	Device	Fixed n	$\Sigma_{ep} \left[\frac{\text{GW}}{\text{K}^n \text{m}^3} \right]$
A.2	W G115*	$n = 5$	1.24 ± 0.06
A.3.1	W TES100x25	$n = 5$	0.264 ± 0.013
	W TES200x50	$n = 5$	0.256 ± 0.015
A.3.2	W TES100x25	$n = 5$	0.299 ± 0.022
	W TES200x50	$n = 5$	0.387 ± 0.022
A.3.4	W TES200x50	$n = 5$	0.379 ± 0.034
	W TES400x100	$n = 5$	0.331 ± 0.028
	W TES800x200	$n = 5$	0.270 ± 0.023
A.4.1	MELANGE Au Pad	$n = 5$	0.531 ± 0.026
	MELANGE Au Pad	$n = 6$	10.7 ± 0.5

thermal conductance measurements of the Au pad system would be useful to understand the difference.

Appendix B

Reduction of Vibrational Noise

The decoupling of sensitive detectors or instruments from environmental vibrations has been pursued by many cryogenic experiments, usually using a three-dimensional elastic pendulum [208] for passive vibration isolation [209–215]. These vibrations can provide additional in the experimental, reducing sensitivity of an experiment. This is especially important in rare event searches that work towards the lowest possible noise environments. For the CPD studied in the main body of this thesis, the excess signals observed within the DM ROI could be removed with a stress-free holding scheme, which would then be sensitive to vibrations, thus motivating the design and fabrication of an internal vibration decoupler. A common source of vibrations in cryogenic experiments that use a cryogen-free dilution refrigerator is the pulse-tube cryocooler. The pulse-tube cryocooler (PT) is critical to achieve sub-10 mK base temperatures and relies on a oscillating compression system to provide cooling power for the dilution refrigerator. The operating frequency of these oscillations is low, with 1.4 Hz as a common frequency. Thus, the goal is to have a system that works as a low-pass filter with a fundamental frequency that is as low as possible. Practically, this will be constrained by the dilution refrigerator dimensions, as we will discuss later in this appendix.

For the system used in the Pyle Lab, we expand upon the designs of two different vibration isolation systems: one developed by R. Maisonobe *et al.* [209] and one developed by L. Gottardi *et al.* [210]. We will call each of these vibration decouplers the “Maisonobe system” and the “Gottardi system,” respectively.

As shown in Fig. B.1, the Maisonobe system is based on a single elastic pendulum placed in the center of the cryostat plate. Their design is simple—they use a single spring pendulum from the 1 K to 10 mK stage. The pendulum ultimately giving the best performance is based on a nylon wire and a single stainless steel spring, together giving fundamental frequencies of 1.8 Hz in the vertical direction and 0.8 Hz in the radial direction. This system allowed the detectors studied to achieve baseline energy resolutions between a factor 2–40 better than without any vibration decoupling.

The Gottardi system is based on the taber vibration isolator proposed in Ref. [211]. This system uses a kevlar wire which is clamped in a triangular configuration, as shown in Fig. B.2, which effectively acts as three springs symmetrically placed at 120° intervals. This allows the




Figure B.1: (Figure from Ref. [209]) Design of the Maisonobe system for passive vibration isolation based on an elastic pendulum where motion is driven by perturbations of the plate.

design to attenuate vibrations along more degrees of freedom than the Maisonobe system. The design also has two stages, as opposed to one, which increases the attenuation power above the fundamental frequencies with the cost of increasing them by a small amount (as opposed to a one-stage system). For their system, the fundamental frequencies were shown to be on the order of tens of Hz. When adjusting this design to the approximate dimensions of the Pyle Lab fridge it was found that the highest frequencies were still about 10 Hz, thus we decided not to use this design.

For the Pyle Lab, we considered a model similar to as depicted in Fig. B.3. This system can be thought of as an in-between for both the Maisonobe [209] and Gottardi [210] systems. With the springs, we should be able to achieve lower fundamental frequencies with the system (as well as in a two-stage version), as well as have attenuation along more degrees of freedom than the Maisonobe system. In order to understand the expected fundamental frequencies in this proposed system, we will derive them using Lagrangian mechanics.




Figure B.2: (Figure from Ref. [210]) (a) Design of the Gottardi system, which uses kevlar wires to act as a elastic double pendulum. (b) The clamping scheme showing how the kevlar wire is installed

B.1 Derivation of Fundamental Frequencies

In deriving the fundamental frequencies of the vibration decoupler, we will consider two versions of our design: one with a single vibrational stage, and one with two vibrational stages.

B.1.1 One Stage

The one-stage derivation is as follows. First, we calculate the equilibrium state with the usual sum of forces

$$-mg + 3k(l_{eq} - l_0) = 0, \quad (\text{B.1})$$




Figure B.3: Our proposed vibration decoupling system with three springs spaced equally apart. Each spring has the same equilibrium length l and the same spring constant k . Note that the coordinate axis we will use is centered on the disk's center of mass, but is translated to the right to reduce overlapping within the figure.

where we assume that the system is such that the center of mass is at the disk's center. If we compare to the Maisonobe system, we see that the change in length of each spring will be a third of that in the Maisonobe system, assuming the same spring constants.

In order to describe how the lengths of the springs change under various movements of the disk, we need to first define the coordinate system. Specifically, we will set the origin of the coordinate system as the center of the tops of the three springs. In other words, the center-of-mass of the disk at equilibrium is set to be

$$\mathbf{x}_{cm} = x\hat{\mathbf{x}} + y\hat{\mathbf{y}} + (z - l_{eq})\hat{\mathbf{z}}, \quad (\text{B.2})$$

where l_{eq} is the equilibrium length of the springs. We also define the x -axis as some angle θ from spring 1, such that spring 2 and spring 3 are placed $\theta + \frac{2\pi}{3}$ and $\theta + \frac{4\pi}{3}$ away from the x -axis. In other words, we define the horizontal axes to be arbitrary. Using our coordinate system, the locations of the top of the three springs are defined as

$$\mathbf{R}_1 = R [\cos(\theta)\hat{\mathbf{x}} + \sin(\theta)\hat{\mathbf{y}}], \quad (\text{B.3})$$

$$\mathbf{R}_2 = R \left[\cos\left(\theta + \frac{2\pi}{3}\right)\hat{\mathbf{x}} + \sin\left(\theta + \frac{2\pi}{3}\right)\hat{\mathbf{y}} \right], \quad (\text{B.4})$$

$$\mathbf{R}_3 = R \left[\cos\left(\theta + \frac{4\pi}{3}\right)\hat{\mathbf{x}} + \sin\left(\theta + \frac{4\pi}{3}\right)\hat{\mathbf{y}} \right]. \quad (\text{B.5})$$

In order to find the vector of each spring, we can follow the vector path of the top of a spring $\mathbf{R}_i$, to the origin, to the location of the center-of-mass (including its displacement), to the bottom of the spring (with some rotation of the disk). Thus, for spring i , its vector is

$$\mathbf{l}_i = \mathbf{x}_{cm} + \mathbf{R}'_i - \mathbf{R}_i, \quad (\text{B.6})$$

where $\mathbf{R}'_i$ is the rotated $\mathbf{R}_i$.

Because the three springs are placed on the edges of the disk, we expect six degrees of freedom, also seen in the Gottardi system. There will be three translational and three rotational modes. The most general way to calculate the fundamental frequencies of these modes would be to take into account all possible rotations of the disk. However, this quickly becomes a mess of Euler angles and is unwieldy to do by hand.

Thus, to simplify our efforts, we will first calculate the modes of the system when we allow rotations of the disk about the horizontal axis, and then calculate the modes of the system when we allow rotations of the disk about the vertical axis. We note that there is an assumption here that the horizontal rotational modes and vertical rotational modes do not couple, which should be a good assumption simply based on intuition of the system.

Rotations about the horizontal—To find all of the fundamental frequencies of the system when allowing for rotations about the horizontal, we will begin by defining the kinetic energy,

potential energy, and Lagrangian equations as

$$T = \frac{1}{2}m(\dot{x}^2 + \dot{y}^2 + \dot{z}^2) + \frac{1}{2}I_x\dot{\phi}^2, \quad (\text{B.7})$$

$$V = \sum_{i=1}^3 \left[\frac{1}{2}k(l_i - l_0)^2 \right] + mgz, \quad (\text{B.8})$$

$$L = T - V. \quad (\text{B.9})$$

The l_i terms are nonlinear, so the Euler-Lagrange equations that we would get out of this Lagrangian would not be analytically solvable. However, we can Taylor expand the potential energy around the equilibrium point to second order to estimate the natural frequencies for small oscillations. As we are looking at rotations of the disk about the horizontal by some angle ϕ , we can relate $\mathbf{R}'_i$ to $\mathbf{R}_i$ by

$$\mathbf{R}'_i = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \phi & -\sin \phi \\ 0 & \sin \phi & \cos \phi \end{bmatrix} \mathbf{R}_i. \quad (\text{B.10})$$

For the zeroth order term of the Taylor expansion, the potential energy is simply $V_0 = \frac{3}{2}k(l_{eq} - l_0)^2$, which is the potential energy stored in each spring at equilibrium. For the first order terms, it can be shown that the first derivative of the potential energy with respect to each variable is zero when evaluated at equilibrium (the point at which we are expanding). For the second order terms, it can be shown that the cross-terms all evaluate to zero at equilibrium, while the second derivatives with respect to each variable are nonzero. Specifically, the potential energy can be approximated as

$$\begin{aligned} V \approx V_0 + \frac{3}{2}k \left(1 - \frac{l_0}{l_{eq}} \right) x^2 + \frac{3}{2}k \left(1 - \frac{l_0}{l_{eq}} \right) y^2 \\ + \frac{3}{2}kz^2 + \frac{3}{4}kR^2\phi^2. \end{aligned} \quad (\text{B.11})$$

Using our equilibrium relation, we can simplify this to

$$V \approx V_0 + \frac{1}{2} \frac{mg}{l_{eq}} x^2 + \frac{1}{2} \frac{mg}{l_{eq}} y^2 + \frac{3}{2}kz^2 + \frac{3}{4}kR^2\phi^2. \quad (\text{B.12})$$

It is clear that the Euler-Lagrange equations will completely decouple, and we do not have any mixing of modes when Taylor expanding up to second order. Applying the Euler-

Lagrange equations, we find that the decoupled equations of motion are

$$\ddot{x} + \frac{g}{l_{eq}}x = 0, \quad (\text{B.13})$$

$$\ddot{y} + \frac{g}{l_{eq}}y = 0, \quad (\text{B.14})$$

$$\ddot{z} + \frac{3k}{m}z = 0, \quad (\text{B.15})$$

$$\ddot{\phi} + \frac{3kR^2}{2I_x}\phi = 0. \quad (\text{B.16})$$

Thus, we have that the fundamental frequencies are

$$f_{x,y} = \frac{1}{2\pi} \sqrt{\frac{g}{l_{eq}}}, \quad (\text{B.17})$$

$$f_z = \frac{1}{2\pi} \sqrt{\frac{3k}{m}}, \quad (\text{B.18})$$

$$f_{x,y,rot.} = \frac{1}{2\pi} \sqrt{\frac{3kR^2}{2I_x}}. \quad (\text{B.19})$$

Note that the x and y translational modes have the same fundamental frequencies due to the symmetry of system. Because we set up the system to have arbitrary x and y axes, the fundamental frequencies about the x and y axes are identical as well.

Rotations about the vertical—To find all of the fundamental frequencies of the system when allowing for rotations about the vertical by some angle ψ , we will again define the kinetic energy, potential energy, and Lagrangian equations as

$$T = \frac{1}{2}m(\dot{x}^2 + \dot{y}^2 + \dot{z}^2) + \frac{1}{2}I_z\dot{\psi}^2, \quad (\text{B.20})$$

$$V = \sum_{i=1}^3 \left[\frac{1}{2}k(l_i - l_0)^2 \right] + mgz, \quad (\text{B.21})$$

$$L = T - V. \quad (\text{B.22})$$

We will again Taylor expand the potential energy, where the relation between $\mathbf{R}'_i$ and $\mathbf{R}_i$ is instead

$$\mathbf{R}'_i = \begin{bmatrix} \cos \psi & -\sin \psi & 0 \\ \sin \psi & \cos \psi & 0 \\ 0 & 0 & 1 \end{bmatrix} \mathbf{R}_i. \quad (\text{B.23})$$

For the zeroth order term of the Taylor expansion, the potential energy is again $V_0 = \frac{3}{2}k(l_{eq} - l_0)^2$, as expected. For the first order terms, it can be shown that the first derivative of the potential energy with respect to each variable is zero when evaluated at equilibrium.

For the second order terms, it can be shown that the cross-terms all evaluate to zero at equilibrium, while the second derivatives with respect to each variable are nonzero. Specifically, the potential energy can be approximated as

$$V \approx V_0 + \frac{3}{2}k \left(1 - \frac{l_0}{l_{eq}}\right) x^2 + \frac{3}{2}k \left(1 - \frac{l_0}{l_{eq}}\right) y^2 + \frac{3}{2}kz^2 + \frac{3}{2}k \left(1 - \frac{l_0}{l_{eq}}\right) R^2\psi^2. \quad (\text{B.24})$$

Using our equilibrium relation, we can simplify this to

$$V \approx V_0 + \frac{1}{2} \frac{mg}{l_{eq}} x^2 + \frac{1}{2} \frac{mg}{l_{eq}} y^2 + \frac{3}{2} kz^2 + \frac{1}{2} \frac{mgR^2}{l_{eq}} \psi^2. \quad (\text{B.25})$$

We again have that the Euler-Lagrange equations will completely decouple, and we do not have any mixing of modes when Taylor expanding up to second order. Applying the Euler-Lagrange equations, we find that the decoupled equations of motion are

$$\ddot{x} + \frac{g}{l_{eq}} x = 0, \quad (\text{B.26})$$

$$\ddot{y} + \frac{g}{l_{eq}} y = 0, \quad (\text{B.27})$$

$$\ddot{z} + \frac{3k}{m} z = 0, \quad (\text{B.28})$$

$$\ddot{\psi} + \frac{mgR^2}{I_z l_{eq}} \psi = 0. \quad (\text{B.29})$$

Thus, the translational modes match those that we derived for when studying the system under horizontal rotations. Thus, we can summarize all of the different fundamental frequencies as

$$f_{x,y} = \frac{1}{2\pi} \sqrt{\frac{g}{l_{eq}}}, \quad (\text{B.30})$$

$$f_z = \frac{1}{2\pi} \sqrt{\frac{3k}{m}}, \quad (\text{B.31})$$

$$f_{x,y,rot.} = \frac{1}{2\pi} \sqrt{\frac{3kR^2}{2I_{x,y}}}, \quad (\text{B.32})$$

$$f_{z,rot.} = \frac{1}{2\pi} \sqrt{\frac{mgR^2}{I_z l_{eq}}}, \quad (\text{B.33})$$

where $f_{z,rot.}$ follows from solving the differential equation describing ψ .

Thus, we have that, using the Taylor expansion of the potential energy, we can analytically solve the system using the Lagrangian formalism. The next Pyle Lab setup to study will be the two-stage vibration decoupler.

B.1.2 Two Stages

As opposed to the derivation of the one-stage system, we will instead go through the various intuitive vibrational modes for the two-stage system case by case. This simplifies the derivation as we are now assuming that the modes decouple, rather than proving this by checking all of the various derivatives. We will also assume that the spring constants for each spring within a stage are the same, but different between the two stages, that the springs for the two stages have different equilibrium lengths, and that the masses of the two stages are also different (but the same disk radius). In Fig. B.4, we show the proposed system, as well as define the coordinate system used.

Case 1—Let's start with vertical translations of the two-stage system. The kinetic energy, potential energy, and Lagrangian can be written as:

$$T = \frac{1}{2}m_1\dot{z}_1^2 + \frac{1}{2}m_2\dot{z}_2^2, \quad (\text{B.34})$$

$$V = \frac{3}{2}k_1z_1^2 + \frac{3}{2}k_2(z_2 - z_1)^2 + m_1gz_1 + m_2gz_2, \quad (\text{B.35})$$

$$L = \frac{1}{2}m_1\dot{z}_1^2 + \frac{1}{2}m_2\dot{z}_2^2 - \frac{3}{2}k_1z_1^2 - \frac{3}{2}k_2(z_2 - z_1)^2 - m_1gz_1 - m_2gz_2, \quad (\text{B.36})$$

where each spring in a stage has identical potential energies (hence the factors of 3). After carrying out the Euler-Lagrange equations and simplifying, we have two coupled differential equations:

$$m_1\ddot{z}_1 + 3(k_1 + k_2)z_1 - 3k_2z_2 = 0, \quad (\text{B.37})$$

$$m_2\ddot{z}_2 - 3k_2z_1 + 3k_2z_2 = 0, \quad (\text{B.38})$$

where we have again neglected constant values (as they can be removed by a change in coordinate system). After solving this system for the eigenvalues, we have that the natural frequencies are

$$\omega^2 = \frac{3[(k_1 + k_2)m_2 + k_2m_1]}{2m_1m_2} \left\{ 1 \pm \sqrt{1 - \frac{4m_1m_2k_1k_2}{[(k_1 + k_2)m_2 + k_2m_1]^2}} \right\}. \quad (\text{B.39})$$

Case 2—Next, let's look at horizontal translations of the disks. This can be thought of as the double pendulum modes for the disks' centers of mass. We start by showing the kinetic and potential energy equations in x and y coordinates:

$$T = \frac{1}{2}m_1(\dot{x}_1 + \dot{y}_1)^2 + \frac{1}{2}m_2(\dot{x}_2 + \dot{y}_2)^2, \quad (\text{B.40})$$

$$V = m_1gy_1 + m_2gy_2. \quad (\text{B.41})$$




Figure B.4: Our proposed vibration decoupling system with three springs spaced equally apart for two different stages. Each spring stage 1 and stage 2 has the same equilibrium length l_1 , l_2 and the same spring constant k_1 , k_2 for those stages, respectively. Note that the coordinate axis we will use is centered on each disk's center of mass, but is translated to the right to reduce overlapping within the figure.

Since this will be oscillations due to the angles that the disk centers of mass make with the vertical, we need to change our variables to be in terms of these angles for each center of mass (α_1 and α_2). That is, we have that

$$x_1 = l_1 \sin \alpha_1, \quad (\text{B.42})$$

$$\dot{x}_1 = l_1 \dot{\alpha}_1 \cos \alpha_1, \quad (\text{B.43})$$

$$y_1 = -l_1 \cos \alpha_1, \quad (\text{B.44})$$

$$\dot{y}_1 = l_1 \dot{\alpha}_1 \sin \alpha_1, \quad (\text{B.45})$$

$$x_2 = l_1 \sin \alpha_1 + l_2 \sin \alpha_2, \quad (\text{B.46})$$

$$\dot{x}_2 = l_1 \dot{\alpha}_1 \cos \alpha_1 + l_2 \dot{\alpha}_2 \cos \alpha_2, \quad (\text{B.47})$$

$$y_2 = -l_1 \cos \alpha_1 - l_2 \cos \alpha_2, \quad (\text{B.48})$$

$$\dot{y}_2 = l_1 \dot{\alpha}_1 \sin \alpha_1 + l_2 \dot{\alpha}_2 \sin \alpha_2. \quad (\text{B.49})$$

We can plug all of these relations into the kinetic and potential energy equations, which gives us

$$T = \frac{1}{2}(m_1 + m_2)l_1^2 \dot{\alpha}_1^2 + m_2 l_1 l_2 \dot{\alpha}_1 \dot{\alpha}_2 \cos(\alpha_1 - \alpha_2) + \frac{1}{2}m_2 l_2^2 \dot{\alpha}_2^2, \quad (\text{B.50})$$

$$V = -(m_1 + m_2)gl_1 \cos \alpha_1 - m_2 gl_2 \cos \alpha_2. \quad (\text{B.51})$$

Using the Euler-Lagrange equations and the small angle approximation, one can find that this gives us two coupled differential equations:

$$(m_1 + m_2)l_1^2 \ddot{\alpha}_1 + m_2 l_1 l_2 \ddot{\alpha}_2 + (m_1 + m_2)gl_1 \alpha_1 = 0, \quad (\text{B.52})$$

$$m_2 l_2^2 \ddot{\alpha}_2 + m_2 l_1 l_2 \ddot{\alpha}_1 + m_2 gl_2 \alpha_2 = 0. \quad (\text{B.53})$$

After solving this system of equations for the eigenfrequencies, we have that the natural frequencies are

$$\omega^2 = \frac{(m_1 + m_2)(l_1 + l_2)g}{2m_1 l_1 l_2} \left\{ 1 \pm \sqrt{1 - \frac{4m_1 l_1 l_2}{(m_1 + m_2)(l_1 + l_2)^2}} \right\}. \quad (\text{B.54})$$

Case 3—Next, we look at the eigenmodes for rotations about a horizontal axis, where each disk rotates by an angle of α_1 and α_2 , respectively. This is shown in Fig. B.5 In this case, we will need to keep track of the locations of the springs along the disk, remembering that each spring is vertically aligned. In this case, we again remind the reader of the below relations for each of the spring locations, where the same holds true for each stage.

$$\mathbf{R}_1 = R [\cos(\theta) \hat{\mathbf{x}} + \sin(\theta) \hat{\mathbf{y}}], \quad (\text{B.55})$$

$$\mathbf{R}_2 = R \left[\cos\left(\theta + \frac{2\pi}{3}\right) \hat{\mathbf{x}} + \sin\left(\theta + \frac{2\pi}{3}\right) \hat{\mathbf{y}} \right], \quad (\text{B.56})$$

$$\mathbf{R}_3 = R \left[\cos\left(\theta + \frac{4\pi}{3}\right) \hat{\mathbf{x}} + \sin\left(\theta + \frac{4\pi}{3}\right) \hat{\mathbf{y}} \right]. \quad (\text{B.57})$$




Figure B.5: An example of how each stage would pitch under rotations of angle α_1 and α_2 about an arbitrarily aligned x -axis.

Thus, we can write the unsimplified version of the kinetic and potential energy equations:

$$T = \frac{1}{2}I_1\dot{\alpha}_1^2 + \frac{1}{2}I_2\dot{\alpha}_2^2, \quad (\text{B.58})$$

$$\begin{aligned} V = & \frac{1}{2}k_1[(R\cos(\theta)\sin(\alpha_1)]^2 \\ & + \frac{1}{2}k_1\left[R\cos\left(\theta + \frac{2\pi}{3}\right)\sin(\alpha_1)\right]^2 \\ & + \frac{1}{2}k_1\left[R\cos\left(\theta + \frac{4\pi}{3}\right)\sin(\alpha_1)\right]^2 \\ & + \frac{1}{2}k_2\{R\cos(\theta)[\sin(\alpha_1) - \sin(\alpha_2)]\}^2 \\ & + \frac{1}{2}k_2\left\{R\cos\left(\theta + \frac{2\pi}{3}\right)[\sin(\alpha_1) - \sin(\alpha_2)]\right\}^2 \\ & + \frac{1}{2}k_2\left\{R\cos\left(\theta + \frac{4\pi}{3}\right)[\sin(\alpha_1) - \sin(\alpha_2)]\right\}^2. \end{aligned} \quad (\text{B.59})$$

Using angle-addition formulae and some algebra, we find that the potential energy equation simplifies to:

$$\begin{aligned} V = & \frac{3}{4}(k_1 + k_2)R^2\sin^2\alpha_1 \\ & - \frac{3}{2}k_2R^2\sin\alpha_1\sin\alpha_2 \\ & + \frac{3}{4}k_2R^2\sin^2\alpha_2. \end{aligned} \quad (\text{B.60})$$

We note that, as expected, there is no dependence on the coordinate system used. Now, we can use the Euler-Lagrange equations and the small-angle approximation to find the coupled differential equations.

$$I_1\ddot{\alpha}_1 + \frac{3}{2}(k_1 + k_2)R^2\alpha_1 - \frac{3}{2}k_2R^2\alpha_2 = 0, \quad (\text{B.61})$$

$$I_2\ddot{\alpha}_2 - \frac{3}{2}k_2R^2\alpha_1 + \frac{3}{2}k_2R^2\alpha_2 = 0. \quad (\text{B.62})$$

Solving for the eigenfrequencies, we find that

$$\omega^2 = \frac{3R^2[(k_1 + k_2)I_2 + k_2I_1]}{4I_1I_2} \left\{ 1 \pm \sqrt{1 - \frac{4I_1I_2k_1k_2}{[(k_1 + k_2)I_2 + k_2I_1]^2}} \right\}. \quad (\text{B.63})$$

Case 4—Lastly, we look at the natural frequencies for rotations of the disks about the vertical axis, i.e. the double torsion pendulum modes. In this case, we will say that each




Figure B.6: An example of how each stage would twist under rotations of angle α_1 and α_2 about the vertical.

disk rotates about the vertical by angles α_1 and α_2 , respectively. In Fig. B.6, we show an example of this type of rotation. We note that in this case, each spring is rotated by an angle θ_1 and θ_2 with respect to the vertical. Thus, the kinetic energy and potential energy equations are:

$$T = \frac{1}{2}I_1\dot{\alpha}_1^2 + \frac{1}{2}I_2\dot{\alpha}_2^2, \quad (\text{B.64})$$

$$V = -(m_1 + m_2)gl_1 \cos \theta_1 - m_2gl_2 \cos \theta_2, \quad (\text{B.65})$$

where we have dropped the constants in the potential energy and neglected the kinetic energy due to vertical movement, as this is very small (also the equations are no longer analytically solvable).

Now, we have that the angles θ_1 and θ_2 are related to the α_1 and α_2 by the following relations:

$$l_1\theta_1 = R\alpha_1, \quad (\text{B.66})$$

$$l_1\theta_1 + l_2\theta_2 = R\alpha_2. \quad (\text{B.67})$$

Next, we carry out the Euler-Lagrange equations, use the small-angle approximation, and use the above relations to arrive at our coupled differential equations:

$$I_1\ddot{\alpha}_1 + \left(\frac{m_1 + m_2}{l_1} + \frac{m_2}{l_2} \right) gR^2\alpha_1 - \frac{m_2}{l_2} gR^2\alpha_2 = 0, \quad (\text{B.68})$$

$$I_2\ddot{\alpha}_2 - \frac{m_2}{l_2} gR^2\alpha_1 + \frac{m_2}{l_2} gR^2\alpha_2 = 0. \quad (\text{B.69})$$

Solving for the eigenfrequencies, we find that

$$\omega^2 = \frac{\left[\left(\frac{m_1 + m_2}{l_1} + \frac{m_2}{l_2} \right) I_2 + \frac{m_2}{l_2} I_1 \right] gR^2}{2I_1I_2} \left\{ 1 \pm \sqrt{1 - \frac{4I_1I_2 \frac{m_1 + m_2}{l_1} \frac{m_2}{l_2}}{\left[\left(\frac{m_1 + m_2}{l_1} + \frac{m_2}{l_2} \right) I_2 + \frac{m_2}{l_2} I_1 \right]^2}} \right\}. \quad (\text{B.70})$$

B.1.3 Spring Mode Intuition

We can gain some intuition on the final vibrational system via these fundamental frequencies. First, we note that the frequencies follow two basic forms: pendulum frequencies (i.e. $f \sim \sqrt{g/l}$) and spring frequencies (i.e. $f \sim \sqrt{k/m}$). Thus, we have that the ideal system will have springs that have the smallest spring rates and pendula that are as long as possible. These will be limited by the design constraints of our system, as will be discussed in Section B.2.2.

However, one constraint that we can use to quickly gain intuition on our spring design is that we have a finite vertical volume for our system. For the one-stage system, we should use the full vertical volume. For the two-stage system, this means that the two stretched lengths




Figure B.7: Showing that the lowest frequency is achieved by having the same stretched length, this is for the translational frequencies in the x, y directions. The values shown assume $m_1 = 2m_2$, but the conclusion is true for any distribution of mass between the two stages.

must add up to some maximum value that is defined by the dilution refrigerator geometry. For our system, we will see that 0.5 m is roughly this maximum. If we take this constraint, we can see from Fig. B.7 that a representative pendulum mode for the two-stage system will have its lowest value with the first stage and second stage springs having the same length.

It is also worth discussing mass distribution between the two stages. For the one-stage system, it is simple, as the spring modes imply to have the single stage be as heavy as possible. For the two-stage system, it will be more complex, as we need to know how to distribute the mass between stages. To do this, we will assume some total mass of 15 kg as a placeholder value for that of the system. In Fig. B.8, we see that a characteristic pendulum mode implies that as much mass as possible should be placed on the upper stage. On the other hand, if we look at a characteristic spring mode, then we see that there will be an upper bound on this value (depending on the spring constants), as shown in Fig. B.9. Nonetheless, the inequality that $m_1 > m_2$, or that the upper stage should have more mass than the lower stage, will hold. Thus, we will design our system with the intuition that we should put more mass on the upper stage than on the lower stage and make a decision on the distribution based on the spring constants we will be able to use.




Figure B.8: Showing that the lowest frequency is achieved by having the majority of the mass on the upper stage, this is for the translational frequencies in the x, y directions. The values shown assume $\ell_1 = \ell_2$.

B.2 Extension Spring Design

We have our expected fundamental frequencies and a small amount on intuition for some of the two-stage system's values. With this in mind, we need to understand some of the fundamental constraints of our system. That is, we must know what the geometric constraints of the dilution refrigerator are (e.g. the maximum vertical space, the horizontal space), as well as constraints on spring design from an engineering point-of-view.

B.2.1 Design Principles

For spring design, the majority of the equations and rules are well-known in standard mechanical engineering textbooks, such as *Shigley's Mechanical Engineering Design* [233], which we will follow when designing our springs.

Spring Rate

We start with the maximum stress in a wire being a superposition of the direct shear stress and the torsional shear stress:

$$\tau_{\max} = \frac{F}{A} + \frac{Tr}{J}, \quad (\text{B.71})$$




Figure B.9: Showing that there may be a raising of the spring mode frequencies if put too much mass on the upper stage, but the $m_1 > m_2$ relation still holds. This is for the translational frequencies in the z direction. The values shown assume $\ell_1 = \ell_2$, and $k_1 = 3k_2$. The inequality holds for any pair k values, while the location of frequency degradation will approach $m_1 = m_2$ as k_1 approaches k_2 from below.

where F is the applied force, A is the cross sectional area of the wire, T is the applied torque, r is the radius of the coiled wire, and J is the polar moment of area. To relate these values directly to the dimensions of the spring, we have that

$$T = \frac{1}{2}FD, \quad (\text{B.72})$$

$$r = \frac{d}{2}, \quad (\text{B.73})$$

$$J = \frac{\pi d^4}{32}, \quad (\text{B.74})$$

$$A = \frac{\pi d^2}{4}, \quad (\text{B.75})$$

where D is the average spring diameter, and d is the wire diameter. Each of these values are shown diagrammatically in Fig. B.10.

Plugging these values in Eq. (B.71), we have that

$$\tau_{\max} = \frac{8FD}{\pi d^3} + \frac{4F\pi}{d^2}, \quad (\text{B.76})$$




Figure B.10: (Figure from Ref. [233]) (a) Diagram of the spring diameter D and the wire diameter d , showing the forces being applied vertically. (b) The applied torque on a given segment of wire given the applied force. Note that this diagram is actually for a compression spring, but the definitions are the same for an extension spring (with the force being in the opposite direction).

which can be further simplified with the definition of the spring index $C \equiv D/d$ and the shear stress-correction factor $K_s \equiv \frac{2C+1}{2C}$, giving

$$\tau_{\max} = K_s \frac{8FD}{\pi d^3}. \quad (\text{B.77})$$

However, this shear stress-correction factor relies on the assumption that the wire is straight, which of course is not true for a spring. Note that a general engineering rule-of-thumb for spring design is that C should fall in the range of $4 < C < 12$, as a spring with index outside of this range will be prone to failure. Generally, in order to be conservative by using the largest empirical stress-correction factor, it is common to replace K_s with the Bergsträsser factor

$$K_B \equiv \frac{4C + 2}{4C - 3}, \quad (\text{B.78})$$

which will be larger for a given C , as compared to K_s . Thus, this gives maximum stress of

$$\tau_{\max} = K_B \frac{8FD}{\pi d^3}. \quad (\text{B.79})$$

At this point, it is of importance for us to relate the total deflection of the spring y to the force F . As shown again in Ref. [233], Castigliano's theorem will result in

$$y = \frac{8FD^3 N_a}{d^4 G}, \quad (\text{B.80})$$




Figure B.11: (Figure from Ref. [233]) (a) Dimensions of a hook in relation to point A on the right hand side of the spring hook. (b) Rotating by 90° , the dimensions of the same hook in relation to point B .

where G is the modulus of rigidity of the spring material, and N_a is the number of active coils (we will define number of active coils in the following paragraphs). Thus, we know that the spring constant is (via Hooke's Law) defined by the ratio of the force to the spring deflection ($k = F/y$), which gives for a spring that

$$k = \frac{d^4 G}{8D^3 N_a}. \quad (\text{B.81})$$

Adding Hooks

The maximum tensile stress at A in Fig. B.11 is given by

$$\sigma_A = F \left[(K)_A \frac{16D}{\pi d^3} + \frac{4}{\pi d^2} \right], \quad (\text{B.82})$$

where $(K)_A$ is a bending stress-correction factor for curvature

$$(K)_A = \frac{4C_1^2 - C_1 - 1}{4C_1(C_1 - 1)} \quad (\text{B.83})$$

and

$$C_1 = \frac{2r_1}{d}. \quad (\text{B.84})$$

The maximum torsional stress at point B is given by

$$\tau_B = (K)_B \frac{8FD}{\pi d^3}, \quad (\text{B.85})$$

where $(K)_B$ is yet another stress-correction factor for curvature

$$(K)_B = \frac{4C_2 - 1}{4C_2 - 4} \quad (\text{B.86})$$

and

$$C_2 = \frac{2r_2}{d}. \quad (\text{B.87})$$

When designing extension springs for loads, both the yield strengths of the coils as well as the hooks must be taken into account to ensure that there is negligible chance of failure.

When making springs, spring manufacturers typically prefer the inclusion of some initial tension such that the free length (i.e. unstretched spring length) can be held more accurately. The deflection of the spring then becomes depending on exceeding the initial tension F_i , such that

$$y = \frac{F - F_i}{k}. \quad (\text{B.88})$$

The free length of a spring including the end hooks is then expressed as

$$L_0 = 2(D - d) + (N_b + 1)d, \quad (\text{B.89})$$

where D is the mean coil diameter, N_b is the number of body coils, and C is the spring index. Note that this assumes that the hook has the same diameter as the coils. For the number of active coils for a spring with ordinary end loops (e.g. Fig. B.11) is given by

$$N_a = N_b + \frac{G}{E}, \quad (\text{B.90})$$

where G is again the modulus of rigidity, and E is the elastic modulus.

Practically, spring manufacturers tend to have a preferred range for what initial tension to lock in for a spring to ensure that the winding process of the wire to create the spring is easily repeatable (i.e. easy to do with the usual equipment available). The preferred range can be expressed in terms of the uncorrected torsional stress

$$\tau_i = \frac{231.0}{\exp(0.105C)} \pm 6.895 \left(4 - \frac{C - 3}{6.5} \right) \text{ MPa}. \quad (\text{B.91})$$

Equivalently, we can plug in the uncorrected torsional stress $\tau = 8F_i D / (\pi d^3)$ to directly calculate the preferred range of initial tension

$$F_i = \frac{\pi d^3}{8D} \tau_i. \quad (\text{B.92})$$

Table B.1: Constants A and m for determining the minimum tensile strength of phosphor bronze spring wires. Values are from Ref. [233], where the measurements are for phosphor bronze wire ASTM B159 and temper CA510.

d [mm]	A [MPa mm]	m
0.1–0.6	1000	0
0.6–2	913	0.028
2–7.5	932	0.064

Safety Factors

To ensure our system is safe (i.e. the spring system is highly unlikely to fail and drop our detectors inside our dilution refrigerator), we need to keep track of the various safety factors for the various stresses in the system (both the body coils and the hooks). Generally, this is the ratio between the maximum allowable stress of the spring material to the expected stress from the hanging load.

For various materials, the minimum tensile strength can be represented by

$$S_{ut} = \frac{A}{d^m} \quad (\text{B.93})$$

where A and m are fitted constants to empirical data. In Table B.1, we show the corresponding values for various wire diameters of phosphor bronze spring wires. For various failure modes, the safety factor will be different, such that we must calculate it for three different cases: failure of the spring body, failure due to bending of the spring hook at point A in Fig. B.11, and failure due to bending of the spring hook at point B in Fig. B.11. For a nonferrous alloy (e.g. phosphor bronze), the maximum allowable stresses in percent of minimum tensile strength are 35%, 55%, and 30%, respectively. Thus, the safety factors n to keep track of when designing extension springs are

$$n = \begin{cases} \frac{0.35S_{ut}}{\tau_{\max}} & \text{failure at body} \\ \frac{0.55S_{ut}}{\sigma_A} & \text{failure at } A \\ \frac{0.30S_{ut}}{\tau_B} & \text{failure at } B \end{cases} \quad (\text{B.94})$$

B.2.2 Geometric Constraints and Expected Load

Starting with the geometric constraints (Fig. B.12), we have that the fridge has 61 cm of vertical space (from mixing chamber to bottom of the 1 K can) and a horizontal space of 30 cm in diameter, and the detector housing that we expect to install to the bottom of the lower decoupler stage has a height of 15.5 cm. With these dimensions, we are restricted to about 45 cm of vertical space. Because we do not want the housing to touch the bottom of the 4 K can, we will design our springs with a vertical constraint of 35 cm. This allows some




Figure B.12: (Left) Cross section of the Pyle Lab dilution refrigerator from the CAD design. The vertical space between the mixing chamber and the bottom of the 1 K can is about 61 cm. Note that the internal payload in this cross section is a placeholder and not what will be installed. (Right) Design of IR-tight detector housing by Yen-Yung Chang and Siddhant Mehrotra. This housing will be attached to the bottom of the lower stage of the vibration decoupler design.

extra space to work with in case we decide to increase the thickness of the various plates, depending on how much vertical space is taken by the spring hooks, and to ensure that there is some vertical space for the system to fall in case a failure (which will then be caught by a safety structure, proposed in Section B.5).

The expected mass of the manufactured detector housing is 18 kg. With detectors, electronics, and other necessities, we will estimate that the total detector payload is about 20 kg when considering spring design. Thus, these constraints in geometry and mass will restrict the length of stretched springs and the radius of the stages, and this must be taken into account.

B.2.3 Proposed Extension Springs

With all of these spring design equations to keep track of, we will design our springs using the following workflow for calculating the various spring dimensions and the expected natural frequencies of the system:

1. Choose a spring material
2. Determine the expected total hanging load and per spring
3. Choose a maximum spring diameter D
4. Solve for the wire d for a given safety factor for failure at the spring body (use 2, as this is generally the least likely failure mode)
5. Determine the low-end value of the preferred F_i range for spring manufacturers
6. Solve for the number of body coils N_b needed to achieve the desired stretched length given the geometric constraints, using the sum of the free length and the expected deflection. Round N_b down (to stay within the geometric constraints) to the nearest half-integer to ensure that the hooks are in-plane with each other
7. Use N_b , D , and d to calculate the spring rate k
8. Calculate the safety factors for the hook failure modes (at points A and B) to ensure they are 1.5 or greater (a reasonable value)
9. Calculate the expected natural frequencies of the system

One Stage

As the one-stage system will be simpler than the two-stage system, we will use it as a relatively in-depth example of the workflow. For our spring material, we will choose phosphor bronze due to its long history of usage in cryogenics for its good thermal conductance. As will be discussed in later sections, this property is important, as the springs themselves should quickly cool down, such that they do not act as a heat source. For the one-stage system, we will assume that the hanging mass is a total of 20 kg, given by our expected load from the IR-tight housing (Section B.2.2). We next set the maximum spring diameter to $D = 25$ mm semi-arbitrarily—given that the stages are up to 27 cm in diameter, the chosen D is about 10% of this, and the springs can be installed close to the perimeter of the stage. Furthermore, we should remember that our intuition on springs was such that the spring should have the smallest possible spring rate $k \propto d^4/D^3$, this implies that we always want the spring diameter D to be as large as possible and the wire diameter d to be as small as possible.

To solve for d given a safety factor of 2 at the spring body, we use Eqs. (B.79) and (B.94), rearranging as

$$\frac{0.35S_{ut}}{n} = K_B \frac{8F_{load}D}{\pi d^3} \quad (\text{B.95})$$

and solving for d , where $F_{load} = \frac{1}{3}mg$. In our case of $m = 20$ kg for the one-stage system, we find that $d = 3.2$ mm is the smallest that the wire diameter can be (any lower and the safety factor would drop below 2). With both D and d known, we use Eq. B.92 to determine

the smallest reasonable initial tension F_i for a spring manufacturer to be able to fabricate. In this case, we find that $F_i = 40$ N for each spring. Comparing to the mass per spring of $20/3$ kg (i.e. $F_{load} = 65$ N per spring), this initial tension is significant and the deflection of the springs under load will be about 2.6 times less than if there were no initial tension in the spring.

The next step is to solve for the number of body coils N_b , which will finish up the spring design parameters. Given our geometric constraints of the fridge, we know that the one-stage system should not have a stretched spring length longer than $l_{max} = 35$ cm. Using Eqs. (B.88) and (B.89), we rearrange the relations to

$$l_{max} = L_0 + y = 2(D - d) + (N_b + 1)d + \frac{8(F_{load} - F_i)D^3(N_b + \frac{G}{E})}{Gd^4} \quad (\text{B.96})$$

and solve for N_b . For our one-stage system, we find that the number of body coils needed to fulfill all of our requirements so far is $N_b = 77$. Note that this was done assuming that the hook radius about point A was the same as the spring radius (i.e. $r_1 = D/2$), so we will keep this assumption and furthermore set the hook radius about point B to the same value ($r_2 = D/2$).

In the last few steps, we use all of our design parameters to calculate the spring rate k and the remaining safety factors related to hook failure. We find that $k = 440.6$ N/m and the safety factors of the hook for failure about point A and point B are 1.63 and 1.82, respectively. Thus, we would expect failure to occur at the hooks before the spring body. Nonetheless, our spring design for the expected load should be safe, as the safety factors are far above 1 (i.e. we would have to overload our system by more than 10 kg to cause a failure).

With $k = 440.6$ N/m, $l_{max} = 35$ cm, a hanging mass of 20 kg, and stage diameter of 27 cm, we finally return to the natural frequencies that we solved for using the Lagrangian method in Section B.1.1. Plugging into Eqs. (B.30)–(B.33), we find that the various natural frequencies are $f_z = 1.29$ Hz, $f_{x,y} = 0.84$ Hz, $f_{z,rot.} = 1.19$ Hz, and $f_{x,y,rot.} = 1.83$ Hz. Above these frequencies, all vibrationally-induced noise should be attenuated. Given that the pulse-tube itself oscillates at 1.4 Hz, these frequencies are not exactly that frequency, and we should be safe from being on a resonance. For this proposed one-stage system, the design and calculated parameters are all stored in Table B.2.

Two Stage

When determining the spring design parameters for the two-stage system, the workflow is equivalent, but must be done twice. That is, we split up the stages into an upper stage with a total hanging mass of $m_1 + m_2$ (the sum of the masses of both stages) and a lower stage with a total hanging mass of $m_2 = 20$ kg (the mass of the IR-tight housing). Again choosing phosphor bronze, we start with the same spring diameter $D = 25$ mm for the springs on each stage. Although m_1 can be any value, we have seen from our intuition that it should be at least as much as m_2 if not more. In this design, we will choose $m_1 = 20$ kg as well

Table B.2: Spring design parameters and expected values for the one-stage system.

System	Parameter	Value
One Stage	Total Hanging Mass m	20 kg
	Spring Diameter D	25 mm
	Wire Diameter d	3.2 mm
	Hook Bend Radius r_1	12.5 mm
	Hook Bend Radius r_2	12.5 mm
	Spring Index C	7.8
	Initial Tension F_i	40 N
	Number of Body Coils N_b	77
	Number of Active Coils N_a	77.4
	Spring Rate k	440.6 $\frac{\text{N}}{\text{m}}$
	Free Length L_0	29.2 cm
	Expected Deflection y	5.7 cm
	Safety Factors n (body, A , B)	2, 1.63, 1.82
	Natural Frequency f_z	1.29 Hz
	Natural Frequency $f_{x,y}$	0.84 Hz
	Natural Frequency $f_{z,rot.}$	1.19 Hz
	Natural Frequency $f_{x,y,rot.}$	1.83 Hz

and will later check that a larger mass would not give significantly greater performance (i.e. significantly lower the natural frequencies of the system).

In this two-stage system, we also will ensure that the stretched lengths of the upper and lower stage springs are not identical, to ensure that the resonances of each spring are nonoverlapping. To do this, we make another semi-arbitrary choice and choose the upper stage springs to have stretched lengths of 17.9 cm and the lower stage springs to have stretched lengths of 16.9 cm. We then follow our workflow to determine the rest of the values, just as we did for the one-stage system. Using the proposed constraints in the previous paragraph, the parameters for the two-stage system are reported in Table B.3.

We now return to our choice of $m_1 = 20$ kg on the upper stage, such that the upper springs together would be holding up $m_1 + m_2 = 40$ kg in mass. In order to show that this was a reasonable choice, we can show how the various natural frequencies change as a functional of upper stage mass m_1 . To do this, we simply follow the same workflow for some arbitrary m_1 and report the various natural frequencies which would correspond to the optimal design, with the results shown in Fig. B.13. In this figure, it is clear that the difference between $m_1 = 20$ kg and $m_1 = 40$ kg is small, such that are original choice is reasonable. We also note that there is a minimum in both of the highest frequencies such that too large of masses actually begin to increase some of the frequencies. This should not be surprising given what we showed the same behavior in Fig. B.9, from which we would

Table B.3: Spring design parameters and expected values for the two-stage system.

System	Parameter	Value
Two Stage, Upper	Total Hanging Mass $m_1 + m_2$	40 kg
	Spring Diameter D	25 mm
	Wire Diameter d	4.1 mm
	Hook Bend Radius r_1	12.5 cm
	Hook Bend Radius r_2	12.5 cm
	Spring Index C	6.1
	Initial Tension F_i	105.5 N
	Number of Body Coils N_b	30.5
	Number of Active Coils N_a	30.9
	Spring Rate k	$3028.5 \frac{\text{N}}{\text{m}}$
	Free Length L_0	17.1 cm
	Expected Deflection y	0.8 cm
	Safety Factors n (body, A , B)	2, 1.64, 1.84
Two Stage, Lower	Total Hanging Mass m_2	20 kg
	Spring Diameter D	25 mm
	Wire Diameter d	3.2 mm
	Hook Bend Radius r_1	12.5 mm
	Hook Bend Radius r_2	12.5 mm
	Spring Index C	7.8
	Initial Tension F_i	40.1 N
	Number of Body Coils N_b	31
	Number of Active Coils N_a	31.4
	Spring Rate k	$1086.0 \frac{\text{N}}{\text{m}}$
	Free Length L_0	14.6 cm
	Expected Deflection y	2.3 cm
	Safety Factors n (body, A , B)	2, 1.63, 1.82
Two Stage (Full System)	Natural Frequencies f_z	1.67, 4.11 Hz
	Natural Frequencies $f_{x,y}$	0.91, 2.21 Hz
	Natural Frequencies $f_{z,rot.}$	1.29, 3.12 Hz
	Natural Frequencies $f_{x,y,rot.}$	2.37, 5.82 Hz




Figure B.13: The dependence of the various frequencies on the mass of the upper stage in the two-stage system. The solid lines denote the higher frequency for the given mode, while the dashed lines denote the lower frequency for the given mode. Note that the jumps in the curves corresponding to spring modes are due to rounding the number of body coils N_b to the nearest half-integer—without this constraint, the curves would be smooth.

expect worse frequencies for larger upper stage masses after some point for the spring modes.

We have spent most of our time with the eigenfrequencies of the system, but it is worth also checking how the eigenvectors change as we increase the mass (while keeping our springs within the safety factors that we have chosen). Taking Eqs. (B.37) and (B.38), we can solve for the ratio of the eigenvector components of the system for the eigenmodes $\omega_{\pm}$ corresponding to vertical translations

$$\frac{\Delta z_2}{\Delta z_1} = \frac{3k_{\text{lower}}}{-m_1\omega_{\pm}^2 + 3(k_{\text{lower}} + k_{\text{upper}})}. \quad (\text{B.97})$$

This ratio of the vertical displacement of the two stages becomes as shown in Fig. B.14 as the ratio of the upper and lower stage masses varies. As we increase the ratio of the masses in the upper and lower stages, the displacement of the lower stage increases considerably in relation to the upper stage. As our detectors will be installed on the lower stage (and be sensitive to frictional rubbing), it is preferable to keep the displacement small, while ensuring that we have excellent attenuation. Thus, this brings us back to the choice of $m_1 = m_2 = 20$ kg, where we have this trade-off of low natural frequencies and small relative displacements of the upper and lower stages.




Figure B.14: Ratio of the vertical displacement (i.e. eigenvector components) of each stage for the corresponding natural frequencies as the ratio of the upper and lower stage masses varies. Note that the jumps in each of the curves are due to rounding the number of body coils N_b to the nearest half-integer—without this constraint, the curves would be smooth.

B.3 Thermal Considerations

We must make sure the heat flow in our vibration decoupler is efficient, such that we do not have any parasitic power sources that will prevent our system from cooling to base temperature (e.g. 8 mK).

B.3.1 Thermal Conductance Between Stages

In our system, we will be using flexible OFHC (Oxygen-Free High thermal Conductivity) Cu braids for thermal links [234] and our designed phosphor bronze springs (which themselves will have some thermal conductance). For these metals, the thermal conductivity can be approximated using the empirical Wiedemann–Franz law, which states that the ratio of the electronic contribution to the thermal conductivity κ to the electrical conductivity σ of a metal is proportional to temperature

$$\frac{\kappa}{\sigma} = LT, \quad (\text{B.98})$$

Table B.4: (Values from Ref. [236]) The thermal conductances of various interfaces at different temperatures and the empirical exponent to be use with Eq. B.102.

Interface Materials	G [W/K] at 4.2 K	G [W/K] at 77 K	y	Ref.
Au/Au	0.2	—	1.3	[237]
Cu/Cu	0.01	0.3	1.3	[237, 238]
Steel/steel	0.005	0.3	—	[238]

where L is the Lorenz number

$$L \equiv \frac{\pi^2}{3} \left(\frac{k_B}{e} \right)^2 = 2.44 \times 10^{-8} \frac{\text{V}^2}{\text{K}^2}. \quad (\text{B.99})$$

For the electrical conductance of a metal, one can generally use a measured value at room temperature and a measured value at a few K. Usually this ratio is defined using the electrical resistivity $\rho = 1/\sigma$, with the Residual-resistance ratio (RRR) defined as

$$\text{RRR} = \frac{\rho_{300\text{K}}}{\rho_{0\text{K}}}. \quad (\text{B.100})$$

Practically, the resistivity of a given sample will be measured at a few K, and the RRR will be determined at this nonzero temperature. Notably, in superconducting metals, the RRR would be defined using the resistivity just above the superconducting transition temperature, as $\rho = 0$ below this. For us, we will use OFHC Cu and CDA 510 phosphor bronze, which each have measured RRRs of 77 [234] and 1.24 [235], respectively.

For determining the thermal conductance G of a metal, we use the Wiedemann–Franz law, measured RRR values, and that the conductance is related to the conductivity via $G = \kappa A/\ell$, where A is the area through which heat flows (e.g. the cross sectional of a wire) and ℓ is the thickness of the metal through which heat is flowing (e.g. the length of a wire). Putting this together, we have that thermal conductance of a normal metal at low temperatures is

$$G = \frac{LA \cdot \text{RRR}}{\rho_{300\text{K}} \ell} T. \quad (\text{B.101})$$

Note that, in the case of Cu braids, there may be some empty space between the braided wires. Thus, the actual cross sectional area A is smaller than what can be measured from the diameter of the braid. This is parameterized by the porosity p , where the true cross section for single braid is given by $A = p\pi r^2$, where r is the radius of the braid. In Ref. [234], the porosity of Cu braid was given to be $p = 0.645$, which we will use in our design.

To thermally connect our Cu braids and springs to the various stages of the passive vibration isolation system, the means of this will be achieved through solid-solid interfaces. Interestingly, the thermal conductance through a solid-solid interface is not dependent on the area of contact, but instead on the force being applied to the interface. This is because the

microscopic roughness of the surfaces means that there are only points of contact, thus adding force creates a better thermal link by increasing the number of points of contact. As discussed by Ekin [236], the thermal conductance of a solid-solid interface can be calculated via the empirical expression

$$G(F, T) = G(F = 445 \text{ N}, T = 4.2 \text{ K}) \frac{F}{445 \text{ N}} \left(\frac{T}{4.2 \text{ K}} \right)^y, \quad (\text{B.102})$$

where F is the force between the two surfaces, T is the temperature of the system, and y is a dimensionless constant with values between 1.3 and 3, depending on the material. For pertinent values for Au/Au, Cu/Cu, and steel/steel interfaces, we have reproduced the table from Appendix A2.3 of Ref. [236] in Table B.4, dropping the row on sapphire/sapphire. Further experimentally-measured values have been compiled in Ref. [239] for various other interfaces.

To create contact between the Cu braids and the different stages, one generally uses screws to create a large contact force. For our springs, we will be using stud anchors, where the weight of the stages themselves will be shown to be sufficient for a good thermal connection. Because these contact forces are in series with the metals, as well as parallel conductances through the springs and the braids, the thermal system will become more complex. In studying this system, we refer back to the relations between thermal conductances in series and in parallel

$$\frac{1}{G_{\text{series}}} = \frac{1}{G_1} + \frac{1}{G_2} + \cdots, \quad (\text{B.103})$$

$$G_{\text{par.}} = G_1 + G_2 + \cdots. \quad (\text{B.104})$$

When keeping track of the various thermal conductances, we will ensure we follow these relations such that our power loads are modeled correctly.

To calculate steady-state power loads after all the thermal conductances have been estimated, we use the heat flow equation between two bodies of different temperatures

$$P = \int_{T_{\text{low}}}^{T_{\text{high}}} dT' G(T'). \quad (\text{B.105})$$

B.3.2 Stiffness of Copper Braids

Because these braids will be a part of the passive vibration isolation system, the stiffness of the braids may have a nonnegligible effect on the spring constants of our springs. The braids will connect between each stage, and thus can be thought of as its stiffness is in parallel in the various spring rates of our phosphor bronze springs. Spring rates in parallel are governed by

$$k_{\text{par.}} = k_1 + k_2 + \cdots, \quad (\text{B.106})$$

which implies that the stiffness of the copper braid must be much less than the spring rates of the phosphor bronze springs in order to negligibly affect the natural frequencies.

As discussed in Section 5.1.6.1 of Ref. [240], the simplest idealization of thermal link stiffness is to model the braid as a fixed-free cantilever beam (i.e. one end of the beam is fixed and one is free). When the braid is bent at with a force F at its free end, the deflection of that end is

$$x = \frac{F\ell^3}{3EI}, \quad (\text{B.107})$$

where E is the elastic modulus and I is the second moment of area. For a bundle of small wires (e.g. a braid), the second moment of area would be given as

$$I_{\text{braid}} = \frac{N\pi d^4}{64}, \quad (\text{B.108})$$

where N is the number of wires and d is the wire diameter. Thus, the stiffness k in this idealized case can be estimated as

$$k_{\text{braid}}^{\text{ideal}} = \frac{3EN\pi d^4}{64\ell^3}. \quad (\text{B.109})$$

Unfortunately, this idealized case often does not apply, as the friction between wires has a significant effect on link stiffness. Furthermore, the Cu in the braids will build up oxide layers on its surface over time (when, e.g., the dilution refrigerator is open to ambient air), which can increase friction and create oxide bonds between wire. However, we can still use this for some intuition, such that the stiffness of a braid should decrease with length, less braids should give lower stiffness, and smaller diameter braids should also give lower stiffness. Knowing that the internal friction of a braid will increase the stiffness, a bent braid should also have a large radius (e.g. less friction between internal strands). Lastly, the braids will need to be attached to a small Cu block, which can be screwed into our stages for thermal connection between surfaces, and the braid should have a minimal bend at the Cu block, or otherwise there would be increased stiffness (and a chance for failure). In Section B.3.3, we will use our geometric constraints of our system to ensure that we design thermal straps that fulfill each of these ideas.

The theoretical studies of metallic strands of wires have been undertaken over the years using Euler-Bernoulli beam theory, e.g. in the textbook *Theory of Wire Rope* [241] and more recently in Ref. [242]. However, the complexity of these analyses and models (as well as the dependence on the specific braid design) prevent us from easily calculating the expectations of the stiffness for different thermal links made of braids. Thus, it is better to empirically measure the stiffness of the braid, which a manufacturer of thermal links (e.g. Technology Applications, Inc. [243]) offer as part of the fabrication process (they have also made many thermal links in the past, and thus have many results from previous measurements).

B.3.3 Geometric Constraints

To understand the allowed lengths of our thermal links, we must understand the various dimensions of our system. As we have estimated spring lengths for the upper and lower




Figure B.15: The various pertinent dimensions of a prototype design, where the spacing of the the stages is determined by the spring lengths in Table B.3. The top plate is the mixing chamber plate (i.e. what the passive vibration isolation system attaches to).

stages, we can mock up a prototype design, where we will assume a 20 kg Cu weight is being used on the upper stage. This Cu weight is not exactly 20 kg, but has a mass of 21.1 kg such that the dimensions are simpler to report. The dimensions of this system are shown in Fig. B.15. Note that the Cu weight has a smaller radius than the stages, which was done to ensure there was room for the 25 mm diameter springs.

For the Cu braids, we are looking for a tradeoff between excellent thermal conductance and low stiffness. Given the geometric constraints of the system, we can find a range of possible values of the lengths, where we assume that the braids attached between each stage via a semicircle and may have extra length the runs parallel to the stages. To keep the braids from physically touching the springs or the outer cans, we can keep the same horizontal constraints of the Cu weight (cannot go further out that 10 cm in radius from the center of a stage. In Fig. B.16, we show the maximum lengths of these braids given the semicircle shape and the maximum extra lengths. The minimum lengths would simply be the semicircle itself, thus giving a range of braid length for the upper stage of 16.5–46.0 cm and for the lower stage of 26.7–49.7 cm. In the next section, we will use these ranges of the lengths to help determine our thermal link design.




Figure B.16: (Left) The maximum length of a thermal braid which would thermally connect the mixing chamber (top plate) to the first stage (connection via the Cu weight). (Right) The maximum length of a thermal braid which would thermally connect the upper stage (top) to the lower stage (bottom). A prototype design of the IR-tight detector housing is shown connected to the lower stage.

B.3.4 Proposed Thermal Links

For a single Cu link, we have that the thermal conductance will given by the conductances of the clamps and the braid added in series

$$\frac{1}{G_{link}} = \frac{1}{G_{braid}} + \frac{1}{G_{clamp}} + \frac{1}{G_{clamp}}, \quad (\text{B.110})$$

which gives the thermal conductance of a single link as

$$G_{link} = \frac{G_{braid}G_{clamp}}{G_{clamp} + 2G_{braid}}. \quad (\text{B.111})$$

Thus, we must have $G_{braid} \ll G_{clamp}$ for the thermal conductance of the braid to dominate the total thermal link conductance. Taking the conductance equations from Section B.3.1,

we have that for a Cu braid

$$\begin{aligned} G_{braid} &= \frac{LA \cdot \text{RRR}}{\rho_{300\text{K}}^{\text{Cu}} \ell} T \\ &= 3.55 \times 10^{-4} \frac{n}{\ell} T, \end{aligned} \quad (\text{B.112})$$

where n is the number of braids and ℓ is the braid length. In this, we have used a porosity $p = 0.645$, braid radius $r = 1.26$ mm, $\text{RRR} = 77$, and $\rho_{300\text{K}}^{\text{Cu}} = 17.1$ nΩm. Thus, we can only vary the length of the braids and how many there are, as these values are either standard for Cu braids [234] or physical characteristics of Cu.

For the clamp that will hold the copper, we will use bolted Cu joints, which have been estimated to have a force of $F = 3$ kN [244]. We will also assume that we have gold-plated the clamps, such that we can use the Au/Au connection values in Table B.4. Thus, for one the clamps in the thermal link, we have that

$$G_{clamp} = 0.2 \cdot \frac{3000}{445} \left(\frac{T}{4.2\text{K}} \right)^{1.3}. \quad (\text{B.113})$$

For the springs, we expect that a single spring would have a thermal conductance that follows

$$\begin{aligned} G_{spring} &= \frac{LA \cdot \text{RRR}}{\rho_{300\text{K}}^{\text{PhBr}} \ell} T \\ &= 0.275 \cdot \frac{\pi(d/2)^2}{(N_b + 2)\pi D} T, \end{aligned} \quad (\text{B.114})$$

where d is the wire diameter, D is the spring diameter, and N_b is the number of body coils, each of which can be found in Table B.3. We can again assume that we have some type of gold-plating for attaching the springs, such that the interfacial conductance will be determined by the weight of the stages

$$G_{anchor} = 0.2 \cdot \frac{m_{stage}g/3}{445} \left(\frac{T}{4.2\text{K}} \right)^{1.3}. \quad (\text{B.115})$$

For each hanging spring, the conductance in series then similarly becomes

$$G_{hang} = \frac{G_{spring}G_{anchor}}{G_{anchor} + 2G_{spring}}. \quad (\text{B.116})$$

Thus, for each of the stages, we have that the effective thermal conductances are given by

$$G_{upper} = G_{link}^{upper} + 3G_{hang}^{upper} \quad (\text{B.117})$$

$$G_{lower} = G_{link}^{lower} + 3G_{hang}^{lower}. \quad (\text{B.118})$$




Figure B.17: (Left) For the upper stage, the thermal conductance of the four braids, a clamp, and the braids and clamps in series. The series conductance is solely determined by the braid. (Right) For the lower stage, the thermal conductance of the four braids, a clamp, and the braids and clamps in series. The series conductance is solely determined by the braid. Because the systems are nearly identical, these two plots are seemingly the same.

Equivalently, we can estimate the thermal flow from the mixing chamber to the bottom stage as characterized by these two thermal conductances in series

$$G_{total} = \frac{G_{upper}G_{lower}}{G_{upper} + G_{lower}}, \quad (\text{B.119})$$

where we have neglected the thermal conductance of the Cu weight on the upper stage, as it will be much larger (a factor of 50,000 as compared to the braid due to its larger cross sectional area) and will be negligible when added in series.

At this point, we need to understand the tradeoffs between stiffness and thermal conductance. Ideally, a larger thermal conductance would allow our system to cool down efficiently and quickly. However, we know that a large thermal conductance in a braid would mean that they would be very short, and we would need many of them. Both of these would serve to increase the stiffness, which would degrade the performance of the passive vibration isolation. As the goal of this system is to achieve the lowest possible fundamental frequencies (i.e. lowest stiffness), we should be prioritizing the stiffness of the braids over excellent thermal conductance. In this way, we will check first the least stiff scenario, and study if the cooling power of the system is sufficient.

The least stiff scenario for the thermal links are to use the longest possible braids and as little as possible. For links from TAI, they appear to manufacture links with four braids as the smallest amount. Thus, we will study the thermal system with the maximum braid lengths (46.0 cm for the upper stage and 49.7 cm for the lower stage) from the geometric constraints and with four braid. In Figure B.17, we plot the thermal conductances of the four braids in the link, a comparison to a clamp, and the series conductance of the braids




Figure B.18: The temperatures of the upper and lower stages for various heat loads given our thermal link design. Note that we would run our detectors on the lower stage.

and the two clamps. We see that the braid conductance dominates and the plots for the two stages look identical due to using nearly the same length of braids for both stages.

Next, we can check the temperature of each stage under various heat loads for this design. To do this, we assume a mixing chamber temperature of $T_{MC} = 8$ mK, and we can integrate the thermal conductances for each stage to calculate the heat load as a function of temperature. That is, we numerically solve the following equations

$$P(T_1) = \int_{T_{MC}}^{T_1} dT G_{upper}(T), \quad (\text{B.120})$$

$$P(T_2) = \int_{T_{MC}}^{T_2} dT G_{total}(T). \quad (\text{B.121})$$

In Figure B.18, we plot the temperatures of the stages for various heat loads. Given that our generally have pW-scale or less bias powers, we would not expect heat from the detectors to affect the temperature of the stages, as a temperature change only becomes significant at about 20 nW.

In this discussion, we did not make specific note of the thermal conductance of the springs and anchors. Calculating the thermal conductance of one of the spring assemblies $G_{hang}(T)$, we find that $G_{hang}(10 \text{ mK}) = 8.5 \text{ nW/K}$ and $G_{hang}(1 \text{ K}) = 0.85 \mu\text{W/K}$. Comparing to Fig. B.17, this is about four orders of magnitude less than the thermal conductance of




Figure B.19: Thermal time of the upper and lower stages as a function of temperature. Both stages have roughly constant time constants on the order of 160 s.

the link G_{link} , such that it is negligible when added in parallel for each stage. Nonetheless, the thermal conductances of the springs were included in the calculation of the heat loads.

We can also check the cooling times of the system to ensure that the stages do not take an unreasonable amount of time to cool down (e.g. a day long cool down time would mean we would have to wait a week for the detectors to fully cool down). The thermal time constant is a good approximation of this, which is given by

$$\tau(T) = \frac{C(T)}{G(T)}, \quad (\text{B.122})$$

where C is the heat capacity and G is the thermal conductance. We can estimate the low-temperature heat capacity of our masses as Cu blocks with

$$C(T) = \frac{m}{M} (\gamma T + AT^3), \quad (\text{B.123})$$

where m is the mass of the body being cooled (e.g. one of the stages), M is the molar mass of the material, γ is the electronic heat capacity coefficient, and A is the lattice heat capacity coefficient. Generally, A is parameterized in terms of the Debye temperature θ_D via $A = 12\pi^4 R / (5\theta_D^3)$, where $R = 8.31 \text{ J}/(\text{mol K})$ is the ideal gas constant. At very low temperatures, the lattice heat capacity term (T^3) is usually negligible, but will have some effect at about 1 K. For Cu, we will use $\gamma = 0.695 \text{ mJ}/(\text{mol K}^2)$, $M = 63.55 \text{ g/mol}$, and

$\theta_D = 344.5 \text{ K}$ [245]. For the masses, we use a $m = 40 \text{ kg}$ for total mass hanging from the upper stage and $m = 20 \text{ kg}$ for the total mass hanging from the lower stage. As we have already calculated the thermal conductances for both stages (G_{upper} and G_{total}), we can calculate the thermal time constants at low temperatures, as shown in Fig. B.19. We see that the time constants are on the order of 160 s, which is fast enough that this thermal linking scheme should not appreciably lengthen the cool down time more than about 15 min (whereas we were worried about it being on the order of a day).

The proposed thermal links are such as follows: the upper stage is thermally connected to the mixing chamber by 4 Cu braids which are 46.0 cm long, and the lower stage is thermally connected to the upper stage by 4 Cu braids which are 49.7 cm long. These designs will keep the thermal links as flexible as possible due to the length and large bend radius, as well as ensure that up to 20 nW of heat will not appreciably affect the various stage temperatures, and that the cool down time is not significantly increased.

B.3.5 Spring Cross-Checks

With our spring designs, we have so far neglected a couple of thermal aspects so far. It will be useful to know how long we expect the phosphor bronze springs to remain warm compared to the stages (i.e. the thermal time constant), as well as how much we expect the lengths to contract as we cool to cryogenic temperatures.

Spring Cool Down Time

For these springs, it is important that they cool down fast enough such that they are not significantly warmer than the surrounding environment at base temperature. If this were not the case, then the springs could act as a parasitic power source that prevents us from quickly reaching base temperature. To check that this is not the case, we will calculate the thermal time constant $\tau = C/G$ for the upper and lower springs.

We will use the thermal conductance of the interfacial contact of the spring to the anchor (Eq. (B.115)) and the heat capacity of the phosphor bronze spring (where we will estimate the heat capacity as that of Cu). The volume of a spring is given by

$$V_{spring} = (N_b + 2)\pi^2 D \left(\frac{d}{2}\right)^2 \quad (\text{B.124})$$

where N_b is the number of body coils, D is the spring diameter, and d is the wire diameter. For the two-stage system, we report the spring volumes and masses in Table B.5, which we will use to calculate the heat capacity.

For the thermal conductance, we will use that each spring on the upper stage carries a third of the total mass of the two stages (i.e. the force on each upper spring is $40g/3\text{ N}$), and each spring on the lower stage carries a third of the mass of the lower stage (i.e. the force on each lower spring is $20g/3\text{ N}$). For the thermal time constant of each spring, we divide the spring heat capacity by the contact thermal conductance of the connection to the

Table B.5: Volume and masses of the springs for the upper and lower stages of the two-stage system. The masses were determined using a phosphor bronze density of 8800 kg/m^3 .

Spring	Parameter	Value
Two Stage, Upper	N_b	30.5
	D	25 mm
	d	4.1 mm
	V_{spring}	33.7 cm^3
	m_{spring}	297 g
Two Stage, Lower	N_b	31
	D	25 mm
	d	3.2 mm
	V_{spring}	20.8 cm^3
	m_{spring}	183 g




Figure B.20: The thermal time constants of the springs in the upper and lower stages of the two-stage passive vibration isolation system. The solid lines represent thermal contact between two Au interfaces, and the dashed lines represent thermal contact between two Cu interfaces. The difference being a factor of 20 is a direct result of the thermal conductance ratios at 4.2 K and the corresponding empirical formula (see Table B.4).

upper plate. In Fig. B.20, the thermal time constants are plotted over the range of 5 mK to 1 K, where the solid lines denote Au/Au contact (gold-plating), and the dashed lines denote Cu/Cu contact (bare Cu). In either case, we see that the cooling down of the springs is less than a minute, with the Au/Au time constant on the order of seconds. As compared to the average dilution refrigerator cool down time over the course of a few days, the springs should not act as parasitic power sources due to these fast time constants.

Thermal Contraction of Springs

As the springs cool down to cryogenic temperatures, both the shear modulus G of the springs will change, as well as the coefficient of thermal expansion α . We return to the definition of the spring rate k

$$k = \frac{d^4 G}{8D^3 N_a}, \quad (\text{B.125})$$

where d is the wire diameter, N_a is the number of active coils, and D is the wire diameter. Thus, D , d , and G will change as we cool down, and so will the spring rate.

For Cu alloys (e.g. phosphor bronze), it has been shown that the elastic modulus E increases by about 9.6% as the temperatures changes from 300 K to 4 K [246]. Empirically, the elastic modulus and shear modulus for Cu alloys have been shown to be proportional by $G = 3E/8$ [247], and we can assume the same percent change is true for the shear modulus.

To calculate the expected thermal contraction of phosphor bronze after reaching cryogenic temperatures, one can use the coefficient of thermal expansion α , which relates an infinitesimal change in temperature dT to an infinitesimal change in length dL via

$$dL = L_0 \alpha(T) dT, \quad (\text{B.126})$$

where L_0 is the initial length of the material. To estimate the percent change in a length dimension for some temperature range, we integrate

$$\frac{\Delta L}{L_0} = \int_{T_1}^{T_2} dT \alpha(T). \quad (\text{B.127})$$

For phosphor bronze, the coefficient of thermal expansion from 300 K to 10 K is shown in Fig. B.21. Integrating across this temperature range, we find that phosphor bronze would contract by 0.34%. Thus, we expect the wire diameter and spring diameters to contract by this same factor.

Returning to the spring rate and using the expected changes in the various parameters, we have that the spring rate at base temperature can be expected to be

$$k_{base} \approx \frac{1.096(0.9966)^4}{(0.9966)^3} k_{300\text{ K}} \quad (\text{B.128})$$

$$k_{base} \approx 1.092 k_{300\text{ K}}, \quad (\text{B.129})$$




Figure B.21: (Figure from Ref. [248]) Coefficient of thermal expansion of phosphor bronze as a function of temperature.

such that the spring rate increases by about 9.2%. As the natural frequencies of our spring modes are proportional to $\sqrt{k/m}$, we expect that these frequencies would increase by about 4.5% (where this increase is the same regardless of the spring design). For example, using the values in Table B.3, the translational spring modes in the z -axis would increase from 1.67 and 4.11 Hz to 1.75 and 4.29 Hz, respectively. For the rotational spring modes about the x and y axes, these frequencies would increase from 2.37 and 5.82 Hz to 2.48 and 6.08 Hz, respectively. These increases are small, so we should expect similar vibration isolation performance at cryogenic temperatures in the spring modes, while knowing that the frequencies will increase by small amount.

For the pendulum modes, we have that they are essentially proportional to $\sqrt{g/\ell}$. As the full extended length depends on both the free length and the deflection via Eq. (B.96), which we have reproduced below and simplified using the definition of k

$$l_{max} = 2(D - d) + (N_b + 1)d + \frac{F_{load} - F_i}{k}. \quad (\text{B.130})$$

Thus, we can plug in our values from Table B.3 with and without the expected thermal contractions of the various parameters and calculate the expected change in l_{max} due to thermal contraction of the springs. We expect the stretched lengths to decrease by 0.5% for the upper stage and by 0.9% for the lower stage. Thus, the pendulum natural frequencies will increase by a factor of less than $\sqrt{1/0.99} \approx 1.005$, or an increase of less than 0.5%. This change is negligible, and we should have effectively the same performance for the pendulum modes.

B.4 Intermission

With our springs and thermal links designed, we can now put together an in-progress design in CAD, the full tentative system being shown in Fig. B.22. Note that the spring hooks in the design make a right angle, which is not the actual proposed design, but was done to save time creating the springs in CAD. To show the proposed connections for the springs and the thermal links, we take a short break from design to help envision what we have actually proposed so far.

For the various connections, we can use off-the-shelf brass bolts from McMaster-Carr Supply Company [249]. In Fig. B.23, we show the proposed designs for creating the thermal connections with both the springs and the thermal links we have designed. Note that the design for the spring hooks (based on brass eye bolts) has clearance holes through the plates, such that the weight is held through the brass hex nut, which is in contact with plate through the mass of the stages. The reason for the clearance hole is to allow alignment of the eye bolts on the top and bottom of each stages, such that the springs are not subject to a twisting torque (which would be possibility if the holes in the plates were threaded and the different eye bolts has some manufacturing variability in their threads).

The first stage Cu plate and its Cu weight have been combined into one part, which avoids an interface between the two, and simplifies alignment of the weight to the center of




Figure B.22: In-progress passive vibration isolation system design as discussed so far in this appendix. Included are the spring designs, the expected plates for each stage, the thermal link designs, and the the various ways of creating connections between interfaces.

mass. Because the design has a few blind holes at the interfaces of the thermal links, there may be a possibility of virtual leaks (i.e. air pockets in the blind holes could be created under vacuum, which would slowly leave, creating the impression of small leak in the vacuum). To mitigate this, one could drill holes through the center of the screws, which would give a simple path for the air to be pumped out. On the other hand, these air pockets will be small and not within the detector housing, such that it may be negligible for overall detector performance.

There is one important feature left to design for this system, which is a safety structure in case of catastrophic failure at one of the interfaces. That is, if a spring fails, then it would be unwise to allow 40 kg of Cu to fall to the bottom of a dilution refrigerator (this would cause immense damage). Thus, we will finish off this design appendix with a proposed safety structure to avoid such a catastrophe.




Figure B.23: (Left) Connection of a spring to a plate via a brass eye bolt, which goes through a clearance hole in the plate. A hex nut is threaded onto the bolt to carry the weight of the stages, also allowing angular alignment of the eye bolts. (Right) Simple connection to the first stage, where two M4 brass screws are screwed into the copper blocks of the thermal links and into blind holes in the copper mass. On the left hand side of the subfigure, the expected flexible braids are seen.

B.5 Safety Structure

The safety structure design for this system needs to be able to catch either the upper or lower stage in case of failure, which can be done by creating a rigid structure below both stages. There are many different possibilities for this type of system, where the principles are less defined than when we were designing the springs and the thermal links. Instead, we will stick to simple principles of ensuring each stage would be caught and would not fall far.

To keep with the cylindrical symmetry of the design, the safety structure itself will also be cylindrically symmetric. For the upper stage, the springs will block any type of circular design. In contrast, for the lower stage, the IR-tight housing constrains the design radially, such that a circular design will be necessary. In Fig. B.24, we show such a proposal, for which we will go into more detail in the following paragraphs. The idea follows that we use three Cu pillars to hold the safety structure for each stage, where the upper stage is held by a Y-shaped structure, and the lower stage is held by an O-shaped safety structure.

At the mixing chamber plate (Fig. B.25), the connection of the Cu pillars is given by brass bolts, which are installed through threaded holes in both the pillar connection and the mixing chamber plate. This allows for the equal spacing of 120° for each pillar. These pillars go out only as far as the largest radial dimension of the mixing chamber plate, such that the design stays within the radial constraints of the dilution refrigerator. We have also added a chamfer to the corner of the pillars in order to improve the strength at that edge,




Figure B.24: The full design of the passive vibration isolation system with the proposed safety structure.

as opposed to having an abrupt right angle. These pillars then continue downwards to hold the safety structure design for the upper stage.

To rigidly hold the Y-shaped safety structure at the upper stage, L-shaped brackets are used with tapped holes, as shown in Fig. B.26. These brackets can be installed at various heights, depending on any variation in weight of the upper stage, or on how much vertical space is desired between the two. The Y-shaped structure itself is made of Y-shaped Cu plates, which have inset regions for holding seven phosphor bronze compression springs. These springs reduce the impulse on the safety structure, such that there is a more gradual stopping of the weight, rather than being instantaneous. These springs can be bought off-the-shelf from McMaster-Carr, such that they are cheap to replace if they break after a failure. To hold the springs in place and in tension, the top Y-plate is threaded, and the bottom has a clearance hole. The bolts are screwed in to the top plate, travel through the clearance hole, and then each spring can be put under tension by tightening with a nut on the bottom. The main reason to keep these in tension is to create a thermal connection




Figure B.25: Connection of the safety structure to the mixing chamber plate, such that it is rigidly and thermally connected at three points.

between the spring and the Cu plates, as the interfacial thermal conductance is related to the force being applied at the interface.

Lastly, for the lower stage, we use an O-shaped circular disk (i.e. an annulus) to catch the lower stage in case of failure, as shown in Fig. B.27. As with the Y-shaped structure for the upper stage, this structure is attached to the pillars through L-shaped brackets, and there are six equally spaced in-tension springs that will compress if the lower stage falls. The mechanism to keep the springs in-tension and thermally connected is identical to the design of the structure for the upper stage, where we have two plates with inset regions and tightened bolts that allow for movement.

For the upper and lower stages, the masses are significant, and the expected compression of the springs must be checked against what is physically allowed. For an arbitrary compression spring of spring rate k , its maximum compression x_f will be given by some mass m being dropped from some height above the uncompressed spring h (or $h + x_f$ above the maximally compressed spring). Thus, we can use conservation of energy to solve for the expected maximum compression for any mass and any height

$$mg(h + x_f) = \frac{1}{2}kx_f^2. \quad (\text{B.131})$$




Figure B.26: The Y-shaped safety structure to catch the upper stage in case of failure. There are seven phosphor bronze compression springs which will reduce the impulse of the falling mass on the structure.

Solving for x_f , we have that

$$x_f = \frac{mg}{k} \left(1 \pm \sqrt{1 + 2 \frac{kh}{mg}} \right), \quad (\text{B.132})$$

where we throw out the answer corresponding to the negative sign (it is nonphysical as it would mean that the spring had extended). For our system, any stage would drop less than 5 cm after failure, so we can use Eq. B.132 to determine the expected maximally compressed lengths as a function of height above the uncompressed spring, as shown in Fig. B.28, where it is assumed that each stage is catching 20 kg of mass. This also takes into account that there are seven springs in series for the upper stage and six springs in series for the lower spring, where the effective spring constant increases additively (for n identical springs, $k_{eff} = nk$).

For the specific springs being used from McMaster (part number 9657K255, the only phosphor bronze compression spring that they have in inventory at the time of writing), the spring rate is $k = 10.4 \text{ N/mm}$, the initial length is 2.5 cm, its length at maximum load is 1.4 cm, and its fully compressed length is 1.3 cm (where fully compressed means that each coil




Figure B.27: The O-shaped safety structure to catch the lower stage in case of failure. Phosphor bronze springs are again used to reduce the impulse on the structure in case of failure, but in this case six springs are equally spaced in a circle.

has full contact and the spring cannot physically compress any more). Both of these latter two lengths are shown in the figure. Thus, if the spring compresses by more than 1.4 cm, then it should not be reused, as it likely has been compressed passed its yield strength and not longer has the same properties. Fortunately, we know these springs are cheap, so if they can only be used once, that is completely permissible. Furthermore, we designed our system to be safe from failure, so this occurrence should be extremely rare. Nonetheless, we should adjust the height of each stage of the safety structure such that the initial height of above the uncompressed spring is minimized, as is the impulse to the safety structure.

B.6 Final Remarks

At this point, we have fully designed a passive vibration isolation system specifically for the Pyle dilution refrigerator. This design will be implemented in the latter half of 2022, and should do an excellent job of reducing the vibrational sensitivity of our devices. It is my hope that the design principles that I have detailed are applicable beyond the Pyle group lab and




Figure B.28: The expected compressed length of the springs as a function of the height above the uncompressed spring, given by Eq. B.132. The values that fall below the solid black line are physically unachievable and represent that the spring will have fully compressed and need to be replaced after such a failure of some part of the passive vibration isolation system.

can act as a starting point for any group pursuing the reduction of vibrational sensitivity for their detectors. For our group, I am excited to see the final results, as well as the insight it will give us on the excess background signals that have been detailed and discussed in this thesis.

Appendix C

Optimal Filtering

The optimal filter (OF) appendix is effectively standard to be included in a CDMS-related thesis, as our energy reconstruction techniques rely on extracting pulse amplitudes from data with known noise. An alternative version of this appendix can be found, e.g., in Thomas Shutt’s thesis [250], Sunil Golwala’s thesis [117], Arran Phipps’ thesis [251], Noah Kurinsky’s thesis [114], and Nicholas Mast’s thesis [252]. Nevertheless, I mention the OF throughout this thesis, so it is useful to include the basic principles to ensure this work is self-contained, as well as include the general ideas behind more complex versions of the OF.

C.1 Optimal Filter Formalism

The OF formalism can be explained as the least-square minimization of the frequency-weighted chi-square between an expected pulse shape and some measured pulse. Here, we will quickly go through the various common results of a single-pulse OF, where we will use the following Fourier transform conventions (previously defined in Section 2.2.2), where we define the transform pair (for some function g) for the continuous case as

$$\tilde{g}(f) = \int_{-\infty}^{\infty} dt g(t) e^{-i\omega t}, \quad (\text{C.1})$$

$$g(t) = \int_{-\infty}^{\infty} df \tilde{g}(f) e^{i\omega t}, \quad (\text{C.2})$$

where $\omega \equiv 2\pi f$ for convenience, and for the discrete case as

$$\tilde{g}_n = \frac{1}{N} \sum_{k=-N/2}^{N/2-1} g_k e^{-i\omega_n t_k}, \quad (\text{C.3})$$

$$g_k = \sum_{n=-N/2}^{N/2-1} \tilde{g}_n e^{-i\omega_n t_k}. \quad (\text{C.4})$$

In this appendix, we will use the continuous approximation for readability, knowing that the calculations are all done in the discrete case.

Before we can extract pulse heights using the OF, we must make a few definitions. We will define $J(f)$ as the two-sided power spectral density (i.e. not summing the positive and negative frequencies) describing the noise of the system (recall the difference between one-sided and two-sided in the discussion at the beginning of Section 2.3.1), $s(t)$ as a template (i.e. expected shape) of the signal normalized to a maximum amplitude (pulse height) of 1, A as the amplitude of the signal, and t_0 as the start time of the signal. When actually recording data, the noise in the system (which can be represented by $n(t)$ in time domain) will mean that we, in practice, will measure for some event

$$v(t) = As(t) + n(t). \quad (\text{C.5})$$

As we have knowledge of $n(t)$ (i.e. $J(f)$) and $s(t)$, we should be able to extract what the best estimate of A is from this noisy signal, which is the purpose of the OF. When Fourier transforming $v(t)$ and $s(t)$, we will denote their frequency space counterparts as $\tilde{v}(f)$ and $\tilde{s}(f)$, respectively.

C.1.1 Extracting Pulse Heights

To extract pulse heights, we use frequency space and calculate the chi-square between the measured signal and the expected signal, weighting the frequencies by the PSD

$$\chi^2 = \int_{-\infty}^{\infty} df \frac{|\tilde{v}(f) - A\tilde{s}(f)|^2}{J(f)}, \quad (\text{C.6})$$

where A is as-of-yet undetermined. To find the best estimate of A , we minimize chi-square and solve for A , which gives the optimal estimator

$$\hat{A} = \frac{\int_{-\infty}^{\infty} df \frac{\tilde{s}^*(f)\tilde{v}(f)}{J(f)}}{\int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2}{J(f)}}. \quad (\text{C.7})$$

From this chi-square minimization, we can also find the variance in this best fit value through the general result of $\sigma_A^2 = \left[\frac{1}{2} \frac{\partial^2 \chi^2}{\partial A^2} \right]^{-1}$, which can be shown to be

$$\sigma_A^2 = \left[\int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2}{J(f)} \right]^{-1}. \quad (\text{C.8})$$

This is the best possible variance achievable given the assumed pulse shape and PSD, and is completely independent on pulse amplitude, as well as the data itself. Thus, given some signal template and known PSD, it is immediately calculable what best possible variance can be expected.

We have not allowed a time degree-of-freedom, so this amplitude estimation is akin to finding the amplitude corresponding to overlapping the expected signal template with the data, and scaling it to have the best fit (lowest chi-square). However, this is only useful if we know the exact location of the event, which will always have some uncertainty (or may be completely unknown). Fortunately, we can add a time degree-of-freedom, which we will do in the following section.

C.1.2 Extracting Pulse Heights and Pulse Times

With a time degree-of-freedom t_0 added to the chi-square, we can take into account the amount of shifting of the template from no shifting (i.e. no delay) to some time of best fit. The inclusion of t_0 is added to the template itself in the chi-square, such that we will minimize

$$\chi^2 = \int_{-\infty}^{\infty} df \frac{|\tilde{v}(f) - Ae^{-i\omega t_0} \tilde{s}(f)|^2}{J(f)}. \quad (\text{C.9})$$

We again minimize chi-square with respect to amplitude, which will result in

$$A(t_0) = \frac{\int_{-\infty}^{\infty} df e^{i\omega t_0} \frac{\tilde{s}^*(f)v(f)}{J(f)}}{\int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2}{J(f)}}, \quad (\text{C.10})$$

with a dependence of the pulse amplitude of the time-shift of the template. At this point, we make the definition of the optimal filter $\phi(f)$

$$\phi(f) \equiv \frac{s^*(f)}{J(f)}, \quad (\text{C.11})$$

such that Eq. C.10 becomes

$$A(t_0) = \frac{\int_{-\infty}^{\infty} df e^{i\omega t_0} \phi(f)v(f)}{\int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2}{J(f)}}. \quad (\text{C.12})$$

Note that the numerator is simply the inverse Fourier transform of $\phi(f)v(f)$, meaning that this is computationally simple to numerically calculate for all time-shifts of the signal template via a fast Fourier transform (a fact that we use repeatedly in our OF algorithms in QETPY [105]).

We can then directly put $A(t_0)$ into the chi-square and simplify the equation to

$$\chi^2(t_0) = \int_{-\infty}^{\infty} df \frac{|v(f)|^2}{J(f)} - A(t_0)^2 \int_{-\infty}^{\infty} df \phi(f)s(f), \quad (\text{C.13})$$

where it is immediately clear that the lowest chi-square has a one-to-one correspondence to the time shift degree-of-freedom that has the largest pulse amplitude. This is extremely




Figure C.1: (Left) The pulse template to be used for this example from simulated noise. (Right) The PSD to simulate noise from and add onto a shifted pulse template.

useful in that if we only need to maximize $A(t_0)$ to find the time delay t_0 , and this will correspond to the minimum chi-square.

On the variance of both A and t_0 , we have that σ_A^2 is again defined by Eq. C.8 (as the time offset is effectively part of the signal template and drops out when taking its modulus). For the variance of the time offset t_0 , we again use the general result from a chi-square that $\sigma_{t_0}^2 = \left[\frac{1}{2} \frac{\partial^2 \chi^2}{\partial t_0^2} \right]^{-1}$, and find that

$$\sigma_{t_0}^2 = \left[A^2 \int_{-\infty}^{\infty} df \omega^2 \frac{|\tilde{s}(f)|^2}{J(f)} \right]^{-1}. \quad (\text{C.14})$$

Where, in this case, there is a dependence on the signal amplitude, where $\sigma_{t_0} \propto 1/A$. That is, as the signal becomes larger, it becomes easier to find its location in time, as expected from an increase in the signal-to-noise ratio.

In this appendix, we kept all of the Fourier transforms in the continuous limit, as it is much easier to handle. However, these formulas that we have (quickly) derived can be converted to the discrete limit and applied to discrete data, which is what we do in practice.

C.1.3 An Optimal Filter Example

To see the actual application of the OF to an event, we can use a known pulse template added to noise generated from a PSD. In this case, we will use the pulse template and PSD as shown in Fig. C.1, where the template has a $20 \mu\text{s}$ rise time and a $66 \mu\text{s}$ fall time.

If we add noise generated from this PSD (using the function `gen_noise` in QETPY [105]) to the pulse template scaled to $1 \mu\text{A}$, we can apply the OF and extract amplitudes. We will also shift this pulse to the right (in time) by $160 \mu\text{s}$ to show the importance of the time-shift




Figure C.2: A simulated pulse, where we show the optimal filtered trace $A(t_0)$, where t_0 corresponds to a shift defined by the top x -axis. The bottom x -axis defines the time since the beginning of the trace. We have extracted both OF amplitude with no delay and OF amplitude with delay, showing the importance of allowing the start time degree-of-freedom for accurate determination of pulse heights. For the no delay fit, $\chi^2_0 = 211734$. For the fit with delay, $\chi^2 = 32767$, matching the expectation of $\chi^2 = 32768$.

degree of freedom. In Fig. C.2, we show the resulting optimal filtered trace (i.e. $A(t_0)$) as compared to the simulated event, as well the OF amplitudes with no time shifting and the with time shifting.

Thus, it is clear that, unless the pulse start time is known perfectly, the time-shift degree-of-freedom is absolutely necessary for accurate calculation of the pulse height of the event in the OF formalism. We also show in this figure both the time delay and the time elapsed within the trace as the top and bottom x -axes, as the two terms can be easily confused. When there is zero time delay in this example, it corresponds to extracting to OF amplitude at a little after 26.2 ms (i.e. the time corresponding to the start time of the pulse in the unshifted template).

C.2 Two-Pulse Optimal Filters

Frequently, when the data being taken has a high rate, there may be two events in a single trace (“pileup”). Extracting the pulse amplitudes and pulse heights for each pulse in a single




Figure C.3: (Left) Two $1\mu\text{A}$ amplitude pulses added to simulated noise 1.44ms apart and fit with an iterative pileup OF. The pulse heights are reconstructed well, as the two pulses are far apart (loosely defined as nonoverlapping). The chi-square for this fit is $\chi^2 = 32767$, matching expectation. (Right) Two $1\mu\text{A}$ amplitude pulses added to simulated noise $64\mu\text{s}$ apart and fit with an iterative pileup OF. The two pulses are overlapping, and the iterative pileup OF does a poor job of reconstructing the amplitudes of the pulses. The chi-square for this fit is $\chi^2 = 64073$, far from the expectation, indicating that this is not a good fit.

event can be numerically quite expensive. However, there are some different techniques that we can use to be able to speed of the numerical calculation, which we will discuss in the follow sections.

C.2.1 Iterative Pileup Optimal Filter

The iterative pileup OF is based on the assumption that the two pulses are far enough away from one another, such that there is no interference in the optimal filtered trace between the two pulses. Thus, one can simply run the single-pulse OF with the time-shifting degree of freedom, find one of the two pulses, subtract out the fit to that pulse, and then run an OF on the subtracted trace. If there are more pulses, this can be repeated indefinitely. In the OF formalism for two pulses, we would first minimize the chi-square in Eq. (C.9), and then use those best fit values again to minimize the iterative chi-square

$$\chi_{iter.}^2(A_1, t_1) = \int_{-\infty}^{\infty} df \frac{\left| \tilde{v}(f) - \hat{A}_0 e^{-i\omega t_0} - A_1 e^{-i\omega t_1} \tilde{s}(f) \right|^2}{J(f)}. \quad (\text{C.15})$$

Within the OF formalism as discussed in Section C.1, the minimization of the chi-square to solve for the best fit values of A_1 and t_1 is the same, but with the substitution of $\tilde{v}(f) \rightarrow \tilde{v}(f) - \hat{A}_0 e^{-i\omega t_0}$.

In Fig. C.3, we show two applications of the iterative pileup OF to events that are spaced far apart and events that are spaced close together. The fit to the closely-spaced pulses poorly reconstructs the pulse amplitudes and demonstrates that a better-performing OF is needed in this regime.

C.2.2 Simultaneous Pileup Optimal Filter

To reconstruct amplitudes of overlapping pulses, the OF needs to simultaneously vary the pulse amplitudes and time offsets. Thus, the chi-square to be minimized becomes (where N is the number of pulse amplitudes to reconstruct)

$$\chi_N^2(\mathbf{A}, \mathbf{t}) = \int_{-\infty}^{\infty} df \frac{\left| \tilde{v}(f) - \sum_{j=1}^N A_j e^{-i\omega t_j} \tilde{s}_j(f) \right|^2}{J(f)}, \quad (\text{C.16})$$

where $\mathbf{A} = (A_1, A_2, \dots)$ is a vector containing the pulse amplitudes, and $\mathbf{t} = (t_1, t_2, \dots)$ is a vector containing the pulse start times.

We start by minimizing this chi-square with respect to an arbitrary amplitude dimension A_k , such that

$$\frac{\partial \chi_N^2}{\partial A_k} = \int_{-\infty}^{\infty} df \left[-\frac{\tilde{s}_k^*(f) e^{i\omega t_k} \tilde{v}(f)}{J(f)} + \frac{\tilde{s}_k^*(f) e^{i\omega t_k} \sum_{j=1}^N A_j \tilde{s}_j(f) e^{-i\omega t_j}}{J(f)} \right] = 0, \quad (\text{C.17})$$

where we have removed an overall factor of 2. This creates a system of equations relating all of the the amplitudes, which we can define more easily in terms of Einstein notation. That is, we define

$$P_{kj}(t_k, t_j) \equiv \int_{-\infty}^{\infty} df \frac{\tilde{s}_k^*(f) \tilde{s}_j(f) e^{-i\omega(t_j - t_k)}}{J(f)}, \quad (\text{C.18})$$

$$q_k(t_k) \equiv \int_{-\infty}^{\infty} df \frac{\tilde{s}_k^*(f) e^{i\omega t_k} \tilde{v}(f)}{J(f)}, \quad (\text{C.19})$$

and can relate these quantities to an amplitude index of A_j by

$$P_{kj}(t_k, t_j) A_j = q_k(t_k). \quad (\text{C.20})$$

Note that we have kept the time degrees of freedom explicit, as we have not yet minimized with respect to them. Solving for A_j means that we need to invert the P matrix, where the solution would be

$$A_j(\mathbf{t}) = [P(t_k, t_j)]_{jk}^{-1} q_k(t_k). \quad (\text{C.21})$$

Thus, we have that every pulse amplitude is dependent on the start times of itself and the other pulses.

At this point, we will now restrict ourselves to the case of $N = 2$ pulses, such that we can directly compare to the iterative pileup OF, as well as set all of the pulse templates to be the same (i.e. $\tilde{s}(f) \equiv \tilde{s}_1(f) = \tilde{s}_2(f)$). In this case, the matrix P is

$$P_{N=2}(t_1, t_2) = \begin{bmatrix} \int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2}{J(f)} & \int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2 e^{-i\omega(t_1-t_2)}}{J(f)} \\ \int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2 e^{-i\omega(t_1-t_2)}}{J(f)} & \int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2}{J(f)} \end{bmatrix}. \quad (\text{C.22})$$

Thus, this matrix must be calculated and then inverted for every combination of t_1 and t_2 in our simulated trace in order to find all of the possible pulse amplitudes as a function of time, and then find the best chi-square, as we did for the single-pulse version. However, if we blindly did this for the entire trace, we would quickly run into a problem: our computer likely does not have a enough RAM to store this. The traces we are using are 32768 bins long, and we can calculate the number of combinations of pairs of pulse times using the binomial coefficient

$$\binom{n}{k} = \frac{n!}{k!(n-k)!}, \quad (\text{C.23})$$

where we find that $\binom{32768}{2} = 536,854,528$ combinations. Thus, to populate this many combinations of a 2×2 matrix of (8B) floating values would take up 17GB of RAM, not to mention that we then have to invert this many matrices. Thus, we must restrict the range of combinations that we will use to create these matrices.

One way to do this is to go back to our P matrix and estimate a range where the off-diagonal components of the matrix are not small, such that the inversion outside of this range is effectively the iterative pileup OF (these large nonzero off-diagonal components represent the interference between pulses when fitting). If we define $\Delta t = t_2 - t_1$, then we see that the off-diagonal component can be written as

$$P_{0,1}(\Delta t) = \int_{-\infty}^{\infty} df \frac{|\tilde{s}(f)|^2 e^{i\omega\Delta t}}{J(f)}, \quad (\text{C.24})$$

which is the inverse Fourier transform of $|\tilde{s}(f)|^2 / J(f)$, and is thus fast to calculate computationally. Plotting this as a function of Δt , and normalizing to 1, we have that the behavior of the off-diagonal component is shown in Fig. C.4. Because of the symmetry of this matrix, this normalized off-diagonal component provides a good estimate of the error of the matrix if we were to assume that it was simply diagonal [253]. We expect that, if two pulses are spaced further than $100 \mu\text{s}$ apart, then the iterative and simultaneous versions of the OF should have similar performance, as the error in P^{-1} is estimated to be less than 7%. Thus, we should restrict the combinations of t_1 and t_2 to the roughly the range of $|\Delta t| < 100 \mu\text{s} \equiv |\Delta t_{\text{max}}|$. To know where to calculate this range, it is useful to run a simple OF first for an estimate for where the pileup pulses are in time t_{start} . Given this time, we have the range of times to calculate the pileup OF as $[t_{\text{start}} - |\Delta t_{\text{max}}|, t_{\text{start}} + |\Delta t_{\text{max}}|]$. Using the 7% cutoff, this reduces




Figure C.4: The off-diagonal component of $P_{N=2}$ as a function of Δt . This gives an estimate of the expected error in the inversion of the matrix as compared to if the matrix were diagonal. For the pulse templates that we are using, the error is less than 7% for pulses that are spaced greater than $100 \mu\text{s}$ apart.

the combinations of times that we need to calculate and store to 7,875, about 5 orders of magnitude as compared to without a time range restriction. With a manageable amount of matrices to store and invert, the simultaneous pileup OF can be run on the closely spaced pulses that the iterative pileup OF could not reconstruct.

In Fig. C.5, we compare the simultaneous and iterative pileup OFs as applied to the closely spaced pileup pulses. The simultaneous pileup OF reconstructs the pulse amplitudes very well, showing the benefit of varying the pulse amplitudes and start times all at once.

C.3 Arbitrary Number of Pulses

The natural next step is to apply the same OF formalism to an arbitrary number of pileup pulses. However, if we were to follow the same method of restricting ranges, then we would run into computational issues extremely quickly. That is, we would need to take the combinations of N pulses of the time range chosen to calculate. If we stay under 10,000 combinations of the start times of $N = 3$ pulses, this corresponds to a time range of 40 bins, or $64 \mu\text{s}$. This is less than the fall time of a single simulated pulse, which would be too small of a time range to extract three possible pileup pulses.

Thus, we must abandon the method of precomputing these matrices and carrying out computationally efficient robust searches along all time degrees-of-freedom. Because of the




Figure C.5: The two $1\ \mu\text{A}$ amplitude pulses added to simulated noise $64\ \mu\text{s}$ apart and fit with both the iterative pileup OF and the simultaneous pileup OF. The simultaneous pileup OF does an excellent job of reconstructing the pulse heights, with goodness-of-fit $\chi^2 = 32463$.

nonlinearity of the time degrees-of-freedom, a simple gradient descent algorithm will not reliably minimize the chi-square. As such, a derivative-free optimization algorithm will be needed, with an excellent option being differential evolution, originally introduced by Storn and Price [254]. The basic idea of this algorithm is that, given some function minimize, it iteratively improves the candidate solution through a random sampling of the variable space. Though it does not guarantee finding the optimal solution, it can be tuned such that it is quite likely in the case of the N pulse OF.

Returning to the N pulse chi-square, we have that we want to minimize

$$\chi_N^2(\mathbf{A}, \mathbf{t}) = \int_{-\infty}^{\infty} df \frac{\left| \tilde{v}(f) - \sum_{j=1}^N A_j e^{-i\omega t_j} \tilde{s}_j(f) \right|^2}{J(f)}, \quad (\text{C.25})$$

for which we have shown the method in Section C.2.2. That is, we need to invert the matrix P in order to find the pulse amplitudes at given times. We know that the pulse amplitudes depend on the start times of all pulses, such that we can calculate them for any given pair of times

$$A_j(\mathbf{t}) = [P(t_k, t_j)]_{jk}^{-1} q_k(t_k), \quad (\text{C.26})$$

where we again are reproducing from Section C.2.2. If we plug this equation into the chi-




Figure C.6: Three $1\,\mu\text{A}$ amplitude pulses added to simulated noise with $64\,\mu\text{s}$ spacing and fit with differential evolution-based pileup OF. This pileup OF reconstructs the pulse heights and times well, with goodness-of-fit $\chi^2 = 32407$.

square, we have

$$\chi_N^2(\mathbf{t}) = \int_{-\infty}^{\infty} df \frac{\left| \tilde{v}(f) - \sum_{j=1}^N A_j(\mathbf{t}) e^{-i\omega t_j} \tilde{s}_j(f) \right|^2}{J(f)}, \quad (\text{C.27})$$

where it is now solely a function of the time degrees of freedom. When discussing the simultaneous OF, this is where we set $N = 2$ and then precomputed the P^{-1} matrices. Because we cannot do this efficiently for $N > 2$, this is where we go to the differential evolution algorithm to minimize the chi-square.

Because differential evolution randomly samples, we can improve the likelihood of finding the global minimum of the chi-square by restricting the range of time values that it considers to where we believe there might be pileup. We can use the same concepts as Section C.2.2 to set it on the estimated errors of the matrices, where we increase the ranges by factors of the number of pileup pulses to extract, in order to ensure we allow all pulses to be spaced by $|\Delta t_{\max}|$.

In Fig. C.6, we show the result of a three pulse fit using differential evolution (via SCIPY [224]) to find the minimum chi-square in the time range around the expected location of the pileup. We see that the algorithm found the pulse heights and times quite well, without needing to do a robust grid search. This algorithm can also scale to any arbitrary number of pulses.

Furthermore, the differential evolution algorithm in `SCIPY` includes an option for creating constraints, based on work by Lampinen [255]. For instance, we might want to constrain the pileup pulses to only have positive amplitudes, if we are in a biasing setup of the TES such that real pulses only go in that direction. Thus, constraints like this or others are easily added, but may significantly slow down the computation time depending on the complexity.

C.4 Optimal Filter Code in QETPY

The various OF codes discussed in this Appendix have been written in `PYTHON` by me and placed in various locations in our open-source package `QETPY` [105]. Although the methods herein are not particularly novel, their existence as usable and publicly available code is an excellent place to start using them without needing to write the algorithms from scratch. Below, I direct the reader to the following classes in `QETPY`

- `qetpy.OptimumFilter`: contains the single-pulse OF and the iterative pileup OF
- `qetpy.PileupOF`: contains the two-pulse simultaneous pileup OF
- `qetpy.PileupDE`: contains the differential evolution–based pileup OF for an arbitrary number of pulses.

For the simultaneous pileup OFs, they do have a built-in assumption that the pulses all have the same template. When analyzing real data, this may not be the case based on the physics involved. This is one possible upgrade to these existing algorithms, which I leave to future `QETPY` collaborators.