

THE LOW-MASS LIMIT: DARK MATTER DETECTORS
WITH EV-SCALE ENERGY RESOLUTION

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF PHYSICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Noah Kurinsky
August 2018

© 2018 by Noah Alexander Kurinsky. All Rights Reserved.
Re-distributed by Stanford University under license with the author.



This work is licensed under a Creative Commons Attribution-Noncommercial 3.0 United States License.

<http://creativecommons.org/licenses/by-nc/3.0/us/>

This dissertation is online at: <http://purl.stanford.edu/xh271zf3698>

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Blas Cabrera, Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Richard Partridge, , Co-Adviser

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Peter Graham

I certify that I have read this dissertation and that, in my opinion, it is fully adequate in scope and quality as a dissertation for the degree of Doctor of Philosophy.

Giorgio Gratta

Approved for the Stanford University Committee on Graduate Studies.

Patricia J. Gumport, Vice Provost for Graduate Education

This signature page was generated electronically upon submission of this dissertation in electronic format. An original signed hard copy of the signature page is on file in University Archives.

Abstract

The SuperCDMS SNOLAB experiment will be a 20-kg scale Si and Ge direct dark matter detection experiment designed to probe down to 300 MeV in dark matter (DM) mass through DM-nucleus scattering and 500 keV in DM electron scattering. In order to reach these low masses with appreciable sensitivity to dark matter, it needs to achieve very low energy resolution (≤ 10 eV) for nuclear recoils in both detector materials, which will be achieved using a new detector design and operating mode, CDMS HV. This detector is designed to operate at a bias of 100V to convert charges liberated in our detector targets into phonon energy in order to resolve individual electron-hole pairs. This has never before been achieved in a kg-scale detector.

In this thesis, I cover three elements of the design of the CDMS HV detectors. I discuss the detector physics controlling how charges and phonons are generated in our detector crystals, comparing theory to results of recent experiments carried out at Stanford. I move on to describe the operating principles of our phonon-mediated charge readout, as well as the design of the CDMS HV detector. I then describe the performance tests of early CDMS HV prototypes in conjunction with the SuperCDMS SNOLAB electronics, and discuss the path towards achieving single electron-hole pair resolving detectors at the kg-scale given the performance obtained thus far. As a result of these tests, we were able to refine our noise and sensor dynamics models, and develop new metrics for diagnosing non-ideal sources of noise to aid in reducing coupling of the external environment to our detectors.

In order to study the microphysics of phonon and charge production in our target crystals, we fabricated a number of gram-scale devices with various sensor designs in order to separate sensor and environmental effects from intrinsic crystal properties. These devices provided the first successful demonstration of using voltage to amplify charge energy by production of phonons (the Neganov-Trofimov-Luke effect) in order to resolve electron-hole pairs, and opened up a new regime of dark matter and photon science at the gram-scale that we are just beginning to explore. A first dark matter search was carried out with one of these gram-scale devices, producing world-leading limits on electron-recoiling dark matter between 0.5 and 5 MeV in dark matter mass for multiple form factors. This device achieved a phonon resolution of 10 eV, allowing a single gram-day of exposure to rival kg-days of exposure in the competing liquid-noble based electron-recoil search.

Preface

This thesis explores the physics of low-temperature crystals, and lays out design and testing principles for cryogenic calorimetry based on CDMS technology. Much of the text was written as the work was done, and many of the small asides or conclusions were reached as a result of discussions with collaborators. The concluding chapter represents my own opinions and conclusions about the next steps for our technology, but reflects, at the time of this writing, a general consensus in the field. Continued development of these precise sensors has reached the point that we can make quick gains in dark matter science, and use these developments to advance other fields of physics such as the study of coherent interactions and applied photonics.

The work presented in this thesis was performed with help from and in collaboration with a large fraction of the SuperCDMS collaboration, and none of the results would have been possible without the fabrication, cryogenics, and physics expertise of my collaborators. In particular, none of these detectors could have been run without the fabrication work done by Matt Cherry, Astrid Tomada, Paul Brink, Mark Platt, and Roger Romani at Texas A&M and Stanford. Testing of devices was done at many test facilities in collaboration with many people:

- University of Minnesota - Matt Fritts, Anthony Villano, Nick Mast, and Allison Kennedy, under the direction of Vuk Mandic and Prisca Cushman), with help from Bill Page from UBC
- UC Berkeley - Suhas Ganjam, Caleb Fink, Sam Watkins, Bruno Serfass, Bernard Sadoulet, and Matt Pyle. This group was also instrumental to taking much of the SLAC data, and there would be no data in this thesis if it weren't for Bruno's help with software tools and DCRC operation.
- Stanford - Jeff Yen, Robert Moffatt, Betty Young, Jon Leyva, Trevor Howarth, James Allen, Steve Yellin, Francisco Ponce, Chris Stanford, and Francesco Insulla, under the direction and with significant help from Blas Cabrera.
- SLAC - Paul Brink, Mike Racine, Tsuguo Aramaki, Mike Kelsey, Gary Godfrey, Pelle Hanssen, Dave Nelson, and Leo Munger, as well as our many SULI and INFN students who helped in the lab or with data taking (Andrea Caputo, Angela Bai, Carlo Gilardi, and Chiara Magliocca), under the direction of Richard Partridge.

- Texas A&M - Jorge Morales and Jon Wilson, who tested and diagnosed DCRC issues and implemented firmware, and Xuji Zhao, who worked on SQUID and Tc testing at SLAC as a visitor.
- University of British Columbia (UBC) - Bill Page, Danika MacDonell, Belina von Krosigk, who helped get the DAQ working at SLAC
- USD - Joel Sander and Amy Roberts, who helped with DAQ setup and software tools

These are only the people actually involved in setting up and running the facilities for the detectors described in this thesis, and does not include those who came previously to help get the facilities running or provided input on the analysis.

Each chapter also includes significant work and insight from many collaborators that I worked with during my years as a graduate student:

- Chapter 1 - I owe a great debt to Dylan Rueter for his patience teaching this experimentalist how to appreciate the finer points of particle physics phenomenology, to Joe DeRose and Sean McLaughlin for references and explanations of the phenomena discussed in the cosmological sections, and to Warren Morningstar for his obsession with gravitational lensing.
- Chapter 2 relied heavily on measurements taken by Robert Moffatt using the device he built for his dissertation [75], and the theoretical treatment of scattering depends largely on work done by Kyle Sundqvist for his dissertation [99]. Robert, Kyle, Betty Young, and Blas all contributed to the analysis. In addition, the intervalley scattering model benefited greatly from discussions with Alexandre Broniatowski during his visit to Stanford. Mike Kelsey and Rob Agnese wrote the code that the simulations were done with and I worked closely with them to do consistency checks and add features in order to enable the code to simulate Si as well as arbitrary crystal orientations. The initial implementation of some the scattering formulae was done by Ramiro Garcio, with follow-up work by Francesco Insulla. The initial investigation of hole anisotropies was done with James Allen, and was the subject of his senior thesis. The final results and data were taken in part by Chris Stanford, who helped me complete this study and is continuing to push the technique to learn much more about charge transport.
- Chapters 3 and 4 were written during the year I was working on the detector design, and I'm entirely indebted to Matt Pyle for his constant guidance and inspiration, as well as Paul Brink and Blas for their insight, feedback, and suggestions for various optimizations to try. The noise analysis in particular relies heavily on some notes Paul put together to help us more easily understand the various features of noise spectra, and we relied on Matt's thesis as well as Kent Irwin and Gene Hilton's seminal paper [50] to put those together. The design process was led by myself but many people in the tower technical meetings provided vital insight and answers to questions about various experimental constraints. Bruce Hines and Martin Huber

helped me appreciate the subtleties of the SQUID readout, and Tsuguo provided key initial measurements of the cold electronics to give us an idea of what was actually feasible in terms of our parasitics. Sunil Golwala, Rich Partridge, and Bernard Sadoulet provided very useful comments and feedback during meetings and individually, and each helped me appreciate the various parts of the phonon and charge readout electronics as I tried to figure out which parts of the design needed to be optimized and how to do so. Betty was always a resource for helping me understand non-idealities of fabricating superconducting thin films.

Many of my students during the first two years worked on various parts of detector design and simulation, including Andrea Caputo, who helped with initial studies of channel layout, and Eduardo Montano who worked on TES simulations. I'm also indebted to Peter Redl and Kristi Schneck who helped me run the detector Monte Carlo for the initial CDMS HV studies that the optimization was based on. Peter also helped me get started with COMSOL simulations, and Angela Bai did some great work on axisymmetric electric field modeling that helped inform the final design; this work was most recently picked up again by Madison Matsen. Jeff Yen's quasi-particle transport data [111] with Robert's analysis also informed the QET efficiency model used for the first time in device design for this thesis. Finally, Paul and Matt Cherry helped coach me through the layout process, suggesting design decisions that improved the fabrication yield of the sensors and made actual operation of the mask easier and less error prone.

- Chapter 5 is a summary of the work of many people I listed in the test facility section of this thesis. Much of the DCRC investigation was led by Tsuguo, with help from myself, Matt Pyle, Caleb Fink, Bruno, Paul, and Rich. Rich helped all of us understand how to track down places where components may not be behaving ideally, and helped solve some of the most egregious noise issues. The resistivity and T_c investigations were done with Matt Cherry, Jeff Yen, and Paul. Leakage investigations were undertaken with Matt Fritts, Nick Mast, Francisco Ponce, Roger Romani, Matt Pyle, Blas, Paul, and with initial work done at Berkeley by Bill Page. Modeling of the leakage was a collaboration between Steve Yellin, Matt Pyle, Bernard, and myself. A lot of work subsequent to my initial complex impedance studies was done by Sam Watkins which added confidence to some of my conclusions, and will be vital to understanding the performance of future detectors. The small detector discussed at the end of the chapter was fabricated and run by Roger Romani for his senior thesis, and the testing setup was largely designed by him, with subsequent improvements by Francisco. Many of later noise studies, and a lot of help processing the data, came from To Chin Yu, who in my last year picked up the BlueFors testing torch and has continually put out very interesting testing results that are shown here and go beyond the testing I describe.
- Chapter 6 presents the results of a collaboration paper, and as a result includes input from

a large section of the collaboration. The core analysis team consisted of myself, Francisco Ponce, Belina von Krosigk, Andrew Scarff, Matt Wilson, and Chris Stanford, with help and input from Rob Calkins, Richard Germond, Blas, Matt Pyle, Rich, Paul, Betty, and Steve Yellin. I appreciate the competition we got from the SENSEI collaboration, which helped us push out a high quality result in a relatively short amount of time; by my count, it was about 2 months from end of data taking to submitting the paper.

The scale of this sort of work is such that I was a cog in a very efficiency machine, as it is not possible for complex devices such as our detectors to be made and run without contributions from many people. It is my hope that this preface conveys the magnitude of contributions to the work presented here from the members of my collaboration, and the degree to which the SLAC, Stanford, and UCB groups collaborate on most studies happening in the bay area. The interaction between the different research groups created an environment where new ideas could be explored constructively, and facilitated much better results than we would have been able to achieve without so much external support. As the saying goes, two heads are usually better than one. I posit that the more heads you have, the better.

Acknowledgments

In the preface I said I owe a huge professional debt to my collaborators, and I also owe massive personal debts to the people who supported me during my grad school career, as I wrote this thesis, and while I tried to figure out where I'd go for the next phase of my career.

First and foremost, I need to thank my wife, Caitlyn, for the support from all the way across the country even though she was buried under her own load of work in medical school. My attending Stanford while she was in at school first in Maine, then New Jersey, then all over the country, meant we were a long distance couple for four very long years, which was difficult for both of us. Thank you for getting on the plane to come visit, for spending your summers in California even though your friends were back east, and for making the time to enjoy life with me. Our trips to wine country, to the city, or just down the road to Bierhaus kept me sane when I lost perspective, especially during the 'deep dark pit of despair' that our group all went through at the beginning of third year, when the end of grad school seemed impossibly far away, and we felt like we had nothing to show for all of our effort. We made some amazing memories all over the country, and in some new parts of the world, in four short years. I don't think I would have even gotten the opportunity to do this work if you hadn't been there to push me to be my best from the first day of college.

To my parents, Lori and Phil, and my brother, Ethan, who endured my years of never being satisfied with simple answers to obvious questions. If there's one thing I've learned from you guys, it's to figure out what your passion in life is, pursue it with all you can, and don't sweat the small obstacles that get in your way. The pragmatism I learned from our family kept me on the right track, and is certainly a big reason I was able to get a PhD in four very short years. And now that I have a PhD in Physics, I can tell you for sure now that the moon does not follow the car because it's tied to the back by a string...or at least we have yet to detect the string to the precision of our current tests. Maybe it's a cosmic string?

To my bay area crew, Dylan Rueter, Joe DeRose, Sean McLaughlin, Warren Morningstar, Krishna Soni, Aly Thompson, Lesje Atkinson, you guys made grad school so much fun when it had no right to be. I will miss the many outings to Fieldwork, Superbowl parties (we didn't really need a keg any of those times), nights out in SF, trips to Tahoe, and BBQs at the Jewell house. To my SLAC crew, Mike Baumer, Chris Davis, Devon Powell, and Dylan (again), I looked forward to our

coffee breaks that always devolved into absurd hypothetical situations (with the exception of our blimp rental service, which will remain unnamed here, and I maintain was a great idea), and I think I would have gone insane in my cube if I didn't have you guys to bother. Chris and Devon, the thesis 'writing' parties made this whole defense process seem fun, if even for a few minutes, and even if we did more drinking than writing. To Chris and Mike, we'll always share that one summer where we all got married because apparently we're followers and insane enough to get married during grad school. To my friends from back home, Broc, Greg, Matt, Colin, who always made time for me and Caitlyn on last minute trips to Boston with less than a day's notice, your company was always a refreshing break from the slog of this PhD process.

To my advisors, Rich and Blas, who let me get away with a four-year dissertation. I will still remember that week at the end of my third year where you both started asking about my career plans, and the panic that ensued when I realized you hadn't just forgotten I was only three years in. I can't thank you enough for giving me the opportunity to participate in such exciting research, and for helping make my four years such an amazing learning experience. Rich, you taught me how to stand behind my ideas and make my arguments airtight, and how to really work as a physicist in a team, and you put me in a position to succeed early by involving me in such a vital part of the experiment. Our group at SLAC was great, if not a bit odd, and showed me the difference being at a national lab makes to the quality of the work you can do as a physicist. I also promise from here on out to remember that 'which' and 'that' are not interchangeable, which is an important lesson that I took from writing this thesis (see what I did there?). Blas, I will never cease to be amazed by the depth of your knowledge and your ability to listen critically to ideas from everyone in your group, regardless of seniority or experience. When I first came to your group meeting, I was intimidated, but I immediately felt that it was a safe space I could feel comfortable to be wrong in, which I was a lot. I see many students come in and have the same reaction, and I think it is rare to find an advisor who can create that sort of productive but comfortable environment. I also intend to keep asking you obscure questions, and I fully expect to get obscure papers from the 50s in response, or hand-written notes that have been sitting on various digital storage media since I was in diapers.

To my unofficial advisors, Paul and Matt Pyle, who I spent so much time with during my PhD and who sat with me for hours making sure we were all on the same page. Paul, I'm sorry I can't barge in on you once a day to interrupt your thought process, but I hope To Chin has taken over that role, I'd hate for you to be able to finish something uninterrupted. Thanks for being my gut check on hard problems, keeping me on track, and for making working at SLAC more fun. Most importantly thank you for making it look so easy to turn a fridge around, and to keep a giant hunk of copper and aluminum sitting at 8 mK for months while I wasted time and electricity failing to measure things. We definitely learned a lot about what does not work. One of my favorite days at SLAC was when we decided it was a good use of our time to find wire in the garbage to make a loop antenna to look for EMI, and then ran around the lab unplugging things to see whether lines

went away. That is now what I aspire to make every day like. Matt, thanks for taking me into your group when nothing at SLAC was working and spending your time teaching me literally everything I know about how our detectors work. We spent so much time on the phone after hours talking about detectors that my room mates knew who you were, and learned to expect to find me debating systematics with you on your walk home. Your work ethic is hard to match, which was a challenge I enjoyed but mostly failed at doing, and your passion for detector physics is infectious. It's been fun working for and with you and I look forward to more years doing so.

To the Stanford group then and now: Betty Young, Steve Yellin, Matt Cherry, Francisco Ponce, Chris Stanford, Roger Romani, Peter Redl, Jeff Yen, and Robert Moffatt. Thanks for taking me under your wing and letting me drool over your data, sorry I never learned how to do a Helium transfer! I was spoiled to be able to walk into meetings, suggest a bunch of runs to do, and have the data show up on SUF only days later. You all told me you were happy you didn't have to analyze it, but I wasn't the one who had to go into work on the weekend. To Francisco especially, thanks for babysitting the fridge for so many weeks while we tried to get our dark matter result out, and thanks to Betty for thanklessly filling in when nobody else could. To Roger: keep doing what you've been doing and you'll go anywhere you want; you blew us all away.

To the SLAC group then and now: Dennis Wright, Mike Kelsey, Tina Cartaro, Tsuguo Aramaki, To Chin Yu, Kristi Schneck, Dave Nelson, Pelle Hansson, Rob Cameron, Gary Godfrey, Ken Fouts and Mike Racine. You are all so different, and it made this one of the most entertaining group meetings to attend. I was lucky enough to spend time one on one with all of you, and got some great career and life advice. I will miss talking to Mike too loud, asking Tina about her experience with the mob in Italy, and making Tsuguo drink beer on shift even though he knew he was going to pass out after one drink (sorry!). Dennis, sorry I spilled coffee all over your office. Ken, I expect to still talk hockey when you visit Fermilab. Kristi, you got me started and taught me how to navigate life as a grad student both at Stanford and within the collaboration; we didn't overlap for long, but a lot of your advice saved me significant pain just knowing I wasn't the only one who had to deal with some of the issues that arose as I worked my way through the PhD. To Chin, I've tried to prepare you to take over - good luck with these guys!

To the Berkeley group, my third home, then and now: Sam Watkins, Caleb Fink, Bruno Serfass, Bernard Sadoulet, Suhas Ganjam, Todd Doughty, and Arran Phipps. I enjoyed my occasional visits to UCB and lunches at the Thai restaurant that I'll never know the name of, and our outing to play that weird French Bocce game that Bernard loves. Thanks for letting me crash your office for the month I lived up there. Sam and Caleb, it was great to commiserate with you guys about being the only bay area grad students, which twice as many PIs as students (meaning we were expected to each do four students worth of work). Thankfully Bruno was there to show us how to do the work of four people and make it seem easy.

To all the other people I worked with so closely in the collaboration over the past few years,

in no particular order: Bill Page, D'Ann Barker, Brett Cornell, Mark Pepin, Matt Fritts, Anthony Villano, Nick Mast, Ziqing Hong, Tom Ren, Jon wilson, Dave Toback, Belina von Krosigk, Rob Calkins, Todd Doughty, Amy Roberts, Elias Lopez-Asamar, Matt Wilson, Tyler Reynolds, Tarek Saab, Rob Agnese, Dan Jardin, Alan Robinson, Lise Wills, Ray Bunker, Sunil Golwala, Lauren Hsu, Andrew Scarff, and many others. I was grateful to be able to work so closely with so many great people. I look forward to being able to continue working with so many of you as we continue getting SuperCDMS SNOLAB up and running and churn out the next generation of exciting dark matter science.

Finally, thanks to my new groups, both at Fermilab and Northwestern, for letting me hang around for the last few months I was writing this thesis, and for giving me the opportunity to continue the work started in this thesis as a postdoctoral fellow. Looking forward to a couple of great years in Chicago!

Contents

Abstract	iv
Preface	v
Acknowledgments	ix
1 The Need for New Physics	1
1.1 The Cosmos in Motion	2
1.2 Missing Mass: The Case for Dark Matter	9
1.2.1 Rotation Curves and the Virial Theorem	10
1.2.2 Modified Gravity	20
1.2.3 Gravitational Lensing	22
1.2.4 Cosmic Microwave Background Radiation	24
1.2.5 Measurements of Local Mass Density	28
1.3 Dark Matter Phenomenology	30
1.3.1 WIMP Dark Matter	32
1.3.2 Sub-GeV Dark Matter	38
1.3.3 Bosonic Dark Matter	41
1.3.4 Other Dark Matter Candidates	44
1.4 Physics Reach of Low-Threshold Detectors	45
2 High Voltage Charge & Phonon Dynamics	49
2.1 Introduction to Band Structure	50
2.2 Electrons: Conduction Band Structure	52
2.2.1 Effective Mass	55
2.2.2 Herring-Vogt Transform	57
2.3 Holes: Valence Band Structure	58
2.3.1 Equal Energy Surfaces	61
2.3.2 Hole Effective Mass	63

2.4	Electron Scattering Processes	66
2.4.1	Impurity Scattering	67
2.4.2	Phonon Scattering	70
2.4.3	Total Scattering Rate	74
2.5	Electron Intervalley Scattering Rate	76
2.5.1	IV Scattering from Phonons	77
2.5.2	IV Scattering from Impurities	78
2.5.3	Total IV Rate Predictions	79
2.6	Results from Stanford Charge Transport Measurements	80
2.6.1	Experimental Setup	80
2.6.2	Intervalley Scattering Rate	81
2.6.3	Hole Anisotropies	88
2.7	Future Charge Transport Studies	90
3	Detector Concepts	93
3.1	Detector Overview	94
3.2	TES Dynamics	97
3.2.1	TES Characteristics and Parameters	99
3.2.2	TES Response Model	104
3.2.3	Electrothermal Oscillation Criteria	114
3.2.4	TES Stability	117
3.2.5	Thermal Phase Separation	119
3.2.6	Additional Internal Degrees of Freedom	120
3.3	TES Noise Modeling	121
3.3.1	Johnson Noise	121
3.3.2	Thermal Fluctuation Noise	123
3.3.3	SQUID Noise	124
3.3.4	Joint Noise Sources	125
3.3.5	Energy Resolution	126
3.4	QET Modeling	126
3.4.1	Energy Conversion Efficiencies	126
3.4.2	Aluminum Fin Length	127
3.4.3	Tungsten-Aluminum Overlap Dimensions	132
3.4.4	Energy Collection Efficiency Summary	135
3.5	Future Modeling Inputs	135

4	CDMS HV Detector Design	137
4.1	Detector Dimensions	138
4.2	CDMS HV QET Design	140
4.2.1	TES Channel Optimization	140
4.2.2	QET Efficiency	142
4.2.3	Energy Resolution	145
4.2.4	Expected QET Performance	150
4.3	HV Detector Simulation	153
4.3.1	COMSOL Electric Field Modeling	153
4.3.2	Charge and Phonon Simulations	160
4.3.3	TES Pulse Simulation	162
4.4	Phonon Channel Layout Optimization	164
4.4.1	Proposed Geometries	165
4.4.2	Vertical Partition Performance	167
4.4.3	Radial Partition Performance	168
4.4.4	Combined Fiducialization	169
4.5	Mask Design Results and Future Studies	171
5	Prototype Detector Performance	174
5.1	Detector Prototype Program	174
5.1.1	TES Resistance Measurements	175
5.1.2	Critical Temperature and Bias Power	179
5.1.3	TES Transition Parameters	181
5.1.4	TES Measurements Summary	188
5.2	Characterizing QET Noise	189
5.2.1	G101c: First Studies at UMN	189
5.2.2	G115 at SLAC and UCB	191
5.2.3	Diagnosing Readout Noise	193
5.2.4	G124 at SLAC: Beating Down the Noise	202
5.2.5	Scaling Noise with Detector Volume	207
5.3	Phonon Energy Resolution	209
5.3.1	Measuring QET Efficiency	209
5.3.2	QET Saturation	211
5.3.3	CDMS-HV Prototypes at UMN	212
5.3.4	G115 at SLAC and UCB	214
5.3.5	G124 at SLAC	216
5.3.6	HVeV Detector at Stanford	220
5.4	High Voltage Performance	223

5.4.1	Leakage Current	223
5.4.2	Muon and Gamma Charge Production	224
5.4.3	First NTL Gain Demonstration: CDMS HVeV	227
5.4.4	Measuring IR Leakage	228
5.4.5	Impact Ionization and Trapping	230
5.5	Future Studies	232
6	First Sub-GeV Dark Matter Results: CDMS HVeV	234
6.1	Experimental Setup	234
6.1.1	Raw Exposure	236
6.1.2	Pulse Processing	239
6.1.3	Calibration	241
6.2	Quality Cuts	245
6.2.1	Livetime Calculation	245
6.2.2	Livetime Cuts	247
6.2.3	Trigger Time Cut	249
6.2.4	Chi-Square Cut	250
6.2.5	Good Noise Cut	252
6.2.6	Final Spectrum and Efficiency	252
6.3	Background Model	254
6.3.1	Impact Ionization and IR	254
6.3.2	Compton Recoils	255
6.3.3	Environmental Backgrounds	257
6.4	New Physics Constraints	258
6.4.1	Ionization Model	260
6.4.2	Dark Photons	262
6.4.3	Sub-GeV Dark Matter	264
6.5	Future Directions	264
7	Future Directions in Low-Threshold Detectors	266
7.1	Lowering Detector Thresholds	266
7.1.1	Alternative Detector Materials	267
7.1.2	Improving QET Resolution	269
7.1.3	Improving Readout Electronics	271
7.1.4	Phonon Imaging and Noise Correlation	273
7.2	Studying Low-Energy Backgrounds	275
7.2.1	High Energy Gammas Producing Low Energy Events	275
7.2.2	Sources of Dark Events	276

7.3 Exploring New Detector Targets	279
7.4 Experimental Outlook	281
A Analytic Detector Partition Functions	283
A.1 Z-Partition	283
A.2 R-Partition	286
B TES Linear Parameterization	288
C Phase Separated TES Dynamics	291
C.1 Setup of Differential Equations	291
C.2 Two-Block Symmetric Solution	294
C.3 Two Block Stable Asymmetric Solution	298
C.4 N-Block Solutions	303
D QET Efficiency	305
E Optimal Filtering	310
E.1 Basic Optimal Filter	310
E.2 Optimum Filter with Time Offset	311
E.3 Two-Pulse Optimum Filter	313
E.4 Joint Channel Optimum Filter	316
E.5 Joint Channel Correlated Optimum Filter	318
E.6 Optimal Filter Resolution	320
E.7 Wiener Filtering	320
E.8 Time-Domain OF: Matched Filtering	321
F Trace Analysis & Model Fitting	324
F.1 Fourier Methods	324
F.1.1 Continuous Transform	324
F.1.2 Discrete Transform	325
F.1.3 Power Spectral Density	326
F.1.4 Correlation, Cross-Spectral Density, and Coherence	327
F.1.5 Discrete Optimum Filter	328
F.2 Time-Domain Statistics	329
F.2.1 Covariance and Correlation	330
F.2.2 Example: Correlation Triggers	332
F.3 Least Squares Fitting	333
F.3.1 Reduced Chi-Square Statistic	334
F.3.2 Generalized Least Squares	335

F.3.3	Linear Least-Squares: Normal Equations	335
F.3.4	Nonlinear Models and Maximum Likelihood	338
G	Single Electron Probability Distributions	339
G.1	Abstract Form	339
G.2	Physical Trapping/Ionization Model	340
G.2.1	Single Charge	340
G.2.2	Two Charges	341
G.2.3	Three or more charges	342
G.2.4	Charge Energy Distributions	344
G.3	IR Loading	345
	Bibliography	347

List of Tables

2.1 Electron Band Structure Parameters	54
2.2 Intervalley Phonon Scattering Parameters	73
2.3 IV Scattering Rate Parameters	87
3.1 TES Parameters	98
3.2 TES Time Constants	115
4.1 QET Coverage Parameters	142
4.2 Detector Design Parameters	151
5.1 Detector Prototype Summary	175
5.2 Measured TES Parameters	185
5.3 Measured TES Bias Power Parameters	185
5.4 Resistivity Summary	188
6.1 Livetime Cuts Summary	249

List of Figures

1.1 Milky Way Models	6
1.2 Shapley and Hubble's Results	9
1.3 Rubin Rotation Curves	16
1.4 Figure 1 of [16], reproduced here for illustration of rotation curve modeling.	19
1.5 Figure 2 of [16], reproduced here for illustration of rotation curve modeling.	21
1.6 Bullet Cluster	24
1.7 Planck CMB Measurement	27
1.8 Local dark matter density	29
1.9 Relic Density	35
1.10 SUSY Models	36
1.11 WIMP Limits	37
1.12 Dark Matter Candidates	39
1.13 Light dark matter prospects	42
1.14 Dark Photon Relic Density	43
1.15 Energy Transfer v. Dark Matter Mass	45
1.16 Detector Bandgaps	46
1.17 SNOLAB Projected Reach	47
2.1 Full Band Structure	53
2.2 FCC Brillouin Zone	54
2.3 Simplified Band Models	55
2.4 Charge Scattering Rates	75
2.5 Intervalley Scattering Rates	79
2.6 Charge Experiment Setup	81
2.7 Electron Data	82
2.8 Electron Data	82
2.9 Example MFP Analysis Figure	83
2.10 Intervalley Scattering Results	85

2.11 Monte Carlo Results	86
2.12 IV Rate Results and Fit	87
2.13 Hole Temperature Trend	88
2.14 Hole Equal Energy Surfaces	89
2.15 Hole Repulsion	91
3.1 NTL Phonon Summary	94
3.2 Schematic of QET	96
3.3 Superconducting Heat Capacity	103
3.4 TES Bias Circuit	105
3.5 TES Time Constants	115
3.6 Phase Separation Points	119
3.7 QET Unit Cell Number	120
3.8 Example Noise Models	125
3.9 Phonon Collection Fraction and Fin Length	129
3.10 QET 2D Efficiency	130
3.11 Phonon Collection Efficiency	132
3.12 Optimized Overlap Region	133
3.13 1D Collection Measurements	134
4.1 Z Partition and Dimensions	138
4.2 Detector Aspect Ratio	139
4.3 TES Normal State Resistance Optimization	141
4.4 Optimized QET	143
4.5 Phonon and Energy Efficiency	145
4.6 Phonon Collection Time	147
4.7 QET Resolution Optimization	149
4.8 First-Pass Collection Efficiency	150
4.9 Expected HV Detector Performance	152
4.10 Bath and SQUID Effects	153
4.11 2D Field Model Geometry	155
4.12 Fringe Field v Aluminum Coverage	157
4.13 Electric Field Edge Effects	158
4.14 3D Field Model Cross-Sections	159
4.15 DMC Input Energy	161
4.16 DMC Yield Plots	162
4.17 TES Current Noise Generation	163
4.18 Simulated Phonon Pulses	164

4.19 DMC Energy Resolution	165
4.20 Proposed Geometries	166
4.21 Initial Phonon Distributions	167
4.22 Z Partition versus Voltage	168
4.23 Radial Partition PCA	170
4.24 BDT Fiducialization Summary	171
4.25 CDMS HV Mask Summary	173
5.1 G115 Ibs Measurements	178
5.2 G124 Bias Power Time Dependence	180
5.3 Complex Impedance Traces	183
5.4 Complex Impedance Fits	186
5.5 Complex Impedance Parameters	187
5.6 UMN Noise Fitting	190
5.7 TES Oscillation and Instability Regions	191
5.8 G115 Mask Design	192
5.9 G115 Noise Comparison UCB v SLAC	193
5.10 SLAC Environmental EMI	194
5.11 SQUID Responsivity and Noise	197
5.12 DCRC D.0 Readout Circuit Diagram	198
5.13 DCRC Noise Improvements	200
5.14 Improved G115 Noise at SLAC	201
5.15 G124 Normal and Superconducting Noise	203
5.16 G124 Transition Noise	204
5.17 Simple Vibrational Noise Model	206
5.18 G124 Noise Coherence	207
5.19 Germanium and Silicon Attenuation Length	210
5.20 Measurements of Si and Ge Prototypes at UMN	213
5.21 G115 Efficiency and Resolution	215
5.22 G124 Pulses and Resolution	217
5.23 G124 Efficiency	218
5.24 HVeV Resolution and IR Filtering	221
5.25 HVeV Absolute Calibration	222
5.26 First CDMS HV Leakage Measurements	225
5.27 Muon Rates and Current	226
5.28 NTL Gain Linearity	227
5.29 Measured Leakage Spectrum	229
5.30 Modeled IR Spectrum	230

5.31 Impact Ionization and Trapping	231
6.1 HVeV Detector Setup	235
6.2 Temperature Time Variation	237
6.3 Time Dependent Trigger Rates	238
6.4 Example Pulse Templates and PSDs	240
6.5 Filtered Pulse Template	241
6.6 Channel Relative Calibration	242
6.7 Temperature Calibration	243
6.8 Energy Scale Calibration	244
6.9 Time-Dependent Calibration	245
6.10 Total Exposure versus Time	246
6.11 Livetime Series Cuts	247
6.12 Time Shifting Noise Blob	249
6.13 Pulse Time Cut	250
6.14 Chi-Square Cut	251
6.15 Spectrum and Efficiency	253
6.16 HVeV Background Model	254
6.17 Leakage Time Dependence	255
6.18 Am-241 Compton Events	256
6.19 Bad Events in the Chi-Square Plane	258
6.20 Background Chi-Square Distributions	259
6.21 Optimal Interval Regions	260
6.22 Fano Factor Signal Variation	262
6.23 HVeV Results	263
7.1 Updated CDMS HV Optimization Space	270
7.2 IR Leakage Rates	276
7.3 Autoionization Rate	278
7.4 Impurity Energies in Si and Ge	279
7.5 New Detector Materials	280
G.1 Higher Order Approximation	343

Chapter 1

The Need for New Physics

“...It seems probable that most of the grand underlying principles have been firmly established and that further advances are to be sought chiefly in the rigorous application of these principles to all the phenomena which come under our notice...An eminent physicist remarked that the future truths of physical science are to be looked for in the sixth place of decimals.”

- Albert Michelson, 1894, dedication of Ryerson Physical Laboratory

The story of modern physics written in textbooks begins with Newton, including prefaces on ancient astronomy, natural philosophy and metaphysics, and continues in an unbroken line to the present day. It is a story of discovery, of trial and error, but ultimately of steady progress towards the modern understanding of how the physical world behaves and the laws it obeys. Those who practice physics as a profession know that it is never quite so clear cut, and every era has its many schools of thought that through time are filtered and combined to produce standard canon. The key insight one can derive from a careful reading of the history of physics, and science in general, is how to seek out phenomena that deviate sharply from expectation as a framing device for looking for new models.

The search for dark matter is just the latest exercise in this practice, which highlights a problem in our understanding of the universe, and proposes a certain class of solutions for it, namely the existence of a new type of matter. My hope is that, even if as you read this chapter the theories discussed have been ruled out, it will still provide a roadmap for discovery that transcends the particular problems it describes. The search for dark matter in particular is driven by a very fundamental mismatch between the nature of gravity and the distribution of matter, and the resolution of this problem will necessarily result in new physics, whether or not new mass is a part of it. At this

moment in time, the dark matter field is an exciting one, with many synergies in technology development and low-background photon detection, and serves as excellent motivation to work on many cross-disciplinary problems for which dark matter is the first stringent and interesting sniff-test.

In this chapter, I will present first a brief chronology of astrophysical observation centered on the study of celestial movements, the discovery of extragalactic objects, and measurements of stellar kinematics in our galaxy as well as those nearby. I will then discuss the key observations that indicated the presence of missing mass, and present the argument for why a dark matter explanation to this problem is particularly attractive. This will be followed by an exploration of the varied dark matter phenomenology, starting with the standard (and now somewhat disfavored) WIMP paradigm and moving into lighter and ultra-light dark matter theories. I will finish by motivating new detectors that can operate at ultra-low energy thresholds in order to study some of these new light dark matter models as a way of contextualizing the work in this thesis.

1.1 The Cosmos in Motion

The earliest surviving ancient work that laid out a systematic explanation for the motions of the planets was Ptolemy's "Almagest"¹, a second century Greek treatise describing the motions of stars in a mathematical framework. This represented the first attempt to apply mathematical reasoning to the observed movement of various heavenly bodies, most importantly the sun, moon, and the known planets². The vision of cosmic structure in this treatise followed the geocentric³ model first proposed by Aristotle, which would remain the prevailing model for 2 millennia after his death, but introduced a mathematical model for explaining the elliptic orbits of the known planets. In this work, he also presented the earliest surviving star catalog, containing the positions of the 1025 known stars visible from Europe.

Ptolemy's work illustrates the limits of what can be studied without telescopes or precision measuring instruments, given the slow progress made over the subsequent 1500 years before the invention of the telescope. Progress in astronomy was slow, requiring dedicated work over a lifetime to make small advances in the understanding of what, to almost all of the astronomers, were points on a fixed sphere that surrounded the earth. Between 100 and 1600 AD, continued observations pushed astronomers away from the simple, fixed-point model of the universe^[19, 8, 40, 28]:

- 6th Century AD: Olympiodorus the Younger argues that the Milky Way, which Aristotle held was an atmospheric effect, must be stellar in nature due to its lack of parallax⁴

¹The English name for this work tells us a lot about its historical importance. 'Almagest' is the westernized version of the Arabic for 'The Greatest', itself translated from the shortened Greek title 'The Great Treatise'. The original name for the work was 'Syntaxis Mathematica', roughly meaning simply a Mathematical Treatise.

²The word planet comes from the ancient Greek for 'wandering stars', or simply 'wanderers' (plantai)

³earth centered

⁴Parallax - change in apparent position on the sky of an object due to the movement of the observer and the finite distance to the object. The effect of parallax can be demonstrated by holding a finger up at arm's length and moving

- 964 AD: Al-Sufi is the first to observe the Large Magellanic Cloud (LMC) and the Andromeda galaxy, the Milky Way's largest satellite and the next nearest galaxy respectfully. He describes them not as new celestial objects, but as 'cloud-like' in appearance. Neither object is visible from Europe, and thus the LMC is not observed by Europeans until Magellan's voyage.
- 1000 AD: Al-Buruni and Ibn Sina (Avicenna) debate whether there is any evidence to support the Aristotelian model of the universe, most importantly the assumption that all natural motion is circular, the basis for modeling the dynamics of the planets.
- 1010's AD: Ibn Al-Haytham (Alhazen) writes the 'Book of Optics', the basis for later European development of the telescope, which replaced Ptolemy's optics as the premier treatise on the subject. Al-Haytham also publicly critiques Ptolemy's mathematical models, objecting to the planets spontaneously changing velocity in order to maintain circular orbits. Al-Haytham also measures the lack of parallax of the Milky Way for the first time, proving Olympiodorus' earlier conjecture in a more precise manner.
- 1272 AD: Al-Tusi explains why the galaxy appears to be continuous despite being made of many stars, writing: "The Milky Way is made up of a very large number of small, tightly-clustered stars, which, on account of their concentration and smallness, seem to be cloudy patches. Because of this, it was likened to milk in color." He also invents the Tusi-couple, which demonstrates how circular motion can give rise to linear oscillation, and is described in Copernicus' work centuries later.
- 1284: Witelo writes the treatise 'Perspectives', partially on optics, building on the work of Al-Haytham and further laying the groundwork for the development of telescopes.
- 1420's AD: Ulegh Beg constructs the observatory at Samarkand (modern day Uzbekistan), similar to the one produced a century later by Tycho Brahe, and produces the largest star catalog since Ptolemy (994 stars) correcting earlier errors and combining knowledge of the northern and middle-latitude stars. He increases his accuracy by making a much larger sextant than those previously used by Arabic astronomers.
- 1543: Copernicus posthumously publishes "On the Revolution of the Celestial Sphere", solving the mathematical problems pointed out by earlier astronomers and making a much simpler cosmological model.
- 1577: Tycho Brahe shows that comets are in fact moving about the sun, and not atmospheric events, by measuring the (non-existent) parallax of a comet. This is considered the first observational challenge to the geocentric model of the solar system.

your head side to side; you observe your finger to move relative to distance objects, telling you it must be closer to you than those objects.

- 1609: Kepler uses Brahe’s observations to construct his mathematically complete model of the heliocentric solar system, measuring distances between the sun and the planets and setting a scale for the solar system. Kepler’s laws lay the groundwork for Newton’s law of gravitation.

In 1608, Galileo Galilei made the first astronomical observations with a rudimentary telescope, proving for the first time that the Milky way was in fact a collection of stars and discovering four of Jupiter’s moons. This began a new era of observations, in which the number of mapped, observable objects in the night sky increased exponentially with each new technological advance. It’s also at this point in history that Kepler’s observations demonstrate the relationship between orbital period and orbital radius from the sun (commonly referred to as Kepler’s 3rd law):

$$\frac{T^2}{a^3} = \frac{1\text{year}^2}{1\text{AU}^3} \quad (1.1)$$

This simple proportionality only works in the heliocentric system, and once demonstrated in this new paradigm, allowed Newton to use this relationship to construct his law of universal gravitation. One can get from Kepler’s third law to a law of universal gravitation by employing Newton’s new laws of motion, published in 1687, obtaining the more general form of Kepler’s third law:

$$\frac{T^2}{a^3} = \frac{4\pi^2}{G(M_1 + M_2)} \quad (1.2)$$

This new equation can be used to measure the mass of any planet using the measured orbital period of its moons, and was the main tool used in the following centuries to relate stars and planets to one another. In the span of less than a century, we see a scientific advance more important than any in the two millennia before it, which can be attributed to both the new technology of the telescope, but also the legacy of dedicated measurement of stars visible to the naked eye. These two developments are most often the reason that we differentiate between Aristotelian (ancient) and Newtonian (classical) physics. What followed was a period of academic exploration enabled by this new understanding of nature and the technology that allowed humans to study the cosmos in finer detail than every before.

Looking Beyond the Solar System

Following the example of the ancients, astronomers during the enlightenment turned their telescopes to the sky in an attempt to find moving objects, hoping to determine the distance to a star and thus set a scale for the universe. While comet hunting, Charles Messier constructed a catalog of 103 (later expanded to 110) bright nebulae, publishing his catalog in 1781 [28]; many of these were in fact nebulae, some stellar associations, and some actually distant galaxies. William Herschel, a musician and astronomer who built more than 400 telescopes during his career, discovered Uranus while conducting surveys of stellar motion. In these surveys, he was also the first to discover the

abundance of stellar binary systems, and was the first to observe that the motion of these systems implied that the stars were orbiting each other. Herschel published his star and nebulae surveys over the span of time starting in 1773 and concluding with catalogs of 800 multiple star systems and 2500 'deep sky objects', mainly Nebulae and galaxies, which greatly expanded on Messier's work. His son would continue his work after his death, and their combined efforts produced the New General Catalog (NGC) containing 8000 of the brightest non-stellar objects in the night sky[28].

Herschel was also the first person to attempt a three-dimensional map of the galaxy. In the mid 18th century, astronomer Thomas Wright and philosopher Immanuel Kant posited that, as the solar system is a rotationally supported disk, it makes sense that the stars we see may also be a rotationally supported disk, which we would thus see as a band across the sky. Herschel took this idea, and by counting the number of stars in each portion of the band visible from his observatory, constructed the first star map of our galaxy in 1785[103], shown in figure 1.1. The limited depth of his telescope and his simplifying assumptions resulted in a very poor map with the sun at the center of the galaxy, but represented the first attempt to map star positions as more than fixed points on a sphere. Herschel's main limitation was his inability to establish distances to stars, and he had to assume they were equally spaced, thus precluding him from finding any evidence of spiral structure or measuring relative stellar densities.

The 1800s saw improvements in technology that led to the first discoveries using gravitational anomalies to find new objects. In 1844, Friedrich Bessel noted that the stars Sirius and Procyon both exhibited odd proper motions, and inferred the existence of as-yet unobserved companion stars, which would later be confirmed by more powerful telescopes[19]. In 1846, Le Verrier and Adams noted the anomalous orbit of Uranus (just discovered by Herschel 60 years earlier) and computed the expected position of a new planet needed to create that orbit. Neptune was thus found the same evening the observatory received their communication suggesting its probable location[19]. Le Verrier also proposed the new planet 'Vulcan' to explain Mercury's anomalous orbit, which was never observed but would later be one of the key predictions that would ensure the success of Einstein's theory of general relativity. Separately, the study of objects beyond the solar system continued, and in 1845 William Parsons observed, through the largest telescope at that time, that some of Herschel's objects had a spiral structure, and seemed to be rotating. Without any means to measure this rotation, however, this was at the time mere conjecture[28].

By the end of the 19th century, photographic plates came into common use, allowing for more detailed and wider ranging study of the night sky. Dark regions were noticed throughout the plane of the milky way (as can be seen in Figure 1.1), and debates began about whether these were empty voids or dark regions obscuring stars behind them. In order to estimate the true number of stars in the Milky way, Lord Kelvin made the first estimate of the mass enclosed in the observed universe (at that time the local stellar neighborhood) by treating it as a gas, and using the measurements of stellar velocities available at the time (by watching the slow movement of stars across the sky) to

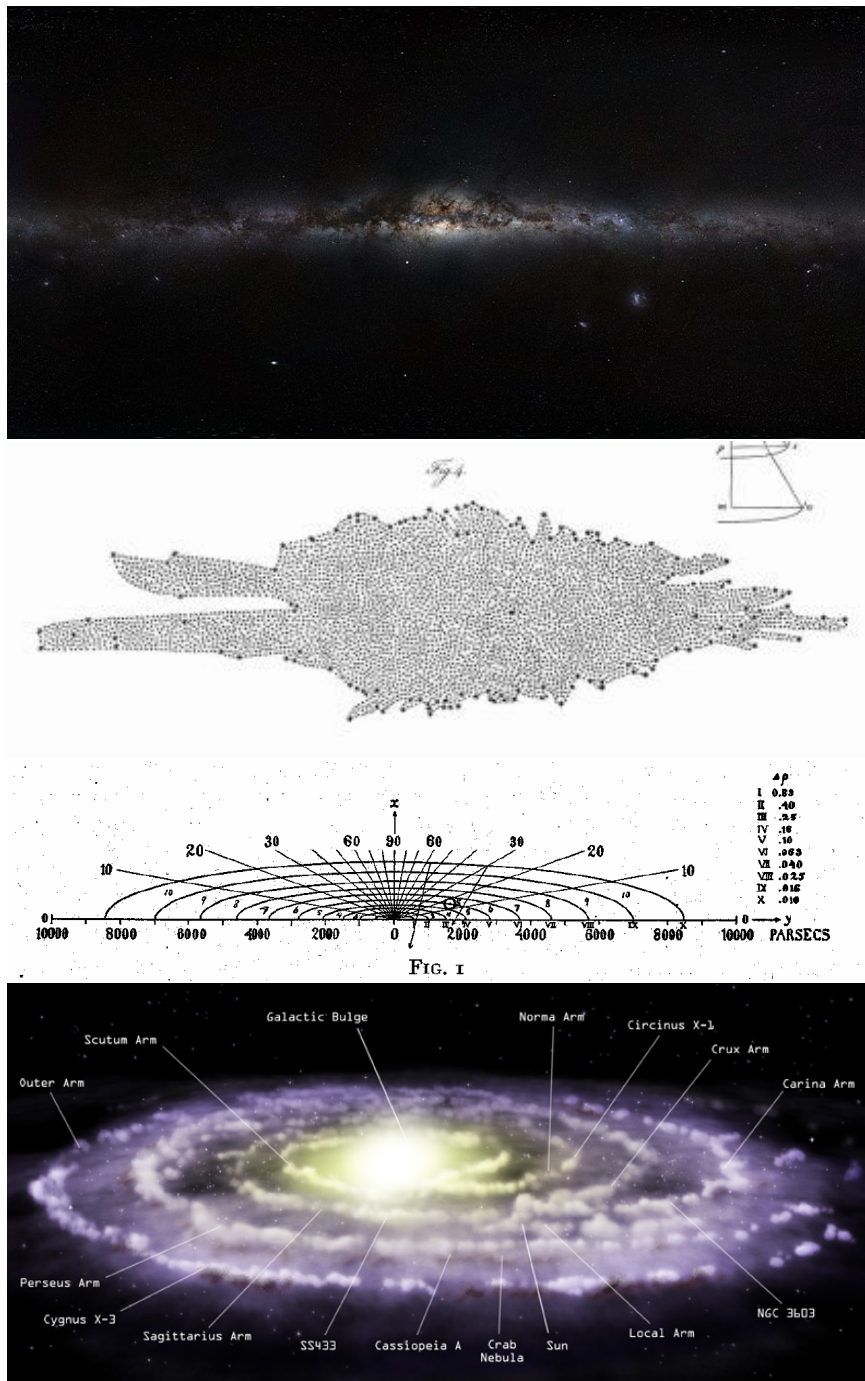


Figure 1.1: Top: Image of the milky way and surrounding space compiled by the ESO^[7]. This is how we observe the Milky way from earth, although with much higher resolution than can be seen by eye or on Earth's surface. Upper Middle: Herschel's reconstruction of the Milky Way based on assumptions of star brightness and placement. Lower Middle: Kapteyn's model of the universe, where the Sun is no longer at the center but the scale and shape are still not quite right. He does not recognize that objects can be outside of this relatively small distribution. Bottom: Modern diagram of the Milky Way^[8].

estimate the upper limit of mass in a gravitationally bound system given these measurements and the accepted size of the galaxy [19]. He estimated that there may be “as many as 10^9 stars [within a sphere or radius $3.09 \cdot 10^{16}$ km]” based on these proper motions. Henri Poincaré (in 1906) and Ernst Öpik (in 1915) compared this estimate to the number of observed stars, and concluded that since the expected number of stars and observed number are comparable, there is no appreciable “dark matter” in the galaxy. The term dark matter at the time generally referred to unseen gas or dead stars, but was very much a part of the scientific lexicon by the 1920’s.

Meanwhile, Jacobus Kapteyn was busy updating the work of Herschel, and in 1922 published a much more rigorously computed model of the galaxy that puts the Sun off-center (though much closer to the center than it should be), and estimates the size and shape of the galaxy, at that time believed to be an ellipsoid. Kapteyn’s model as originally shown can be seen in Figure 1.1. As part of this work, Kapteyn (and his student Jan Oort) rigorously calculated an upper limit for the local mass density of dark matter in the solar neighborhood, finding the lowest limits then predicted of around 1% of the mass of the sun per cubic parsec⁹. There is roughly one star per 7 cubic parsec, meaning that the dark matter density is at most less than 10% of the mass in the stellar neighborhood; given that stars are not the only known mass in the galaxy, this is therefore not a surprising number. [19]

A Larger Picture

This well-behaved system would have worked had it not been for major observations made by Harlow Shapley, Vesto Slipher, and Edwin Hubble. In 1915, Vesto Slipher used spectra produced on photographic plates to demonstrate the doppler shifting of these spectra across the spiral nebulae first observed by Parsons, demonstrating observationally that the nebulae did appear to rotate. Meanwhile, between 1915 and 1919, Shapley estimated the distances to 93 globular clusters¹⁰ by observing variable stars, which have a known absolute luminosity for a given variation period and thus allow a straightforward distance measure. Shapley realized these clusters were not located isotropically, but instead were centered on a point towards a spot in the galactic plane roughly 15 kpc from the Sun, and that the furthest globular cluster should be 70 kpc from the Sun [28].

Shapley’s measurements were obviously in tension with Kapteyn’s model, putting the center 5 times farther away and implying that the galaxy should be a factor of 10 larger than what Kapteyn observed. On the other hand, Shapley argued that all of the observed objects seen thus far must be contained in the Milky Way, while those who had been studying the spiral Nebulae (led by Heber Curtis) maintained that what they were observing were actually distinct galaxies, given the distances implied by stellar novae¹¹ seen in these objects compare to those observed locally. This

⁹1 parsec is equal to 2.36 light-years, or $3 \cdot 10^{13}$ km (30 trillion km).

¹⁰A globular cluster is a tight grouping of stars that formed at the same time and have a much higher stellar density than the surrounding space

¹¹A nova is a bright flare-up of a star caused by ignition of fusion on its surface.

was the subject of the 1920 “Great Debate” on the nature of the galaxy and universe, both the scale of each and whether Milky Way was the universe or whether the universe contained many ‘island universes’, to use Kant’s original term. This argument was obviously complicated by the fact that many of the measurements being used had large errors due to extinction by interstellar gas and dust and the limited precision of observing tools available at the time. A sketch of the discrepancy between models of the galaxy can be seen in Figure 1.2

Enter Edwin Hubble¹² In 1923, Hubble used the new 100-inch Mount Wilson telescope to observe variable stars in Andromeda, the next closest galaxy¹³, finding a distance of 285 kpc, definitively proving that the spiral nebulae were in fact new galaxies²⁸. This is actually smaller than the modern value of 770 kpc, but is still large enough to prove the point. In 1925, Slipher observed that out of a sample of 40 spiral nebulae, now thought to be distinct galaxies, almost all of the galaxies exhibited a distinct redshift in their spectrum, implying they were moving away from the Milky Way. In 1929, Hubble followed this up with more variable star measurements from a handful of nearby galaxies, establishing for the first time that the universe is expanding, and showing that the receding velocity of galaxies is proportional to their distance:

$$v = H_0 d \tag{1.3}$$

The figure from his seminal 1929 paper is reproduced in Figure 1.2. This relationship is now called Hubble’s law, and H_0 is Hubble’s constant, which he estimated to be about 100 km/s/Mpc based on his observations⁴⁹. The modern value, though still in flux, is close to 72 km/s/Mpc, although its measurement as of this writing is still hotly debated²¹.

Hubble’s 1929 paper marks the beginning of the field of observational cosmology, that is, the study of the universe, as it showed for the first time that our universe consisted of more than just the Milky Way. It is also where the study of dark matter begins to become a distinct pursuit, and will be described in the next section. Many important scientific advances were made as a result of Hubble’s initial work, most notably the study of this expansion that led to the even more surprising discover of dark energy, which causes this expansion rate to increase as a function of time. The technology that enabled us to probe the universe beyond our stellar neighborhood was driven by advances in imaging, telescopes, and signal processing, and we will come back to some of the lessons from this story at the end of the chapter when we discuss the way forward in sensitive astro-particle detectors.

¹²This story illustrates Hubble’s foundational role in both astrophysics and cosmology, and why the Hubble space telescope, which imaged further than humans had ever seen before, was named after him.

¹³Andromeda is commonly referred to as M31, from its designation in Messier’s catalog, and is one of the brightest non-stellar objects visible through any telescope. The Hubble Space Telescope image of Andromeda was also (until recently) the default background on all AppleTM computers.

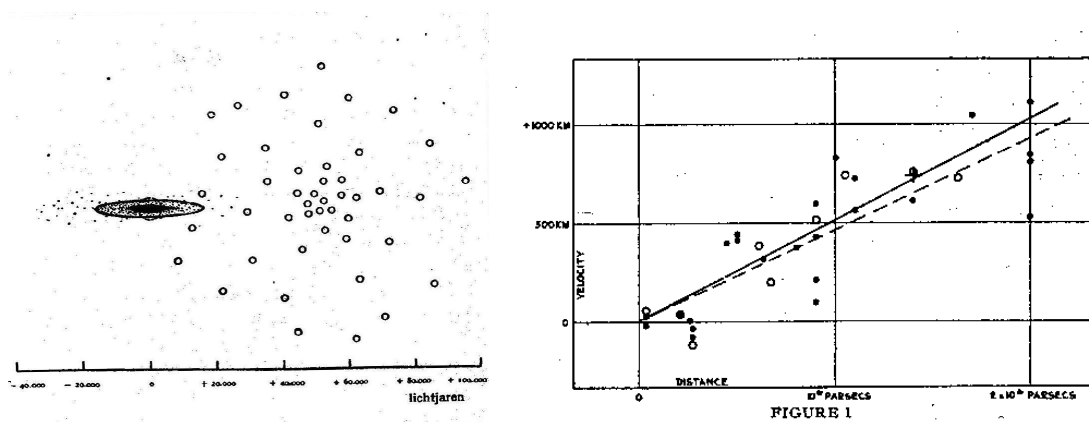


Figure 1.2: Left: Illustration of the Kapteyn v Shapley galaxy model problem, showing Kapteyn’s model of the Milky way as a collection of ellipses, and Shapley’s measured distributions of globular clusters [104]. Right: Figure 1 from Hubble’s 1929 Paper [49] showing the distance-velocity relation. You can see that some of the closest galaxies, including M31, are actually moving towards us. You can also compare the distance scales between objects in our galaxy in the left image and extragalactic nebulae, now known to be other galaxies, in the right image.

1.2 Missing Mass: The Case for Dark Matter

Over the past few decades, the historical case for dark matter has achieved a sort of orthodoxy, from which I don’t deviate much here except to insert derivations where I feel they add strength to the arguments, and editorialize a bit when it comes to the comparison between modified dynamics and particle dark matter. Much of the historical content in this section comes without explicit citation because it is now common knowledge in the field, however I refer the reader to some excellent review articles, in particular Refs [20, 18, 19], which take slightly differing takes on the history of the field but do an excellent job documenting their sources.

What differentiates the study of dark matter from its ‘sister’ field of dark energy, described at the end of the last section, is what makes this chapter so important, namely that the first conversation we have to have about dark matter is to justify that we’re not just fooling ourselves with shoddy observation. The New York Times “Book of Physics and Astronomy” has a time-line of physics that has 3 dark-energy related entries but none referring to dark matter, precisely because the existence of dark matter came to be known gradually as we eliminated possible systematics and narrowed down the explanation for an odd observational truth: there doesn’t seem to be enough matter anywhere we look outside of the solar system. This is an issue that crept up on us as we continued to log entries in our big book of mass discrepancies, and is marked, for the most part, not by groundbreaking observations, like those of Reiss, Schmidt & Perlmutter [80, 86] for dark energy, but by the lack thereof.

It is my goal in this section that by the end, the reader will appreciate that the simplest explanation for all of these phenomena is not a new theory of gravity, but a fundamental misunderstanding of the nature of matter. That we tend to prefer one theory of what that matter is will be left to later sections, but at this moment in time, there is no doubt that simple fixes to our theory of gravity can not properly explain all of the phenomena discussed here.

1.2.1 Rotation Curves and the Virial Theorem

The initial indication that galaxies were more complicated than expected came through studies of the rotational velocities of stars in a galaxy measured as a function of distance from the galactic center. A galaxy is a complex system of millions to billions of stars, but to a very good approximation we can employ the virial theorem to model the relationship between kinetic and gravitational potential energy. The virial theorem states that the time-averaged kinetic energy of a bound system obeys the relation

$$\langle KE \rangle = -\frac{1}{2} \sum \langle F \cdot r \rangle \quad (1.4)$$

and for a potential that only depends on the distance r between particles, giving the form $V(r) \propto r^n$, we find the much simpler (and more often seen) statement

$$2\langle KE \rangle = n\langle V \rangle. \quad (1.5)$$

For a gravitational potential with $n = -1$, we thus have

$$2\langle KE \rangle = -\langle V \rangle \quad (1.6)$$

where V is the average gravitational potential of a particle in the system, and KE is the average kinetic energy of the particle. This theorem allows us to use a general framework to relate mass and energy in a gravitationally bound system where calculation of net force on a given particle is more cumbersome than calculating its gravitational potential. In very regular systems, such as the solar system, we can show that the results of this calculation are the same as using the equations of centripetal acceleration, but an equivalent calculation for dispersion supported systems is not quite so simple, thus allowing us to more confidently use this equation to model dynamics of less ordered gravitationally bound systems.

Gravitational Potentials

There are two general cases we'll consider here, a sphere and disk each of constant mass density. In the case of the sphere, we know that the total mass within a given radius is given by the formula

$$M(R) = \frac{4}{3}\pi R^3 \rho \quad (1.7)$$

which gives us the differential equation

$$\frac{dM(R)}{dR} = 4\pi R^2 \rho \quad (1.8)$$

We also know that the gravitational potential energy from a shell of mass $M_{shell}(R) = \frac{dM(R)}{dR} dR$ is

$$V_{shell} = -\frac{GM_{shell}M(R)}{R} \quad (1.9)$$

so to bring an additional shell onto a sphere with radius R , the change in potential energy is

$$dV_{sphere}(R) = -\frac{16\pi^2 G \rho^2}{3} R^4 dR \quad (1.10)$$

giving a binding energy of

$$V_{sphere}(R) = -\frac{16\pi^2 G \rho^2}{15} R^5 = -\frac{3}{5} \frac{GM(R)^2}{R} \quad (1.11)$$

Note that this is not the simple case of $V = GM^2/R$, as the system is diffuse and gravitationally bound, and the particle in question also forms part of the gravitational mass. We can do the same calculation for a disk of uniform mass density:

$$M(R) = 4\pi R^2 \sigma \rightarrow \frac{dM(R)}{dR} = 8\pi R \sigma \quad (1.12)$$

which gives the differential equation

$$dV_{disc}(R) = -G32\pi^2 \sigma^2 R^2 dR \quad (1.13)$$

giving a binding potential of

$$V(R) = -\frac{32\pi^2 G \sigma^2}{3} R^3 = -\frac{2}{3} \frac{GM(R)^2}{R} \quad (1.14)$$

The key finding is that for a shape of constant density, the potential at some point inside the shape has the form

$$V(R) = -\alpha \frac{GM(R)^2}{R} \quad (1.15)$$

with the general condition that $\alpha \leq 1$. The shape and density profile of the system can only produce smaller potentials, given that the constant density case gives $\alpha < 1$ and the case of a shell, ring, or point all have $\alpha = 1$.

Velocity Profiles

In our case, where particles are stars with mass m_p orbiting an enclosed mass M_{tot} , we know that for a spiral galaxy with essentially circular motion in a 2D plane, we have

$$\langle V \rangle \approx -\alpha \frac{GM_{tot}m_p}{R} \quad (1.16)$$

and

$$\langle KE \rangle = \frac{1}{2}m_p v_p^2 \quad (1.17)$$

which gives the simplified form of the Virial theorem

$$v^2 = \alpha \frac{GM_{tot}}{R} \quad (1.18)$$

For a dispersion supported galaxy, like an elliptical galaxy, this needs to be modified slightly. In the case of highly elliptical motion, where there is a constant trade-off between kinetic and potential energy, we know that statistically, a star observed at radius R will have mean kinetic energy

$$KE = \frac{3}{2}m_p \sigma_v^2 \quad (1.19)$$

giving the Virial relation

$$\sigma_v^2 = \frac{\alpha}{3} \frac{GM_{tot}}{R} \quad (1.20)$$

In both cases, then, we find that measurements of a given velocity distribution as a function of distance from the orbital center gives a measure of the mass enclosed at that distance. By comparing these mass measurements at various distances, we can build up an understanding of the radial mass distribution, and compare to that we measure for various components through complementary methods.

To demonstrate the implications of this formula, let's apply it to the solar system to verify that it does indeed predict the correct motion for a very well studied system. We use the case of circular motion here, and given that the vast majority of the mass in the solar system is concentrated in the sun, we predict that the velocity of a planet orbiting at radius R should be (using $\alpha = 1$ for a compact system)

$$v = \sqrt{\frac{GM_{sun}}{R_{planet}}} \quad (1.21)$$

We don't typically measure planetary velocity, but more commonly talk about orbital semi-major

axis a and orbital period t . Using the relation $v = \frac{2\pi R}{T}$, we find that

$$\frac{4\pi^2 R_{planet}^2}{T_{planet}^2} = \frac{GM_{sun}}{R_{planet}} \quad (1.22)$$

$$R_{planet}^3 = \frac{GM_{sun}}{4\pi^2} T_{planet}^2 \quad (1.23)$$

This relation is just Kepler’s third law of planetary motion expressed in Newtonian gravity; taking the ratio of different periods and orbital radii gives the proportionality explicitly stated in this law.

Early applications of this argument in its Newtonian form resulted in the discovery of Neptune, due the anomalous motion of Saturn relative to the motion it should have exhibited had it been the last massive planet in the solar system [20]. Application of this same logic to the orbit of mercury predicted yet another planet, however the true explanation for the precession of mercury turned out to be explained by corrections from General Relativity, and thus deviations from the Newtonian prediction produced new discoveries in the theoretical, rather than experimental, domains. The utility of comparing observed celestial motion to these scaling relations has proven a useful tool for discovering gravitational anomalies, and it was natural to apply this technique not only to our own solar system, but also to galaxies as we began to observe them.

In the 1930’s, Fritz Zwicky used measurements from Edwin Hubble’s seminal paper on the first measurement of universal expansion to attempt to measure the mass of the Coma cluster, a cluster of galaxies close to our own [115]. In this and a later work that refined the calculation, Zwicky observed that the Virial mass and luminous mass (calculated using the number of galaxies and a mass-to-light ratio conversion) differed by a factor of 500, implying that the vast majority of mass in the cluster was some sort of “dark matter” [19]. In 1936, Sinclair Smith performed a similar measurement for the Virgo cluster and obtained a mass-to light ratio of 200, smaller than Zwicky’s but still way more massive than the expectation of a ratio of 1-3 [19]. These results were recognized as important at the time, but due to observational limitations as well as some obvious geopolitical distractions in the 40’s and 50’s, the observational investigations stagnated, with researchers assuming either the assumption of stability in the clusters was misfounded, or normal matter was missed in the observations.

The Mass to Light Ratio

A key weakness in these early arguments, which might explain their relatively low importance at the time, lies in the number of assumptions necessary to calculate the luminous mass of a galaxy. Suppose we observe a new galaxy, and measure its total luminosity in the optical regime. We can get a preliminary estimate of its mass by making a simple assumption that all of the mass comes from sun-like stars, and we know the mass and luminosity of a sun-like star. In units of solar luminosity and solar mass, then, the mass to light ratio is just 1. This sets a baseline from that we can say

naively that the mass to light ratios are too large in clusters, and is approximately what was assumed in these early papers.

We can now make a more sophisticated model. Our knowledge of modern astrophysics allows us to break down stellar mass by type, and use the HertzsprungRussell (HR) diagram for a typical galaxy to get a more refined estimate of stellar mass. For a main-sequence star, which are the vast majority of stars in a galaxy, we have the mass-luminosity relation [28]

$$L \propto M^{3.5} \rightarrow M_{star}^{3.5} = M_{sun}^{3.5} * \frac{L_{star}}{L_{sun}}. \quad (1.24)$$

This means that for a star of mass M_{star} and luminosity L_{star} , the mass to light ratio is

$$\frac{M_{star}}{L_{star}} = \frac{M_{sun}}{L_{sun}} \left(\frac{M_{sun}}{M_{star}} \right)^{2.5} = \frac{M_{sun}}{L_{sun}} \left(\frac{L_{sun}}{L_{star}} \right)^{5/7}. \quad (1.25)$$

We can make the simple assumption that all stellar mass is in main-sequence stars, and we can use the form of the typical galaxy luminosity function to compute the weighted mean mass to light ratio:

$$\left\langle \frac{M}{L} \right\rangle = \frac{\int \left(\frac{L_{sun}}{L_{star}} \right)^{5/7} \Phi(L) dL}{\int \Phi(L) dL} \quad (1.26)$$

where $\Phi(L)$ is the luminosity function, usually calculated as the number density of stars with a given luminosity (or mass) per cubic parsec. This function is normally a double exponential with a shallow slope below 1 solar mass, and a steeper slope out to high mass. We can thus bracket the range of M/L ratios by using either a flat distribution or a step function above 1 solar mass. Integrating with respect to L in either case gives

$$\left\langle \frac{M}{L} \right\rangle \approx \int \left(\frac{L_{sun}}{L_{star}} \right)^{5/7} dL = \frac{7L_{sun}}{2} \frac{\left[\left(\frac{L_{max}}{L_{sun}} \right)^{2/7} - \left(\frac{L_{min}}{L_{sun}} \right)^{2/7} \right]}{L_{max} - L_{min}} \quad (1.27)$$

Including larger mass stars, $\langle M/L \rangle$ decreases, while for majority light stars, $\langle M/L \rangle$ increases; it is also essentially around 1 regardless of the relative composition for this model of stellar abundance, and has a very weak dependence on the range of masses allowed in the galaxy. This model matches well with the observed mass to light ratio in the Milky Way's disk of around 3 [28], based on a stellar mass of $6 * 10^{10}$ solar masses, a gas mass of $0.5 * 10^{10}$ solar masses, and an observed B-band¹⁴ luminosity of $1.8 * 10^{10}$ solar luminosities.

Barring a significant period of star formation (as is found in high-redshift galaxies), most galaxies will have a very small range of mass to light ratios, and given modeling of their star-formation history, we can very accurately estimate their stellar mass. What's more, because the majority of all stellar

¹⁴B band refers to a blue filter centered around 430–450 nm depending on the particular instrument used.

mass in a galaxy is concentrated in low-mass stars, and they nominally have lifetimes much longer than the age of the universe, which means very little error in the measured mass can be attributed to high-mass stars; if anything we will over-estimate stellar mass, not underestimate it, if there are more high-mass stars than expected.

We can be confident that the above reasoning will hold for converting light to stellar mass, but a reasonable question to ask as an immediate follow-up is how well we can constrain the non-stellar mass of a galaxy, and how much higher the mass to light ratio can get before gas and dust cannot be invoked in order to bring measurements into a reasonable range. Studying the mass distribution in our own galaxy, as well as in nearby clusters, allows us to put upper limits on the fraction of the missing mass implied by these high mass to light ratios than can be explained by normal matter.

One of the first examples of such studies was the estimate of intra-cluster gas mass (that is, gas between galaxies in a cluster) due to measured X-ray emissions. Let's imagine that the majority of the unseen mass in galaxy clusters is gas, predominantly hydrogen. It is being bound by gravity, but it is being prevented from gravitational collapse by radiation pressure. The stable solution to this hydrodynamical system gives an expected gas temperature as a function of mass and radius of a spherically symmetrical gas cloud, roughly following the relation [20]

$$k_b T \approx 1.5 \text{ keV} \left(\frac{M(R)}{10^{14} M_{sun}} \right) \left(\frac{1 \text{ Mpc}}{R} \right) \quad (1.28)$$

The measured gas temperature for the Coma cluster is around 10 keV, implying a gas mass of $7 \cdot 10^{15}$ solar masses (assuming $R \approx 10$ Mpc). This is in extreme tension with Zwicky's measurement, and implies that intra-cluster gas cannot account for the missing mass inferred from Zwicky's virial mass estimate.

Galaxy Rotation Curves

A general confusion about discrepancies in mass to light ratios was the approximate state of the field of dark matter through the 1960s. During the 1970s, Vera Rubin and Kent Ford [94] began to catalog observations of the observed radial velocity of side-on spiral galaxies as a function of distance from their galactic centers, and showed that across galaxy size and morphology, all galaxies showed a flat rotation curve out to high radii (see Figure 1.3). We can use the earlier virial relations to make sense of the implications of their observations.

We expect, from equation 1.18, the orbital velocity to go as $1/r$ for stars at the edge of the galaxy, where the enclosed mass has reach an asymptotic value. The observations by Rubin and Ford were done using molecular hydrogen, with velocities obtained by the Doppler shift in the $H\alpha$ emission lines [94] and references therein]. This allows observations to extend as far from the galactic center as hydrogen can be found, to regions where the hydrogen is diffuse enough to inhibit star formation.




Figure 1.3: Composite plot of the rotation curves measured from Rubin et. al. [94], showing that all curves quickly reach their asymptotic flat value.

A flat rotation curve thus implies that the mass within the observed radius R follows the scaling law

$$M(R) = \frac{v_{flat}^2}{\alpha G} R \approx \frac{3v_{flat}^2}{2G} R \quad (1.29)$$

where $\alpha \approx 2/3$ assuming axial symmetry for spiral galaxies. This implies that, for example, if the mass is spherically distributed, then the density goes as

$$M(R) = \int 4\pi R^2 \rho(R) dR \approx \frac{3v_{flat}^2}{2G} R \quad (1.30)$$

$$\rho(R) \approx \frac{3v_{flat}^2}{8\pi G R^2} \quad (1.31)$$

In other words, the majority of the mass in the galaxy seems to be a distribution with more total mass further from the galactic center, with a mass density that falls as $1/R^2$ far beyond the last observable stars and vastly outweighs the observable gas beyond the end of the stellar disk. This was the first evidence that the dark matter problem was more than just a problem of accounting for missing mass; in order to explain these observations, we have to rethink our models of galactic structure, where the majority of mass seems to be spatially separated from the luminous matter.

This measurement of mass density within a given radius at multiple distances from a galaxy's center was an important part of establishing that the gravitational and stellar masses had to be different, and that the discrepancy couldn't just be due to intergalactic extinction or Baryonic matter. To see why, let's consider a model where extinction is the primary cause of this discrepancy. In this model, we consider each line of sight to have roughly equivalent amounts of extinction. If we have a density that follows the given scaling law, we find that we have a line of sight mass density

$\Sigma(r)$ of

$$\Sigma(r) = \int_{-\infty}^{\infty} \rho(\sqrt{r^2 + z^2}) dz = \frac{3v_{flat}^2}{8\pi Gr} \quad (1.32)$$

where r is the distance from the center of the galaxy perpendicular to the line of sight. This is to say that even if there is some extinction, we should see the surface brightness follow roughly the same trend as the rotation curve, peaking at the center but falling off much slower than $1/r^2$ and extending far beyond the outer edges of the visible galaxy; we see instead that the star light has a much smaller length scale than the rotation curve.

The obvious next step in this line of reasoning is to suggest that maybe the main mass density is not stellar in origin or is primarily composed of stars that have a much lower luminosity than those on the main sequence. The latter is easier discussed conceptually; the likelihood of many low-mass stars being undetected in this abundance is essentially 0 given their long main-sequence lifetimes. This does not preclude more exotic astrophysical objects (e.g. black holes or neutron stars), which we'll discuss more in the section on gravitational lensing. Suffice it to say that due to these optical measurements, we can limit the candidates to non-stellar origins.

The solution to the problem of differing mass and rotation curve measurements came before Rubin & Ford's work with the detection of the 21-cm hyperfine hydrogen transition by Purcell & Ewan [36, 19]. This transition is the result of a slight increase in stored energy when the proton and electron in the hydrogen atom are misaligned, such that for very diffuse cold hydrogen with randomly populated spins, half of the atoms are initially capable of emitting a photon with wavelength $\lambda = 21.106$ cm. We can calculate both the energy and the probability of this transition to very high precision from first principles; this calculation can be shown to exactly match the observed energy [36]. The probability of an isolated atom undergoing this transition is found by application of Fermi's golden rule and results in a transition probability of [37]

$$\tau_{HI}^{-1} = \frac{64\pi^4 \mu_B^2}{3h\lambda^3} \approx 2.85 \cdot 10^{-15} \text{ Hz} \quad (1.33)$$

where μ_B is the Bohr Magnetron:

$$\mu_B = \frac{e\hbar}{2m_e} \quad (1.34)$$

and λ is the wavelength defined above. This is obviously a very small probability, but for a hydrogen gas mass of millions of solar masses, the emission of this line is a fairly common occurrence, and it has been used for over 60 years to simultaneously map gas mass and its rotational velocity both in the Milky way and in remote galaxies.

Given a known line energy, we can compute the mass of hydrogen emitting that line as

$$M_{HI} \approx \frac{P_{HI}\tau_{HI}}{E_{21cm}} m_H \quad (1.35)$$

and the maximum velocity of that cloud of hydrogen relative to the observer as

$$\frac{\lambda_{max}}{\lambda_{emit}} = \gamma \left(1 + \frac{v_{\parallel}}{c}\right) \rightarrow v_{\parallel} = c \left[1 - \gamma^{-1} \frac{\lambda_{max}}{\lambda_{emit}}\right]. \quad (1.36)$$

We can thus simultaneously determine, for an edge-on galaxy, rotational velocity and mass density as a function of radius as far out as the neutral hydrogen is detectable. The only systematic on this measurement is the relative probability of the 21cm light reaching the observer, explored in detail by Field [\[37\]](#). In studying the systematics of this measurement, it becomes abundantly clear that any local conditions can only increase the emission strength for a given mass of atomic hydrogen. Without being exhaustive:

- Temperature increases emission rate by increasing the population of higher energy states and the likelihood that adjacent atoms stimulate emission
- Magnetic field causes line splitting, but on average leaves the rate unchanged as it increases the rate for half of the orientations and decreases it for the other half.
- Incident radiation will either not affect the hydrogen atoms or further stimulate emission.
- The interstellar medium (ISM) is largely transparent to this line, and emission and absorption will always happen at equal rates in hydrogen, thus intergalactic extinction is fairly negligible for nearby galaxies.

Playing devil's advocate, the only way to increase mass estimates of gas given known 21-cm line emissions is to adjust the relative abundance of elements and molecular hydrogen in the ISM. These abundances are highly constrained by stellar physics and early universe cosmology, thus getting more than a factor of 2-3 for a given galaxy is highly disfavored, and the same factor systematically across all observed galaxies would result in a much different star formation history than what we observe.

To summarize the argument in this section, I refer you to Figure 1 of [\[16\]](#), reproduced here as Figure [1.4](#). This figure shows the rotation curves inferred from the measured luminosity of the galaxy, as well as those for the measured gas content, compared to the rotation curve measured from the 21cm line Doppler shifts. In each case, one can see that for the gas to account for the entirety of the rotation curve, the abundance would need to be scaled up by between 3 and 10 depending on the galaxy, and that the inferred rotation curve from dark matter is dominant across the board and has a fairly consistent shape across the galaxies, with the exception possible of NGC 7331 (this could just be a difference in scale factor). If we accept the gas measurements, then we're left with two dark matter hypotheses: exotic stellar material, or new matter entirely. First, however, we need to discuss the third distinct possibility that gained some traction from the 1970s to the mid-late 90's: modified gravity and its first iteration, Modified Newtonian Dynamics (MOND).

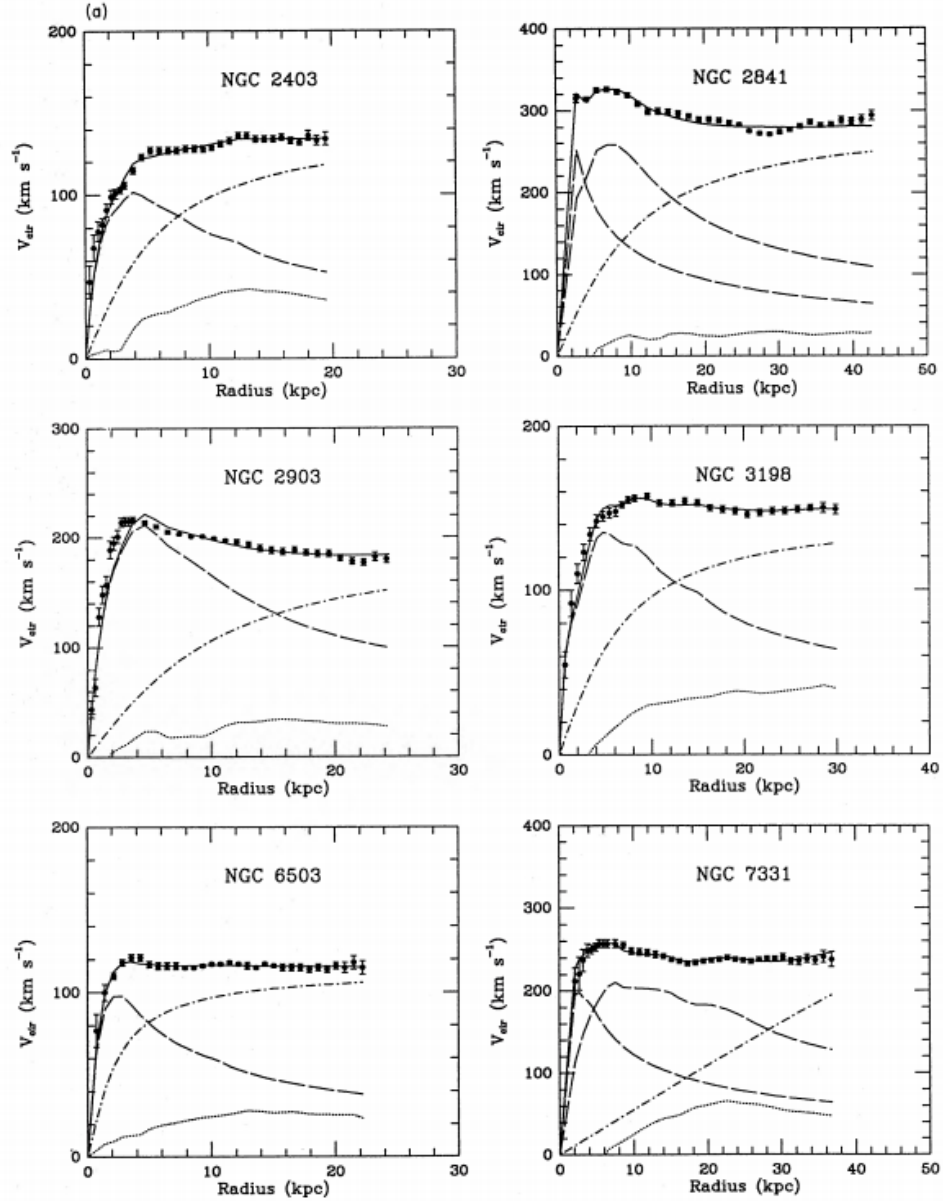


Figure 1. Three-parameter dark-halo fits (solid curves) to the rotation curves of sample galaxies. The rotation curves of the individual components are also shown: the dashed curves are for the visible components, the dotted curves for the gas, and the dash-dot curves for the dark halo. The fitting parameters are the mass-to-light ratio of the disc (M/L), the halo core radius (r_c), and the halo asymptotic circular velocity (V_h). The galaxies from the sample of Begeman are shown in (a) and the lower luminosity galaxies in (b). Best-fit values for the free parameters are given in columns 2, 3 and 4 of Table 2.

Figure 1.4: Figure 1 of [16], reproduced here for illustration of rotation curve modeling.

1.2.2 Modified Gravity

We can see from Figure 1.4 that an immediate solution to the rotation curve problem is one we've encountered in the past: that our theory of gravity is wrong. It's fairly apparent that if the gravitational force fell off as $1/r$ instead of $1/r^2$, then we could easily accommodate flat rotation curves without invoking 'spooky' dark matter. Assuming we're further away from the galaxy than any of its mass constituents, then we can modify equation 1.20 by multiplying by R giving the equation

$$v^2 = \alpha GM_{tot} \quad (1.37)$$

which is clearly consistent with a flat rotation curve. In a phenomenological modified gravity scenario, one only needs to fit the mass to light ratio and scale length over which the gravitational force goes from $1/R^2$ to $1/R$ to fit the rotation curves given the same measurements. This is exactly what is done in 16 for the exact same rotation curves from the previous section, and the results are reproduced in Figure 1.5.

This seems like a much simpler explanation than invoking an entirely new type of matter that has to weigh at least 10x (sometimes 100x) the known mass of a galaxy. As a result, the authors conclude 16:

The overall conclusion is that MOND is currently the best phenomenological description of the systematics of the [rotation curve] discrepancy in galaxies.

In the interest of completeness, I want to summarize the basic argument made by proponents of modified gravity before discussing some of the more recent theoretical and experimental advances that have led to it falling out of favor.

Theories of modified gravity in their original form were referred to as MOdified Newtonian Dynamics (MOND), first put forward by Mordehai Milgrom in 1982 19 as an intellectual exercise following the logic in the previous paragraph. The argument went that in the limit of very weak acceleration, if $F = m \frac{a^2}{a_0}$, then the natural prediction in this limit is that of a flat rotation curve 72:

$$F = \frac{ma^2}{a_0} = \frac{GMm}{r^2} \quad (1.38)$$

$$\frac{(v^2/r)^2}{a_0} = \frac{v^4}{a_0 r^2} = \frac{GM}{r^2} \quad (1.39)$$

$$v^4 = a_0 GM \quad (1.40)$$

In Milgrom's original formulation, Newton's second law was written as

$$F = m\mu\left(\frac{a}{a_0}\right)a \quad (1.41)$$

where $\mu(x)$ was said to be some function, coming from a complete theory, with the limiting cases

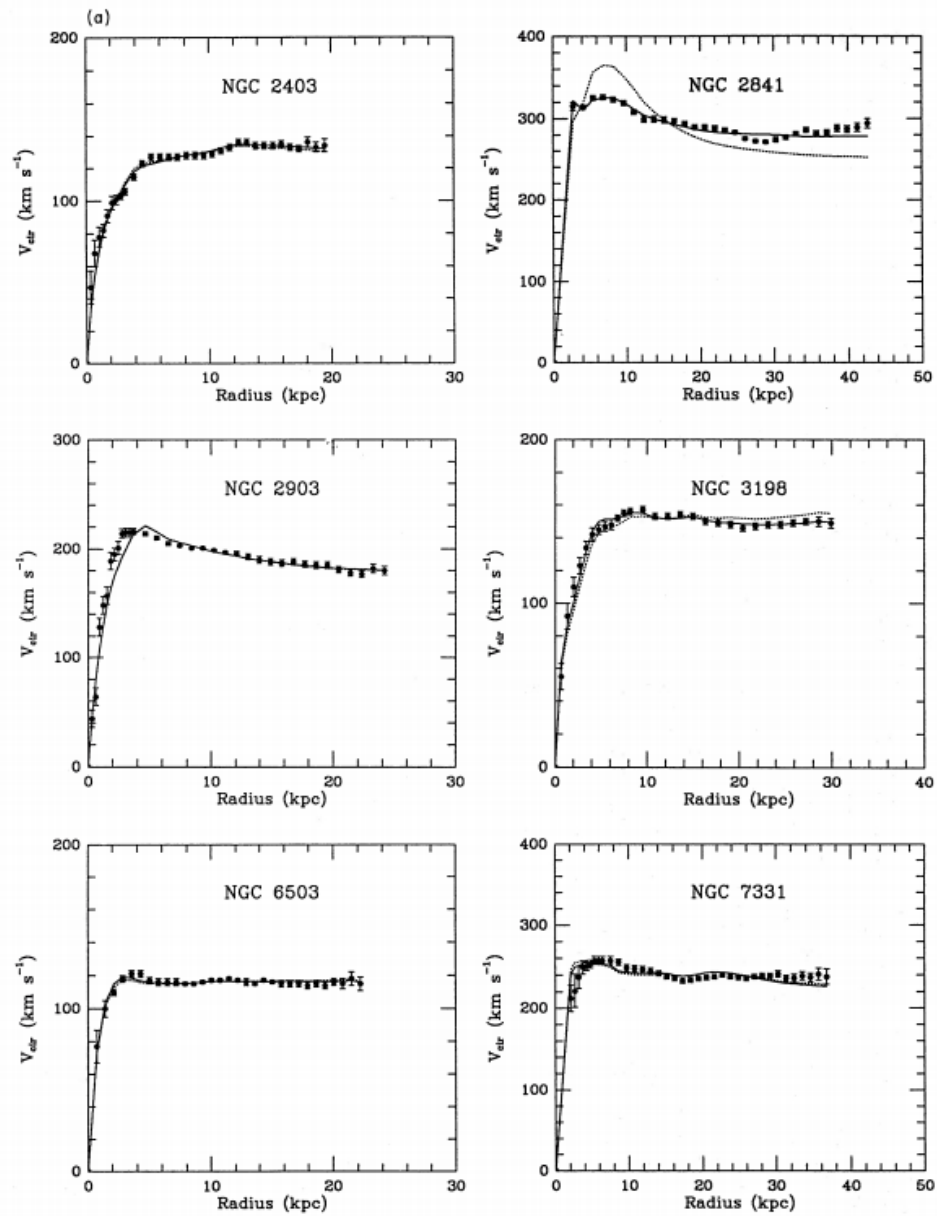


Figure 2. MOND fits to the rotation curves of the sample galaxies. The value of a_0 is fixed at $1.21 \times 10^{-8} \text{ cm s}^{-2}$. The dotted curves show the one-parameter fits (M/L) and the solid curves are the two-parameter fits (M/L and distance). In the gas-rich galaxies, NGC 1560, DDO 154, and DDO 170, M/L effectively disappears as a fitting parameter because of the dominant contribution of the gas to the total mass. Best-fit values of M/L are given in columns 2 and 5 and the best-fit distance in column 6 of Table 3.

Figure 1.5: Figure 2 of [16], reproduced here for illustration of rotation curve modeling.

$\mu \rightarrow 1$ for $x \gg 1$ and $\mu \rightarrow x$ for $x \ll 1$, with $a_0 \sim 1.2 \cdot 10^{-10} m/s^2$.

The immediate problem with this initial proposal as a theory is that it violates essentially all physical conservation laws, and in a few years Milgrom and Bekenstein produced a theory starting from the Lagrangian of Newtonian mechanics (to ensure conservation of energy and momentum), which allowed them to reproduce in a more rigorous way the scaling law shown above for spherical symmetry within some limits[17]. The modern theory, Tensor-Vector-Scalar gravity (TeVeS), now incorporates relativity and can account for much more complex systems than the original version of modified gravity, but at the expense of a handful of new parameters. In addition, despite this additional freedom, TeVeS has thus far failed to match modern cosmological measurements, described in the next few sections, which has led to the particle theory of dark matter steadily gaining traction over the past 20 years[19].

A recent development worth mentioning at this point, despite being chronologically out of order, is the work of Eric Verlinde on emergent gravity[106]. If the main criticism of past theories of modified gravity were that they had become too convoluted and were not rooted in any deep laws of physics, then Verlinde's elegant use of thermodynamics to allow gravity to become an emergent phenomenon is just the opposite. As it is a new and as-yet still developing theory, I will not comment further, other than to note that as of this writing it has successfully accounted for some rotation curves, but has yet to be accurately applied to the evidence we will discuss in the next sections, and therefore nothing more precise can be said than it is perhaps a more elegant and natural form of the Bekenstein-Milgrom theories that came before it. In addition, despite the initial success, there has already been pushback by groups claiming that in highly gas-rich galaxies, emergent gravity cannot match observed rotation curves[62].

1.2.3 Gravitational Lensing

In response to the debate pitting modified gravity against particle dark matter, an additional means of measuring gravitational potential becomes a powerful tool to add to optical and radio rotation curves that is independent of the light emitted by the galaxy. General Relativity (GR) predicts that a gravitational potential well should cause light to bend around a large mass, creating a gravitational lens.

Consider the most basic example of gravitational lensing in GR, the Schwarzschild lens[96]. Using the simple metric around an isolated mass distribution, the angular deflection due to the point mass is

$$\Theta = \frac{4GM}{rc^2} = \frac{2R_s}{r} \quad (1.42)$$

where M is the mass of the lens and r is the impact parameter of the lens, or distance to the lens in the plane perpendicular to the observer, and R_s is the Schwarzschild radius

$$R_s = \frac{2GM}{c^2} \quad (1.43)$$

This formula is a good approximation for small angles in the limit that $r \gg R_s$, which for very large masses is an excellent approximation on cosmological scales. In the simplest lensing case, we can consider a source directly behind the lensing object; at some critical radius r_c we will observe the light from this object for all angles about the lensing object, forming a classic 'Einstein ring'. In this case, we can use the geometry of the problem to determine a desired unknown as follows. We know that by symmetry, for us to actually observe the object, it must be as far behind the lens as the lens is from us. Thus the triangle of the lens distance d , impact parameter r , and observed angle θ_{obs} give

$$\theta_{obs} = \arctan \frac{r}{d} \quad (1.44)$$

and thus the deflection angle is

$$\alpha = 2\left(\frac{\pi}{2} - \theta_{obs}\right) \quad (1.45)$$

giving the relations

$$\pi - 2\theta_{obs} = \frac{4GM}{rc^2} = \frac{4GM}{d \tan(\theta_{obs})c^2} \quad (1.46)$$

and

$$M_{lens} = \frac{c^2 d \tan(\theta_{obs}) (\pi - 2\theta_{obs})}{4G} \quad (1.47)$$

Thus if we know the distance to either the lens or the distance object, we can determine the mass enclosed within a sphere with the radius of the impact parameter. In many cases for these lenses it is possible to determine the distance to either the lens or lensed object through the cosmological distance ladder, allowing us to map the gravity well of a given lens.

This simple treatment breaks down when we depart from simple lens models or attempt lensing measurements within a continuous mass distribution, but is a good illustration of the principle and can be used without modification in a small number of cases where the assumptions are met. Most famously, one of the first confirmations of general relativity came from the measurement of gravitational lensing around the sun, which of course is highly symmetric and contained [\[96\]](#).

The key example of lensing measurements employed to address the dark matter problem comes in the form of the bullet cluster [\[30\]](#), shown in [Figure 1.6](#). Measurements of this galaxy cluster with Hubble Space Telescope (optical) and Chandra (X-ray) observations allow us to trace the mass distribution through the lensing profile, shown in the left panel of the figure, and the gas mass of the cluster, which is the majority of the system's baryonic mass, is shown in the right panel. Due to the recent collision of the two clusters, much of the interstellar gas has been heated and stripped by tidal forces, but the majority of the mass has proceeded through the collision seemingly unaffected. Thus we find that the majority of the gas is spatially separated from the majority of the mass, which does not seem to interact.

The bullet cluster observations show both that dark matter has a much weaker self-interaction, if a particle, than hydrogen, and that simple modifications to gravity, which nonetheless preserve




FIG. 1.— Shown above in the top panel is a color image from the Magellan images of the merging cluster 1E0657–558, with the white bar indicating 200 kpc at the distance of the cluster. In the bottom panel is a 500 ks Chandra image of the cluster. Shown in green contours in both panels are the weak lensing κ reconstruction with the outer contour level at $\kappa = 0.16$ and increasing in steps of 0.07. The white contours show the errors on the positions of the κ peaks and correspond to 68.3%, 95.5%, and 99.7% confidence levels. The blue +s show the location of the centers used to measure the masses of the plasma clouds in Table 2.

Figure 1.6: Figure reproduced from [30].

the central potential, cannot explain this phenomenon while dark matter can. This and subsequent cluster lensing studies have shown the same pattern repeats in a handful of other clusters that have recently undergone collisions, and together with constraints from the CMB is the main line of evidence used to argue against modified gravity. MOND in its simpler form can do better than Newtonian gravity, but still differs by a factor of 2 [19], while very exotic gravity can get closer to a full explanation, but invokes new fields that need to act much like new particles to be compatible with early universe cosmology.

This is an ongoing field, and since the first publication in 2006 [30] many papers with new lens measurements have been published [96]. Over the intervening decade, consensus has built that the most natural explanation for these phenomena are non-Baryonic particle dark matter. Most recently, work has gone into using mass contours from gravitational lensing to try to put lower limits on the dark matter self-interaction cross-section [53]. Due to the rarity of cluster collisions, it is unlikely that we will ever have a sample of more than a few dozen recently collided clusters, but it is possible that as more work is done on the upcoming series of dark energy campaigns (DES, DESI, LSST and others) more lensing candidates will be identified and a more statistically representative sample will arise that will begin to settle questions of cosmic variance versus naturalness.

1.2.4 Cosmic Microwave Background Radiation

The most compelling evidence for particle dark matter, and the reasoning that guides us towards the best candidates, are tied to the standard cosmological model as measured through effects only tractable through the application of general relativity [15]. The best tools we currently have to understand the constituents of the universe are general relativity, as a mathematical framework, coupled

¹⁵This section largely follows that found in the particle dark matter review by Bertone, Hooper, and Silk [20]

to Λ CDM, our current model for how the universe evolved. The Λ in this acronym stands for dark energy, which at present we believe to constitute about 75% of the energy density of the universe in the present epoch, and CDM is for cold dark matter, which controlled the evolution of the universe for the majority of its history until relatively recently.

We start with the Einstein field equation:

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = -\frac{8\pi G_N}{c^4}T_{\mu\nu} + \Lambda g_{\mu\nu} \quad (1.48)$$

where $R_{\mu\nu}$ and R are the Ricci tensor and scalar, $g_{\mu\nu}$ is the metric, G_N is the Newtonian gravity constant, $T_{\mu\nu}$ is the stress-energy tensor, and Λ is the cosmological constant. As pointed out in [20], this equation relates the geometry of the universe (on the left side) to the energy content of the universe (on the right). If we assume an isotropic and homogeneous universe, we get the line element

$$ds^2 = -c^2 dt^2 + a(t)^2 \left(\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right) \quad (1.49)$$

which defines the metric in the previous equation given $ds^2 = g_{\mu\nu}x^\mu x^\nu$. Here, $a(t)$ is the scale factor (the relative size of the spatial and time dimensions) and k is the curvature of space-time, where $k = 0$ is flat space. We can solve the Einstein equation for this metric to get the Friedmann equation

$$H(t)^2 + \frac{k}{a^2} = \frac{8\pi G_N}{3} \rho_{tot} \quad (1.50)$$

where we have introduced the time-dependent Hubble parameter

$$H(t) = \frac{\dot{a}}{a} \quad (1.51)$$

and ρ_{tot} is the total average energy density of the universe. This gives us a way to relate the relative expansion of the universe parameterized by $H(t)$ to the energy density of its constituents. We typically re-write this in terms of the dimensionless parameters

$$\Omega_i = \frac{\rho_i}{\rho_{crit}} \quad (1.52)$$

where

$$\rho_{crit} = \frac{3H(t)^2}{8\pi G_N} \quad (1.53)$$

and the Friedmann equation becomes simply

$$\Omega - 1 = \frac{k}{H(t)^2 a^2} \quad (1.54)$$

In a flat universe, where $k = 0$, this means that

$$\Omega = \sum_i \Omega_i = 1 \quad (1.55)$$

and for a universe with curvature Ω will then either be greater than or less than 1 depending on the sign of k .

These relations are useful for relating the energy density at earlier times in the universe to the current epoch; for Λ CDM with $k = 0$, we find that

$$\frac{H(t)^2}{H_0^2} = \left[\Omega_\Lambda (1+z)^{3(1+w)} + \Omega_M (1+z)^3 + \Omega_R (1+z)^4 \right] \quad (1.56)$$

where $\Omega_M = \Omega_b + \Omega_c$ is the mass density (in Baryons and cold dark matter respectively), Ω_Λ is the dark energy density, and Ω_R is the radiation energy density, all in the current epoch. z is the redshift, defined as

$$z = \frac{\lambda_{obs}}{\lambda_{emit}} - 1 = \frac{a_{then}}{a_{now}} - 1. \quad (1.57)$$

Thus we can use measurements of $H(t)$ in different epochs to constrain the relative energy density of different components of the universe as a function of time.

Measurements of the cosmic microwave background (CMB) capture the energy densities as they were at recombination ($z \approx 1100$) and allow us to use Λ CDM cosmology to project primordial densities forward in time, in essence measuring the current mass density in the universe as compared to dark energy and radiation. The CMB is essentially an initially constant temperature surface that, as it evolved forward in time, carried with it the local structure of the gravity well in which it last froze out. If we model the early universe as a baryon-radiation fluid and include additional damping dynamics from non-baryonic dark matter, we can determine the power spectrum of local over-densities and under-densities due to the relative strength of gravity and radiation pressure. The relative coupling of the radiation to the matter thus allows us to use these dynamics to determine the fraction of the matter that is baryonic. The relationship between the location of matter anisotropies and energy densities is non-analytic but can be modeled numerically, and by fitting simulations with given cosmologies to the observed CMB power spectrum we can constrain the relative densities of the constituent parts of the universe at recombination. For a more detailed discussion, I refer the reader to [20, and references therein].

The most stringent current constraints from large-scale CMB measurements come from the *Planck* satellite [82], which measured the CMB temperature anisotropy shown in Figure 1.7. The best-fit parameters for the nominal 6-parameter model (assuming equations of state for the dark matter and dark energy) give the matter density

$$\Omega_M = 0.313 \pm 0.013 \quad (1.58)$$




Figure 1.7: CMB temperature anisotropy measured by the *Planck* experiment [82]. The relative height of the first few peaks and their locations help to determine the relative fraction and importance of dark matter, Baryons, and radiation in the early universe.

and the relative baryon and cold dark matter fractions

$$\Omega_b h^2 = 0.02222 \pm 0.00023, \quad \Omega_c h^2 = 0.1197 \pm 0.0022 \quad (1.59)$$

where $h = H_0/100$, giving the co-moving cold dark matter to baryon density ratio

$$\frac{\Omega_c}{\Omega_b} = \frac{0.1197}{0.02222} \approx 5.4$$

This roughly matches the ratio found from modern analysis of galaxy clusters discussed earlier in the chapter; the dark matter outweighs hot intra-cluster gas by a factor of 5-6.

Due to the fact that the oscillations seen in the CMB are frozen at the time of recombination, observations made for $z < 1100$ can also implicitly measure aspects of these same oscillations seen in the CMB anisotropy. One particularly powerful method is to measure 3-d clustering of galaxies as a function of redshift; given that the mean correlation lengths should be the same throughout the universe at a given redshift, the evolution of this parameter traces the expansion of the universe and helps constrain the cosmological parameters. These Baryon Acoustic Oscillation (BAO) constraints are typically used as a complementary dataset to increase the precision of CMB measurements, along with ground-based observations of the higher-order modes of the CMB [82]. The addition

of this complementary data highlights non-trivial covariances in the cosmological parameters that affect the constraint of the matter fraction in the early universe.

Despite the success of using the CMB and BAO to match local observations historically, the complexity of fitting a very high-dimensional and highly covariant set of cosmological parameters has led to some tension between high redshift and local measurements [100, 21, 87]. Fortunately, the dark matter/baryon ratio is largely invariant to these uncertainties given that h cancels in this ratio, as shown above, and thus this is one of the results most robust to discrepancies between local and high-redshift measurements [82]. Nevertheless, these high precision cosmological studies are an incredibly powerful tool for understanding the role of dark matter in our universe, and it behooves the dark matter physicist to stay apprised of recent developments in these fields. The most recent paper discussing local/high- z discrepancies comes by way of the Dark Energy Survey, which demonstrates that by relaxing some of the less constrained priors on the fits, discrepancies in H_0 discussed by e.g. the SH0ES and H0LiCOW projects [87, 100, 21] can be resolved [31]. Relaxing these constraints represent interesting statements about physics either locally or in the early universe, and time will tell whether these discussions fulfill their promise of using cosmology to make fundamental physics measurements comparable to those possible in particle physics experiments.

1.2.5 Measurements of Local Mass Density

While the best evidence for the existence of dark matter comes from measurements on the galactic and cosmological scales, the local dark matter density is most relevant to our ability to study dark matter in earth-bound laboratories. As mentioned earlier in this chapter, the first estimates of the local dark matter mass density were done by Kapteyn and Oort in the 1920's using local stellar kinematics. A great review of the history of these measurements can be found in [84], and Figure 1.8 shows a summary of these measurements reproduced from that work.

The simplest way to measure the local dark matter density is through kinematics of the local stellar population. For gravity in the weak-field limit, we can relate the gravitational potential and mass distribution using Poisson's equation:

$$4\pi G\rho = \nabla^2\Phi = \frac{\partial^2\Phi}{\partial z^2} + \frac{1}{R} \frac{\partial}{\partial R} \left(R \frac{\partial\Phi}{\partial R} \right) \quad (1.60)$$

Suppose we have a disk that has a scale height z_0 much smaller than its radius scale R_0 , such that, in the local stellar vicinity,

$$\rho(R, z) \approx \rho(R_{sun}) \exp(-z/z_0), \quad (1.61)$$

which also implies that the radial term in equation 1.60 is necessarily small. This gives us a one-dimensional equation in z :

$$4\pi G\rho \approx \frac{\partial^2\Phi}{\partial z^2}. \quad (1.62)$$




Figure 1.8: History of measurements of local density from taken from [84]. We can see that few measurements were possible until the advent of modern surveys in the 90's but the initial measurements were fairly close to the modern expected value of $\sim 0.3 \text{ GeV}/c^2$.

Integrating once in z gives us an equation that will begin to look familiar:

$$4\pi G\Sigma(z_{max}) \approx \frac{d\Phi}{dz} = -F_z(z_{max}) \quad (1.63)$$

where $\Sigma(z_{max})$ is the surface mass density of the disk integrated out to the disk height in question. For stars far from the disk, this is just the equation of simple harmonic motion! This implies that for this geometry, we should observe stars of a given mass far from the disk to oscillate through the disc with angular frequency

$$\omega = \sqrt{\frac{4\pi G\Sigma}{m_{star}}} \quad (1.64)$$

and thus by measuring the distribution of stellar velocities of tracer stars, we should be able to infer the effective surface density of the disk for small enough volumes. This method is similar to those used originally by Kapteyn and Oort, where the inferred local dark matter density is thus

$$\rho_{dm}(z_{max}) = \frac{\Sigma(z_{max}) - \Sigma_b(z_{max})}{2z_{max}}. \quad (1.65)$$

The density is highly degenerate with the baryonic mass density $\Sigma_b(z)$, which is generally one of the larger systematics in these analyses.

As pointed out by [84], this method suffers from the many assumptions we needed to make along the way, and more sophisticated methods (including those used by Jeans) make fewer assumptions at the expense of mathematical complexity, and the constraint of the local density is in itself a very involved field of astrophysics. As shown in Figure 1.8, the modern accepted value for the local dark matter density is $\sim 0.3 \text{ GeV}/c^2$, with a significant error bar extending from 0.2 to 0.6 depending on the methods involved and the data used. Data being taken currently by the Gaia satellite promises to greatly improve the precision of this measurement in the coming years [84, 19, 20].

1.3 Dark Matter Phenomenology

In the previous sections, I reviewed the astrophysical motivation for the existence of a new type of matter that behaves differently than baryonic matter. In this section I will review the remaining candidates for this dark matter framed in the context of directly detecting interactions of dark matter particles χ with standard model particles in an earth-bound experiment, and then touch upon other candidates that have been ruled out by indirect means or by searches at colliders.

In the context of a direct detection experiment, the approach to studying dark matter can be reduced to the measurement of a differential event rate as a function of time in the presence of known backgrounds. For an isolated nucleus, we know that the event rate will be

$$\Gamma_n = n_\chi \sigma \langle v \rangle = \frac{\rho_\chi}{m_\chi} \sigma \int v f(v) dv \quad (1.66)$$

and so for a detector of mass M_{det} and nuclear mass m_n we have the total detector rate

$$\Gamma_{det} = \frac{M_{det}}{m_n} \frac{\rho_\chi}{m_\chi} \sigma \int v f(v) dv. \quad (1.67)$$

Taking the derivative of this with respect to energy transferred $E = q^2/2m_n$ and velocity v gives the differential rate

$$d\Gamma_{det} = \frac{M_{det}}{m_n} \frac{\rho_\chi}{m_\chi} (d\sigma/dE) v f(v) dE dv \quad (1.68)$$

This equation allows us to input the physics for our particular model. For the case of spin-independent scattering, we calculate the differential cross-section:

$$\frac{d\sigma}{dE} = 2m_n \frac{d\sigma}{dq^2} = \frac{m_n}{2\mu^2} \frac{\sigma_0}{v^2} F^2(|q|) \quad (1.69)$$

where

$$\mu = \frac{m_n m_\chi}{m_n + m_\chi} \quad (1.70)$$

and $F^2(|q|)$ encodes the momentum-dependence in the coupling of the dark matter to the standard matter particle it scatters off. This gives the final differential rate equation

$$d\Gamma_{det} = \frac{M_{det}}{m_n} \frac{\rho_\chi}{m_\chi} \frac{m_n}{2\mu^2} \sigma_0 F^2(|q|) \frac{f(v)}{v} dE dv \quad (1.71)$$

$$= M_{det} \frac{\rho_\chi}{2\mu^2 m_\chi} \sigma_0 F^2(|q|) \frac{f(v)}{v} dE dv \quad (1.72)$$

$$\frac{d\Gamma_{det}}{dE} = M_{det} \frac{\rho_\chi}{2\mu^2 m_\chi} \sigma_0 F^2(|q|) \int_{v_{min}}^{\infty} \frac{f(v)}{v} dv \quad (1.73)$$

where to conserve energy, we find for a given energy transfer there is a minimum allowed velocity

$$v_{min} = \sqrt{\frac{m_N E}{2\mu^2}}, \quad (1.74)$$

and the velocity distribution has an implicit cutoff at v_{esc} , the escape velocity of the galaxy. This is the rate equation for direct detection of dark-matter nucleon scattering given in the often cited dark matter reviews [\[64, 54\]](#).

I have provided the derivation again here to allow us to see how this is modified when we change from a single nucleon to a coherent atomic scattering model later in this section, and to illustrate how the different aspects of the calculation (velocity distribution, target mass, differential cross-section) can be changed to compute rates for different models. We can generalize a bit more by explicitly splitting these components. First, we have the inverse velocity expectation value (following the

notation of [33])

$$\eta(v_{min}) = \int_{v_{min}}^{\infty} \frac{f(v)}{v} dv = \int_{v_{min}}^{\infty} \frac{g_{\chi}(\mathbf{v})}{v} d^3\mathbf{v} \quad (1.75)$$

where we've explicitly used the 3-dimensional velocity distribution rather than an isotropic 1-dimensional velocity distribution to make the formula more general (here using the convention that $\int f(v)dv = \int f(\mathbf{v})d^3\mathbf{v} = 1$). We can make this a dimensionless quantity of order unity by multiplying by $v_0 = 220\text{km s}^{-1}$, giving the rate equation [54]

$$\frac{d\Gamma_{det}}{dE} = \left(M_{det} \frac{\sigma_0}{2\mu^2 v_0} \frac{\rho_{\chi}}{m_{\chi}} \right) F^2(|q|) (v_0 \eta(v_{min})) \quad (1.76)$$

The first term in parentheses sets the baseline rate, and the second two dimensionless terms will often be of order unity, so the leading term can be used to estimate the differential rate below the velocity cut-off and set the rate scale for a given detector mass, dark matter mass, and target composition.

In the following sections, I will present different theories for the composition of particle dark matter, and conclude with a section discussing the non-particle (astrophysical) candidates that have not been ruled out. I will use this section as a reference when discussing science results later in this thesis, but also to illustrate how differing kinematics between these classes of theories allow us to explain dark matter relic abundance, and how we can modify our scattering formalism to predict energy spectra given new DM-SM couplings.

1.3.1 WIMP Dark Matter

For the first two decades of direct detection experiments, the focus of these searches was on 'Weakly Interacting Massive Particles' (WIMPs), motivated by the growing prominence of Supersymmetric models with a stable lightest Supersymmetric particle (LSP), which is a natural dark matter candidate [54, 64, 18, 20]. In this section I will review the arguments for WIMP dark matter, and summarize the studies that have now largely ruled out the best-motivated class of WIMPs. These models are still viable but the parameter space has shrunk as new searches have excluded larger areas of parameter space, and as the LHC has failed to turn up any evidence of Supersymmetry (SUSY).

The first place to start when considering particle dark matter is to determine which candidates can reproduce the cosmological abundance measurement discussed in the previous section. Here I follow the derivation of Ref [54]. Suppose we have a new unknown particle χ that was in equilibrium with standard model particles in the early universe, such that $\chi\bar{\chi} \rightarrow l\bar{l}$ and $l\bar{l} \rightarrow \chi\bar{\chi}$ occur at equal rates, where l are some other lighter standard model particles. The equilibrium number density at some temperature T is given by

$$n_{\chi} = \frac{g}{2\pi^3} \int f(\mathbf{p}, T) d^3\mathbf{p} \quad (1.77)$$

where $f(\mathbf{p}, T)$ is the density of states (either Fermi-Dirac or Bose-Einstein) and g is the number of internal degrees of freedom of the particle. For $T \gg m_\chi$ (in units of energy) we find that

$$n_\chi \propto T^3 \quad (1.78)$$

while for $T \ll m_\chi$ we find the Boltzmann limit

$$n_\chi \approx g \left(\frac{m_\chi T}{2\pi} \right)^{3/2} \exp(-m_\chi/T) \quad (1.79)$$

This second equation will thus predict that the density always goes to 0 as the universe cools, but when the expansion rate exceeds this annihilation rate, the particles are no longer in thermal equilibrium with the universe and their annihilation will stop, leaving a relic density defined by the given mass and freeze-out temperature T_f at which this occurs.

The annihilation rate of the dark matter into standard model particles is given by the equation

$$\Gamma(T) = n_\chi(T) \langle \sigma_A v \rangle \quad (1.80)$$

and the relic density at freeze-out is given by setting $\Gamma(T_f) = H(T_f)$. We thus find that

$$n_\chi(T_f) = \frac{H(T_f)}{\langle \sigma_A v \rangle} \quad (1.81)$$

and substituting the radiation dominated form $H(T) = 1.66g_*^{1/2}T^2/m_{pl}$ [54] we find

$$n_\chi(T_f) = \frac{1.66g_*^{1/2}T_f^2}{m_{pl}\langle \sigma_A v \rangle} \quad (1.82)$$

where m_{pl} is the Planck mass, and g_* is the effective number of degrees of freedom. Dividing through by the entropy density ($s \approx 0.4g_*T^3$) gives

$$\frac{n_\chi}{s}(T_f) = \frac{4.15}{g_*^{1/2}T_f m_{pl} \langle \sigma_A v \rangle} \quad (1.83)$$

For weak-scale processes we find that solving the boltzmann equation at this point yields $T_f \approx m_\chi/20$, which gives the equation

$$\frac{n_\chi}{s}(T_f) \approx \frac{100}{g_*^{1/2} m_\chi m_{pl} \langle \sigma_A v \rangle} \quad (1.84)$$

putting this in terms of measured quantities, and recognizing that the ratio on the left-hand side is

constant in time, we have

$$\langle\sigma_{Av}\rangle = \frac{100}{g_*^{1/2} m_{pl}} \frac{s_0}{\rho_\chi} = \frac{100}{g_*^{1/2} m_{pl}} \frac{s_0}{\rho_c \Omega_\chi h^2} \quad (1.85)$$

where s_0 is the current entropy density, ρ_c is the critical density, Ω_χ is the fraction of energy composed of dark matter in the universe, and h is the reduced Hubble constant. This is typically written (with standard assumptions) as

$$\langle\sigma_{Av}\rangle \approx \frac{3 * 10^{-27} cm^3 s^{-1}}{\Omega_c h^2} \quad (1.86)$$

Given that the denominator is of order unity, we thus have an estimate to an order of magnitude of the cross-section, assuming thermal velocity, which is conveniently the same scale as the weak force [54]. Note the assumptions that went into this calculation; we'll re-visit some of them in the next section.

For a numerically complete solution, the equilibrium density is found by solving the Boltzmann equation [54, 95]

$$\frac{dn_\chi}{dt} + 3Hn_\chi = -\langle\sigma_{Av}\rangle [n_\chi^2 - (n_\chi^{eq})^2] \quad (1.87)$$

the results of which can be seen in Figure 1.9. This figure is an excellent illustration of the robustness of the approximate result shown above; for the standard inflationary model, the relic density is entirely determined by the velocity averaged annihilation cross-section $\langle\sigma_{Av}\rangle$. This provides an important input to selecting theories with stable dark-matter like particles based on expectations for this cross-section, but (aside from logarithmic corrections) does not specify the dark matter mass.

When this result came out, two facts were apparent: that the predicted cross-sections are what would be expected for a weak-scale interaction, and that SUSY, which was expected to be seen at the LHC, easily accommodated dark matter candidates of the right mass and interaction scale. Any complete description of SUSY is beyond the scope of this thesis, and I refer the reader to Ref [54, 18, 20] and references therein as a starting point for a more in-depth discussion. What is necessary for this thesis is to recognize that SUSY is a framework that describes an entire class of models, and was not proposed to explain dark matter. Given that SUSY predicts an entirely new spectrum of uncharged neutralino particles, SUSY theories provide natural and well-motivated dark matter candidates. Figure 1.10 shows a range of WIMP candidates that arise naturally in SUSY and are consistent with the measured relic density.

For direct detection, the fact that we have a predicted present-day annihilation rate and a predictive means to relate this to the nucleon-scattering cross-section means that we can predict, for a given neutralino, what the observed differential scattering rate should be in an earth-bound




Figure 1.9: The canonical figure showing how relic density for the full numerically solved Boltzmann equation is entirely determined by the velocity-averaged annihilation cross-section at freeze-out, taken from Ref [57].

detector. Refs [20, 54, 64] give the WIMP-nucleus scattering coefficients σ_0

$$\sigma_{0spin} = \frac{32}{\pi} G_F^2 \mu^2 \Lambda^2 J(J+1) = \frac{32}{\pi} G_F^2 \mu^2 (J+1) [a_p \langle S_p \rangle + a_n \langle S_n \rangle] \quad (1.88)$$

$$\sigma_{0scalar} = \frac{4\mu^2}{\pi} [Zf_p + (A-Z)f_n]^2 \quad (1.89)$$

where G_F is Fermi's constant, J is the nuclear angular momentum, $\langle S_p \rangle$ ($\langle S_n \rangle$) is the proton (neutron) spin expectation value, Λ encodes the expected spin, Z is the number of protons, and A is the number of nucleons in the target detector. The constants a and f encode the coupling of the WIMP to the nucleons, thus allowing the rates to be tied to the model being tested. It is in general much harder to probe spin-dependence, given the limited number of targets with appreciable nuclear spin, and




Figure 1.10: Scatter plot of relic neutralino density versus neutralino mass for a set of SUSY models from Ref [54]. The highest density of models consistent with astrophysical observations is found in the 20-100 GeV range, but masses up to 300 GeV are viable candidates.

the standard has become to quote the nucleon cross section assuming spin-independent (or scalar) scattering for better comparison between experiments with different target nuclei. We can see from $\sigma_{0,scalar}$ that the overall event rate scales as A^2 in the case that $f_p \sim f_n$, so limits are put on σ_0 and converted to σ_n as

$$\sigma_n \approx \frac{\sigma_{0,meas}}{A^2} \quad (1.90)$$

For these measurements, the form factor is taken from experimentally measured nuclear form factors for given materials, and event rates are thus numerically computed, though they can be well fit with approximate forms; see Refs [54, 64] for more details. In most cases, $F(Q) \sim 1$ is a good order of magnitude approximation. Most of the shape of the differential event rate comes from momentum and energy conservation of elastic scattering and the shape of the WIMP velocity distribution, rather than significant changes in cross-section with momentum.




Figure 1.11: Current spin-independent dark matter-nucleon cross-section limits from Ref [91]. Overlaid are contours for dark matter candidates in SUSY models not excluded by LHC run 1.

The current state of direct-detection constraints for nuclear-recoil dark matter can be seen in Figure 1.11 compared to the current SUSY candidates not excluded by recent results from the LHC. The limits (and the few claimed discovery contours) in this plot show the progress made in direct-detection over the past decade, and show that the 'natural' weak-scale WIMP ($\sigma \sim 10^{-43} \text{cm}^2$, $m_\chi \sim 20 - 400 \text{GeV}/c^2$) has been almost entirely ruled out. The success of Xenon-based experiments show the sensitivity gained from the coherent enhancement of event rates in massive nuclei, and the future of high-mass dark matter searches will be driven by a new generation of ton-scale Xenon detectors in low-background environments. When the mass of the dark matter particle drops below the Nucleon mass, these technologies rapidly lose dark matter sensitivity, and the lower masses are dominated by low-threshold Si/Ge experiments, which also benefit from a higher dark matter number density and thus do not need to have such a large fiducial mass. These considerations will be discussed in much more detail later in this thesis.

The limits in Figure 1.11 also show results from the DAMA/LIBRA collaborations, which have shown an annual modulation in event rate at increasingly higher significance for two decades [14]. Since the velocity of the earth relative to the galactic center changes as the earth revolves about the sun, there will be an annual modulation in the velocity distribution of WIMPs encountered by earth-bound detectors. While this is true, it has proven difficult to show that any observed annual modulations were due to dark matter and not other seasonal backgrounds, and experiments that

are much more sensitive than the initial DAMA experiments have failed to see either an annual modulation or an unknown background consistent with dark matter. As a result the community has moved away from looking for annual modulations as a discovery approach, and the DAMA result stands as a curiosity that is inconsistent with other results. As of this writing, there are a series of planned follow-up experiments that aim to try to reproduce the DAMA experiment and understand this yearly modulation. For more details see Ref [14] and references therein.

In summary, while WIMP dark matter was very promising, the lack of evidence for SUSY and the null results coming out of the large liquid Noble detectors has made supersymmetric WIMP dark matter no more attractive than the other proposed models, and the community has begun to branch out into other mass ranges motivated by non-supersymmetric theories, including WIMP masses down to and below that of a proton motivated by asymmetric dark matter and other theoretic models. In the rest of this section I will consider the other non-WIMP candidates for dark matter in the context of their relation to the established field of WIMP dark matter searches.

1.3.2 Sub-GeV Dark Matter

In the previous section I showed that the relic density calculation was mass-independent, and the WIMP miracle suggested a mass range to start the search motivated by naturalness and SUSY. Now that the most obvious place to look has been ruled out, the field is starting to spread out to cover a larger mass and cross section range, and considering a wider range of theories (which solve other problems in particle physics) that could accommodate dark matter candidates. A subset of these models is summarized in Figure 1.12.

One particularly interesting class of models are those in the keV-GeV range, which are most generally described by the simplified model of a hidden sector, where dark matter (DM) interacts with the standard model (SM) via a new force mediator. The simplicity of this picture allows us to relate the dark matter and mediator masses to the annihilation cross section and determine where the interesting mass/cross section range would lie for a given model. There are 3 cases generally considered in the literature [2, 51]:

- 'Secluded' Freeze-out ($m_\chi > m_{med}$) - The DM is not the lightest particle in the dark sector, and is frozen out to the mediator, which subsequently continues to interact with SM particles. In this case the annihilation cross-section has the form

$$\langle \sigma_{Av} \rangle \propto \frac{g_D^4}{m_\chi^2} \quad (1.91)$$

- 'Direct' Freeze-out ($m_\chi < m_{med}$) - The DM is the lightest dark sector particle and freezes out




Figure 1.12: Summary of dark matter candidates from Ref [2].

from SM particles through a heavy mediator

$$\langle \sigma_{Av} \rangle \propto \frac{g_D^2 g_{SM}^2 m_\chi^2}{m_{med}^4} \tag{1.92}$$

- Freeze-in - If the dark sector is initially empty and does not thermalize, then the thermal relic paradigm is not applicable; instead the relic can be produced by SM particles freezing out into light dark-sector mediators, some of which decay into the DM.

These mechanisms allow us to determine regions of parameter space in which to look for these simplest models, and motivate tailored searches for lower mass particles. Lower-mass DM candidates also imply higher number densities, given that we know the mass density, and can be probed by smaller experiments.

The challenging aspect of searching for these lighter DM particles, however, is that the scattering kinematics result in a very small momentum transfer, and thus require much more sensitive detectors.

To see why this is the case, consider the maximum energy transfer in an elastic two-body collision:

$$\Delta E_{max} = \frac{2\mu^2 v_\chi^2}{m_T} \quad (1.93)$$

$$\approx \frac{2m_\chi^2 v_\chi^2}{m_T} \quad (m_\chi \ll m_T) \quad (1.94)$$

$$\approx 2m_T v_\chi^2 \quad (m_\chi \gg m_T) \quad (1.95)$$

where m_T is the target mass and v_χ is the velocity of the incoming WIMP. Considering a DM particle moving near the galactic escape velocity $v_\chi \sim 600$ km/s gives the numerical bounds

$$\Delta E \lesssim 70 \text{ eV} \left(\frac{m_\chi}{500 \text{ MeV}} \right)^2 \left(\frac{m_T}{28 \text{ GeV}} \right)^{-1} \quad (m_\chi \ll m_T) \quad (1.96)$$

$$\lesssim 4 \text{ eV} \left(\frac{m_T}{511 \text{ keV}} \right) \quad (m_\chi \gg m_T) \quad (1.97)$$

where I've used numbers for Si and a free electron in order to give the most favorable estimates. These energy scales (both below 100 eV) are well below the current energy threshold of the most sensitive liquid noble dark matter experiments, and are only starting to become accessible with lighter target materials.

One fortunate caveat is that these equations hold for elastic collisions, but not for inelastic collisions between bound electrons and dark matter. Because the electron is in a bound state, its momentum is not definite, and in principle its momentum can be arbitrarily high with non-zero probability. Ref [33] shows that the momentum transfer bound is

$$\Delta E \leq \frac{1}{2} \mu v_\chi^2 \approx 1 \text{ keV} \left(\frac{m_\chi}{500 \text{ MeV}} \right) \quad (m_\chi \ll m_T) \quad (1.98)$$

This is a much larger energy scale, and should dominate any experimental spectra produced by DM in detectors if the electron and nuclear interaction cross sections are within a few orders of magnitude of each other. For this reason, limits on DM-SM cross sections are given in terms of the electron-scattering interaction cross section instead of the nucleon cross section, as is done for DM in the WIMP regime.

The mathematical framework for light dark matter scattering on electrons, especially in semiconductors, is not quite as standardized as it is for WIMP-nucleon scattering, so here I will defer to the formalism of Ref [33] when there is a question of which definition to use. We start with a decomposition similar to earlier, where we separate a normalized form factor $F(q)$ from a cross-section scale $\bar{\sigma}_e$ (similar to σ_0) to get the bound-state cross-section scale [33]

$$\bar{\sigma}_e = \frac{|M_{free}(\alpha m_e)|^2}{16\pi(m_\chi + m_e)^2} \quad (1.99)$$

where $\overline{M_{free}}$ is the matrix element for free DM-electron scattering averaged over all initial and final spin states. The differential event rate can be written in a similar but slightly different form than for the WIMP [33]:

$$\frac{d\Gamma_{det}}{dE} = \left(M_{target} \frac{\rho_\chi}{m_\chi} \frac{\bar{\sigma}_e \alpha}{2\mu_{\chi e}^2 v_0} \frac{m_e}{m_T} \right) \int dq \left(\frac{m_e}{q^2} v_0 \eta(v_{min}(q, E)) \right) F_{DM}(q)^2 |f_{crystal}(q, E)|^2 \quad (1.100)$$

where $f_{crystal}$ is the crystal form factor determined from transition rates between electron eigenstates in the crystal lattice. This rate is obviously more complicated to compute (owing to the lattice dynamics), and electron scattering is much simpler in detector materials where electron states are more continuous. This means the electrons are closer to being free-electrons, and hence get less of a kinematic boost from their binding nucleus.

This equation is re-factored so that the integral is a dimensionless function of the variables and the leading term contains the rate normalization. Comparing this to equation [1.76] we find that the leading term is very similar, except for the replacement of σ_0 with $\sigma_e \alpha$ and the addition of a scale factor $\frac{m_e}{m_T}$. The major difference is that, since we're in a crystal lattice, we need to integrate over q and E separately rather than separating the mean inverse speed from the form factor. It is simple to show that in the limit of continuous states, the form factor reduces to delta functions in energy and momentum, and this equation becomes much closer to that in equation [1.76].

The current limits as well as future prospects for select semiconductor targets can be seen in Figure [1.13]. Unlike the WIMP limits shown in Figure [1.11], the majority of the simplest relic density parameter space has not yet been probed, and light dark matter is in a sense low-lying fruit for the upcoming series of small low-threshold experiments.

Light dark matter (defined as dark matter in the keV-GeV range) is thus a viable candidate to explain the current dark matter density, and requires the development of a class of detectors distinct from those design to search for WIMP dark matter. Light dark matter is well motivated and can be probed by experiments much smaller and less complex than those needed for the remaining WIMP parameter space. This is an ever-growing field, and I refer the reader to Refs [51, 2, 1] for more in-depth discussion and the references therein.

1.3.3 Bosonic Dark Matter

Figure [1.12] clearly shows a lower limit for hidden-sector DM in the keV range, but allows for other models below that range under the heading 'Ultralight Dark Matter'. The lower limit for fermionic dark matter is around a few keV and is set by constraints from Lyman-alpha forest measurements on substructure formation [15]; below this mass the dark matter is numerous enough to allow its degeneracy pressure at early times to have a measurable effect on substructure formation in the early universe [2]. This does not preclude ultra light bosonic dark matter, however, which does not produce such degeneracy pressure, and thus the lower-limit for ultra-light dark matter has a much

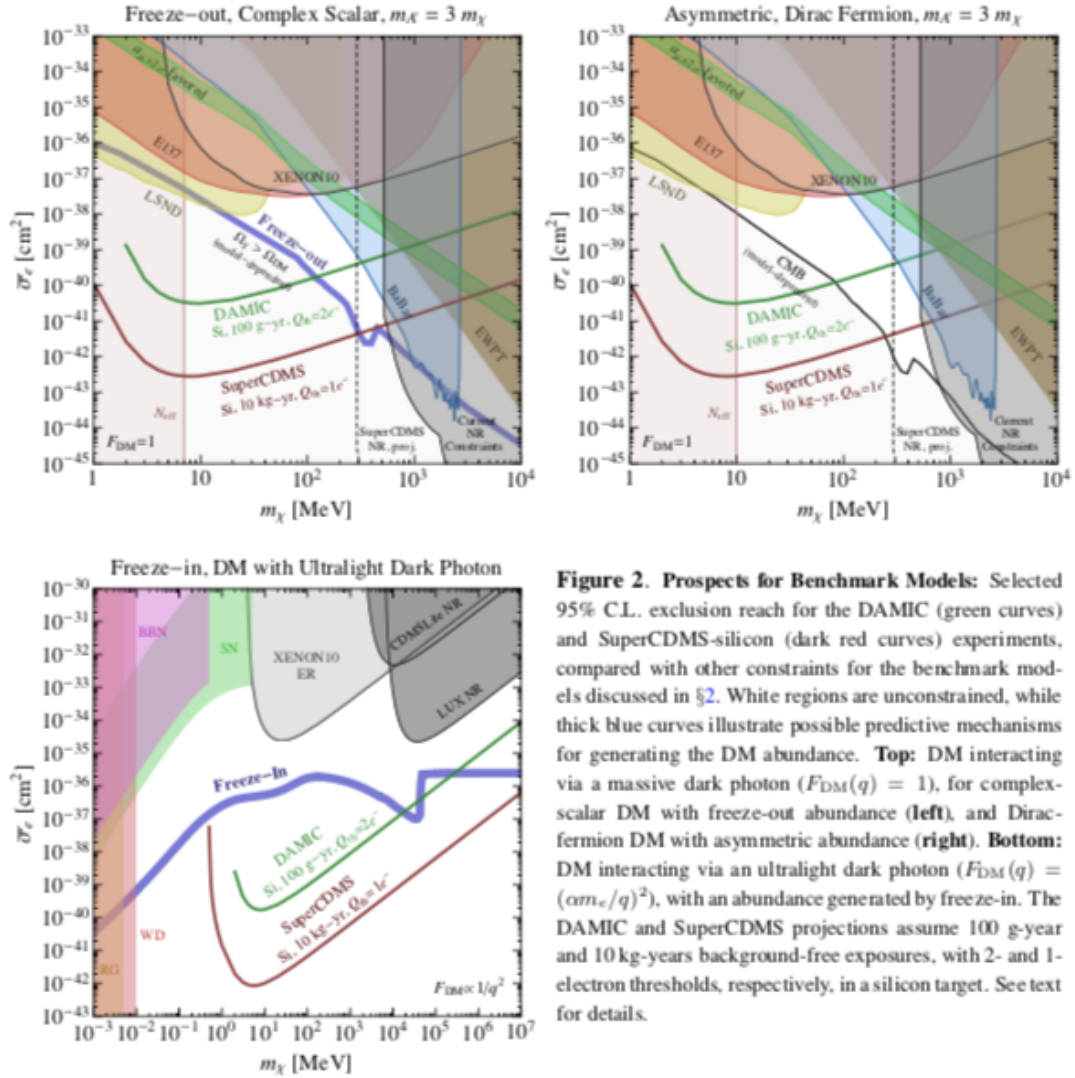


Figure 2. Prospects for Benchmark Models: Selected 95% C.L. exclusion reach for the DAMIC (green curves) and SuperCDMS-silicon (dark red curves) experiments, compared with other constraints for the benchmark models discussed in §2. White regions are unconstrained, while thick blue curves illustrate possible predictive mechanisms for generating the DM abundance. **Top:** DM interacting via a massive dark photon ($F_{\text{DM}}(q) = 1$), for complex-scalar DM with freeze-out abundance (left), and Dirac-fermion DM with asymmetric abundance (right). **Bottom:** DM interacting via an ultralight dark photon ($F_{\text{DM}}(q) = (\alpha m_\chi/q)^2$), with an abundance generated by freeze-in. The DAMIC and SuperCDMS projections assume 100 g-year and 10 kg-years background-free exposures, with 2- and 1-electron thresholds, respectively, in a silicon target. See text for details.

Figure 1.13: Figure 2 from Ref [33] showing current limits and reach of upcoming experiments compared to the thermal relic targets for some simple models of hidden sector light DM. Note that the freeze-in scenario can only be probed by direct detection experiments, and is otherwise only limited by astrophysical constraints.




Figure 1.14: (Reproduced from Figure 7 of Ref [34]). Parameter space for hidden-photons with mass $m_A < 1$ MeV. Colored regions are: experimentally excluded regions (dark green), constraints from astronomical observations (gray) or from astrophysical, or cosmological arguments (blue), and sensitivity of planned and suggested experiments (light green) (see Ref [34] for more details)

lower bound in mass.

Consider, for example, a freeze-out scenario where the lightest dark sector particle is the bosonic force mediator and the thermal relic is a sea of dark photons. The simplest detection scenario is through kinetic mixing with the SM photon and results in a complete absorption in the target medium. For these events, the rate in a perfect detector would look like a delta-function in energy, and would present the event rate[47]

$$\Gamma = \frac{M_{det}}{\rho_{det}} \frac{\rho_\chi}{m_A} k_{eff}^2 \sigma_1(m_A) \quad (1.101)$$

where m_A is the dark photon mass, σ_1 is the real part of the conductivity in the given target for a photon with an energy equal to the dark photon mass, and $k_{eff} \propto \kappa$ is the effective kinetic mixing parameter κ including in-medium effects that alter the kinetic mixing properties from their vacuum expectations[47].

The majority of bosonic dark matter models, which cover a very broad mass range and include the well-motivated QCD axion, are beyond the scope of this thesis, and I refer the reader to Refs [34, 11, 2] for a thorough review. A summary of the current constraints on bosonic dark matter can be seen in figure 1.14

1.3.4 Other Dark Matter Candidates

Coming back to Figure [1.12](#), I have discussed dark matter candidates in the range from well below an eV and up to a few hundred TeV. There is a large break above the WIMP mass range to macroscopic-sized objects. The few natural candidates above this bound have been largely ruled out over the past two decades; most of the baryonic candidates are heavily disfavored by the excellent fit to the CMB power spectrum, which clearly implies that dark matter is non-baryonic. Nevertheless, there remain a few possibilities for very heavy dark matter:

- **MAssive Compact Halo Objects (MACHOs)** - Shortly after the WIMP paradigm was suggested, the possibility that there were a multitude of compact objects that constituted the dark matter was proposed. The MACHO collaboration searched for MACHOs in the Milky Way halo by looking for micro-lensing events, and put an upper limit on the DM fraction of 8%. The CMB measurement of Baryon fraction makes it highly likely that this is an over-estimate and precludes MACHOS from being a dominant fraction of the dark matter (see Ref [19](#) and refs. therein).
- **WIMPZillas** - There is a strong upper limit on the WIMP mass that comes from the unitary bound; depending on the assumptions made this limit is $\sim 30\text{--}300$ TeV [20](#). This limit only applies for thermal relic DM, so it is possible that much heavier particle DM exists, but it would need to be produced in a freeze-in scenario. The motivation for WIMPZillas comes from unexplained high-energy cosmic rays that would be naturally produced by a WIMPZilla decay [20](#).
- **Primordial Black Holes** - One of the few still viable non-particle dark matter candidates is primordial black holes, created in various modified inflation scenarios. Many of the potential mass-ranges have been ruled out, and current constraints are discussed in detail in Ref [27](#). The modifications needed to early universe cosmology to accommodate these primordial black holes are significant, but given current inflationary constraints they cannot be ruled out [19](#).

This section is much shorter than it would have been even 10 years ago, as increasingly precise cosmological measurements constrain the baryonic content of the universe, and local observations continue to rule out compact objects. The particle candidates represent the most compelling explanation for the host of observations we attribute to dark matter, and modified gravity cannot yet explain the CMB and lensing phenomena. Thus where once the ‘dark matter’ was thought to just be undiscovered objects, the field has converged on the idea that new particles and potentially new forces are increasingly needed to explain both dark matter and a number of standard model anomalies. As cosmology continues to increase in precision, and gravitational wave observatories operate for longer and at better sensitivities, we will continue to constrain early-universe cosmology and further narrow the field of candidates, potentially ruling out all non-particle dark matter or

discovering that primordial black holes should indeed be expected in a certain mass-range. The upcoming CMB-S4, LSST, and SKA experiments will make complementary measurements that will vastly expand our understanding of the high-redshift universe and promise to shed more light on the nature of dark matter.

1.4 Physics Reach of Low-Threshold Detectors

In this chapter I've laid out the case that 1) dark matter exists and 2) particle dark matter in the eV-TeV range is well motivated from an early-universe cosmological perspective. In this section I will briefly motivate the work presented in this thesis in terms of the discovery potential of new detectors for dark matter across this mass range. I will focus primarily on the reach of semiconductor and superconductor technologies, and summarize the full discovery potential of all current methods under development to achieve sensitivity to lower-mass dark matter.

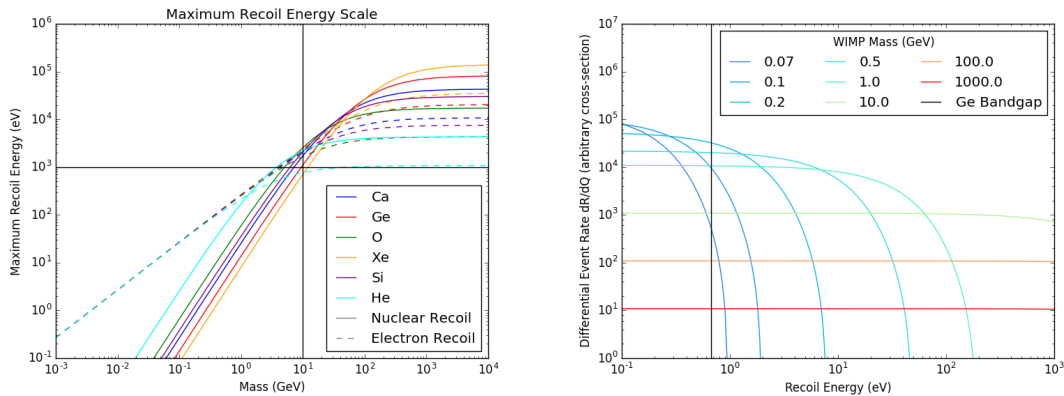


Figure 1.15: Left: Maximum energy transfer to a target through nuclear (solid) or inelastic electron (dashed) scattering for a variety of target and dark matter masses, assuming the dark matter is moving at $v_\chi \sim 600$ km/s, roughly the galactic escape velocity. Higher energies are possible if the earth's velocity increases the center of mass velocity of the DM-Detector system, but the energy limit is captured by the figure to better than an order of magnitude. Right: Differential recoil spectra for nuclear recoils at various WIMP masses in Ge, showing that as we go to lower mass, the maximum energy drops but the overall rate increases due to the increased number density. The bandgap of Ge is shown to indicate the minimum energy at which electron-hole pairs can be produced, important for voltage-mediated semiconductor detectors.

The challenge of detecting dark matter masses below that of the target nucleus is summarized by equations [1.96](#)–[1.98](#), where it was shown that energy transfer is inefficient for elastic scattering but can be boosted by inelastic electron-mediated atomic scattering. Figure [1.15](#) shows the recoil energy transfer at $v \sim 600$ km/s for a variety of WIMP masses and targets for both elastic nuclear and inelastic electronic scattering, as well as a series of nuclear recoil spectra for WIMPs at masses

in the MeV-TeV range. The crosshairs in the left-hand panel show that the mass range is cleanly bifurcated into masses above 10 GeV, which are accessible by experiments with keV thresholds, and masses below 10 GeV, which need eV-scale thresholds for nuclear and electronic recoils. There is also a tradeoff where the high-mass is dominated by nuclear recoil while the low mass is dominated by inelastic scattering.

The majority of detectors overcome low energy backgrounds (IR and cosmogenic radioactivity) by employing charge production either as the primary means of measuring energy deposition or as a means of discriminating between event types, and thus detectors designed with eV-scale energy thresholds also require eV-scale work-functions, such as those found in semiconductors. As can be seen from Figure I.15, liquid nobles (with the exception of Helium) pay the largest nuclear-recoil penalty at low DM mass; coupled with their relatively high ionization energies (~ 12 eV in Xe, ~ 16 eV in Ar) and propensity for non-gaussian charge leakage, these detectors face significant challenges achieving sub-keV energy thresholds. Semiconductors on the other hand have eV-scale gaps, and are crystalline, so they have much more predictable behavior. Superconductors can achieve even lower bandgaps in the meV range. A range of these materials is shown in Figure I.16.

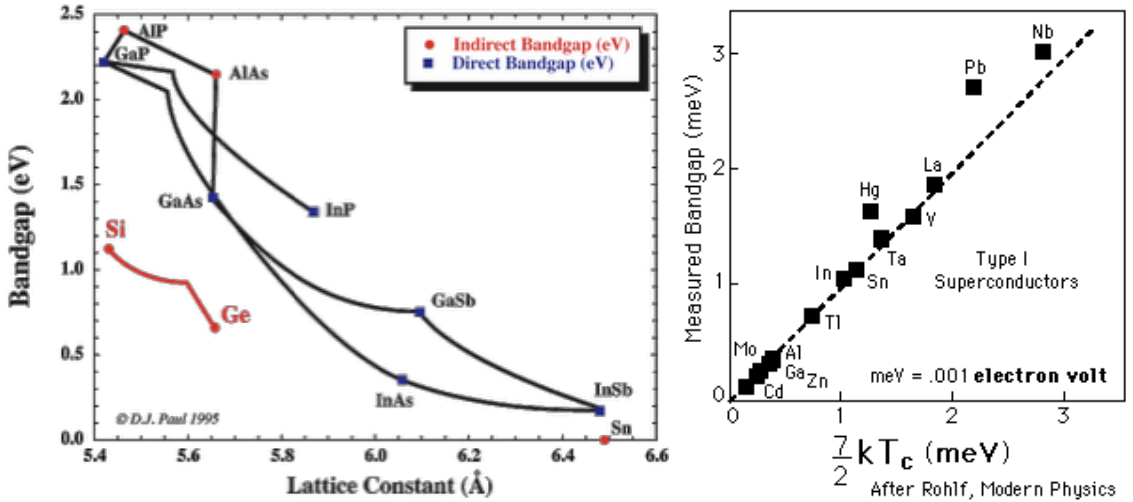


Figure 1.16: Semiconductor (left) and superconductor (right) bandgaps for a variety of materials in the meV-eV range. It is clear that for semiconductors, the bandgap (depending on the nature of the transition) is dependent primarily on the lattice spacing, while the superconducting bandgap is determined very well by the critical temperature.

In this thesis I focus on the development of high-sensitivity readout for Si/Ge detectors, which employ superconducting sensors to achieve sensitivity to excitations at the energy-scale of their bandgaps. In particular, the detectors discussed in this thesis are designed to achieve sensitivity to single electron-hole pairs in Si and Ge, and thus should be sensitive to energy deposits above the bandgap in each material. The sensitivity of these detectors to DM through nuclear recoils and

inelastic scattering (for the exposure and background expected for SuperCDMS SNOLAB [4]) are shown in Figure 1.17

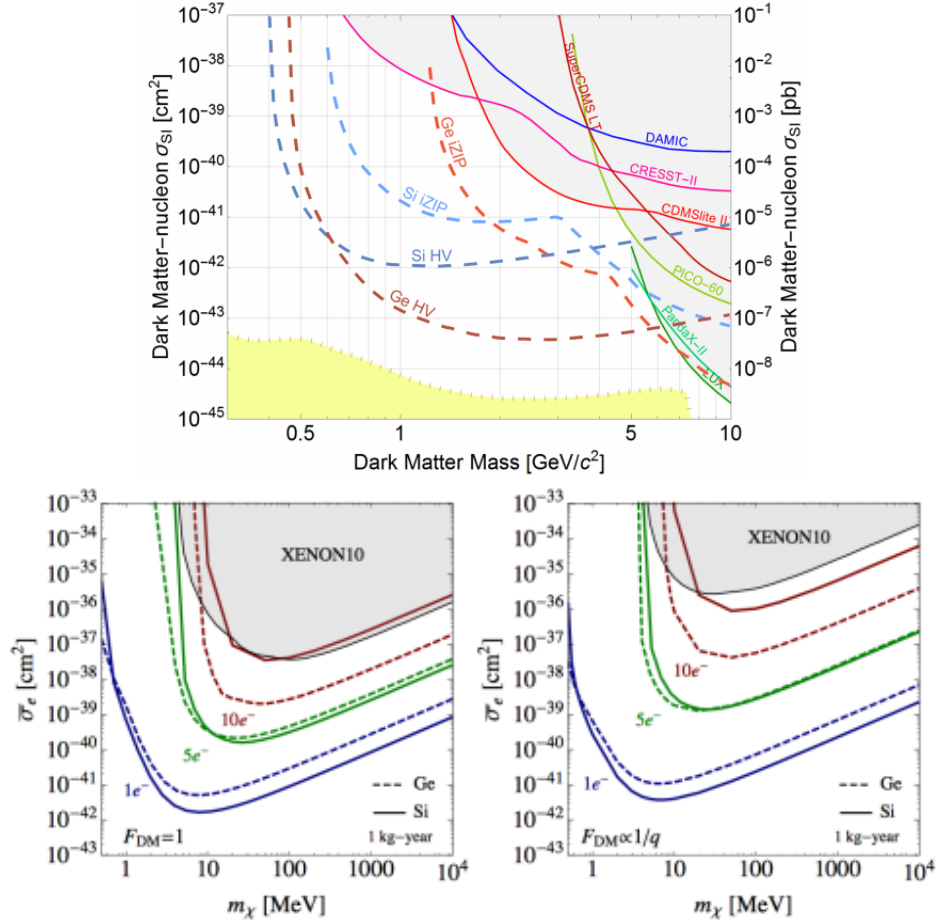


Figure 1.17: Top: Projected sensitivity curves for SuperCDMS SNOLAB detectors through the nuclear recoil channel [4]. Bottom: Projected sensitivity of single-charge sensitive Si/Ge detectors for two different form factors for an exposure without background similar to what is expected for SuperCDMS SNOLAB [33]. Similar limits are also shown in Figure 1.13

The structure of the following chapters is as follows. In chapter 2 I will discuss charge and phonon dynamics in the crystal relevant to the phonon-mediated gain operating of these detectors. In chapters 3 I will present the response and noise model for transition edge sensors in the context of our prototype devices, and in chapter 4 I will discuss the design of the SuperCDMS SNOLAB HV detector readout, including the optimization of the phonon sensors and other design constraints related to readout electronics. In chapter 5 I will present results from various prototype detectors, both full-size SNOLAB designs and scaled down designs meant to probe various potential failure modes. I will then present in chapter 6 the first electron-scattering results produced by SuperCDMS

with a gram-scale prototype detector constructed according to the design principles described in chapter [3](#). I will conclude with a chapter discussing future prospects and R&D for pushing beyond the limitations of the generation of detectors discussed in this thesis, as well as potential imaging applications for the current and near-future versions of this technology.

Chapter 2

High Voltage Charge & Phonon Dynamics

“Quantum mechanics is certainly imposing. But an inner voice tells me that it is not yet the real thing. The theory says a lot, but does not really bring us any closer to the secret of the “old one.” I, at any rate, am convinced that He does not throw dice.”

- Albert Einstein, Letter to Max Born (4 December 1926)

My foray into charge transport simulations stemmed from the joint need to improve the detector Monte Carlo being developed for SuperCDMS SNOLAB and to analyze the data from the first silicon charge transport experiment conducted at Stanford. The results described here were published in Ref [74]. This section summarizes the aspects of charge transport relevant to those efforts, and is by no means a complete description or derivation of the laws contained here. An excellent introduction to these topics in a CDMS context can be found in Kyle Sundqvist’s thesis [99], and a description of the experiment from which the data in this section are taken can be found in Robert Moffatt’s thesis [75]. For further discussions of impact ionization, see Arran Phipps’ thesis [81]. An excellent review of Monte Carlo methods which derives and discusses many of the formulae presented here can be found in Ref [52]. For specifics on the CDMS detector Monte Carlo, I refer the reader to Rob Agnese’s thesis [3] and references therein. For complete pedagogical reviews I refer the reader to Ashcroft and Mermin [10] for an introduction and Ridley [85] for an advanced treatment of semiconductor physics.

2.1 Introduction to Band Structure

Band structure is a phenomenon that arises naturally from periodic potentials in quantum mechanics. In its full glory, the 3D band structure of a material can be computed from overlap integrals of the bound electronic orbitals of the free atoms which comprise the material, but in many cases the form is remarkably similar to a 1D approximation of periodic potentials. An even simpler example comes to us by way of just applying periodic boundary conditions to an otherwise unbound state. In this section I will review this basic example, and then present the results of periodic finite square well potentials in an attempt to give the reader some context for why we can treat a complex lattice structure as modified free-space. The rest of this section will be much less pedagogical but I feel this is an important example which makes the remaining infrastructure of solid state physics much less mysterious.

Let's consider the Dirac comb, a series of delta function potentials of height α spaced a distance a apart. Starting with the 1D time-independent Schrodinger equation

$$E\psi = \frac{\hbar^2}{2m_e} \frac{d^2\psi(x)}{dx^2} + V(x) \quad (2.1)$$

we solve for the case of a periodic potential

$$V(x) = \alpha \sum_{n=-\infty}^{\infty} \delta(x - na) \quad (2.2)$$

For periodic potentials, we know from Bloch's theorem that the solutions will be of the form

$$\psi(x + a) = e^{iKa} \psi(x) \quad (2.3)$$

where a is the spacing between sites and K is the crystal momentum, distinct from k , the wave vector. In later sections k will be the crystal momentum, the lattice wave-vector, but we make this distinction here to be clearer. First, we know the solution to the Schrodinger equation in free-space:

$$\psi(x) = A \cos(kx) + B \sin(kx) \quad (2.4)$$

where $k = \sqrt{2mE}/\hbar$. We now apply our boundary conditions that 1) Bloch's theorem holds and 2) the wave-function is continuous across lattice sites.

The first condition gives

$$\psi(a) = e^{iKa} \psi(0) \quad (2.5)$$

$$A \cos(ka) + B \sin(ka) = e^{iKa} A \quad (2.6)$$

$$e^{-iKa} [A \cos(ka) + B \sin(ka)] = A \quad (2.7)$$

The second condition we find by integrating the Schrodinger equation around a lattice site. We find

$$\frac{\hbar^2}{2m_e} \int_{a-\epsilon}^{a+\epsilon} \frac{d^2\psi(x)}{dx^2} dx = \int_{a-\epsilon}^{a+\epsilon} (E\psi(x) - V\psi(x)) dx \quad (2.8)$$

$$\frac{\hbar^2}{2m_e} \left[\frac{d\psi(x)}{dx} \Big|_{a+\epsilon} - \frac{d\psi(x)}{dx} \Big|_{a-\epsilon} \right] = E \int_{a-\epsilon}^{a+\epsilon} \psi(x) dx - \alpha\psi(a) \quad (2.9)$$

for $\epsilon \ll a$. Taking ϵ to zero gives the boundary condition

$$\frac{d\psi(x)}{dx} \Big|_{a+} - \frac{d\psi(x)}{dx} \Big|_{a-} = -\frac{2m\alpha}{\hbar^2} \psi(a) \quad (2.10)$$

For our wave equation, this evaluates to the condition

$$kB - e^{-iKa} [B \cos(ka) - A \sin(ka)] = -\frac{2m\alpha}{\hbar^2} A \quad (2.11)$$

so that equations [2.7](#) and [2.11](#) allow us to solve for the unknown coefficients A and B.

Solving these equations yields the transcendental equation

$$\cos(Ka) = \cos(ka) - \frac{\alpha m}{\hbar^2 k} \sin(ka) \quad (2.12)$$

which is our primary result. We find that the crystal momentum K is not a monotonic function of k unless $\alpha \rightarrow 0$. In this case, in the small k and K limit we have

$$1 - (Ka)^2 \approx 1 - (ka)^2 - \frac{\alpha m}{\hbar^2} (1 - (ka)^2/6) \quad (2.13)$$

$$(Ka)^2 \approx \left(1 - \frac{\alpha m}{6\hbar^2}\right) (ka)^2 + \frac{\alpha m}{\hbar^2} \quad (2.14)$$

$$\hbar^2 K^2 \approx \left(1 - \frac{\alpha m}{6\hbar^2}\right) (2mE) + \frac{\alpha m}{a^2} \quad (2.15)$$

$$E \approx \frac{\hbar^2 K^2}{2m^*} - \frac{\alpha}{2a^2} \frac{m}{m^*} = \frac{\hbar^2 K^2}{2m^*} + E_0 \quad (2.16)$$

where m^* is the effective mass of the particle

$$m^* = \left(1 - \frac{\alpha m}{6\hbar^2}\right) m \quad (2.17)$$

and E_0 is the zero crystal momentum energy (the work function). This is our secondary result; the periodic potential scales the free-space momentum to the crystal momentum, which can be thought of as changing the particle's mass! This is a nice illustration of where the concepts of effective mass and crystal momentum come from. It's also a very clean illustration of how dispersion relations in crystals can be assumed quadratic in the small K limit. Notice that effective mass can be negative for potential barriers ($\alpha > 0$), but not for potential wells ($\alpha < 0$) and reduced to the free-space

dispersion relation in the limit $\alpha \rightarrow 0$.

We finish this derivation by applying periodic boundary conditions to the whole system, assuming it is a unit of N cells which wrap around on each other such that

$$\psi(x + Na) = \psi(x) \quad (2.18)$$

which by Bloch's theorem means that

$$\psi(x + Na) = e^{iKNa}\psi(x) = \psi(x) \rightarrow K = \frac{2\pi n}{Na} \quad (2.19)$$

giving us the final equation

$$\cos\left(\frac{2\pi n}{N}\right) = \cos(ka) - \frac{\alpha m}{\hbar^2 k} \sin(ka) \quad (2.20)$$

In the limit of infinite lattice sites (infinite free-space) we thus have continuous ranges of allowed values for $ka = \sqrt{2mE}a/\hbar$, or energy bands. Our third result is thus that periodic potentials give rise to band structure!

Consider now a slightly more complicated model, where instead of delta functions we have finite square wells of depth α and width b spaced distance a apart. This commonly referred to as the Kronig Penney model, and is used as a toy model for one-dimensional lattices. When we follow this same procedure, we have free-electron wavefunctions with wave-vectors k_1 and k_2 and we arrive at the transcendental equation [58](#)

$$\cos(ka) = \cos\left(\frac{2\pi n}{N}\right) = \cos(k_2b) \cos(k_1(a-b)) - \frac{k_1^2 + k_2^2}{2k_1k_2} \sin(k_2b) \sin(k_1(a-b)) \quad (2.21)$$

This equation makes it much clearer that we're producing a new momentum state k which is some combination of k_1 and k_2 , and numerically solving this transcendental equation gives us the dispersion relation $E(k)$ given that k_1 and k_2 are both a function of energy, and the rest of the values depend on constants of the problem. One can imagine the 3-dimensional problem contains of these equations, approximately orthogonal, allowing for asymmetric dispersion relations, and therefore asymmetric effective mass. This is a very important aspect of charge propagation in semiconductors, and while we cannot analytically express the full band structure, it provides nice intuition for the concepts of dispersion relations and effective mass.

2.2 Electrons: Conduction Band Structure

The full band-structure of semiconductor lattices is obtained by numerical integration of the interference of free-atom electronic orbitals, and has a complex dispersion relation, as shown in [Figure 2.1](#).




Figure 2.1: Band structure computed numerically as shown in Ref [33] for Si and Ge.

The k -vector axis rotates through the set of valleys in the Brillouin zone where minima can exist (by symmetry arguments, minima can only exist along a handful of unit vectors in k -space) [85]. In diamond cubic lattices (where each lattice site has four covalent bonds to adjacent sites), the nomenclature for face-centered cubic bravais lattices is used. The k -space minima are found along the following directions:

- $\Gamma - k = 0$, the center of the Brillouin zone. All valence bands are centered at Γ , and direct-gap semiconductors are as well, such as GaAs.
- $X - \hat{k}$ is along the coordinate axes, meaning there are 6 valleys per lattice site. This is the lowest valley in Si.
- $L - \hat{k}$ is equally shared along the coordinate axes; in the first Brillouin zone, $\hat{k} = \frac{1}{\sqrt{3}}\langle 1, 1, 1 \rangle$, meaning there are 8 valleys per lattice site. This is the lowest valley in Ge.
- $K - \hat{k}$ is equally shared between two coordinate axes, for example $\hat{k} = \frac{1}{\sqrt{2}}\langle 1, 1, 0 \rangle$, meaning there are 12 valleys per lattice site.
- $W - \hat{k}$ is at 30° from each of the coordinate axes in each coordinate plane, for example $\hat{k} = \frac{1}{2}\langle 1, \sqrt{3}, 0 \rangle$, meaning there are 20 valleys per lattice site.

A schematic of these valleys, along with the full k -space unit cell, can be seen in Figure 2.1.

The energy scale of the full-band structure shown in Figure 2.1 is much larger than the typical charge carrier energies even at room temperature, so an effective band structure is an excellent approximation to this full numerical model in most cases. For Si and Ge, the dispersion relations




Figure 2.2: Brillouin zone of face centered cubic lattice, showing the path through the Brillouin zone taken to produce the band structure figures shown in this section.

Substrate	ϵ_{gap} (eV)	m_{\perp} (m_e)	m_{\parallel} (m_e)	α eV $^{-1}$
Ge	0.67	0.08	1.59	0.7
Si	1.107	0.19	0.98	0.5
Diamond	5.49	0.36	1.4	

Table 2.1: Dispersion relation parameters reproduced from Ref [52] for Si, Ge and diamond. At the time of the cited review, the band structure of diamond was much less well measured than for the other two substrates, hence the absence of a good measure of non-parabolicity.

are very well approximated by the form

$$E - \epsilon_{gap} = \frac{k_x^2 + k_y^2}{m_{\perp}} + \frac{k_z^2}{m_{\parallel}} \quad (2.22)$$

where m_{\perp} is the mass transverse to the k-vector minimum, m_{\parallel} is the mass along the k-vector minimum, and k_i is with respect to the k-vector minima for a given crystal. The simplified band structure shown in Figure 2.3 shows that Si has a minimum along the X-valleys and Ge has a minimum along the L-valleys. We can also see from Figure 2.3 that the electron bands are well approximated by a quadratic form for $E \ll 100$ meV [52]. The parameters for Si, Ge, and diamond are shown in Table 2.1

One modification to this simple form to maintain accuracy to higher energy ranges is to introduce

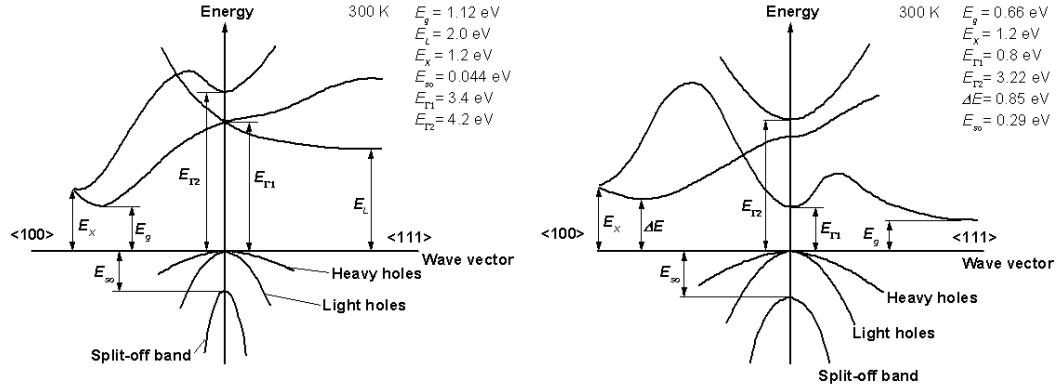


Figure 2.3: Simplified models of the band structure in Si (left) and Ge (right) with indirect and direct transitions labeled.

a non-parabolicity parameter α [52], such that

$$E \rightarrow E(1 + \alpha E) \quad (2.23)$$

where the product αE is dimensionless. Values for α for Si and Ge are also found in Table 2.1

The practical implications of this valley structure in semiconductors is that when charges have a large mean-free path, there are highly preferred regions of momentum space, and thus propagation is largely anisotropic. Measuring this anisotropy as a function of applied voltage in various substrates was the primary goal of the studies shown later in this chapter. The degree of anisotropy is thus a good measure of the mean-free path in a high-purity crystal, and therefore a sensitive probe of electron-phonon interactions.

2.2.1 Effective Mass

One odd side-effect of the anisotropic propagation of electrons in a crystal is that charge carriers no longer have a well-defined mass, and our definition of mass will depend on the context in which it's used. A very thorough discussion of this is found in Kyle Sundqvist's thesis [99], and I summarize the conclusions here and refer the reader there for the more thorough treatment.

There are many ways of pedagogically deriving the effective mass formula, but the simplest approach is just to try to write the dispersion relation as a tensor product, recognizing that in free-space the dispersion relation is just

$$E(k) = \frac{\hbar^2 \mathbf{k}^2}{2m} = \frac{\hbar^2}{2} \mathbf{k}^\dagger \mathbf{m}^{-1} \mathbf{k} \quad (2.24)$$

where $\mathbf{m} = m\mathbf{I}$ and \mathbf{I} is just the identity matrix. To get the effective mass tensor, then, we can take

derivatives of the dispersion relation with respect to k_i , giving the form [99](#)

$$[m^{-1}]_{ij} = \frac{1}{\hbar^2} \frac{d^2 E(k)}{dk_i dk_j} \quad (2.25)$$

If we apply this to equation [2.22](#) we get the effective mass tensor

$$\mathbf{m}^{-1} = \begin{bmatrix} m_{\perp}^{-1} & 0 & 0 \\ 0 & m_{\perp}^{-1} & 0 \\ 0 & 0 & m_{\parallel}^{-1} \end{bmatrix} \quad (2.26)$$

for a valley along the z-axis.

A practical upshot of the effective mass tensor not being a simple multiple of the identity matrix is that the kinematics in a crystal are not the same as in free-space. For example, the velocity of an electron is given by the equation [99](#)

$$\mathbf{v} = \hbar \mathbf{m}^{-1} \mathbf{k} \quad (2.27)$$

and thus the velocity is not parallel to the crystal momentum. That is not necessarily surprising, but consider acceleration by a force:

$$\mathbf{a} = \mathbf{m}^{-1} \mathbf{F} \quad (2.28)$$

due to this mass asymmetry, the acceleration vector is not strictly parallel to the force vector! We should note that the momentum change is still proportional to the electric field:

$$\delta \mathbf{p} = \mathbf{m} \frac{d\mathbf{v}}{dt} = q \mathbf{E} \quad (2.29)$$

This is primarily what produces asymmetric charge propagation. Thus to model charge dynamics in a crystal, we can find the derivatives of the band structure about the minimum with respect to $k_i k_j$, and use these to determine the effective mass, which in turn determines the charge dynamics in the crystal.

This effective mass is obviously useful for equations which can be written as vector products or vector sums, but how do we compute a particle's scalar effective mass? The two most commonly used scalar effective masses are the conductivity effective mass

$$\frac{1}{m_c} = \frac{1}{3} \left[\frac{2}{m_{\perp}} + \frac{1}{m_{\parallel}} \right] \quad (2.30)$$

and the density of states effective mass

$$m_d = [m_{\parallel} m_{\perp}^2]^{1/3} \quad (2.31)$$

These effective masses have different use cases; the conductivity effective mass is useful for mean-free

path calculations, as it is a harmonic average, while the density of states effective mass is useful for calculations involving volumes of phase space. In this thesis I will only use the tensor effective mass or the density of states effective mass, as the conductivity effective mass is used primarily in mobility calculations, which I do not cover here.

2.2.2 Herring-Vogt Transform

A final result from the formulation of the effective mass tensor in the previous section has to do with a trick used to greatly simplify scattering Monte Carlo simulations, the Herring-Vogt transform, which is employed in the CDMS detector Monte Carlo. It is described in full detail in Ref [52] and references therein, and I summarize it here because of its central nature in our simulations as well as the impact it has on using our simulations to understand our data.

The non-spherical nature of the dispersion relations in Si and Ge makes integrals over k-space much more complicated, as the energy for a given k-vector magnitude is angle-dependent. In order to use isotropic scattering approximations, we'd like to use an isotropic effective mass, which means we'd like to solve for the transformation

$$m_e \mathbf{m}^{-1} = \mathbf{T}^\dagger \mathbf{T} \quad (2.32)$$

where m_e is the mass of the free electron. Solving this equation for \mathbf{T} gives the Herring-Vogt transform [52]

$$\mathbf{T} = \begin{bmatrix} \sqrt{\frac{m_e}{m_\perp}} & 0 & 0 \\ 0 & \sqrt{\frac{m_e}{m_\perp}} & 0 \\ 0 & 0 & \sqrt{\frac{m_e}{m_\parallel}} \end{bmatrix} \quad (2.33)$$

This allows us to rewrite the dispersion relation:

$$E(k) = \frac{\hbar^2}{2} \mathbf{k}^\dagger \mathbf{m}^{-1} \mathbf{k} \quad (2.34)$$

$$= \frac{\hbar^2}{2m_e} \mathbf{k}^\dagger \mathbf{T}^\dagger \mathbf{T} \mathbf{k} \quad (2.35)$$

$$= \frac{\hbar^2}{2m_e} (k')^2 \quad (2.36)$$

where we have the transformed k-vector $\mathbf{k}' = \mathbf{T} \mathbf{k}$. This allows us to re-write integrals over k-space by substituting $\mathbf{k} \rightarrow \mathbf{k}'$ and $d\mathbf{k} \rightarrow \left(\frac{m_d}{m_e}\right) d\mathbf{k}'$, where m_d is the density of states effective mass defined in the previous section.

The downside of employing this transform is that it is not unitary, and therefore it does not obey simultaneous energy and momentum conservation. This is not a deal-breaker for us, as there is always some numerical error in Monte Carlo simulations, but it does require periodic corrections

to maintain these conservations separately. The net result of this transform is to sacrifice some simulation accuracy for speed. In the case of CDMS simulations, this transformation has been shown to produce broadening of electron valleys compared to the full vector-based treatment of charge transport [73].

2.3 Holes: Valence Band Structure

Kane [55] derives the perturbed theory of valence band structure in p-type semiconductors (note that the doping does not seem to affect the band structure calculations). In the small k limit, it will turn out to have the parabolic, warped form

$$\epsilon(k) = Ak^2 \pm [B^2k^4 + C^2(k_x^2k_y^2 + k_x^2k_z^2 + k_y^2k_z^2)]^{1/2} \quad (2.37)$$

which can also be converted to a polar form using the relations

$$\frac{k_x}{k} = \cos(\theta) \sin(\phi) \quad (2.38)$$

$$\frac{k_y}{k} = \sin(\theta) \sin(\phi) \quad (2.39)$$

$$\frac{k_z}{k} = \cos(\phi) \quad (2.40)$$

giving

$$\frac{\epsilon(k)}{k^2} = A \pm [B^2 + C^2 \sin^2(\phi) (\cos^2(\theta) \sin^2(\theta) \sin^2(\phi) + \cos^2(\phi))]^{1/2} \quad (2.41)$$

where the effective masses of some key directions are

$$m_{100} = m_{010} = m_{001} = \frac{A \pm B}{\hbar^2}, \quad m_{111} = \frac{A \pm [B^2 + \frac{4}{9}C^2]^{1/2}}{\hbar^2} \quad (2.42)$$

which gives the mass anisotropy

$$\eta = \frac{m_{111}}{m_{100}} - 1 = \frac{A \pm [B^2 + \frac{4}{9}C^2]^{1/2}}{A \pm B} - 1 \quad (2.43)$$

$$= \pm \left(\frac{[B^2 + \frac{4}{9}C^2]^{1/2} - B}{A + B} \right) \quad (2.44)$$

$$= \pm \left(\frac{[b^2 + \frac{4}{9}c^2]^{1/2} - b}{1 + b} \right) \quad (2.45)$$

where $b = \frac{B}{A}$ and $c = \frac{C}{A}$. As we could have deduced from the earlier form, the parameter that controls the degree of anisotropy is the c parameter, however the sensitivity of the anisotropy to c is set by b . This is the anisotropy measure is explored in Ref [78].

Kane's more complete solution is to solve for the band structure by perturbation of the $k \cdot p$ Hamiltonian about $k = 0$. We know that solutions to the Hamiltonian

$$\left(\frac{p^2}{2m} + V(r)\right) \psi(k, r) = E_k \psi(k, r) \quad (2.46)$$

have the form

$$\psi(k, r) = e^{ik \cdot r} u_k(r) \quad (2.47)$$

where $u_k(r)$ is cell periodic in the same manner as $V(r)$. We can re-write the Hamiltonian equation without the Bloch exponential as

$$\left(\frac{p^2}{2m} + V(r) + \frac{\hbar k \cdot p}{m} + \frac{\hbar^2 k^2}{2m}\right) u_k(r) = E_k u_k(r) \quad (2.48)$$

The insight of the $k \cdot p$ perturbation theory [26] is that we can find the $k = 0$ exact solutions easily, and then find solutions for arbitrary k by finding the perturbations in $k \cdot p$ to high enough order. We thus re-write this equation as

$$\left(H_0 + \left(\frac{\hbar}{m}\right) k \cdot p\right) u_k(r) = E'_k u_k(r) \quad (2.49)$$

where $E'_k = E_k - \frac{\hbar^2 k^2}{2m}$ and $H_0 = \frac{p^2}{2m} + V(r)$. Kane [55] reduces the perturbation Hamiltonian equation to the form

$$H_{kp} c = E'_k c \quad (2.50)$$

where

$$H_{kp} = \begin{bmatrix} Lk_x^2 + M(k_y^2 + k_z^2) & Nk_x k_y & Nk_x k_z \\ Nk_x k_y & Lk_y^2 + M(k_x^2 + k_z^2) & Nk_y k_z \\ Nk_x k_z & Nk_y k_z & Lk_z^2 + M(k_x^2 + k_y^2) \end{bmatrix} \quad (2.51)$$

Including the spin-orbit coupling, the second order energy of these states has the eigenvalue equation

$$\begin{aligned} & H'_{11} H'_{22} H'_{33} + 2H_{12} H_{23} H_{13} - H'_{11} H'^2_{23} - H'_{22} H'^2_{13} - H'_{33} H'^2_{12} \\ & - (\Delta/3) (H'_{11} H'_{22} + H'_{11} H'_{33} + H'_{22} H'_{33} - H'^2_{12} - H'^2_{13} - H'^2_{23}) = 0 \end{aligned}$$

where

$$H'_{ii} = H_{ii} + (\hbar^2/2m)k^2 - E_k \quad (2.52)$$

This equation is just the eigenvalue equation of the perturbed Hamiltonian with a spin-orbit correction. If we set $k = 0$, this reduces the the equation

$$E_k^3 + \Delta E_k^2 = (E_k + \Delta)E_k^2 = 0 \quad (2.53)$$

which has two degenerate solutions at $E_k = 0$ and one split by $E = -\Delta$, as expected.

The original matrix had L, M, and N as coefficients of overlap integrals, and Kane cites the DKK paper [32] for this parameterization, as they attempt to measure these parameters. We can see that the energy eigenstates are functions of these parameters, but a re-factoring can be done to simplify the equations. We can re-write H_{kp} to get

$$H_{kp} = \begin{bmatrix} L'k_x^2 + M'(k_y^2 + k_z^2) & Nk_xk_y & Nk_xk_z \\ Nk_xk_y & L'k_y^2 + M'(k_x^2 + k_z^2) & Nk_yk_z \\ Nk_xk_z & Nk_yk_z & L'k_z^2 + M'(k_x^2 + k_y^2) \end{bmatrix} \quad (2.54)$$

where $L' = L + \frac{\hbar^2}{2m}$ and $M' = M + \frac{\hbar^2}{2m}$.

The general form for nonzero k is somewhat more complicated than the $k = 0$ solution, but setting two of the k components to 0 gives the energies

$$E_{100}(k) = M'k^2 \quad (2.55)$$

$$= \frac{1}{2} \left((L' + M')k^2 - \Delta + \sqrt{k^4(L' - M')^2 + \frac{2}{3}k^2(L' - M')\Delta + \Delta^2} \right) \quad (2.56)$$

$$= \frac{1}{2} \left((L' + M')k^2 - \Delta - \sqrt{k^4(L' - M')^2 + \frac{2}{3}k^2(L' - M')\Delta + \Delta^2} \right) \quad (2.57)$$

while for k vectors along the [111] directions gives the energies

$$E_{111}(k) = \frac{(L' + 2M' - N)}{3}k^2 \quad (2.58)$$

$$= \frac{1}{2} \left(\frac{(2L' + 4M' + N)}{3}k^2 - \Delta + \sqrt{k^4N^2 + \frac{2}{3}k^2N\Delta + \Delta^2} \right) \quad (2.59)$$

$$= \frac{1}{2} \left(\frac{(2L' + 4M' + N)}{3}k^2 - \Delta - \sqrt{k^4N^2 + \frac{2}{3}k^2N\Delta + \Delta^2} \right) \quad (2.60)$$

Comparing these solutions tells us that, in general, none of the bands are spherical, but that one of the bands is parabolic. The other two bands are split from a central parabolic band at $k = 0$ by Δ and at higher k by all four parameters. The [110] energies do not have simple closed forms, so it's possible that all three bands are somewhat non-parabolic even in the small k limit.

2.3.1 Equal Energy Surfaces

In order to plot the equal energy contours for this band structure, we need to first rotate the 001 k basis that the structure is specified in into the desired crystal orientation. We use the rotation matrix

$$R = \begin{bmatrix} -\frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{3}} \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{6}} & \frac{1}{\sqrt{3}} \\ 0 & \sqrt{\frac{2}{3}} & \frac{1}{\sqrt{3}} \end{bmatrix} \quad (2.61)$$

to rotate vectors in the 111 basis to the 001 basis, such that we can insert them into the energy surface equation in manner which is meaningful for our 111 crystal orientation. We will use this implicitly in the following logic.

We want to relate the theoretical band structure in some analytic way to the patterns we measure on the surface of our crystals. The key issue is that we cannot necessarily just relate k vectors to this pattern, as the whole idea of a crystal momentum is that it does not correlate necessarily to spatial dimensions; this is the idea behind the effective mass. We can relate the spatial and crystal spaces through the equation

$$v(k) = \hbar^{-1} \nabla_k E(k) \quad (2.62)$$

where we assume we can reconstruct the angular dependence of $v(k)$ by assuming $v(k)$ is the average carrier velocity (carriers propagate at constant energy in a constant electric field), and thus

$$x_i \approx v_i(k) \Delta t \quad (2.63)$$

If we had timing information this might be the extent of our logic, however we do not measure the arrival time of the charges, so we need to eliminate Δt from this relation. Suppose instead of the position we fit the fractional components of the position (the position unit vector):

$$\frac{x_i}{|x|} \approx \frac{v_i(k) \Delta t}{|v(k) \Delta t|} = \frac{v_i(k)}{|v(k)|} \quad (2.64)$$

So we have a way to roughly relate position and velocity, which also allows us to relate position and the energy surface equation, as long as we forfeit knowledge of the k -vector magnitude. So be it. There are obviously some assumptions in here about the degree to which scattering affects spatial distribution of charge carriers, and for a single charge we shouldn't expect it to following in a straight line. A better way to visualize this argument is to realize that the effective mass tensor just means that, in some other space, the charges propagate isotropically, and we then rescale the spatial dimensions according the mass tensor to get the true spatial distribution; the degree of warping between the dimensions is independent of scattering and energy aside from the energy scale helping to determine the available scattering phase space.

So lets normalize this equation by the vector magnitudes:

$$\frac{x}{|x|} \approx \frac{v(k)}{|v(k)|} = \frac{\nabla_k E(k)}{|\nabla_k E(k)|} \quad (2.65)$$

we see that constants will start going away, reflecting the fact that we can't constrain overall magnitudes in the band structure without velocity information, at least not with our position fits. Lets now calculate the gradient of the energy surface in 001 to see how that turns out before launching down the road for the 111 case. We have

$$\frac{d}{dk_i} E(k) = \frac{d}{dk_i} \left[Ak^2 + (B^2 k^4 + C^2 (k_x^2 k_y^2 + k_x^2 k_z^2 + k_y^2 k_z^2))^{1/2} \right] \quad (2.66)$$

$$= 2Ak_i + \frac{1}{2} (B^2 k^4 + C^2 (k_x^2 k_y^2 + k_x^2 k_z^2 + k_y^2 k_z^2))^{-1/2} \frac{d}{dk_i} (B^2 k^4 + C^2 (k_x^2 k_y^2 + k_x^2 k_z^2 + k_y^2 k_z^2)) \quad (2.67)$$

$$= 2Ak_i + \frac{(2B^2 + C^2)k_i k^2 - C^2 k_i^3}{(B^2 k^4 + C^2 (k_x^2 k_y^2 + k_x^2 k_z^2 + k_y^2 k_z^2))^{1/2}} \quad (2.68)$$

$$= 2Ak_i + \frac{(2B^2 + C^2)k_i - C^2 k_i^3/k^2}{(B^2 + C^2 (k_x^2 k_y^2 + k_x^2 k_z^2 + k_y^2 k_z^2)/k^4)^{1/2}} \quad (2.69)$$

$$= 2k_i \left[A + \frac{(B^2 + \frac{1}{2}C^2) - \frac{1}{2}C^2 k_i^2/k^2}{(B^2 + C^2 (k_x^2 k_y^2 + k_x^2 k_z^2 + k_y^2 k_z^2)/k^4)^{1/2}} \right] \quad (2.70)$$

$$= 2k_i \left[A + \frac{(B^2 + \frac{1}{2}C^2) - \frac{1}{2}C^2 k_i^2/k^2}{D^{1/2}} \right] \quad (2.71)$$

so the vector has the magnitude

$$|\nabla_k E(k)| = 2 \left(\sum_i k_i^2 \left[A + \frac{B^2 + \frac{1}{2}C^2 - \frac{1}{2}C^2 k_i^2/k^2}{D^{1/2}} \right]^2 \right)^{1/2} \quad (2.72)$$

$$= 2 \left(\sum_i \frac{k_i^2}{D} \left[AD^{1/2} + B^2 + \frac{1}{2}C^2 - \frac{1}{2}C^2 k_i^2/k^2 \right]^2 \right)^{1/2} \quad (2.73)$$

$$= 2 \left(\sum_i \frac{k_i^2}{D} \left[E - \frac{1}{2}C^2 k_i^2/k^2 \right]^2 \right)^{1/2} \quad (2.74)$$

$$= 2 \left(k^2 \frac{E^2}{D} - (k_x^4 + k_y^4 + k_z^4) \frac{EC^2}{k^2 D} + (k_x^6 + k_y^6 + k_z^6) \frac{C^4}{4k^4 D} \right)^{1/2} \quad (2.75)$$

$$= \frac{2k}{D^{1/2}} \left(E^2 - (k_x^4 + k_y^4 + k_z^4) \frac{EC^2}{k^4} + (k_x^6 + k_y^6 + k_z^6) \frac{C^4}{4k^6} \right)^{1/2} \quad (2.76)$$

where

$$D = B^2 + C^2 \left(\frac{k_x^2 k_y^2 + k_x^2 k_z^2 + k_y^2 k_z^2}{k^4} \right) \quad (2.77)$$

$$E = AD^{1/2} + B^2 + \frac{1}{2}C^2 \quad (2.78)$$

So we have (in terms of all of the replacement constants)

$$\frac{(\nabla_k E(k))_i}{|\nabla_k E(k)|} = \frac{k_i}{k} \frac{AD^{1/2} + (B^2 + \frac{1}{2}C^2) - \frac{1}{2}C^2 k_i^2 / k^2}{(E^2 - (k_x^4 + k_y^4 + k_z^4) \frac{EC^2}{k^4} + (k_x^6 + k_y^6 + k_z^6) \frac{C^4}{4k^6})^{1/2}} \quad (2.79)$$

This is no simple expression, and it's still a function of the k-vectors, but now the k-vectors are normalized in all expressions. As a sanity check that we're really just probing anisotropy, let's set C to 0. We find that $D = B^2$, $E = AB + B^2$, and we get

$$\frac{(\nabla_k E(k))_i}{|\nabla_k E(k)|} = \frac{k_i}{k} \frac{AB + B^2}{\sqrt{(AB + B^2)^2}} = \frac{k_i}{k} \quad (2.80)$$

and we just find that the k-vectors should map proportionally to the velocity vectors. These expressions don't naturally simplify beyond this form, but we can see that the anisotropy does not strictly depend on k-vector magnitude, but the degree to which one direction in k-space is preferred. The observed anisotropy is dependent not only on the over-all energy, but also the distribution of energy in the steady state. That being said, we can use the expected mean carrier energy and k-vector to predict the degree of anisotropy seen in the charge collection pattern using this relation.

2.3.2 Hole Effective Mass

For valence bands, effective mass is a confusing concept. Whereas for the electrons our roughly elliptical valleys gave a k-independent effective mass, making it a tensor with a physically intuitive interpretation, for the holes the effective mass tensor is dependent on the direction in which the hole propagates. There are a few ways to define effective mass which each give us a different way of reformulating the problem, each of which aims to make an aspect of the problem more independent of the exact state of the hole.

The traditional definition defines mass as a tensor to convert force applied to acceleration observed. Using the relation

$$a = \frac{d}{dt} v_g = \hbar^{-1} \frac{d}{dt} (\nabla_k \varepsilon(k)) = \hbar^{-1} \nabla_k \frac{d\varepsilon(k)}{dt} = \hbar^{-1} \nabla - k \left(\frac{dk}{dt} \cdot \nabla_k \varepsilon(k) \right) \quad (2.81)$$

and conservation of crystal momentum

$$F = \frac{dp_{crystal}}{dt} = \hbar \frac{dk}{dt} \quad (2.82)$$

gives

$$a = \hbar^{-1} \nabla_k \left(\frac{F}{\hbar} \cdot \nabla_k \varepsilon(k) \right) = [m^*]^{-1} F \quad (2.83)$$

where the mass tensor is traditionally written as

$$[m^*]_{ij}^{-1} = \hbar^{-2} \frac{\partial^2 \varepsilon(k)}{dk_i dk_j} \quad (2.84)$$

We can also use vector identities to distribute the gradient:

$$\nabla_k (F \cdot \nabla_k \varepsilon(k)) = (\nabla_k \varepsilon(k) \cdot \nabla_k) F + (F \cdot \nabla_k) \nabla_k \varepsilon(k) + \nabla_k \varepsilon(k) \times (\nabla_k \times F) + F \times (\nabla_k \times \nabla_k \varepsilon(k)) \quad (2.85)$$

$$= (F \cdot \nabla_k) \nabla_k \varepsilon(k) \quad (2.86)$$

which makes more explicit the assumption we're making above, namely that

$$\frac{dF}{dk_i} = 0 \quad (2.87)$$

in addition to the fact that F is conservative ($\nabla_k \times F = 0$). These are obviously equivalent, but distributing the vector products makes it easier to see how F comes out of the integral, and what we assume in order to do that. So the effective mass is determined by the outer product of a gradient and a divergence of the dispersion relation.

We could also formulate the effective mass using the energy relation, given that we're interested primarily in the anisotropy in the effective mass and care less about acceleration than about deformation of the equal energy surface. Suppose that we try to re-write the dispersion relation as

$$\varepsilon(k) = \frac{\hbar^2}{2} k^T [m^*]^{-1} k \quad (2.88)$$

do we retain the effective mass if we try to apply the previous definition? Let's try:

$$[m^*]_{ij}^{-1} = \hbar^{-2} \frac{\partial^2 \epsilon(k)}{dk_i dk_j} = \frac{1}{2} \frac{\partial^2}{dk_i dk_j} [m^{*'}]_{ij}^{-1} k_i k_j (2 - \delta_{ij}) \quad (2.89)$$

$$= (2 - \delta_{ij}) [m^{*'}]_{ij}^{-1} \frac{1}{2} \frac{\partial^2}{dk_i dk_j} k_i k_j + (2 - \delta_{ij}) k_i k_j \frac{\partial^2}{dk_i dk_j} [m^{*'}]_{ij}^{-1} \quad (2.90)$$

$$= [m^{*'}]_{ij}^{-1} + (2 - \delta_{ij}) k_i k_j \frac{\partial^2}{dk_i dk_j} [m^{*'}]_{ij}^{-1} \quad (2.91)$$

$$[m^*]^{-1} = [m^{*'}]^{-1} + k^T \left[\frac{\partial^2}{dk_i dk_j} [m^{*'}]^{-1} \right] k \quad (2.92)$$

This tells us that the two masses are equivalent if they're k-independent, which for the holes is not the case, so the expression of the hole energy in terms of the effective mass is an approximation good in the limit that this derivative term is small. If we consider that the quadratic order is also the limit to which the band-structure is accurate, it's possible to see how this isn't a huge problem; if we're limited to low energies, the approximate form of the effective mass dependent on k will be a close match to the true energy effective mass.

This form as a tensor product in energy is useful because it allows us to reformulate our Monte-Carlo rules in a much simpler way. If the energy surface is spherical, this means both that we are equally probable to emit a phonon in all directions, as a unit-vector in k-space with a given magnitude has the same energy regardless of its orientation. This allows us to run an isotropic monte-carlo with an additional step to transform into this isotropic space when the propagator is applied, and to transform back to the original space to compute momentum and direction.

We can break the effective mass tensor down into the linear combination of two tensors, one isotropic, and one anisotropic. We have thus that

$$[m^*]^{-1} = \frac{1}{m_e} \left(A\mathcal{I} \pm f(B, C, \hat{k}) \right) \quad (2.93)$$

where the \pm changes sign depending on the band, the minus sign corresponding to the heavy band, and \mathcal{I} being the $N = 3$ identity matrix. The anisotropic term is only dependent on the direction of k , and has both a scalar and tensor anisotropy. We break down this function into the scalar λ and tensor Λ , such that

$$\lambda = \frac{1}{2} \left(B^2 + \frac{C^2}{k^4} (kx^2ky^2 + kx^2kz^2 + ky^2kz^2) \right)^{-3/2} \quad (2.94)$$

which is just the anisotropic term from the energy surface cubed; note that it can be written in

polar form as described earlier. The diagonal elements of the tensor Λ are

$$\Lambda_{ii} = B^2(2+C^2) + \sum_j \sum_k (1 - \delta_{ij}\delta_{ik}\delta_{jk}) \left[C^2(1 + 2B^2) \frac{k_j^2 k_k^2}{k^4} \left(1 - \frac{ki^2}{k^2} \right) - B^2 C^2 \frac{k_i^2}{k^6} (ki^4 + 3(k_j^4 + k_k^4)) \right] \quad (2.95)$$

and the off-diagonal elements are

$$\Lambda_{ij} = C^2 \frac{k_i k_j}{k^2} \sum_k (1 - \delta_{ij}\delta_{ik}) \left[(2B^2 + C^2) \frac{k_k^2}{k^2} \left(\frac{2k_k^2}{k^2} - 1 \right) - (4B^2 + C^2) \frac{k_i^2 k_j^2}{k^4} \right] \quad (2.96)$$

We can further separate this into three main components, an isotropic scalar, and anisotropic scalar, and an anisotropic tensor, by redefining Λ_{ii} as

$$\Lambda'_{ii} = \sum_j \sum_k (1 - \delta_{ij}\delta_{ik}\delta_{jk}) \left[(1 + 2B^2) \frac{k_j^2 k_k^2}{k^4} \left(1 - \frac{ki^2}{k^2} \right) - B^2 \frac{k_i^2}{k^6} (ki^4 + 3(k_j^4 + k_k^4)) \right] \quad (2.97)$$

$$\Lambda'_{ij} = \frac{k_i k_j}{k^2} \sum_k (1 - \delta_{ij}\delta_{ik}) \left[(2B^2 + C^2) \frac{k_k^2}{k^2} \left(\frac{2k_k^2}{k^2} - 1 \right) - (4B^2 + C^2) \frac{k_i^2 k_j^2}{k^4} \right] \quad (2.98)$$

and writing the effective mass as

$$[m^*]^{-1} = \frac{1}{m_e} (A\mathcal{I} \pm \lambda [B^2(2 + C^2)\mathcal{I} + C^2\Lambda']) \quad (2.99)$$

shows us that there are various approximations we can make, but that this is essentially a tensor expansion.

In summary, the hole effective mass is a much more complicated structure than the electron effective mass, and a proper simulation treatment of the holes is beyond the scope of this thesis. The intent of this section was to build intuition for the concept of warped bands, and how their behavior is a much more complicated function of hole momentum and k -vector. Later in this chapter, we will show that in certain limits, the anisotropic hole propagation appears to manifest itself but recede as a function of temperature, which can be understood in the context of a momentum and energy dependent effective mass.

2.4 Electron Scattering Processes

In this section I will present the theoretical scattering rates and cross-sections for the various electron scattering processes as found in the literature, and show some order of magnitude estimates for our chosen operating regime ($T \sim 50$ mK, $E \sim 0.01$ -100 V/cm).

2.4.1 Impurity Scattering

Impurity scattering can be broken up into charged and neutral impurity scattering, due to the different ranges of interactions between these mechanisms as well as the spectral dependence. In addition, despite having some knowledge of the main impurity type, the contributions from charged and neutral impurities may not be proportional to their relative abundance due to the different nature of impurities that are in either state at cryogenic temperatures. Here I describe the scattering rates as a function of energy, integrating over angle for simplicity.

Neutral Impurities

Scattering of hot electrons from neutral impurities is a complicated quantum process which many have attempted to approximate through various simplifying assumptions, though a consensus has been reached that this neutral scattering is the dominant determinant of mobility in very low-temperature crystals at low to moderate field strengths. The most recent theoretical efforts have used a full quantum mechanical treatment, calculating transition probabilities between initial and final free-electron states given their interaction with bound electron orbitals in a neutral impurities to determine the band-warping effect this impurity will have. I will summarize first the simple model of a spherical square well, from which the simplest scattering cross-section is derived, and contrast this with the results of the full quantum model to justify the use of the approximate form for our implementation.

The most successful analytic solution for the energy dependence of the neutral scattering cross-section comes by way of Sclar and Anselm [98, 85], though the formula is attributed to Sclar due to his complete theoretical framework. These authors treat a neutral hydrogen atom as a spherical square well potential, due to the propensity for neutral impurities to create weakly-bound states with electrons at low temperatures and field strengths. The full derivation of the scattering probability is beyond the scope of this work, but can be found in Ref [97], and references therein.

The basic derivation follows the method of calculating the net phase shifts in the initial and final wavefunctions, typically given as

$$\psi_{in} = A j_0(\alpha r), \quad \psi_{out} = B [j_0(kr) \cos(\delta) + n_0(kr) \sin(\delta)] \quad (2.100)$$

where

$$j_0(x) = \frac{\sin(x)}{x}, \quad n_0 = -\frac{\cos(x)}{x}, \quad \alpha = \left[\frac{2m^*(E + V_0)}{\hbar^2} \right]^{1/2} \quad (2.101)$$

V_0 is the depth of the well, and δ is the resulting phase shift from some simple potential. For this form, the cross-section is given by the expression

$$\sigma(\theta) = \frac{\sin^2(\delta)}{k^2} \quad (2.102)$$

We can thus solve for $\sin(\delta)$ by application of the corresponding boundary conditions, which yields the condition at the edge of the square well $r = a$

$$\sin(\delta) = \frac{k \cot(ka) - \alpha \cot(\alpha a)}{(k^2 + \alpha^2 \cot^2(\alpha a))^{1/2} (1 + \cot^2(ka))^{1/2}} \quad (2.103)$$

If we restrict $ka \ll 1$ (low energies) we obtain the approximate form

$$\sin(\delta) \approx \frac{k(1 - \alpha a \cot(\alpha a))}{(k^2 + \alpha^2 \cot^2(\alpha a))^{1/2}} \quad (2.104)$$

If we have, additionally, that $\alpha a \cot(\alpha a) \ll 1$ (the bound state is fairly shallow), we get

$$\sin^2(\delta) \approx \frac{k^2}{k^2 + \alpha^2 \cot^2(\alpha a)} \quad (2.105)$$

giving the cross-section

$$\sigma(\theta) \approx \frac{1}{k^2 + \alpha^2 \cot^2(\alpha a)} \quad (2.106)$$

If the well is much deeper than the bound state we're interested in, we can make the replacement [\[85\]](#)

$$\alpha^2 \cot^2(\alpha a) \approx - \left(\frac{2m_D E_T}{\hbar^2} \right)^{1/2} \quad (2.107)$$

such that

$$\sigma(\theta) \approx \frac{\hbar^2}{2m_D(E + E_T)} \quad (2.108)$$

and thus we have the differential scattering rate

$$\frac{d\Gamma(E)}{d\Omega} = n_I \sigma(E) v(E) = n_I \sqrt{\frac{2E}{m_D}} \frac{\hbar^2}{2m_D(E + E_T)} = \frac{n_I \hbar^2}{\sqrt{2} m_D^{3/2}} \frac{E^{1/2}}{E + E_T} \quad (2.109)$$

Here and below, m_D is the “density-of-states” effective mass of the electron, $(m_L \times m_T^2)^{1/3}$.

The total scattering rate is thus obtained by multiplying this rate by 4π steradians, however we want the intervalley (IV) and intravalley (V) rates separately. We can make the simplifying assumption that the scattered electron will settle into the nearest valley after scattering, such that we can divide the final scattering scale by the number of valleys, and we have

$$\Gamma(E)_{IV} = 4\pi \left[\frac{N_V - 1}{N_V} \right] \frac{n_I \hbar^2}{\sqrt{2} m_D^{3/2}} \frac{E^{1/2}}{E + E_T} = \left[\frac{N_V - 1}{N_V} \right] \frac{2\sqrt{2}\pi n_I \hbar^2}{m_D^{3/2}} \frac{E^{1/2}}{E + E_T} \quad (2.110)$$

$$\Gamma(E)_V = \left[\frac{1}{N_V} \right] \frac{2\sqrt{2}\pi n_I \hbar^2}{m_D^{3/2}} \frac{E^{1/2}}{E + E_T} \quad (2.111)$$

such that, as expected, this effect is stronger when there are more valleys in which to scatter.

The last detail of the modeling is how to determine the threshold energy (in eV). This is assumed to be due to hydrogenic centers, and is given by

$$E_T = 0.75 \frac{m_D}{m} \frac{\epsilon_0}{\epsilon} \quad (2.112)$$

which is $\sim 5 * 10^{-4}$ for Ge and $\sim 2 * 10^{-3}$ for Si [85]. In principle, this energy threshold is a free parameter, and will depend on the type of impurity. This is an important point which will help explain why scattering rates are not just proportional to impurity concentration. In general, the way to account for some number of bound state energies $E_{T,i}$, and bound state centers $n_{I,i}$, is to sum over these scattering rates, accounting for each, giving the total rate

$$\Gamma(E)_{IV} = \left[\frac{N_V - 1}{N_V} \right] \frac{2\sqrt{2}\pi\hbar^2}{m_D^{3/2}} \sum_i n_{I,i} \frac{E^{1/2}}{E + E_{T,i}} \quad (2.113)$$

In this way, different impurities with different bound state energies can produce arbitrarily complex scattering rate trends, and result in different overall rates of inter-valley scattering. The weaker the bound state, however, the weaker the low-energy scattering will be; so measurements of the IV scattering rate, both in terms of energy dependence and overall rate, place both lower and upper bounds on the relevant bound-state energies given the expected impurity concentration.

This is the simplest treatment which produces the correct energy dependence. It is possible to extend this approach in the future by extending scattering to $l > 0$ harmonics, which would predict some angular dependence in the scattering rate, or through a full anisotropic treatment using the non-spherical valleys as is done in Ref [60]. In the cross-sections calculated with the full phase-shift treatment in this reference, we can see that above a threshold value the cross-section resembles that of the Sclar formula, and could easily be produced using a normal Sclar prescription with very low and intermediate energy bound states. There is, however, an anisotropy in incoming and outgoing angle, however the first we could integrate over, and the second represents most likely a 10-20% correction to our predicted rates, less than the typical uncertainty in our impurity concentrations. For these reasons, the simple treatment, tuned using one low-voltage point, should be sufficient for this simulation.

Charged Impurities

We assume that charged impurities contribute negligibly to transport at very low temperatures, given that they have a much larger binding energy than the neutral impurities and thus, by analogy, will be much less likely to become ionized in the thermal limit.

2.4.2 Phonon Scattering

There are two distinct categories of phonons in a crystal, optical and acoustic, where acoustic have a linear and increasing dispersion relationship with wave-vector near $k = 0$, and optical phonons have a finite and often large energy and velocity at $k = 0$. In the ohmic regime, acoustic phonons are the dominant scattering mechanism, which means that this scattering rate should be dominant in the intermediate energies where ohm's law is applicable at our low temperatures, while optical phonons will mediate higher energy transitions (e.g. between valleys). Here I summarize the scattering rates, so that they can be related to intervalley scattering in the next section, and used in mobility calculations.

Acoustic Phonon Scattering

The 'room-temperature' acoustic phonon scattering rate as a function of energy is normally given as [\[12\]](#)

$$\Gamma_{ac}(E) = \frac{\sqrt{2} k_b T m_D^{3/2} D_{ac}^2}{\pi \hbar^4 \rho u^2} \sqrt{E(1 + \alpha E)} (1 + 2\alpha E) \quad (2.114)$$

where these are all strictly intravalley scatters. The average sound speed, $u = (u_L + 2u_T)/3$. This scattering rate will not contribute to intervalley scattering due to the low energy of the phonons emitted, it will however limit the mobility of electrons in intermediate field regimes. This assumes equipartition of energy into emission and absorption and fully elastic acoustic scattering, and thus does not account for the inelastic cooling effect we see from acoustic scattering at low temperature. We need to derive the true low-temperature form, which will also guide the derivation of the optical rate as well.

The differential scattering probability in wave-vector k has the form

$$\Gamma(k, k') = \frac{\pi q \mathcal{E}_{ac}^2}{V \rho u_l} \left[\frac{N_q}{N_q + 1} \right] \delta(\epsilon(k) - \epsilon(k') \mp \hbar q u_l) \quad (2.115)$$

where the top is absorption and the bottom is emission. In our limit, we're going to be only concerned with emission, so we'll limit ourselves to considering the bottom part of the expression and ignore N_q . This gives

$$\Gamma(k, q) = \frac{\pi q \mathcal{E}_{ac}^2}{V \rho u_l} \delta(\epsilon(k) - \epsilon(k') - \hbar q u_l) \quad (2.116)$$

and we get the rate as a function of k -vector by integrating over q and k' (expressing k' as $k - q$):

$$\Gamma(k) = \frac{V}{(2\pi)^3} \int \Gamma(k, q) q^2 d(\cos(\theta)) d\phi dq \quad (2.117)$$

$$= \frac{V}{(2\pi)^3} \int \frac{\pi q \mathcal{E}_{ac}^2}{V \rho u_l} q^2 \delta(\epsilon(k) - \epsilon(k') - \hbar q u_l) d(\cos(\theta)) d\phi dq \quad (2.118)$$

$$= \frac{\mathcal{E}_{ac}^2}{4\pi \rho u_l} \int q^3 \delta(\epsilon(k) - \epsilon(k') - \hbar q u_l) d(\cos(\theta)) dq \quad (2.119)$$

To satisfy the delta function and get rid of the integral over $d(\cos(\theta))$, we get

$$\Gamma(k) = \frac{\mathcal{E}_{ac}^2}{4\pi \rho u_l} \int q^3 \frac{m}{\hbar^2 k q} \delta(\cos(\theta) - \dots) d(\cos(\theta)) dq \quad (2.120)$$

$$= \frac{\mathcal{E}_{ac}^2 m}{4\pi \hbar^2 \rho u_l k} \int_{q_{min}}^{q_{max}} q^2 dq \quad (2.121)$$

$$= \frac{\mathcal{E}_{ac}^2 m}{12\pi \hbar^2 \rho u_l k} (q_{max}^3 - q_{min}^3) \quad (2.122)$$

where the extra terms in the delta function are related to energy conservation. For emission in a parabolic band we have

$$q_{min} = 0, \quad q_{max} = 2k \left(1 - \frac{k_l}{k}\right) \quad (2.123)$$

giving

$$\Gamma(k) = \frac{\mathcal{E}_{ac}^2 m}{12\pi \hbar^2 \rho u_l k} \left(2k \left(1 - \frac{k_l}{k}\right)\right)^3 \quad (2.124)$$

$$= \frac{2\mathcal{E}_{ac}^2 m}{3\pi \hbar^2 \rho u_l} k^2 \left(1 - \frac{k_l}{k}\right)^3 \quad (2.125)$$

$$= \frac{u_l}{3l_0} \frac{k^2}{k_l^2} \left(1 - \frac{k_l}{k}\right)^3 \quad (2.126)$$

$$= \Gamma_0 M^2 (1 - M^{-1})^3 \quad (2.127)$$

where $\Gamma = \frac{u_l}{3l_0}$ is the rate when the remainder of the equation is 1, $M = \frac{k}{k_l}$ is the mach number, and

$$l_0 = \frac{\pi \hbar^4 \rho}{2m^3 \epsilon_{ac}^2}, \quad k_l^2 = \frac{m^2 v_l^2}{\hbar^2} \quad (2.128)$$

This is the same as the rate shown in Ref [25], and is the rate law we use for low-energy NTL phonon production. We expect to find that under this rate law, drift velocity is a very weak power law as a function of electric field.

Optical Phonon Scattering

The optical phonon scattering rate to zeroth order (in which a phonon is emitted) as a function of energy [52, 112] is

$$\Gamma_{op,0}(E) = \frac{Z_f m_D^{3/2} D_{op,0}^2}{\sqrt{2} \pi \hbar^2 \rho \hbar \omega_{op}} \sqrt{E'(1 + \alpha E')(1 + 2\alpha E')} \quad (2.129)$$

where $E' = E - \hbar\omega_{op}$ and $\hbar\omega_{op}$ is the energy of the optical phonon which is emitted by the carrier. Here we restrict to emission due to low temperature being considered, meaning very few phonons will be available for the absorption process, and the N_q that appears in the formula is set to 0. This formula only applies to scattering between equivalent valleys, as we only consider propagation in the lowest valleys, and does not consider the contribution of valleys with higher bandgaps. In Si, we also have to consider the first-order transition rate equation due to selection rules which forbid intervalley transitions for the low-energy phonons. This process takes on the scattering rate

$$\Gamma_{op,1}(E) = \frac{\sqrt{2} Z_f m_D^{5/2} D_{op,1}^2}{\pi \hbar^4 \rho \hbar \omega_{op}} \sqrt{E'(1 + \alpha E')(1 + 2\alpha E')(E'(1 + \alpha E') + E(1 + \alpha E))} \quad (2.130)$$

where a new deformation potential has been defined for the first order process. We can see that the equations are related by the expression

$$\Gamma_1 = \Gamma_0 \frac{2m_D}{\hbar^2} \frac{D_1^2}{D_0^2} (E'(1 + \alpha E') + E(1 + \alpha E)) \quad (2.131)$$

which tells us that the deformation potentials here do not have the same units; D_0 will have units (typically) of eV/cm, while D_1 has units of eV; they differ by a multiplicative factor with units of inverse length, the wave-vector. So what we're really saying here is that D_1 is the quartic term in our deformation potential expansion in k-space, while D_0 is the quadratic term. Near the threshold for a given transition, we have

$$\Gamma_1 \approx \Gamma_0 \frac{2m_D}{\hbar^2} \frac{D_1^2}{D_0^2} \hbar\omega_{op} \quad (2.132)$$

Plugging in some values to convert the D_1 values given in Ref [112], we find comparable D_0 values as those found in Ref [52]. As these processes are first-order, they are subdominant far above threshold, so their energy dependence is of less consequence. For simplicity, without changing the total scattering rate, we will stick with the 0th order form for all transitions and use the constants from Ref [52] for Si. There no such discrepancy for germanium. These constants can be seen in table 2.2.

We would also like this scattering rate in terms of k-vector for compatibility with the equation

Substrate	Energy (meV)	D_0 (10^8 eV/cm)	Transition
Si	12.0	0.5	TA (g1)
	18.4	0.8	LA (g2)
	61.8	11.0	LO (g3)
	18.9	0.3	TA (f1)
	47.2	2.0	LA (f2)
	58.8	2.0	TO (f3)
Ge	27.6	3.0	LA, LO (LL1)
	10.3	0.2	TA (LL2)

Table 2.2: Phonon intervalley scattering parameters, reproduced from table VI of Ref [52](#).

from the previous section. We can write the scattering matrix element as

$$\Gamma(k, k') = \frac{\pi D_{op}^2}{\rho V \omega_{op}} \left[\begin{array}{c} N_q \\ N_q + 1 \end{array} \right] \delta(\epsilon(k') - \epsilon(k) \mp \hbar \omega_{op}) \quad (2.133)$$

and if we use the zero-point approximation, then we have simply

$$\Gamma(k, k') = \frac{\pi D_{op}^2}{\rho V \omega_{op}} \delta(\epsilon(k') - \epsilon(k) + \hbar \omega_{op}) \quad (2.134)$$

To get the integrated scattering rate as a function of k , we have

$$\Gamma(k) = \frac{V}{(2\pi)^3} \int \Gamma(k, k') q^2 d(\cos(\theta)) d\phi dq \quad (2.135)$$

$$= \frac{D_{op}^2}{8\rho\omega_{op}\pi^2} \int \delta(\epsilon(k') - \epsilon(k) + \hbar\omega_{op}) q^2 d(\cos(\theta)) d\phi dq \quad (2.136)$$

$$= \frac{D_{op}^2}{8\rho\omega_{op}\pi^2} \int \frac{m}{\hbar^2 k q} q^2 d\phi dq \quad (2.137)$$

$$= \frac{D_{op}^2 m}{4\pi\rho\hbar^2 k \omega_{op}} \int_{q_{min}}^{q_{max}} q dq \quad (2.138)$$

$$= \frac{D_{op}^2 m}{8\pi\rho\hbar^2 k \omega_{op}} (q_{max}^2 - q_{min}^2) \quad (2.139)$$

$$= \frac{D_{op}^2 m}{4\pi\rho\hbar^2 \omega_{op}} \frac{\sqrt{8m}}{\hbar} (E(k) - \hbar\omega_{op})^{1/2} \quad (2.140)$$

$$= \frac{\pi D_{op}^2 (2m)^{3/2}}{\rho\omega_{op} 4\pi^2 \hbar^3} \left(\frac{\hbar^2 k^2}{2m} - \hbar\omega_{op} \right)^{1/2} \quad (2.141)$$

$$= \frac{\pi D_{op}^2}{\rho\omega_{op}} \frac{m^2}{\pi^2 \hbar^2} \left(k^2 - \frac{2m\omega_{op}}{\hbar} \right)^{1/2} \quad (2.142)$$

$$= \frac{D_{op}^2}{4\pi\rho\omega_{op}^3} k_{op}^4 (k^2 - k_{op}^2)^{1/2} \quad (2.143)$$

$$= \Gamma_0 \frac{k}{k_{op}} \left(1 - \frac{k_{op}^2}{k^2} \right)^{1/2} = \Gamma_0 M_{op} (1 - M_{op}^{-2})^{1/2} \quad (2.144)$$

where

$$\Gamma_0 = \frac{D_{op}^2}{4\pi\rho\omega_{op}^3} k_{op}^5 = \frac{D_{op}^2}{\pi\rho} \sqrt{\frac{2}{\omega_{op}}} \left[\frac{m}{\hbar} \right]^{5/2} \quad (2.145)$$

and

$$k_{op} = \sqrt{\frac{2m\omega_{op}}{\hbar}} \quad (2.146)$$

Here we see that this rate law turns on at a higher energy than the acoustic rate (and hence is not important at low field strength) but has a weaker momentum dependence, so there is an intermediate field strength region where this becomes more important before being overtaken again by the acoustic phonon scattering rate. This rate law is also responsible for intervalley scattering at high field strengths where there is a strong energy dependence. This is not yet implemented in the DMC, but will be important for accurately simulating high field behavior.

2.4.3 Total Scattering Rate

The neutral impurity, phonon, and total scattering rates are shown in figure [2.4](#). We can see that these rates break transport into acoustic and optical dominated regimes, with large changes in




Figure 2.4: Impurity and phonon scattering rates (both intravalley and intervalley) as a function of carrier energy for Si and Ge. A large energy range is shown in the top panels, while the energy range of interest at our electric fields and operating temperature is shown as a zoomed-in region in the bottom panels. Note that the point where intervalley optical phonons become important is where the acoustic rate becomes sub-dominant to the optical rate, and neutral scattering is never important from a mobility standpoint. The acoustic scattering is entirely intra-valley, so the dominant intervalley scattering at low energy is the neutral impurity scattering.

scattering rate, but also changes to the steady-state charge energy distributions. As optical phonons turn on, they provide a significant kinematic cooling effect which should alter the steady-state energy distribution of the charges. In addition, the acoustic phonon rate will not contribute to intervalley scattering, meaning at low energy we expect that to be dominated by neutral impurities, which in turn are always a negligible effect on the mobility. This assumes typical impurity densities of 10^{11}cm^{-3} for Ge and 10^{13}cm^{-3} for Si with typical binding energies of hydrogenic (shallow single-binding) impurity sites.

While these plots show the scattering rates as a function of energy, and can be thought of as roughly the model we're looking for, we need to estimate the rate as a function of electric field, which is discussed in the next section. We do this to get a feel for how the rate laws change as a function of electric field, and how the energy dependence factors into their average scattering rates.

2.5 Electron Intervalley Scattering Rate

In this section I derive approximate curves for intervalley scattering rate as a function of effective carrier temperature as a proxy for electric field strength. I will later relate effective carrier temperature and electric field strength to plot intervalley scattering rate versus field strength, but it will be a much more approximate curve; the goal here is to show that through some relation between carrier temperature and field strength, the observed scattering rate can be reproduced with the appropriate trends. If this model were incorrect, we would see peaked behavior or the need to scale the scattering rate vertically, rather than just reforming the x-axis.

To make this conversion, I will employ the displaced Maxwellian distribution as the approximate distribution function in energy space. This distribution is essentially a Maxwell-Boltzmann distribution with nonzero mean velocity. This takes the form

$$f(\epsilon, T, v_d) = \frac{1}{v_d \sqrt{2\pi k_b T m_D}} \left[\exp\left(-\frac{(\sqrt{2m_D\epsilon} - m_D v_d)^2}{2m_D k_b T}\right) - \exp\left(-\frac{(\sqrt{2m_D\epsilon} + m_D v_d)^2}{2m_D k_b T}\right) \right] \quad (2.147)$$

where drift velocity can be approximated by the power law

$$v_d = v(E = 1V/cm)E^\alpha \quad (2.148)$$

For Silicon, we find from DMC drift curves that $v(E=1V/cm) \approx 1.95 * 10^4$ m/s, and $\alpha \approx 0.205$. Note that this is a steep departure from the ohmic limit where velocity is linearly dependent on field strength, but is consistent with previous findings. This still leaves electron temperature, which is also a power-law function of electric field, and needs to be measured, or extracted from the DMC.

We thus calculate the scattering rate as a function of electric field as

$$\Gamma_{iv_i}(E) = \int_{\hbar\omega_{iv}}^{\infty} f(\epsilon, T(E)) \Gamma_{iv_i}(\epsilon) d\epsilon \quad (2.149)$$

and rely on the Monte Carlo to prescribe the relations between velocity and effective carrier temperature for the electrons.

2.5.1 IV Scattering from Phonons

The most thorough reference on scattering processes of electrons in Si/Ge and alloys is Ref [52], which thoroughly derives, in a systematic manner, the transition probability for an electron to scatter off of both acoustic and optical phonons. A percentage of these transitions can lead to intervalley scattering, with the largest component of the scattering rate derived from the optical phonon scattering rate. As the change in k vector and energy are nearly exact, emission of either acoustic or optical phonons is treated using the optical phonon emission rate equation given in the previous section.

Let's ignore the non-parabolicity parameter for now, for the sake of simplicity. We can write each intervalley scattering process (per transition) as

$$\Gamma_{iv_i}(T) = \int_{\hbar\omega_{iv}}^{\infty} f(E, T) \Gamma_{iv_i}(E) dE \quad (2.150)$$

$$= \sqrt{\frac{2m_D^3}{\pi^3 k_b T} \frac{D_{iv_i}^2}{\hbar^2 \rho \hbar \omega_{iv_i}}} \int_{\hbar\omega_{iv_i}}^{\infty} \exp\left(-\frac{E}{k_b T}\right) \sqrt{E(E - \hbar\omega_{iv_i})} dE \quad (2.151)$$

$$= \sqrt{\frac{2m_D^3}{\pi^3 k_b T} \frac{D_{iv_i}^2}{\hbar^2 \rho \hbar \omega_{iv_i}}} \exp\left(-\frac{\hbar\omega_{iv_i}}{k_b T}\right) \int_0^{\infty} \exp\left(-\frac{E}{k_b T}\right) \sqrt{E(E + \hbar\omega_{iv_i})} dE \quad (2.152)$$

$$= \sqrt{\frac{m_D^3 k_b T}{2\pi^3} \frac{D_{iv_i}^2}{\hbar^2 \rho}} \exp\left(-\frac{\hbar\omega_{iv_i}}{2k_b T}\right) K_1\left(\frac{\hbar\omega_{iv_i}}{2k_b T}\right) \quad (2.153)$$

So for high temperature we expect the rate to go as \sqrt{T} , and be heavily suppressed for low temperature around temperatures of the energy scale of the intervalley phonon. What's interesting is that the magnitude of the rate ends up being independent of the phonon energy. Here then, the total intervalley scattering rate from all intervalley phonons is

$$\Gamma_{iv,p}(T) = \sum_i \Gamma_{iv_i}(T) \quad (2.154)$$

The overall coefficient comes outside the sum, and we sum over squared deformation potentials times phonon-energy dependent suppression factors.

2.5.2 IV Scattering from Impurities

Here we will again ignore band non-parabolicity and consider the case of a simple power-law for the charged impurity scattering which can approximate the regime we are concerned with in the charged impurity case, and the full formula for the neutral impurity case. Considering charged impurities first, we find

$$\Gamma_{iv,c} = \int_0^\infty f(E, T) \Gamma_c(E) dE \quad (2.155)$$

$$= \frac{2n_I \Gamma_0}{\sqrt{\pi k_b^3 T^3}} \int_0^\infty \exp\left(-\frac{E}{k_b T}\right) E^{-1/2} dE \quad (2.156)$$

$$= \frac{2n_I \Gamma_0}{\sqrt{k_b^3 T^3}} (k_b T)^{1/2} = \frac{2n_I \Gamma_0}{k_b T} \quad (2.157)$$

while for the neutral impurities, we find

$$\Gamma_{iv,n} = \int_0^\infty f(E, T) \Gamma_n(E) dE \quad (2.158)$$

$$= \int_0^\infty 2 \sqrt{\frac{E}{\pi k_b^3 T^3}} \exp\left(-\frac{E}{k_b T}\right) \frac{4\sqrt{2}\pi n_I \hbar^2}{m_D^{3/2}} \frac{E^{1/2}}{E + E_T} dE \quad (2.159)$$

$$= 8 \sqrt{\frac{2\pi}{k_b^3 T^3}} \frac{n_I \hbar^2}{m_D^{3/2}} \int_0^\infty \exp\left(-\frac{E}{k_b T}\right) \frac{E}{E + E_T} dE \quad (2.160)$$

$$= 8 \sqrt{\frac{2\pi}{k_b^3 T^3}} \frac{n_I \hbar^2}{m_D^{3/2}} \left[k_b T - E_T \exp\left(\frac{E_T}{k_b T}\right) \Gamma\left(0, \frac{E_T}{k_b T}\right) \right] \quad (2.161)$$

$$= 8 \sqrt{\frac{2\pi}{k_b T}} \frac{n_I \hbar^2}{m_D^{3/2}} \left[1 - \frac{E_T}{k_b T} \exp\left(\frac{E_T}{k_b T}\right) \Gamma\left(0, \frac{E_T}{k_b T}\right) \right] \quad (2.162)$$

where Γ here is not a rate but the incomplete Gamma function. The portion of this function in brackets is $\approx \frac{k_b T}{E_T}$ for $k_b T \ll E_T$ and 1 for the opposite limit. We can simplify this to a good degree of accuracy (better than 15% across all regimes, better than 1% everywhere but close to the turn-over point) using the function

$$\Gamma_{iv,n} \approx 8 \sqrt{\frac{2\pi}{k_b T}} \frac{n_I \hbar^2}{m_D^{3/2}} \left[\frac{k_b T}{k_b T + E_T} \right] = \frac{n_I \Gamma_{0,n}}{\sqrt{k_b T}} \left[\frac{k_b T}{k_b T + E_T} \right] \quad (2.163)$$

which looks very similar to the original scattering rate function. These trends both allow us to determine effective temperature for a given electric field strength and rule out certain processes. For example, if scattering rate increases monotonically with electric field strength, and we can assume effective temperature can only increase with field strength, then for that range of fields we can rule out charged impurity scattering as the main driver of this scattering. That being said, given just these trends, it will be hard to make a predictive model of intervalley scattering because they are

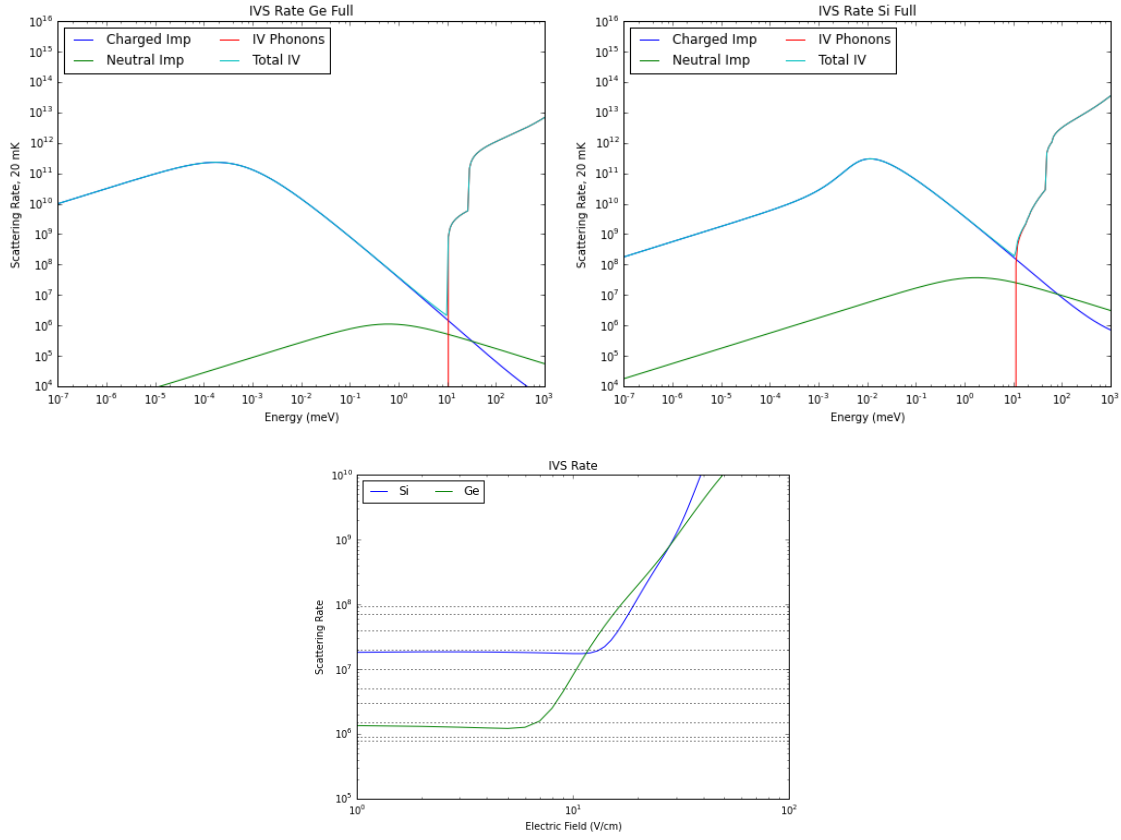


Figure 2.5: Top: Scattering rates relevant to calculating the intervalley scattering rate as described in this section, with the charged impurity scattering rate from [52] shown for comparison with the other two rates described in this chapter. Bottom: Predicted intervalley scattering rates, using the electric field to carrier velocity conversions found in Ref [99].

not the limiting processes for mobility, and there is not a simple power-law or scaling relationship between effective carrier temperature and electric field strength. That being said, we can model it by calculating the mobility including all sources of intravalley scattering, which will give us an approximate conversion.

2.5.3 Total IV Rate Predictions

Application of these equations gives us a mapping of scattering rate as a function of carrier effective temperature, which we can convert to electric field using the Monte Carlo calculations from Ref [99]. Plots of the relevant scattering rates, as well as the final calculation for Si and Ge as a function of electric field, can be seen in Figure 2.5. These plots do not show uncertainties due to impurity density and depth as well as carrier effective temperature, but show that the convolved distribution

results in the prediction of a flat scatter rate at low field strength and a strongly exponential rate at high field strength. The relative number density of impurities (weighted by their depth) is given by the cross-over point between a flat and exponential rate.

This figure shows that charged impurities would produce a rate higher than that of the phonons, inconsistent with our data, and decrease rapidly with energy, also inconsistent with our measurements and those of e.g. Ref [24]. For this reason, we have not included charged scattering in this model. What we are left with is a gentle, roughly constant scattering near the mean carrier energy of 1-10 meV, and a much higher scattering rate which turns on above a moderate threshold. This threshold is larger for Si than Ge, and we note the larger overall impurity scattering rate in Si due to the larger number of scattering centers.

2.6 Results from Stanford Charge Transport Measurements

The previous sections laid down the theoretical framework for understanding some of the results from the charge transport imaging experiments which I will describe in this section. The first iteration of this experiment probed electron and hole transport in Ge as a function of bias voltage for a crystal oriented along the $\langle 001 \rangle$ crystal axis, measuring the intervalley scattering rate and comparing various simulation methodologies to data to determine their accuracy. The full details can be found in Refs [75, 73]. In this section I describe the same study applied to a Si crystal oriented along the $\langle 111 \rangle$ axis, presenting the measurement of intervalley scattering rate versus electric field and discussing some surprising hole transport phenomena.

2.6.1 Experimental Setup

The design of the experiment used to measure transport properties in cryogenic semiconductor crystals was first described in Ref [73], and a schematic of the measurement technique can be seen in Figure 2.6. The basic technique is to raster scan a laser across the back of the test sample, and use the two charge electrodes to determine the amount of charge reaching the small central charge electrode relative to the total charge as a function of initial position, building up an image of the spatial charge distribution as a function of voltage.

In addition, data were taken at a range of base temperatures in order to probe temperature effects between 500 mK and 5K, which in principle should both be closer to the zero-point (non-thermal) approximation than the adiabatic approximation, but do show some interesting differences. A range of data at these base temperatures can be seen in Figures 2.7 and 2.8. Note that as expected, the electron patterns are similar, differing only in over-all charge collection efficiency and slightly in intervalley scattering rate. The holes, however, exhibited an asymmetry not seen in the Ge experiment, which required us to dig into the modeling of the valence bands shown earlier in this chapter.




Figure 2.6: Schematic of the experimental setup, with a laser reflected off of a mirror at the base temperature stage which can be used to raster scan the crystal face to produce charges at different initial positions on the backside of the crystal. These charges are drifted through a variable electric field to the electrodes on the front side, and the fraction of charge collected by the front electrode is used to produce a map of the spatial charge distribution arriving at the front face as a function of initial position. See Refs [75, 73] for more details.

2.6.2 Intervalley Scattering Rate

The intervalley scattering rate as a function of voltage was measured by employing our detector Monte Carlo to simulate acoustic phonon scattering, allowing for determination of the effective carrier temperature, and by simulating a large range of effective intervalley scattering rates in order to match observed and simulated charge collection patterns. The acoustic scattering parameters were taken from Ref [52], and the resulting drift velocity curve can be seen in Figure 2.11

For comparison to Refs [12, 11, 24], we would like to measure the intervalley scattering rate, but our experimental setup is much more conducive to measuring the intervalley scattering length, and using the Monte Carlo to convert the inferred scattering length into a scattering rate. The procedure for determining the intervalley scattering mean free path is best illustrated by the example analysis plots in Figure 2.9. We employed two methods to estimate mean-free path from both the data and Monte Carlo simulations, one using only the zero-scatter peaks (localized in position space and roughly Gaussian) and one using only the central Gaussian blob.

The first method uses the fact that the ratio of zero scatters to multiple scatters can be used as a measurement of the intervalley scattering mean-free path. The fraction of charges which undergo

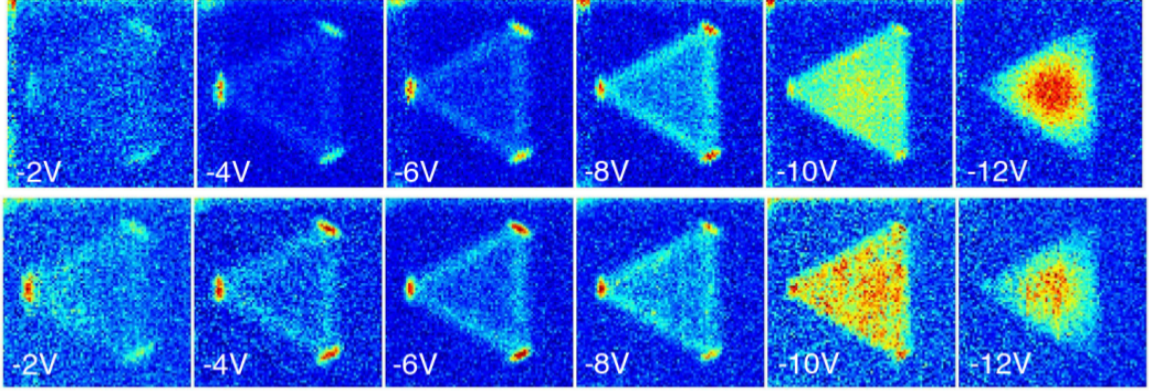


Figure 2.7: Electron charge collection patterns as a function of voltage for 500mK (top) and 5K (bottom) in a ~ 4 mm thick crystal.




Figure 2.8: Hole charge collection patterns as a function of voltage for 500mK (top) and 5K (bottom) in a ~ 4 mm thick crystal.

no scatterings is given by the equation

$$T = \frac{N_0}{N_{total}} = \exp(-\eta/\lambda) \quad (2.164)$$

where η is the crystal thickness and λ is the scattering mean-free path. If we can get a robust estimate of T , then we determine the mean-free path using the equation

$$\lambda = \eta \left[\ln \left(\frac{N_{total}}{N_0} \right) \right]^{-1} \quad (2.165)$$

We determine η by calibrating the Monte Carlo pattern to the crystal pattern; the relative spread of the valleys is geometrically related to the crystal thickness at low voltage, and we find based on simulations that the crystal is $\eta \approx 3.4$ mm thick. We determine the number of zero-scatters N_0 in

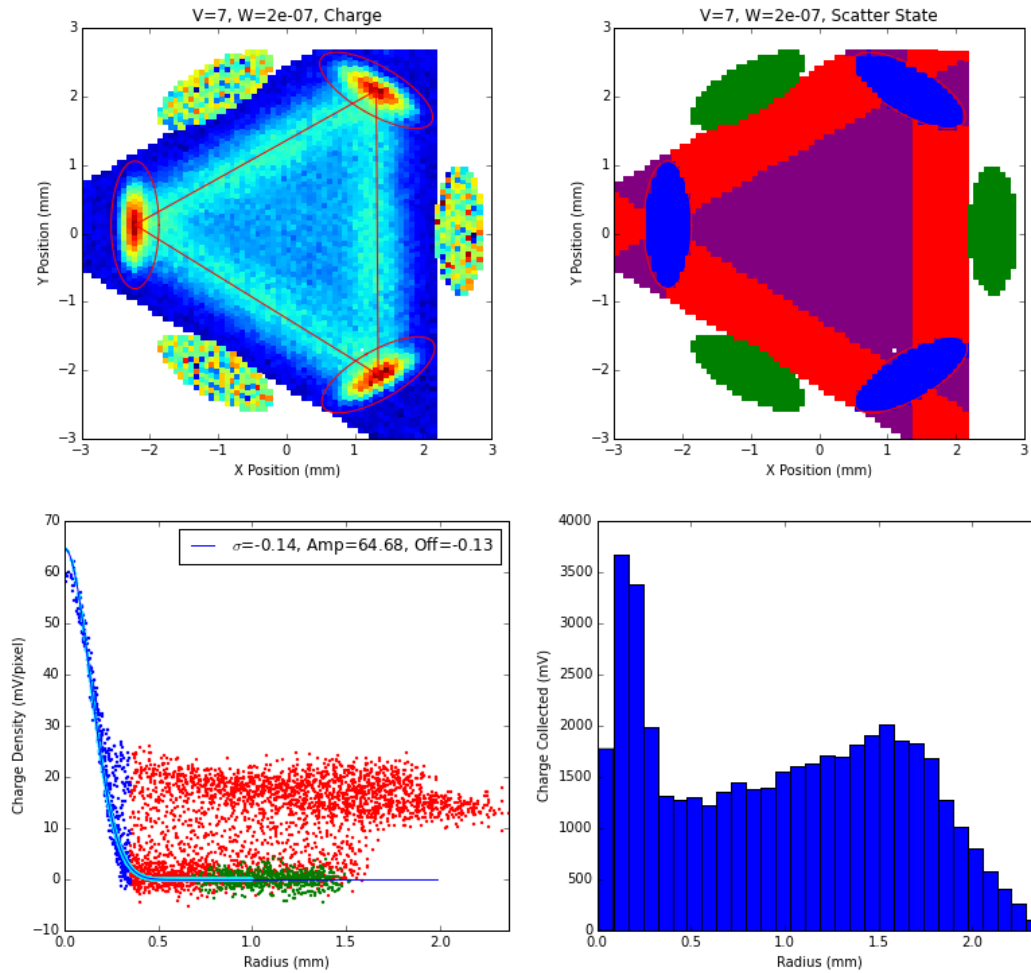


Figure 2.9: This figure shows an example of how information was extracted from both data and Monte Carlo in order to estimate the intervalley scattering mean-free path. The top left panel shows the measured charge distribution, overlaid with lines circling the location of the zero-scatter peaks and following the single-scatter lines connecting the peaks. The top-right panel shows how the zero, single, and multiple scatter peaks are determined, and also the green regions used to estimate the noise floor. These regions are plotted in the lower two panels, and a gaussian fit to the zero-scatter peaks is shown as a cyan line.

two ways:

- Integrate the charge within the identified zero-charge region.
- Fit the outer half of the zero-charge region to a Gaussian (distorting the ellipse into a circle) and using the measured amplitude and variance to estimate the charge in the zero-scatter peak.

The second method is more robust, as charges which undergo scatters close to the electrode may still be included in the nominal zero-scatter region, but they have a low probability of making it to the far half of the zero-scatter ellipse. The impact of this single-scatter distribution can be seen in Figure 2.9 as blue points which do not align with the Gaussian fit. The results of these two methods are both plotted in Figure 2.10 in both upper panels. We can see that the integration method is biased relative to the fitting method, as in the high mean-free path case, it includes extra charges, and in the low mean-free path case, it has worse signal to noise. Thus for voltages with significant zero-scatter peaks, we use the fitted method, and the integration method gives us a measure of the maximum systematic error on the measurement.

The second technique, which we use for short mean free paths (high voltage) is the variance technique. For many intervalley scatters, the observed charge distribution is roughly Gaussian, and we know that for a random walk with mean step size λ , the variance of the resulting probability distribution is proportional to that mean-free path:

$$\sigma^2 = \frac{t}{\delta t} \lambda^2 = \frac{\eta/v}{\lambda/v} \lambda^2 = \eta\lambda \quad (2.166)$$

The mean-free path using this exact expression was attempted, and the resulting measurement differed from the zero-scatter measurements by a factor of order unity, related to the fact that this is not a true free-space random walk and the mean free path is directionally dependent. We thus used some points with good zero-scattering measurement to calibrate the relative scale factor for this technique, and then used the technique to extend the mean-free path measurement out to higher voltage, where the zero-scatter technique was statistics limited.

The mean-free path measurements can be seen in the top-left panel of Figure 2.10 and the agreement between variance and zero-scatter techniques is seen to be very good. Also shown in the bottom-left panel of Figure 2.10 is the model prediction for mean-free path given different effective neutral impurity densities using the scattering rates derived earlier in this chapter, and using tabulated values for the carrier temperature and drift velocity as a function of voltage from Ref 99. This is a much better prediction than those shown in 11, where intervalley transitions below the optical mode energy was modeled as charged impurity scattering. This result thus lends support to an intervalley scattering model dominated by high-energy phonons at high voltage and neutral impurities at low voltage.

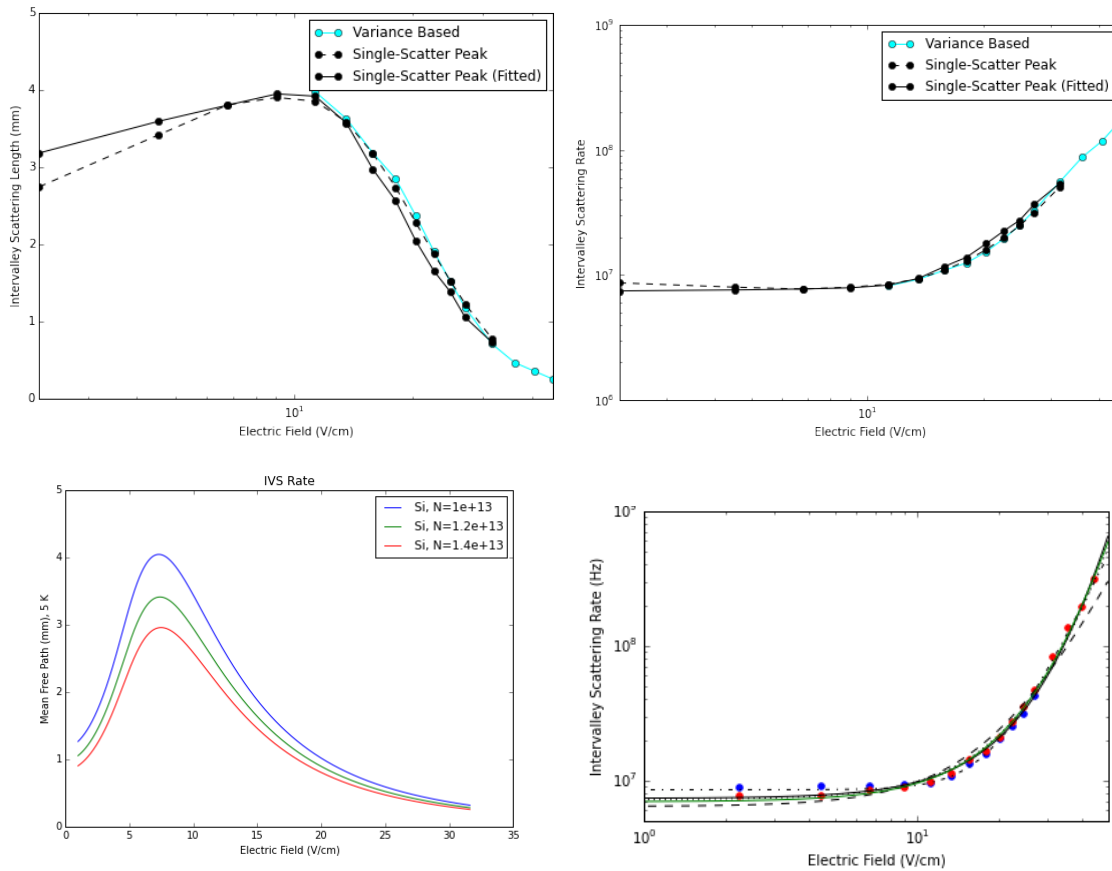


Figure 2.10: Top left: Intervalley scattering length as a function of applied electric field based on three different estimation techniques described in the text. Bottom left: predicted intervalley scattering length as a function of electric field based on the scattering theory presented earlier in the chapter for different neutral impurity densities. Top right: IV scattering length inferred from Monte Carlo simulations given mean free path estimates in the top left panel. Bottom right: final intervalley scattering rate points fit to different functional forms as described in the text.




Figure 2.11: Left: Electron drift velocity as a function of applied electric field using the acoustic scattering parameters from the literature (see the text for more details). This drift curve is used to convert the measured mean-free path into a scattering rate for comparison with literature predictions and other measurements. Right: Rate determination using matching DMC mean free path versus rate to data-measured mean free path at each measurement voltage. Both methods were used to determine the rate independently as an estimate of systematic error where possible.

The second half of this analysis entailed converting these mean free paths into voltage-averaged scattering rates, which was done by running a grid of Monte Carlo simulations at different voltages and IV scattering rates, applying the same measurement techniques to estimate the effective mean-free path, and matching the scattering rate figure at a given voltage to the corresponding measurement at that voltage. The detector Monte Carlo employed is the G4CMP package, a GEANT extension for cryogenic semiconductor charge and phonon transport simulations being developed by our collaboration to simulate SuperCDMS detectors (see Ref [3] and references therein). An example of this correspondence procedure can be seen in Figure 2.11 along with the drift curves for all simulations as a function of voltage. This figure also demonstrates the independence of the drift velocity and IV scattering rate, given that it is a minor process in determining overall electron mobility relative to acoustic scattering. The resulting measured IV scattering rates can be seen in Figure 2.10 in the right-hand panels.

The trend of IV scattering rate versus voltage was fit to two function forms. The first is from [24]:

$$\Gamma_{IV} = \Gamma_0 [E_0^2 + E^2]^{a/2} \quad (2.167)$$

The second is based on the addition of the phonon and impurity scattering mean free paths based on Matthiessen's rule:

$$\Gamma_{total} = \sum_i \Gamma_i \quad (2.168)$$

which for a two-component model, where the low-field scattering is impurity dominated and the

Functional Form	Γ_0	E_0	Γ_1	α
Non-Linear (Eq 2.167)	$3.5 \cdot 10^{-20} \text{ Hz} \cdot (m/V)^\alpha$	3395 [V/m]	-	7.47
Linear (Eq 2.169)	$9.8 \cdot 10^6 \text{ Hz}$	-	$3.11 \cdot 10^{-7} \text{ Hz} \cdot (m/V)^\alpha$	4.02

Table 2.3: Intervalley scattering rate parameters for Si fit to the data presented in this section.




Figure 2.12: Intervalley scattering rate fit to the equations presented in this section, including residuals for the rate fits. The linear equation is shown as a solid line, and can be seen to have consistently small residuals, while the non-linear form cannot simultaneously fit all regimes.

high-field scattering is dominated by a phonon mean free path, we have

$$\Gamma_{IV} = \Gamma_I + \Gamma_{Ph} \approx \Gamma_0 + \Gamma_1 E^\alpha \quad (2.169)$$

In the low and high-field limit, these equations have the same field dependence, but in the intermediate limit where $E \sim E_0$, they behave differently. Our data were a much better fit to the second form, as shown in Figure 2.10, where the earlier form which is added in quadrature is not capable of fitting the data across all measured electric fields. The fit parameters for Si for both functional forms can be seen in Table 2.3.

We can compare these functional form using the residuals, as shown in figure 2.12 reproduced from Ref [74]. We find that the uncorrelated model based on adding mean free paths is a consistently better fit than the non-linear form across the range of electric fields based on the consistently lower and unbiased residuals of the solid line (the linear model) as compared to the dashed and dotted lines (the non-linear model fit to either the high or low energy points). We also notice, in comparison to the Ge measurements and the predictions shown in figure 2.5, that intervalley scattering turns on at a higher than expected field strength, but the power laws are similar between Si and Ge as

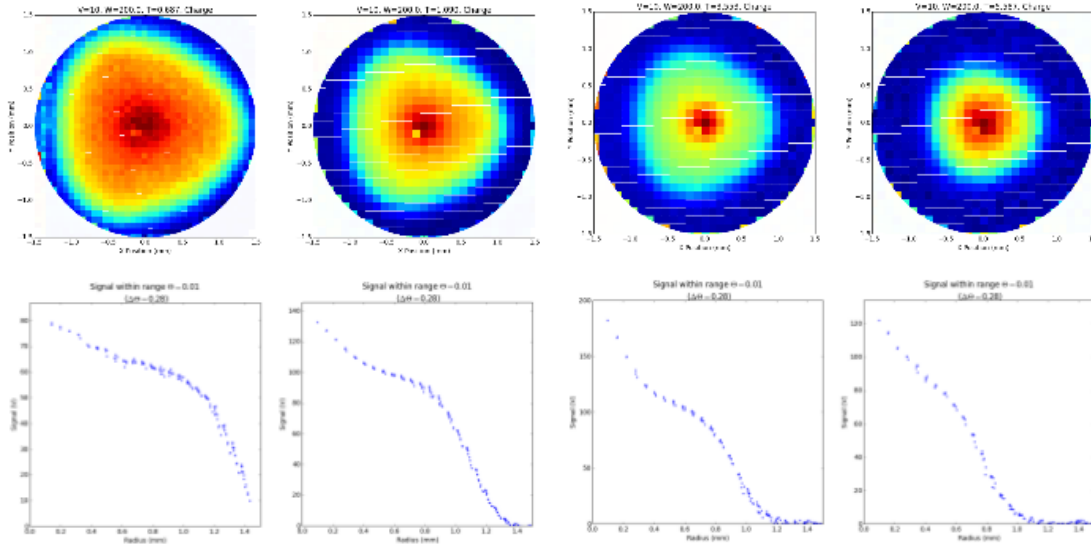


Figure 2.13: Hole collection pattern as a function of temperature for 500mK, 1.4K, 3K, and 5K from left to right for fixed laser intensity. The top panel shows the two-dimensional collection pattern, and the bottom shows the charge density as a function of distance from the center of the pattern. See text for discussion.

predicted. We can attribute this to a difference in effective carrier temperature as a function of electric field due to different acoustic coupling and optical phonon energies.

2.6.3 Hole Anisotropies

Subsequent to the intervalley scattering measurements, we began to analyze the hole data and quickly realized that there was temperature dependence in the hole patterns, and at low temperature the hole collection was highly anisotropic. I explored these anisotropies with James Allen, and the detailed study of the character of the anisotropies is discussed in his undergraduate thesis (Ref [9]). The initial observations are summarized in Figure 2.13, where you can see the tri-lobe shape become less prominent as a function of temperature and the hole pattern become more Gaussian.

As discussed in section 2.3, in the small k limit, the hole bands are highly warped, with light and heavy hole bands exhibiting asymmetric and complementary equal energy surfaces. This warping is most pronounced for $E \sim \epsilon_{so}$, the spin-orbit coupling energy. This is plotted in figure 2.14 for a relatively stationary hole as well as a hole with significant crystal momentum along the electric field. We can see that for the second case, an anisotropy very similar to that seen in our hole transport data arises in the equal energy surfaces. We saw in the previous section that the symmetry of these surfaces is the same as the spatial anisotropy, though the shape can be significantly different depending on the energy distribution, and is based on the gradient of these surfaces.




Figure 2.14: Brillouin zone for the holes rotated so that k_z is along the $\langle 111 \rangle$ direction. **Top:** Dispersion relation at $k_y = 0$ for $E(k_x = 0, k_z) = 1$ meV and $k_z = 0$ at the band minimum. The equal energy surfaces for both the light and heavy holes as a function of crystal momentum become anisotropic for charges at non-zero drift velocity. For drifting charges, the heavy hole band will be preferentially filled by inter-band phonon emission. **Bottom:** Contours of constant kinetic energy for heavy (light) holes are shown on the left (right) where the center of the pattern shows contours for carrier energy of ~ 3 meV for the heavy holes and the mean energy of the contours is ~ 10 meV, comparable to the mean hole energy for low electric field strength. The upper plot is a cutline through the lower plots for $k_y = 0$. Reproduced from Ref [74].

Our initial hypothesis to explain the effect seen in figure 2.13 was that as we lower the temperature, hole acoustic scattering is suppressed, and holes achieve higher kinetic energies. Given that the spin-orbit coupling energy is 44 meV in Si [52], and that is also characteristic of typical carrier energies at low temperature, we should expect to see a lot of warping as the mean free path of the carriers increases. In contrast, we do not see this effect in Ge, which has a spin orbit coupling energy of ~ 200 meV, much larger than the mean carrier energy at any temperature studies in our previous paper (see Ref [73]).

Later data taken for a few targeted experimental conditions showed that, while the anisotropy does seem to be present, the temperature dependence we're seeing has less to do with intrinsic charge propagation than the spatial distribution of initial charge produced. This data, shown in figure 2.15 shows that the shape of the charge distribution is preserved but the overall size and flatness depends more on the initial charge produced in the crystal than the temperature or field strength. This data, along with higher voltage electron data at various temperatures, indicates that small-scale transport is affected by charge repulsion, which needs to be modeled using various initial laser intensities to isolate its effects away from anisotropies due to transport. Our current detector Monte Carlo implements the holes as a spherical, parabolic band, and we are in the process of implementing anisotropic hole transport and charge repulsion in order to test this hypothesis.

2.7 Future Charge Transport Studies

The charge transport work presented in this chapter was a driving force for improvements to G4CMP, which was based on an earlier MATLAB Monte Carlo built assuming any crystal under study was Ge aligned along the X-valley, and was meant to simulate cryogenic Ge at low electric fields. As of this thesis, we have now run Ge $\langle 001 \rangle$ [73, 75] and Si $\langle 111 \rangle$ (this thesis and publication in preparation), which motivated the development of miller index rotation and generalization of the code to an arbitrary substrate. We still need to run the complementary crystals (Ge $\langle 111 \rangle$ and Si $\langle 001 \rangle$) to verify that the acoustic and intervalley scattering rates are correctly implemented as a function of rotation angle and voltage.

We are also in the process of operating Si at higher electric fields (up to 500 V/cm) and we have observed tentatively that some of the anisotropy in the electron pattern re-emerges at field strengths above those included in this chapter. The detector Monte Carlo does not predict this, so we need to determine whether this is the result of some detector effect, or whether this is due to higher order phonon processes not implemented in the current Monte Carlo. Many of the effects we observe in our charge transport studies are outside of the realm of the existing literature, which does not focus on spatial charge propagation at low temperature, and thus we need to carefully implement additional physical processes to determine the origin of the signal.

There are many improvements which need to be made to the Monte Carlo in order to produce




Figure 2.15: Two-dimensional charge density patterns for electrons in $\langle 1, 1, 1 \rangle$ Si as a function of bias voltage. **First and second row:** Data recorded at 5 K and 500 mK. **Third row:** A simulation of the patterns in the zero-temperature limit. For the rows of data, each pixel is assigned a color according to the measured pulse height, with red indicating larger pulses (and thus more charge collection). The white numbering shows the maximum pulse height (in meV) and the normalized integrated intensity (relative to -12 V 500 mK) for that panel. For the MC, a conversion was used such that peak intensity for the -12 V simulation agrees with the -12 V 500 mK data. Reproduced from Ref [74].

something useful for detailed charge transport studies:

- Merge the 'Luke Emission' process with charge transport code to tie acoustic and intervalley scattering to electron-phonon amplification. There is no need to separately calculate the phonon emission rates and intervalley scattering rates given that they are known processes, and tying intervalley scattering to optical/acoustic phonon scattering will make the code more accurate and predictive.
- Finish the implementation of impurity scattering, including a more physical calculation of impurity scattering rates. The scattering rate in this thesis assumes purely elastic scattering, but does not do the proper solid angle calculations, and a full theory of neutral impurity scattering at low temperature should be properly quantum mechanical.
- Allow for valley minima along vectors with non-unit length, to allow for exploration of crystal deformation on charge transport properties.
- Implement anisotropic hole propagation in order to simulate the hole collection patterns seen in the data.

These studies are not necessary for using G4CMP in SuperCDMS SNOLAB detectors, as we will always be able to tune relative scattering rates to match data, but will be important if G4CMP is to be used to predict behavior in new materials which have not been as thoroughly tested in cryogenic settings, such as diamond.

Chapter 3

Detector Concepts

“If you take a bale of hay and tie it to the tail of a mule and then strike a match and set the bale of hay on fire, and if you then compare the energy expended shortly thereafter by the mule with the energy expended by yourself in the striking of the match, you will understand the concept of amplification.”

- William Shockley, as quoted by Fred Warshofsky, 1989

This chapter documents the work that went into designing the HV detectors for SuperCDMS SNOLAB (summarized in Ref [59]), and comments on various design choices in light of the testing results shown in the following chapter. It should be seen as a snapshot of SuperCDMS detector technology at this time of writing, in the same way that Matt Pyle’s thesis [83] was a snapshot of detector design for the last generation of SuperCDMS Soudan detectors. At the beginning of the work shown here, it seemed clear to us that bigger was better, but our recent work has resulted in a bifurcation of detectors into smaller, ultra-high sensitivity detectors for calibration and surface testing and kg-scale detectors for higher thresholds and lower backgrounds. This is driven by some of the non-ideal background processes that were encountered during testing, but also by the recognition that for sub-GeV dark matter and photon detection, running kg-scale detectors amounts to trying to tie a shoe with a steam shovel...things move too quickly to get the job done right.

This chapter is organized as follows. Section [3.1] describes the basic flow of energy through the detector and contextualizes the following sections. Section [3.2] derives transfer and Green’s functions for the transition edge sensor arrays we use for our detectors in terms of the free design parameters of the QET. Section [3.4] describes the efficiency model for a single QET, and section [4.2] then describes the bandwidth matching and efficiency optimization of the QETs with the detector given their segmentation into channels and the wiring needed to read them out. Section [4.3] describes the electric field and detector response modeling used to simulate the operation of the detector and




Figure 3.1: Schematic overview of the separate prompt and NTL phonon emission for high voltage detector operation.

optimize the detector-housing separation, and section 4.4 describes the use of the Monte Carlo to optimize the mask layout for the best position resolution in the detector. I conclude with some summarizing thoughts and suggestions for future work based on work since this design process and test results shown later in this thesis.

3.1 Detector Overview

SuperCDMS detectors are phonon and charge calorimeters; that is, they measure total energy collected by summing measured power over time, calculating this total energy using various measures of deposited energy in the charge and phonon energy systems. When any interaction occurs in the detector crystal, it produces a large population of phonons, which we call 'prompt' phonons, as well as a number of charge carriers. Immediately after the event, the energy is divided according to the energy conservation formula

$$E = E_{ph} + \epsilon_{gap} n_{eh} \quad (3.1)$$

where E_{ph} is the energy of the prompt phonons, ϵ_{gap} is the bandgap energy (discussed in chapter 2), and n_{eh} is the number of electron-hole pairs produced by the event.

The detectors are operated with a voltage between the two faces in order to drift the charge carriers generated by events to the instrumented detector faces. As the charges are drifted, they generate an additional phonon energy due to production of additional athermal phonons we refer to as Neganov-Trofimov-Luke [69, 77] (NTL) phonons. Due to energy conservation, the total amount of phonon energy produced is just $e \cdot V$, so that the total phonon energy produced during the event is just

$$E_{ph} = E_r + n_{eh} V \quad (3.2)$$

We can parameterize the charge production as

$$n_{eh} = y(E_r) \frac{E_r}{\epsilon_{eh}} \quad (3.3)$$

where $y(E_r)$ is the charge yield (normalized to 1 for electron-recoils, and taking on a value between 0 and 30% for nuclear recoils) so that the total phonon energy is

$$E_{ph} = E_r \left(1 + V \frac{y(E_r)}{\epsilon_{eh}} \right) \quad (3.4)$$

Here ϵ_{eh} is the energy needed to produce an electron-hole pair on average, which differs from the energy gap for indirect band-gap semiconductors. This shows that we have two nominal detector modes, a low-voltage mode (where there is minimal degeneracy between phonon and charge energy) and a high-voltage mode, where the phonon energy is proportional to the charge energy. These two extremes are approximately the case for the iZIP (low-voltage) and HV (high-voltage) detector designs. This detector model is shown schematically in Figure 3.1.

As an aside, we should note that it is guaranteed that all of the energy that enters the detector will end up as phonon energy eventually. If we generate an electron-hole pair with a large photon (100 eV for example) it will very quickly emit high-energy athermal phonons which will produce additional charge carriers. It is an empirical fact that for electron-recoils above $\sim 10\text{eV}$, the mean charges produced is linear in initial energy, taking the form

$$n_{eh} = \frac{E_r}{\epsilon_{eh}} \quad (3.5)$$

where ϵ_{eh} is roughly the energy of the direct bandgap. Each of these charges can then continue to emit athermal phonons to reach the bandgap energy ϵ_{gap} , immediately re-emitting phonons with total energy $\epsilon_{eh} - \epsilon_{gap}$ for every charge carrier. For these electron-hole pairs to reach the detector surfaces, they need to retain this bandgap energy, but once at the surface there is some probability they can re-emit that energy as a phonon and recombine. If that happens during the normal time-scale of the event, the total energy is given by equation 3.4, but if that happens at some time later, we measure an event with total energy

$$E_{ph} = E_r \left(1 + (V - \epsilon_{eh}) \frac{y(E_r)}{\epsilon_{eh}} \right) \quad (3.6)$$

and a handful of secondary events with energy equal to the indirect bandgap. This secondary population would grow as a function of time, and would be removed by neutralization and/or pre-biasing of the detector. This doesn't necessarily change our understanding of the detector signals but does represent both a systematic on the determination of total energy efficiency (we don't know what the true input energy was) and a potential background we need be aware of when employing




Figure 3.2: Schematic of energy transport in a QET. The phonons break cooper pairs in the superconducting aluminum fin, which diffuse to the trap region. The trap region has a lower bandgap due to the proximity effect with the tungsten, and the quasiparticles are trapped there and drop their energy as heat into the TES. The TES then amplifies this input power as a change in current, read out by upstream amplifiers.

crystals with moderate carrier lifetimes (high enough to be metastable but not stable).

Figure 3.2 shows an overview of the phonon sensors, called Quasiparticle-assisted Electrothermal-feedback Transition Edge Sensors (QETs). The operating principle of the QET is to use a material with a small bandgap to collect phonon energy, and diffuse this energy towards a smaller volume with a lower bandgap which then traps and absorbs this energy (see Figure 1.16 for the energy scale of superconductor and semiconductor bandgaps). Each conversion step in this process has an associated efficiency which has either been measured or estimated, and the design of the QET-crystal system has become a mostly solved optimization problem, as will be described in the next few sections. The fundamental limitation comes from the process of converting phonons to quasiparticles, which can be shown to be at most $\sim 60\%$ efficient [23].

To give the reader some pedagogical context for why we're using cryogenic technology to achieve such low energy resolution, let's consider the model of the ideal calorimeter, which is a block of material with specific heat c and volume V such that $C = cV$. The block is attached to some heat bath of infinite heat capacity by a thermal conductance G . This system is thus governed by the equations

$$\frac{dT}{dt} = -\frac{G}{C}(T - T_b) \rightarrow T(t) = (T_0 - T_b)e^{-t/(C/G)} + T_b \quad (3.7)$$

which responds to power fluctuations slower than the time constant $\tau = \frac{C}{G}$. In fourier space this transfer function is

$$S(f) = \frac{\epsilon}{1 - i\omega\tau} \rightarrow |S(f)|^2 = \frac{\epsilon^2}{1 + \omega^2\tau^2} \quad (3.8)$$

where ϵ is the energy transfer efficiency. We find the energy resolution of a device with noise power

$|N(f)|^2$ and signal power $|S(f)|^2$ as (see section [E.6](#))

$$\sigma_E^2 = \left[4 \int_0^\infty \frac{|S(f)|^2}{|N(f)|^2} \right]^{-1} \quad (3.9)$$

and for the signal found above, and a noise power $|N(f)|^2 = 4k_b T^2 G$, we find

$$\sigma_E^2 = \frac{k_b T^2 G}{\epsilon^2} \tau = ck_b \frac{VT^2}{\epsilon^2} \quad (3.10)$$

This is a fundamental limitation of calorimeters that are thermal noise limited. If, for example, you want to design a calorimeter using a known material at a known efficiency, the design trade-off is thus between volume and temperature; so large calorimeters with very low energy resolution need to be operated at low temperature. The caveat to this argument is of course that athermal signals can overcome the thermal noise, but there will almost always be a critical thermal-noise limited stage that obeys this scaling relation.

Unlike past CDMS detectors, the CDMS-HV detectors do not read out the ionization signal, and thus instead of balancing charge and phonon resolution, we can attempt to maximize phonon resolution in our mask designs. We also need to balance resolution with phonon fiducialization, however, which adds additional degrees of freedom to the design of our QET, as the size of the unit cell will depend on the number of unit cells per unit area, set by the phonon channel size. In this chapter I will describe the elements of detector modeling incorporated into the optimization for the CDMS HV detector, the results of which are the subject of chapter [4](#).

3.2 TES Dynamics

The QET unit cell consists of a tungsten TES connected to aluminum fins via small tungsten connectors, and a phonon channel consists of N TESs voltage biased in parallel. The effective resistance of a given TES is

$$R_T = \rho_w \frac{l_T}{w_T h_T} \quad (3.11)$$

and the normal state parallel resistance of N TES is

$$R_n = \frac{R_T}{N_T} \quad (3.12)$$

so we can relate the dimensionality of the TES to number per phonon channel:

$$N_T = \frac{R_T}{R_n} = \frac{\rho_w}{R_n} \frac{l_T}{w_T h_T} \quad (3.13)$$

Quantity	Symbol	Adopted Value
Normal State Phonon Channel Resistance	R_n	$\sim 150 \text{ m}\Omega$
Operating Point Phonon Channel Resistance	R_o	$R_n/3$
TES Circuit Inductance (effective)	L_{eff}	$\sim 145 \text{ nH}$
TES Critical Temperature	T_c	45-60 mK
TES Thickness	h_T	40 nm
TES Width	w_T	$2.4 \mu\text{m}$
TES Length	l_T	100 – 300 μm
TES Fractional Volume Coefficient	ζ_T	~ 0.7
Tungsten Square Resistance	R_{sq}	3.3 Ω
Tungsten Resistivity	ρ_W	$1.32 \cdot 10^{-7} \Omega\text{m}$
Wiedermann-Franz Coefficient	β_{wf}	$2.44 \cdot 10^{-8} \frac{W\Omega}{K^2}$
Electron-Phonon Coupling Constant [45]	Σ	$0.4 \cdot 10^9 \frac{W}{K^5 m^3}$
Thermal Conduction Power-Law Exponent	n	5
Tungsten Specific Heat	c_W	$108 \frac{J}{K^2 m^3}$
Superconducting Heat Capacity Increase	f_{sc}	1.0 - 3.0
Crystal Temperature	T_{bath}	< 30 mK
Logarithmic Temperature Sensitivity	α_I	100-150
Logarithmic Current Sensitivity	β_I	< 0.3

Table 3.1: TES design parameters adopted for geometry determination, with lower limits on dimensions set by photolithography and continuity considerations. ρ_w , Σ , and n from Table 3.1 in Ref [83] (page 24). Crystal temperature taken from chapter 5 of the CDR. See text for more details.

where these quantities are as defined in Table 3.1. Here we see that we can maximize the number of TESs per channel by minimizing TES width and thickness and maximizing TES length. The minimum TES width and thickness are process dependent, with values shown in Table 3.1, and so the only parameters we are able to optimize further are the normal state resistance of the phonon channel and the length of the TES. The resistance of a given phonon channel is set by constraints of our readout system as well as electrothermal oscillation, which is described in section 3.2.3. The maximum length of TES is limited by the need to keep the TES in thermal equilibrium, and the phase separation length, described in section 3.2.5, sets the upper limit of TES length for which thermal equilibrium can be maintained.

The resistance of a given TES channel is bounded on either side by signal-to-noise considerations of our readout electronics. We expect the non-TES resistance of the resistive load on the input-coil circuit to be on the order of $5m\Omega$ for SNOLAB cold electronics, and thus we desire that the operating point resistance of the TES circuit be much larger than this value. We would also like to maximize the current range of our voltage biased TES circuit to maximize signal-to-noise in the SQUID readout and thus want to use as low a resistance as possible for the TES operating point. In addition, the time constant of the LRC circuit constructed by the inductor, TES, and thermal heat capacitance sets the bandwidth of our TES readout, and given some inductance we can increase the bandwidth by increasing resistance and/or capacitance of the TES. These considerations are covered in section 3.2.2. In the following section, I discuss the assumptions which go into QET optimization,

and where some of the parameters taken as constants come from.

3.2.1 TES Characteristics and Parameters

The parameters and parameter ranges chosen for the design of the TES in our QETs can be seen in Table 3.1 along with some physical constants which allow us to relate heat capacities and conductances to physical dimensions and critical temperature. Many of these are quantities derived from more fundamental physical principles, summarized here for completeness, while others are measured properties of our devices which are process specific; still others are values chosen to ensure safety margins away from critical failure regimes.

Tunable Parameters

For our purposes, we have chosen the normal-state resistance of the entire phonon channel to be $\sim 150\text{ m}\Omega$, corresponding to an operating point resistance of $\sim 50\text{ m}\Omega$. This operating point is roughly twice as high as for past devices, driven in part by the fact that our inductance is higher, and to first order we want to maintain an inductance to resistance ratio which we've known to be stable in the past. This is an assumption which we'll re-visit at the end of this section once we have established stability and oscillation criteria, but it is within a factor of two of our final value. Note that decreasing this value, we'll see, increases TES number, but will also increase Johnson noise. This will be explored later in the section; I use it as a reference point to explore the response equations derived in this section.

Many of the parameters in Table 3.1 are set to the smallest values achievable with current fabrication abilities:

- The width of our TESs are limited mainly by their photolithography and are currently assumed to be $2.4\text{ }\mu\text{m}$; in reality this number drifts between 2 and 3 microns, and is thus a geometric average. It is limited by the diffraction limit of the alignment stage of photolithography and the shear resistance of the photoresist used.
- The height of the TESs is similarly limited by fabrication, as we can't consistently make thinner films without risking breaks in continuity. In addition, tungsten has a non-trivial phase boundary between the bottom and the top of the film, thus a change in thickness usually produces a change in transition temperature and thermal conductivity (results relevant to this point will be explored in the next chapter).
- The inductance listed in the table is an estimate based on the initially measured inductance of our tower wiring, the input coil, and assumes the mask will not have an inductance exceeding $\sim 20\text{ nH}$.

- The T_c is primarily set by that which we can consistently fabricate without a significant critical temperature gradient across the TESs in a given channel, as well as the requirement that T_c be more than 10% higher than bath temperature for optimal temperature flow and predictable gain. It is nominally limited by the magnetic impurity content of the tungsten target used during fabrication [113].

To expand on this last point, tungsten has three crystal phases. α phase tungsten, which has a T_c of 15mK, and β phase, which has a T_c of 1-4K [92]. By mixing these phases or perturbing them, one can tune the T_c to have an intermediate value. In thin films, the crystal phase grown at the bottom of the film will depend on the substrate, and as the film grows the ambient environment will determine whether this phase is maintained or a different phase develops. Thus the T_c of a film is the average T_c of the layers due to the proximity effect [23], and if there is a gradient of crystal morphology, we will expect geometric effects to alter the T_c of the film, not to mention the physical properties (e.g. thermal conductance, heat capacity, and resistance, which are all linked at low temperature by electron-phonon dynamics).

In addition, the following sections show some dependence on intrinsic parameters as well as V_{TES} , the TES conduction volume:

$$V_{TES} = l_{TES} h_{TES} w_{TES} \quad (3.14)$$

This volume is obviously important for an individual TES; more relevant to our design goals is the total TES channel volume

$$V_{channel} = V_{TES} N_{TES} = \frac{\rho W}{R_n} \frac{l_{TES}}{w_{TES} h_{TES}} l_{TES} h_{TES} w_{TES} = \frac{\rho W}{R_n} l_{TES}^2 \quad (3.15)$$

So the heat capacity of the channel increases as l_{TES}^2 . This is one justification for opting for longer TESs all else being equal, as you'll have a larger TES dynamic range, but of course in principle a worse resolution by a similar amount.

As a note of caution, I would like to stress to the reader that superconductivity is a bulk phenomenon, and the dimensions of the TES (as well as the morphology of the film) will affect all of the parameters described in this section, including heat capacity, resistance, critical temperature, and thermal conductivity. Many of our designs which did not perform as expected resulted from the assumption that ρ and C are independent of the critical temperature of the material, and that T_c is independent of film dimensions. In practice both have been shown to be poor assumptions. This will be discussed further in the next chapter, but I urge the reader to keep it in mind when considering any specific numbers considered in this section.

Intrinsic Resistance Parameters

Other parameters are measured from our devices, and we have less control over them. Most notably, when we talk about dynamics in the next section, we'll need to linearize resistance in temperature

and current changes as

$$R(I, T) \approx R_0 + \left. \frac{\partial R}{\partial T} \right|_{I_0} \delta T + \left. \frac{\partial R}{\partial I} \right|_{T_0} \delta I \quad (3.16)$$

As in Ref [50], which references papers by both McCammon and Mather for this substitution, we use the dimensionless logarithmic temperature sensitivities α and β into the equation to finalize the linearization, where

$$\alpha = \left. \frac{\partial \log R}{\partial \log T} \right|_{I_0} = \frac{T_0}{R_0} \left. \frac{\partial R}{\partial T} \right|_{I_0} \quad (3.17)$$

and

$$\beta = \left. \frac{\partial \log R}{\partial \log I} \right|_{T_0} = \frac{I_0}{R_0} \left. \frac{\partial R}{\partial I} \right|_{T_0} \quad (3.18)$$

to give

$$R(I, T) \approx R_0 \left(1 + \alpha \frac{\delta T}{T_0} + \beta \frac{\delta I}{I_0} \right) \quad (3.19)$$

As we want changes in resistance to be sensitive to temperature variations, and minimally sensitive to current variation, we see that our ideal device would have very high α and $\beta \sim 0$. α is in this case the slope of the super-conducting transition, which is very large and thus explains why TESs are such sensitive thermal energy sensors, however it is very hard to control α and β through fabrication, and β tends to be non-zero in any real device.

These are two parameters which I quote ranges for because we have an idea of values measured in past devices, but we won't know the true values until the first mask is fabricated and characterized. We expect α to be in the 100-150 range, which tells us that a 1% change in temperature leads to a doubling of resistance in the case that $\beta \sim 0$, but for a large β , current would decrease, reducing the overall increase in TES resistance. The connection of these parameters to a phenomenological resistance curve is shown in section B. The ideal critical temperature is a first-order phase transition, so the closer a film is to its bulk transition temperature, the sharper we can expect the transition to be [92]. The less bulk-like a film is, the more one should expect deviations from ideal and shallower transitions [65].

Intrinsic Heat Transfer Parameters

Finally, there are some constants which have to do with thermal storage and transfer. The heat capacity for the TES at critical temperature is given by

$$C_{T_c} = f_{sc} c_W V_{TES,eff} T_c \quad (3.20)$$

where c_W is the specific heat capacity of tungsten, and f_{sc} is the correction to the normal curve due to added heat capacity of cooper pair breaking in the super-conducting transition, which can vary from 1 to ~ 2.4 (also given in Table 3.1 [13, 66, 67]. This will be discussed further in a moment.

The effective TES volume is given by

$$V_{TES,eff} = V_{TES} + \epsilon V_{trap} \quad (3.21)$$

where $V_{TES} = l_{TES} h_{TES} w_{TES}$ is the volume of the conductive length of the TES. The epsilon factor here represents the fact that the trap only contributes a fraction of its volume to the heat capacity, as a portion which is in close contact with the aluminum has its effective heat capacity greatly reduced. This can also be written as

$$V_{TES,eff} = V_{TES} \left(1 + \epsilon \frac{V_{trap}}{V_{TES}}\right) = \frac{V_{TES}}{\zeta} \quad (3.22)$$

where

$$\zeta = \frac{V_{TES}}{\epsilon V_{trap} + V_{TES}} \quad (3.23)$$

is the TES fractional volume coefficient. $\zeta = 1$ implies that only the conductive length of the TES contributes to the heat capacity, while low ζ implies that the trap volume is more important to the heat capacity.

We also recognize that heat capacity is a function of temperature, and parameterize it similarly to how the resistance change was parameterized:

$$C(T) = C_{T_c} + \left. \frac{\partial C(T, I)}{\partial T} \right|_{I_0, T_0} \delta T = C_{T_c} + C_{T_c} \gamma_C \frac{\delta T}{T_c} \quad (3.24)$$

where γ_C is the logarithmic derivative quantity

$$\gamma_C = \frac{T_c}{C_{T_c}} \left. \frac{\partial C(T, I)}{\partial T} \right|_{I_0, T_0} \quad (3.25)$$

which describes the relative change in heat capacity with change in temperature. Ref [66] suggests that for super-conductors, this has a value around 3. We are prevented from simply expanding about T in the equation for the steady-state capacitance because f_{sc} is also a function of temperature, so γ mainly captures this dependence rather than scaling with the original constants.

As an aside, the fact that heat capacity changes dramatically is actually a primary prediction of BCS superconductivity, and the BCS prediction is that the heat capacity will jump discontinuously between the normal metal value C_n and an increased heat capacity $C_s = 2.43C_n$, after which is expected to fall rapidly to 0 as an exponential ($C_s \sim$ [13, 10]). The form for the superconducting heat capacity for temperatures lower than T_c follows $C_s(T) \sim 8.5 \exp(-1.44T_c/T)$ [13], and is linear close to the transition (as shown in Figure 3.3). BCS specifies, however, that this transition is infinitely sharp, but we know that proximity effects will produce a smooth transition between normal and superconducting states, so we rely on empirical means to determine the rate of change of heat

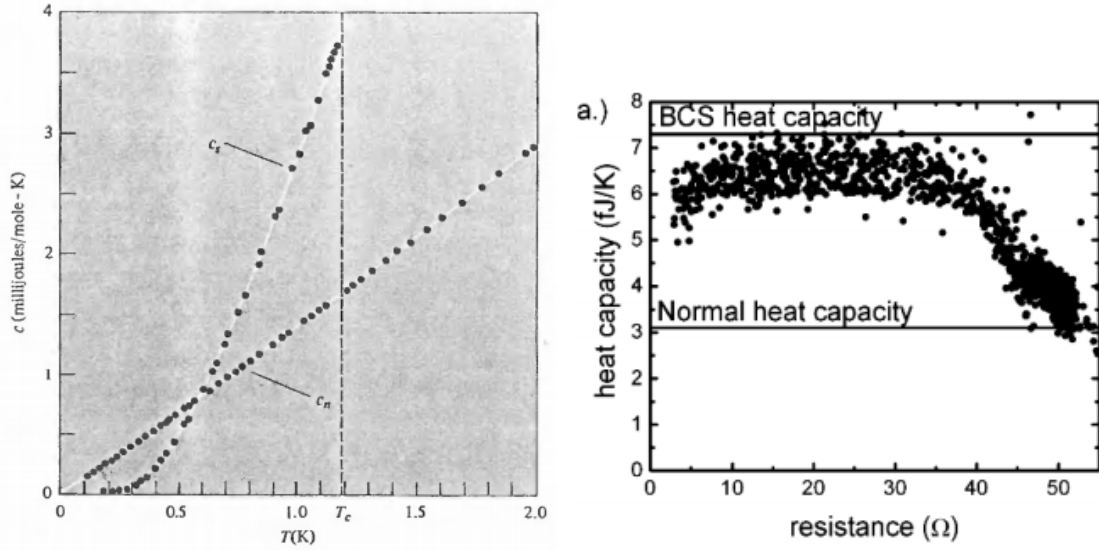


Figure 3.3: Left: Measurement of heat capacity above/below the transition temperature in a sample of bulk Al, reproduced from Ref [10]. The data points are an excellent fit to an exponential in the superconducting state and a line in the normal state, as discussed in the text. Right: Measurement of heat capacity as a function of resistance in the transition for a TES with normal resistance of $R_n \sim$, reproduced from Ref [67]. This shows that for a TES-sized structure, there is a non-trivial change in heat capacity between the expected normal and superconducting values, and the slope of that change is dependent on the operating point in the transition.

capacity through the transition. An example measurement of heat capacity versus resistance is also shown in Figure 3.3 reproduced from Ref [67]. This demonstrates the non-linear transition from the normal to superconducting heat capacity. The first measure of heat capacity in a SuperCDMS TES will be discussed in Chapter 5.

Getting back to the response model, we write the power transferred between the TES and the bath as

$$P_{bath} = K(T^n - T_{bath}^n) \approx KT^n \quad (3.26)$$

where K is decomposed into the TES volume and the Electron-Phonon coupling constants Σ :

$$K = V_{TES,eff}\Sigma \rightarrow P_{bath} \approx \Sigma V_{TES,eff}T^n \quad (3.27)$$

Here Σ is used as measured by Ref [45], however this is one of our most uncertain parameters, and will need to be re-measured for new devices to decrease the uncertainty.

3.2.2 TES Response Model

We can derive the green's function impulse response of a TES to determine how TES parameters, and the components of the bias circuit, affect the TES bandwidth, and determine the stability conditions for TES parameters in terms of fixed circuit constraints. Many circuit parameters are fixed, and we need to design the TES to optimize dynamics in light of their values. We can simplify the circuit part of the problem significantly by reducing the number of circuit elements, and the circuit complexity. In the process we'll end up with a model defined only in terms of realistic inputs and measurable outputs.

We voltage bias the TES circuit indirectly by supplying a stable current source across two parallel branches. Each branch can be considered to have a complex impedance consisting of a resistor and inductor (they should have negligible capacitance), with the TES branch also routed through the TES itself. The ideal circuit has only the shunt resistor in the left branch, which we could refer to as the bias branch, and only the TES and SQUID input coil in the right branch, the signal branch.

In reality, there is parasitic resistance and inductance in both branches. As we can simply add inductances and resistances in series, we can start from the model shown in the center of figure [3.4](#) with our only assumption being that $R_{bias} \gg Z_{loop}$. Given that the inductances are on the order of nH, and resistances on the order of m Ω , this is a valid assumption over all frequencies of interest (< 10 MHz, at which point a 1 μ H inductor as a reactance of 100 Ω), and we won't consider it further.

To get from the central to the right diagram in figure [3.4](#) we just need to do some algebraic manipulation of the loop equation for the central circuit. The current around the loop follows the equation

$$\left(R_{sb} + L_{bias} \frac{d}{dt} \right) I_{sh} = \left(R_p + L_{TES} \frac{d}{dt} \right) I_s + R_{TES}(I, T) I_s \quad (3.28)$$

where $I_{sh} + I_s = I_b$, and I_b is the current from the current source. This equation also allows us to use the relation $I_{sh} = I_b - I_s$ to get

$$\left(R_{sb} + L_{bias} \frac{d}{dt} \right) (I_b - I_s) = \left(R_p + L_{TES} \frac{d}{dt} \right) I_s + R_{TES}(I, T) I_s \quad (3.29)$$

$$\left(R_{sb} + L_{bias} \frac{d}{dt} \right) I_b = \left(R_p + L_{TES} \frac{d}{dt} \right) I_s + \left(R_{sb} + L_{bias} \frac{d}{dt} \right) I_s + R_{TES}(I, T) I_s \quad (3.30)$$

$$\left(R_{sb} + L_{bias} \frac{d}{dt} \right) I_b = \left(R_{loop} + L_{loop} \frac{d}{dt} \right) I_s + R_{TES}(I, T) I_s \quad (3.31)$$

At this point we can just read off the DC relationship between the bias and TES currents, and we can see why this is not equivalent to the dI_b/dI_s relationship we'll discuss later; the TES has a complex impedance which is not captured here and that will be derived shortly. But without knowing anything about the TES, we can use this relationship, with very slowly changing I_b and I_s ,

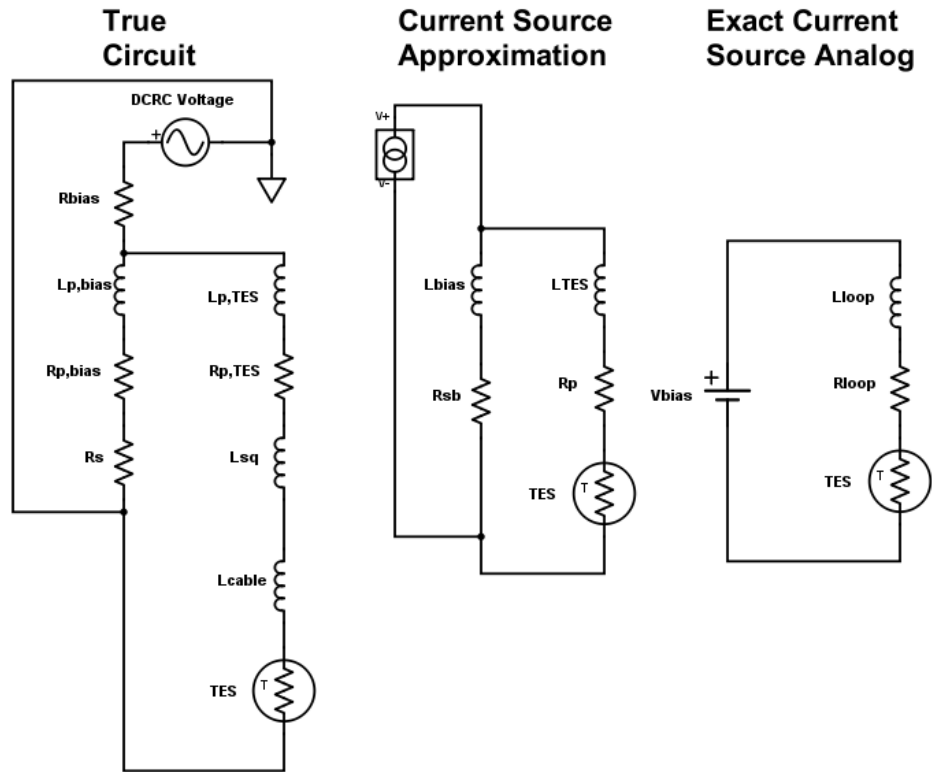


Figure 3.4: TES circuit analogs. Left is the actual circuit, arranged top to bottom by temperature. Middle is a very close approximation, assuming $R_{bias} \gg Z_{loop}$, which is true for $1 M\Omega$ bias resistor in the SNOLAB electronics. Right is exactly equivalent to the middle circuit; resistance and inductance are moved to the right-hand side, and an effective impedance is used to calculate the voltage supplied by the voltage source.

to show that

$$\frac{I_b}{I_s} = \frac{R_{loop} + R_{TES}(I, T)_{DC}}{R_{sb}} = 1 + \frac{R_p + R_{TES}(I, T)_{DC}}{R_{sb}} \quad (3.32)$$

which will allow us to measure the TES transition curve without knowing anything about its frequency response.

This also tells us that to an excellent approximation, we can write

$$V = \left(R_{sb} + L_{bias} \frac{d}{dt} \right) I_b = Z_{sb} I_b \approx R_{sb} I_b \quad (3.33)$$

which will be valid as long as I_b is well controlled, or we meet the condition

$$\left| \frac{Z_{sb}}{R_{sb}} \right| - 1 < \epsilon \quad (3.34)$$

$$\left| \frac{(R_{sb} + i\omega L_{bias})}{R_{sb}} \right| - 1 < \epsilon \quad (3.35)$$

$$|1 + i\omega L_{bias}/R_{sb}| < 1 + \epsilon \quad (3.36)$$

$$1 + \frac{1}{2} (\omega L_{bias}/R_{sb})^2 < 1 + \epsilon \quad (3.37)$$

$$\omega L_{bias}/R_{sb} < \sqrt{2\epsilon} \quad (3.38)$$

$$f < \frac{\sqrt{2\epsilon}}{2\pi} \frac{R_{sb}}{L_{bias}} \quad (3.39)$$

For SNOLAB electronics, to 10% precision, this is a good approximation below frequencies of 100 kHz. If we weren't using a stable power source, this would be a problem, but given that the high frequency components should be much smaller than the DC value, we can ignore inductor effects in the bias voltage until we discuss measurements of complex impedance.

With our simplified circuit, we can now begin the modeling including thermal TES response. We begin by recognizing that, for a TES cooled by a bath with power P_{bath} , heated by Joule heating P_J , and subject to some signal power P , we have the differential equation

$$C \frac{dT}{dt} = -P_{bath} + P_J + P \quad (3.40)$$

and electrically, our circuit obeys the differential equation

$$L \frac{dI}{dt} = V - IR_L - IR(T, I) \quad (3.41)$$

where L is the self-inductance of our squid input coil, R_L is the resistance of the input coil, and $R(T, I)$ is the TES resistance as a function of I and T . Before going further, we see that we'll have some highly non-linear terms here, so to obtain any solution we'll have to approximately linearize a few of these terms.

The bath power is typically written [50](#)

$$P_{bath} = K(T^n - T_{bath}^n) \approx KT^n = \Sigma V_{TES,eff} T^n \quad (3.42)$$

for the normal case that T_{bath} is even 5 mK colder than the TES. We can thus linearize power about some nominal temperature T_0 (which will also be roughly T_c) to be

$$P_{bath} \approx P_{bath,0} + n\Sigma V_{TES,eff} T^{n-1} \delta T = P_{bath,0} + G\delta T \quad (3.43)$$

so we have

$$G = nKT^{n-1} = n\Sigma V_{TES,eff} T^{n-1} \quad (3.44)$$

where G is the differential thermal conductance of the TES/QET system.

Linearizing the resistance about nominal values of I_0 and T_0 as in the previous section gives

$$R(I, T) \approx R_0 + \alpha \frac{R_0}{T_0} \delta T + \beta \frac{R_0}{I_0} \delta I \quad (3.45)$$

which allows us to fully linearize our differential equations. The joule heating of the TES at a steady state current I_0 and operating resistance R_0 is then given by

$$P_J = P_{J0} = I_0^2 R_0 \quad (3.46)$$

and the full linearization in R and I is thus

$$P_J \approx P_{J0} + 2I_0 R_0 \delta I + I_0^2 \left(\alpha \frac{R_0}{T_0} \delta T + \beta \frac{R_0}{I_0} \delta I \right) \quad (3.47)$$

$$= P_{J0} + I_0 R_0 (2 + \beta) \delta I + \alpha \frac{P_{J0}}{T_0} \delta T \quad (3.48)$$

We thus have the power differential equation

$$C(T) \frac{d\delta T}{dt} = I_0 R_0 (2 + \beta) \delta I - \left(G - \alpha \frac{I_0^2 R_0}{T_0} \right) \delta T + \delta P \quad (3.49)$$

where δT has come in as we assume equilibrium at some T_0 , giving $I_0^2 R_0 = P_{bath0} + P_0$ and allowing them to cancel. If we also consider the linearization in heat capacity, we find for constant R (before

substitution of specific forms):

$$\frac{dT}{dt} = \frac{1}{C(T)}(P_{J0} - P_{bath0}) \quad (3.50)$$

$$\approx \frac{1}{C_{T_c}}(P_{J0} - P_{bath0})\left(1 - \gamma \frac{\delta T}{T_c}\right) \quad (3.51)$$

$$C_{T_c} \frac{d\delta T}{dt} = -\gamma \frac{\delta T}{T_c}(P_{J0} - P_{bath0}) \quad (3.52)$$

$$= -\gamma \frac{\delta T}{T_c} \left(I_0^2 R_0 - \frac{GT_c}{n} \right) \quad (3.53)$$

$$= \gamma \frac{G}{n} \delta T - \gamma I_0^2 R_0 \frac{\delta T}{T_c} \quad (3.54)$$

which tells us that the change in heat capacity counter-acts the normal ETF effects from the previous equation. We can include this capacitance change by the substitutions of effective values for G and α_I :

$$G_{eff} = G \left(1 - \frac{\gamma}{n} \right) \quad (3.55)$$

and

$$\alpha_{eff} = \alpha_I - \gamma \quad (3.56)$$

which agrees with the results found in Ref [66], but has not been included in past attempts to model TES dynamics by CDMS. We note that this does not greatly affect the α parameter, but will significantly change the effective conductance, and create a longer time constant in our TES, which is something we'd like to be as short as possible.

Often, the dimensionless parameter \mathcal{L} is introduced to simplify these equations:

$$\mathcal{L} = \frac{\alpha I_0^2 R_0}{GT_0} = \frac{\alpha}{n} \left(1 - \frac{T_{bath}^n}{T_c^n} \right) \quad (3.57)$$

This allows us also to introduce the natural TES time constant

$$\tau = \frac{C}{G} \quad (3.58)$$

and we can then simplify to find

$$\frac{d\delta T}{dt} = \frac{I_0 R_0 (2 + \beta)}{C} \delta I - \frac{1}{\tau} (1 - \mathcal{L}) \delta T + \frac{\delta P}{C} \quad (3.59)$$

This tells us that, with the C and G derived earlier,

$$\tau = \frac{f_{sc} C W}{n \Sigma} T_c^{2-n} \quad (3.60)$$

and the TES cooling timescale is thus independent of geometry, and just dependent on intrinsic

parameters. Including the capacitance effects just mentioned, we find

$$\tau_{eff} = \frac{C}{G(1 - \frac{\gamma}{n})} = \frac{f_{sc} c_W}{(n - \gamma) \Sigma} T_c^{2-n} \quad (3.61)$$

and we leave \mathcal{L} unaffected as we assume $\alpha \gg \gamma$. In a perfect superconductor we would expect $\gamma \sim -3$, but for our devices, which have non-negligible currents running through them, we expect this slope to be less dramatic. Setting γ to zero allows us to put an upper limit on the TES response time, and a non-zero value will thus increase the bandwidth of a TES.

Moving now to the current equation, we find for steady state I_0 at V_0, R_0, T_0 , we have

$$\frac{d\delta I}{dt} = \frac{\delta V}{L} - \frac{(R_L + (1 + \beta)R_0)}{L} \delta I - \frac{\alpha R_0 I_0}{T_0 L} \delta T \quad (3.62)$$

which, again substituting \mathcal{L} , gives the final coupled differential equations

$$\frac{d\delta T}{dt} = \frac{I_0 R_0 (2 + \beta)}{C} \delta I - \frac{1}{\tau} (1 - \mathcal{L}) \delta T + \frac{\delta P}{C} \quad (3.63)$$

$$\frac{d\delta I}{dt} = \frac{\delta V}{L} - \frac{R_L + (1 + \beta)R_0}{L} \delta I - \frac{G\mathcal{L}}{I_0 L} \delta T \quad (3.64)$$

Here we find, in the decoupled limit, the natural decay times

$$\tau_{LR} = \frac{L}{R_L + (1 + \beta)R_0} \quad (3.65)$$

$$\tau_{TES} = \frac{\tau}{1 - \mathcal{L}} \quad (3.66)$$

which tell us that there are timescales determined by the normal LR circuit behavior, where the TES appears to have a resistance $R_{eff} = (1 + \beta)R_0$.

This also determines for us how we need to bias our TES to optimize performance. For semiconductor thermistors, α was negative, and thus current bias was chosen to decrease decay time and increase bandwidth. For our positive α devices, we thus need voltage bias for the sign of the decay constant to not create a positive feedback effect.

These parameters allow us to write these coupled equations in a tidy matrix form:

$$\frac{d}{dt} \begin{bmatrix} \delta T \\ \delta I \end{bmatrix} = \begin{bmatrix} -\tau_{TES}^{-1} & \frac{I_0 R_0 (2 + \beta)}{C} \\ -\frac{G\mathcal{L}}{I_0 L} & -\tau_{LR}^{-1} \end{bmatrix} \begin{bmatrix} \delta T \\ \delta I \end{bmatrix} + \begin{bmatrix} \frac{\delta P}{C} \\ \frac{\delta V}{L} \end{bmatrix} \quad (3.67)$$

This suggests solutions in either the fourier or time basis, both of which can be found in Ref [50]. Here I will discuss the time basis solution, as it is relevant to the later TES simulation and helps determine our stability criterion.

Fourier Basis Solution: Frequency Response

We can easily solve this in the fourier basis, for $\delta V = 0$. We obtain the solution

$$i\omega \begin{bmatrix} \delta T \\ \delta I \end{bmatrix} = \begin{bmatrix} -\tau_{TES}^{-1} & \frac{I_0 R_0 (2+\beta)}{C} \\ -\frac{G\mathcal{L}}{I_0 L} & -\tau_{LR}^{-1} \end{bmatrix} \begin{bmatrix} \delta T \\ \delta I \end{bmatrix} + \begin{bmatrix} \frac{\delta P}{C} \\ 0 \end{bmatrix} \quad (3.68)$$

$$\begin{bmatrix} \frac{\delta P}{C} \\ 0 \end{bmatrix} = \begin{bmatrix} \tau_{TES}^{-1} + i\omega & -\frac{I_0 R_0 (2+\beta)}{C} \\ \frac{G\mathcal{L}}{I_0 L} & \tau_{LR}^{-1} + i\omega \end{bmatrix} \begin{bmatrix} \delta T \\ \delta I \end{bmatrix} \quad (3.69)$$

$$(3.70)$$

which gives

$$\delta P = \left[C(\omega_{TES} + i\omega) \frac{\delta T}{\delta I} - I_0 R_0 (2 + \beta) \right] \delta I \quad (3.71)$$

and

$$\frac{\delta T}{\delta I} = -\frac{(\omega_{LR} + i\omega) I_0 L}{G\mathcal{L}} \quad (3.72)$$

so substitution gives

$$\frac{\delta I}{\delta P} = -\frac{1}{I_0} \left(\frac{(\omega_{TES} + i\omega)(\omega_{LR} + i\omega) C L}{G\mathcal{L}} + R_0 (2 + \beta) \right)^{-1} \quad (3.73)$$

We can substitute back in for the τ terms to find

$$\frac{\delta I}{\delta P} = -\frac{\mathcal{L}}{I_0} \frac{1}{(1 - \mathcal{L} + i\omega\tau)(R_L + (1 + \beta)R_0 + i\omega L) + \mathcal{L}R_0(2 + \beta)} \quad (3.74)$$

To get some intuition, in the limit that $\mathcal{L} \gg 1$, $\beta = 0$, and $R_0 \gg R_L$, we find the simpler function

$$\frac{\delta I}{\delta P} = -\frac{\mathcal{L}}{I_0} \frac{1}{(i\omega\tau - \mathcal{L})(R_0 + i\omega L) + 2\mathcal{L}R_0} \quad (3.75)$$

and we have the DC gain

$$\frac{\delta I}{\delta P}(\omega \rightarrow 0) = \frac{1}{I_0 R_0} = V_0^{-1} \quad (3.76)$$

which we could have guessed, given that an resistor with instantaneous response has

$$P = IV \rightarrow \frac{\delta P}{\delta I} = V_0 \quad (3.77)$$

and thus lowering the bias current or the operating resistance will raise the DC gain. Unfortunately, we'll see later that these will also increase the current noise, so we'll need a full analysis including dominant noise to determine how best to set these values.

We can by a similar method obtain the transfer function for voltage to current fluctuations. We

obtain the solution

$$\begin{bmatrix} 0 \\ \frac{\delta V}{L} \end{bmatrix} = \begin{bmatrix} \tau_{TES}^{-1} + i\omega & -\frac{I_0 R_0 (2 + \beta)}{C} \\ \frac{G\mathcal{L}}{I_0 L} & \tau_{LR}^{-1} + i\omega \end{bmatrix} \begin{bmatrix} \delta T \\ \delta I \end{bmatrix} \quad (3.78)$$

$$(3.79)$$

which gives

$$\delta T = \frac{I_0 R_0 (2 + \beta)}{C(\omega_{TES} + i\omega)} \delta I \quad (3.80)$$

and thus

$$\delta V = \left(L(\omega_{LR} + i\omega) + \frac{G\mathcal{L}R_0(2 + \beta)}{C(\omega_{TES} + i\omega)} \right) \delta I \quad (3.81)$$

and inversion gives

$$\frac{\delta I}{\delta V} = \frac{(\omega_{TES} + i\omega)}{L(\omega_{LR} + i\omega)(\omega_{TES} + i\omega) + \tau^{-1}\mathcal{L}R_0(2 + \beta)} \quad (3.82)$$

which simplifies further to

$$\frac{\delta I}{\delta V} = \frac{1 - \mathcal{L} + i\omega\tau}{(R_L + R_0(1 + \beta) + i\omega L)(1 - \mathcal{L} + i\omega\tau) + \mathcal{L}R_0(2 + \beta)} \quad (3.83)$$

which in the DC limit, taking the same limiting cases again, we find to be the expected value of R_0^{-1} . We can complete the triangle of power, voltage, and current linearization:

$$\frac{\partial P}{\partial V} = \frac{\partial P}{\partial I} \frac{\partial I}{\partial V} = I_0(1 - \mathcal{L}^{-1})(1 + i\omega\tau_{TES})$$

and we see the inductor pole go away. This makes some intuitive sense, as the inductor does not store voltage or contribute voltage power, and thus would not factor into the voltage to power conversion. What is also interesting to note is that high frequency voltage fluctuations actually couple into the TES power more, which is due to the electrothermal feedback pole.

We further note, however, that due to the electrothermal feedback, we need to modify this derivation to have $\delta P = -I_0\delta V$ for fluctuations in the TES voltage, which tells us that for the TES, we have

$$\frac{\delta I}{\delta V_{TES}} = \frac{1 - \mathcal{L} + \mathcal{L} + i\omega\tau}{(R_L + R_0(1 + \beta) + i\omega L)(1 - \mathcal{L} + i\omega\tau) + \mathcal{L}R_0(2 + \beta)} \quad (3.84)$$

$$= \frac{1 + i\omega\tau}{(R_L + R_0(1 + \beta) + i\omega L)(1 - \mathcal{L} + i\omega\tau) + \mathcal{L}R_0(2 + \beta)} \quad (3.85)$$

thus the feedback allows the zero to be much more powerful in the total transfer function, damping low-frequency voltage signals.

Finally, if we defined the complex impedance of the TES in terms of the complex impedance of

the circuit:

$$Z_{loop}(\omega) = \frac{\delta V}{\delta I} = i\omega L + R_L + Z_{TES}(\omega) \quad (3.86)$$

In the absence of other circuit elements, we find the complex impedance of the TES to be

$$Z_{TES}(\omega) = \frac{(R_L + R_0(1 + \beta) + i\omega L)(1 - \mathcal{L} + i\omega\tau) + \mathcal{L}R_0(2 + \beta)}{1 - \mathcal{L} + i\omega\tau} - i\omega L - R_L \quad (3.87)$$

$$= \frac{R_0(1 + \beta)(1 - \mathcal{L} + i\omega\tau) + \mathcal{L}R_0(2 + \beta)}{1 - \mathcal{L} + i\omega\tau} \quad (3.88)$$

$$= \frac{R_0(1 + \beta)(1 + i\omega\tau) + \mathcal{L}R_0}{1 - \mathcal{L} + i\omega\tau} \quad (3.89)$$

$$= R_0 \frac{(1 + \beta)(1 + i\omega\tau) + \mathcal{L}}{1 - \mathcal{L} + i\omega\tau} \quad (3.90)$$

This suggests that we should be able to measure both \mathcal{L} and β by applying voltage bias excitations much higher than and much lower than the Pole, where the high frequency measurement gives β , and the combination gives \mathcal{L} .

Time Basis Solution: Green's Function Response

We can solve the homogeneous equation ($\delta P = \delta V = 0$) by finding the eigenvalues and eigenvectors of the matrix; for eigenvalues λ_{\pm} and eigenvectors Λ_{\pm} , we find the general solution

$$\begin{bmatrix} \delta T \\ \delta I \end{bmatrix} = A_+ \exp(-t\lambda_+) \Lambda_+ + A_- \exp(-t\lambda_-) \Lambda_- \quad (3.91)$$

and we can thus recognize that $\tau_{\pm} = \lambda_{\pm}^{-1}$. Using any method, one finds the eigenvalues

$$\frac{1}{\tau_{\pm}} = \frac{1}{2\tau_{LR}} + \frac{1}{2\tau_{TES}} \quad (3.92)$$

$$\pm \frac{1}{2} \sqrt{\left(\frac{1}{\tau_{LR}} - \frac{1}{\tau_{TES}}\right)^2 - 4\frac{R_0}{L} \frac{\mathcal{L}(2 + \beta)}{\tau}} \quad (3.93)$$

$$= \frac{1}{2} \left(\frac{1 \pm \gamma}{\tau_{LR}} + \frac{1 \mp \gamma}{\tau_{TES}} \right) \quad (3.94)$$

where

$$\gamma = \left(1 - 4\frac{R_0}{L} \frac{\mathcal{L}(2 + \beta)}{\tau} \left(\frac{1}{\tau_{LR}} - \frac{1}{\tau_{TES}} \right)^{-2} \right)^{1/2} \quad (3.95)$$

For this to be non-oscillating, then, we need γ to be real, and it is thus in the range (0,1), making it a correction factor. This will be expanded upon in the next subsection.

If we assume ideal operation, where, $\delta V=0$, and phonons with energy E couple into the system

as a delta function:

$$\delta P(t) = E_{phonon}\delta(t), \quad \delta I(0) = 0, \quad \delta T(0) = 0 \quad (3.96)$$

we can solve for the unknown parameters A_{\pm} to obtain the green's function response to a single phonon of energy E . Ref [50] obtains the result ($\gamma < 1$)

$$\delta I(t) = C_N(\tau_i) \frac{E_{ph}}{(2 + \beta)I_0 R_0} \left(e^{-t/\tau_+} - e^{-t/\tau_-} \right) \quad (3.97)$$

where

$$C_N(\tau_i) = \frac{(\tau_{TES}/\tau_+ - 1)(\tau_{TES}/\tau_- - 1)}{\tau_{TES}^2 (1/\tau_+ - 1/\tau_-)} \quad (3.98)$$

is the normalization constant, with units of s^{-1} . Note also that we can calculate the desired operating current as

$$I_0 = \sqrt{\frac{\Sigma V_{TES,eff} T_c^n}{R_0}} \quad (3.99)$$

from the equilibrium condition between the bath and joule heating.

With this green's function in hand, we can recognize that for the limiting cases where $\gamma \sim 0$ and $\gamma = 1$, we have varying levels of over-damping, with $\gamma = 0$ being the critically damped solution, with a slightly differing green's function (the limit must be evaluated, as the numerator and denominator are nominally zero). For $\gamma = 1$, we find that the rise time becomes the LR transient time-scale, and the fall time becomes the TES cooling time, as one might expect.

Poles and Time Constants

We should take a second to stop and summarize the various poles and time constants affecting each of the quantities derived above, and quote values corresponding to the constants given for our proposed TES geometry and known device constants.

The time domain time constants (which one can check match the equivalent Fourier domain poles in both Matt Pyle's and Paul Brink's derivations) are given in equation [3.94]. The constants therein, and in particular those within γ , given in equation [3.95] were defined in the previous sections, but I will summarize them again here.

τ , the intrinsic response time of the TES without electro-thermal feedback, is given by equation [3.60]

$$\tau = \frac{C}{G} = \frac{f_{sc} c_W}{n \Sigma} T_c^{2-n} = \frac{f_{sc} c_W}{5 \Sigma T_c^3} \approx 592 * f_{sc} \mu s \quad (3.100)$$

using the quantities shown in Table [3.1]. Here, f_{sc} is left explicit to show that this is a minimum fall time, and stress the high variability of this quantity within the transition between 1.0 and ~ 2.5 when the TES is nearly superconducting. It is more straightforward from this point to quote rise/fall times at the start of the transition, as this is better defined, but it is important to keep in mind that this is shorter than the true fall time. Setting $f_{sc} = 1.0$ thus gives $\tau \approx 592 \mu s$, while $f_{sc} \sim 2.5$ gives

$\tau \approx 1.48ms$.

The rest of the time constants are derived from τ , aside from the LR time constant

$$\tau_{LR} = \frac{L}{R_L + (1 + \beta)R_0} \approx 2 \mu s \quad (3.101)$$

We have the electro-thermal feedback time constant

$$\tau_{TES} = \frac{\tau}{1 - \mathcal{L}} \approx -24.0 f_{sc} \mu s \quad (3.102)$$

which again scales with the position in the super-conducting transition. The rise and fall times have a much more complicated scaling given γ , which we can also express in terms of this scaling constant:

$$\gamma^2 = 1 - 4 \frac{R_0}{L} \frac{\mathcal{L}(2 + \beta)}{\tau} \left(\frac{1}{\tau_{LR}} - \frac{1}{\tau_{TES}} \right)^{-2} \quad (3.103)$$

$$= 1 - 4 \frac{R_0}{L} \frac{\mathcal{L}(2 + \beta)\tau_{LR}^2}{\tau} \left(1 - \frac{\tau_{LR}}{\tau_{TES}} \right)^{-2} \quad (3.104)$$

$$\approx 1 - 4 \frac{R_0}{L} \frac{\mathcal{L}(2 + \beta)\tau_{LR}^2}{\tau} \left(1 + 2 \frac{\tau_{LR}}{\tau_{TES}} \right) \quad (3.105)$$

$$\approx 1 - \frac{0.578}{f_{sc}} \left(1 - \frac{0.169}{f_{sc}} \right) \quad (3.106)$$

This shows us that increasing the heat capacity through the transition causes an increase in γ , which drives rise time to be smaller and fall time to be longer. The limit of much larger heat capacity corresponds to the diagonal basis, where rise time is the LR time constant and fall time is the TES time constant. Figure 3.5 shows the rise and fall time as a function of f_{sc} , compared to the TES and LR time constants.

We'll see in the noise section that these time constants are relevant as poles for the TES and electronic (non-squid) noise terms, so we should also find the frequencies associated with these time constants. Simply inverting the time constants should give a measure of the angular frequency, so the frequencies of interest are

$$f = \frac{\omega}{2\pi} = \frac{1}{2\pi\tau} \quad (3.107)$$

The time constants and poles for $f_{sc} = 1.0$ and $f_{sc} = 2.5$ are summarized in table 3.2

3.2.3 Electrothermal Oscillation Criteria

From the above, we see that for critically or under-damped behavior, we have the condition

$$1 \geq 4 \frac{R_0}{L} \frac{\mathcal{L}(2 + \beta)}{\tau} \left(\frac{1}{\tau_{LR}} - \frac{1}{\tau_{TES}} \right)^{-2} \quad (3.108)$$




Figure 3.5: Rise and Fall time as a function of the ratio of superconducting to normal heat capacity, as described in the section. Note that both rise and fall time are larger than their asymptotic values, and that fall time gets larger linearly with f_{sc} while rise time gets smaller with the inverse of f_{sc} , to first order.

Quantity	$\tau_{1.0}$ (μs)	$\omega_{1.0}$	$\tau_{2.5}$ (μs)	$\omega_{2.5}$
τ	593	268 Hz	1481	107 Hz
LR	2.0	79.6 kHz	2.0	79.6 kHz
TES (abs)	23.7	6.7 kHz	59.1	2.7 kHz
Rise	2.4	67.2 kHz	2.1	74.9 kHz
Fall	28.2	5.6 kHz	78.5	2.0 kHz

Table 3.2: TES Rise and Fall time constants

which expands, without taking any explicit limits, to become

$$\left(\frac{1}{\tau_{LR}} - \frac{1}{\tau_{TES}} \right)^2 - \frac{4\mathcal{L}(2+\beta)R_0}{\tau L} \geq 0 \quad (3.109)$$

$$\left(\frac{R_0(1+\beta) + R_L}{L} + \frac{\mathcal{L}-1}{\tau} \right)^2 - \frac{4\mathcal{L}(2+\beta)R_0}{\tau L} \geq 0 \quad (3.110)$$

This is a quadratic equation, which when solved for the loop gain yields a range of values where oscillation will occur. The region of oscillation is bound by the loop-gain values

$$(\mathcal{L}_{\pm} - 1) = \frac{\tau R_0}{L} \left[3 + \beta - \frac{R_L}{R_0} \pm 2\sqrt{(2+\beta) \left(1 + \frac{L}{R_0\tau} - \frac{R_L}{R_0} \right)} \right] \quad (3.111)$$

$$= \frac{R_0}{R_{\tau}} \left[3 + \beta - \frac{R_L}{R_0} \pm 2\sqrt{(2+\beta) \left(1 + \frac{R_{\tau} - R_L}{R_0} \right)} \right] \quad (3.112)$$

where I've introduced the derived parameter

$$R_{\tau} = \frac{L}{\tau} \quad (3.113)$$

to put all of the units into easily comparable units. For the SNOLAB design, we have $R_\tau \approx 250nH/600\mu s \approx 4 \text{ m}\Omega$, and for a design operating point around $50 \text{ m}\Omega$ this is obviously sub-dominant. If we create a device with a much higher T_c , as was done for the first UMN test devices in older electronics, however, $R_\tau \approx 400nH/80\mu s \approx 25 \text{ m}\Omega$, and we have a much lower oscillation threshold. In addition, older electronics have higher values of R_L , further reducing the location of the oscillation region.

One quantity of interest might be the minimum resistance where oscillation will be observed. We see that we can find this without much work, as it will occur when the last factor in the earlier equations is 0:

$$0 \leq 2\sqrt{(2 + \beta) \left(1 + \frac{R_\tau - R_L}{R_0}\right)} \quad (3.114)$$

$$R_0 \geq R_L - R_\tau \quad (3.115)$$

Given that all of these numbers are positive, this tells us something important; that it is not possible to eliminate the possibility of oscillation in any design where the TES is voltage biased if inductance and loop gain cannot be controlled. It will never be possible to have a stable voltage biased region free of oscillation for any set of temperature response and inductance.

In the past, we've been used to considering the lower bound an upper limit and using the loop gain to put restrictions on inductance, but given measured properties of an electronics setup, we can determine the range of transition slopes which will produce oscillation. In the limiting case normally applicable to our detectors ($\mathcal{L} \gg 1$, $R_\tau, R_L \ll R_0$), this tells us that we avoid oscillation for

$$\mathcal{L} \leq \frac{R_0\tau}{L} \left(3 + \beta - 2\sqrt{2 + \beta}\right) = \frac{R_0}{R_\tau} \left(3 + \beta - 2\sqrt{2 + \beta}\right) \quad (3.116)$$

assuming $\beta \geq 0$. We benefit from increasing β because it decreases the L/R time constant, thus allowing more bandwidth for a given inductor-TES combination, and increases the allowance of inductance before the TES begins to oscillate. This is the bound derived in Matt Pyle's thesis. It's clear that for older devices this is a good limit, but as we move to lower resistance and T_c , we need to be careful not to take this limit without explicit evidence that this is the case.

We'll see later on that we'd like the TES fall time to be much shorter than the timescale of a phonon pulse. Since these dynamics can shorten the overall fall-time, let's say that we set a minimum fall time τ_{min} of

$$\tau_{min} = \frac{\tau}{c} = -\frac{\tau_{TES}}{c} \quad (3.117)$$

This gives the condition

$$\frac{1}{\tau_{LR}} + \frac{1 + 2c}{\tau_{TES}} \geq \sqrt{\left(\frac{1}{\tau_{LR}} - \frac{1}{\tau_{TES}}\right)^2 - 4\frac{R_0}{L} \frac{\mathcal{L}(2 + \beta)}{\tau}} \quad (3.118)$$

which, solved, gives

$$-4 \frac{1+c}{\tau_{LR}\tau_{TES}} \leq 4 \frac{R_0}{L} \frac{\mathcal{L}(2+\beta)}{\tau} + \frac{4c(1+c)}{\tau_{TES}^2} \quad (3.119)$$

$$(1+c) \frac{\mathcal{L}(R_0(1+\beta) + R_L)}{L\tau} \leq \frac{R_0}{L} \frac{\mathcal{L}(2+\beta)}{\tau} + \frac{\mathcal{L}^2 c(1+c)}{\tau^2} \quad (3.120)$$

$$(1+c)(R_0(1+\beta) + R_L) \leq R_0(2+\beta) + \frac{L\mathcal{L}c(1+c)}{\tau} \quad (3.121)$$

$$(1+c)R_L \left(1 - \frac{L\mathcal{L}c}{R_L\tau}\right) \leq R_0(1 - c(1+\beta)) \quad (3.122)$$

and thus

$$R_0 \geq R_L \frac{1+c}{1-c(1+\beta)} \left(1 - \frac{L\mathcal{L}c}{R_L\tau}\right) \quad (3.123)$$

In the limit that the second term on top is negligible (which we can certainly design by making L small enough), we find

$$R_0 \geq R_L \frac{1+c}{1-c(1+\beta)} \quad (3.124)$$

This tells us two things; firstly that the minimum fall time is the natural fall time of the TES under ETF, and secondly that non-zero beta increases the minimum fall-time above the natural TES falltime. These both come from the constraint that R_0 can't be infinite. We thus see that for us to achieve 1.2 times the natural TES fall-time, we need to have ;

$$R_0 \geq R_L \frac{1.83}{0.17 - 0.83\beta} \quad (3.125)$$

which for $\beta = 0$ is about a factor of 10, and for any value of R_0 to achieve this, we need $\beta < 0.2$. This shows that having a β as close to zero as possible is hugely advantageous. Even for a value of $\beta \sim 0.1$, we need $R_0 \approx 20R_L$ to achieve this bandwidth.

3.2.4 TES Stability

We also need to ensure that the TES is stable. The impulse response equation suggests that this will be the case as long as both τ_{\pm} are positive, which we can write as the inequality

$$0 \leq \text{Re} \left[\frac{2}{\tau_{\pm}} \right] = \text{Re} \left[(\tau_{LR}^{-1} + \tau_{TES}^{-1}) \pm \gamma (\tau_{LR}^{-1} - \tau_{TES}^{-1}) \right] \quad (3.126)$$

Here we recognize that γ is very limited in real space, and only takes values between 0 and 1. We can also see that for $\mathcal{L} > 1$, which is a very conservative criterion (a device which does not satisfy this condition would be useless for us), we see that the prefactor of γ in this expression is always

positive. The difference between falltimes is

$$\frac{1}{\tau_+} - \frac{1}{\tau_-} = 2\gamma \left(\frac{1}{\tau_{LR}} - \frac{1}{\tau_{TES}} \right) \geq 0 \quad (\mathcal{L} \geq 1) \quad (3.127)$$

$$\rightarrow \tau_+^{-1} \geq \tau_-^{-1} \quad (\mathcal{L} \geq 1) \quad (3.128)$$

This tells us that for a given τ_{LR} and τ_{TES} , under the condition that $\mathcal{L} > 1$, the fall-time will go negative first, and solving for this point will tell us the stability criteria for the TES.

We thus see that the TES is stable if the fall-time is positive; a negative fall-time is indicative of thermal runaway. This gives the condition (in the overdamped case)

$$\frac{1}{\tau_{LR}} + \frac{1}{\tau_{TES}} \geq \sqrt{\left(\frac{1}{\tau_{LR}} - \frac{1}{\tau_{TES}} \right)^2 - 4 \frac{R_0}{L} \frac{\mathcal{L}(2+\beta)}{\tau}} \quad (3.129)$$

squaring both sides and simplifying gives

$$-\frac{4}{\tau_{LR}\tau_{TES}} \leq 4 \frac{R_0}{L} \frac{\mathcal{L}(2+\beta)}{\tau} \quad (3.130)$$

$$\frac{(\mathcal{L}-1)(R_0(1+\beta) + R_L)}{L\tau} \leq \frac{R_0}{L} \frac{\mathcal{L}(2+\beta)}{\tau} \quad (3.131)$$

$$(\mathcal{L}-1)(R_0(1+\beta) + R_L) \leq \mathcal{L}R_0(2+\beta) \quad (3.132)$$

$$(\mathcal{L}-1)R_L \leq (\mathcal{L}+1+\beta)R_0 \quad (3.133)$$

$$R_0 \geq \frac{\mathcal{L}-1}{\mathcal{L}+\beta+1} R_L \quad (3.134)$$

and thus we have a properly decaying TES as long as the operating resistance is greater than the Thevenin equivalent resistance of the bias circuit. If \mathcal{L} is much larger than 1, we have the simpler condition that $R_0 \geq R_L$, and if \mathcal{L} is less than 1, the TES will always be stable.

For the underdamped case, we set the square root term to 0, and we find the bound:

$$\frac{1}{\tau_{LR}} + \frac{1}{\tau_{TES}} \geq 0 \quad (3.135)$$

$$\frac{R_0(1+\beta) + R_L}{L} \geq \frac{\mathcal{L}-1}{\tau} \quad (3.136)$$

$$R_0(1+\beta) \geq \frac{L(\mathcal{L}-1)}{\tau} - R_L \quad (3.137)$$

$$R_0 \geq \frac{R_\tau(\mathcal{L}-1) - R_L}{1+\beta} \quad (3.138)$$

The good news about both of these stability criteria is that they can be remedied beyond initial design to a certain point by increasing the bias current (and thus the bias voltage), giving R_0 a higher value, but to maximize resistance of the TES to either thermal runaway or oscillation, we

should design the TES with a large margin in mind for both of these conditions. We would also like to see that the fall time is not infinite, so a stricter condition might be to determine a minimum fall-time (given bandwidth considerations) and determine the resulting resistance and $\frac{L}{R_0}$ ratios.

3.2.5 Thermal Phase Separation

The maximum length of a TES is set by the critical temperature T_c of the device, which determines the phase separation length, or the length scale below which the TES remains in thermal equilibrium (and can thus be treated as a homogeneous device). We need to construct our TES as long as possible without causing them to phase separate. The maximum length before appreciable phase separation is given by the formula ([83], page 102):

$$l_{max} = \sqrt{\frac{\pi^2 \beta_{wf} \zeta T}{n \Sigma T_c^{n-2} \rho_w \left(\frac{\alpha}{n} \left(1 - \frac{T_{bath}^n}{T_c^n} \right) - 1 \right)}} \quad (3.139)$$

where most constants have been defined above. We gain one new constant, β_{wf} (Weidemann-Franz coefficient), which relates temperature, electrical conductivity σ and thermal conductivity κ as

$$\frac{\kappa}{\sigma} = \beta_{wf} T \quad (3.140)$$

This is also referred to as \mathcal{L} , the Lorenz number. The exact constant of proportionality at 0 temperature is known to be L_0 , the Lorenz number.




Figure 3.6: TES phase separation length as a function of critical temperature (left), where the dashed lines illustrate the effect of increasing the bath temperature, and as a function of transition sharpness and volume fraction (right), with width and thickness as set in Table 3.1

For $n = 5$, the ratio $\left(\frac{T_{bath}}{T_c} \right)^5$ is small (less than 10%) for deviations of 5% of T_{bath} below T_c , so if we assume $T_{bath} < 0.95T_c$ (which is in line with design of the SNOlab fridge below 30 mK) we

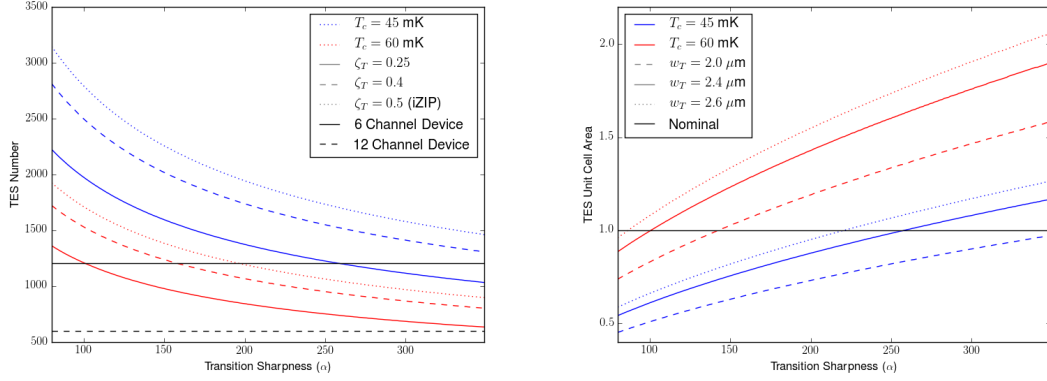


Figure 3.7: Number of QET unit cells as a function of alpha and TES width (left), and as a function of transition sharpness and volume fraction (right), with width and thickness as set in Table 3.1. The third plot shows area per unit cell (assuming 1/6 or 1/12 the area of the crystal surface per channel) as a function of transition sharpness and TES width, using zeta=0.25. The horizontal lines denote the number of unit cells assuming a nominal area of one square millimeter for a 6 and 12 channel face.

have

$$l_{max} \approx \sqrt{\frac{\pi^2 \beta_{wf} \zeta_T}{\Sigma T_c^3 \rho_w (\alpha - 5)}} \quad (3.141)$$

and given that β_{wf} , Σ , and ρ_w are constants, we can write them as an overall constant to show that

$$l_{max} \approx C_W \sqrt{\frac{\zeta_T}{T_c^3 \alpha}} \quad (3.142)$$

(here we absorb n into our margins given that $\alpha \geq 80$) where

$$C_W = \pi \sqrt{\frac{\beta_{wf}}{\Sigma \rho_w}} \quad (3.143)$$

This shows the freedom of design we have when designing our TES length; we want to maximize ζ_T and minimize T_c , while choosing the smallest value of α which gives us good energy resolution.

3.2.6 Additional Internal Degrees of Freedom

In this section we have assumed that the TES can be described by a simple block in internal thermodynamic equilibrium at all frequencies of interest, which in general is not truly the case. I have provided an exploration of two-block TES models with both elements of the TES participating in electrothermal feedback in Appendix C. I refer the reader to Ref [70] for a detailed exploration of 2 and 3 block TES models where only a single block participates in electro-thermal oscillation,

and to Ref [83] for a continuous treatment of phase separated TES dynamics. In all cases the reader will find that any additional internal degrees of freedom degrade expected TES performance, so in the spirit of optimal detector design I have given the general criteria for avoiding phase separation and allowed us to assume the 1-block model is approximately correct. It should be noted that even for a temperature uniform TES, additional performance degradation can occur for the cases of a hanging heat capacity or intermediate thermal conductance which are explored in Ref [70], and will be discussed in Chapter 5

3.3 TES Noise Modeling

Given the readout scheme for the TES, the noise sources we care about are those which add current noise to the TES circuit. In principle, we also care about noise in the readout circuit, but as that noise is expected to be sub-dominant to noise in the TES circuit, and is more readily fixed, we will assume here that it is negligible.

The three main sources of current noise in the TES circuit are the Johnson noise in the passive components, Johnson noise in the TES, and thermal fluctuation noise across the thermal link between the TES and the bath. We separate the noise terms this way because we have more ability to tweak various aspects of our passive components to reduce their noise if necessary, but the characteristics of the Johnson noise in the TES are much more constrained by the other aspects of our optimization, and by chosen operating points.

3.3.1 Johnson Noise

The voltage noise for a resistor at temperature T due to statistical charge fluctuations is white, and is given as

$$\bar{V}_j^2(f) = S_V = 4k_bTR \quad (3.144)$$

For the passive components, we find that these simply add, so we have

$$\bar{V}_{j,passive}^2(f) = \sum_{passive} S_V = 4k_b \sum_{passive} T_i R_i \quad (3.145)$$

and we can convert to a current noise by dividing by the squared impedance of the current loop containing the TES, Z_{loop} :

$$\bar{I}_{j,passive}^2(f) = \frac{4k_b}{|Z_{loop}|^2} \sum_{passive} T_i R_i \quad (3.146)$$

We can simplify this by defining an effective passive temperature T^* (following Ref [83]) as

$$T^* = \frac{1}{R_0} \sum_{passive} T_i R_i \quad (3.147)$$

such that

$$\bar{I}_{j,passive}^2(f) = \frac{4k_b R_0 T^*}{|Z_{loop}|^2} \quad (3.148)$$

In the DC case where $Z_{loop} \approx R_0$, the TES operating resistance, we find that

$$\bar{I}_{j,passive}^2(f) \approx \frac{4k_b T^*}{R_0} \quad (3.149)$$

and we can reduce passive noise if necessary by changing the resistance of the shunt resistor and modifying the bias current to maintain the same voltage but decrease the passive Johnson noise by modifying the effective temperature. This assumes the parasitic resistance is not a factor we can control, and is comparable in magnitude to the shunt resistance.

The TES Johnson noise is slightly more complicated, as we cannot simply divide by the loop impedance to convert voltage noise to current noise. We find first that

$$\bar{V}_{j,TES}^2(f) = 4k_b T R(T) \approx 4k_b T_0 \left[R_0 + \alpha \frac{R_0}{T_0} \delta T + \beta \frac{R_0}{I_0} \delta I \right] \quad (3.150)$$

so building intuition, we notice that for a constant current, voltage noise will increase with small temperature changes proportional to the nominal gain, and current noise translates to voltage noise directly with strength proportional to the β term.

What we inevitably want to probe however is the current noise for a Fourier voltage fluctuation of a given strength. The current noise in this case has been shown to be (by Irwin)

$$I_{j,TES}^2(f) \approx 4k_b T_0 R_0 (1 + 2\beta + \dots) \left(\frac{\partial I}{\partial V_{TES}} \right)^2 \quad (3.151)$$

where the differential voltage to noise expression comes from the dynamics as derived earlier. While this noise term is correct under the assumption that $\delta P = 0$, to be able to simulate noise on pulses accurately, we need to account for the fact that the TES resistance is changing in this noise term as suggested in the previous expression.

We can factor operating resistance out of the voltage to current response function to give

$$\frac{\delta I}{\delta V_{TES}} = \frac{1}{R_0} \frac{1 + i\omega\tau}{\left(\frac{RL}{R_0} + (1 + \beta) + i\omega \frac{L}{R_0}\right)(1 - \mathcal{L} + i\omega\tau) + \mathcal{L}(2 + \beta)} \quad (3.152)$$

which gives

$$I_{TJN}^2(f) \approx \frac{4k_b T_0 (1 + 2\beta)}{R_0} \left(\frac{1 + i\omega\tau}{\left(\frac{RL}{R_0} + (1 + \beta) + i\omega \frac{L}{R_0}\right)(1 - \mathcal{L} + i\omega\tau) + \mathcal{L}(2 + \beta)} \right)^2 \quad (3.153)$$

For completeness, the full form of the passive Johnson noise is

$$I_{PJN}^2(f) \approx \frac{4k_b T^*}{R_0} \left(\frac{1 - \mathcal{L} + i\omega\tau}{\left(\frac{R_L}{R_0} + (1 + \beta) + i\omega\frac{L}{R_0}\right)(1 - \mathcal{L} + i\omega\tau) + \mathcal{L}(2 + \beta)} \right)^2 \quad (3.154)$$

For our full energy resolution estimate, we'll also need the power noise, found by multiplying the current noise by

$$\left| \frac{\delta P}{\delta I} \right| = \left| \frac{I_0}{\mathcal{L}} [(1 - \mathcal{L} + i\omega\tau)(R_L + (1 + \beta)R_0 + i\omega L) + \mathcal{L}R_0(2 + \beta)] \right| \quad (3.155)$$

which gives

$$N_{TJN}^2(f) \approx 4k_b T_0 (1 + 2\beta) R_0 I_0^2 \left(\frac{1 + i\omega\tau}{\mathcal{L}} \right)^2 \quad (3.156)$$

and

$$N_{PJN}^2(f) \approx 4k_b T^* R_0 I_0^2 \left(\frac{1 - \mathcal{L} + i\omega\tau}{\mathcal{L}} \right)^2 \quad (3.157)$$

This is useful in that our input power is much simpler in the Fourier domain than the TES current response, and thus the noise model encodes all of the TES physics while the signal model only incorporate phonon physics. We will revisit the current and power noise relationship at the end of this section when discussing TES channel operating resistance.

3.3.2 Thermal Fluctuation Noise

Thermal fluctuation noise is a power noise across the TES-bath interface, constituting a white power noise of the form [\[83, 50\]](#)

$$N_{TFN}^2 = 4k_b T_0^2 G_0 F_{tfn} \quad (3.158)$$

where

$$F_{tfn} = \frac{1}{2} \left(1 + \left(\frac{T_{bath}}{T_0} \right)^6 \right) \approx \frac{1}{2} \quad (T_{bath} < 0.6 * T_0) \quad (3.159)$$

This can also be re-written using the power balance equation, which tells us that

$$P_0 = I_0^2 R_0 = \frac{G_0 T_0}{5} \left(1 - \left(\frac{T_{bath}}{T_0} \right)^5 \right) \quad (3.160)$$

so we can write

$$N_{TFN}^2 = 20k_b T_0 I_0^2 R_0 F_{tfn} \left(1 - \left(\frac{T_{bath}}{T_0} \right)^5 \right) \approx 10k_b T_0 I_0^2 R_0 \quad (3.161)$$

The current noise can be found by simply multiplying this by the power to current transfer

function:

$$I_{TFN}^2 = 10k_b T_0 I_0^2 R_0 \left| \frac{\partial I}{\partial P} \right|^2 \quad (3.162)$$

$$= \frac{10k_b T_0}{R_0} \left(\frac{\mathcal{L}}{(1 - \mathcal{L} + i\omega\tau) \left(\frac{R_L}{R_0} + (1 + \beta) + i\omega \frac{L}{R_0} \right) + \mathcal{L}(2 + \beta)} \right)^2 \quad (3.163)$$

3.3.3 SQUID Noise

The squid readout noise consists of a white high-frequency noise with a rising tail at low frequencies. The high frequency component is essentially determined by johnson noise across the bias junction of the SQUID, while the low frequency behavior is determined by the temperature of the squid; at high temperature (above the energy of the macroscopic quantum state), the noise flattens out, and in the low-temperature limit, the noise becomes pure 1/f, characteristic of quantum fluctuations between two closely spaced states.

We can generalize the squid noise with the function

$$I_{SQ}(\omega) = I_{SQ,\infty} \left(1 + \left(\frac{\omega_{sq}}{\omega} \right)^{n_{sq}} \right) \quad (3.164)$$

where 1/f noise is the case of $n_{sq} = 1$, and white noise the case of $\omega_{sq} = 0$. For lower operating temperature, the white noise level decreases (the bias johnson noise decreases), but n_{sq} and ω_{sq} both increase. We can also define a SQUID temperature noise T_{SQUID} such that

$$I_{SQ,\infty}^2 = \frac{4k_b T_{SQ}}{R_0} \rightarrow T_{SQ} = \frac{R_0 I_{SQ,\infty}^2}{4k_b} \quad (3.165)$$

which allows us to write

$$I_{SQ}^2(\omega) = \frac{4k_b T_{SQ}}{R_0} \left(1 + \left(\frac{\omega_{sq}}{\omega} \right)^{n_{sq}} \right)^2 \quad (3.166)$$

and as before, we can convert this to a power noise by using the current to power transfer function

$$N_{SQ}^2 = I_{SQ}^2 \left| \frac{\partial I}{\partial P} \right|^2 \quad (3.167)$$

$$= \frac{4k_b T_{SQ}}{R_0} \left(1 + \left(\frac{\omega_{sq}}{\omega} \right)^{n_{sq}} \right)^2 \left| \frac{I_0}{\mathcal{L}} [(1 - \mathcal{L} + i\omega\tau)(R_L + (1 + \beta)R_0 + i\omega L) + \mathcal{L}R_0(2 + \beta)] \right|^2 \quad (3.168)$$

$$= 4k_b T_{SQ} I_0^2 R_0 \left| \left(1 + \left(\frac{\omega_{sq}}{\omega} \right)^{n_{sq}} \right) \left[\left(\frac{1 - \mathcal{L} + i\omega\tau}{\mathcal{L}} \right) \left(\frac{R_L}{R_0} + (1 + \beta) + i\omega \frac{L}{R_0} \right) + (2 + \beta) \right] \right|^2 \quad (3.169)$$

This has the additional behavior of rising with decreasing frequency, to become the dominant low-frequency noise, compared to Johnson noise which is flat at DC.




Figure 3.8: Left: Noise model for UMN cold electronics with significant shunt and parasitic resistances at 1K which are the dominant noise contribution. Right: Noise model for SNOLAB cold electronics with minimal parasitic resistance and 5 m Ω shunt resistor for the same device, showing the improvement in noise performance for electronics designed around a low R_n TES. SQUID noise is not included in these models due to the fact that different SQUIDs were used and the resulting total noise comparison would not be very informative, but in principle the SQUIDs will dominate at high and low frequency.

3.3.4 Joint Noise Sources

As a short note before continuing to the next section, it should be pointed out that both current and power noise models are very easy to combine, but that one needs to be careful that the noise power, not amplitude, is what is added. The noise power N^2 adds linearly, but noise amplitude N adds in quadrature; yet it is usually amplitude spectral density (not really power spectral density) that we plot. In other words, what in CDMS is referred to as the noise PSD is really the amplitude spectral density, and when making the noise model one should plot

$$I_{tot} = \sqrt{I_{sq}^2 + I_{PJN}^2 + I_{TJN}^2 + I_{TFN}^2} \quad (3.170)$$

which has the normal units A/\sqrt{Hz} . I make this note to clear up confusion people normally have when encountering noise analysis for the first time. A good way to remember this is that, in time domain, people like to think about the standard deviation, which adds in quadrature, rather than the variance, which adds linearly, even though the variance is really the parameter they're talking about. This is actually a direct analog, as you can see in Appendix [F](#).

Two examples of noise models for warm and cold passive components (based on the SuperCDMS SNOLAB cold electronics as well as the legacy electronics at UMN) can be seen in Figure [3.8](#). For more discussion of subtleties of noise modeling in non-ideal TES arrays, see Matt Pyle's thesis [\[83\]](#).

3.3.5 Energy Resolution

The amplitude resolution for a signal with normalized template $s(f)$ and noise PSD $N^2(f)$ is given in Appendix [E.6](#) to be

$$\sigma_E^2 = \left[4 \int_0^\infty \frac{|s(f)|^2}{N^2(f)} \right]^{-1} \quad (3.171)$$

Using the TFN noise, we find that the optimum filter resolution is (see also Ref [\[83\]](#))

$$\sigma_E^2 \approx \frac{4k_b T_c^2 G F_{tfn}}{\epsilon^2} \left(\frac{1}{\omega_{pulse}} + \frac{1}{n F_{tfn} \omega_{TES}} \right) \quad (3.172)$$

$$\approx \frac{2k_b T_c^2 G}{\epsilon^2} \left(\frac{1}{\omega_{pulse}} + \frac{1}{\frac{n}{2} \omega_{TES}} \right) \quad (3.173)$$

$$= \frac{2k_b T_c^2 G}{\epsilon^2} \left(\tau_{pulse} + \frac{2\tau_{TES}}{n} \right) \quad (3.174)$$

as $F_{tfn} \approx \frac{1}{2}$ in the low bath temperature limit, in all cases.

Using equation [3.44](#), we find for the ω_{pulse} case that

$$\sigma_E^2 \approx \frac{2k_b n \Sigma V_{TES,eff} T_c^{n+1}}{\epsilon^2 \omega_{pulse}} \quad (3.175)$$

and thus the largest overall impact on the energy resolution comes from T_c , which can dramatically decrease resolution for small T_c decreases. This is the main design driver for SuperCDMS detectors which are mostly phonon fall-time dominated. For the TES fall-time limited case, we get to the ideal calorimeter resolution equation at the beginning of the chapter.

3.4 QET Modeling

With a complete model of TES dynamics, we can tackle modeling the performance of individual QETs, and full channels of QETs, which are the primary phonon collection systems which channel energy to the TES at their center. We need to develop a qualitative model of energy transport in a QET to quantify the energy efficiency of a certain geometry, and phonon channel, to determine how much energy makes it into the TES to be measured, which factors into the energy resolution.

3.4.1 Energy Conversion Efficiencies

There are two critical interfaces where phonons are absorbed which affect the energy transport efficiency in the QET. The first is the crystal/aluminum interface, when phonons are converted into quasiparticles through the Kaplan downconversion process. This has a calculated (and monte carlo validated) efficiency of $\epsilon_{PP} \sim 52\%$ [\[23\]](#).

The second is the aluminum/tungsten interface, which was uncertain even in the iZIP (it was back-calculated, given other efficiencies, from the measured phonon efficiency) and is even more-so in our current devices given the change of device geometry to the inverted interface. It is also unclear how much of this conversion efficiency has been factored into the measured efficiency of energy transport in the fins in the form of the overlap efficiency. The back-converted, and adopted, efficiency for this interface from the iZIP4 (see Matt's thesis) is $\epsilon_{W-Al} \sim 62\%$. These, combined with the QP transport efficiency (next subsection) and energy collection efficiency (from finite aluminum coverage, described in the next section) comprise the total efficiency of the QET.

3.4.2 Aluminum Fin Length

The second aspect of the QET design to consider would be the aluminum fin length. The aluminum fins collect Phonons from the substrate and transport them to the TES, and so we would naively expect that longer fins would be more ideal, as a larger fraction of the energy would be collected in the fins, and the ratio of live to dead aluminum would be higher; aluminum is also used to bias the TES. A QET with ideally efficient fins (which transport 100% of the energy to the TES) would have tiny TES and cover the rest of a phonon channel with Aluminum.

In reality, there are many non-idealities which make long fins undesirable, most important of which is that there is a finite diffusion length of quasi-particles in aluminum. The fraction of quasiparticles collected in the TES is thus a function of fin length, and it turns out also to be a function of fin thickness. The derivation that follows was done by Robert Moffatt, reproduced here for completeness.

1-D QP Diffusion

We can model the 1-D diffusion of QPs in aluminum fins with diffusion equation

$$\frac{\partial n(x, t)}{\partial t} = D_{al} \frac{\partial^2 n(x, t)}{\partial x^2} - \frac{n(x, t)}{\tau_{al}} + s(x, t) \quad (3.176)$$

where $n(x, t)$ describes the number and density of QPs as a function of position and time, D_{al} is the diffusivity of QPs, and τ_{al} is the QP trapping time; s is assumed to be some delta source $s(x, t) = \delta(x - x_0)\delta(t - t_0)$.

We model the absorption rate I_{abs} as

$$I_{abs} = n(0, t)\nu_{abs} \quad (3.177)$$

where ν is defined as the per-qp absorption probability with units length/time. For our 1-d model, the probability of being at an absorption interface at any given time, in a tungsten overlap of length

$l_{overlap}$, for a fin of thickness h_{fin} , for a qp traveling at speed v_{qp} , is

$$\nu_{perfect} = v_{qp} \frac{l_{overlap}}{h_{fin}} \quad (3.178)$$

and for imperfect absorption with probability p , we have

$$\nu_{abs} = p_{abs} \nu_{perfect} = p_{abs} v_{qp} \frac{l_{overlap}}{h_{fin}} \quad (3.179)$$

We'd like also to write our diffusion variables in terms of the physical dimensions of our system, including lengths and velocities. We can write D_{al} in terms of fin thickness and velocity in thickness limited scattering as

$$D_{al} = h_{fin} v_{qp} \quad (3.180)$$

and we make the assumption that W dominates absorption time, and thus $\tau = \nu_w^{-1} h_{fin}$, where ν_w is the empirically measured tungsten absorption rate with units meters/second.

We can further simplify the problem by constructing the dimensionless parameters Λ_d , λ_d , and Γ_d , where

$$\Lambda_d = \frac{\sqrt{D_{al} \tau_{al}}}{L} = \frac{\sqrt{h_{fin} \nu_{qp} \tau_{al}}}{L} = \sqrt{\frac{\nu_{qp}}{\nu_w}} \frac{h_{fin}}{L} \quad (3.181)$$

$$\lambda_d = \frac{D_{al}}{\nu_{abs} L} = \frac{h_{fin}}{v_{qp} l_{overlap} p_{abs}} \frac{h_{fin} \nu_{qp}}{L} = \frac{1}{p_{abs}} \frac{h_{fin}^2}{l_{overlap} L} \quad (3.182)$$

$$\Gamma_d = \frac{\lambda_d}{\Lambda_d} = \frac{\sqrt{D_{al}}}{\nu_{abs} \sqrt{\tau_{al}}} = \sqrt{\frac{h_{fin} \nu_w \nu_{qp}}{h_{fin}}} \frac{h_{fin}}{v_{qp} l_{overlap} p_{abs}} = \frac{1}{p_{abs}} \sqrt{\frac{\nu_w}{\nu_{qp}}} \frac{h_{fin}}{l_{overlap}} \quad (3.183)$$

giving a solution with two geometric parameters, set by the aspect ratio of thickness to length and thickness to overlap.

Robert's solution (with this re-parameterization) gives the relative collection fraction f_c as

$$f_c = \frac{\Lambda_d}{\coth(\Lambda_d^{-1}) + \Gamma_d} \quad (3.184)$$

and we find, from real TES data, that

$$\sqrt{\frac{\nu_{qp}}{\nu_w}} \approx \frac{600}{0.9} \approx 660 \quad (3.185)$$

and

$$\frac{1}{p} \approx \frac{1600 * 5}{0.9^2} \approx 10^4 \quad (3.186)$$

and thus the relative fraction of Phonons collected in our aluminum fins as a function of fin length,




Figure 3.9: Phonon collection fraction as a function of 1-d fin-length, where h_{fin} is fin thickness and γ is W-Al overlap length. According to this model, we should maximize our film thickness and overlap as much as possible, with the caveat that we should do so only as long as the heat capacity of the device does not dramatically increase when doing so.

thickness, and tungsten overlap can be calculated. The trends for various values within our design range can be seen in figure [3.9](#).

2-D QP Diffusion

It should be noted that the above diffusion relations are for 1-D diffusion, which is the valid case for iZIP geometries, but not so for "stadium" geometries, which are oblong and have more surface area further from the TES. True 2D diffusion should behave somewhat differently from 1D due to two non-orthogonal degrees of freedom for QP propagation, and the larger fraction of QPs collected further from the QET.

A similar calculation can to the previous section can be performed for 2-D diffusion, with a more complicated result¹:

$$F = \frac{2\rho_i}{\rho_o^2 - \rho_i^2} \frac{I_1(\rho_o)K_1(\rho_i) - I_1(\rho_i)K_1(\rho_o)}{I_1(\rho_o)(K_0(\rho_i) + \lambda_a K_1(\rho_i)) + (I_0(\rho_i) - \lambda_a I_1(\rho_i))K_1(\rho_o)} \quad (3.187)$$

¹As derived by Robert, will document full derivation at a later time

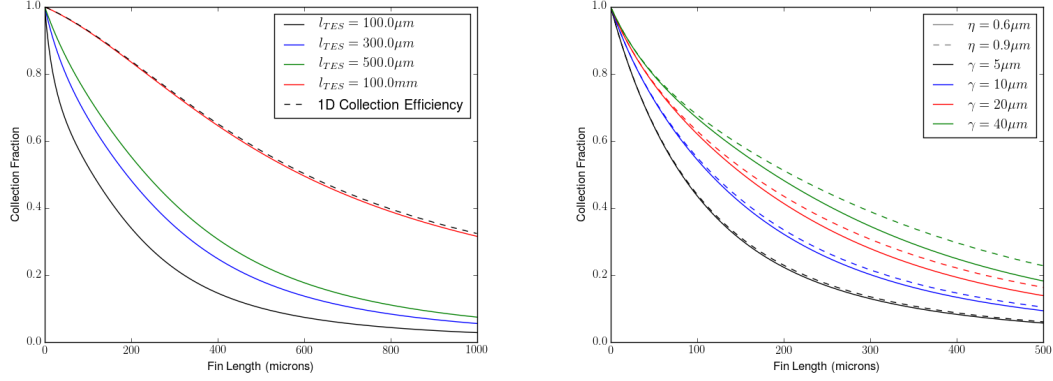


Figure 3.10: Left: Phonon efficiency, moving from 2-D to 1-D geometries, illustrating the fact that in the infinite TES limit, the 2D and 1D collection fractions agree. Right: Phonon collection fraction as a function of 2-d fin-length, where h_{fin} is fin thickness and γ is W-Al overlap length. Notice that the overall efficiency for a given fin length is much worse than expected from the 1-D case.

where

$$\rho_i = \frac{l_{tes}}{\pi l_d} = \frac{l_{TES}}{\pi l_{fin}} \Lambda_d^{-1} \quad (3.188)$$

$$\rho_o = \rho_i + \frac{l_{fin}}{l_d} = \left(\frac{l_{TES}}{\pi l_{fin}} + 1 \right) \Lambda_d^{-1} \quad (3.189)$$

$$\lambda_a = \frac{l_a}{l_d} = \frac{\sqrt{D_{al} \tau_{al}}}{D_{al} / \nu_{abs}} \quad (3.190)$$

and λ_d is as defined in the 1-D case. We can further parameterize this by setting $r_d = l_{tes} / \pi l_{fin}$, giving

$$\rho_i = \frac{r_d}{\Lambda_d} \quad (3.191)$$

$$\rho_o = \frac{r_d + 1}{\Lambda_d} \quad (3.192)$$

and the collection fraction is thus described in terms of the dimensionless parameters Λ_d , λ_d , and r_d , adding an additional degree of freedom to the 1-D case.

The comparison of this solution with the solution for 1-D collection efficiency can be seen in figure 3.10, as well as the various efficiency curves for a 300 micron TES, which is the nominal design length, as discussed in the next section. Note that the overlap and thickness have a less dramatic effect on the collection efficiency.

We can re-write the pre-factor in terms of just Λ_d and r_d :

$$\frac{2\rho_i}{\rho_o^2 - \rho_i^2} = \frac{2\rho_i}{(\rho_i + \Lambda_d^{-1})^2 - \rho_i^2} \quad (3.193)$$

$$= \frac{2\rho_i}{\Lambda_d^{-2} + 2\rho_i\Lambda_d^{-1}} \quad (3.194)$$

$$= \frac{\Lambda_d}{\frac{\Lambda_d^{-1}}{2\rho_i} + 1} \quad (3.195)$$

$$= \frac{\Lambda_d}{r_d^{-1} + 1} \quad (3.196)$$

and so in the 1-D limit ($l_{fin} \ll l_{TES}$) this reduces to the same large l_{fin} scaling as the previous 1-D solution:

$$\frac{2\rho_i}{\rho_o^2 - \rho_i^2} \approx \Lambda_d = \frac{l_d}{l_{fin}} \approx \frac{l_d l_{TES}}{A_{QET}} \quad (3.197)$$

In the 2-D limit ($l_{fin} \gg l_{TES}$), we find that

$$\frac{2\rho_i}{\rho_o^2 - \rho_i^2} \approx \frac{l_d l_{TES}}{\pi l_{fin}^2} = \frac{l_d l_{TES}}{A_{QET}} \quad (3.198)$$

and so regardless of geometry, both find that collection efficiency scales inversely as fin area and proportional to diffusion length and TES length. The fact that the efficiency depends on TES length will be immensely important in scaling energy resolution with TES dimensions.

QP Efficiency and Flux Trapping

A final consideration in QP fin design is the mitigation of trapped flux vortices, which trap quasiparticles in the fin and lower the efficiency of the QPs. Aluminum is nominally a type 1 superconductor, meaning it does not permit bulk magnetic flux lines, however thin films (on the order of hundreds of nanometers) have been shown to develop stable or metastable flux vortices which become very persistent [23]. The transition from type 1 to intermediate type for this film thickness occurs for square films on the order of 40 microns square, and it is likely that increasing the film thickness in aluminum increases the size of the contiguous superconductor which is able to fully repel flux lines. This is one area of R&D that should be explored further by our collaboration as we move to thicker Aluminum films due to their increased QP collection efficiency.

Figure 3.11 shows the impact of these flux vortices on the QP collection efficiency in our QETs. A prototype iZIPv6 was operated at UMN, and was subjected to different magnetic fields by positioning the mu-metal magnetic shield at various heights relative to the detector to create different ambient magnetic fields. The detector phonon collection efficiency was measured for each cool-down, and the resulting efficiency is seen to be a strong function of ambient magnetic field, implying that even for this detector (for 300nm film thickness) there is a significant likelihood of forming flux vortices for




Figure 3.11: Left: Phonon collection efficiency for a prototype iZIP detector measured at UMN by Matt Fritts, showing that as the magnetic flux around the detector is increased during cooldown the efficiency decreases. This indicates that flux traps do persist in that design despite the films being ~ 300 nm thick. Right: QET unit cell showing fin slits separating each major fin (defined by each fin connector) as well as thinner slits meant to reduce the area of contiguous superconducting regions, and allow magnetic flux to pass through the fin without producing vortices. This design ensures that a circle of radius ~ 50 microns cannot be found anywhere in the design. For 600nm films this should be enough to ensure that any vortices in the fins are more easily expelled.

the iZIP detector design.

For the CDMS HV design, in order to be conservative, we decided on a maximum fin width of 50 microns, decreasing larger widths by adding small slits to larger fin areas to disrupt large areas of potential vortex formation. This can be seen on the right in Figure 3.11 and we can differentiate between vortex (minor) slits and major slits needed to prevent shorting the TES by the presence of the green layer in the figure, which is the insulating amorphous layer in the detector mask. We don't particularly care about continuity issues in those slits, so they have been made narrower than the major slits. Though these were initially included to prevent flux vortices, the slits also serve to make diffusion in the fins more quasi-1D, and should in principle slightly improve the QP collection efficiency for phonons absorbed on the edge of the fins. This is not taken into account in the optimization discussed in the next section but was a secondary design philosophy for this QET.

3.4.3 Tungsten-Aluminum Overlap Dimensions

Figure 3.12 shows the tungsten-aluminum overlap interface dimensions used for the HV design. As described in the previous section, our QP diffusion model suggests that an increased overlap should improve QP collection efficiency out beyond 40-50 microns. At the time this device was designed, only data for 5 and 10 micron overlaps had been acquired, and the design value chosen was 20 microns




Figure 3.12: Tungsten (black) and aluminum (grey/blue) overlap region for v0 (top) and v1 (bottom left) compared to a microscope image of the fabricated overlap region for the inverted geometry for v2 (bottom right). You can see that the overlap geometry isn't changing but the design of the TES connector includes features to avoid connectivity issues near corners in the design.

assuming the gains would continue to be seen for larger overlap values. The design philosophy was to minimize the mean free path of a given carrier through the overlap to the TES, which is how the rivet-like geometry seen in Figure 3.12 was decided on.

The additional dimensions and features of the overlap region were designed to minimize the fraction of tungsten not covered by aluminum. This was driven by the assumption (based in part on the measurements from Ref [45]) that all tungsten in contact with aluminum is highly proximitized, and thus does not contribute to either the heat capacity or thermal conductance of the TES, while the part of the fin connector not covered by aluminum will affect both quantities. This is captured by the ζ parameter introduced earlier in this chapter; a ζ of ~ 0.7 corresponds to including all 8 fin connectors in the total tungsten volume, which affects both the heat capacity and thermal conductance of the device (in principle leaving τ invariant).

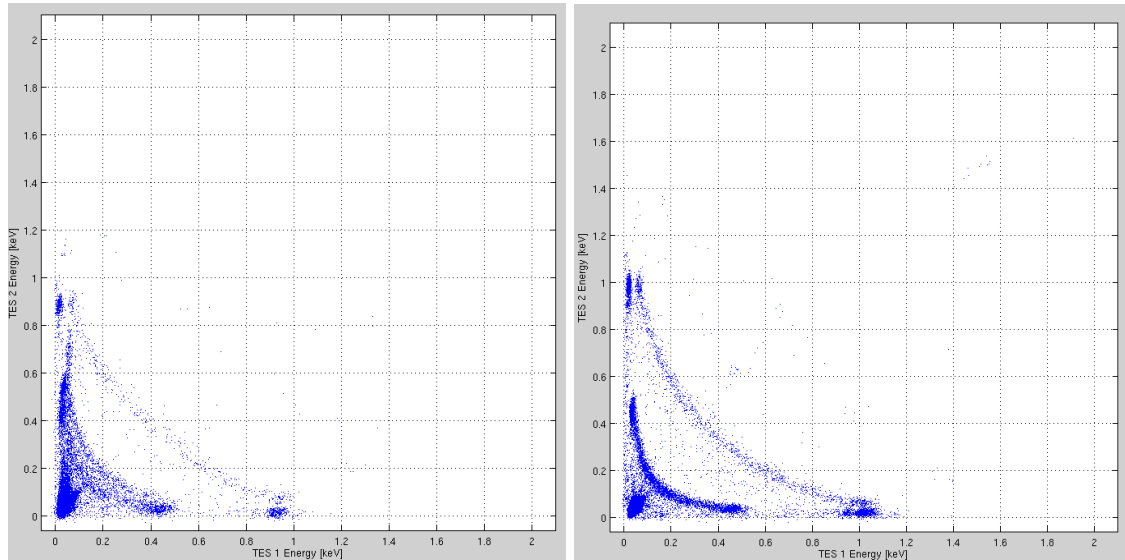


Figure 3.13: Energy collection in two TES attached to the same aluminum fin for (left) a 30 micron overlap and (right) a 40 micron overlap. This data suggests that continuing to increase W/Al overlap length does increase collection efficiency, at the banana reaches higher energies in these plots on both axes and there is a smaller conversion penalty as seen by the spacing of the shared energy band from the axes in the two plots.

Stanford QP Collection Measurements

One piece of information that has become clearer since the design of the HV mask as finalized is that the projected gains with overlap do seem to be realized, at least in our test devices. Figure [3.13](#) shows data taken with the same test device geometry described in Ref [\[111\]](#).

These data qualitatively confirm that QP collection efficiency increases with overlap length for small geometries, leading us to believe that we could in principle increase overlap length for future designs. There are a few potential problems which arise from much larger overlap regions which are not accounted for in our models however:

- We have not quantified what fraction of the tungsten under the aluminum fin contributes to the heat capacity and thermal conductance of the TES (as we'll see in the next chapter, it's very possible that it's a non-negligible amount).
- The test devices utilized thus far are one-dimensional; it is unclear whether these same gains should necessarily be realized in two-dimensional traps.
- At some point we expect that the finite diffusion length of phonons through the overlap region should cause efficiency to drop again, as was seen in previous designs with meandering fin connectors (this is summarized nicely in Matt's thesis [\[83\]](#)).

Some of these questions will be addressed in the testing section, but some are (as of this writing) under studying in small test devices, and another set have yet to be tested. The energy conversion at the overlap is by far the least mature aspect of the QET design process, and a future design that fully matches the overlap to the fin and the fin to the substrate should achieve the optimal energy resolution for a QET with a given geometry.

3.4.4 Energy Collection Efficiency Summary

In this section, we've seen that most steps of the phonon measurement process have known efficiencies which we can, in principle, use to optimize the QET design. To assist the reader in compiling this model, I will summarize the various components:

- Phonon to QP Conversion Efficiency - Kaplan downconversion limits to 52% [23]
- QP Collection Efficiency - $f_c(\eta, \gamma)$ is a function of the aspect ratio of the fin and its thickness and is related to measured film properties as shown in this section. It also depends on the amount of flux trapping in the fins, which can in principle be made negligible.
- Trapped QP to TES Efficiency - Not yet modeled, but measured in IZIP4 to be $\sim 62\%$

If we can perfectly collect all phonons in our fins and transport them perfectly to the TES, then, we predict an energy efficiency of $\sim 32\%$; if we measure an efficiency which exceeds this, then we have either messed up the measurement or made a significantly observation about fundamental QET efficiencies. As the fin connector interface is further refined, it may be that efficiencies in the low-mid 40% range are possible, but even the perfect design cannot exceed the $\sim 50\%$ down-conversion limit according to the design principles presented here. This is a good metric to compare measurements to as a very first sniff test.

3.5 Future Modeling Inputs

In this chapter, we have seen a very mature model of SuperCDMS QETs, and I have tried to highlight areas of the modeling which have progressed recently as well as areas which are in need of further study. Many of the results shown in the next few chapters bring alot of the assumptions made in our detector designs into question. Most notably, we implicitly assume that changing the T_c of our detectors changes the resistance, heat capacity, and thermal conductivity of our films according to the prescribed power laws.

One crucial assumption that we have not explicitly verified (but is one part of the model which can be addressed through an ongoing TES test program) is the accuracy of the Wiedemann-Franz law at low temperature:

$$\kappa = \beta_{wf} \sigma T = \beta_{wf} \frac{T}{\rho} \quad (3.199)$$

While in general this result does hold, it has not been explicitly verified at low temperature, and deviations from this law abound in the literature for different experimental conditions [10]. It is more likely than not that thermal conductance is higher than predicted by this scaling relation due to phonon dynamics, but smaller values in highly-correlated systems (e.g. superconductors) have also been observed. This adds a large degree of uncertainty to our phase separation boundary, and given that resistivity can in general be a complex function of T_c , it's unclear whether our estimates of κ based on other measurements will be systematically low or high. As we'll see in the next chapter, a longer TES allows for much more aggressive mask designs for large detectors, so if the phase separation boundary is much further out this would change our detector optimizations.

In general, a better understanding of the resistance, thermal conductance, electron-phonon coupling, and proximity effects in our QETs will improve our ability to predict the noise performance of a given design. In this chapter we explicitly assume that all tungsten underneath the Al is proximitized, and does not contribute to the heat capacity of the TES, but we have begun to believe this is not the case, as will be discussed in Chapter 5. Other effects we have observed but do not understand include

- Dependence of heat capacity on transition point (as well as slope of this change)
- Dependence of thermal conductance (to the bath) and heat capacity as a function of T_c and film morphology, as well as the aforementioned effect of proximitizing the overlap regions
- Dependence of resistivity on film morphology and T_c
- Resilience of various QET designs to flux trapping as a function of fin feature size, and further studies of QP diffusion as a function of fin shape
- Better modeling of QP trapping and diffusion of energy through the overlap interface; current modeling assumes no dependence on trap shape and size
- Improved understanding of phonon losses at uninstrumented surfaces (currently our model assumes perfect reflection at all non-aluminum interfaces)

Many of these questions will be answered in the near future by small device and test pattern tests at SLAC and UCB, but others will require more dedicated fabrication and testing programs, and it is likely that trends can be understood on a more fundamental level once measured.

Chapter 4

CDMS HV Detector Design

“If you find that you’re spending almost all your time on theory, start turning some attention to practical things; it will improve your theories. If you find that you’re spending almost all your time on practice, start turning some attention to theoretical things; it will improve your practice.”

- Donald Knuth, quoted in: Arturo Gonzalez-Gutierrez (2007) Minimum-length Corridors: Complexity and Approximations

In this chapter I apply the concepts described in the previous chapter to the optimization of the first 100 mm CDMS HV detector. The general optimization procedure is the following:

1. Select a detector size and form factor
2. Construct a readout noise model to determine the right operating conditions for the TES channel
3. Given fixed TES channel properties, optimize the QET design
4. Given a fixed QET design, determine an optimal channel layout
5. Figure out how to make the optimized design work in practice

The first 4 steps will be described in detail and related to the models described in the previous chapter, and I will provide some commentary on the last step given the experience with this detector design to help inform future SuperCDMS mask and housing designs. While most of the cerebral part of mask design lies in the optimization, mistakes made during the last step can render a detector susceptible to critical failures, and a non-functional detector, regardless of the degree of optimization, is useless. This chapter provides a framework for making optimized CDMS detectors which are robust




Figure 4.1: Left: Z partition (energy fraction absorbed on first pass) as a function of the aspect ratio of the phonon signal. If the signal were entirely ballistic, then R is the detector radius, but if it is diffusive, then R is the diffusion radius. Right: Difference in z-partition for diffusion limited phonons for differences into the diffusive radius on either side of the detector (e.g. for NTL emission by holes v. electrons)

to the realities of fabrication constraints, human error, and non-idealities in detector operation and installation. It is not exhaustive but should provide a useful guide of things to consider during the design process.

4.1 Detector Dimensions

An often overlooked design choice when we make CDMS detectors is the form factor of the detector crystal. The normal considerations generally involve the detector mass and the dimensions of the available stock material, as well as fabrication limitations. CDMS has historically chosen detector substrates which are nominally cylindrical, and have aspect ratios of 3:1 as seen for the Soudan (3 inch diameter, 1 inch thick) and SNOLAB (100 mm diameter, 33.3 mm thick) detector crystals.

There are different philosophies for determining the correct form factor of the detector, which depend on the detector properties you expect to limit its performance most critically. If you have bad surface treatment on the side-walls, for example, you might want to opt for a thin detector, but if you believe your surfaces are all equally lossy (or non-lossy as may be the case), then you might opt to maximize the fiducial volume of your detector. In this case, the degree of fiducialization depends on your surface to volume ratio as well as the accuracy of your reconstruction methods. For a phonon-only detector the position resolution achievable depends on the phonon dynamics; the optimization scheme will depend on whether the phonons are ballistic (traveling directly from the point of origin) or diffusive.

The characteristic length scale of phonon propagation is energy dependent, but in general we find that mean free paths on the order of a few mm are about right for Ge^[99, 83], and we can




Figure 4.2: Left: Fraction of the volume able to be fiducialized at the optimal aspect ratio for a range of r and z resolutions as a function of the ratio of σ_r and σ_z . Right: Optimal aspect ratio for a given resolution ratio, demonstrating the independence of the optimum on other detector factors. This is an idealized model, as we should expect that sufficiently large detectors will lose all position resolution in certain limits.

expect they're comparable in Si. The fundamental limitations of ballistic phonon propagation are anharmonic decay and impurity scattering, meaning the purity of the crystal is often important, but even a perfect crystal will not permit purely ballistic phonon propagation. A rough measurement of the effective phonon diffusion radius will be discussed later in this chapter, but for now we'll work in orders of magnitude.

Figure 4.1 shows the z partition quantity derived in Appendix A for different effective phonon diffusion radii. For a detector that is much wider than it is thick, the partition of energy is linear in z , and the position resolution is only limited by the energy resolution. For a diffusion limited ball of phonons, or equivalently, a detector that is much thicker than it is wide, there is only appreciable z -dependence near the detector surfaces. If linear performance is desired, then a detector with dimensions on the order of mm which will collect all phonons during their initial propagation and has a very flat aspect ratio will be ideal. For larger detectors, the phonon diffusion is the limiting radial scale, so we do not necessarily benefit from the aspect ratio in this partition quantity.

The radial partition is also limited by phonon diffusion, and will only improve in terms of absolute fiducialization with increased detector radius. In the diffusive limit, channel sizes on the order of 5–10 mm in radial rings, with multiple radial channels, will generally give good radial reconstruction performance, with the caveat that the radial partition is complicated by events which occur near channel splits.

Given reconstructed position resolutions σ_r and σ_z , assuming these are independent of the detector form factor (as they are in the purely ballistic limit), we can actually find the optimum detector form factor independent of total detector volume. If we assume the detector is cylindrical, and that the fiducial volume consists of all points more than 3σ from a surface, then we have the fiducial

fraction

$$f_{fid} = \frac{V_{fid}}{V_{tot}} = \frac{(r - 3\sigma_r)^2(z - 6\sigma_z)}{r^2z} = \left(1 - 3\frac{\sigma_r}{r}\right)^2 \left(1 - 6\frac{\sigma_z}{z}\right) \quad (4.1)$$

If we fix the total detector mass, then we fix the relation between r and z to be

$$z = \frac{V_{tot}}{\pi r^2} \quad (4.2)$$

which gives

$$f_{fid} = \left(1 - 3\frac{\sigma_r}{r}\right)^2 \left(1 - 6\pi r^2 \frac{\sigma_z}{V_{tot}}\right) \quad (4.3)$$

This equation has the analytic solution

$$\frac{r}{z} = \frac{1}{2} \frac{\sigma_r}{\sigma_z} \rightarrow \frac{d}{z} = \frac{\sigma_r}{\sigma_z} \quad (4.4)$$

as shown in Figure [4.2](#). If we can be reasonably confident that position resolution is insensitive to changes in aspect ratio then we can use measured resolutions to determine the optimal detector form factor.

More likely, however, is that both are also a function of the aspect ratio, and what we have is actually a transcendental equation; for any given measurement we can only say how far from the optimal configuration we are. For the SuperCDMS SNOLAB detectors, the same aspect ratio as the Soudan detectors was chosen based on source material and fabrication constraints. If we find that the demonstrated position resolution of these detectors is very far from a 3:1 ratio, then we need to re-evaluate how to either improve position resolution with the same form factor, or change the form factor of the detector crystals. This was not done for these detectors due to the complexity of position resolution calculations but is worth considering for a future upgrade phase of detectors.

4.2 CDMS HV QET Design

The previous chapter described the response model of each of the individual parts of the QET. In this section I will use the results of the previous chapter to describe the design and optimization of the QETs independent of the rest of the detector design for SuperCDMS SNOLAB CDMS HV. Table [3.1](#) contains input parameter values assumed for this optimization; some of these assumptions turned out either to be overly confident or overly optimistic as will be discussed later in this chapter.

4.2.1 TES Channel Optimization

Much of the TES dynamics can be removed from the QET optimization based on the characteristics of the readout circuit; we can assume an infinitely fast phonon response, which will only be more sensitive to noise than our phonon signal with a finite fall-time which in general is slower than the

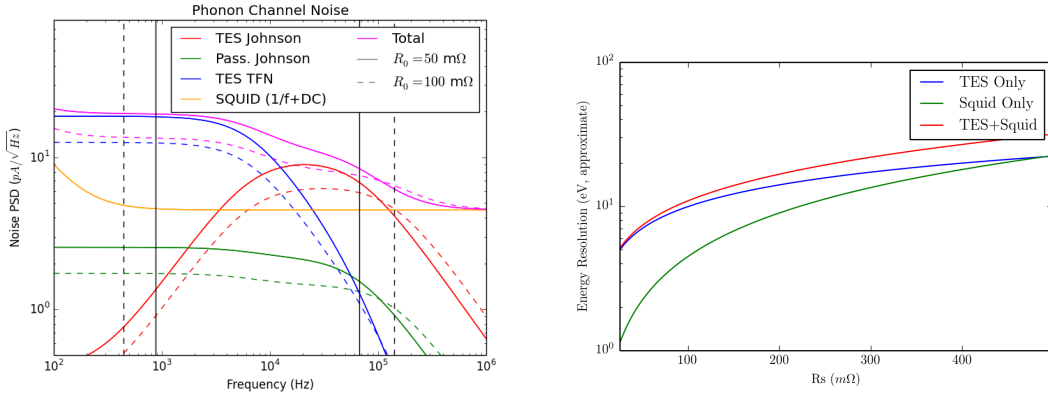


Figure 4.3: Left: Current noise amplitude power spectra for the original normal state resistance target (300 mΩ) and the final normal state resistance target (150 mΩ). We can see that in both cases, passive noise is sub-dominant, but the SQUID noise is a larger contribution at the higher normal state resistance. Right: Energy resolution (assuming nominal efficiency and phonon fall-time of the CDMS-HV design described later in the chapter) as a function of normal state resistance, showing the SQUID only case versus the TES only case. We can see that the SQUID has a large impact on resolution for higher normal state resistances, and lowering the normal state resistance both lowers the overall resolution as well as the impact of the SQUID relative to the TES.

TES response.

Figure 4.3 shows the noise model for our TES readout circuit, using the functional forms from the previous chapter. The passive noise is seen here to be completely subdominant to the TES noise, by design. This is achieved by assuming both R_{sh} and R_p are 5 mΩ, and are operated at 1K (where the HEMT card is in the fridge)¹. The SQUID noise shown in Figure 4.3 is based on the measured noise from SNOLAB SQUIDS tested at both UC Denver and SLAC, which show a baseline noise floor around 4.5 pA/√Hz. The exponent in equation 3.166, n_{sq} , is assumed to be nominally 1, which is in a sense the worst case scenario. Later in this chapter I will explore the impact of varying the SQUID noise exponent and low-frequency baseline on the expected energy resolution.

We can see from Figure 4.3 that the initial choice for TES normal state resistance would have put the TES noise floor near the SQUID noise floor, leading to sub-optimal energy resolution and a high dependence on individual squid performance. Given that the iZIP design already assumed a normal state resistance of 150 mΩ, we decided to also lower the HV normal state resistance to this point. We can see for this normal state resistance, the TES noise is further above the SQUID noise floor, and the resolution is much less dominated by SQUID performance, though we will still be dependent on the SQUID noise to the 10% level or so. This drives home the importance of doing full channel characterization at SNOLAB to determine the optimal bias point for the TES channel as a

¹Table 4.2 contains all of the assumptions for variables not optimized as well as the results of the optimization described in this section.

fraction of normal state resistance. Too high in the transition and we run the risk of either being SQUID dominated or in a part of the transition which is not sharp enough for ideal electrothermal feedback. Too low in the transition we risk either oscillation or large nonlinearity due to being too close in resistance to $R_l = R_p + R_s h$. It's likely that the optimal bias point will be in the range of 30-60 m Ω for each channel, but will depend on the individual electrical characteristics of each channel.

4.2.2 QET Efficiency

Quantity	Symbol	Adopted Value
Crystal Radius	R_{xtal}	50 mm
Sensor Radius	R_{ph}	48 mm
Sensor Area	A_{ph}	7238 mm ²
Phase Separation Length	l_{ps}	348 μ m
TES Design Length	l_{TES}	200 μ m
TES Connector Length	l_c	2.5 μ m
QET Fin length	l_{fin}	240 μ m
Fin Channel Width	w_{chan}	10 μ m

Table 4.1: QET parameters used to determine efficiency, aluminum coverage, and limits on optimization domain. Input TES parameters are also shown in Table 3.1 and optimized QET parameters are given in Table 4.2

The layout of the QET comes from the considerations in the last chapter, where we noted that the 2D geometry was more ideal than 1D from an efficiency perspective for the right optimization. We also want as much instrumented ('live') aluminum coverage as possible, as phonons absorbed in non-instrumented ('dead') aluminum (bias lines) will not be detected, and the ratio of live to dead aluminum sets the base-line efficiency. In addition, the QET is segmented into 8 fins, such that the TES may be properly biased, mitigate magnetic flux traps[23], and to channel QPs into the tungsten traps rather than letting them travel parallel to the TES. The final QET design can be seen in figure 4.4. In this section I'll describe the optimization process that led us to this design.

For this geometry, we eventually want to calculate the aluminum coverage, both live and dead, and so we need to first calculate the unit cell area, as this sets the bias line dimensions and scales to the final percentage. We use 50 mm radius Germanium and Silicon crystals, but leave the outer 2mm uncovered to protect against accidental shorts from touching the detector casing, so only 48mm of radius are instrumented. Those outer 2 mm will contain small guard rings attached to ground to allow for more uniform electric field at high radius. The area per QET is thus

$$A_{cell} = \frac{A_{ph}}{N_{TES}} = \left(\frac{A_{ph} R_n}{\rho_w} \right) \frac{w_T h_T}{l_{TES}} \quad (4.5)$$




Figure 4.4: Optimized QET design, where blue is aluminum, green is amorphous germanium, and black is tungsten. See text for more details. Note that this shows an earlier version of the QET with more amorphous Ge etched in the fin connector region than in the final design.

and the area of the QET is roughly

$$A_{QET} \approx \frac{\pi}{2}(l_{TES} + 2l_{fin})l_{fin} \quad (4.6)$$

We recognize here that, due to removal of some material, the true area is slightly smaller:

$$A_{QET} = \frac{\pi}{2}(l_{TES} + 2l_{fin})l_{fin} - \frac{\pi}{4}l_{TES}(w_{TES} + 2\sqrt{3}l_c) \quad (4.7)$$

$$- 6l_{fin}w_{chan} - n_{squares} * A_{square} \quad (4.8)$$

where the quantities are as defined in table 4.1. The channels and holes remove some live aluminum area, as does the TES. These are second order corrections I will ignore in the scaling relations below, but they are included in the numerical optimization.

The dead aluminum area is just as dependent on TES length as the live aluminum area. The top and bottom connectors depend on the side length of the unit cell, as does the length of bias line per cell. Here I'm going to assume that each TES shares each of its bias lines with its upper and lower neighbors, so that for a given unit cell we need only include the area from the top bias line. For the limit of large N_{TES} , this will be accurate. The dead aluminum is then

$$A_{dead} \approx (\sqrt{A_{cell}} - l_{TES} - 2l_{fin} - w_{hbias}) * w_{vbias} \quad (4.9)$$

$$+ \sqrt{A_{cell}}w_{hbias} \quad (4.10)$$

where w_{hbias} and w_{vbias} are the widths of the horizontal and vertical bias lines respectively. If we

want to probe the amount of dead aluminum at fixed TES length as a function of l_{fin} , we can write this as

$$A_{dead} \approx \sqrt{A_{cell}}(w_{vbias} + w_{hbias}) - l_{TES} * w_{vbias} - 2l_{fin} * w_{vbias} \quad (4.11)$$

$$\approx w_{bias}(2\sqrt{A_{cell}} - l_{TES} - 2l_{fin}) \quad (4.12)$$

$$\approx 2w_{bias}(L_0 - l_{fin}) \quad (4.13)$$

where $L_0 = \sqrt{A_{cell}} - \frac{1}{2}l_{TES}$ is the half-aluminum length for 0 fin length, and we've assumed the wire lengths are equal and minimal due to photolithography constraints.

Using these areas, we can begin to calculate phonon collection efficiencies. It is important to note that for each phonon interaction with aluminum instrumented surface, there is a non-trivial probability of reflection rather than absorption. On a given bounce, then, the fraction of phonons absorbed by a given QET can be written

$$f_{abs} = p_{abs} \frac{A_{QET}}{A_{cell}} \quad (4.14)$$

$$\approx p_{abs} \frac{\pi}{2} (l_{TES} + 2l_{fin}) l_{fin} \left(\frac{\rho_w}{A_{ph} R_n} \right) \frac{l_{TES}}{w_T h_T} \quad (4.15)$$

$$= \left(\frac{\pi p_{abs} \rho_w}{2 A_{ph} R_n w_T h_T} \right) (l_{TES}^2 l_{fin} + l_{fin}^2 l_{TES}) \quad (4.16)$$

$$= C_{TES} (l_{TES}^2 l_{fin} + l_{fin}^2 l_{TES}) \quad (4.17)$$

where it is clear we have a TES growth regime and a fin length growth regime. It will turn out that, given that fin trapping length is on the order of our TES length, that we're not in a limiting case, so I won't bother with any expansion here. Similarly, the phonons lost in the dead aluminum can be calculated as

$$f_{lost} = p_{abs} \frac{A_{dead}}{A_{cell}} \quad (4.18)$$

which will decrease with increasing TES length.

Finally, to calculate the phonon collection fraction in the TES as a function of fin length, we need to include the collection efficiency described in the previous section. This gives

$$f_{abs} = f_c(l_{fin}) p_{abs} \frac{A_{QET}}{A_{cell}} \quad (4.19)$$

and if we recall the large-fin limits from the previous section, we find that

$$f_{abs} \approx p_{abs} \frac{l_d l_{TES}}{A_{cell}} \quad (4.20)$$

in the large fin limit. Thus beyond the point where the bessel pre-factor goes to 1, around 1 diffusion




Figure 4.5: Left: 2D Phonon collection efficiency plane as a function of TES and fin length. Note that the small fin limit is dead-aluminum limited, and the large-fin limit is limited by collection efficiency. Right: 2D energy efficiency plane as a function of fin length and TES length. Both plots show the phase separation boundary for a TES at 45 mK.

length, we find that the energy collected on a given bounce stops increasing. The lower collection fraction means energy was lost, so this is not beneficial.

We need also to include the dead aluminum in our consideration, to compute the total phonon collection efficiency integrated out to infinite time after the initial pulse. On each pass, f_{abs} phonons are absorbed, however those which are absorbed are not collected with 100% efficiency, and some of the phonons will be absorbed in dead aluminum. Thus the maximum energy resolution is the fraction of collected phonons to total absorbed phonons on a given pass:

$$\epsilon = \frac{f_c A_{QET}}{A_{QET} + A_{dead}} \quad (4.21)$$

$$= \frac{f_c}{1 + \frac{A_{dead}}{A_{QET}}} \quad (4.22)$$

We see then that at short fin length, the efficiency scales approximately as the live to dead aluminum ratio, and at large fin length (when dead aluminum is negligible), the efficiency scales as the collection efficiency. This can be seen in figure 4.5, where the efficiency for a range of TES sizes is plotted as a function of fin length.

4.2.3 Energy Resolution

Given the model of efficiency as a function of our design parameters, we can now do the full energy resolution optimization. Remembering that efficiency depends on signal bandwidth and efficiency, from equation 3.175, we first need to determine signal bandwidths before calculating energy resolutions as a function of QET dimensions.

Signal Bandwidths

The two quantities which are important in determining this resolution are the TES bandwidth and the pulse bandwidth. The TES bandwidth we define as

$$\omega_{TES} = \frac{1}{\tau_{TES}} = \frac{G(1 + \mathcal{L})}{C} = \frac{n\Sigma(1 + \frac{\alpha}{n})}{f_{sc}c_W T_c^{2-n}} \quad (4.23)$$

We can estimate the pulse bandwidth ω_{pulse} given the aluminum coverage f_{al} , absorption probability p_{abs} , and thickness transit time $t_{transit}$ as

$$\omega_{pulse} = \frac{1}{\tau_{pulse}} = \frac{f_{al}f_{abs}}{t_{transit}} = \frac{f_{abs}}{t_{transit}} \frac{A_{Al}}{A_{puck}} \quad (4.24)$$

$$= \frac{f_{abs}}{t_{transit}A_{puck}} N_{TES}(A_{QET} + A_{dead}) \quad (4.25)$$

$$= \frac{f_{abs}v_{phonon}}{\mathcal{V}_{puck}} N_{TES}(A_{QET} + A_{dead}) \quad (4.26)$$

and so we see that the TES bandwidth is set by device parameters, while the pulse bandwidth is mainly geometric.

Given this relation, and the pulses seen for the iZIP v5, we can parameterize ω_{pulse} in terms of the known value for the iZIP. We saw from Soudan data a pulse decay time of $\sim 750\mu s$, for an aluminum coverage of 4.8%, giving us the relation

$$\frac{f_{abs}}{t_{transit}} = \frac{f_{abs}v_{phonon}}{h_{puck}} \approx \frac{1}{0.048 * 750 * 10^{-6}} \approx 2.7 * 10^4 \quad (4.27)$$

and in terms of intrinsic constants, we have, for $h_{puck} \approx 25$ mm,

$$f_{abs}v_{phonon} \approx 705m/s \quad (4.28)$$

and if the DMC measured value of $f_{abs} \sim 0.4$ is to be believed, then we find $v_{phonon} \approx 1760$ m/s.

This bandwidth helps us determine a signal template to use in our optimal filter calculations. We can imagine that the phonon pulse has a similar shape to the TES impulse response, i.e. a double exponential function, which has the normalize fourier transform

$$N_{pulse,ideal}(\omega) = \frac{1}{\tau_{pulse} - \tau_{phonon}} \left(\frac{\tau_{pulse}}{1 + i\omega\tau_{pulse}} - \frac{\tau_{phonon}}{1 + i\omega\tau_{phonon}} \right)$$

where the falltime is τ_{pulse} , calculated above, and the rise time is

$$\tau_{phonon} \sim \frac{t_{transit}}{2} \quad (4.29)$$




Figure 4.6: Phonon collection time as a function of TES and and fin length for Ge (Si) on the left (right). The iZIP, which only has an aluminum coverage of $\sim 4\%$, is much slower. Si is faster than Ge due both to the faster sound speed and the better impedance match between Si and Al. The Si HV detector is expected to be fully bandwidth matched, i.e. the phonon and TES falltimes should be about the same.

The phonons arrive at the surface on the timescale of a single crystal transit (half the time from the center) with dispersion of the same timescale as the phonon transit time. This is at first glance conservative estimate, as if the phonons were concentrated in a delta in time, isotropic, and traveled at the same velocity equal to the group velocity, we would see effectively 0 rise time. What we see, however, is that phonon channels away from the initial vertex see a much longer effective rise time, can in principle only improve the energy resolution of the measurement by restricting the signal to a smaller volume of frequency space. As we don't expect the rise time to be longer than the fall time, we thus see that estimating our pulse with just the fall-time will be a conservative energy resolution estimate.

If we add in the fact that only a certain fraction of the signal gets through, according to the efficiency, ϵ , then we can represent our full signal template as

$$N_{pulse}(\omega) = \frac{\epsilon}{\tau_{pulse} - \tau_{phonon}} \left(\frac{\tau_{pulse}}{1 + i\omega\tau_{pulse}} - \frac{\tau_{phonon}}{1 + i\omega\tau_{phonon}} \right)$$

This function can now be used to create an optimum filter for the data, and thus allow us to compute the energy resolution of our optimally filtered signal using the power noise estimated earlier.

In summary, figure 4.6 shows the expected phonon and TES falltimes as a function of TES and fin length. An immediate difference between the Si and Ge plots is that Si is faster by a factor ~ 4.5 ; this is an experimentally determined fact, determined by comparing the phonon collection time in a Si iZIP5 to that in the Ge iZIP5. This is due to a better impedance match between Si and Al [56] and a faster phonon sound speed [52], which are both corrections of ~ 2 as compared to Ge. Work is in progress to try to use the DMC to predict fall-times and transmission/reflection coefficients from

first principles as a function of phonon frequency and incidence angle based on Ref [56].

Optimum Filter Energy Resolution - TFN Dominated

We can now optimize the QET and TES dimensions assuming the TFN dominated TES resolution given in equation 3.175. Plugging in the pulse bandwidth gives the single channel resolution

$$\sigma_E^2 \approx \left(\frac{2k_b n \Sigma T_c^{n+1} \mathcal{V}_{puck}}{f_{abs} v_{phonon}} \right) \frac{\mathcal{V}_{TES}}{\epsilon^2 N_{TES} (A_{QET} + A_{dead})} \quad (4.30)$$

It should be noted that this TES volume is the volume of the TES in a channel, which is actually $n_{tes} V_{TES}$ where $V_{TES} = \zeta^{-1} l_{TES} w_{TES} h_{TES}$. This gives

$$\sigma_E^2 \approx \left(\frac{2k_b n \Sigma T_c^{n+1} \mathcal{V}_{puck}}{f_{abs} v_{phonon}} \right) \frac{V_{TES}}{\epsilon^2 n_{chan} (A_{QET} + A_{dead})} \quad (4.31)$$

Note that ζ is a properly function of the TES parameters as well, but has a fairly weak dependence on l_{TES} compared the QET area and efficiency. In addition, we're trying to optimize fin length for a given TES length, so we're not very concerned about this dependence. Plugging in the volume gives the final relation

$$\sigma_E^2 \approx \left(\frac{2k_b n \Sigma T_c^{n+1} \mathcal{V}_{puck} w_{TES} h_{TES}}{f_{abs} v_{phonon} n_{chan}} \right) \frac{l_{TES} \epsilon^{-2} \zeta^{-1}}{A_{QET} + A_{dead}} \quad (4.32)$$

Here we have terms which are close to independent of each other; the amount of dead aluminum in a QET will always depends on TES number, unfortunately, because it depends on the size of the unit cell, and thus there is still some dependence of A_{dead} on TES length. We do have terms separating the effect of fin length from the other effects, and we can call σ_0 the energy resolution of entire QET detector array in the case where we have perfect fins and maximal QET areal coverage, covering the puck surface:

$$\sigma_{0,single}^2 = \left(\frac{2k_b n \Sigma T_c^{n+1} \mathcal{V}_{puck} w_{TES} h_{TES}}{f_{abs} v_{phonon} n_{chan}} \right) \quad (4.33)$$

This is for a single channel; for the entire detector, the total resolution will add in quadrature, giving

$$\sigma_0^2 = \left(\frac{2k_b n \Sigma T_c^{n+1} \mathcal{V}_{puck} w_{TES} h_{TES}}{f_{abs} v_{phonon}} \right) \quad (4.34)$$

where we see that the overall energy resolution is independent of channel number for a perfect QET. This is not in fact true for a real QET, as more channels will allow more efficient energy collection.

We thus have a quantity independent of fin length which we can optimize first, and optimize fin length given the TES length which comes from minimizing this quantity in light of other experimental constraints. In addition, if we can place some sort of acceptable figure on this number, we can even determine how much TES length we're willing to give up to increase robustness of phase separation,




Figure 4.7: Left: Energy resolution in the TFN dominated limit as a function of TES length and fin length, showing the phase separation boundary shaded in grey. Right: Optimized fin length from energy resolution as a function of TES length.

which hurts our overall gain.

Retaining the fin-related quantities, we thus have

$$\sigma_E^2 = \frac{l_{TES}\sigma_0^2}{\zeta\epsilon^2(A_{QET} + A_{dead})} \quad (4.35)$$

$$= \sigma_0^2 \frac{l_{TES}}{\zeta} \frac{A_{QET} + A_{dead}}{f_c^2 A_{QET}^2} \quad (4.36)$$

or, written as signal to noise,

$$\frac{S}{N} \propto \sqrt{\sigma^{-2}} \propto \sqrt{\frac{\zeta}{l_{TES}}} \frac{f_c A_{QET}}{\sqrt{A_{QET} + A_{dead}}} \quad (4.37)$$

We can expect this to rise in the small fin limit where $f_c \sim 1$, and in the large fin limit, we substitute the earlier forms to find

$$\frac{S}{N} \propto \sqrt{\frac{l_a \zeta l_{TES}}{A_{QET}}} \quad (4.38)$$

From this we expect that maximizing signal to noise entails maximizing TES length, and choosing an optimal fin length between extremes where over-all signal to noise is maximal.

The relative signal to noise of QETs with different TES lengths can be seen in figure 4.7. The optimal fin length for signal to noise can also be seen, and what is remarkable is that over the range of TES lengths considered, the optimal fin length is roughly constant. This means we can, to first order, consider the two dimensions separately. Figure 4.7 also shows that the chosen optimum is not strictly the energy resolution optimum, which can be understood by referring back to figure 4.5. The chosen QET dimensions have a much larger overall aluminum coverage (35%) than the strict




Figure 4.8: Left: First-pass phonon collection efficiency in Ge as a function of fin length for different TES lengths. In the limit of very long fins, the collection is diffusion limited, while in the short-fin limit it is linear, reflecting an approximately perfect fin collection. Right: Signal to noise versus first-pass collection efficiency showing that large gains in position sensitivity can be achieved by choosing slightly higher energy resolution.

optimum, and as a result have a much larger total phonon collection fraction as a result of the higher live/dead ratio. This will lead to a much larger first-pass collection efficiency, as shown in figure [4.8](#).

4.2.4 Expected QET Performance

The optimized QET dimensions arrived at by the process just described can be found in table [4.2](#). Also shown are the parameters for the iZIP7 for comparison. The final design is expected to have an energy resolution of about 15% and an aluminum coverage of 35%, meaning it will have a much larger degree of position dependence than the iZIP. The trade-off between this optimization and the iZIP optimization can be seen clearly in the table; one sacrifices signal bandwidth to regain the same efficiency for a larger energy resolution when capacitance has to be taken into account.

Given the ideal modeling of this optimized QET, we can now explore how this design can be expected to work under various operating conditions, and as a function of the T_c we're able to achieve during fabrication. This will inform our intuition about likely failure modes and give us a sense for how much fabrication flexibility we have for the CDMS HV detector.

Figure [4.9](#) shows the noise equivalent power for both the Ge and Si HV detectors, which are the same due to the fact that both designs have the same R_n and readout. The difference comes in the form of the signal fall-time, with the Si detector integrating out to higher frequency due to the faster phonon collection time. The other panel shows the expected resolution as a function of T_c for both detectors, demonstrating that the Si detector has better resolution for all T_c values and is impacted less by the other noise sources than is the Ge design. The result of this plot is to show that while the Si detectors should achieve around 7 eV at our target 45 mK T_c , the Ge detectors

Parameter	Design Value	
	CDMS-HV	iZIP
Crystal Temperature	<30 mK	
TES Parameters		
Length	200 μm	155 μm
Normal State Resistance	150 m Ω	
Operating Resistance	50 m Ω	
Loop Inductance	\ll 500 nH	
Shunt Resistance	5 m Ω	
Parasitic Resistance	< 5 m Ω	
$\alpha \left(\left. \frac{R_0}{T_c} \frac{dR}{dT} \right _{I_0} \right)$	\sim 150	
$ \beta \left(\left. \frac{R_0}{I_0} \frac{dR}{dI} \right _{T_c} \right)$	< 0.3	
T_c	40-45 mK	40-60 mK
Risetime (L/R)	2-3 μs	2-4 μs
Falltime (τ_{TES})	30-40 μs	10-40 μs
QET Parameters		
Geometry	“Stadium”	“Linear”
Fin Length	240 μm	80-110 μm
Trap Geometry	“Semicircle”	“Rectangle”
Trap Length	20 μm	5 μm
QET Number	\sim 1800	\sim 1400
Energy Efficiency (ϵ_E), Ge	15%	13%
Energy Efficiency (ϵ_E), Si	22%	19%
Aluminum Coverage	35%	4%
Phonon Falltime (τ_{phonon}), Ge	200 μs	1400 μs
Phonon Falltime (τ_{phonon}), Si	40 μs	300 μs
Charge Input Capacitance	N/A	\leq 300 pF
Charge Channel	100-180 pF	
HEMT Input	100 pF	
Parasitic	20 pF	
Charge Collection Efficiency	N/A	95%

Table 4.2: Detector design parameters table reproduced from the SuperCDMS SNOLAB Design Report. This also includes numbers relevant to iZIP for comparison purposes to the CDMS HV design.




Figure 4.9: Phonon energy resolution expectations for optimized CDMS HV QETs. Left: Noise equivalent power by source, including thermal fluctuation noise, TES Johnson noise, bias circuit (passive) Johnson (PJN) noise sources, and SQUID noise. Given the similarity in TES parameters, the noise sources are similar, but the iZIP phonon pulse is slower, hence the lower-frequency pole in the (arbitrarily scaled) signal curve. Pulses are ≈ 4.5 faster in Si due to higher transmission probability and faster phonon velocity, which is the main driver of improved energy resolution between the two substrates. Right: Expected energy resolution as a function of sensor T_c , including the effect of various noise sources, for both Ge (solid lines) and Si (dashed lines). “Cold Tower” includes only TES noise, “Real Tower” includes passive Johnson noise, and “white SQUID” includes the SQUID white noise of $4.5 \text{ pA}/\sqrt{\text{Hz}}$ without any rising low-frequency component. The “SQUID + $1/f$ ” line includes low-frequency noise rising as f^{-1} , normalized to $10 \text{ pA}/\sqrt{\text{Hz}}$ at 100 Hz . The up-turn in the energy resolution plots at low T_c is the result of approaching the bath temperature, assumed to be 30 mK . Achieving the bath temperature goal of 15 mK would ensure the T_c^3 trend continues (in the context of the assumed model). The gray range in the figures indicates the T_c targets for each detector type, determined by the energy resolution goals as shown in the plots.

do not quite meet the goal of 10 eV for nominal values. We’ll explore in the next chapter what the actual performance of initial prototypes turned out to be.

In addition, figure [4.10](#) shows the impact of different SQUID noise parameters for the Ge CDMS HV detector. In a sense this is actually a study of the impact of low-frequency noise of different characteristics on the detector resolution, both in the energy and timing resolutions. We see that an increase of 50% at 100 Hz in SQUID noise does not matter for a $1/f$ power law, but for a shallower slope directly results in a 50% increase in resolution. The timing resolution is more robust to SQUID noise changes because it is primarily sensitive to high frequency noise. The timing resolution for the Ge detector is limited by the TES bandwidth, which decreases as temperature is lowered.

As a final note, I want to point out that these resolutions assume that all channels are added to produce energy and timing estimates, and that the channels are fit independently. CDMS has historically not achieved much improvement from joint channel fits, but in a detector with extreme position dependence, it will likely benefit us to begin to implement position-dependent joint energy




Figure 4.10: Impact of different SQUID noise parameters for energy (left) and timing (right) resolutions for the Ge HV detector. The impact on the Si HV design is less given the intrinsically larger bandwidth. These plots also assume $T_{bath} = 20$ mK, and show the impact of a lower bath temperature on both quantities as compared to Figure 4.9.

estimators. A simple illustration of the impact of this change is to imagine that all of the energy is in a single channel. If we only use that channel, then the resolution decreases compared to the channel sum by $\sqrt{12}$, resulting in a significant improvement. If we add all channels, we're just adding 11 channels worth of noise. I'll explore this more in the next chapter in the context of the G124 results, but this is an important point that I felt needed to be put here as a caveat on these projections.

4.3 HV Detector Simulation

Even with an optimized QET, we have some additional design freedom in terms of how we arrange the bias lines on the mask, how the voltage bias is applied, how the detector is housed, and how the channels are arranged. To this point in this chapter, we were able to perform analytic calculations for our optimization, but these remaining questions are best answered through simulation. In this chapter I will first discuss COMSOL electric field modeling which addresses the first three points, and then the Detector Monte Carlo (DMC) studies of detector performance to help determine the best phonon channel segmentation².

4.3.1 COMSOL Electric Field Modeling

For all CDMS detectors, the electric field shape and critical features are important for predicting the behavior and failure modes of a given detector design. For the iZIP, the scalloped fields near the surface define the charge fiducial volume, and for the high-voltage detector the uniformity of the

²Many more plots can be found in the ebook note at http://titus.stanford.edu/cdms_restricted/kurinsky/HVMask/index.html

field is important to ensure that all events in the fiducial volume undergo full NTL amplification. In general, there will always be regions of the crystal which do not behave according to our idealized model, so having an accurate electric field map is crucial to understanding non-ideal detector performance.

COMSOL is a toolkit setup to solve finite element models for a variety of physical systems. Finite element modeling is one numerical tool for solving a system defined by a differential equation at a given point. In our case we want to solve the electrostatic Maxwell's equations for the electric field:

$$\nabla \cdot E(x, y, z) = \frac{\rho(x, y, z)}{\epsilon_r(x, y, z)} \quad (4.39)$$

$$\nabla \times E(x, y, z) = 0 \quad (4.40)$$

where each electrode in the geometry is held at a known voltage, and the detector housing is held at ground relative to the electrodes. We want to include the relevant real-world properties of each material used (including conductivity and relative permittivity) where, in our case, we just case whether a material is a conductor or an insulator, since we assume all components of the model are in their time-dependent equilibria. For this reason, we simulate the detector mask and copper housing as perfect electrical conductors to simplify our calculations.

The most challenging aspect of creating a finite element model with our geometry is the relative scale of the features involved in our designs. If we were to simulate all features, we'd need a mesh that could resolve micron-scale features while creating an electric field model for a cm-scale object, a 4 order of magnitude scale. The meshing utilities in COMSOL only allow a mesh to vary by an order of magnitude in scale at the most. To resolve an entire CDMS HV detector will all mask features with a tetrahedral mesh would require at minimum a 30 Tb file for just the crystal bulk if only the very surface layer could resolve all electrode features. This is obviously a prohibitive calculation, but that degree of accuracy is also unnecessary. Given an energy resolution of even 10 eV, a 100 V map precise to the 1 V level will look identical to the exact voltage map for all events up to 100 eV if systematics are completely one-sided; if we want to go to 100 keV, then we know we need a map precise to $\sim 10^{-3}\%$, or 1 mV. This very detailed map (assuming very irregular field) would be around 100 Gb, still prohibitively large.

We can reduce the field to a manageable size by employing a few tricks to justify using a coarser grid. First, we know that the electric field from two wires looks like that from a single wire for distances much further than their separation. In kind, we can also show that the electric field from a set of infinite parallel wires looks like an infinite conducting plane for distances much further from the wires than their separation. For the HV detector, we can thus expect that the field more than a few times the mean QET separation is the same as that for a parallel plate capacitor, so it may be possible to assume a large fraction of the inner portion of the detector adheres to this expectation, and only the surfaces need to be simulated.




Figure 4.11: Geometry use to study the axisymmetric (azimuthal symmetry) COMSOL field model, given that an accurate 3D field which fully resolves the electrodes is not possible. The models studies with this geometry are also simplified by assuming a uniform plane on one side of the detector to reduce the total mesh size by a factor of two.

Azimuthally Symmetric Field

To begin to explore symmetries in the HV detector field, we constructed an approximate HV detector model assuming an azimuthally symmetric geometry of large and small electrodes on the top and bottom of the detector, with a copper housing separated by variable distance from the edge of the detector, and electrodes on the detector extending to 2 mm from the crystal edge. We focused only on Ge for these studies as it has a larger relative permittivity ($\epsilon \sim 16$) than Si ($\epsilon \sim 11.2$) [10]. To approximate the distribution of the electrodes on the surface, the rings were alternately thin (10 microns) to simulate bias lines and thick to make the total aluminum coverage match the HV design (35%). An detailed overview of the geometry used is shown in figure 4.11³. We used this model to study the impact of housing spacing and fractional aluminum coverage on the maximum electric field along the surface of the crystal, in order to determine how to mitigate surface breakdown risk. We then used these models to study the impact of different design choices on field uniformity.

The first result, shown in figure 4.12, is that the fringing field between two electrodes at the detector surface is a function *only* of the total aluminum coverage. It's straightforward to see why this is the case; if we imagine the detector as a parallel plate capacitor, we know that the total charge is going to be Q and for a 100% filling factor we know that $\sigma = Q/A$. Gauss's law tells us that as long as the charge inside a Gaussian rectangular prism of dimensions $d^2 \times l$ through the plane of the plate is constant, the electric field will also be constant. So we can show that for wires

³An extensive exploration of these models and more information can be found at <https://confluence.slac.stanford.edu/display/CDMS/COMSOL+Grid+Models>

spaced a distance d apart of width w , we have

$$d^2 E = \frac{\sigma dw}{\epsilon} \quad (4.41)$$

$$\sigma = \frac{\epsilon E}{w} \quad (4.42)$$

If we write the filling factor as

$$f_a = \frac{A_{wires}}{A_{plane}} = \frac{wd}{d^2} = \frac{w}{d} \quad (4.43)$$

then we can write the charge density σ as

$$\sigma = \frac{\epsilon E}{f_a d} = \frac{\sigma_0}{f_a} \quad (4.44)$$

where $\sigma_0 = Q/A$ is the parallel plate limit. So the charge density is only dependent on the average aluminum coverage, and thus the maximum field, which occurs near every wire, will only be dependent on the fractional aluminum coverage. This is of course only valid if there exists a limit $l \gg d$ to make this argument exact. The finite element solutions agree perfectly with this functional form as shown in figure [4.12](#). The upshot is thus that if we know a given detector technology will experience a breakdown at a certain field strength, we can determine the desired operating voltage and adjust the aluminum coverage to ensure that maximum electric field is well under that near a wire with the given charge density. Alternatively stated, if a given detector breaks down at voltage V , doubling the aluminum coverage should also double the breakdown voltage to $2V$ to a good approximation.⁴

We also experimented with the effects of housing spacing and geometry on the fringe fields and fiducial radius of the detector. The latter term is defined as the detector radius within which all events experience the full NTL gain. Figure [4.13](#) shows detector cross-sections where the electric field meets the side-wall midway down the detector side, but ends up far from the edge of the detector near the face. This effect was also seen in CDMSlite data, and is a known systematic that leads to reduced fiducial radius. The figure shows that fiducial volume is greatly increased by increasing the housing separation even by a few hundredths of an inch, and that the maximum electric field near the crystal edge is also decreased. The comparison of the current housing geometry and the proposed upgrade geometry can be seen in the lower two panels of figure [4.13](#).

Coarse 3D Field

We were able to get a simple version of the 3D detector model to run, though not in an upgraded SNOLAB housing. The difficulty lies in the size requirements discussed at the beginning of this

⁴Blas pointed out that for an infinitely thin film of width w , the field and charge density are technically infinite at the edges. The analytic solution assumes cylindrical wires with finite radius, and the numerical solution modeled the actual film thickness. The match between the two thus demonstrates that the scaling relations hold even if there is a relative geometric scale factor between the solution.

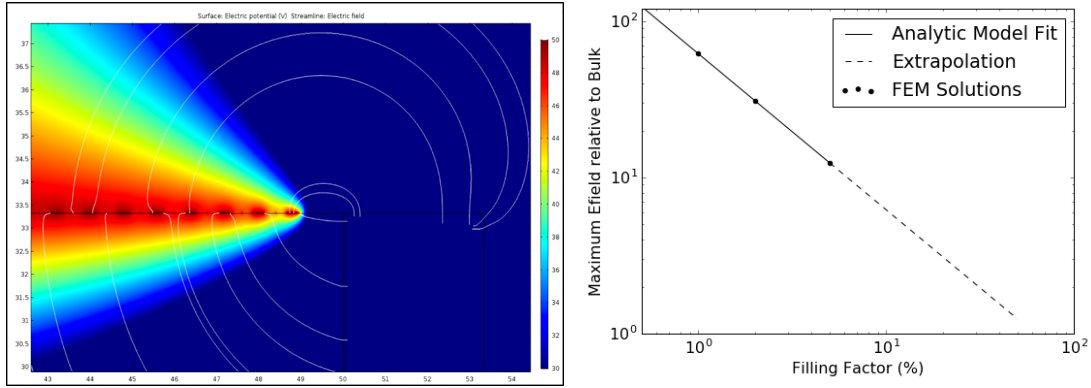


Figure 4.12: Left: Electric field in axisymmetric simulation around a close-up image of two wires. Right: Maximum electric along the detector surface as a function of the fraction of the surface covered in aluminum.

section; the discrepancy between the crystal and mask scales is too large for the SNOLAB detectors to generate anything close to a scale-accurate model. While the field near the mask may not be entirely accurate, a very good approximation can be made by constructing the mask out of concentric rings with the correct surface coverage fraction. Even this model could not be fully resolved near the bias lines, which were made (as in the 2D model) to approximate the QET rings in terms of aluminum coverage. It was also made using an outdated housing model, which does not curve following the edges (as it does in the SNOLAB design) but has a hexagonal inner profile.

Cross-sections of the 3D models can be seen in figure 4.14 for models made of concentric rings and grids, both set to 20% aluminum coverage to simulate the coverage of an early HV prototype. Figure 4.14 explores the meshing limitations near the electrodes, as well as the impact of two-sided versus one-sided bias, and the effect of biasing the housing lids (or equivalently, the adjacent detector faces) to the same bias as the sensors. From these simulations, we can conclude that two-sided symmetric bias will produce the most regular field, with the largest fraction of events that experience full NTL gain, and that allowing the adjacent surfaces to float separately from the housing or biasing them to the same voltage and the detector face leads to much less extreme surface fields. It's likely that if we're dominated by field-dependent surface leakage, this will be an important factor in achieving higher voltage bias than for the case of a fully grounded housing.

Future Work

While the studies above were very informative, it still remains that a full HV detector field model has not been constructed due to the inherent limitations of the COMSOL meshing tools. As of this writing studies are ongoing to develop a standalone FEM code to simulate these fields with more meshing freedom, as well as to directly build detector models from mask files and CAD housing

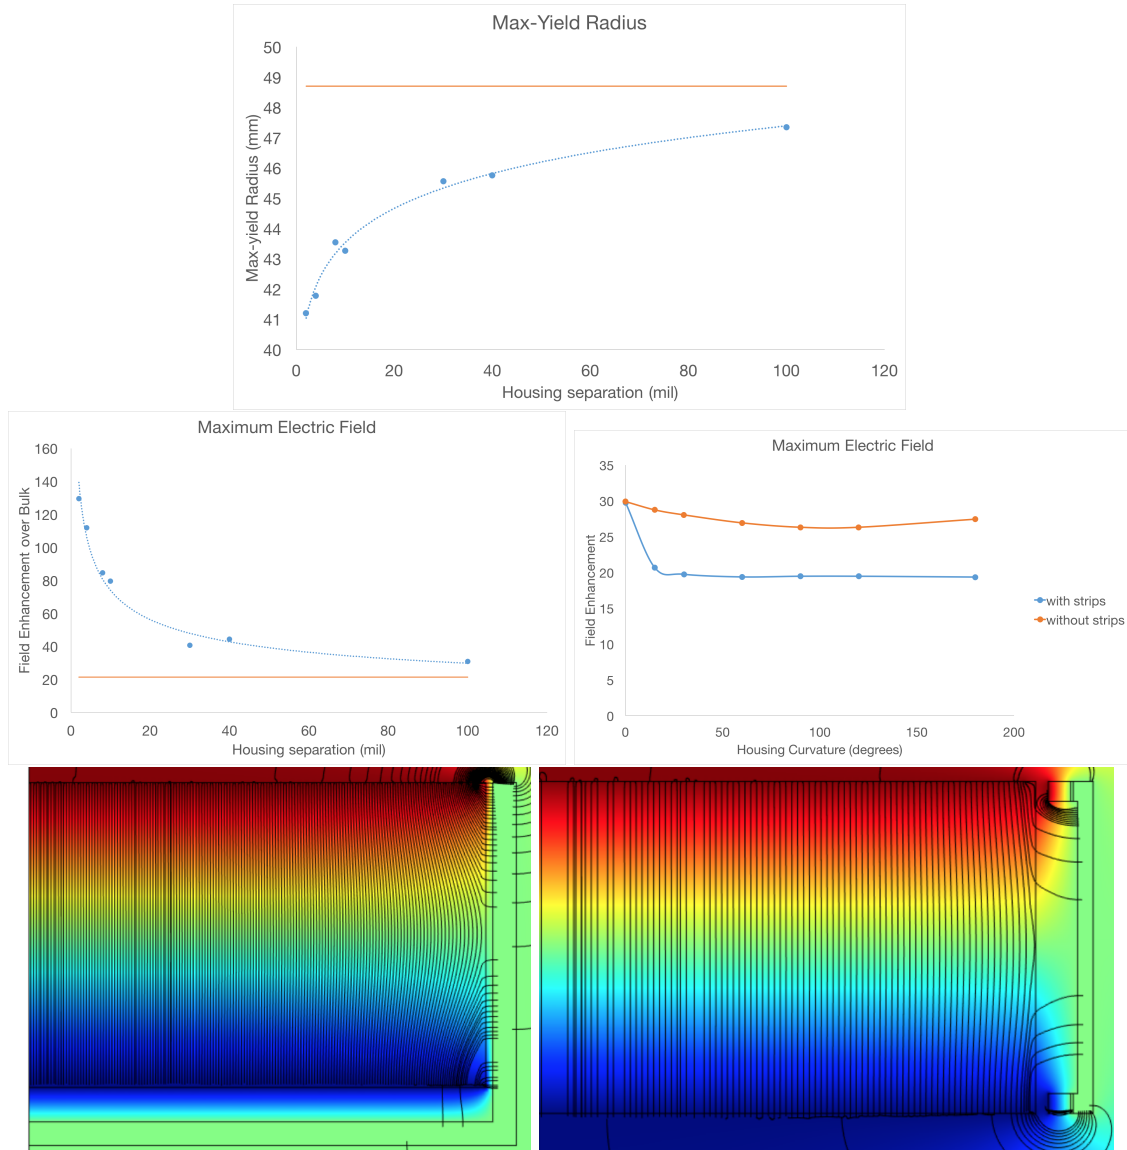


Figure 4.13: Top: Radius within which all collected events will experience full NTL gain as a function of housing separation. Center: Maximum electric field as a function of housing spacing (left) and housing wall curvature (right). The right plot also includes the effect of adding strips near the detector faces with the same voltage bias (e.g grounded to the crystal face) to shield against high electric fields at the crystal edge. Bottom: Comparison of the nominal housing design (left) and the suggested improved housing design with large dielectric buffers pushing the grounded housing side-wall away from the biased edge strips.

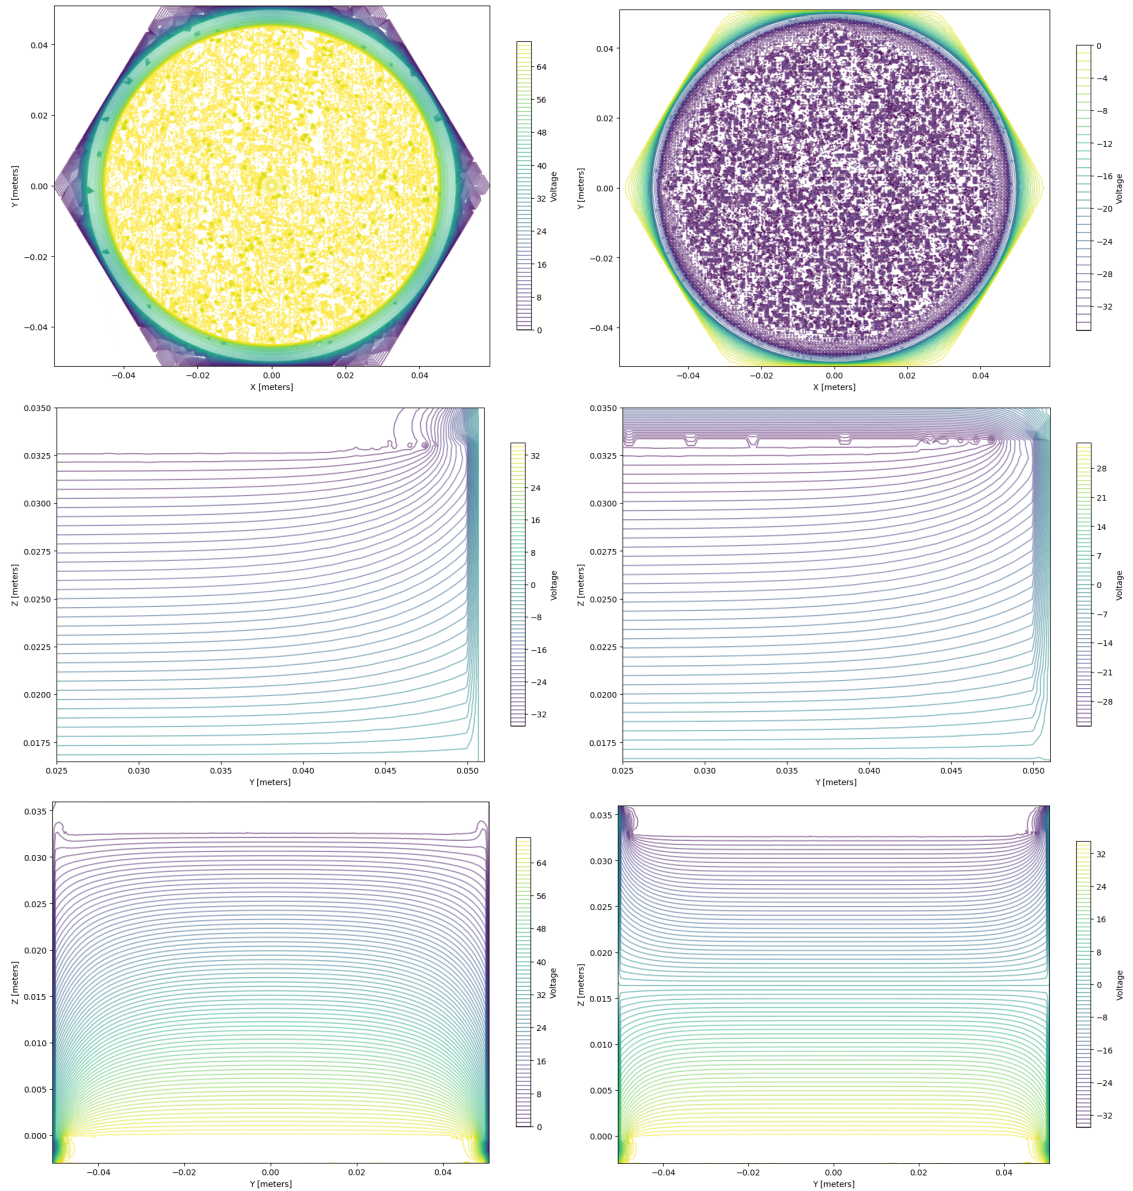


Figure 4.14: Top: XY slice through the top of the detector for a 1-sided 'Ring' model (left) and 2-sided 'Grid' model (right). Neither model can be fully meshed even in the finest mesh setting due to the size of the model, but the 'Grid' model has fewer mesh errors. Middle: Comparison of yz plane slice, zoomed into the outer corner, for two-sided bias where the housing lids are either held at the same bias as the detector face (left) or held at ground (right). Bottom: Comparison of one-sided and two-sided bias showing the field symmetry in the two-sided case, and that a larger fraction of the detector will have events with full NTL gain due to the effect the housing has of bringing lines to the sidewall.

drawings. We have been exploring the <https://fenicsproject.org/> python finite element package as well as the open-source <http://gmsh.info/> meshing tool in order to determine how detailed we can make our models in a more flexible development environment. The benefit of these tools is many more meshing options, and easier customization of boundary conditions to avoid meshing regions outside of the crystal. Though we are attempting detailed 3D models, it is likely that we will either generate a very simple model similar to a parallel plate capacitor to approximate the HV detector, and study deviations from this ideal field structure using axisymmetric models similar to those shown here. With precisely determined ring spacing and aspect ratios, we can approximate the true field very well. For more details see Madison Matsen's work on this topic.

4.3.2 Charge and Phonon Simulations

The DMC discussed in the introduction to chapter 2 has historically existed in two forms, one implemented in MATLAB [63] and a second implemented as an extension to GEANT4 (G4CMP [22, 3]). In this section I present results derived from the MATLAB DMC, but in the rest of this thesis all DMC results come G4CMP. The main difference between the two and speed and flexibility; the MATLAB DMC laid the groundwork for G4CMP but was slower and specific to Ge in a certain crystal orientation with a fixed set of valleys. The MATLAB DMC was sufficient for simulation new designs in a Ge detector, which is the focus of this section.

The MATLAB DMC takes as input an electric field map, and the position, recoil energy, and recoil type of an event. It then creates an initial population of charge carriers and phonons. The phonons are propagated until they either generate new phonons through down-conversion or are absorbed by a sensor. The charges are propagated, and as they propagate radiate new NTL phonons, simulating the detector voltage gain. For more information on the specific detector physics implemented in the MATLAB DMC, see ref [63]. For this study the final charge positions were discarded, and the map of phonon hits on the top and bottom detector faces as a function of time were recorded in order to build phonon traces, which were then converted to TES responses, with TES noise added using the analytic noise PSDs. In this section I will describe the simulations done and some initial results, and then I will show the studies performed using these simulations to optimize the detector mask in the following section.

Two sets of DMC inputs were generated, one to establish energy yield as a function of initial event position, and one to do the fiducialization studies of the next section with a variety of energies. 10,000 events were distributed evenly through the crystal, the first sample with energy 100 eV, the second with the energy spectrum seen in Figure 4.15. All simulations were run with the two-sided bias, without housing top and bottom bias, corresponding to the E-field in the lower left of figure 4.14. The full voltage drop in these simulations was 70V, as simply scaling the potential files up to 100V caused the DMC to fail inexplicably. The aluminum coverage was set to 35% uniformly on either side of the crystal, and the aluminum absorption probability was set to 0.33.




Figure 4.15: Histograms of input event energies for a 1 GeV (left) and 10 GeV (right) spin-independently coupled dark matter particle assuming standard DM kinematics and halo profile.

The MATLAB DMC has a built in limit of 400,000 individual phonons, which causes a job to fail if the number of generated phonons exceeds this limit. The DMC does not actually simulate each phonon, but limits the total phonon energy simulated to a preset number and scales the unit phonon energy accordingly. Given that the high-voltage detector will produce a much larger number of NTL phonons, we need to determine the appropriate scaling to keep the DMC phonon number below this limit.

The nominal energy scale (the `phononEnergy` parameter) is 240 eV, so the range 60-480 was chosen and simulation of 1000 events distributed over the range of energies of a 3 GeV WIMP and throughout the crystal were run. I found that failure due to excessive phonon number begins to occur sharply above 180 eV, and thus chose this value for my simulations. Given that the high voltage detectors need to be sensitive to WIMPS below this mass, and lower energy NRs will produce fewer NTL phonons, this should suffice for all of my WIMP simulations.

The energy distribution bottoms out at the bandgap energy of Germanium, however the DMC will not produce events below 3 eV. This is due to the fact that the high energy limit for EH pair generation was employed, which intersects a single pair produced at 3eV incident energy. The true curve allows a single pair to be produced for energies down to the band-gap energy threshold, so for us to generate the lowest energy events, we simply need to boost the number of 3eV events generated. We would additionally need to subtract off the additional energy imparted to the phonon system, however the NTL gain makes this phonon energy mostly negligible, so to first order it isn't necessary to generate true minimum energy events. We thus probe the lowest energy events available to us by generating extra 3 eV events, but in reality gain little by this boosting, and we merely cap energy at 3 eV instead of the bandgap energy as the relative proportion of these events is already very high.

The first result of the simulation, using the monoenergetic sample, is to determine the energy yield as a function of initial scatter position in the crystal. The plot of fractional energy yield as a

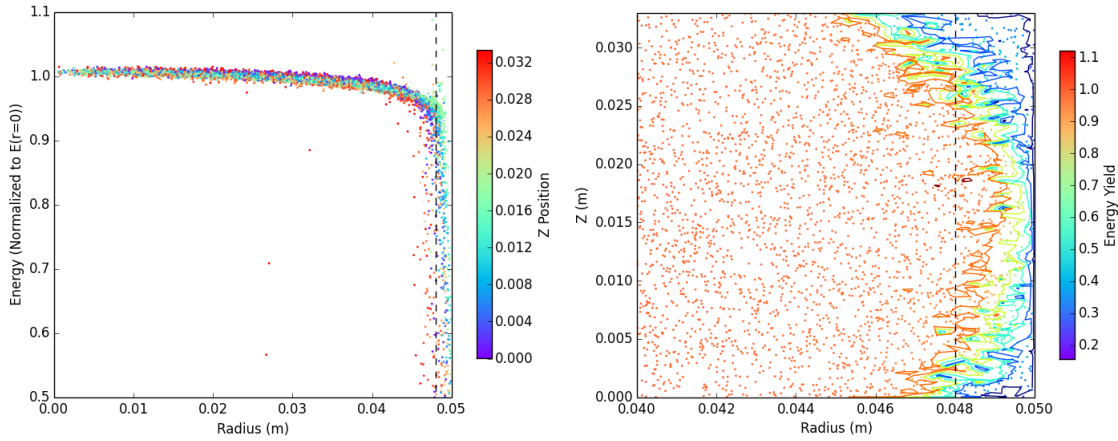


Figure 4.16: Energy yield of events as a function of their initial position in the crystal. Left: Yield versus radius colored by z -position. Right: Radius versus z position near the edge of the crystal with yield shown in color scale. The vertical black line shows where the electrode ends. The result is that the NTL gain is a function of radius and there is radial- z covariance at high radius, limiting the full gain region to about 5 mm from the crystal sidewall.

function of radius, with z position colored according to the scale, can be seen in figure 4.16. Here, the vertical line denotes the point at which the aluminum coverage ends on the surface of the crystal, and thus where the electric field becomes non-uniform. The nice effect is that yield in the middle of the crystal is largely independent of z position. If we change the view to look at the R-Z plane, and color energy yield, we see a non-trivial correlation in the plane which follows the non-idealities in the electric field very closely. The zoomed in view shows us the drop in yield more finely. We can use this plot to define a radial fiducial volume which excludes the last 5mm of the crystal, with the remaining events ensured to have at least 95% energy yield.

It is apparent from these results that the more uniform we can make our field, the larger the volume we can fiducialize in the crystal, but only to the limit that we remain a certain distance from the surface. The mean free path of high energy gammas, the largest surface background, is roughly 3mm [111], so even with the housing at ground, the two-sided bias allows the majority of the region inside 3 mm in the crystal to be fiducialized. Any improvements that increase the uniformity inside this radius are worthwhile, but aside from that, additional E-field shaping is not necessary for greater fiducialization. More uniform fields would, however, increase the gain uniformity across the crystal.

4.3.3 TES Pulse Simulation

We implemented our own TES simulations instead of using the DMC TES simulator for these studies for a few reasons. Primarily, this allowed us the flexibility to easily and quickly modify the phonon




Figure 4.17: Sample of random noise generated according to the expected TES noise profile in the time domain (left) and frequency domain (right).

sensor layout. It also allowed us to ensure the proper TES response was employed, as the at this point DMC had not been validated for the CDMS HV TES response, and this allowed for a much simpler and streamlined simulation process. It also allowed us to easily integrate the projected SNOLAB TES noise into the TES pulse generation before optimal filtering was applied. The expected and simulated noise PSDs, as well as a time domain random noise sample, can be seen in figure [4.17](#). Note that these are normalized to white noise with magnitude 1; the true current noise picks up a factor of

$$\sqrt{\frac{4k_B T}{R_0}} \approx 5 - 10 \text{pA} \quad (4.45)$$

depending on the operating resistance of the TES. The use of an accurate noise model is crucial to understanding how the noise impacts the TES pulses, and to measuring the effective phonon resolution from these simulations.

The TES pulse simulation process is illustrated in figure [4.18](#) for a large pulse in a channel above the event dominated by NTL phonons, and a small pulse primarily containing ballistic phonons. We histogram the phonons into bins of width 0.5 microseconds. We then convolve this histogram with the TES impulse response corresponding to our TES parameters, show in blue in the figure. Note that overall normalization is removed to be able to show all of these pulses in the same scale. We then add TES noise, and apply an the wiener filter described in section [E.7](#) with an explicit 100 kHz roll-off to remove phonon shot noise.

The result of this filtering is that most pulses can be accurately reproduced after deconvolving the TES response out of the optimal filtered TES pulse. This is stand-in method for the eventual process of fitting a template library to the optimal filtered phonon pulses, but for this study give a nice approximation to the phonon pulse we obtain after removing electronics and shot noise. It also allows us to estimate the energy resolution using a method that approximates the best version of




Figure 4.18: Simulated phonon signals given an input phonon energy arrival histogram showing the convolution the the TES response, adding noise, filtering with a wiener filter, and deconvolution to compare the estimated and input pulses.

the optimal filter, for an infinite template library with perfect noise and signal characterization.

From the two optimal filtered pulses, we generate analysis quantities to be fed into the subsequent fiducialization processes. The four quantities measured for each pulse are time and height of the pulse peak, as well as total pulse energy (the pulse integral) and "Luke Phonon" (NTL phonon in this thesis) energy, defined as the amount of energy collected during the initial 15 microseconds of the pulse. This time is approximately the time it would take for a phonon to traverse the longest dimension in the crystal, so in this way we approximately count only the phonons collected upon their first reflection. To remove the energy dependence from the individual channel quantities, we sum the quantities across all channels and divide by the sum, to give "fractional" values (pFrac, pLukeFrac), and include also the total as an analysis quantity. The time and peaks are not normalized at this stage of the simulation.

The measured resolution for the total energy measurement can be seen in figure 4.19 for both DM samples. We can see that they are consistent with one another and slightly below our 10 eV expectation due to the fact that the wiener filter can account for position dependence, and effectively de-weight channels without much signal. This is an indication that in these detectors we're going to benefit from algorithms that can jointly fit all channels in a position dependent way.

4.4 Phonon Channel Layout Optimization

Using these simulations, we can test, for the optimized CDMS HV QET, a variety of choices for how to lay out the 12 phonon channels on the detector face. In this section I describe the different fiducialization studies done for the various geometries. The selection of geometries to test is described in section 4.4.1. I then explore the position dependence in the detectors in sections 4.4.2 and 4.4.3




Figure 4.19: Left: Energy resolution for all events in the samples for the different mass dark matter particles. Right: Energy resolution for the 1 GeV DM for the configurations described in the text of section 4.4. X-axes and distribution moments are in eV, histograms are normalized.

and compare the total fiducialization power of the different configurations in section 4.4.4.

4.4.1 Proposed Geometries

Despite having selected a channel resistance and QET design, we have a lot of freedom in our choice of phonon partition given that we have twelve channels to cover two detector faces. A very early decision was made to have 6 channels on both faces so as not to introduce more position dependence than we already expect to be present. We also imposed the requirements that all phonon channels cover equal crystal area (so that the aluminum coverage is homogeneous) and that both radial and azimuthal information is retained. This eliminates the configuration of six concentric rings as well as any configuration with rings only split in half. The remaining options are summarized in figure 4.20, and are referred to by the numbers in the figure⁵

These options essentially boil down to whether there is an innermost spot, and how thick the outer ring is. Configurations 3-6 all have the central spot, but change how the remaining channels are arranged. Configuration 0 was the one originally proposed in the SuperCDMS SNOLAB conceptual design, while configuration 3 is the same as the phonon channel layout in the iZIP. In all cases, the opposite face would look similar but have the channel splits rotated by 60 degrees for a 3-channel ring design and 45 degrees for a 4-channel ring design to ensure that if a signal lies on the channel split for one face, it will be the center of the channel on the other face, increasing the probability that it's accurately reconstructed.

The first result from simulating these configurations was already shown, namely that in the small signal limit, the resolution is unaffected by the layout of the channels. As long as we can

⁵All figures in this section for all configurations and more voltages can be found at http://titus.stanford.edu/cdms_restricted/kurinsky/HVMask/index.html

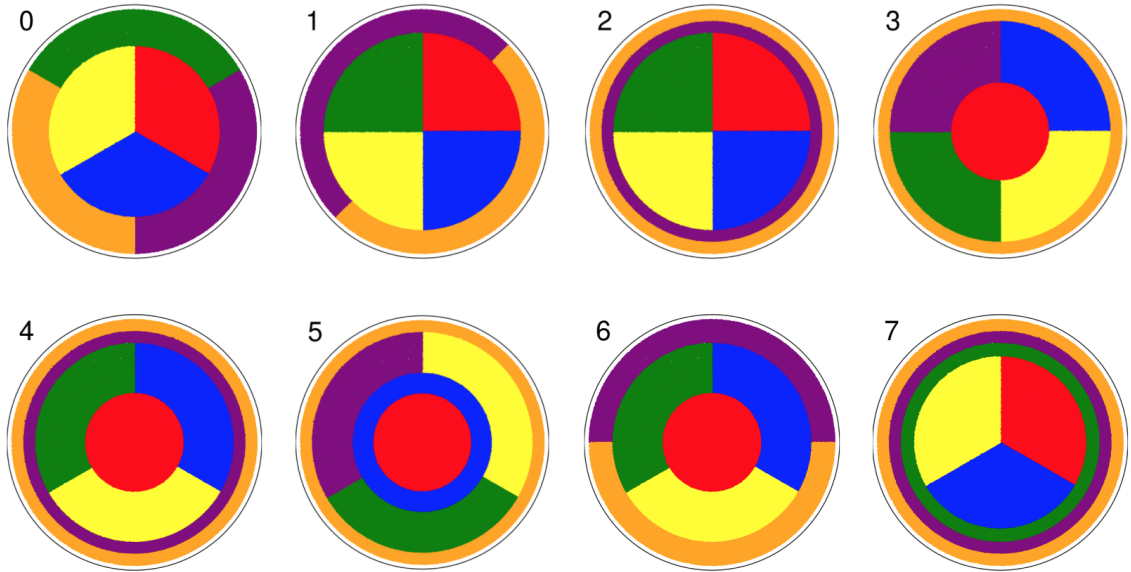


Figure 4.20: Proposed detector geometries explored for the HV mask using the DMC. See text for more details.

perform a good optimal filter reconstruction of the pulses, this will be true for signals close to the detector threshold, but the TES simulation used does not account for partial saturation of a TES channel (where some but not all of the TES unit cells are driven normal), which decreases the power to current conversion and drives the resolution up for higher energy events. These DM masses are low enough that saturation is not of primary concern, and the events are high enough in energy (1 keV or larger) that they're very far above threshold, but the effects of saturation should not be underestimated for these detectors, which are much more position dependent than previous SuperCDMS detectors.

We can also look at how the NTL phonons are distributed in the channels for each configuration. Figure 4.21 shows the phonon hits colored by channel number on each side of the detector for configuration 0 for an event of moderate energy, at high radius and midway through the crystal. We immediately notice two things: first, the electron pattern is larger than the hole pattern (a result of the oblique propagation of the electrons), and second, that the electron pattern is of the scale of the outer ring, about 10 mm. This implies that narrower rings should share more of the phonon distribution, and be more sensitive to the precise initial position of high-radius events. We'll see if this is the case shortly, but it's an interesting observation that makes intuitive sense to us.




Figure 4.21: Distribution of phonon hits on the top (left) and bottom (right) faces of the crystal in the DMC for phonons which arrive within the first 30 μs after the event, roughly the time it takes a phonon to traverse the crystal. These are the phonons with the most position dependence, so their distribution gives us the most intuition about how to maximize

4.4.2 Vertical Partition Performance

We can construct the vertical partition as defined in appendix [A](#) from the different energy quantities determined from the phonon pulses by summing all of the quantities from the top and bottom together, with the results shown in figure [4.22](#). The variable plotted is the z-partition quantity

$$p_z = \frac{E_{top} - E_{bottom}}{E_{top} + E_{bottom}} \quad (4.46)$$

where E_{top} and E_{bottom} are the respective channel sums for the energy collection in the first dynamic timescale. The models shown assume different diffusive radii, and the data at 1V (dominated by prompt phonon energy) are consistent with a slightly asymmetric diffusive radius of 10 mm.

We can model the z-partition of the higher voltage data well by assuming that the NTL phonons are emitted in a column equally spaced along the axis of propagation. This means that the z-partition will progressively get less sensitive as we raise the voltage, as each side will receive half of the total NTL energy concentrated in the phonon channels above and below the event. The model used has the functional form

$$p_z = \frac{\Delta\Omega}{\Sigma\Omega + 4\pi f_{NTL}} \quad (4.47)$$

where

$$f_{NTL} = \frac{y(E_R)V}{\epsilon_{eh}} \quad (4.48)$$

We can thus see from this equation that in the limit that $V \rightarrow 0$, there is identically zero z sensitivity in this partition quantity.

The trend seen in the one-sided partition mirrors similar trends seen in CDMSlite data^{[6](#)} though a z-partition was not constructed for that data due to the minimal gain from only removing half

⁶http://titus.stanford.edu/cdms_restricted/Soudan/R133/ebook/151208_BP/T2Z1VScan.html

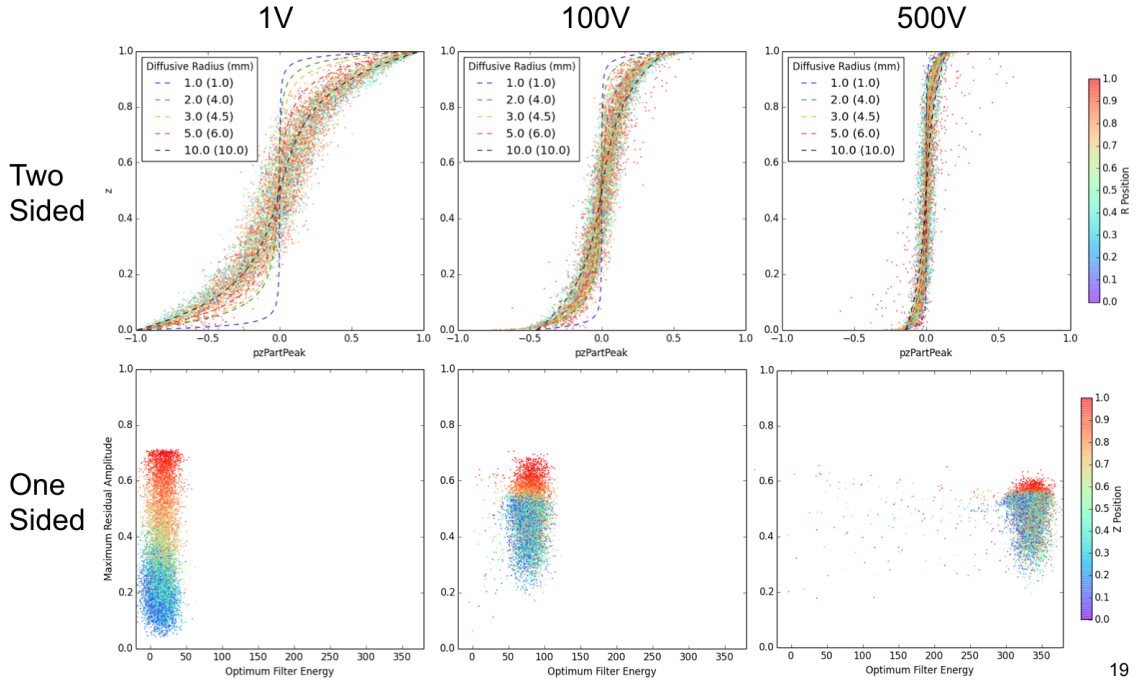


Figure 4.22: Z partition quantity for the 10 GeV events as a function of voltage for a two-sided readout (top) and a one-sided readout (bottom) for comparison to CDMSlite data. As voltage is increased, there is less z-dependence in the signal, but there is always a surface layer which can be removed by a z-partition either defined as the side difference (top) or difference between adjacent channels (bottom).

of the detector surface. This does suggest, however, that the DMC results represent a real effect, and what we’re seeing is not a DMC artifact. The apparent z-sensitivity of the HV mask was not expected due to the relative strength of the luke phonons, and the accuracy of the luke phonon simulations has still yet to be fully validated. An attempt to reproduce these effects with real data will thus be a high priority for the first CDMS HV pathfinder tests being planned at the time of this writing.

4.4.3 Radial Partition Performance

Constructing a radial partition quantity in a geometry independent way, especially accounting for z-dependence in the radial signature of events, is not as straightforward as the z partition. In order to construct radial quantities without considering each geometry individually, we perform a principal component analysis on the matrix of 12 energy or time variables and true radius, and use the principal component (assuming the highest variance will be a covariance with r) as a proxy for the radial partition. This is obviously not the ideal quantity necessarily for any geometry, but gives

us a natural means to characterize intrinsic radial dependence in each proposed mask layout.

All partition quantities can be seen in the note linked earlier. The partition plot for configuration 0 can be seen in figure [4.23](#). The figure shows that all ring splits are captured in this partition quantity, and that there is some residual z-dependence; in addition, the best partitions are found from the quantities defined based on the initial part of the signal, which we expect to be most position dependent. To characterize the performance of these quantities, we draw a Receiver Operating Characteristic curve, which draws a line through the space of 'bad' events rejected versus 'good' events accepted parametrically as a function of the cut-value employed. If we define the outer 5 mm in radius to contain the bad events we intend to reject with this partition, we find that all configurations can reject all high radius events to some extent in the peak-height defined quantity, but that the configurations with two thin outer rings perform the best. This of course does not account for z-dependence, but given that the simple z partition does not depend on configuration, we can start to narrow our choices to, for example, configurations 1,2,4, and 6. These all achieve 100% bad event rejection at 70% fiducial event efficiency in this quantity, and they are all patterns with two outer rings and four inner rings. Interestingly, configuration 7 performs worse than these despite having 3 outer rings, but at a fairly low level; it is on par with the other configurations to the 1% leakage level.

4.4.4 Combined Fiducialization

The final study done to choose the mask layout was to train a boosted decision tree on each of the variables to reject all events within the top or bottom 3 mm of the detector and the outer 5 mm in radius. A summary of the input and training results of this BDT fiducialization summary are shown in figure [4.24](#). This BDT used python's sklearn package, and specifically its AdaBoost function, and is the exact same algorithm employed by the BDT implemented in ROOT's TMVA. A boosted decision tree is given truth information (good event or bad event) as well as inputs, and attempts to combine a large number of weak binary classifiers into a strong non-linear classifier to generate a non-linear cut maximizing the ability to reconstruct truth events and exclude "fake" (surface or low-yield) events.

Training a BDT, like most multivariate tools, is highly non-trivial. The free parameters to set include learning rate, number of boosted classifiers, and fraction of events used for training and testing. This separation is necessary to ensure that the BDT is not over-trained. I observed that only 200 boosted classifiers were necessary for a convergent test error, and that the ideal fraction of the sample reserved for testing was about 30%; this gave the most convergent test error. The BDT results shown were trained on 70% of the data, and all cuts and ROC curves are as computed for the 30% test sample. Initial cross-validation results showed no dependence of performance on the specific test sample selection, so the last 3000 events were reserved for testing.

The BDT results shown in figure [4.24](#) demonstrate that configurations having at least one thin

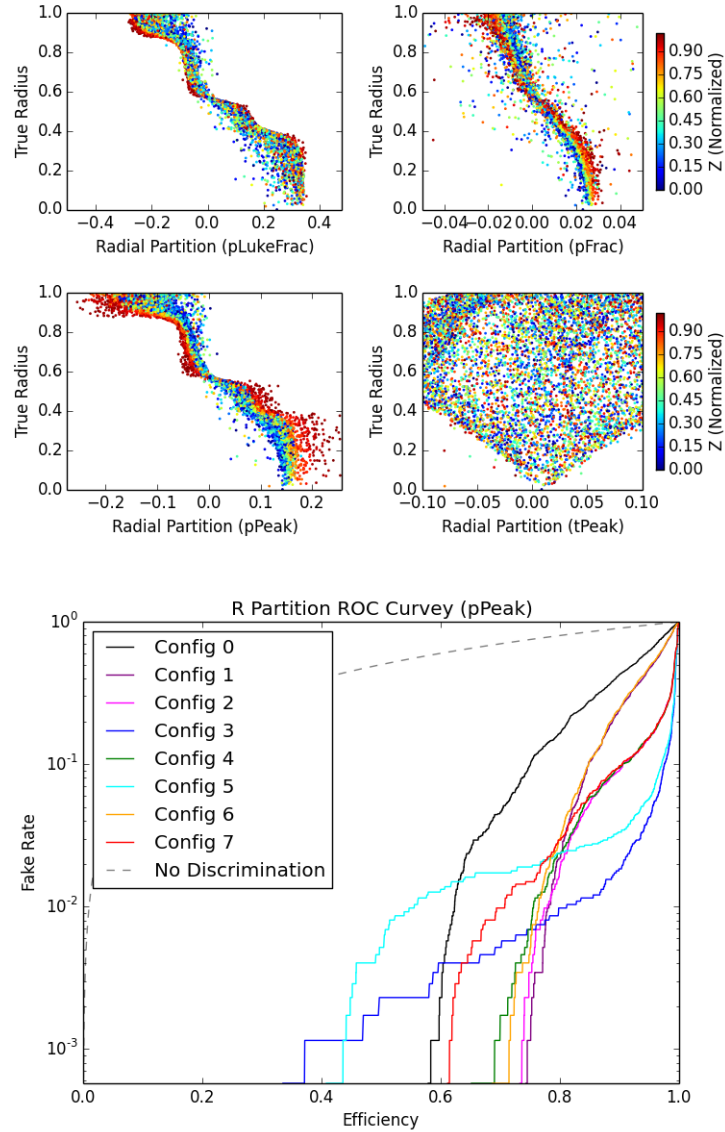


Figure 4.23: Top: First principal component for each of the four variables extracted from DMC data for configuration 5, showing that this principal component picks up all ring splits. The quantities sensitive to the height of the initial peak (the most position dependent aspect) are also the most sensitive to position; we even see some residual z-dependence, even though it is not included in the PCA training. In general this method doesn't produce a meaningful quantity for the timing information. Bottom: Receiver operating characteristic (ROC) curve for rejecting 'non-fiducial' radial events as a function of fiducial efficiency. For this study, non-fiducial events were defined as those events in the outer 5 mm in radius. All designs with thin outer rings have good rejection at higher efficiency, and those designs with two outer rings can reject all 'bad' events at high fiducial efficiency.

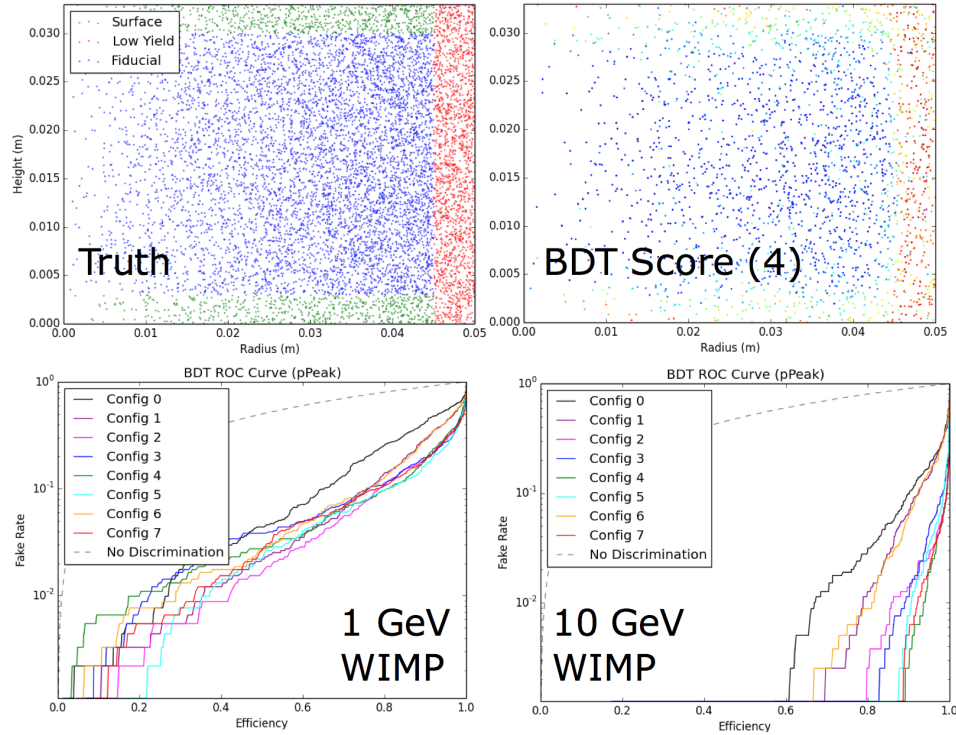


Figure 4.24: Top: All events with their truth labels (left) and the resulting BDT score for configuration 4 for the test data, with blue representing high-confidence fiducial events and red high confidence surface events. Bottom: ROC curves for the low-mass (left) and high-mass (right) samples showing that the leakage rate decreases for higher energies as is usually the case. Even for the 1 GeV DM particle, a 90% reduction in background can be achieved at 70% fiducial acceptance, while full fiducialization can be achieved for the higher energy sample at 90% fiducial acceptance. Based on the 10 GeV ROC curve, we selected the fourth proposed geometry as the final CDMS HV channel layout.

outer ring perform better regardless of the internal structure, but suggests that by a narrow margin, the configuration with two thin outer rings, one innermost ring, and three intermediate sections has the best performance (configuration 4). This is likely due only to radial effects. As a result of this analysis and the studies presented in this section, configuration 4 was chosen as the most robust and position-rich configuration with the highest likelihood of providing a good degree of fiducialization.

4.5 Mask Design Results and Future Studies

In this chapter, I described the entire optimization process for the CDMS HV detector in the order in which all CDMS detectors need to be designed. We first discussed how to determine channel properties for a TES channel based on readout noise, then how the QETs are optimized for energy

resolution, and finally how to optimize QET layout on the mask. The result is a detector which is as sensitive to phonon energy and position as we know how to design within the limitations of our current technology. The rest of the mask development process entails fabrication and wiring considerations, and tweaks to various design aspects to reduce initially observed failure modes. These include the addition of repair pads, bonding pads, and readout wiring. The outer 2 mm is left instrumented due to decreased fabrication yield and contact with the detector clamps, but an outer aluminum and tungsten grid was included in the design with minimal aluminum coverage in order to extend the electric field closer to the edge of the crystal. A breakdown of the optimized mask image can be seen in figure [4.25](#).

In the next chapter, I will discuss the testing program undertaken with various types of CDMS HV prototypes and the lessons learned from these initial tests. A key takeaway of this chapter, however, is the number of new detector features that have yet to be validated. An important next step when the first fully-functional CDMS-HV detector is run is to try to validate the DMC and show that the position dependence and phonon distributions are correct, and to further study the TES response in order to get the DMC to be predictive of real events. This will allow us to use the DMC to train more advanced event reconstruction algorithms (such as convolutional neural nets) so that we are not limited by the accuracy of our TES templates in an optimal filtering framework. This chapter should provide some guidance on what we expect to see in these detectors, but we should expect new subtleties to arise which will add to our understanding of the underlying physics in the detectors. These should be taken into account in the next round of detector designs.




Figure 4.25: Top: Top side detector mask with zoom-in regions showing a fraction of the channel and a subset of QETs. Even in the large image the channel splits and wiring can be clearly seen. Some dead metal is saved by having all QETs share bias lines with those adjacent. QET placement was done to make each channel have equal area and equally space the QETs across the detector face. Bottom: 3D rendering of CDMS HV detector showing the relative rotation of the top and bottom masks.

Chapter 5

Prototype Detector Performance

“An expert is a person who has found out by his own painful experience all the mistakes that one can make in a very narrow field.”

- Niels Bohr, as quoted by Edward Teller (6 September 1954)

In the previous chapters, I laid out the theory of operation for SuperCDMS detectors with historical measurements justifying some extrapolations made to design the CDMS HV detector. Once the design was complete, work was begun to validate each aspect of the detector performance models, both in full-scale prototypes and smaller detectors meant to test scaling relations and simplify some aspects of prototype testing. The structure of this chapter is meant to organize detector testing into types of study, focused either on noise, resolution, or TES characterization, as these are largely independent studies. In each section I will compare and contrast results from the different detector prototypes, which are summarized in section [5.1](#). I will end with a section summarizing lessons learned about backgrounds and noise, and some concluding remarks about ongoing work to continue to characterize backgrounds found and measurements made during these early tests.

5.1 Detector Prototype Program

The SuperCDMS SNOLAB project began a program of fabricating test detectors to demonstrate the HV detector performance and test the new SNOLAB cold hardware designs in late 2016. These devices are summarized in table [5.1](#). The first detector (G101c) was fabricated with a previously used detector crystal at Texas A&M in mid 2016 and was run at UMN in October-November of that year, and the subsequent findings are documented in this chapter. I will briefly recount the general discoveries of this program before going into more analysis details.

Detector	Substrate	Dim. (mm)	Mass	Institution	Comments
G101c	Ge	100x33.3	1.4 kg	UMN	CDMS HV v0
S101	Si	100x33.3	600 g	UMN	CDMS HV v1
G115	Ge	76x4	100 g	UCB, SLAC	CDMS HV QET on custom 1-sided mask
G124	Ge	100x33.3	1.4 kg	SLAC	CDMS HV v1
G147	Ge	100x33.3	1.4 kg	SLAC	CDMS HV v1
HVeV	Si	10x10x4	0.9 g	Stanford	Test Device Mask

Table 5.1: Summary of detector prototypes discussed in this chapter. This covers many of the detectors tested at UMN, UCB, SLAC, and Stanford between 2016 and 2018^[2]. Not included are the TES test chips and photon detectors also tested at UCB, SLAC, and Stanford during this period.

One of the main results of the testing program was that our devices have systematically lower R_n , higher G , and highly variable T_c than we anticipated in the design detailed in chapter 4. The latter two measurements are systematics limited, with both having a high dependence on how well the fridge base temperature is known. Normal resistance, on the other hand, can be measured reliably either in situ (with the normal TES bias circuit) or via dedicated four-wire measurements, as was done for G124 at SLAC. The in situ measurements can also be made in absolute terms (what we call IbIs measurements) or in relative terms (what we call dIdV measurements), though these are also limited by systematics on the shunt resistance. The fact that these also varied significantly from expected confirms to a certain extent that the change in TES properties is really a result of fundamental physics. In the next few sections I will describe how the measurements were made, and contrast the results from different detectors.

5.1.1 TES Resistance Measurements

The simplest way to measure the normal and parasitic resistance of the TES circuit is to capture points very low and very high in the IbIs curve, where the TES is superconducting or normal. For any point on the transition, the steady state bias current produces a steady state source current according to the relation

$$I_{s,0} = \left[1 + \frac{Z_0}{R_{shunt}} \right]^{-1} I_{b,0} \quad (5.1)$$

where for $I_b \ll I_{op}$, $Z_0 = R_p$, for $I_b \gg I_{op}$, $Z_0 = R_{normal} + R_p$, and for $I_b = I_{op}$, $Z_0 = R_0 + R_p$. This is only true if the excitation frequency is well below any of the impedance poles.

For testing purposes, we can re-arrange this expression to allow for calculations based on IbIs data or complex impedance measurements. We write $I_{\{b,s\}} = I_{\{b,s\}0} + \delta I_{\{b,s\}0}$ and take the case

with and without changing current to find that

$$I_{s,0} = \left[1 + \frac{R_0 + R_p}{R_{shunt}} \right]^{-1} I_{b,0} \quad (5.2)$$

$$\delta I_s = \left[1 + \frac{Z_{loop}}{R_{shunt}} \right]^{-1} \delta I_b \quad (5.3)$$

where $Z_{loop} = \frac{\delta V_s}{\delta I_s}$. Re-arranging in terms of measured quantities gives

$$R_0 + R_p = R_{shunt} \left[\frac{I_{b,0}}{I_{s,0}} - 1 \right] \quad (5.4)$$

$$Z_{loop} = R_{shunt} \left[\frac{\delta I_b}{\delta I_s} - 1 \right] \quad (5.5)$$

The impedance in the loop is affected by the dynamics of the transition, and we have

$$Z_{loop} = \frac{(R_p + R_0(1 + \beta) + i\omega L)(1 - \mathcal{L} + i\omega\tau) + R_0\mathcal{L}(2 + \beta)}{1 - \mathcal{L} + i\omega\tau} \quad (5.6)$$

$$= R_p + i\omega L + R_0 \left[(1 + \beta) + \frac{\mathcal{L}(2 + \beta)}{1 - \mathcal{L} + i\omega\tau} \right] \quad (5.7)$$

$$= R_p + i\omega L + R_0 \left[(1 + \beta) + \frac{\mathcal{L}}{1 - \mathcal{L}} \frac{2 + \beta}{1 + i\omega\tau_{eff}} \right] \quad (5.8)$$

where β and \mathcal{L} are as defined in section [3.2](#). In the limit that $\mathcal{L} \gg 1$, we have

$$Z_{loop} \approx R_p + i\omega L - R_0 \left[\frac{1 - (1 + \beta)i\omega\tau_{eff}}{1 + i\omega\tau_{eff}} \right] \quad (5.9)$$

such that crossing through the pole, to first order, just flips the sign of R_0 and thus the phase of the response. Measuring either side of the pole (once the transfer function is flat) yields a measurement of β . Measuring the pole compared to the natural response (no electro-thermal feedback) gives a measurement of \mathcal{L} , which once measured can be used to find the corrected β , which should be slightly smaller than the initial measurement.

If we make measurements far below the low-frequency pole, we can set $\omega = 0$ to give

$$Z_{loop,DC} = R_p + R_0 \left[(1 + \beta) + \frac{\mathcal{L}}{1 - \mathcal{L}}(2 + \beta) \right] \quad (5.10)$$

$$= R_p + R_0 \left[\frac{1 + \beta + \mathcal{L}}{1 - \mathcal{L}} \right] = R_p + \kappa R_0 \quad (5.11)$$

where we have defined the quantity

$$\kappa = \frac{1 + \beta + \mathcal{L}}{1 - \mathcal{L}} \quad (5.12)$$

This is great; this means that by doing low-frequency complex impedance and IbIs measurements

(taking the peak-to-peak and DC values of the resulting source current), we can measure R_p , R_0 , β and \mathcal{L} . To summarize, then, we find

$$R_0 = R_{shunt} \left[\frac{I_{b,0}}{I_{s,0}} - 1 \right] - R_p \quad (5.13)$$

$$\kappa R_0 = R_{shunt} \left[\frac{\delta I_b}{\delta I_s} - 1 \right] - R_p \quad (5.14)$$

For instance, if we have a sharp transition ($\mathcal{L} \gg 1 > \beta$), we will have $\kappa \approx -1$ in the transition, and will expect a phase shift in the response of source current to bias current, meaning we can use the equations

$$R_0 = R_{shunt} \left[\frac{I_{b,0}}{I_{s,0}} - 1 \right] - R_p \quad (5.15)$$

$$R_0 \approx R_{shunt} \left[\left| \frac{\delta I_b}{\delta I_s} \right| + 1 \right] + R_p \quad (5.16)$$

as long as we are confident in this criterion for the transition point we are biased at.

An example IbIs measurement can be seen in figure [5.1](#) for G115 from the second run at UCB and the first run at SLAC. These data were both taken by varying the TES bias from normal down to superconducting, with the SLAC measurements made in the opposite direction as well. At all points the operating current was measured and dIdV data were collected; these will be discussed further in the next subsection. What is interesting to note is that while the resistance curves look very similar, the current scale is vastly changed. This is the first indication that the parasitic power environment in each setup was very different even after some initial improvements at UCB, either due to bath temperature differences or true parasitic power sources.

Figure [5.1](#) also shows a comparison with the second way of estimating TES resistance. If a high energy event saturates the TES, then we can directly measure the difference in current ΔI_s between R_0 and R_n . During the pulse, we've fixed $V_b = I_b R_{sh}$, and we know that for the normal state (where $R_l = R_{sh} + R_p$)

$$I_{s,normal} = \frac{V_b}{R_l + R_n} \quad (5.17)$$

while for the operating point,

$$I_{s,0} = \frac{V_b}{R_l + R_0} \quad (5.18)$$

which gives the expression

$$\Delta I_s = \frac{V_b}{R_l + R_n} - \frac{V_b}{R_l + R_0} \quad (5.19)$$

solving for R_0 , we find

$$R_0 = R_{sh} \left[\frac{R_{sh}}{R_l + R_n} - \frac{\Delta I_s}{I_b} \right]^{-1} - R_l \quad (5.20)$$




Figure 5.1: Resistance measurements for G115 made at UCB (top, 4 QET channels) and SLAC (bottom, 2 QETs channels, A and C). The left-hand side shows IbIs measurements, and the right-hand side shows the resistance measurements derived from the IbIs. The ‘dIdV’ label in the top right figure refers to measurements taken with dIdV data, for comparison to resistance inferred from saturated muon pulses, shown as dots in the figure. These data taken from http://titus.stanford.edu/cdms_restricted/detector_physics/HV/ebook/170511c/index.html and http://titus.stanford.edu/cdms_restricted/detector_physics/HV/ebook/170628/index.html.

This is an alternate measurement which can help identify systematic errors, or can be used as a substitute when full IbIs data are not available. We can see that for this measurement, the muon-based R_0 measurement disagrees with the IbIs measurement due to additional (non-TES related) resistance which appears at $60 \mu\text{A}$, most likely a wire-bond going normal.

The normal state resistance values for all prototype detectors were actually much lower than anticipated. These measurements are summarized in Table 5.2 with the other results from this section. The largest implication is obviously that we can’t operate the TES as low in the transition as anticipated, which limits the dynamic range of the TES more. It also implies certain changes in both the heat capacity and thermal conductance. As we saw in the previous chapter, we know that

in general, the Wiedemann-Franz law implies that

$$\kappa = \beta_{wf} \frac{T}{\rho} \quad (5.21)$$

so a reduction by a factor of 2-3 in the resistivity implies an increase by that same amount in thermal conductance, which also corresponds to an increase by $\sqrt{2}$ - $\sqrt{3}$ in phase separation length. This implies that a fraction of the lost resistance can be made up for by increasing TES length in a modified HV design, which will have a large impact on collection efficiency as well. This of course assumes Σ is invariant to T_c , which also does not seem to be the case. We'll discuss this a bit later in this chapter.

The observed reduction in resistivity is not unexpected. Tungsten has two superconducting crystal phases, the α phase ($T_c \sim 15$ mK) and β ($T_c \sim 1$ K) phase, which can be mixed to produce a film with intermediate T_c . These films also have an order of magnitude difference in resistivity (β phase is about 10x higher resistivity) [44] and different lattice constants ($a_\alpha \sim 3.2$ Å, $a_\beta \sim 5.1$ Å). The first layers of a tungsten film tend to grow in the β phase, with subsequent layers in thicker films forming α phase structures as residual oxygen in the growth chamber gets depleted by gettering. Thicker films are therefore more α -like, having both lower resistance and lower T_c [61]. These films are very sensitive to material purity [112, 113], but for comparable growth conditions on similar substrates, films of thickness 40 nm such as ours should have resistance and T_c only weakly scale with thickness. This will ideally be the subject of future work as we continue to experiment with fabrication of lower T_c devices.

5.1.2 Critical Temperature and Bias Power

Measurements of T_c are fairly simple but are the most dominated by systematics. The easiest way to measure T_c is to put a small square wave (with very small bias power) on the TES line with the TES normal, and raise and lower the temperature of the fridge with a heater to observe where the TES goes through its transition. This is obviously affected by the power systematic from the TES bias, as well as any lag between the fridge and crystal temperature, but will generally give a measurement good to within 10 mK of any other method.

The most precise way to measure T_c is to measure the bias power at a few different base temperatures, and then extrapolate the power measurements to find the temperature bias power crosses zero to infer T_c . Given the TES power equation

$$P = K(T^5 - T_b^5) \quad (5.22)$$

if we use the joule power to maintain the TES at T_c , we get the relation

$$P = K(T_c^5 - T_b^5) = I_0^2 R_0 \quad (5.23)$$




Figure 5.2: Drift in the bias power for a channel on G124 as a function of time since the first observation of the channel going superconducting. We see that the crystal is cooling with a time constant of a few days, and it’s unclear if the final bias power is actually in the $T_b \sim 0$ limit.

so a simple linear fit to P versus T_b^5 given the measurements at multiple bath temperatures will yield T_c^5 as the y-intercept. Plenty of examples of this technique can be found in Jeff Yen’s thesis [111] and references therein.

The largest systematic on this technique is still the accuracy with which the crystal temperature is known. An interesting observation made early on with low T_c HV prototypes at SLAC was that, despite our thermometer sitting stably well below T_c , the bias power of the devices was slowly creeping up over a matter of days (as shown in figure 5.2). Subsequent follow-up measurements indicated that the cirlex clamps had anomalously low thermal conductance below 40 mK, resulting in the crystal remaining much hotter than the base temperature of the fridge. In fact, any parasitic power is completely degenerate with a higher T_c ; if we add a constant power term, we find that

$$K(T_c^5 - T_b^5) = I_0^2 R_0 + P_{parasitic} \rightarrow KT_c^5 - P_{parasitic} = I_0^2 R_0 + KT_b^5 \quad (5.24)$$

So if parasitic power is 0, we will measure the true T_c , but if it’s non-zero we’ll measure an artificially deflated T_c . This is the other major cause of time-dependent T_c changes. For these reasons, one of the largest anticipated challenges for SuperCDMS detectors will be ensuring that the crystals are cold enough that we can operate the QETs in the $T_b \sim 0$ limit where they will perform optimally, and that there is no parasitic power; to do this significant modifications will have to be made to the clamp design and careful design of all components in the bias loop.

Even if you’re sure your crystal is at the same temperature as your fridge, and you know you have no parasitic power, if you truly want to accurately characterize your sensors, a precise temperature

calibration is vital. Significant issues arose in the early fabrication studies for lower T_c devices due to misunderstandings arising from thermometer mis-calibration. In addition, accurate thermometry at very low temperature is in itself an art, and careful design of future systems with multiple redundant thermometers well-coupled (at low frequency) to the TES under study will improve our ability to characterize various fabrication runs and build accurate models of detector performance.

Given power measurements, the other quantity of interest which comes out of the fits (in the form of the slope and exponent fit to the power versus bath data) is the measurement of the thermal conductance $G = nKT_c^{n-1}$, for the parameterization

$$P = K(T^5 - T_b^5) \approx KT_c^5 + G(T - T_c) \quad (5.25)$$

If we assume our highest power measurements are consistent with the limit $P_{parasitic}, T_b \rightarrow 0$, we find that

$$G \approx nKT_c^{n-1} = \frac{nP_0T_c^{n-1}}{T_c^n} = \frac{nP_0}{T_c} \quad (5.26)$$

and if we know the TES volume, we can thus measure the electron-phonon coupling Σ as

$$\Sigma = \frac{K}{V_{TES,eff}} = \frac{P_0}{T_c^n V_{TES,eff}} \quad (5.27)$$

For the same TES design, then, if P_0 does not scale as T_c^5 , this implies a non-constant Σ or $V_{TES,eff}$. The latter could be a function of the degree of proximitization of the overlap regions, but should in principle be a small effect. The results of bias power measurements, and derived parameters, can be found in table [5.3](#).

5.1.3 TES Transition Parameters

If we look back at equation [5.8](#), we can actually derive further insight into the TES response from looking at either the frequency or time-dependence in the complex impedance. Suppose instead of $\omega = 0$, we take $\tau_{eff} < \omega^{-1} < \tau_{LR}$. In this limit, we get

$$Z_{loop,peak} = R_p + R_0(1 + \beta) \quad (5.28)$$

or in terms of the measurement, we find

$$R_0(1 + \beta) = R_{shunt} \left[\frac{\delta I_b}{\delta I_s} - 1 \right] - R_p \quad (5.29)$$

I call this $Z_{loop,peak}$ because if we put a square wave signal into our TES, we'll have the shape shown in figure [5.3](#) with an initial rise due to the β parameter. Let's consider the case of large loop gain ($\mathcal{L} \gg 1$). If we analyze the step response, we'll find that the complex impedance and IbIs both

tell us the same R_0 ; this tells us implicitly that we're in the large loop gain limit. The amount that the pulse rises before falling to its DC value will tell us β in combination with our previous R_0 measurement.

Given that we can only confirm whether the loop gain is large, how do we measure it when we find no difference between complex impedance and absolute bias measurements? The only solution in this case is to measure the fall-times, either in the time domain or the Fourier domain. The latter will allow us to first determine the inductance in the loop:

$$\tau_{LR} = \frac{L}{R_l + (1 + \beta)R_0} \quad (5.30)$$

This is only true *if we are not limited by the SQUID feedback loop bandwidth*, which is a concern for the SNOLAB electronics. Sensitivity to inductance should be better in the normal state than the superconducting state, and plotting the reconstructed τ_{LR} as a function of bias point should make it obvious when the measurement is bandwidth limited. Determining the inductance ahead of time gives one fewer free parameter for the remaining complex impedance fit.

The longer falltime in fourier space is given by the equation

$$\tau_{eff} = \frac{\tau}{1 - \mathcal{L}} \quad (5.31)$$

which implies that if we have a high-confidence estimate of $\tau = C/G$, and we're sure it doesn't change through the transition, we can also obtain a high-confidence estimate of \mathcal{L} . The problem here is getting a high confidence estimate of τ , and that crucial assumption that it doesn't change; remember from the TES dynamics chapter that we expect a factor of two change in heat capacity between normal and superconducting states based on BCS superconducting theory. One possibility is to operate the detector at T_c to measure τ without electro-thermal feedback. Another would be to use the bias power to measure G , and direct hits on the TES to measure C . These approaches also have significant systematics, but are in principle possible.

The best measurements of the transition parameters (α, β , and τ) come from simultaneous fitting of the complex impedance, shown in the time-domain in figure [5.3³](#). For the complex impedance shown above, we can re-parameterize this as a function of four unknown parameters (to remove the majority of the parameter covariance) as

$$\frac{dV}{dI} = Z_{loop} = A(1 + i\omega\tau_{el}) + B[1 + i\omega\tau_I]^{-1} \quad (5.32)$$

³Figure [5.3](#) also shows a nice diagnostic to use when determining the best TES bias from complex impedance in real time. Starting from normal data, we can determine when the TES has gone into electrothermal feedback by observing the phase shift in the complex impedance, and by maximizing the post-phase shift DC impedance, we can find the lowest bias point. This is generally done in small steps until the TES goes superconducting, and the stable bias point is then chosen to be slightly above the last observed TES bias point.

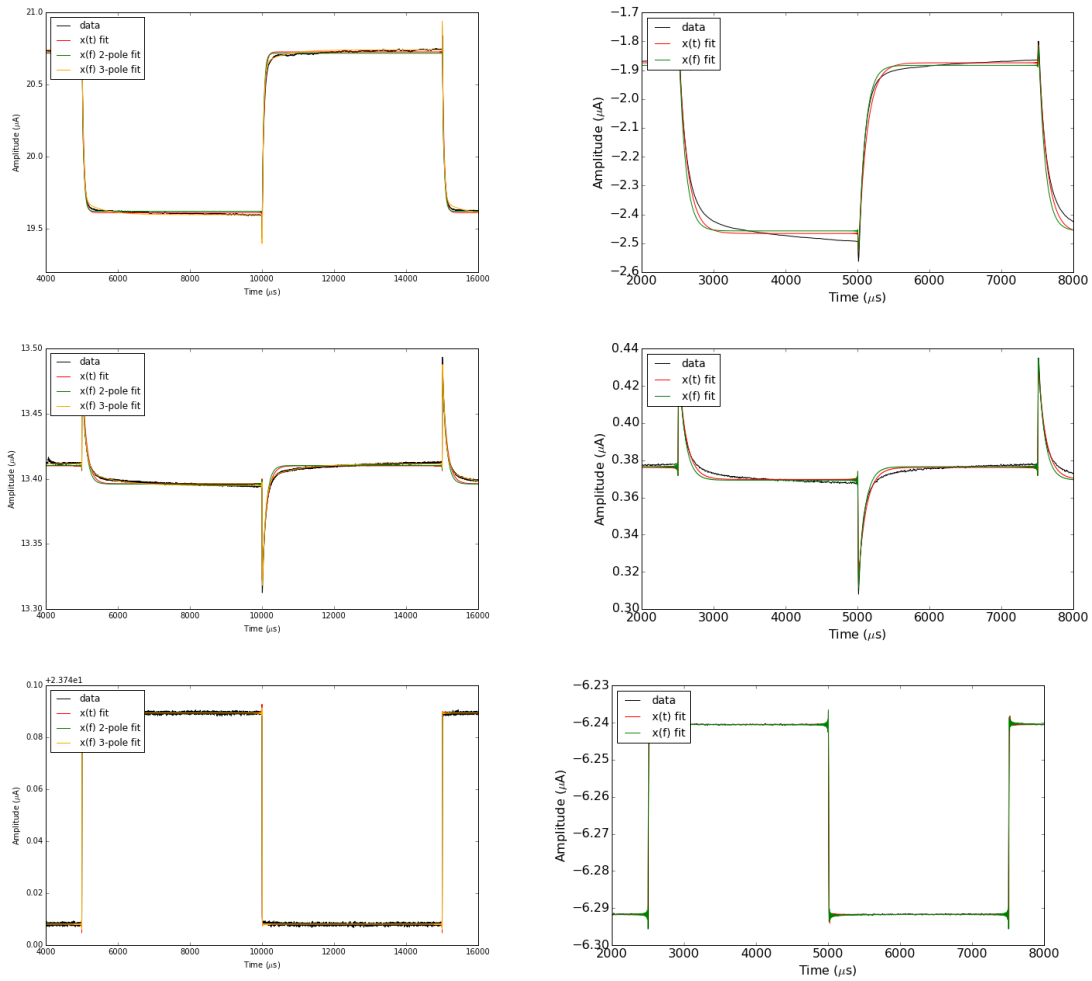


Figure 5.3: Complex impedance measurements of G115 TES channels at the TES operating point (top), entering into the electro-thermal feedback regime (middle) and normal (bottom). The left column is measurements done in the BlueFors at SLAC during Run 23, the right is for UCB run 482. The main difference, aside from the bias circuit and readout being different, has to do with the fridge base temperature. For the bias point far down in the transition, the lower base temperature at SLAC allows for a much faster fall time (larger degree of electrothermal feedback) than at UCB, where the base temperature was near T_c .

where

$$A = R_l + R_0(1 + \beta) \quad (5.33)$$

$$B = R_0(2 + \beta) \frac{\mathcal{L}}{1 - \mathcal{L}} \quad (5.34)$$

$$\tau_I = \frac{\tau}{1 - \mathcal{L}} \quad (5.35)$$

$$\tau_{el} = \frac{L}{A} = \frac{L}{R_l + R_0(1 + \beta)} \quad (5.36)$$

This also allows us to easily extract all of the unknown transition parameters sequentially, solving first for β and then for \mathcal{L} :

$$L = \tau_{el} A \quad (5.37)$$

$$\beta = \frac{A - R_l}{R_0} - 1 \quad (5.38)$$

$$\mathcal{L} = \frac{B}{R_0(2 + \beta) + B} \quad (5.39)$$

$$\tau = \tau_I(1 - \mathcal{L}) \quad (5.40)$$

This is actually a fairly clean way to fit the complex impedance, and by performing the fit in the Fourier domain we can use the real and imaginary information to get a properly weighted fit as well as to accurately fit both magnitude and phase of the complex impedance. We performed a least squares fit, computing the residual of the real and imaginary components as independent components as

$$F(Z_{meas}|A, B, \tau_I, \tau_{el}) = \sum_i \left[\frac{\text{Re}(Z_{meas}(\omega_i) - Z_{loop}(\omega_i|A, B, \tau_I, \tau_{el}))}{\text{Re}(\sigma_{Z_{meas}}(\omega_i))} \right]^2 \quad (5.41)$$

$$+ \left[\frac{\text{Im}(Z_{meas}(\omega_i) - Z_{loop}(\omega_i|A, B, \tau_I, \tau_{el}))}{\text{Im}(\sigma_{Z_{meas}}(\omega_i))} \right]^2 \quad (5.42)$$

and employing the least-squares fitting routine in scipy to minimize the residual with the appropriate loss and tolerance given the allowed range of parameters.

As can be seen in figure [5.3](#) we also fit 3-pole complex impedance curves to the data. The motivation for these fit was the observed departure of the complex impedance measurements from the simple two-pole model. As discussed in chapter [3](#), a two-block TES model of either an intermediate or hanging heat capacity can both be represented by the modified complex impedance response [70](#)

$$\frac{dV}{dI} = Z_{loop} = A(1 + i\omega\tau_{el}) + B \left[1 + i\omega\tau_I + \frac{C}{1 + i\omega\tau_3} \right]^{-1} \quad (5.43)$$

Det.	Fab Institution	T_c (mK)	R_n (m Ω)	$R_{n,exp}$ (m Ω)	ρ/ρ_{exp} .
G101c	TAMU	80	90–110	150	0.6–0.8
S101	TAMU	70–75	75–85	150	0.5–0.7
G115	TAMU	45	180	\sim 400	0.45
G124	SLAC	45–60	30–80	150	0.2–0.5
G147	SLAC	60–70	77–100	150	0.5–0.65
HVeV	SLAC	52	600–700	2.4–2.75	0.2–0.25

Table 5.2: Summary of TES critical temperature (T_c), normal resistance (R_n), and bias power (P_0) measurements. The R_n can be compared to the expectation to build up a picture of resistivity change as a function of T_c . Also included is the fabrication facility, which reflects real differences in film quality.

Det.	P_0 (pW)	G (nW/K)	K (μ W/K ⁵)	Σ (GW/K ⁵ m ³)
G101c	90–130	5.6–8.1	27–40	0.44–0.63
S101	50–70	3.6–4.7	30	0.47
G115	5	0.5–0.6	26–28	0.8–1.5
G124	10–30	1.1–2.5	40–54	0.61–0.86
G147	35–40	2.8–2.9	24–45	0.38–0.72
HVeV	0.05	4.8×10^{-3}	0.13	0.31

Table 5.3: Summary of TES bias power (P_0) measurements and derived quantities. The large uncertainty in G115 stems from an uncertainty in the effective tungsten volume.

where

$$f_g = C(1 - \mathcal{L}) \quad (5.44)$$

is the fraction of the total heat capacity in the hanging block, and τ_3 is the falltime of the internal degree of freedom between the TES and the block.

The results of these complex impedance fits for the UCB and SLAC data can be seen in figures 5.4 and 5.5. We find that for all cases where we do not have strong electro-thermal feedback, the single-block TES model is not sufficient to characterize the observed behavior, and multiple internal degrees of freedom are needed to accurately model the TES response. We also see large differences in performance between the two test facilities largely on the basis of bath temperature; the SLAC data has a larger bias power, loop gain, and a more meaningful trend in the fits for τ . The three-pole fits results also indicated that regardless of how well the two-pole model fit the data, the three pole model was always a better fit, and indicated that the additional block had a heat capacity fraction of 20–30% that of the total system. This is likely the fin connector based on the fact that measurements of the TES without fin connectors done at UCB indicate an excellent fit with a single two-pole complex impedance. Future work will be directed towards characterizing the impact of the fin connectors on TES response. This will also be discussed further in the next section.




Figure 5.4: Fits to the complex impedance in admittance space (left) and the complex impedance plane (right) for G115 at UCB (top) and SLAC (bottom). Note that the two-pole fit (dashed line) is a better fit to the data for the SLAC data, where the base temperature was lower, but the three-pole fit is needed to fit the UCB data. This is likely due to increased sensitivity to T_c variation given the higher base temperature at UCB.

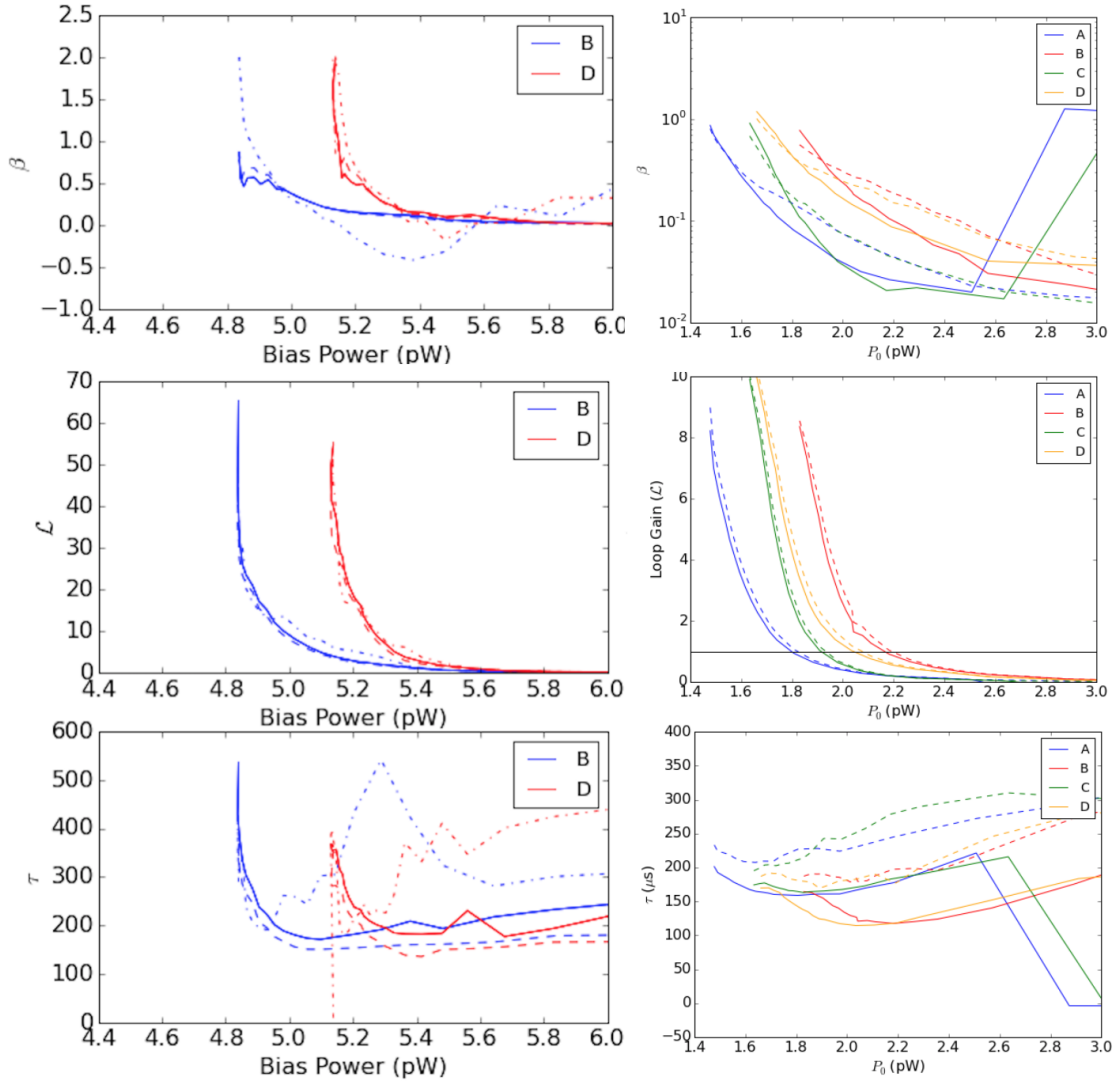


Figure 5.5: Complex impedance parameters for G115 run at SLAC (left) and UCB (right) as a function of bias power. The fully electro-thermal feedback regime should have constant bias power, as seen at SLAC, and high loop gain. The solid lines are the 3 parameter fits while the dashed lines are 2 parameter fits and the dot-dashed are time-domain 2 parameter fits. These measurements show that, as long as we can operate low in the transition, our expectation of high loop gain and $\beta \sim 1 - 2$. We also see the fits of τ increase by a factor of ~ 2.5 (close to the BCS expectation) as we bias lower in the transition.

Sample	Thickness (nm)	T_c (mK)	ρ (n Ω m)
Soudan iZIP	~ 40	80–120	130
SLAC Test Chip	40	40	48
TAMU Test Chip	30	60	72
Bulk Expectation (α Tungsten) [44, 61]	-	15	26.5

Table 5.4: Comparison of resistivity expected for the design of CDMS HV compared to the resistivities each fabrication facility is achieving based on test chips measured at SLAC. Note that the reduced thickness at TAMU is associated with both an elevated T_c and resistivity, implying a larger fraction of the film which is β tungsten.

5.1.4 TES Measurements Summary

Tables 5.2 and 5.3 summarize the TES characteristics measured from the range of prototypes tested during this detector performance measurement campaign. Taken together, these measurements suggest some non-trivial trends in both resistance and thermal conductance, but for the higher T_c detectors we find that Σ is close to the nominal value used in the previous chapters to determine the TES design. Looking at detector systematics, let’s play devil’s advocate to see how the Σ measurements might change if various assumptions are relaxed. If we are not in the infinite bath limit, then T_c and bias power are artificially low. Both will result in an increase in Σ . If we’ve underestimated the tungsten volume in a given channel, however, we will have over-estimated Σ . These are the two main systematics that could be tweaked to bring measurements in line. Resistivity would also be impacted by a comparison of true and expected TES thickness and linewidth.

So taking the measurements at face value, we find that the lower T_c detectors tend to have upward trends in Σ as critical temperature falls, but this is tenuous at best. Dedicated studies of G as a function of T_c and sensor geometry have been undertaken at SLAC and UCB by Sam Watkins, coupled with noise and complex impedance measurements which improve on the techniques shown here. These studies should provide a better handle on how the Σ parameter scales with film properties. What does seem to be the case is that both fabrication facilities are producing higher quality films than in the past, resulting in lower resistivity, which may partially account for the higher values of Σ as T_c is reduced. This is backed up by the test samples for which ρ was directly measured at SLAC, summarized in table 5.4.

In summary, in this section I reviewed measurements which put to rest one design concern (TES transition sharpness) while highlighting many we hadn’t necessarily considered (T_c variation and resistivity changes) as well as problems we did anticipate, that turned out to be just as challenging as expected (parasitic power and bath temperature uncertainty). The measurements in this section help justify further study with the small test devices being run at SLAC by the UCB group, which should be used for a near-term CDMS-HV redesign. It is likely that the existing CDMS HV design will still work given our good fortune with how the preproduction cables performed (in some respects), but the fact remains that these measurements show that the optimized design is not the true optimum

for the films we're producing. CDMS HV v2 will likely look substantially different than CDMS HV v1 both due to these measurements as well as the measurements described later in this chapter.

5.2 Characterizing QET Noise

The measurements described in the previous section, coupled with the noise modeling theory presented in section 3.3, allow us to study detector noise in these prototype detectors to determine whether we're dominated by intrinsic TES noise (on which we based our detector optimizations) or whether either bias circuit or environmental noise is substantially degrading performance. In this section I summarize the initial noise studies performed at UMN, UCB, and SLAC and the lessons learned from these initial tests. This will be presented somewhat historically, updating initial conclusions based on later findings but giving a sense for the thought process that went into our noise characterizations.

5.2.1 G101c: First Studies at UMN

The first HV detector test, done with G101c, was focused on demonstrating initial performance of the HV detectors was not limited by unforeseen complications. At the time, UMN was the only test facility able to run 100 mm detectors, but was equipped with legacy hardware that meant we were expected to be limited by both electrothermal oscillation and warm electronics noise. It nevertheless provided a testbed for early characterization of HV detector QETs and warned of fabrication and readout considerations taken into account for later detectors. The resistance measurements and bias power for this detector was already discussed in the previous section. What was new for this detector was the systematic study of detector noise where the detector was the dominant source of this noise.

Fits to the noise for one channel of G101c can be seen in figure 5.6. The fitting procedure used the previously determined parasitic resistance (and the estimated temperature of the shunt and parasitics) to fit the inductance to the PSD, and used the low and high frequency discrepancy between the expect and observed noise to fit the components of the SQUID noise. The SQUID noise was then refined based on the normal noise. Note that in both cases the passive voltage noise is just modified by the TES resistance, and both poles are well fit by this inductance.

We then move to the transition noise. We can determine an estimate of the bias resistance from IbIs and the DC component of the complex impedance, but for this detector we didn't do complex impedance fitting, so we weren't able to simply predict the noise. Based on our assumptions for the thermal conductance, we can estimate α and β in order to fit the electro-thermal oscillation peak, as shown in the figure. The first thing we notice is that, as expected, passive noise dominated the noise budget of this detector, but the oscillation peak is fairly well modeled by this set of parameters with transition values fairly in line with what was shown in the previous section. We noted at the




Figure 5.6: Fits of UMN noise in the superconducting state (top), normal state (middle), and transition state (bottom). Each plot adds a noise source, allowing for each noise source to be fit quasi-independently (only one or two noise parameters at each step). Note the peak characteristic of electrothermal oscillation; the central frequency of this peak is useful to determine the inductance in the TES bias loop.




Figure 5.7: Left: Regions of TES stability as a function of operating resistance and loop gain for fixed $\beta = 0.2$ and an inductance of 400 nH. Right: TES falltimes and electro-thermal oscillation frequency as a function of R_0 for the same β and inductance, and a fixed \mathcal{L} .

time, as I will now, that the low-frequency noise seems to rise, which was not expected, and this is a feature we’ve since observed in other test facilities. We’ll revisit it shortly but it’s worth noting that it was also present at UMN.

The confusing behavior we saw at UMN was that electro-thermal oscillation began to appear midway through the transition, but would disappear if the TES was biased low enough in the transition. In order to understand why this was the case, we need to remember that the condition for electro-thermal oscillation depends on all TES parameters including \mathcal{L} and β , as illustrated in figure 5.7 and that the disappearance of the oscillation could also be thought of as the oscillation moving outside of the frequency of interest or being sub-dominant to other noise features. If we take into account the fact that \mathcal{L} and β are a function of R_0 , we can draw different lines through the parameter space of Figure 5.7 which would see oscillation appear, grow stronger, and then recede as the bias is lowered. It doesn’t entirely go away, but the frequency shifts enough that it’s functionally absent, and the amplitude should correspondingly decrease.

Due to these oscillations and the large amount of passive voltage noise, there wasn’t much more we can learn from G101c (or S101) at UMN from a noise standpoint. In addition, both detectors had higher T_c targets, as shown in the table, due to the UMN fridge being limited to operating at 35–40 mK base temperatures, so we did not expect stellar performance from these detectors from a resolution standpoint. What was interesting was how the noise changed with bias; this will be discussed further in section 5.4.

5.2.2 G115 at SLAC and UCB

In parallel with the CDMS HV work at UMN, a 3 inch diameter 4 mm thick one-sided detector was developed by Matt Pyle and Suhas Ganjam at UCB using the SNOLAB QET design but with

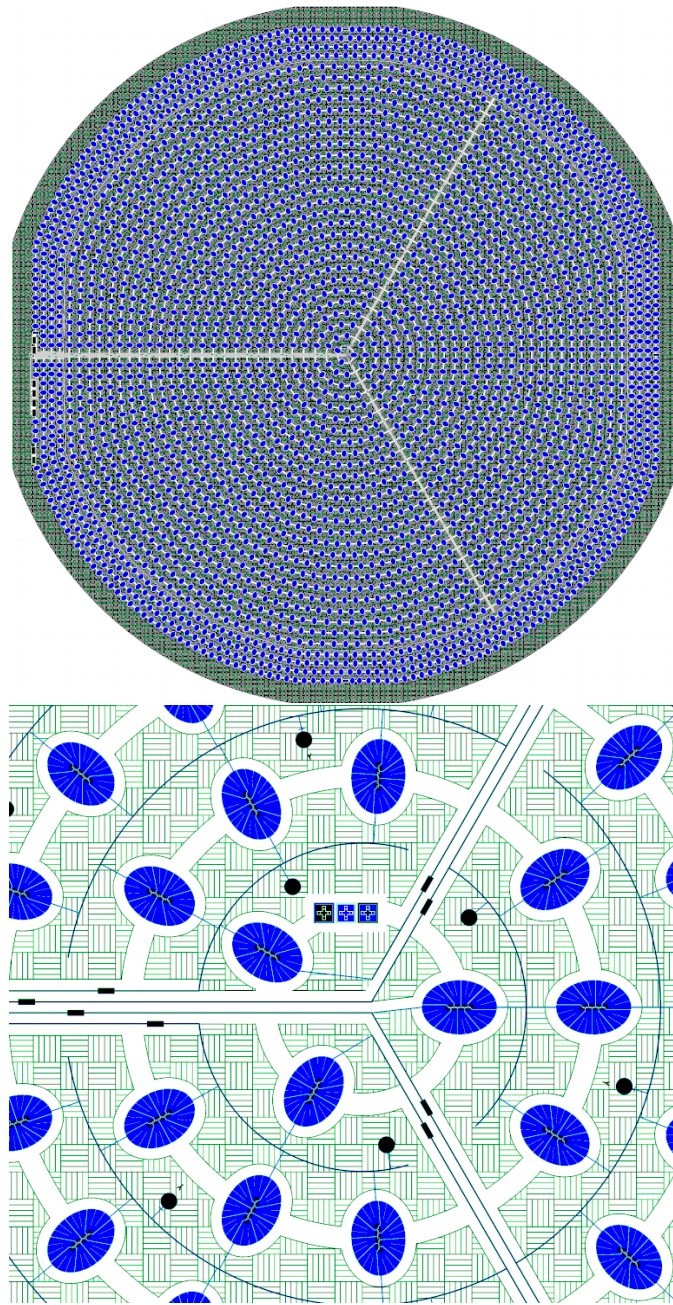


Figure 5.8: Top: Full image of G115 design with the same color scheme as the HV mask (aluminum is blue, tungsten is black, amorphous Si is green). There are four channels, with the inner 3 channels taking up 75% of the detector area and the outer channel plus the outer 2 mm guard ring taking up the last 25% of the area. Bottom: Zoom-in on the central region showing the parquet pattern in more detail and that the QETs and wiring are the same as for the CDMS HV mask. Phonon losses due to the the W-only grid are minimized by its small fill factor (6%), reduced thickness and no Al. The other side has just the parquet pattern done with all three layers, again with a 6% fill factor, but the grid includes aluminum and so is an appreciable phonon sink.

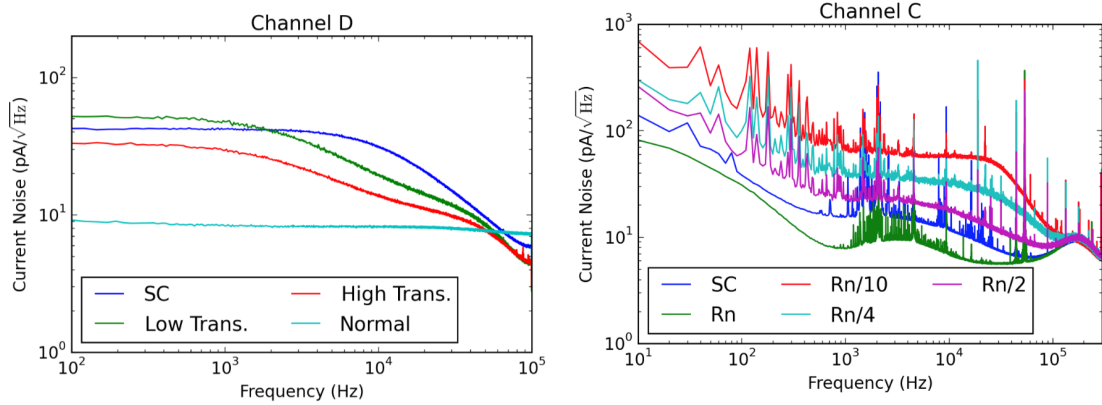


Figure 5.9: Comparing G115 noise for the UCB data taken with the FEBS (left) to the SLAC data taken with the DCRC D.0 before any on-board modifications were made. The superconducting noise at UCB was much higher due to the larger shunts at high temperature and the larger parasitics. The SLAC noise is about what was expected qualitatively - the TES dominates the noise - but there are many features that are not what they should be, including forests of lines, the low-frequency rise, and a hump in the kHz region. For the lowest bias, point, which is ~ 20 m Ω , we see noise that is clearly TES dominated and finally flat.

channels which were ~ 2.3 times smaller ($R_n \sim 0.35\Omega$, $N_{TES} \sim 700$) and a parquet pattern in between the QETs (see figure 5.8). It was a four-channel device with three central channels and an outer ring similar to the phonon channel layout for the Soudan iZIP detectors. A smaller design was chosen in order to focus on QET performance, and it allowed us to test theories we had developed early on about the role of muons and high energy gamma backgrounds on high voltage noise and efficiency in a detector with a different form factor. This is discussed in more detail in Section 5.4.

The measured noise at UCB and the first run at SLAC is shown in figure 5.9. The good news was that the superconducting and normal noise were both lower than that seen at UCB in the regions consistent with voltage noise. The bad news was that obviously there is a forest of lines, a lot of low frequency noise not seen at UCB, and a high frequency hump. In addition, we can't fit the UCB noise in a way that's consistent with the G measurement derived from the bias power. That's alright in light of the fact that we have low loop gain, and we're near the bath temperature, at UCB, whereas at SLAC the sources of additional noise need to be understood and reduced.

5.2.3 Diagnosing Readout Noise

Here I take a quick diversion from discussion detector properties to diagnosing noise seen in the PSDs at SLAC in order to explain how the DCRC readout was improved as well as how new noise sources can be traced to different parts of the circuit. This exercise helped isolate potential areas noise sources could originate and helped us determine where our dominant noise sources are at SLAC.




Figure 5.10: EMI measured in the SLAC Building 33 clean room with a 16 turn 100 ft length coil attached to a signal analyzer with the light on (off) shown on the left (right). We see both a low-frequency noise drop as well as a large 40 kHz line largely disappear.

The first noise feature we noticed at SLAC even with the DCRC powered off was a very strong time-correlated EMI noise all over the board, and very strong at the SQUID pre-amplifier input. We decided to grab some spare wire from the electronics shop, made a 16 loop coil out of 100 feet of wire, and plugged it into a signal analyzer to measure the EMI environment in the lab. Moving the cart around showed that though the features were consistent, they were stronger near the lights. Figure 5.10 shows the EMI measurement with the lights on and off, showing the huge EMI reduction from having the lights off. The features can be attributed to the ballasts of the fluorescent bulbs, which run nominally in the kHz region to obtain higher energy efficiency. The low frequency was attributed to 60Hz and harmonics from the ballast wiring.

You can see that even though all of the lights in the lab are off, the 40 kHz feature is still seen due to lights elsewhere in the building, leading us to believe that the ground wiring was acting as an antenna for this EMI. We also could never get the remaining features to disappear regardless of what we turned on or off in the lab. We tried both the floating and grounded outer connector, with the other connector floating, and found that we did get better noise with grounding, which is what is shown in the figure. Remaining sources of noise were either the few pieces of equipment we couldn't turn off or sources from outside the clean room. The features we found did correlate fairly well with PSD features, so this suggests repeating these measurements as we install new equipment will help us diagnose EMI at SNOLAB.

After this point we always ran with as much equipment off as possible, and moved on to characterizing intrinsic readout noise. Here we'll pause for some preliminaries about SQUIDS to help differentiate some terminology (see also refs [46, 50] and refs therein). Consider a SQUID as a black box, where there are two inputs and one output. One input is the TES input coil, one is the feedback input coil. The output is simultaneously used to determine the SQUID signal and introduce the SQUID bias current, so it's almost an unintentional third input. Let's ignore that for a second.

Considering the inputs, each has a mutual inductance with the SQUID such that the flux can be written

$$\Phi_i = M_i I_{in} \quad (5.45)$$

and the SQUID is a device that turns magnetic flux into voltage. Given that there is some function for the SQUID (roughly looking like a sinusoid) that relates output voltage to flux, we can thus say that the SQUID voltage is

$$V_{sq} = V(M_i I_i) \quad (5.46)$$

and we can thus calculate the SQUID responsivity

$$Z_{sq} = \frac{dV_{sq}}{dI_i} \approx \frac{dV_{sq}}{d\Phi} \frac{d\Phi}{dI_i} = M_i \frac{dV_{sq}}{d\Phi} \quad (5.47)$$

which has units of Ohms. This also tells us that the responsivity for the feedback coil is just a constant factor of the mutual inductance times the responsivity of the input coil:

$$Z_{sq,fb} = \frac{M_{fb}}{M_{in}} Z_{sq,in} \quad (5.48)$$

which means as long as we measure the $V(\Phi)$ curve for the SQUID with one coil, and we know the mutual inductance ratios, we can model the whole system fairly well. The first step of tuning SQUIDS is thus to measure the $V(\Phi)$ curve, and find the point in the curve with the largest slope with the correct sign (so that feedback is stable) and adjust the SQUID feedback offset to sit at that voltage point. The sign of the slope is also important, but the closed-loop feedback will choose the slope based on the polarity of the feedback circuit (it can be positive or negative).

The slope we choose does make a significant difference for the noise performance; here we bring the the third input/output line into our analysis. If I apply no current to the SQUID, there's no voltage for the flux to modulate, so obviously the bias I apply is important. In addition, I can measure the SQUID $V(I_{bias})$ curve. During SQUID tuning, we normally observe the $V(\Phi)$ curve as I_{bias} changes, and the best SQUID point is an optimization of both these parameters. This does imply, however, that the SQUID acts both as a transimpedance amplifier and as a resistor, so we're sensitive to the SQUID dynamic impedance

$$R_{sq} = \frac{dV_{sq}}{dI_{bias}} \quad (5.49)$$

at the chosen operating point. This is important because if there is some current noise on the SQUID bias line, then we have the SQUID voltage noise

$$N_{I,sq} = \frac{N_{V,SQ}}{Z_{SQ}} = \frac{N_{V,ISQ0} + N_{I,bias} R_{sq}}{Z_{SQ}} \quad (5.50)$$

Here I've assumed that SQUID noise is intrinsically a voltage noise which could either be driven by flux noise or Johnson noise across the Josephson junctions. For our purposes it doesn't really matter. We can see that our optimal noise performance will be for the largest SQUID responsivity and the smaller SQUID dynamic impedance. This is demonstrated for G115 data taken at SLAC along with SQUID characterization done in a following run in figure [5.11](#).

Assuming we have our SQUIDS optimally configured, we can start looking at the readout circuit, which is a fairly general closed-loop SQUID readout; figure [5.12](#) shows the DCRC D.0 phonon readout circuit. There are two programmable gain amplifiers in this circuit which will help us try to diagnose the dominant sources of noise in this circuit, and a switch to allow us to look at open loop (no feedback) noise performance to further isolate different components.

The first thing to consider is whether the closed loop is truly closed by looking at the transfer function to see what the frequency dependence looks like. We know that the transfer function at low enough frequency (let's say 10 kHz) is white, with the form [46](#)

$$H = R_f \frac{M_{TES}}{M_{FB}} \left[1 + \frac{R_f}{G} \frac{M_{TES}}{M_{FB}} \right]^{-1} \quad (5.51)$$

where the open-loop current to voltage gain G is

$$G = Z_{sq} A_{LNA} A_{PFA,FB} \quad (5.52)$$

where Z_{sq} is the SQUID responsivity, and the A terms are the gain for the Low Noise pre-Amplifier (LNA) and closed-loop feedback Programmable Gain Amplifier (PGA).

For the minimum gain of 100 from the LNA, with the PGA at 1 and a SQUID responsivity of 100 Ohms, we find that for us to be approximately at infinite gain, we need

$$R_{fb} \ll \frac{M_{FB}}{M_{TES}} Z_{sq} A_{LNA} A_{FB} \sim 8 \text{ k}\Omega \quad (5.53)$$

Which tells us that for low PGA settings, the 5k resistor is allowing us to operate in an infinite gain limit. When we set the FB gain to 2, we get closer though by no means are we in a strong feedback limit. This suggested trying different feedback resistors below this limit to see how it impacted the noise performance. The result was that we achieved the closed-loop limit with less gain but reduced the bandwidth of the feedback circuit due to the maximum slew-rate of the PGA. This was a longer study that I'll leave out the details of, but suffice it to say that varying the feedback resistor was not helpful.

Assuming then that we're in the infinite gain limit, we can write the TES referenced current noise as follows:

$$N_{ITES,out}^2 = N_{ITES,in}^2 + N_{ISQ}^2 + N_{IDCRC}^2 = N_{ITES,in}^2 + N_{ISQ}^2 + \frac{N_{VDCRC}^2}{Z_{SQ, TES}^2} \quad (5.54)$$

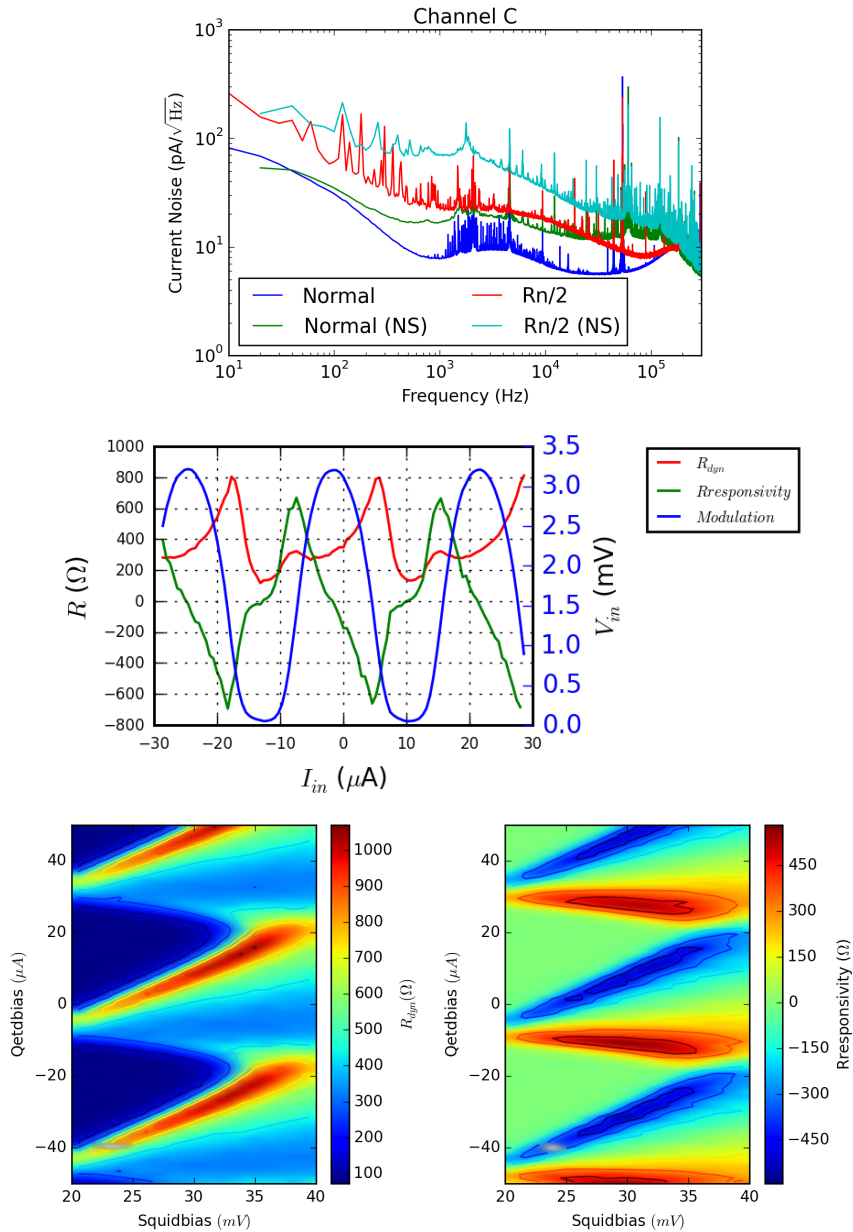


Figure 5.11: Top: Noise PSDs for different SQUID slopes when we accidentally flipped the polarity of the SQUID feedback circuit. The increased noise is due to current noise on the SQUID bias line which is amplified by the increased dynamic resistance. The ‘NS’ in the caption refers to no RF source present; this was supposed to be a source on/off test. Middle: Modulation curve for a SNOLAB SQUID measured for the SLAC setup including the responsivity and dynamic resistance measurements as a function of feedback coil current (input flux). We can see that responsivity is symmetric but dynamic resistance is not, so we want to bias on the positive responsivity slope where dynamic resistance is low. Bottom: Two-dimensional bias versus flux measurements for a SNOLAB SQUID showing that the low dynamic resistance region corresponds to a responsivity region that lines along a constant input flux. This produces better noise response and less susceptibility to input bias variation, for a stabler SQUID feedback circuit.

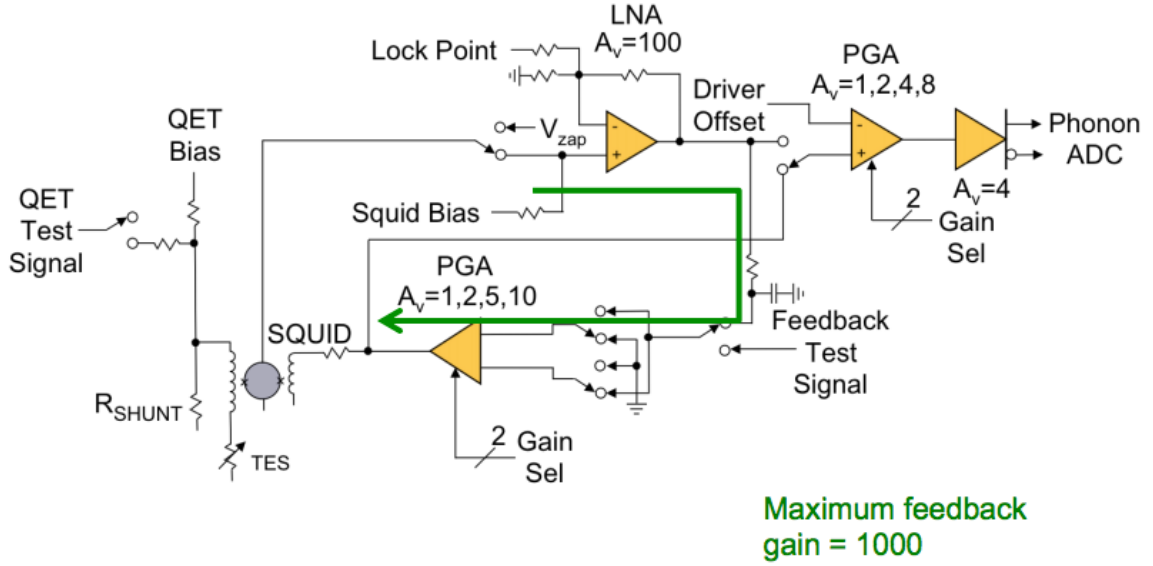


Figure 5.12: DCRC Rev D.0 SQUID and TES bias circuit schematic, used to construct the noise model discussed in this section.

where $N_{ITES,in}$ is the intrinsic TES current noise, N_{ISQ} is the SQUID noise, and N_{VDCRC} is all of the DCRC voltage noise referenced to the squid output. We can further decompose DCRC voltage noise by its location in the feedback circuit:

$$N_{V,DCRC}^2 = N_{LNA}^2 + \frac{N_{FBPGA}^2}{A_{LNA}^2} + \frac{N_{RFB}^2}{A_{LNA}^2 A_{FBPGA}^2} \quad (5.55)$$

We can put this all together by referencing ADC input noise (including the PGA and ADC) to TES current noise, for an effective current noise

$$N_{ITES,tot}^2 = N_{ITES,out}^2 + \frac{N_{VAMP}^2}{R_f \frac{M_{TES}^2}{M_{FB}^2}} \quad (5.56)$$

where we separate the noise into pre and post PGA noise:

$$N_{VAMP}^2 = N_{VAMPPGA}^2 + \frac{N_{VAMPADC}^2}{A_{AMPPGA}^2} \quad (5.57)$$

Chaining this all together gives the following expression for the measured current noise at the DAQ:

$$N_{ITES,DAQ}^2 = N_{ITES,in}^2 + N_{ISQ}^2 + \frac{N_{LNA}^2 + \frac{N_{FBPGA}^2}{A_{LNA}^2} + \frac{N_{RFB}^2}{A_{LNA}^2 A_{FBPGA}^2}}{Z_{SQ,TES}^2} \quad (5.58)$$

$$+ \frac{N_{VAMP PGA}^2 + \frac{N_{VAMP ADC}^2}{A_{AMP PGA}^2}}{R_f^2 \frac{M_{TES}^2}{M_{FB}^2}} \quad (5.59)$$

$$= N_{ITES,in}^2 + N_{ISQ}^2 + \frac{1}{Z_{SQ,TES}^2} \left[N_{LNA}^2 + \frac{N_{FBPGA}^2}{A_{LNA}^2} \right] \quad (5.60)$$

$$+ \frac{N_{RFB}^2}{Z_{SQ,TES}^2 A_{LNA}^2 A_{FBPGA}^2} + \frac{N_{VAMP PGA}^2}{R_f^2 \frac{M_{TES}^2}{M_{FB}^2}} + \frac{N_{VAMP ADC}^2}{A_{AMP PGA}^2 R_f^2 \frac{M_{TES}^2}{M_{FB}^2}} \quad (5.61)$$

$$= N_{ITES,in}^2 + N_{ISQ}^2 + \frac{N_{RFB}^2}{Z_{SQ,TES}^2 A_{LNA}^2 A_{FBPGA}^2} \quad (5.62)$$

$$+ \frac{N_{VAMP PGA}^2}{R_f^2 \frac{M_{TES}^2}{M_{FB}^2}} + \frac{N_{VAMP ADC}^2}{A_{AMP PGA}^2 R_f^2 \frac{M_{TES}^2}{M_{FB}^2}} \quad (5.63)$$

where I have re-arranged noise terms to suggest a way to diagnose which is the leading source of noise. We can vary both PGA gains, the feedback and SQUID resistances, the current noise on the TES, and toggle open and closed feedback and thus manipulate each of these noise sources independently.

We subsequently performed the following tests, with the results summarized below:

1. Vary driver gain, fixing all other terms, for highest usable feedback PGA gain. No noise change is observed, so $N_{VAMP,ADC}$ is sub-dominant
2. Reduce feedback resistor, fixing all other terms, for highest usable feedback PGA gain. Changing this also brings us closer to the infinite gain case. Only change seen is worse noise performance at high frequency, so changing feedback resistor is not helpful as we lose high-frequency noise performance.
3. Noise does not appreciably change as we adjust the feedback gain beyond 1, so we can assume the noise there is also sub-dominant, once we're appreciably in the infinite gain limit. Changing the feedback resistors did increase closed-loop gain but also reduced the feedback loop bandwidth, which reduced the overall performance of the feedback circuit.
4. Can see noise change at high frequency above some pole when bias condition is changed, suggesting that above a cutoff frequency we really are dominated by TES noise, but not as much as we should be.
5. The noise floor below 100 kHz seems to be unchanged by most of these tests; leads to open-loop tests which suggest that SQUID noise (or noise at the input of the LNA) is dominant.

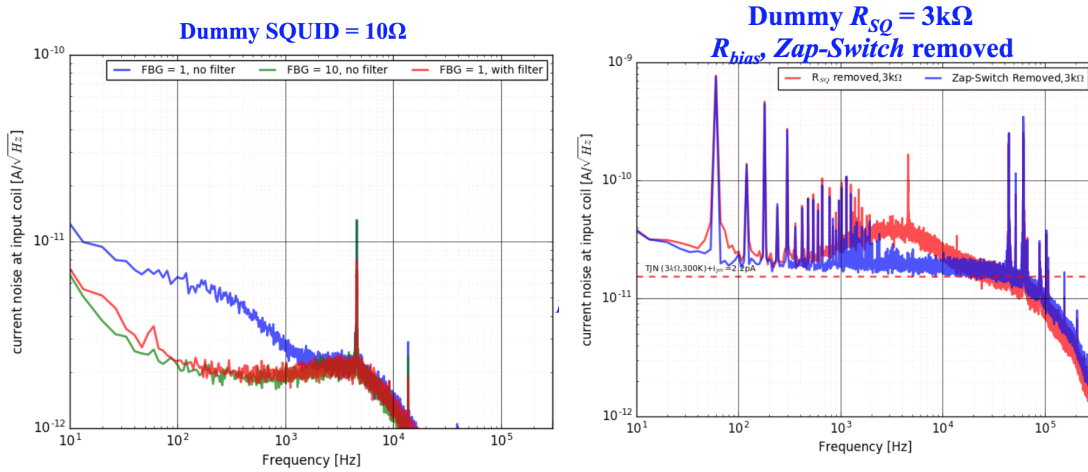


Figure 5.13: Left: Impact of filtering SQUID lock-point amplifier power supply. Right: impact of removing the zap switch, which coupled power line noise into the signal.

The observed SQUID noise is higher than expected, and the observed noise is dependent on the responsivity of the SQUID. With the SQUID disabled, we still see troubling low-frequency and mid-frequency noise trends.

We thus determined that most of the large noise features were coming in on the DCRC before the LNA input, and Tsuguo was able to diagnose these noise sources by attaching resistors in place of the SQUID at room temperature, with the results shown in figure 5.13. We had current noise leaking onto the SQUID through the SQUID zap switch power supply, which was fixed by adding power line filtering. We also had voltage noise coming in from the amplifier for the SQUID lock-point offset, which was also solved with additional filtering. The resulting noise from G115 after these modifications were made can be seen in figure 5.14.

We can draw a large number of conclusions from the plots using the modified DCRCs. We see that the DCRCs exceed the noise measured with a custom analog amplifier, suggesting that there may still be room for improvement in the DCRC, but we note that both D1 and D0 are capable of meeting spec with these modifications. We can also see that D1 has fewer low-frequency noise lines than D0 due to some shielding improvements, as well as a lower bandwidth due to the addition of filtering in the feedback loop. This is largely above the bandwidth of our signal, though we'd like to see as much of the TES rise as possible, so we should be careful to get the bandwidth up to around 500 kHz. More work is needed to solve some of the oscillation problems we've seen when trying to improve the high bandwidth performance of the DCRC.

After these modifications, the 'dIdV' coupled noise (voltage noise on the TES bias line) became the focus of our investigations, as can be seen in the top right and bottom left panels of figure 5.14. We see that there is a voltage coupled noise which moves with the TES bias point in the normal and

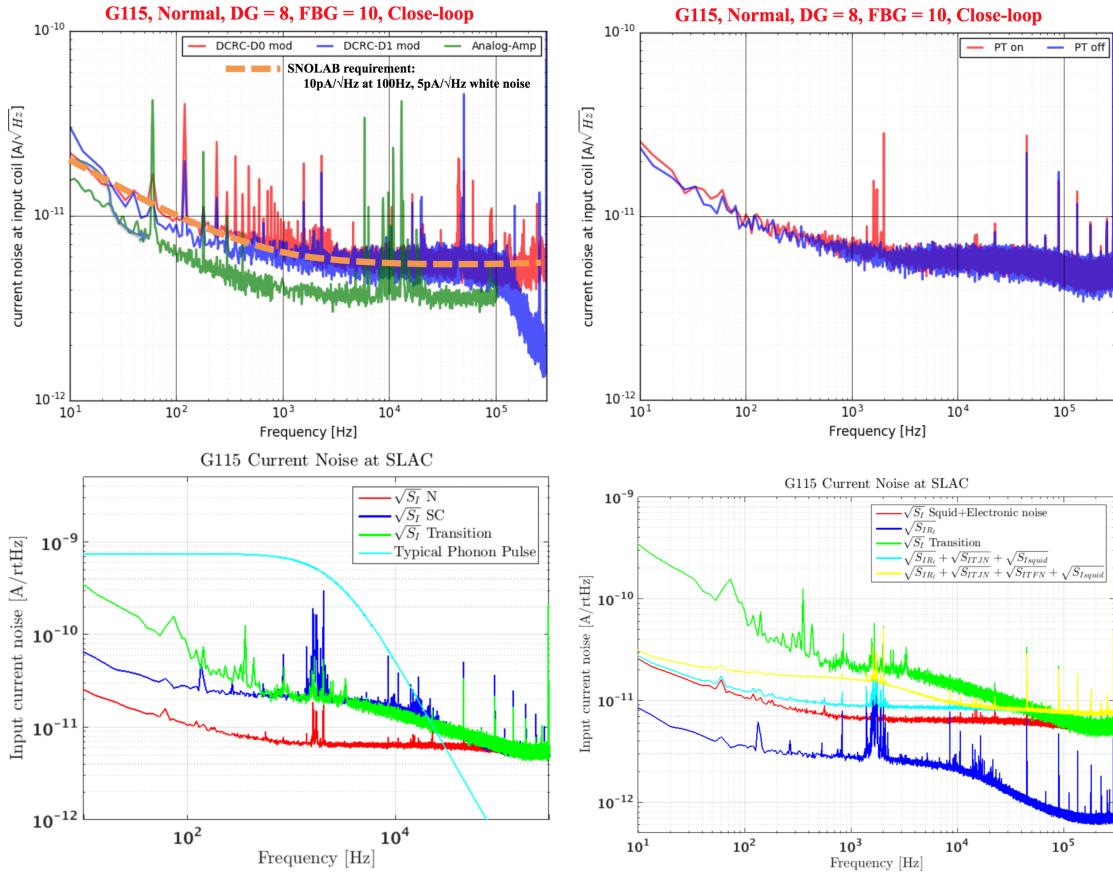


Figure 5.14: G115 Noise for various readout and TES bias conditions at SLAC from Run 26 in the BlueFors. Top: Normal noise for DCRC compared to the analog SQUID amplifier (left) and for the pulse tube on/off (right). We see that both D0 and D1 have an acceptable noise floor after the modifications, but that they're still DCRC limited given that we can achieve better noise with other electronics. They meet the spec we set for SNOLAB. See text for more details. Bottom: noise in the superconducting, normal, and transition states for G115 (left) as well as Caleb's noise model fit (right) showing that we do achieve TES-limited noise above 1 kHz. This also demonstrates that there is remaining noise both down the TES bias line as well as excess power noise on the TES.

superconducting states largely driven by the pulse tube, suggesting either EMI from the controller or some vibrational coupling to the wiring. We also see though that in transition, the dominant noise we have is a low-frequency noise which is demonstrably not a result of either small pulses leaking through our cuts (it does not have the right high frequency pole) and is not residual muon tails. We'll explore this more in the next section with G124, which was much more sensitive to both noise sources.

These studies suggest some design principles for CDMS electronics going forward, and possible improvements for the existing electronics for SNOLAB:

- Every input or output to the SQUID and TES should be as filtered as possible to reduce power line noise coupling into the SQUID loop. In the future, moving the pre-amp inside the fridge would mitigate a lot of the issues we're seeing.
- Any switches which don't need to be high speed should be made electromechanical to eliminate the possibility of capacitive coupling.
- We should always opt for LED lights run (if possible) from a DC supply. Any high-frequency supplies will be large EMI sources.
- A differential D.0 board was tested but ultimately a single-ended design was selected. If a lower noise differential amplifier was found that would vastly reduce our susceptibility to on-board EMI and ground loops.

These are just based on the DCRC and SQUID noise tests, and do not include lessons learned from TES transition noise, which at this point in the chapter we haven't dug into in detail. We'll draw a few more conclusions about readout noise from the next chapter.

5.2.4 G124 at SLAC: Beating Down the Noise

After these noise studies, we moved on to further characterizations of G124 at SLAC with the goal of validating the noise model for the SNOLAB HV detectors given that G115 seemed to have anomalously high, correlated power noise (as shown in figure 5.14, the measured noise is higher than the model across the entire frequency range). For these data from Run 33-35, the DCRC has been modified, so the normal and superconducting noises are about what we expect though there are lines due to EMI pickup and a slow rise at low frequency from residual (but mostly uncorrelated) DCRC noise. These PSDs are shown in figure 5.15.

Measurements of transition noise in G124, however, show another host of problems, some still in the readout but most related to external noise sources coupling to the detector inside the fridge. These noise sources are summarized in figure 5.16. The largest problem which is immediately noticeable by eye in the traces is due to telegraph noise, which is RF EMI from the DCRC during readout of trace data which is aliased to lower frequency due both to temporal variation and the




Figure 5.15: PSDs for G124 side 2 ($R_n \sim 50 \text{ m}\Omega$) for the normal (left) and superconducting (right) states. We can see that most of the lines scale with resistance, meaning they're due to EMI pickup on the bias line and are predominantly on the TES side, although the lines seem to scale faster than the white noise. We can note now that there's still a small low-frequency rise below 1 kHz, but it's much less distinct than before the DCRC modifications.

burst-like nature of the additional RF noise it generates. This problem is in itself a significant issue for our readout scheme and required us to artificially slow the DCRC during acquisition so that we don't read out traces while the DCRC is transmitting data. This was the subject of a ton of work by the UCB group on the UCB fridge and calls for fitting all of the fridge wiring access points with RF Pi filters, which will be done for SNOLAB. The UCB group showed that when properly RF shielded, this noise goes away, but despite multiple attempts the Vacuum Interface Board (VIB) which brings signals into the fridge still allows this noise in at SLAC. This was the subject of RF surveys done by Caleb in runs 35-37, the results of which are still in process.

With the trigger rate suppressed, we can probe the telegraph-free noise as shown in the remaining panels of figure 5.16. The remaining major noise sources are the turbo maintenance signal (which is EMI coupled dIdV noise that also contributes bias power) and pulse tube noise which seems primarily vibrational. The turbo noise consists of a 1.6 kHz narrow peak that turns on and off periodically to control the turbo. It's seen in the normal and superconducting state as well, and scales with the TES resistance, so we can conclude that the noise is primarily a current noise down the TES bias line. It can be seen in figure 5.16 for a time period where it was on for all traces rather than being average out or removed by quality cuts during a small period of time when the pulse tube was off, so we can see that whenever this noise source is present we have a large degradation in resolution.

The pulse tube on turbo off data shown in figure 5.16 is very telling, as the turbo line disappears but not the low-frequency noise. Turning the pulse tube off as well allow the noise to become largely flat down to 100 Hz, demonstrating that the cause of this noise is very likely the pulse tube. The fact that this feature only shows up in the transition state implies it's a power noise, which means it's

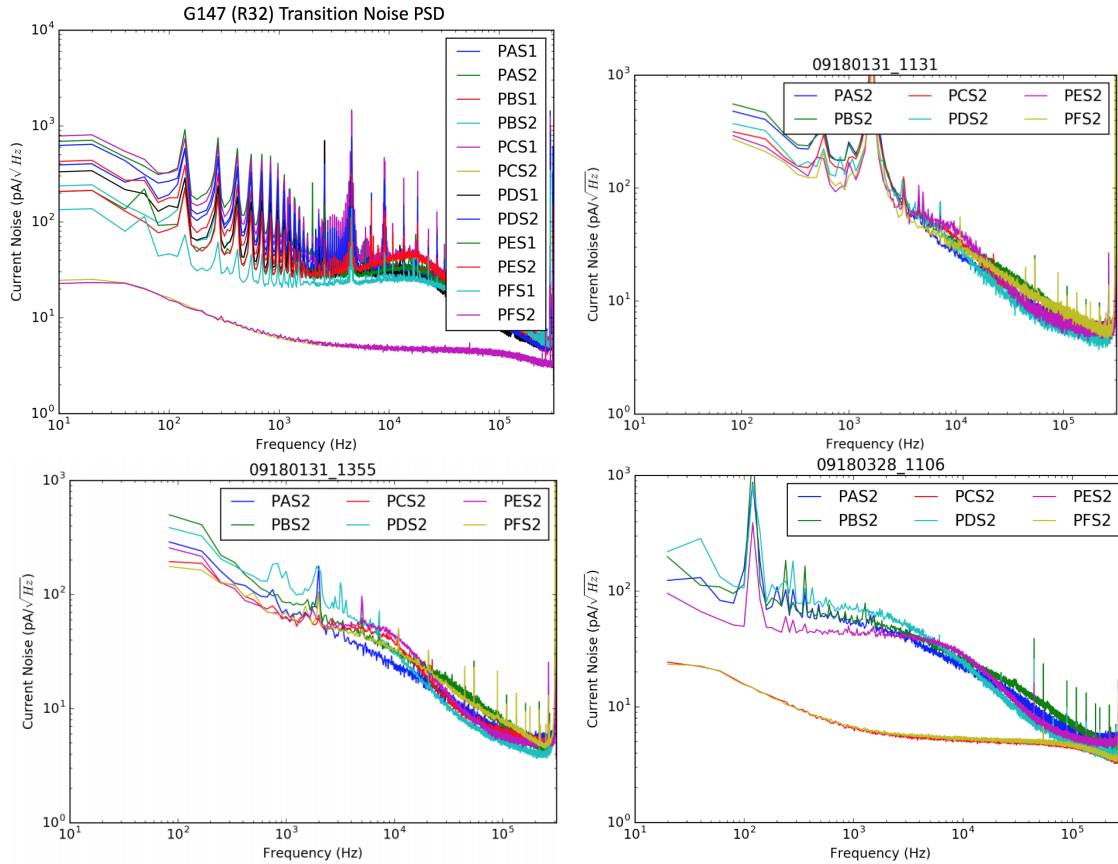


Figure 5.16: Transition noise for different operating conditions. Top left: Telegraph noise on G147, G124 not shown because the telegraph noise was bad enough to drive all of the PSDs off-scale. Top Right: transition PSD with pulse tube off but turbo on during a period of high turbo EMI. Bottom Left: PSD for pulse tube on, turbo off, showing a reduction in kHz-frequency noise but all of the same low-frequency noise. Bottom Right: PSDs for turbo and pulse tube off; PSDs are largely flat down to 100 Hz. Two channels are disconnected, so here you can also see the effective DCRC input noise which has a low-frequency increase due to pre-amp input noise. See text for more details.

most likely vibrationally coupled to the detector. This data motivated a larger study of vibrations from the pulse tube into the cryostat by To Chin Yu, using various vibration damping techniques.

To understand how vibrational noise can coupling nonlinearly into the detector, consider the vibrational velocity ansatz

$$v(t) \approx \sum_{i=-\infty}^{\infty} \cos(\omega(t - Ti)) \exp(-(t - Ti)/\tau) \Theta(t - Ti) \quad (5.64)$$

This is a driven damped oscillation with characteristic frequency ω , damping decay time constant τ , and periodicity T . Two example cases of this model are shown in figure 5.17 for different decay constants. The basic idea is that vibrational power couples into the detector either through voltage noise, which we'll see in all bias states, or through friction.

For the frictional case, we can use a simple kinetic friction model

$$F = \mu F_n \geq \mu m_{det} g \quad (5.65)$$

where the inequality depends on how tightly held the detector is; increasing the clamp tension will also increase the frictional force and thus vibration coupling. The energy dissipation is thus

$$P = \frac{dE}{dt} = \frac{d}{dt} |F * x| = F |v| \quad (5.66)$$

Here we get an absolute value because power will never be negative. This is a very non-linear transfer function, which means we'll alias and frequency shift much of the vibrational power. This is what is illustrated in the bottom panels of figure 5.17. We have begun a program of vibration studies to try to measure the vibration velocity spectrum for comparison to the detector noise and determine whether we can use this model to help mitigate or reduce the noise. We do notice that the amount of vibrational noise seen does depend both on detector mass and mounting scheme so we expect there is a lot we can do to reduce vibrational coupling to the detectors.

A final note on noise studies of G124 comes from looking at the coherence, a statistic described in the signal processing appendix which essentially measures the causal linkage between two signals in the frequency domain. We find that for the large samples of G124 noise, the vast majority of the noise is highly correlated below 1 kHz, around 60-80% or more for almost all channels. This precludes most mechanisms which are unrelated to the detector crystal given that opposite sides are separated on the readout cabling and run from two different DCRCs, and supports the conclusion that this low-frequency noise is vibrational.

As described in the next section, we can use this coherence to perform a degree of noise correction when the correlated and uncorrelated noises are comparable in size, as shown in figure 5.18. When we do this correction we can see that the corrected noise resembles the noise we see with the turbo and pulse tube off, suggesting that these correlations are in fact the noises we've identified. This

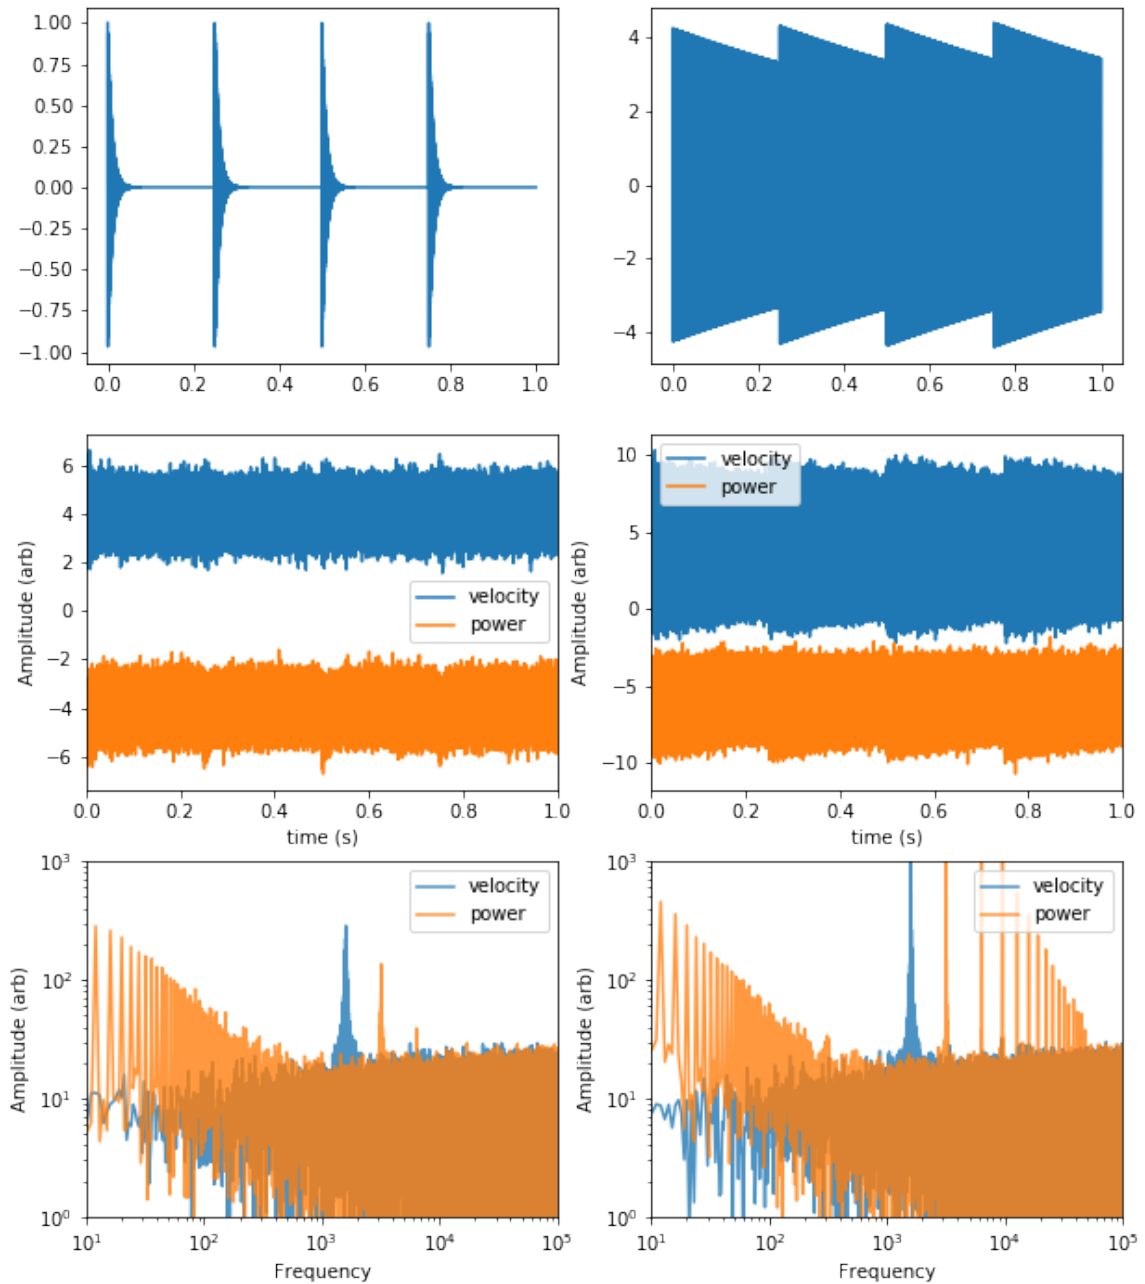


Figure 5.17: A simple vibrational noise model based on the ansatz equation [5.64](#). Both columns use an oscillation frequency of 1.6 kHz and period of 0.25 seconds, the left column uses a 10 ms decay right and the right column a 1 second decay time. The top row shows the pure signal, the middle shows the signal with noise and the rectified signal with noise, and the bottom shows PSDs for these two signals. We see that both have similar low-frequency noise, but the longer decay time allows for much stronger high-frequency noise.

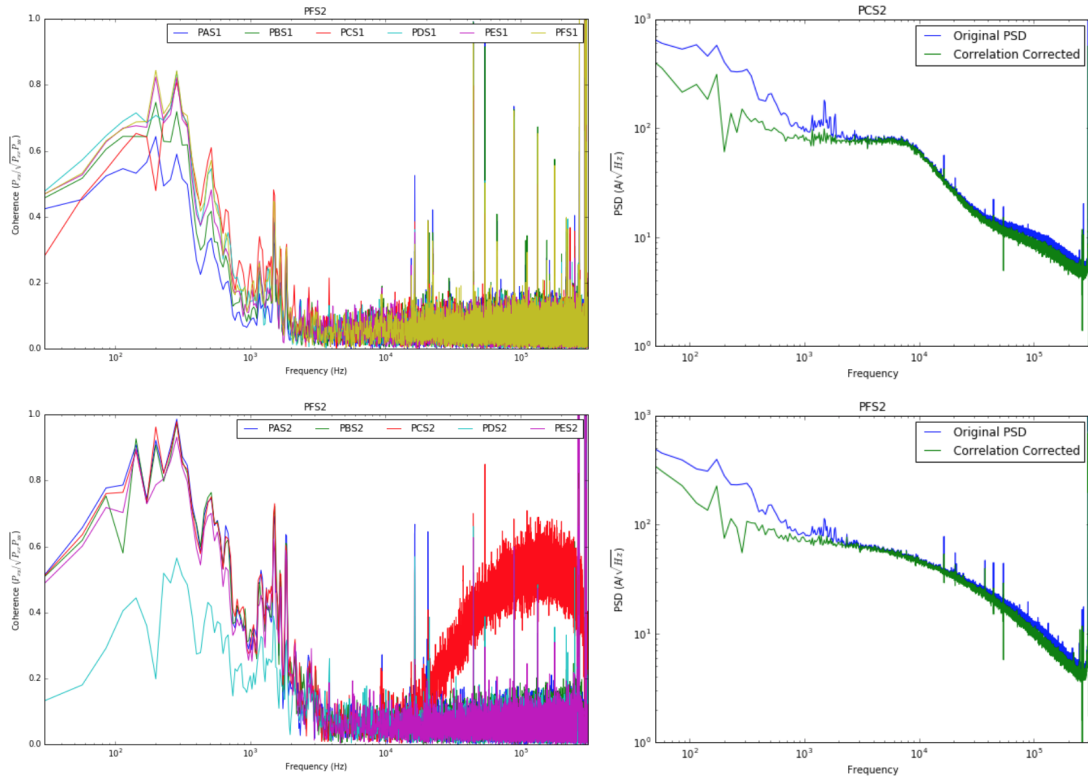


Figure 5.18: Coherence of channel PFS2 with the other 11 channels on each side of the detector, showing that the correlation is higher on the same side than the adjacent side but is always at least 60% correlated. This suggests a crystal-driven noise source. On the right we see two channels corrected for correlated noise according to the methods in the next section. See text for more details.

is both a promising way of dealing with correlated noise as well as a hopeful sign that these noise sources can be controlled to get to our ideal TES-limited noise level. It also motivated more advanced OF techniques detailed in the appendix as well as in the next section that can use multi-channel information to exceed the projected noise performance given a joint-channel signal model.

5.2.5 Scaling Noise with Detector Volume

Some useful techniques to identify sources of crystal noise were identified during our noise measuring/mitigation campaign that deserve a mention at this point, though the work to investigate this is still ongoing. The fact that G115 and G124 had such different volumes and surface areas, and were different from the smaller detector chips and test devices, allowed us to rule out some sources of noise and suggested additional things to study. Here I will briefly lay out these cases without providing any of the ongoing analysis as a suggestion for future work.

The main comparisons we made were between volume, mass, and surface area dominated noise sources. Imagine an extra power noise is found in a detector. We can think of a few possible sources and posit tests to rule each out:

- RF in the fridge - This should scale (as a power) with absorbing area for a detector in the same RF environment, which for Si should just depend on the overall absorption cross-section on the detector face. One suggestion for why G115 had such worse power noise than other detectors was because it had a parquet grid on one side *and* in the negative space between QETs on the other side, meaning any incident RF on the face would be absorbed as current. Since it is very high in frequency, it will just be aliased in the TES current noise as an additional flat power noise. Similar trends were also seen with larger test chips that had high cross-sectional areas for the same resistance.
- Vibrations - this should scale with mass and volume. If the clamps have the same amount of torque, then the Si detectors should have less vibrational noise than the Ge, assuming both surfaces have similar surface roughness (which should be roughly the case). Detectors which are rigidly connected to the fridge (i.e. those that are attached to the MC directly via glue or mechanical clamp) should not show this power noise, although it may be hard to differentiate between sinking the thermal signal (which would also ruin the collection efficiency) and true reduction in vibrational coupling
- Thermal Noise from Muons or Gammas - one theory that has yet to be fully tested was that much of this noise may be due to residual charge generated by muons or gammas in the detector bulk. The idea is that since we have long carrier lifetimes, then the carriers should spontaneously re-combine or trap on impurities with some characteristic decay time for them to re-emit that energy which is relatively independent of the energy of the initial event. This means that detectors may have residual power noise proportional to the input power integrated from all sources (including muons, gammas, etc). So for the same lab environment, the smaller the cross-section of the detector, the less of an effect this should have (less power in and less power out). For sensors of the same volume and T_c , we should then see the power noise decrease with crystal volume, and the power noise should be reduced for deeper or better shielded sites for the same detector.

This is not a complete list of things that could be tested but it does capture the main points we're in the process of investigating. As the CUTE, NEXUS, and UCB test facilities come online as well as SLAC and Stanford, the variety of test conditions and detectors should combine to allow us to separate out these effects. That will depend on having very consistent readout electronics (with no additional parasitic power) and a consistently low base temperature to ensure that any effects we're seeing are not due to base temperature drifts or electronics noise described previously.

5.3 Phonon Energy Resolution

In appendix [E](#), I show that the energy resolution depends on the detector power noise and energy efficiency:

$$\sigma^2 = \left[4\epsilon^2 \int_0^\infty \frac{\tilde{s}^2(f)}{J(f)} df \right]^{-1} \quad (5.67)$$

where \tilde{s} is the signal template normalized to 1 (I've explicitly separated out the efficiency ϵ). Because the HV detectors have so much position dependence, however, it's hard to simply use a calibration source to determine resolution, so for simplicity we separated the resolution measurement into two components: measuring first efficiency, and then using the templates for pulses used to measure efficiency to estimate baseline resolution using measured noise.

5.3.1 Measuring QET Efficiency

We determine QET efficiency using known calibration sources, and matching the integrated pulse spectra (or optimal filter spectra) to known energies from our sources. The equations for determining absorbed QET energy are derived in appendix [D](#) both for the $T_b = 0$ limit as well as the finite bath temperature of small \mathcal{L} limit, and given again here for simplicity. In the $T_b = 0$ limit, we use the equation

$$E_{abs} \approx - \int_{-\infty}^{\infty} \Delta P_J(t) dt \approx \int_0^\infty \left[\left(2 \frac{R_p + R_{sh}}{R_p + R_{sh} + R_{TES,0}} - 1 \right) I_b R_{sh} \Delta I_{TES} + 2(R_p + R_{sh})(\Delta I_{TES})^2 \right] dt \quad (5.68)$$

and for the finite bath power limit, we find the second-order expansion

$$E_{abs} \approx \int_0^\infty \left[\left(2 \frac{R_p + R_{sh}}{R_p + R_{sh} + R_{TES,0}} - 1 - \frac{1}{\mathcal{L}} \right) I_b R_{sh} \Delta I_{TES} \right. \quad (5.69)$$

$$\left. + \left(2(R_p + R_{sh}) + \frac{R_{sh} - R_p - R_0}{\mathcal{L}} \right) (\Delta I_{TES})^2 \right] dt \quad (5.70)$$

which reduces to the first equation in the limit $\mathcal{L} \rightarrow \infty$. These are exact only in the small signal limit, which is approximately the limit

$$\left| \left(2 \frac{R_p + R_{sh}}{R_p + R_{sh} + R_{TES,0}} - 1 \right) I_b R_{sh} \Delta I_{TES} \right| \gg 2(R_p + R_{sh})(\Delta I_{TES})^2 \quad (5.71)$$

$$\left| \left(2 \frac{R_p + R_{sh}}{R_p + R_{sh} + R_{TES,0}} - 1 \right) \frac{I_b R_{sh}}{2(R_p + R_{sh})} \right| \gg |\Delta I_{TES}| \quad (5.72)$$

$$\left| \left(\frac{1}{R_p + R_{sh} + R_{TES,0}} - \frac{1}{2(R_p + R_{sh})} \right) I_b R_{sh} \right| \gg |\Delta I_{TES}| \quad (5.73)$$

$$\left| I_0 - \frac{I_b R_{sh}}{2(R_p + R_{sh})} \right| \gg |\Delta I_{TES}| \quad (5.74)$$

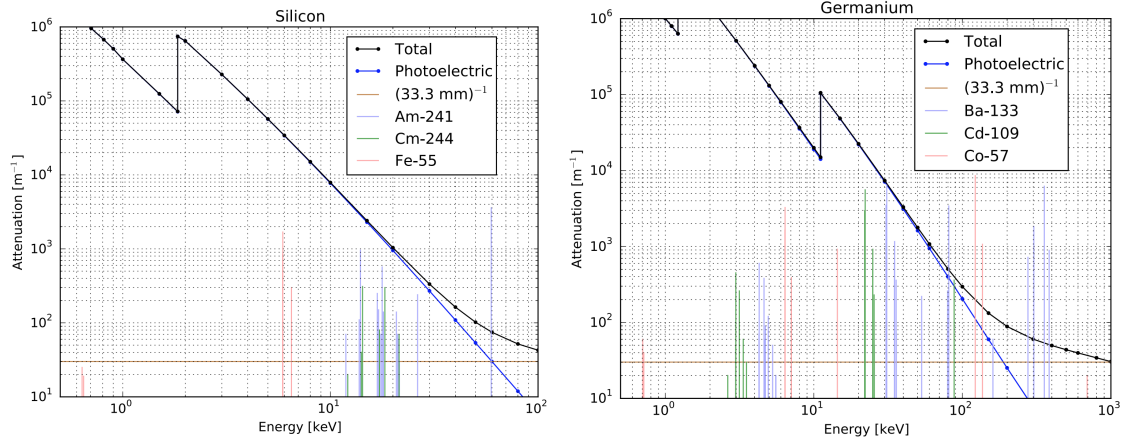


Figure 5.19: Attenuation length for a photon of known energy in Si (left) and Ge (right) compared to the nominal thickness of the SNOLAB detector crystals. Also shown are lines from sources which span the shown ranges, indicating which sources are capable of penetrating a significant distance into the detector bulk. Figures courtesy Paul Brink.

The second term here is the minimum stable TES voltage bias operating point, so what this tells us is that the closer the TES is to its minimum stable point, the less linear the integral response will be. This is obviously good justification for minimizing both R_{sh} and R_p as we've done for the SNOLAB design as well as keeping the operating resistance of the TES at a higher bias than the minimum; it motivates a further study of optimal bias points in a detector of given T_c and bias loop properties.

Assuming that the quadratic correction expands the dynamic range sufficiently, the other part of the efficiency measurement we need to consider is the source to use, which is also complicated by our energy scale. We're trying to measure the baseline energy resolution of a 10 eV device, but the attenuation length of any photon less than ~ 60 keV in Si (~ 120 keV in Ge) is much shorter than the thickness of the crystal, so these events are very localized (see figure 5.19). If we use sources with energies far below these, nearer the expected threshold, we're bound to see surface events which will not be characteristic of the typical events we'd like to be sensitive to. This leaves us with a few options, some of which have been tried, and some of which are planned:

- **Laser (Photon) Calibration:** We can either use LEDs or a fiber-coupled laser to pump small energy deposits into the detector at known times and with relatively well-known total power. These are all surface events, however, and the actual power reached in the detector is systematically limited unless we can directly measure the number of photons either with electron-hole pairs or via some other method.
- **Gamma Source Absorption:** Am-241, which has lines at 3-60 keV, and Fe-55 with lines at ~ 6 keV, are high enough in energy to penetrate into the Ge or Si detectors but low enough

that they're not massively saturating the QETs. They are still very clearly in a non-linear regime even using the quadratic corrections, which means local saturation is a problem. This is show in section [5.3.4](#).

- **Compton Scatters:** There are many features in the Compton spectrum; for example, a 60 keV photon can produce a continuum of energies below 11 keV, and at very low energy well-defined steps can be seen in the spectrum and used for calibration. This requires a lot of statistics (more than can reasonably done for a short run with a test detector) or generally a very hot source, which will heavily saturate the detector. In addition, the relative fraction of photons which escape the crystal after a Compton event is much smaller for large detectors, so this is an inefficient way to do this for our large SNOLAB crystals. If resolution is good enough to resolve electron-hole pairs, this could be a very clean low-energy calibration.
- **Activation Lines:** We can activate Ge with Ge-71 which produces a series of well-defined calibration lines in the detector with energies of 140 eV, 1 keV, 10 keV, and higher. This is, however, not an option in Si, but should be sufficient for Ge if we can get high enough statistics, which we were able to do for CDMSlite. It does, however, require us to radioactively contaminate our detectors, so it's not something we can do often, certainly not during the run.
- **Neutron Calibrations:** Using a monoenergetic neutron beam, or a beam of known and tunable energy profile, should allow us to use simulations along with the measured spectrum to infer an efficiency. This is not as easy or necessarily as clean as using a well-defined line, but it may be our only option for Si HV.

In the rest of this chapter we'll focus on the first two cases, and mention the prospect of using Am-241 Compton recoils in the small HVeV detectors. This is certainly not a solved problem and will be one of our main challenges at SNOLAB.

5.3.2 QET Saturation

The equations above are still limited by the approximation used for bath power, as we also expect local QET saturation to have a dominant effect on our pulses. To get a sense of the energy scale of saturation, we just need to find the heat necessary to drive the TES normal. Given the specific heat c_W of tungsten around 1K, we find that the heat capacity of a TES of volume V_{TES} is

$$C = c_W V_{TES} T_c \quad (5.75)$$

We can define a bias to normal temperature width of ΔT_N which drives the TES normal:

$$\Delta R = R_n - R_0 \approx \alpha R_0 \frac{\Delta T_N}{T_c} \rightarrow \Delta T_N \approx \frac{T_c}{\alpha} \frac{\Delta R}{R_0} = \frac{T_c}{\alpha} \left[\frac{R_n}{R_0} - 1 \right] \quad (5.76)$$

This gives us a saturation energy of

$$E_{sat} = c_W V_{TES} T_c \Delta T_N \approx c_W V_{TES} \frac{T_c^2}{\alpha} \left[\frac{R_n}{R_0} - 1 \right] \quad (5.77)$$

This is an important relation to take note of as we go to lower resistance and lower T_c ; we achieve lower energy resolution but also smaller dynamic range. In addition, the sharper the transition, the lower the saturation energy. Taking the design values for the SNOLAB detectors, we can calculate the expected saturation energy.

$$E_{sat} \approx 108 * (1800) * 1.92 * 10^{-17} * \frac{0.045^2}{150} [3 - 1] = 10^{-16} J \approx 630 \text{ eV} \quad (5.78)$$

This is assuming energy is uniformly distributed across the entire TES channel. If, as in the case of HV operation or thin detectors, energy is location in some fraction f of the channel, we find the saturation point is much lower by simply multiplying by that fraction. As we'll see in this chapter, this is a much lower energy than we observe in our full-scale detectors (they're capable of reading out 60 keV without being fully saturated) but it gives us an order of magnitude estimate of the energy scale. We will see that these 60 keV events are well outside of the linear regime of our detectors.

5.3.3 CDMS-HV Prototypes at UMN

As mentioned earlier, two CDMS HV prototypes were tested at UMN. The first detector was G101c, which was also the first CDMS HV prototype tested. I showed in an earlier section that the noise was entirely limited by the readout electronics, so we're not concerned necessarily with showing whether the resolution is good; the electrothermal oscillations will also severely degrade resolution. Both G101c and S101 were also higher T_c devices, so it should be easier to measure phonon collection efficiency due to their higher TES heat capacity even though they would have a degraded resolution. Given that the phonon dynamics have nothing to do with the detector T_c , the efficiency and fall time measurements can be used to validate our resolution projections for lower T_c devices aside from the assumptions we make about noise scaling.

A selection of results from these two detectors can be seen in figure [5.20](#). Both detectors show the expected phonon collection times, and are consistent with efficiency at or slightly above the 15% expectation. We also see that effectively efficiency decreases as a function of total phonon energy, as expected both from non-linearity and local saturation effects. The figure also illustrates why it is hard to use integral quantities to estimate detector resolution. The resolution of G101c is likely close to 100–200 eV, while the 1.5 keV line indicates something closer to 0.5 keV. This is likely due to uncorrected position dependence, and could be improved on by use of new OF techniques. This has since been done, and shows some of the expected improvement. The upshot is that there is still a large uncertainty on what the actual efficiency will be for low- T_c detectors in a position independent

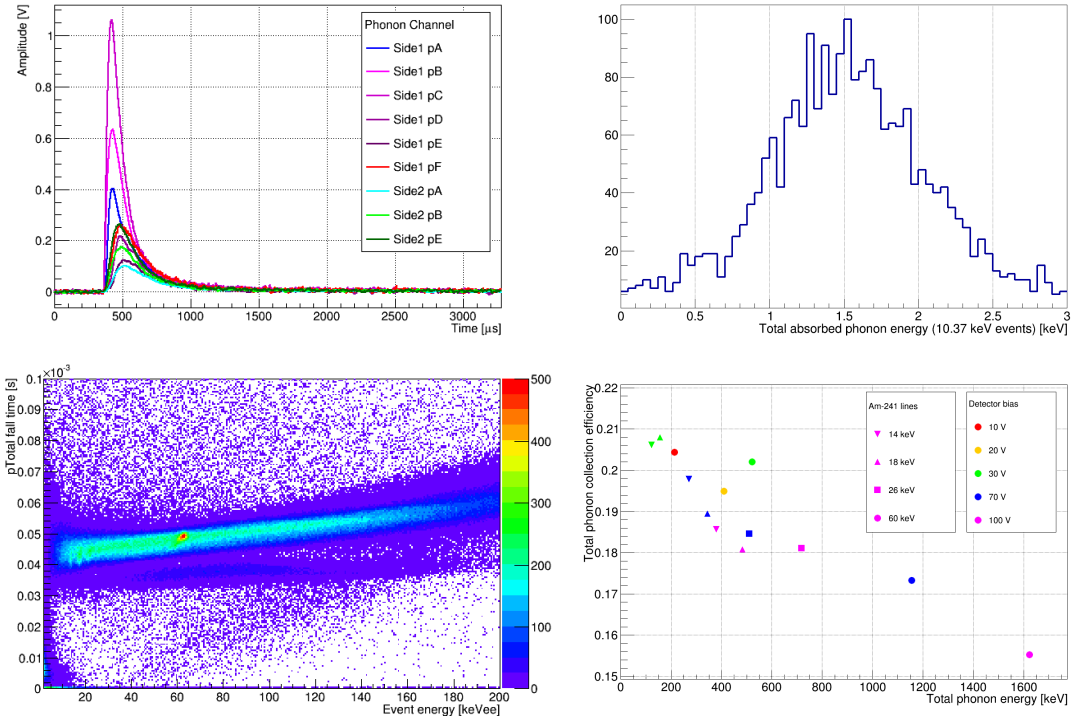


Figure 5.20: Various test results for G101c (top) and S101 (bottom) at UMN. The major results were that the phonon fall-times match those predicted for the Si and Ge CDMS HV designs, and that the efficiency of the prototypes meets and in some cases exceeds the projected targets. We can also see that there is a lot of smearing on the 10 keV line in the top right figure, so that the resolution shown has more to do with signal position dependence than intrinsic resolution. We can also see that the observed efficiency in the lower right panel depends on the absolute energy in the detector, and we already lose signal for 100 keV events. These results are not as severe as expected given that these detectors had a higher T_c and therefore larger heat capacity.

way.

This figure also illustrates some of the methods which were initially tried for detector calibration. The 10 keV Ge activation line was used in the standard way for a uniform calibration source after detector activation, while the Si HV detector was calibrated using the Am-241 source which largely produced surface events. The 60 keV line is likely the best indication of the true efficiency in the bulk, as it penetrates further into bulk, but is also high enough in energy to cause concerns about saturation. The difficulties encountered with these detectors thus motivates looking at G115 to see if we can make more sense of the detector performance and reduce systematics on some of these measurements.

5.3.4 G115 at SLAC and UCB

After operating the first HV detector, we ran G115 at UCB along with a handful of other devices in two runs between the end of 2016 and early 2017. The goal of these runs was to use a thinner detector with less muon and gamma backgrounds, and fewer channels, to try to confirm some of the detector models and study the CDMS HV QET on a simpler device. At UCB the Fe-55 source was used to calibrate another detector, so an Am-241 source with a collimator was used to calibrate G115.

For the initial run at UCB (R480), the collimator wasn't properly assembled, resulting in a much larger incidence of 60 keV events than intended and a much higher source power. The initial efficiency measured for the first run was less than 2%, much lower than the anticipated 10–15%, and as a result of the excess power noise and low efficiency, the resolution came out around 200 eV, when we had hoped for 20 eV. When we ran the second time (R482), the collimator was fixed and the source was much lower activity. As a result, the efficiency rose to around 4%, and the resolution dropped to 50 eV for our best noise data. These data are shown in figure [5.21](#). For both runs we can see that efficiency drops for higher energies, and as we don't have any calibration lines below 3 keV, we're not sure whether the fitted function is correct or whether there's some non-linear transition from 5% or higher at low energy to under 4% by the time the 60 keV event occurs.

The data at UCB being anomalous, both in terms of noise and efficiency, we decided to run G115 at SLAC starting with Run 23. We've already discussed the initial noise issues which occurred during this and subsequent runs in the previous section. The best data taken with G115 at SLAC comes from Run 27 can be seen in figure [5.21](#), where we used an Fe-55 source. The resolution is obviously not great; we can't differentiate between the 5.9 keV and 6.4 keV lines very cleanly, and for this run we probably achieved a resolution of around 200 eV, similar to the original UCB run. Given time constraints and limited manpower at the time, we didn't explicitly measure efficiency aside from verifying that it was no better than at UCB, in the 2–3% range. This was obviously disappointing but not unexpected, and by this point we had accepted a few reasons why G115 was not going to work very well at SLAC either.

So why can we expect G115 to fare any better at another test facility? There are a few possible explanations for why G115 has lower collection efficiency due to environmental limitations:

- Both SLAC and UCB weren't able to cool the detector much below 40 mK for the majority of the runs, Run 23 being an exception due to its length.
- UCB was always bath power limited, but SLAC was not during Run 23. SLAC was always limited by the very high background in the lab from K-40 decays and had no overburden to shield against muons, so the majority of the events in all runs were very high energy compared to the Fe-55. The majority of all data taken in any detector at SLAC had, for one reason or another, remnant thermal slopes which indicate residuals from high energy events. The level

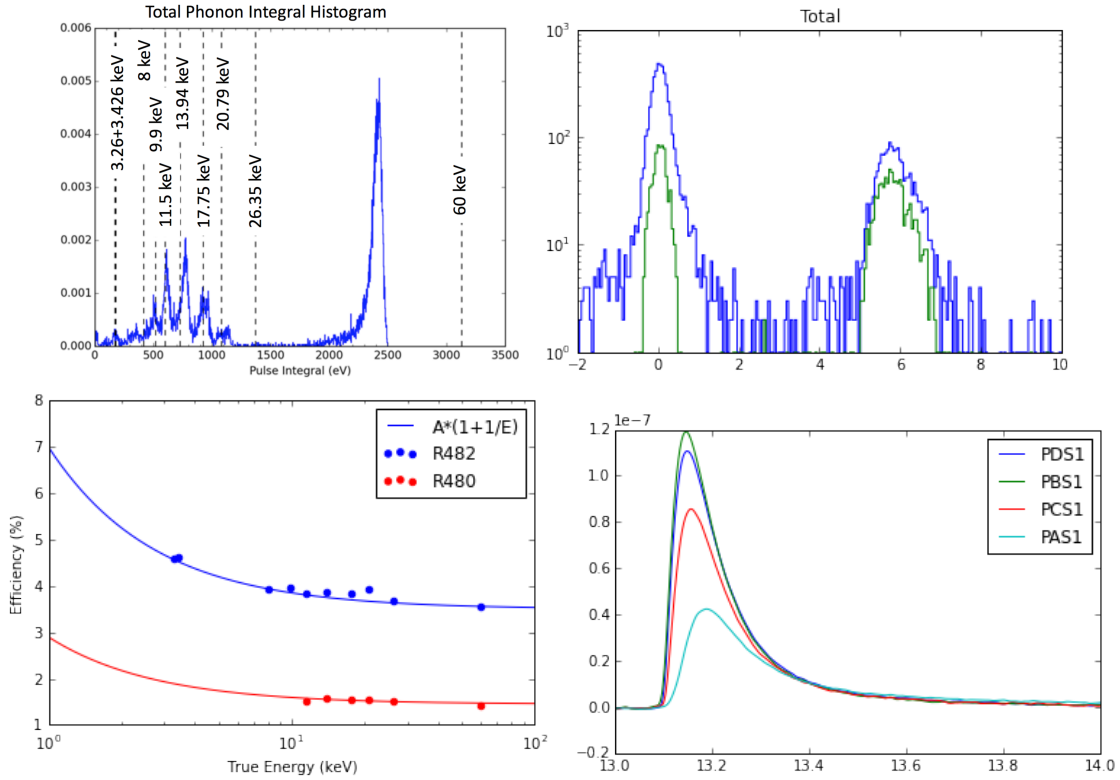


Figure 5.21: Top: G115 spectrum for Am-241 at UCB (left) and Fe-55 at SLAC (right) in keV, showing that both have resolutions for these data on the order of 100 eV. The Am-241 calibration has much more non-linearity seen at higher energy, showing why Am-241 might not be a good calibration choice for our detectors. In the SLAC data, the green points are those that pass the quality cuts. Bottom left: Efficiency for UCB runs where the source presented a much high parasitic power in the first run, lowering the effective efficiency of the device. We can also see the nonlinearity plotted as a function of energy, likely showing how the local saturation loss is reduced at higher energies. Bottom right: pulses from G115 at SLAC (x-axis is time in ms). The long fall time (compared to a much faster estimated TES fall time) is also indicative of strong saturation lengthening the pulse duration.

of pileup was just too high.

- SLAC had demonstrable RF leakage issues that UCB did not; a fridge with low background, low RF leakage, and low temperature may be able to get G115 to work as intended.

A final reason G115 did not work as well as anticipated had to do with some miscommunication during fabrication, in which the backside Tungsten grid was also patterned with aluminum which acts as a large signal sink. G115 also had the parquet pattern on the front-side, which may have acted as an additional phonon sink. It's possible that if the rear electrode were etched off, or a modified version were made, a similar detector could ultimately achieve 10–15% efficiency...but there seem to be some generic issues with thin detectors that we'll discuss at the end of the chapter.

5.3.5 G124 at SLAC

The Fe-55 calibrations with G124 at SLAC at the end of 2017 and into early 2018 were our last efforts to get a CDMS HV detector to meet the SNOLAB project requirements. The goal was to use the detector with the lowest T_c , and put the Fe-55 source over the best channel, to demonstrate that < 50 eV resolution was at least possible. We allowed the detector to cool for a few weeks, watching the bias power drift up, before taking calibration data. The resulting data (for the noise shown earlier) can be seen in figures [5.22](#) and [5.23](#).

The first thing to note is the huge difference in pulse-height between the primary channel and the rest of the detector. This largely comes from the very localized nature of the source; the Fe-55 events should all be very close to the surface, meaning most of the phonon energy will end up in that single channel. In addition, side 1 has much longer fall-times than side 2, due to the lower T_c and the fact that we could not operate in the low bath limit for that side of the detector. There are many possible explanations, but we saw the bias power for those channels increase much more slowly than for the channels closer to the source. Also shown in the figure is the spectrum and resolution measurement showing that the detector channel achieved 29 eV resolution. This also means that for our nominal operating methods we would get 100 eV resolution for the full detector. This is obviously not great news, but there is an upside to this that I'll come back to in a minute.

Figure [5.23](#) shows the efficiency measured for G124 in this run for different channel combinations. The fact that we have a low energy line somewhat close to the baseline, and the fact that this is not done with the optimum filter, means our estimate has pretty large error bars; something like 15–30% resolution. It's very possible that we may have enhanced efficiency for this surface event, or that the position dependence somehow inflates the efficiency, but even conservatively we can say that this detector probably has 20% efficiency. This is also in agreement with UMN, and gives a bit of relief in light of the low G115 result.

So if we have good efficiency, and the pulse falltimes (for the pulses in full ETF) are what we expect, why are we missing our target by a factor of (at best) 10, at worst 20–30? The answer

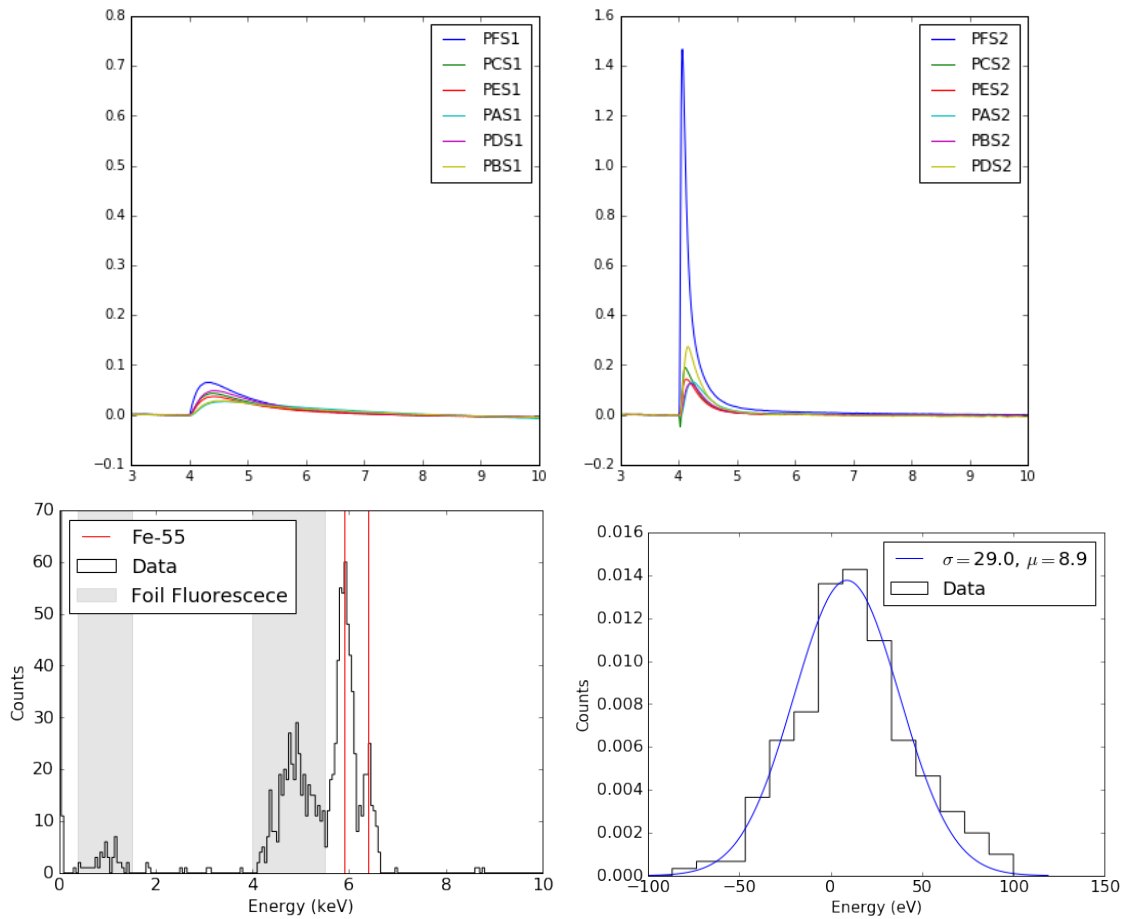


Figure 5.22: Top: G124 pulse template for Fe-55 events selected based on their pulse height in PFS2 (the channel directly below the source). The channel directly below the events is clearly absorbing the majority of the energy, and the remaining pulses still show residual position dependence. Much of the pulse shapes are however distorted by the fact that side 1 had a lower T_c and was not in full electro-thermal feedback. Bottom: spectrum of the Fe-55 source using only PFS2 after optimal filtering (left) and a zoom in near 0 showing the 29 eV resolution of that channel (right).




Figure 5.23: Energy integral for only the central channels on each side, all channels except the channel under the source, and the full detector sum. The lack of optimum filter means the full detector resolution is around 200 eV, and some additional position-dependence brings that up to around 500 eV. Based on these measurements, the detector is 15–30% efficient; likely somewhere around 20% or higher. We were not able to get a better estimate than this due to the noise and base temperature limitations of this run.

lies in a few places, but mostly in the fact that we have elevated and correlated noise between the channels at low frequency which is driving our resolution much higher than it would be if we were really TES limited. These devices do also seem to have a somewhat higher power noise than we expected, but not by more than a factor or two, definitely not a factor of 10. With the introduction of a more robust optimum filter, described in appendix [E](#), we can use the channel templates shown here and the noise covariance matrix to subtract off correlated noise and produce a single energy for all detector channels which is more precise than doing the OF separately and adding the channels.

We can estimate the uncorrelated PSD from noise traces using the formalism of the joint-channel correlated optimum filter. The χ^2 as shown for the single-channel OF under the assumption of uncorrelated noise is

$$\chi_{i,0}^2 = \int_{-\infty}^{\infty} \frac{|v_i(f)|^2}{J_i(f)} \quad (5.79)$$

while for the correlated noise, if we set all channels except for one to 0 (we only compute the

amplitude using a single channel) we find the nominal chi-square

$$\chi_{i,0}^2 = \int_{-\infty}^{\infty} |v_i(f)|^2 \mathbf{\Sigma}_{ii}^{-1}(f) \quad (5.80)$$

where $\Sigma(f)$ is the CSD matrix for all of the channels as described in the appendix. This suggests that the best estimate of the correlation-corrected PSD $\bar{J}(f)$ is found by inverting the diagonal elements of the inverse:

$$\bar{J}_i(f) = [\mathbf{\Sigma}^{-1}]_{ii}^{-1} \quad (5.81)$$

To be clear, we're inverting the entire matrix over all sampled frequencies, and then taking the inverse of a given diagonal matrix element as an estimate of the uncorrelated PSD.

To see whether this makes sense, let's consider a simple model of two channels with their own uncorrelated noises $u_i(f)$ and a correlated noise $c(f)$, such that the noise for channel i is

$$n_i(f) = u_i(f) + a_i c(f) \quad (5.82)$$

where a_i is some complex phase representing a transfer function of the correlated noise into a given channel. This makes the autocorrelations

$$J_i(f) = T^{-1}(u_i(f)^* + a_i^* c(f)^*)(u_i(f) + a_i c(f)) = T^{-1}(|u_i(f)|^2 + |a_i c(f)|^2) \quad (5.83)$$

and the cross-correlations (for $i \neq j$)

$$\Sigma_{ij} = T^{-1}(u_i(f)^* + a_i^* c(f)^*)(u_j(f) + a_j c(f)) = T^{-1} a_i^* a_j |c(f)|^2 \quad (5.84)$$

In the two-channel case, we find that the inverted first diagonal element of Σ^{-1} is

$$[\mathbf{\Sigma}_{11}^{-1}]^{-1} = J_1(f) - \frac{\Sigma_{12}\Sigma_{21}}{J_2(f)} \quad (5.85)$$

$$= T^{-1} \left[|u_1(f)|^2 + |a_1(f)|^2 |c(f)|^2 - \frac{|a_1(f)|^2 |a_2(f)|^2 |c(f)|^4}{|u_2(f)|^2 + |a_2(f)|^2 |c(f)|^2} \right] \quad (5.86)$$

$$= T^{-1} \left[|u_1(f)|^2 + |a_1(f)|^2 |c(f)|^2 - \frac{|a_1(f)|^2 |c(f)|^2}{1 + \frac{|u_2(f)|^2}{|a_2(f)|^2 |c(f)|^2}} \right] \quad (5.87)$$

we see that if channel 2 is dominated by correlated noise ($|u_2(f)|^2 < |a_2(f)|^2 |c(f)|^2$), then we have

$$[\mathbf{\Sigma}_{ii}^{-1}]^{-1} \approx T^{-1} |u_1(f)|^2 = J_1(f) \quad (5.88)$$

which is to say, if we have a second channel which samples the correlated noise, it can be used to remove the correlated noise from the first channel. If at a given frequency channel 2 is dominated

by uncorrelated noise, then we have

$$[\Sigma_{ii}^{-1}]^{-1} \approx T^{-1}|u_1(f)|^2 + |a_1(f)|^2|c(f)|^2 = J_1(f) \quad (5.89)$$

which simply says that if there is no correlated noise in channel 2, you can't use it to remove correlated noise in channel 1; in that case the correlated noise is sub-dominant in channel 2, and it will be a better way to characterize your signal.

For G124 in Run 33, I estimated a single channel resolution of 29 eV without this technique. If we estimate the resolution of a channel using the correlation-corrected PSD (which can't remove all of the noise, but can remove some of it) we get 24 eV. For the full detector, we'd expect this to be a larger reduction, as we add this in quadrature rather than partially linearly. For the full detector, using the NxM optimal filter described in appendix E I estimated a resolution of 17 eV. So the lesson is that the better we know our signal and our noise, the more options we have to achieve a better resolution, though using a joint-channel optimum filter does require us to have a large template library and do a position-dependent OF reconstruction.

Since this analysis was done, we've spent a significant amount of time investigating vibrations and RF interference, and we've made some headway. Caleb re-did this study with data from Run 37 and was able to collection PSDs for the pulse tube off. Using an estimated efficiency of 20%, he estimated that for the pulse-tube off data, with these same templates, we should be able to get channel resolutions of 4–6 eV and detector resolutions of around 20 eV. With the pulse tube on, we've improved a bit to 22 eV per channel and 200 eV total. This validates the idea that our remaining noise is almost entirely vibrational, or at least related to operating the pulse tube. The hope is that the SNOLAB fridge, which has a much more robust vibration isolation scheme, will not have these vibration problems, but in the meantime we're investigating ways to de-couple the system from pulse tube vibrations as a backup plan.

5.3.6 HVeV Detector at Stanford

The last detector we'll discuss in this section is the gram-scale HVeV detectors [89] that have allowed us to study HV behavior without all of the background issues that plagued larger detectors at the surface. The majority of the discussions of this detector will be in the next section, and are the subject of the next chapter, but the efficiency measurement gives us a hint about why there's such a large difference between G115 and the larger CDMS HV detectors.

The HVeV detectors are $1 \text{ cm}^2 \times 4 \text{ mm}$ Si and Ge chips fabricated (at least at the time of this writing) on 4 mm thick 3" high-purity wafers, then sliced from the wafers into individual chips. As a result, their surfaces are not as intrinsically smooth as the larger 100 mm detector crystal which receive the heavy etch before fabrication. They have two channels with 29 QETs in the inner channel and 25 in the outer channel, with a design resistance of around 2Ω . More details of the detector and

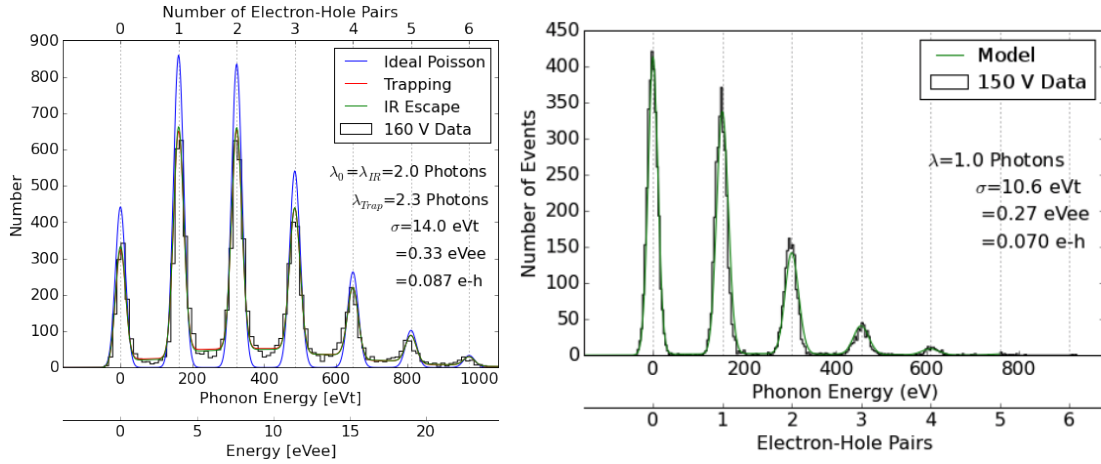


Figure 5.24: Left: Resolution of the first testing run for 160 V data, showing a phonon resolution of 14 eV and a charge resolution of 0.09 e-h pairs. The fill-in between the peaks was attributed to IR photons from the laser, and is described in the next section.

calibration setup are given in Refs [89, 5].

The calibration of the TES for these detectors was much easier than for any of the larger detectors. We used a red (650nm, 1.9 eV) laser connected via a fiber (first wide-band, then single-mode) to the back of the chip in bursts to produce calibration datasets of Poisson distributed photons. Once we are able to resolve single electron-hole pairs, there is a very unambiguous energy calibration we can do, shown in figure 5.24, which gives the detector resolution at 10–14 eV. We predicted a resolution of 1 eV, so this is obviously far from expectation. Luckily it held voltage high enough for us to begin to study our quantized backgrounds. Part of this resolution degradation is due to our noise environment, and part was due to fridge stability, both of which were improved between the first measurement (at 14 eV) and the second measurement (at 10 eV). The other part was that we assumed an efficiency of 10%, which turned out to be an over-estimate. It’s likely that, even without improving efficiency, we can achieve 4–6 eV with this device in the near future.

The very clean laser calibration also allowed us to measure, for the first time, a very fine-grained saturation correction as a function of input energy, shown in both panels of figure 5.25. The panel on the left shows the amplitude versus known energy (based on the number of electron-hole pairs) for the same poisson distribution of electron-hole pairs at different voltages, while the right panel shows the data points at a fixed voltage for varied input laser energies (different poisson photon means). The two calibrations are both consistent below 2 keV, but by going to higher energy we can see that we initially under-estimated the effect of non-linearity. This is the first measurement showing that, for a low- T_c device, we have a demonstrably non-linear result well below 1 keV. This also means that, had we estimated the efficiency for 6 keV Fe-55 lines, we would get an effective efficiency closer to 1–2%. The nonlinearity is thus important to measure down as close to threshold

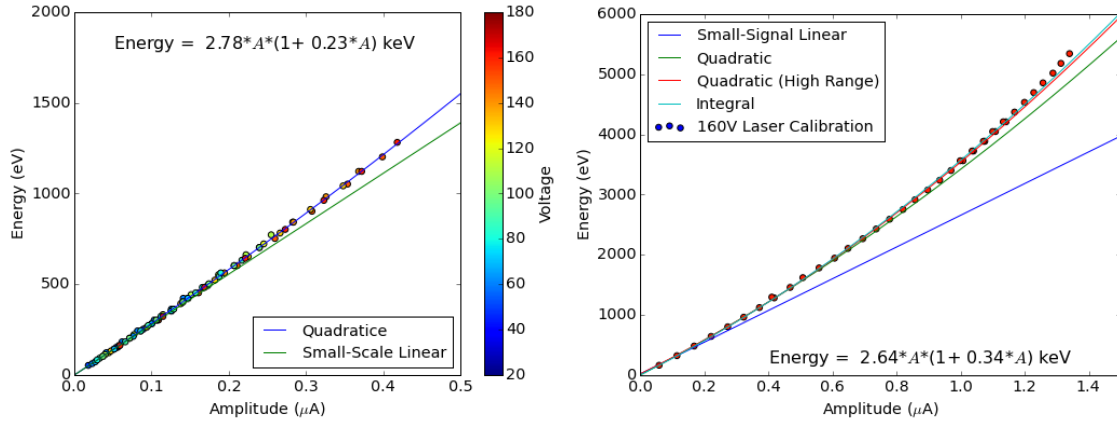


Figure 5.25: HVeV detector calibration for a low- λ dataset at different voltages (left) and a varying λ dataset at a fixed voltage of 140 V. Both calibrations agree up to 2 keV, and are demonstrably non-linear. This supports a conclusion that for a 10 eV device, saturation becomes important at 1 keV, which is roughly a linear dynamic range of 100; if the detector does achieve 6 eV resolution, we have shown we can use a quadratic correction out to 1000 in dynamic range. It is important, however, to explicitly calibration every order of magnitude in energy scale. It's possible that the calibration uncertainty would be reduced by using pulse integrals rather than amplitudes, but the net effect is the same.

as possible, and means that until we achieve quantization sensitivity, we'll have large calibration systematics below our lowest energy data point.

The fact that these devices weren't running in the same fridge where we were using the DCRC for noise and complex impedance characterization means it was hard for us to do any meaningful complex impedance studies until fairly recently, much after these results were published. I was, however, able to estimate the efficiency assuming some R_{sh} and R_p values taken from Jeff Yen's early measurements of the device and with some basic IbIs studies was able to estimate also the operating resistance of the device. For $R_p = R_{sh} = 30 \text{ m}\Omega$ and $R_0 \sim 0.2 \pm 0.1 \Omega$, I found an efficiency of around $5 \pm 1\%$. It's possible these were over-estimated, however, and it's more likely that the detector has an efficiency around 3–4%. In any case, this is in agreement with G115, which was also 4 mm thick.

Both G115 and HVeV detectors were not given the same level of surface treatment as the larger detectors, and they were both 1-sided. One can imagine that rougher surfaces may have contributed to phonon down-conversion and allowed phonons to thermalize faster, suppressing phonon efficiency. We could argue that the position dependence of the G124 data doesn't really support that theory, but then again that data has its own issues as well. Whatever the reason, the fact that all 100 mm detectors achieved 20% efficiency despite their issues, while G115 and HVeV detectors (both different substrates) have limited efficiency suggests a route for elucidating how to improve efficiency going forward.

5.4 High Voltage Performance

Finally, we come to tests of the feasibility of the NTL gain method. I will briefly review the results from UMN demonstrating that the detectors can hold 100 V across the bulk, and discuss the limitations and implications of these initial measurements. We'll discuss the impact of gamma and muon events on the minimum attainable leakage current, and then I'll discuss the HV test results from the HVeV detector as published in Refs [89, 5].

5.4.1 Leakage Current

In this section, we need to define some terms and see how they'll impact our detector performance before jumping into the results. We'll differentiate between two types of 'leakage' currents, 'surface' and 'bulk' leakage. They have varied origins and different behavior depending on the source, but they are the only two meaningful categories for these studies.

'Surface' leakage refers to a leakage through the bulk of the detector from one electrode to the other, and in its simplest form should depend on the voltage across the detector as if it were an ohmic resistor or a diode. In the first case, the leakage current will increase linearly with voltage, and in the second it will increase exponentially with voltage above some cutoff. We can thus define the leakage as a function of voltage as

$$I_{surf} = \frac{V_{bias}}{R_{det}} + I_{surf,diode} \exp(V/V_{diode}) \quad (5.90)$$

The second case clearly will quickly become breakdown, and we'd like to avoid it at all costs...if we can measure it that's bad news. Let's assume that's negligible for now. Realistically, the detector resistance is also not really a constant, because at these temperatures our substrates are non-ohmic. The better way to think of this is really as the resistance of the interface layer..if there's a finite resistance, even voltage dependent, it has to at least increase linearly if not faster.

The other leakage is a bulk leakage, i.e. it results from some event in the bulk (either a decay of a metastable bound state or IR excitation) and it produces only one charge species, but it's roughly voltage independent. We know the voltage independence can be the case because the auto-ionization probability of even a meV-scale bound state at our temperatures is essentially 0 until we get above kV/cm field strength (see Kyle Sundqvist's thesis[99] for more details on this). This means that the bulk could be

$$I_{bulk} = e^- \Gamma_{bulk} = e^- [\Gamma_{IR} + \Gamma_{Bound}] \quad (5.91)$$

Essentially this will be fixed by crystal properties, and will increase in both cases proportional to volume. If introducing IR filtering or neutralizing the crystal helps, we should expect that the leakage is bulk driven for the most part.

In our HV detectors, the additional complication is that we don't measure charge, we measure

phonon energy. For a surface event, charges drifted to the opposite surface always experience the full NTL gain, so the leakage power is

$$P_{surf} = V_{bias} I_{surf} = \frac{V_{bias}^2}{R} + O(V_{bias}^3) \quad (5.92)$$

whereas for a bulk event, which is equally probably throughout the crystal, the *average* event only sees half of the voltage, such that

$$P_{bulk} = \frac{V_{bias}}{2} I_{bulk} = \frac{e^{-\Gamma_{bulk}} V_{bias}}{2} \quad (5.93)$$

These are both important results. This tells us that even if we aren't able to resolve the leakage, we can determine its cause by how the charge leakage power scales with voltage (by measuring changes in TES bias power or power noise). If we fit the power versus voltage curve to a polynomial, then the linear component measures the bulk rate and the quadratic and higher components measure the surface rate.

The first measurements of CDMS HV power versus voltage are seen in Figure 5.26 made by Matt Fritts at UMN with G101c and S101. In the Si detector, the story is very clean: there's a 1 pA leakage which is voltage independent to the extents tested. In Ge, the leakage is clearly interface-limited and polarity dependent. Our insulating layers are amorphous Si on both detectors, so it's likely that they create a better charge leakage barrier in Si while acting like a diode in Ge. From this point on, we can assume either that this particular Ge detector was a dud (possible but not likely) or that the Ge HV detectors are going to be much more difficult to operate at high voltage, and we're going to take a resolution hit (more likely). Given our 0V reach with Ge, and that Si has better reach for electron-recoil dark matter, this isn't killer, but it's certainly an area where it's important for future R&D.

5.4.2 Muon and Gamma Charge Production

Let's pause for a minute to consider all of the sources of bulk leakage to make more sense of the leakage rate in Si. As a very rough rule of thumb, we expect one electron-hole pair per ~ 3.8 eV of recoil energy for an electron recoil, less by some quenching factor for a nuclear recoil at high energy. If we have some rate of events of energy E_R , then we can find the deposited power as

$$P = \Gamma_{E_R} E_R \left(1 + \frac{y(E_R)V}{3.8} \right) \quad (5.94)$$

This gives us the same power scaling as a bulk leakage, so we can always assume that any number of unresolved events of a high enough rate, or the residual fraction of larger events, can be interpreted as an elevated bias power or power noise. For a background of known differential rate $d\Gamma/dE$, we

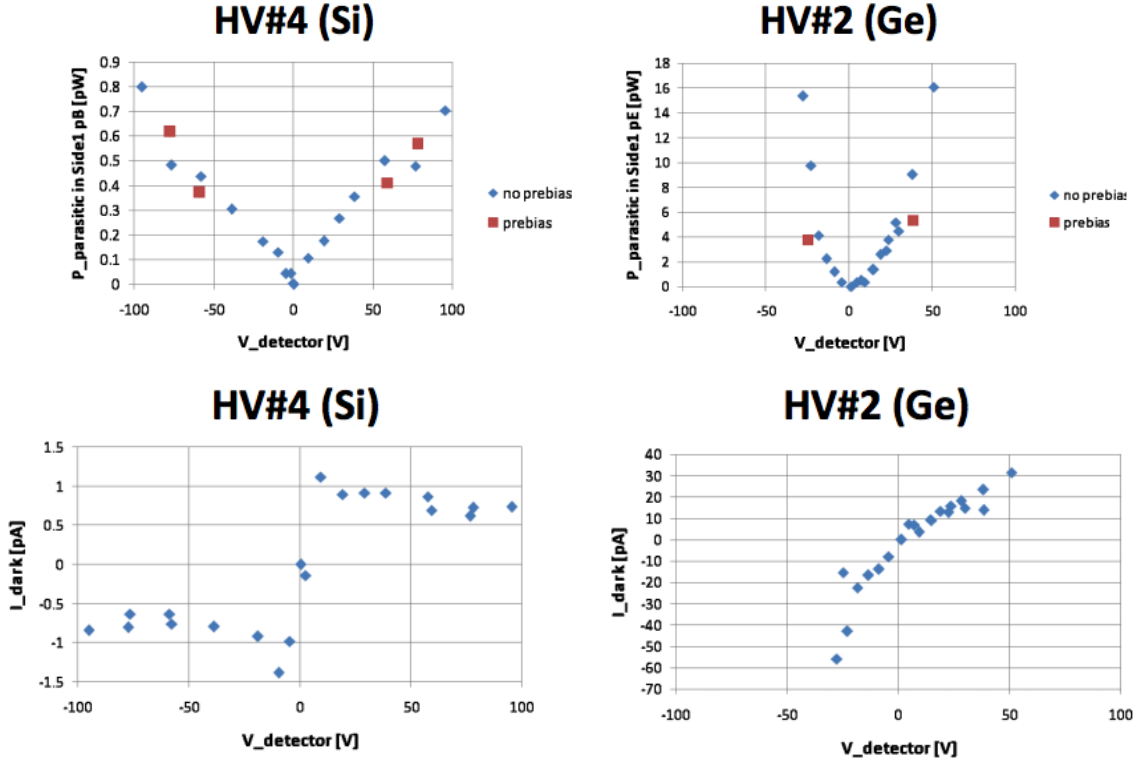


Figure 5.26: Top: Bias power shift measurements as a function of one-sided bias voltage (one face operated at 0V, the other biased but not readout at voltage) for Si and Ge. Bottom: Top plots divided by bias voltage to infer the effective leakage current as a function of bias voltage. We see a constant leakage current in Si, and an asymmetric leakage in Ge, likely due to polarity-dependent charge breakdown at the interface. We can see that the positive bias is linear (probably ohmic leakage) but the negative bias has a diode-like increase. Figures courtesy Matt Fritts.

find

$$P = \int_0^\infty \left(\frac{d\Gamma}{dE}(E_R) \right) E_R \left(1 + \frac{y(E_R)V}{3.8} \right) dE_R \quad (5.95)$$

This can be broken into electron and nuclear recoils to get the expected power for a total spectrum composed of known components of both, but our backgrounds are likely to be dominated by electron recoils, so we can just assume for simplicity that $y(E_R) = 1$.

The background spectrum in most labs is dominated by K-40, Uranium decay products, and muons; the shielding smears out the radioactive backgrounds so they look mostly like a decaying exponential with a mean energy around 150 keV. Assuming this occurs at a rate of 30 Hz (as measured in G124) we can see an order of magnitude estimate for the power would be ~ 4.5 MeV/s or 0.7 pW at 0V, equivalent at high voltage to a leakage current of 0.35 pA. We can remove some fraction of these events, but we can already see that a high background rate is not conducive to

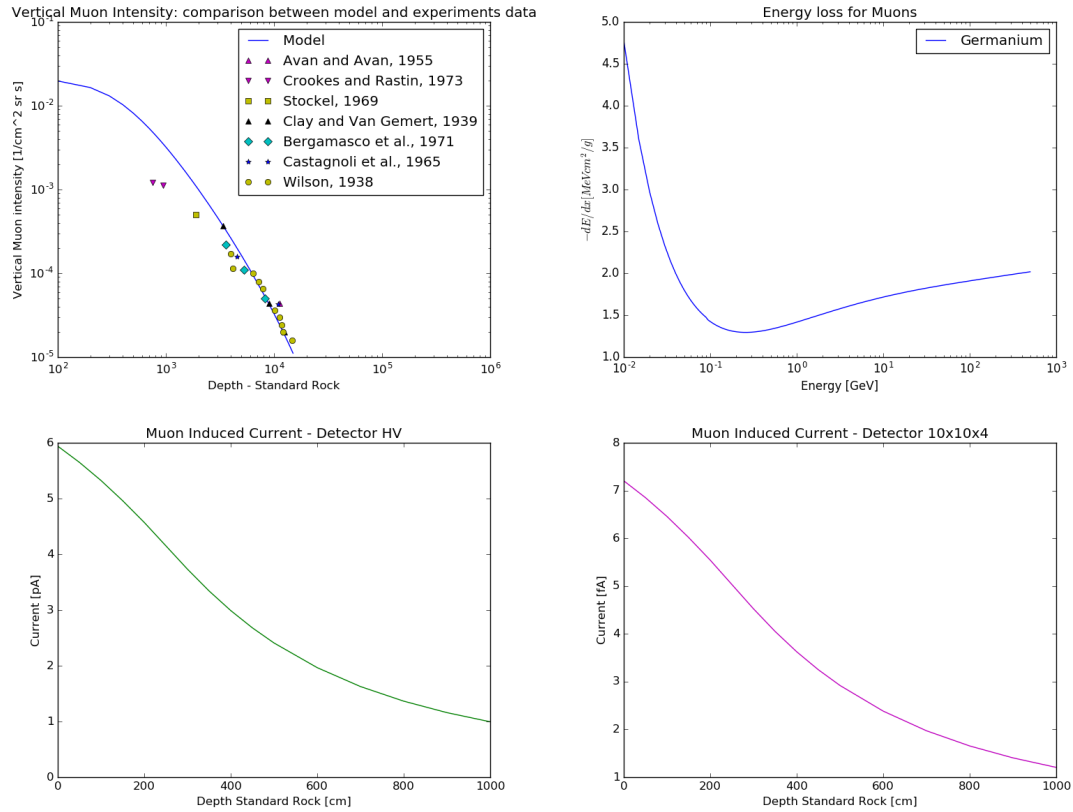


Figure 5.27: Top left: muon flux as a function of overburden from standard rock. Top right: Bethe-Bloch equation for germanium showing energy loss as a function of muon energy. Bottom: Muon-induced current for the Ge HV detector (left) and Ge HVeV chip detector (right). For Si, the rates are the same but the energy deposition is reduced by a factor of about 3 (the ratio of the densities). Figures courtesy Carlo Gilardi.

these low-power devices with pA-scale bias power.

Muons are a much worse problem, however. At the surface, the rule of thumb for muons is that you'll get about 1 'per hand' per second, or about 4 Hz at the surface for the HV detectors. Muons are minimizing ionizing particles and deposit tracks in the detector with about 2 MeV/g/cm of energy, which at 0V corresponds to a power flux of 110 MeV/s (in discrete events) or a leakage current of 6 pA. We see these events, and we can reject the initial events, but there are two problems we'll have regardless of what we try. At high voltage, the power is increased; at 100V, the power is ~ 1 nW, and the decay time of a thermal pulse is around a quarter of a second. This means on average a few percent of that power remains in the detector, looking like a parasitic power.

The second problem is that this is a lot of electron-hole pairs, and we'll see later that there's a relatively high probability that some fraction trap during transit, allowing them to be spontaneously

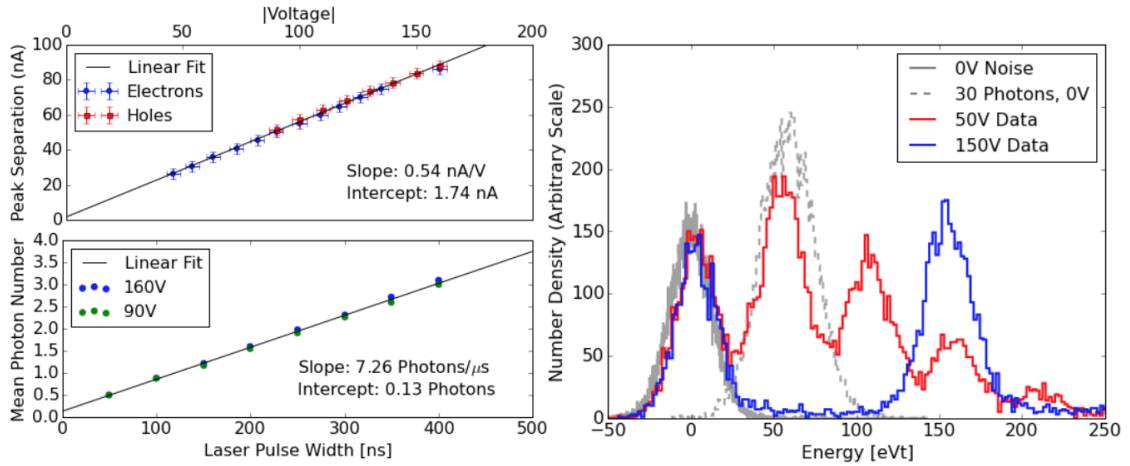


Figure 5.28: Left: NTL gain linearity (top) and laser linearity (bottom) tested by varying bias voltage and laser intensity respectively. Right: Demonstration of NTL gain linear signal/noise scaling, showing that the 150 V data has the same energy noise but a massive electron-hole pair gain resulting in an ideal resolution that scales inversely with voltage. See text for more details, figures from Ref [89].

released or IR ionized at a later point. Both of these factors can easily conspire to generate a pA-scale parasitic power effect at high voltage, and are the primary reason we needed to run smaller detectors to study HV behavior at the surface. Muon currents and rates for large and small detector form factors can be seen in figure 5.27. The figure also shows that switching to a smaller detector form factor should make this a significantly more tractable problem.

5.4.3 First NTL Gain Demonstration: CDMS HVeV

The first NTL gain demonstrator, documented in Ref [89], was the Si HVeV detector, which was discussed briefly in the previous section. I already showed that we were able to demonstrate 10–14 eV phonon resolution, but that was only possible because the device was able to harness the NTL gain without any increase in baseline resolution. The gain and linearity tests performed for the HVeV detector are shown in figure 5.28. We show first that the photon number fit to the laser spectra is linear in laser power, and very stable with voltage, so that we can trust the laser power as an energy scale. We also show that for both biases the position of the first electron-hole pair peak is linear in voltage, showing that this gain method works and there is not appreciable impact ionization which would cause an upward trend in this plot.

The best illustration of this technique, however, is the right side of figure 5.28, where we show four run conditions. The first condition has no laser data, and produces only a Gaussian noise distribution about 0. The second condition is at 0V as well, but with 60 eV mean photon energy,

resulting in a slightly larger distribution centered around 60 eV; there's some broadening due to Poisson statistics in the laser spectrum. We then reduce the mean photon number to around 1, and we see that the 50 V data just overlaps the same Gaussian noise at 50, 100, and 150 eV, and the 150 V data simply moves the single e-h pair peak up by a factor of 3.

The immediate problem with this measurement, however, is that there is a non-trivial population of events between the 0 and 1 electron-hole pair peaks, which isn't consistent with our story of quantized electron-hole pairs. As shown in figure 5.24, our calibration data is consistent with a quantized event population either coincident with IR photons, which are able to free sub-gap bound states and produce a fraction of an electron-hole pair, or produced by impact ionization or trapping during propagation. These can be tested individually with varying laser inputs, and in the initial paper we hypothesized that for the laser data, these were primarily IR events with 3% impact ionization and trapping [89]. The additional IR filtering did reduce the fill-in to the 3% level, validating this hypothesis, but our backgrounds are still dominated by both processes.

5.4.4 Measuring IR Leakage

One of the first studies we did to characterize the population of quantized events in our detectors was to use the delay optimum filter (see appendix E) to search random traces of pulses. The downside of using the delay OF is that the longer the trace is, the higher the probability of finding a random fluctuation away from 0, which produces a binomial noise distribution. By reflecting the spectrum around 0, we can subtract off the noise to infer the distribution as close to 0 as possible, with the results for each bias polarity shown in figure 5.29. We see that for most biases, there is a voltage-independent flat background, which is what we'd expect for IR or spontaneous emission events which produce a single charge instead of an electron hole pair (see appendix G for more details). Note that for a pure IR or spontaneous background, we'd expect a tail of events, but as shown in figure 5.30, our timing information allows us to get essentially 0 pileup.

The addition of the IR filters and single-mode fiber did reduce the IR background by a factor of 10 as measured by this method, and the modeled charge leakage from the long-exposure science run as seen in figure 5.29 (2 Hz above 0.8 e-h pairs, assumed flat, gives us 10 Hz). Let's compare the leakage we obtained to that of competing experiments. We obtain, with our filtering, an IR leakage rate of 10 Hz (assuming all non-quantized events are IR-driven). SENSEI and DANAE achieve dark rates at the surface of $1e^-$ per pixel per day, which for SENSEI is around 8–9 Hz/g and for DANAE is uncertain but at least 2 orders of magnitude larger. We achieve something like 11 Hz/g at the surface, so the difference between the HVeV detector and the SENSEI detectors really does come down to timing. That being said, DAMIC have demonstrated a dark rate of $< 10^{-3} e^-$ per pixel per day, or about 10^3 better than any surface detector, at 120 K [6]. This was for a device with worse charge resolution ($1.8 e^-$) underground and with significant shielding and IR mitigation. Whether the improved dark rate can be attributed to a lower muon background (and thus less filled

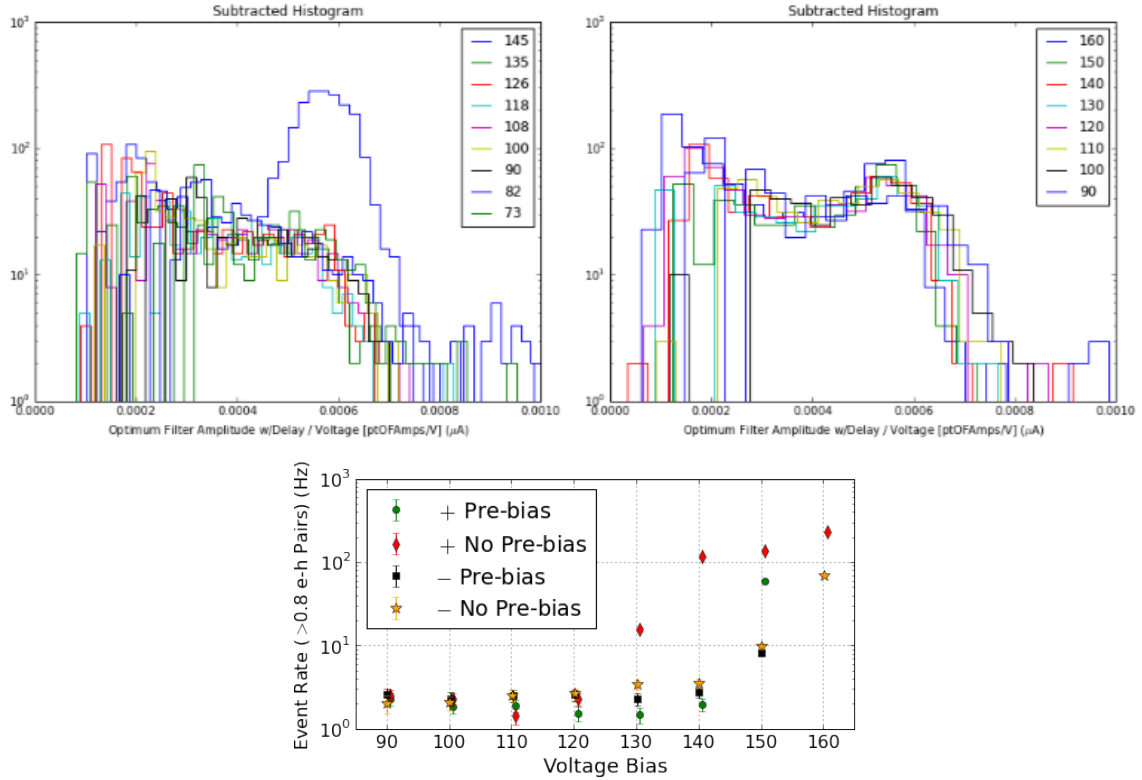


Figure 5.29: Top: Leakage events from randoms with the 0-signal noise subtracted by reflection as a function of voltage for the positive (left) and negative (right) voltage bias. We see the beginning of quantized charge breakdown around -140V, which was dependent on pre-bias, but that the overall rate and shape is largely voltage independent and continuous; a bulk leakage consistent with IR. Bottom: Background rate above 0.8 e-h pairs, which had an identical shape but was at a 10 times lower rate after IR filtering, suggesting that it is still IR or bulk dominated and again is independent of voltage. Here we again see that the leakage rate for the positive data increases asymmetrically with the negative bias, but this can be mitigated by pre-biasing the detector at 10V higher than the intended operating bias, suggesting that it's possible this is not a fundamental limitation but depends on the neutralization state of the crystal.

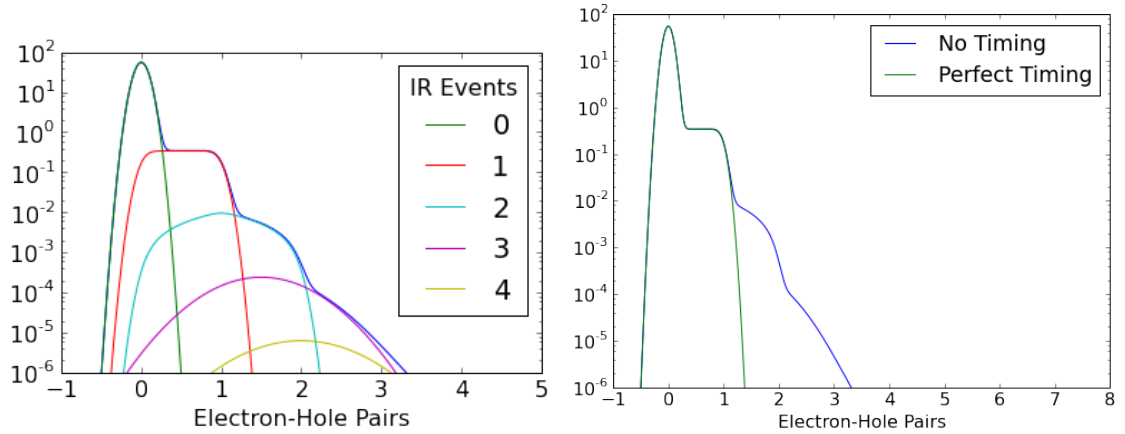


Figure 5.30: Left: IR Spectrum integrated over an hour without any fine-grained timing rejection showing how coincidences over a 1 ms period create the compound shape with a high energy tail. Right: IR spectrum assuming either no pileup rejection or that all multiple events can be separated, showing that even with a high IR rate, good timing allows for background-free operating at 2 e-h pairs.

IR targets) is an open question.

5.4.5 Impact Ionization and Trapping

The assumption made by most experiments when making sensitivity projections based on dark rates is that all dark events are uncorrelated; that is, that dark rates themselves cannot produce additional dark rates. This is not actually the case, as high energy charges can free additional sub-gap charge states to produce more free carriers, or carriers can become bound or even re-combine with a finite probability, and we lose a charge from an already small charge population. For a dark matter search, impact ionization is a larger concern, as it allows us to have correlated backgrounds above a single electron-hole pair which are inseparable from the true charge carriers. Worse yet, these carriers can then generate additional charge carriers, giving us a completely time-coincident high-energy tail that fills in the >1 e-h pair bins. An approximate form is discussed in appendix G, but an accurate spectrum for probabilities larger than $\sim 1\%$ needs to be obtained through Monte Carlo.

As will be discussed in the next chapter, we were somewhat limited by impact ionization in our first science run, and we made some attempts in pre-run calibration to quantify impact ionization and trapping from laser data and randoms. The results of these studies can be seen in figure 5.31. The first method involved using a period of high surface leakage, where only single leakage events occurred, to precisely measure the impact ionization probability. This came out to around 4% as expected from fits to laser calibration data, and the shape is an excellent match to the model; this also allowed for some estimate of IR rate but the threshold was too close to the single e-h pair peak

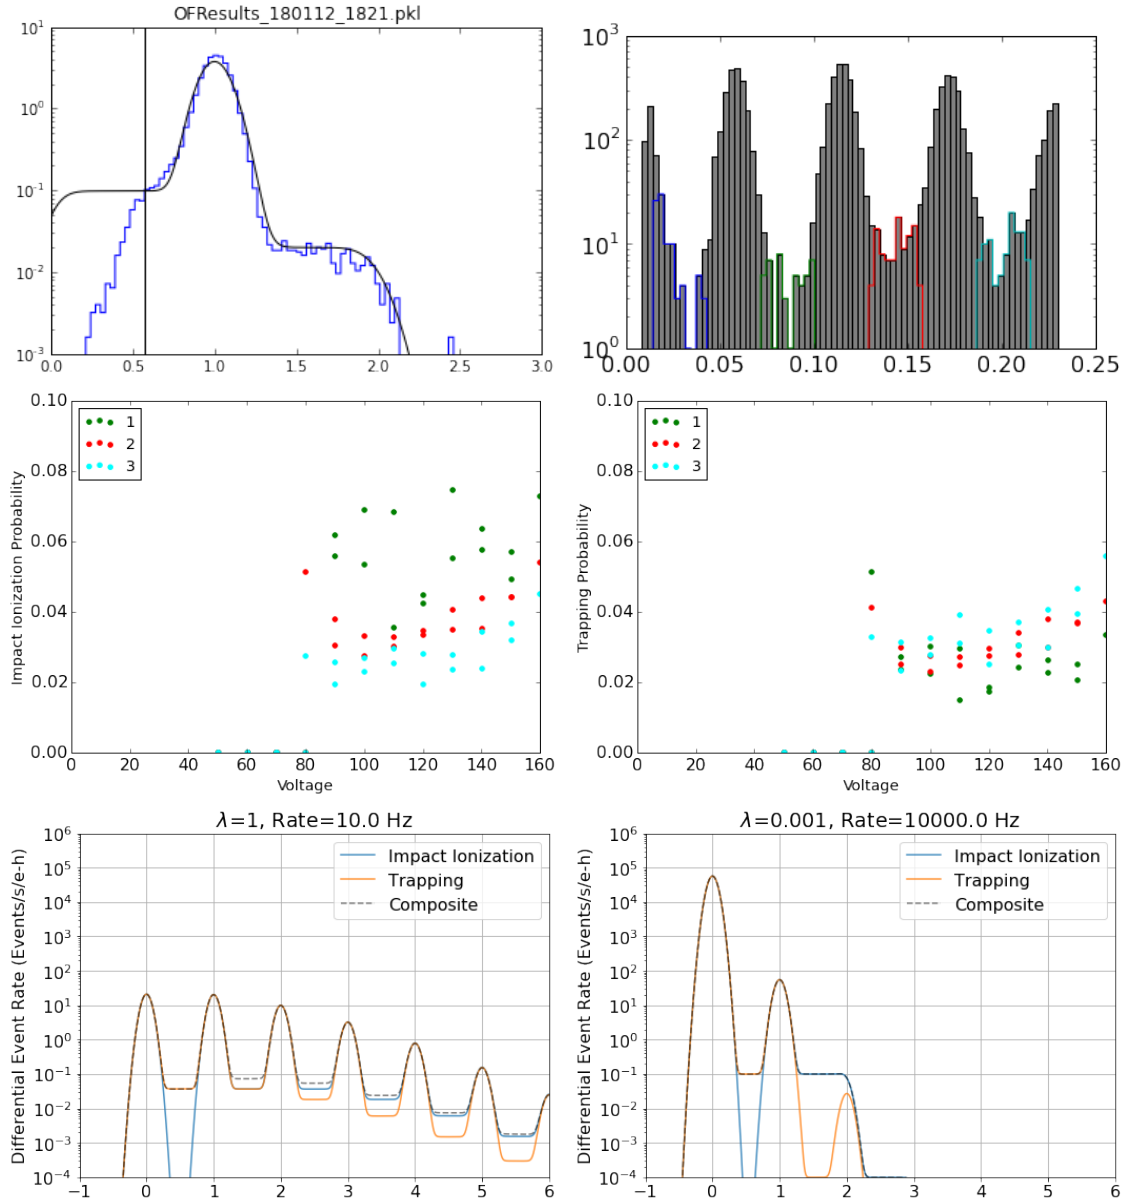


Figure 5.31: Top Left: Leakage data used to estimate impact ionization probability, consistent with 3% impact ionization (x-axis is electron-hole pairs). Top Right: Laser calibration data used to measure impact ionization and trapping by selecting inter-peak regions (x-axis in OF amplitude). Middle: Impact ionization (left) and trapping probability (right) inferred from inter-peak regions; see text for more detail. Bottom: Propose laser calibration showing impact ionization and trapping at the 1% level for laser data with a normal λ (left) and a very low λ at a high rate (right).

due to the high rate.

The other method involved using laser bias data at each voltage and finding the ratio of events between the peaks to the photon peak on either side to estimate upper limits on the impact ionization or trapping from each peak, assuming all events in that region resulted from each process. This was done by selecting all events $> 2\sigma$ from a peak as fill-in events, and events within 2σ of a peak as quantized events. This is a very systematics limited way of measuring this, but we can see that impact ionization is limited to $\sim 3\%$ independent of voltage bias, and trapping is limited to $\lesssim 2\%$. Taken together, the fact that we see so few events between the 0 and 1 electron-hole pair peaks likely means trapping is negligible at these voltages, and we're dealing with a voltage-independent impact ionization mechanism, but more precise measurements are needed to eliminate statistical systematics and the possibility of confusing IR events with the other mechanisms, assuming we can still get IR into the fridge from the laser.

Figure 5.31 shows in the bottom panel modeled detector response for a proposed calibration to more accurately and independently measure impact ionization and trapping by running the laser at a high rate with very low mean photon number. This allows for the trapping and impact ionization plateaus to come through above other backgrounds, but assumes that the laser doesn't appreciably emit IR. The limit of this method is set by how cleanly the laser can be operated at this high rate and as always how well our IR filtering in the fiber operates. Ideally this will allow us to determine whether either process is highly voltage dependent, which might suggest it's better to run at a lower voltage for a cleaner background.

5.5 Future Studies

This was a long chapter full of initial studies that could and should be followed up on. In particular, there are a few questions I've highlighted here that are crucial to investigate further:

- The detector crystals do not cool nearly fast enough, likely due to clamp conductance near base temperature; better clamps will drastically improve our ability to operate detectors below 50 mK. Much of the issues which arose during operation can be attributed to temperature differences between the crystal and the bath.
- We need to determine how to better mitigate low-frequency noise in dry dilution refrigerators. Pulse tubes make for more stable operation, and really should be the way forward, but at the time of this writing we have not operated a detector in a dry fridge and attained good transition noise performance with the pulse tube on.
- We also need to solve our EMI problems; the telegraph noise is in particular the worst problem, but RF should not be able to penetrate into the experimental space. If this means limiting the bandwidth of all lines into the fridge, and putting the pre-amp and SQUID loop inside,

the fridge, then it's likely worth it. This is the setup in the Northwestern ADR and it seems to be more robust to telegraph and RF noise, though any hole for RF to leak in becomes the weak link.

- Low-energy calibrations are also a weak point for our collaboration...it's not going to be feasible to use a laser with a fiber long-term, nor is it reasonable to put a dark matter detector in a beamline and use it for a low-background search. If we can think of clever ways to produce 100 eV-scale events in the bulk in Si and Ge without lots of activation that would be a big step forward.
- Understanding power noise - we don't have a consistent way to model our power noise which relates to the QET design. The 3-pole complex impedance, which is also seen by other low- T_c TES experiments such as CRESST, may represent something fundamental, but it's likely that our detectors have such low heat capacity that parasitics are becoming more important.
- We need to understand the character of leakage current and how neutralization and pre-biasing, as well as backgrounds and crystal purity, impact it. We also need to determine a way to test our IR shielding structures to see if they're really IR tight.

A lot of these are conclusions we couldn't begin to reach until we started testing these detectors, and solving any of them is a big step in the right direction. It is likely that many of them will be addressed in the coming year and hopefully inform final design elements of SNOLAB, or at least allow for upgrades to solve these problems during the first shutdown. They will also enable huge strides forward in the R&D programs which currently struggle with many of these problems while trying to get prototypes to work.

Chapter 6

First Sub-GeV Dark Matter Results: CDMS HVeV

“We shall see what we shall see. We have the start now; the developments will follow in time.”

- Wilhelm Rontgen, 1896

This chapter describes the dark matter search presented in Ref [\[5\]](#), which demonstrated the use of NTL amplification to achieve single-charge resolution and used this device to set new limits on Sub-GeV dark matter. I will discuss the analysis done to produce the final limits, the resulting spectrum, the signal model, and then the limits. I will contrast them to similar experiments and discuss where we go next with these detectors in terms of dark matter discovery potential.

6.1 Experimental Setup

The HVeV detector was a $0.91 \text{ g } 1 \text{ cm}^2 \times 4 \text{ mm}$ Si chip cut from a 3-inch high-purity ($\sim 20k\Omega\text{cm}$) Westinghouse wafer. The detector design and holder setup can be seen in figure [6.1](#) and is described in detail in Refs [\[89, 5, 90\]](#). The QETs are square with 6 fins that each have about 200 micron length, and have holes in them to confine magnetic vortices. There are 24 QETs in the outer channel and 25 QETs in the inner channel; each channel is given roughly equal area, though there is not uniform surface coverage.

In the previous chapter we explored the performance of the HVeV detector, and for the science run we needed to determine how to calibrate the detector, in what mode to operate, and how to bias the crystal and TES channels. We used the measurements from figure [5.29](#) to choose -140 V as our operating bias with 5 minute pre-bias periods of 10 V every 10,000 events. We set the laser to fire at




Figure 6.1: Top Left: Initial setup of detector and HV holder on the dilution refrigerator probe, with the front of the detector showing. The two halves of the holder are biased independently, and the fiber is glued to the back of the holder. Top Right: IR filter upgrades between the fiber and detector as well as a side-view of the detector holder. Bottom: Image of the detector before mounting with zoomed in images of the QETs and the backside parquet bias pattern.

a rate of 1 Hz, and recorded the laser TTL signal in a third analog input with a DC offset tracking the fridge thermometer resistance to get event-level information about the fridge temperature. We also recorded the trigger channel, which was the sum of the two channels after putting them through a shaping amplifier, and used a level trigger on this channel as our primary trigger. Recording the signal used for triggering allowed for a very clean estimate of trigger efficiency during the run as a function of pulse energy, and the level trigger was set to get 50% efficiency at 0.5 e-h pairs (in order to lower the overall trigger rate) and achieved $> 95\%$ efficiency at 1 e-h pair.

We also intended to look at resolution as a function of TES bias for both channels, but without an automated way to change bias this test fell to the wayside unfortunately. We can see that the current noise is fairly constant but the separation between peaks does change during the transition (there is a correlation between thermometer resistance and amplitude shown in figure [6.7](#)). We chose the lowest stable operating point for each TES based on early calibration runs and maintained this same bias point throughout the science run.

6.1.1 Raw Exposure

We began acquiring science data on February 9 (2018) with the goal of getting to 48 hours of total exposure, assuming that about half would have noise or temperature fluctuation issues. We assumed we would run for a couple of days but ended up acquiring data on and off for a week, until February 16, due to more noise and equipment issues than anticipated. The best period for data taking was usually between midnight and 8 am, which indicated to us that activity in the building contributed to some of the noise issues we saw during the day.

During the run, we noticed significant changes in the trigger rate and the character of the leakage (all data was pre-processed in an automated way for data quality monitoring). The trigger rate variations were largely correlated with temperature fluctuations (or at least with the measured resistance of the thermometer), but there were also changes in the leakage which were very sudden and not obviously correlated with any change in the detector's environment. The temperature fluctuations (and thermometer resistance) during the run are shown in figure [6.2](#), and the leakage as a function of time (as well as an example of time-dependent leakage) can be seen in figure [6.3](#).

To try to mitigate the increased leakage seen around 50–70 hours into the data taking, we neutralized for a few hours. This had the effect that it did seem to get rid of the burst of time-dependent leakage, but it also seemed to elevate the bulk leakage rate. This will be discussed in more detail later in the chapter. After about a day, some of the surface leakage started to re-appear, so we tried to flip the polarity of the voltage, which did not seem to affect the character of the leakage, but it did coincide with a large increase in temperature variability, as shown in figure [6.2](#). The corresponding increase in the trigger rate is just due to the fact that the pulse amplitude correlates with voltage, so as temperature fluctuated the fraction of the bulk background below 1 e-h pair (and the fraction of noise triggers) increased in a correlated manner.




Figure 6.2: Top: Temperature as a function of time during the science run. Neutralization occurred at 70 hours, and the polarity flip occurred at 90 hours. Bottom: Thermometer resistance as a function of time. The calibration from resistance to temperature is highly non-linear and includes large systematics, so we opted to use thermometer resistance instead of temperature as the independent variable for temperature-dependent calibrations. The thermometer resistance should also have a similar temperature dependence to the TES bath power, so it should allow a linear correction.

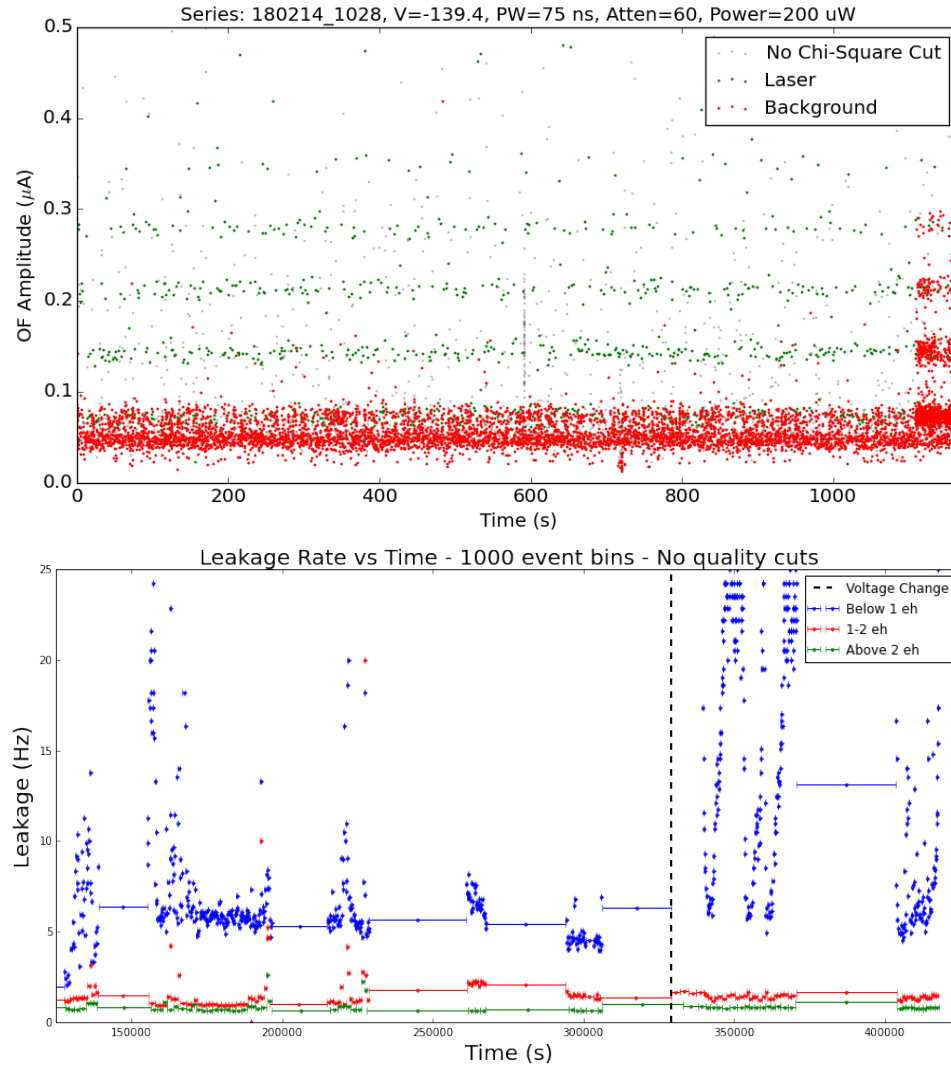


Figure 6.3: Top: All events during a series separated by whether they were coincident with the laser (green), anti-coincident (red), or failed the chi-square cut (grey). We can see that during this run, a quantized leakage appears towards the end of the run. After pre-biasing and starting the next series, this leakage is gone, but seems to re-occur randomly during the subsequent series. Bottom: Event rate as a function of time for different energy ranges. The blue points trace noise and IR leakage triggers, which can be seen to explode when the temperature gets very low due to the increased signal to noise from the lower crystal temperature.

Some other issues during data taking included the channels being swapped accidentally on the 11th, which required reprocessing all of the data with a flag to swap the channels after this date, and some gain variation during the first 20 hours in one of the SRS amplifiers that made it necessary to re-scale one of the channels during that time even before calibration. Due to these issues, we ended up abandoning half of the full exposure, opting to include only the negative bias data and excluding the initial 20 hours of data due to this gain change. This left us with about 24 hours of initial raw exposure before further cuts.

6.1.2 Pulse Processing

As discussed in the previous section, during data taking an initial processing was done on each series by using a single trace to generate both noise and templates, and then applying a time-fixed and time-shifting optimal filter to all events in the series. The upper plot in figure 6.3 is an example of a data quality plot generated with this method, and example pulses and noise PSDs for two series at different times during the run can be seen in figure 6.4.

After the initial processing during the science run, we re-processed the data a number of times with improvements to the OF algorithms and additional quality variables. The largest differences had to do with the part of the trace used for fitting and the templates used. We constructed a template from the automatically generated templates made during the run by averaging all of the templates from good data series and applying an additional low-pass filter above 100 kHz as shown in figure 6.5. In addition, the initial processing applied the OF to the second half of the trace, but for the time-shifting OF, this cut off some of the lower energy pulses, which were not triggered on until later in the pulse, so we used the same template but used the trace offset by an additional 500 microseconds to give more pre-trace time for the time-shifting OF. This improved our reconstruction efficiency for low amplitude pulses, as will be discussed in the next section.

Finally, we computed a handful of different goodness of fit metrics for each pulse, including both frequency and time-domain chi-squares for the fixed and time-shifting OF. We found that as a general property, the OF de-weights high frequency information in Fourier space, so that while for a given trace it is still able to find the optimum offset and amplitude, it is not sensitive to pileup. In contrast, the time-domain fit is not filtered and is highly sensitive to pileup, and was a much better discriminator of pileup even though it was more sensitive to low-frequency noise. The chi-square cuts will be discussed in more detail in the next section, but for the initial analysis a very rough cut of ~ 10 was put on the time-domain chi-square to remove the worst pulses, so we were only looking at good pulses with some small fraction of bad events to be removed later in the analysis.

We also implemented a pileup optimum filter, described in appendix E which was used to remove IR pileup when that was our dominant background. After the improved IR filtering, we found little evidence that the pileup OF was finding any real pulses, and it is much more computationally intensive to compute than either of the single-pulse optimum filters due to the fact that you can't

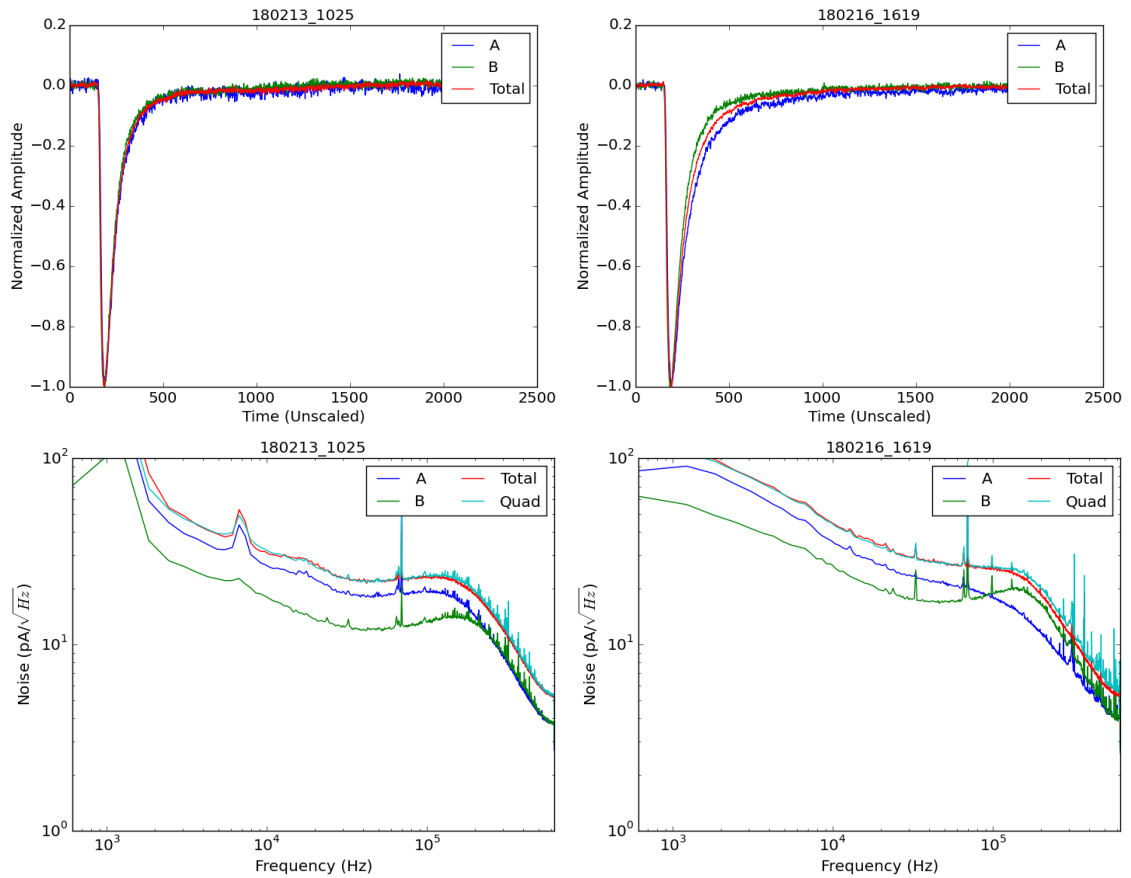


Figure 6.4: Top: Templates automatically generated for two example series from the second half of the traces in the series, after some automatic cuts have pre-selected events most characteristic of the average pulse in the series. Bottom: PSDs generated from the first half of the traces for each series after similar automated cuts are made. Note that the templates are very similar but the noise is very different; for this reason we used the PSDs for each series but a single template for all data with additional post-processing.




Figure 6.5: Left: Filtered pulse template (black) compared to the unfiltered pulse template (green). Right: Two of the individual series pulse templates (green and brown) compared to the averaged template (black) and the averaged template after filtering (blue).

use the same Fourier-transform trick as you can for the one-pulse OF. In order to apply it to the entire science dataset we had to down-sample in time by a factor of 10, and when the initial studies found no significant pileup (that wasn't also removed by a time-domain chi-square cut) we decided to leave it out of the later rounds of processing. This is an area where this analysis can be improved as more robust analysis tools are developed for SuperCDMS SNOLAB.

6.1.3 Calibration

Before the energy calibration, we need to find the relative calibration between two channels. The assumption here is that there is a variation in the amount of absorbed by each channel as a function of event position. If this is only a difference in the fraction of energy absorbed, then we can solve for the relative power to current conversions and make the total of the energy in each channel position independent. If there are other effects at play (loss of energy in the crystal as a function of position) there will still be residual impacts on overall resolution, but this will be seen as a nonlinearity in this correction. This could be due to reduced efficiency at high radius due to phonon down-conversion, or changes in dI/dP due to different passive resistance or bias point. Ideally, all channel properties would be the same, but in practice it's hard to achieve consistency to better than 10%.

The system of equations then for the total energy is

$$I_{tot} = c_A I_A + c_B I_B = E_{abs} \frac{dI}{dE} \quad (6.1)$$




Figure 6.6: Relative calibration using a time series with a large population of quantized leakage with a different position dependence than the laser data. Data along the lines, which have a slope of -1, indicate the lines of equal reconstructed energy. This relative calibration required channel A to be scaled by 1.31 relative to the initial value.

If we select a fixed energy, then we should have

$$c_B I_B = \epsilon E - c_A I_A = B - c_A I_A \quad (6.2)$$

$$I_B = B' - \frac{c_A}{c_B} I_A \quad (6.3)$$

where

$$B' = \frac{\epsilon E}{c_B} \quad (6.4)$$

is set by the known energy of the line. This also shows us that the value of c_B is degenerate with ϵ and can just be absorbed into the measured dI/dE for the weighted channel sum. We will set this to one, and get the modified equations

$$I_{tot} = c_A I_A + I_B = E_{abs} \frac{dI}{dE} \quad (6.5)$$

and

$$I_B = B' - c_A I_A \quad (6.6)$$

so we can determine the relative channel weighting by a linear fit to I_B versus I_A for a given known input energy. This method avoids having to assume a distribution of events, just that there is some position dependence in the signal.

The application of this method is shown in figure [6.6](#), where the data from the time series in




Figure 6.7: Calibration to remove the dependence of reconstructed amplitude on temperature. The first three peaks are fit to lines to get a correction factor as a function of temperature. See text for more details. It’s significant that when the fridge is colder, the current separation between the peaks increases, while the noise remains about the same. This suggests that the low efficiency of the detector is more a temperature effect than an intrinsic limitation. We do not naively expect such linearity in thermometer resistance versus phonon gain; it is possible that if both resistance and thermal conductance are phonon driven, they both obey the same temperature dependence of T^5 , making them linearly proportional.

figure 6.3 has been used as a proxy for position dependence in the crystal. I found that a 33% correction to channel A allowed for the most position-independent energy reconstruction. After determining the relative weighting of the channels, the data was all re-processed to take this correction into account so that all of the summed quantities would reflect the correct channel weighting.

Next, the temperature correction had to be applied given the high correlation between pulse amplitude and mixing chamber thermometer resistance. The resulting correction, done by Andrew Scarff, can be seen in figure 6.7. This procedure fit a line to each of the laser peaks for the calibration data across the entire science run as a function of thermometer resistance, in order to remove the temperature correlation. The idea here is that what is changing is the dI/dP fraction as a function of operating temperature (or equivalently as a function of thermometer resistance), so that

$$y_{amps} = \frac{dI}{dP}(T)n_{eh} \approx \left[\frac{dI}{dP}(R_0) + A(R - R_0) \right] \quad (6.7)$$

We should find that by measuring the slope of current versus thermometer resistance (assuming both R_{temp} and dI/dP have the same temperature dependence) for a few electron-hole pair lines,




Figure 6.8: Left: Example calibration spectrum with lines drawn at the peak points found using the continuous wavelet transform method. This is a fairly standard peak-finding algorithm and requires some tuning for a given case in order to reduce sensitivity to noise. By properly tuning this algorithm I was able to compute a relative scale factor for every series in the data. Right: Laser calibration data with different laser points showing how calibration data points can be found up to the limit of around $1.2 \mu\text{A}$, when the non-linearity causes the detector to lose single electron-hole pair resolution. The same method was applied to this data to get the points shown in figure 5.25 to set the absolute energy scale.

we should consistently find that the slope of the line divided by the number of electron hole pairs is a constant. The offset here just allows us to normalize all of the data to the same current as at 35 mK. This can be generalized to a polynomial as well, as long as you're sufficiently below saturation, because dI/dP is only a function of temperature, not of pulse height, until you have local saturation effects.

Finally, after these corrections, we can perform a non-linear current to energy calibration using this same laser calibration data. This was the calibration shown in the previous chapter, in figure 5.25. The method for this calibration can be seen in figure 6.8. A continuous wavelet transform was used to algorithmically find the position of the first 10 e-h pair peaks in current, and these data were then fit to a quadratic form, as shown in figure 5.25 to produce a linear energy scale. After this calibration was done to all of the data at once, we could still see residual variation as a function of time in the laser calibration data, so the continuous wavelet transform method was used to calibrate up to 6 keV for all of the data, and then each series was individually corrected by a linear scale factor to produce the final, calibrated dataset. All of the science calibration before and after this last calibration pass can be seen in figure 6.9.

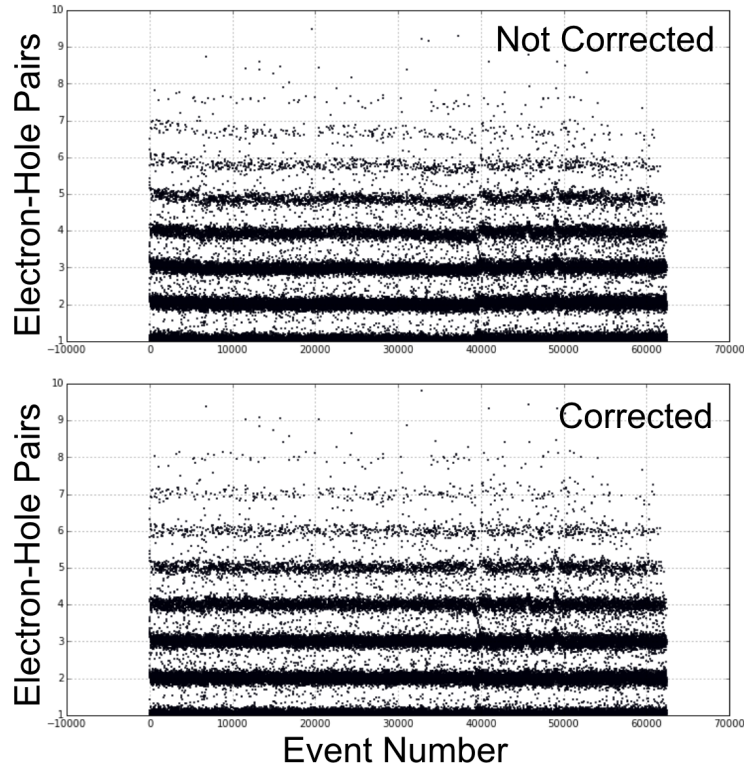


Figure 6.9: Top: Events after temperature calibration before time-dependent calibration. Bottom: All events after time-dependent calibration, showing much more consistent calibration across the entire dataset.

6.2 Quality Cuts

In this section I will discuss the livetime calculations and the cuts applied to the data, and summarize by showing the final spectrum and efficiency model used in subsequent limit calculations. This section is much shorter than for most analyses as we decided to be as conservative as possible for this first data run, given that the detector and calibration schemes are new, and that there were so many issues while running the detector, reflected by the amount of time it took to reach our target exposure.

6.2.1 Livetime Calculation

For our analysis, we'll need to compute livetime both for livetime cuts (which just remove blocks of livetime) as well as quality cuts, which remove individual events. We can estimate the impact on livetime of this second class of events, as well as the total livetime of the final dataset, using the laser data, as described below.




Figure 6.10: Exposure as a function of wall-time, showing the large blocks during which we were debugging the setup or running calibrations. Both the neutralization at 70 hours and the polarity flip can be seen to occur shortly before resuming data taking.

First we consider livetime cuts. For each series, we bin the events into time windows of Δt , and make a cut decision for each window's average trigger rate compared to the background series. Suppose I remove n of N time periods from my data; the resulting livetime for that series is then

$$T_{live} = [N - n] \Delta t$$

This is a very straightforward correction, but to make it we need to be able to calculate Δt in an unbiased way.

Calculating the livetime of a given selection of data is actually fairly straightforward to calculate directly from the data, given that the laser was firing at 1 Hz for the duration of the science run. Given a known rate Γ_{laser} and mean photon number λ , we have the total exposure

$$T_{exposure} = \frac{N_{laser}}{\Gamma_{laser}} [1 - e^{-\lambda}]^{-1}$$

where N_{laser} is the number of laser events retained in the data, and the last term is a correction given that we will not trigger on the 0 photon events. We estimated λ in turn using the measured mean photon number λ_m for each dataset, calculating λ as

$$\lambda = W(-\lambda_m e^{-\lambda_m}) + \lambda_m \quad (6.8)$$

where W is the Lambert W function.

Given that we have not yet applied cuts on the pulse amplitude, this is an unbiased livetime estimate as we're just counting the number of laser TTL pulses which generated a threshold trigger, and our threshold was around 0.5 e-h pairs, well below the first peak. We also spot-checked this

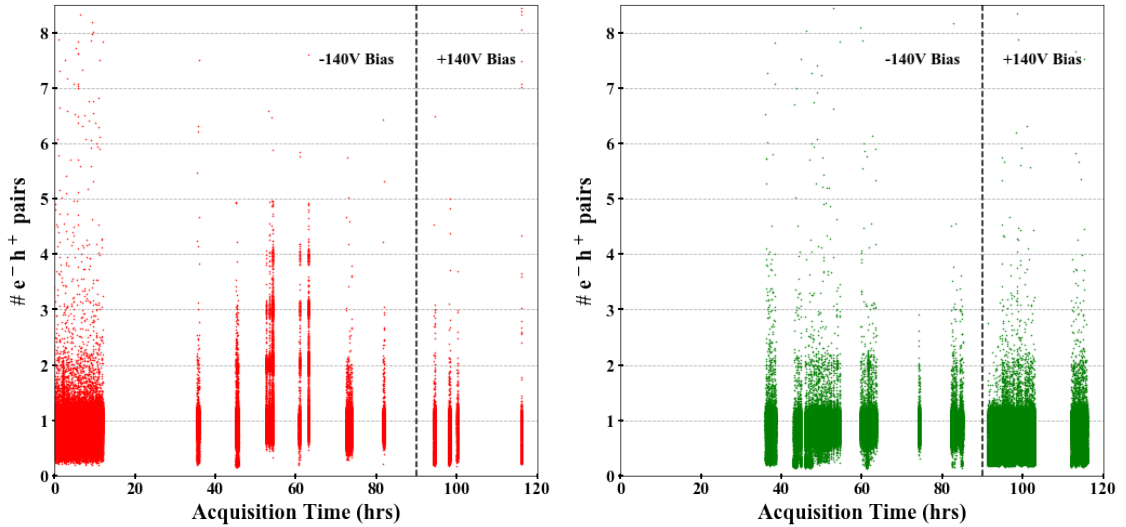


Figure 6.11: Left: Series removed by initial livetime cuts based on equipment or stability issues, or sustained periods of excessively high trigger rate. Right: Data retained after these cuts, including the +140V data which is not included as part of the science exposure but is used to study leakage current during the run. See text for more details.

livetime by adding up the retained time in a few series and comparing to this value to verify that these assumptions held. The total exposure before cuts can be seen in figure 6.10 as a function of wall-time, broken down by crystal bias and including the total exposure without crystal bias. The expected 48 hours of data turned out to be an over-estimate due to a mis-estimation of the deadtime associated with higher trigger rate data which occurred towards the end of the run; we estimated about 16 hours of positive bias data, but only obtained about 9.

6.2.2 Livetime Cuts

I mentioned in the previous section that we had a lot of data taking issues during parts of the run, and that we only include the negative bias data as part of the science exposure. Figure 6.11 shows the data retained versus cut, which includes data taken in the first day with pre-amp gain issues, data which included the spontaneously leakage events¹, periods with sustained high trigger rates, and series during which the temperature was changing. These were fairly severe cuts but were always done by series number or in larger blocks of time so as to avoid biased selection of data.

We did try to retain the data during the first day, which required re-calibrating the gain of one channel and swapping the channels relative to one another, as well as a separate energy calibration scheme, but the calibration data was never near as well behaved as in the rest of the data, leading us to believe that the gain variation had non-trivial time dependence. We believe something was

¹Bias referred to these as burp events, which made me laugh every single time we talked about this data.

malfunctioning in one of the SRS amplifiers, and we decided that the data was too unstable for us to be able to calibrate. Given that we are background limited in this analysis after only a few hours, it does not impact the limit substantially to remove this block of time.

During data taking, we noticed that there was a significant amount of time-variation in the trigger rate due both to changing threshold as well as true spikes in event rate associated with higher than normal backgrounds. We developed a cut to remove time periods that deviated by more than 7 sigma from the expected trigger rate within a given run to get rid of these bad time periods, based on both total events and only those events that pass a rough chi-square cut and are above threshold. The framework for this cut is described below.

Suppose we have a nominal trigger rate Γ . In some time window t , we can expect a mean of $\lambda = \Gamma t$ events, and the distribution of these events should be roughly Gaussian with standard deviation $\sigma = \sqrt{\lambda} = \sqrt{\Gamma t}$. Suppose now that we want to reject time periods which are an m -sigma fluctuation from our expected trigger rate, and we want this to roughly correspond to rejecting trigger rates n times the nominal rate. This corresponds to the condition

$$n\Gamma t = n\lambda = m\sigma = m\sqrt{\Gamma t}$$

This gives us a condition on the amount of time we should bin events into so that we have a high-confidence rejection at a given level, according to:

$$t = \frac{m^2}{n^2\Gamma}$$

For example, suppose we have a nominal trigger rate of 5 Hz, and we want to reject instantaneous rates above 10Hz at the 10-sigma. This means that $n = 2$, $m = 10$, and $\Gamma = 5$, such that

$$t = \frac{100}{4 * 5} = 5s$$

So if we bin events every 5 seconds and apply a cut at 10 Hz, we will have 10-sigma rejection of time-series that do not match our expectation.

For these data we binned data into 5 minute time bins, and applied two ‘burst’ cuts independently of one another. The first burst cut rejected time periods based on trigger rate alone, while the ‘leakage’ burst cut rejected these periods based on events above a single electron-hole pair. The majority of time-periods removed by these cuts were redundant, and the overall loss from both of these cuts was less than 10%.

Table [6.1](#) summarizes the impact of each cut on the total livetime. In addition to the cuts discussed here, the laser events do remove a small fraction of the total livetime, less than a percent, due to the fact that we discard the laser triggered time periods from the science exposure. The final exposure before event cuts was about 19.5 hours or 17.5 gram-hours.

Cut	Passing Livetime	Passage Fraction
No Laser	36.4 hours	$\sim 100\%$
Time Period	21.7 hours	59.6%
Trigger Rate	33.5 hours	92.2%
Leakage Rate	35.4 hours	97.4%
Livetime	19.5 hours	53.5%

Table 6.1: Livetime passage efficiency for the selection of good time ranges, trigger rate and leakage rate cuts as described in the text.




Figure 6.12: Comparison of reconstructed pulse height for an event without a pulse based on an initial noise variance (the blue curve) as we increase the number of search windows. When we get to 20 search windows, we see that the probability of estimating 0 energy for our truly 0 energy event is 0; this is therefore a biased measurement. The noise for the data in this science run is consistent with 100 independent search windows; this is a weighting between the number of time bins in the template and the length of the trace. X-axis is OF amplitude in μA .

6.2.3 Trigger Time Cut

The science data was all drawn from the optimal filter with fitted time offset, which provided a more linear initial detector response but also allows for some additional artifacts to leak into the spectra. For low-energy pulses, the trigger time actually occurs after the start of the pulse in a fairly well-correlated manner, but by searching the entire trace, the probability of finding a statistically significant noise fluctuation shaped like a pulse becomes large, such that the noise blob for the time-delayed OF has no events with zero amplitude. Figure 6.12 illustrates this effect.

Luckily, we have multiple ways to combat this. We can either estimate the new noiseblob with the shifting window and make a threshold cut, or we can use the reconstructed time as a cut to remove random noise events inconsistent with the trigger time. We opted for this second method in order to maintain as low an effective trigger threshold as possible. We used the laser data (which is a population of events known to be real) as a proxy for good events, and set the cut to have high efficiency at 1 electron-hole pair. The power of this cut is that events should be within the time resolution of their true location, but traces without real pulses will have uniformly distributed start




Figure 6.13: Pulse time for the laser events as a function of reconstructed pulse energy, showing the events which pass the pulse timing cut in green. See text for more details.

times, so the rejection factor is roughly the timing resolution divided by the length of the trace. The size of the noise blob increases as the square root of 'sampling windows' (the number of sampling windows is proportional to trace length), so the timing cut fares better for longer traces while the threshold cut fares worse.

The result of this cut can be seen in figure [6.13](#). Note that the pulse start time gets closer to zero for larger pulses, and that the variance in pulse start time decreases with increased pulse size; both of these effects are to be expected. Consider, for example, a very small pulse. If it experiences a downwards noise fluctuation, the trigger will occur later, after a subsequent upwards noise fluctuation resulting in a net negative template offset relative to trigger time. In addition, we expect timing resolution to be inversely proportional to pulse energy, so the better timing resolution matches our expectation.

Based on these arguments, the natural question is why we didn't choose a form for the timing cut which more closely follows the laser data. The reason for this was because we do not have an estimate for how position dependence factors into reconstruction time. The laser events all occur at a very small spot on the rear face of the detector; if position dependence also manifests as a time offset, being overly aggressive with this cut would reduce our DM signal efficiency and bias our result artificially low. In the future, a position independent calibration (such as low energy Compton recoil deposits) should allow a more stringent time cut where we can be sure that we're not reducing our sensitivity to different parts of the fiducial volume.

6.2.4 Chi-Square Cut

The time-domain χ^2 for the optimum filter with time offset had a very rich structure as a function of energy, and we initially intended to do a very complex cut involving both time and Fourier domain chi-square for both channels as well as the sum. This was greatly hampered by noise variation within

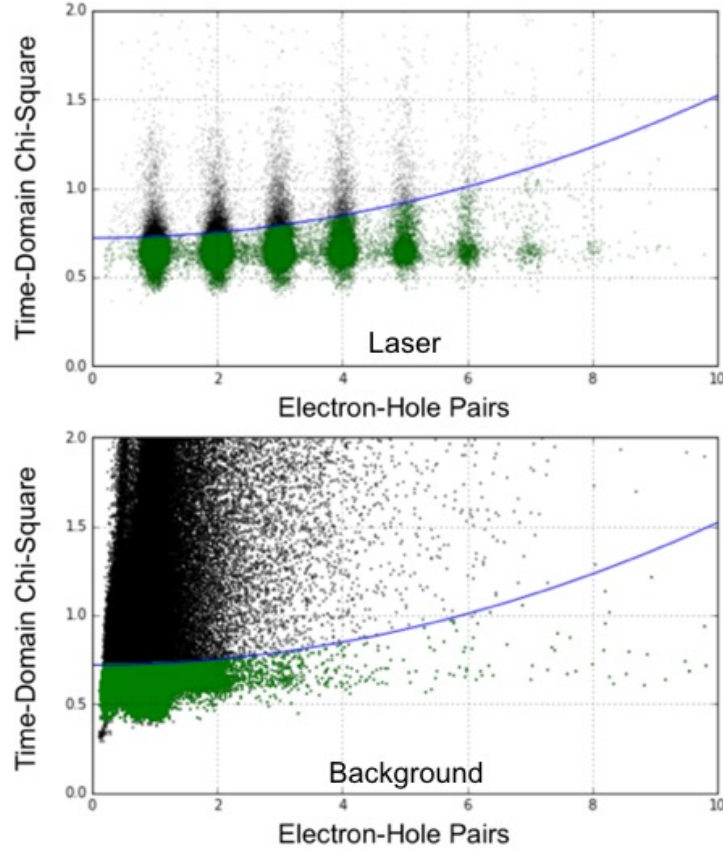


Figure 6.14: Top: χ^2 versus energy plane with the cut-line drawn for the laser data. Bottom: Same plane but for the events without a laser coincidence. The χ^2 is below one due to an error in the scale factor by a power of 2, as we only sum over half of the trace but use the variance for the full trace. See text for more details.

a series. To understand why this would impact the χ^2 , consider the formula

$$\chi^2 = \frac{\sum_t (v(t) - as(t - t_0))^2}{\sigma^2} \quad (6.9)$$

where σ^2 is the variance of a trace when there is no residual; for a trace without a pulse, the mean χ^2 for non-stationary noise which is uncorrelated with the signal is 1. If the noise changes during the run (as a result of new lines showing up or temperature changing), then what we see is that the expected value of σ^2 changes, causing a time-dependence in the mean χ^2 and an increased χ^2 variance.

For stable series, the χ^2 distributions were well behaved, but over the course of the run there was enough temperature and noise variation that the chi-square variance increased and washed out some of the structure. As a result, a pulse that would be a bad chi-square in one series is within

the good range in another, and we lose some discrimination power. We decided to simply set the low-energy chi-square cut at the point where the laser and background distributions diverged, and fit a quadratic cut-shape to maintain the same efficiency above threshold as a function of pulse energy. We toyed with the idea of doing a time-dependent chi-square cut, but initial studies didn't indicate that it affected the spectrum appreciably, and again we were cautious not to bias the result or over-complicate the first analysis with a new detector.

The resulting chi-square cut is shown in figure [6.14](#) for both the laser and physics search events with livetime and timing cuts included. The smearing of the laser data in chi-square is apparent from the laser data, and there is an upward trend as a function of energy indicative of the fact that the pulse shape is changing slightly due to local saturation as energy increases. In contrast, we don't see any meaningful structure in the background data aside from a population of events which is a good fit out to high energy, and another population with small amplitude and rapidly increasing chi-square. The majority of these events seem to be RF-induced trains of pulses as well as tails of high energy events. We'll discuss this a bit more in the next section.

6.2.5 Good Noise Cut

The final cut applied was the good noise cut, based on the mean and variance of the pre-trigger portion of the trace. Given that this data was AC-coupled, a non-zero mean is indicative of an event on a high energy event tail, which would be mis-reconstructed. A large increase in variance is indicative of a bad noise environment, allowing for the possibility that a larger than normal fluctuation may produce a false event. This cut was set to include 98% of the laser events regardless of their energy, and in the end had a negligible impact on the final spectra. It was included in the final dataset for robustness but was almost entirely redundant with the chi-square and pulse time cuts.

6.2.6 Final Spectrum and Efficiency

The laser and background data after livetime cuts, applying event cuts sequentially, can be seen in figure [6.15](#). The laser data does not change appreciably aside from the threshold being increased by the pulse time cut, but otherwise all cuts are flat as a function of energy by design. The background spectrum initially contained a large population of events above 1 e-h pair that were dominated by bad events, and after cuts has a shape which is less smooth, and has a much lower rate above 3 e-h pairs.

There were some significant surprises when we first obtained this spectrum, but all follow-up investigations did not have any significant impact on the shape, only reducing the overall rate. We explored all possible variables to see if the events above a single electron-hole pair were inconsistent with laser data. We did not find any significant differences in pulse shape, arrival, time, goodness of fit, time-correlations, or any other potential factors we could think of. In the next section we




Figure 6.15: Top: Laser spectrum after cuts are applied cumulatively, only noticeable shape change in spectrum is the noise below 1 e-h pair being rejected by the χ^2 and arrival time cuts. Middle: Science spectrum after cuts applied cumulatively. The χ^2 cut is responsible for the majority of the reduction in differential event rate, with threshold differences seen below 1 e-h pair. Bottom: Efficiency estimated by dividing the number of observed events in each electron-hole pair half bin by the initial amount. We see that the χ^2 cut brings the overall efficiency down to about 85%, but it is tuned to have a flat efficiency across our energy range. The arrival time cut sets the effective analysis threshold. The good noise cut removes a relatively flat 2–3% of events above threshold.




Figure 6.16: Data from the science exposure compared to the best-fit impact ionization and IR background model in units of electron-hole pairs. This shows that we can not model the background above 2 e-h pairs as impact ionization related unless we use a non-physical correlation model.

will explore the origin of the events in three distinct regions of the spectrum, which provide future avenues for improving the sensitivity of our next science run.

6.3 Background Model

We expected that our dataset would look like the IR background estimated from randoms plus some Compton background, but we ended up with something much more continuous and were background limited earlier than expected. We can still explain most of this background, but the regions from 2 to 5 electron-hole pairs, which was the biggest surprise, is still somewhat uncertain. In order to add confidence to our data, we constructed a basic background model to show that we understand the detector response near threshold and well above threshold. In this section I'll discuss the known backgrounds and also present our theory for the residual unknown backgrounds.

6.3.1 Impact Ionization and IR

In section [5.4.5](#) I showed that for data taken before this science run, we estimated that we had 2–4% impact ionization probability for events originating in the bulk. We know that the randoms below 1 e-h pair are due to a flat bulk background (which we assume is due to IR, but that's not guaranteed, so we'll just call it a volume leakage) at 10 Hz. If we model the background as a 10 Hz volume leakage with 3% impact ionization, we find we can match the background almost exactly up to 2 e-h pairs, as shown in figure [6.16](#).

This model shows that we can't explain the background above ~ 2.1 e-h pairs with a simple impact ionization model, and given that it's not quantized, it's unlikely to be a known background. If it's




Figure 6.17: Time dependence of the event rate in the bulk leakage region (0.8–2.5 e-h pairs) as well as events above this energy after quality and livetime cuts. The bulk leakage rate spiked after neutralization and seems to decay with a time constant of about a day. It doesn’t seem to be affected by reversing the crystal bias.

above the Si bandgap, it should largely be quantized, and if it’s below the bandgap it should already contribute to the bulk rate measurement, and coincidences are already included in our model. The model also includes a surface leakage component which is below the IR background at 1 e-h pair but peaks up above the impact ionization background at 2 e-h pairs. This contributes slightly more impact ionization but is an order of magnitude shy in explaining the origin of this continuous higher energy background.

Finally, we did see significant time-dependence in the leakage rate, as shown in figure [6.17](#), but only in the data consistent with the bulk leakage. We expected that neutralization would improve our leakage, but it made it demonstrably worse. We surmised that the leakage may have been due to charge buildup at the detector surface, so we flipped the polarity, but that didn’t appreciably affect the leakage rate. It’s likely that allowing the detector to run allows the shallow traps to slowly decay, and as long as we’re at high voltage the crystal remains neutralized. By re-neutralizing we re-filled those traps, resetting the leakage clock. My current theory is that the steady-state leakage is set by the equilibrium sub-gap state population freed by IR and re-filled by a small fraction of charges which become trapped in the bulk. This will be tested in upcoming calibrations by attempting to raise the crystal temperature to empty sub-gap states while maintaining overall neutralization. It’s likely though that the leakage rate, if initially low, will rise with time if I’m correct.

6.3.2 Compton Recoils

Let’s consider the other end of the energy spectrum, Compton scattering events. For a very high energy gamma, given that the crystal is very small, there’s a good probability of the gamma Compton




Figure 6.18: Long exposures taken at the end of Run 159. Left: background only, showing that our background is consistent with IR, impact ionization, and Comptons. Right: Adding an Am-241 source leaves the lowest energies unchanged but elevates the flat Compton spectrum. This data had much lower statistics due to the high trigger rate and lower lifetime which was obtained as a result of running for a fixed number of events, rather than a fixed amount of time.

scattering from an electron in the detector and depositing a small fraction of its energy before escaping from the detector. Because we are well below any escape peak energy, we can model the Compton background as flat in energy, and assuming a simple model, we can compute some rough Compton rates to see if they agree with our data.

Let's scale from another known test facility. At UCB, without a lead shield, they measure an event rate of ~ 0.06 Hz/gram over all energies. If we assume that's spread evenly between 0-100 keV, then we have

$$6 * 10^{-7} \text{ Hz/g/eV} = 2 \mu\text{Hz/g/neh} = 0.2\text{events/g/day} \quad (6.10)$$

So if for 0.5 gram-days, we only expect one Compton event during the exposure. Looking at the spectrum, we have around 5 events/eh, or about 50 events in our energy range, which is almost 2 orders of magnitude higher than expected. We do have significantly less fridge mass surrounding the detector, and likely a much more radioactive internal environment, so this is not entirely out of the question. SLAC has a larger background rate, and we could argue that since the rate outside fridge is 180 Hz/kg, without any shielding we could have 3 events, so we're only about a factor of 20 higher than expected. The fact that the spectrum flat in energy suggests these events are consistent with a flat Compton background.

The one measurement we have to test this hypothesis was from a previous run, where we tried to use Am-241 as a calibration source to measure the detector response to events isotropically distributed in the crystal. The spectrum with and without the Am-241 source, located outside the cryostat, can be seen in figure 6.18. Both exposures have the same low-energy background, with noticeably smaller leakage in the 2-eh pair bin relative to the lower energy bin due to stabler operating conditions but also due to the higher overall IR rate. They also have the same population

of events which generate a continuous tail out to the edge of the energy window. The difference is that with the Am-241, the flat background above 3 e-h pairs is much larger, as is the population of high energy events (mostly 60 keV from the source).

This calibration suffered from a high multiple trigger rate, resulting in lost livetime, but when it is repeated in the near future it should be a very rich dataset for understanding gamma backgrounds in this detector. One large uncertainty in our background modeling that we will have until we can get this Am-241 calibration working is the relative importance of coherent nuclear scattering of gammas, given that the energy scale is somewhat uncertain but the ionization yield is entirely unknown [88]. The next run should take advantage of this calibration to study position dependence in the detector, which was a large uncertainty for the present analysis and prevented us from doing any fiducial cuts that would have reduced residual high-radius surface leakage.

6.3.3 Environmental Backgrounds

I've just shown that we can understand our background below 2 e-h pairs and above 5 e-h pairs; what about the intermediate region? We are able to make a contrived charge leakage model to explain the background, but past experience and early data runs give us another possible explanation for these events, which is reinforced by analyzing the goodness of fit distributions near our χ^2 cut. Figure 6.19 shows events taken from different regions of the energy versus χ^2 plane after calibration. We see that the region trending towards low energy and good fit has a large population of both high-energy event tails and trains of pulses, which are not physics and are likely the result of RF power down the TES bias line.

The second class of events is more of a concern to us, as they will not be rejected easily by the χ^2 cut in the frequency domain. We can reject all of these events with 2 or more pulses with a time-domain χ^2 cut, which made our pileup OF not necessary as discussed earlier, but if these events occur as a single pulse with some probability we will have no way to reject them. To see whether some background could be leaking through our χ^2 cut, we looked at the χ^2 distributions as a function of energy, shown in figure 6.20.

This figure seems to suggest that while the backgrounds below 2.1 e-h pairs have a similar distribution to the laser events, the majority of events above 2.5 e-h pairs seem to be drawn from a distribution with χ^2 that has a different shape than the laser events, which we know are good. There are few enough of these events which pass all of our cuts (a few hundred) that we were able to look at them individually; the events that pass our cuts look the same as laser events by eye, and thus we could not reject them. It seems likely however that the majority come from the same source that produces this population of events. If this is RF driven, then either better RF shielding or the addition of a second detector will allow us to remove or veto these events respectively in a future run.

We do have precedent for these events from Soudan, but the change in detector optimization has

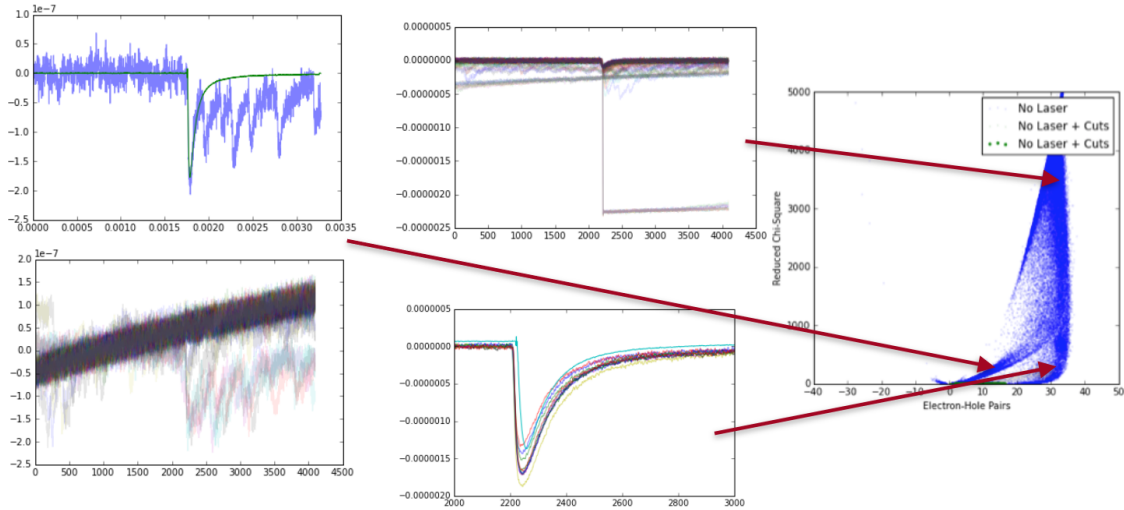


Figure 6.19: Bad events as determined by χ^2 as a function of energy with examples from a few sections of parameter space. Top left: trains of pulses, each of which are consistent with our pulse template, occurring at regular intervals. Bottom left: high-energy pulse tails, and additional pulse trains. Top right: saturated high-energy pulses. Bottom right: good events with varying degrees of local saturation.

changed how they appear in our data. The ‘glitch’ events seen in the Soudan data looked like TES green’s functions, but we were able to reject them based on pulse shape due to the fact that our physics pulses had falltimes driven by phonon physics rather than TES physics. In these very thin detectors, the pulse looks identical to a TES green’s function, meaning any external power, if short, will look identical to a good event. This is perhaps a good reason to allow the pulse to be longer than the TES fall-time, to have additional rejection power for these events. This wouldn’t be a problem, at least in theory, if we had multiple detectors or an additional power-sensitive channel in the cryostat, as we’d have another way to veto the events as long as they had the same coupling to this power noise.

6.4 New Physics Constraints

Given the uncertainty in our background models discussed in the previous section, and the fact that the background is non-quantized (meaning we have little discovery potential for this run) we set limits on electron-recoil dark matter and dark photons using an optimal interval limit setting technique [109, 110] with a slight modification. Both of our signal models are quantized, so we only consider background events within 2σ of a quantization peak, as shown in figure 6.21. Otherwise the limit setting is the standard for a basic dark matter exclusion curve, though in this case we are background limited and thus our limits would improve with a likelihood analysis. In this section I’ll

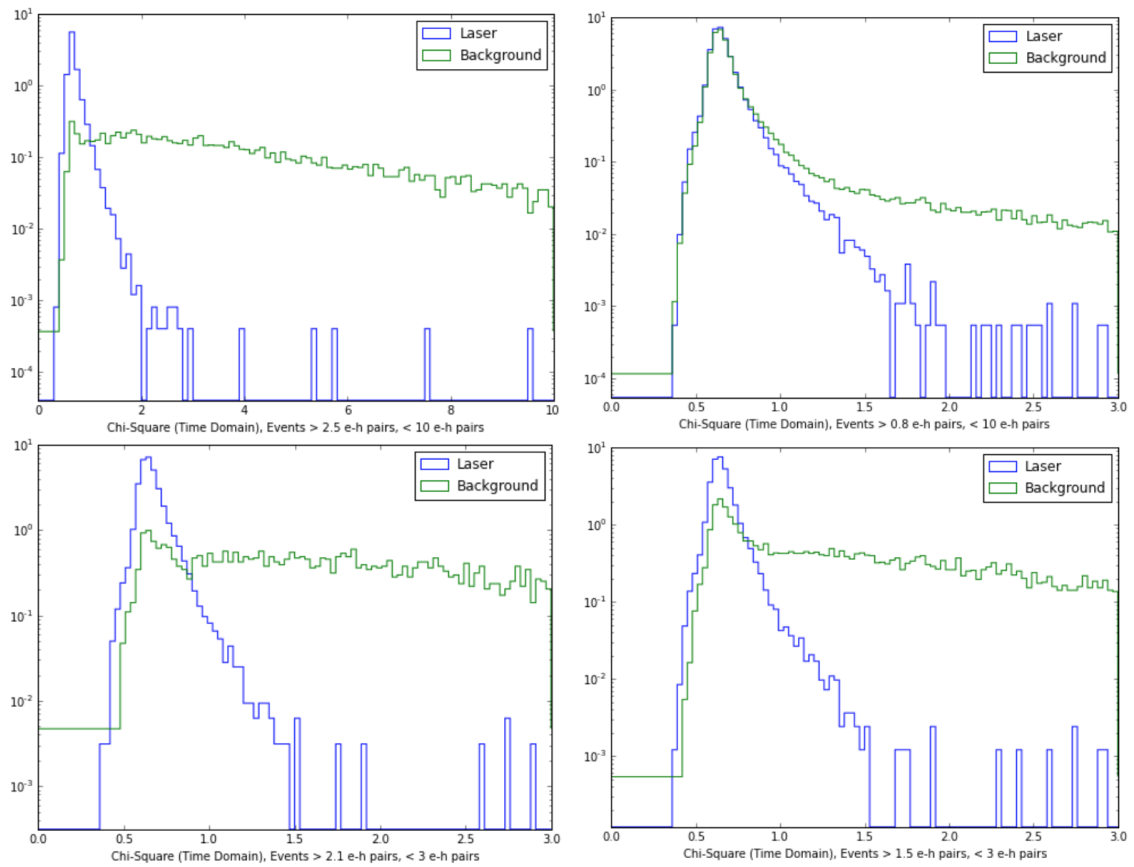


Figure 6.20: χ^2 distributions for events with different energy cuts separated into foreground and background distributions. The top-right distribution shows that the majority of low-energy events are consistent with laser events, but the top left and bottom panels show that above the known leakage-driven background there is a significant population of events likely leaking into the signal region. See text for more details.




Figure 6.21: Regions of the background considered by the optimal interval are within 2σ of the quantized signal peaks in order to reject parts of the background inconsistent with signal.

briefly review the signal models before presenting and discussing the limits set with this data.

6.4.1 Ionization Model

In order to project an event of known energy into our measured signal space, we adopted an ionization production model that is consistent with experimental measurements [108, 29, 107] and has the following mean n_{eh} :

$$\langle n_{eh}(E_\gamma) \rangle = \begin{cases} 0 & E_\gamma < E_{gap} \\ 1 & E_{gap} < E_\gamma < \epsilon_{eh} \\ E_\gamma/\epsilon_{eh} & \epsilon_{eh} < E_\gamma \end{cases} \quad (6.11)$$

where $E_{gap} = 1.12$ eV and $\epsilon_{eh} = 3.8$ eV [105]. This is a non-trivial uncertainty in our analysis, as there are no recent measurements of the ionization yield in Si for crystals in our operating conditions. This is the model which is being used by all similar analyses, so we should have a consistent systematic as a result, and we plan to measure this in the near future with these devices by using light sources of variable energy.

The probability distributions in the first two cases are delta functions, necessary in order to conserve energy. In the third case, we generated discrete distributions with an arbitrary Fano factor, F , by interpolating between binomial distributions with the same $\langle n_{eh} \rangle$, but different integer number of trials, allowing all energetically accessible numbers of e-h pairs. The Fano factor is defined as

$$F = \frac{\sigma^2}{\mu} \quad (6.12)$$

where $\mu = \langle n_{eh} \rangle$ is the mean of the distribution, and σ^2 is the variance. A completely uncorrelated (poisson) process has a Fano factor of 1, but in most radiation detectors we find Fano factors on the order of 0.1-0.2 due to the fact that large deviations from the mean are kinematically suppressed.

We generated probability distributions for a given mean and Fano factor using the binominal distribution with n trials of probability p . The binomial distribution has the its variance σ^2 and mean μ obey the relations

$$\mu = np \quad (6.13)$$

$$\sigma^2 = np(1-p) = \mu(1-p) \quad (6.14)$$

which allows us to calculate the n and p values from the Fano factor and mean number of electron-hole pairs as

$$F = \frac{\sigma^2}{\mu} = (1-p) \rightarrow p = 1-F \quad (6.15)$$

$$n = \frac{\mu}{p} = \frac{\mu}{1-F} \quad (6.16)$$

The caveat to these equations is that the binomial distribution is quantized, and thus n is an integer; this means that we have to interpolate between the distributions for the integers directly above and below the fractional mean value given by our mean and Fano factor combination. This is done according to the following procedure:

$$n_l(\mu, F) = \text{floor} \left(\frac{\mu}{1-F} \right) \quad (6.17)$$

$$n_h(\mu, F) = \text{ceil} \left(\frac{\mu}{1-F} \right) \quad (6.18)$$

$$F_l(\mu, F) = \mu/n_l(\mu, F) \quad (6.19)$$

$$F_h(\mu, F) = \mu/n_h(\mu, F) \quad (6.20)$$

$$\Delta F(\mu, F) = \frac{F - F_l(\mu, F)}{F_h(\mu, F) - F_l(\mu, F)} \quad (6.21)$$

$$P_l(x|\mu, F) = \text{Binomial}(x|n_l(\mu, F), 1 - F_l(\mu, F)) \quad (6.22)$$

$$P_h(x|\mu, F) = \text{Binomial}(x|n_h(\mu, F), 1 - F_h(\mu, F)) \quad (6.23)$$

$$P(x|\mu, F) = P_l(x|\mu, F)(1 - \Delta F(\mu, F)) + P_h(x|\mu, F)\Delta F(\mu, F) \quad (6.24)$$

where $P(x|\mu, F)$ is the final probability distribution. This is essentially the weighted mean of two binominal distributions given a non-integer mean with weights defined by how close the Fano factor is to the intended Fano factor. You can verify that this weighting gives the correct mean and Fano factor.




Figure 6.22: An excluded dark photon signal for the different Fano factors considered as part of our ionization model, described in the text. The science spectrum is shown in the background for comparison.

For the sensitivities shown in the next sections we use the measured high energy F of 0.155 [79]. We also vary the F used in the ionization model from its lowest mathematically possible value to 1 to estimate our sensitivity to the unmeasured ionization distribution width at low energies. This is another parameter that has not actually been measured, and that we can measure, using this detector in the near future.

Finally, we convolved the predicted e-h pair spectrum with the experimental resolution of 0.1 e-h pairs to produce a simulated spectrum for each model. An example dark photon model that is excluded by our analysis can be seen in figure 6.22 with the different Fano factors mentioned above, showing how the signal is redistributed based on the Fano factor with mean determined by our ionization model. For all limits in this analysis the Fano factor uncertainty as shown as a solid band, and it contributes a very small difference due to the threshold-driven nature of these experiments; if anything, larger Fano factors produce more competitive limits.

6.4.2 Dark Photons

Limits for dark photon absorption as a function of mass for the energy range considered here can be seen in figure 6.23 compared to the solar limits, DAMIC (CCD) limits [7], and Xenon10/100 [35] limits. Our analysis is competitive with DAMIC down to 5 eV, and achieves a lower threshold due to the fact that DAMIC had a charge resolution of $\sigma \sim 1.8 e^-$. This limit is based on a high-precision measurement of their 0-charge resolution where they infer that their dark rate is $< 10^{-3}$ electrons per pixel per day, as mentioned in the previous chapter. They use a likelihood analysis of their data, which is just a 0-signal Gaussian without a statistically significant tail, and use that to set limits based on this dark rate.




Figure 6.23: Top: Limits on dark photon kinetic mixing compared to the results from DAMIC, XENON10 and XENON100 [7] and references therein]. Bottom Left (Bottom Right): Limit on DM interacting with electrons via a heavy dark photon ($F_{DM} = 1$) (ultra-light dark photon ($F_{DM} \propto 1/q^2$)) compared to the XENON10 results [35]. The red line is the limit curve with a Fano factor of 0.155. The salmon colored region indicates the systematic uncertainties due to varying the Fano factor in the ionization model between the lowest mathematically possible value and 1, as well as from uncertainties in the photoelectric cross section for dark photon absorption. For signal models as well as additional astrophysical constraints, see Ref. [2]

The criticism of that analysis is that it has no discovery potential. If they had measured a dark rate, it would also be consistent with a handful of different noise sources; they also do not have any calibration to show that they would have been sensitive to a dark rate. In contrast, all dark photon masses considered in our limit are well above our detection threshold, and we are background limited. As a result, our limit suffers at low dark photon mass. One could argue that if there really was a dark photon at this mass and cross-section, DAMIC would likely be insensitive to it. We are also unsure of their ionization model, and despite some communications with the corresponding authors of Ref [7] we still do not know whether, if we were to use their ionization model, we would get a similar limit.

Also shown is the solar limit, which is set by the solar cooling bound given that the core temperature of the sun is around 5 keV in energy. This sets a limit at that mass and also on all lower

masses, with the proportionality

$$\epsilon^2 \propto m_V^{-2} \tag{6.25}$$

given that lower mass dark photons obey the same limit, but are emitted relativistically, and their mass only enters the rate equation in the coupling term. This is still the strongest limit on dark photons in this mass range. It's possible that our limits in their current form could be used to set an indirect limit on dark photons due to earth cooling, which is an analysis in progress at the time of this writing in collaboration with Peter Graham's group at Stanford.

6.4.3 Sub-GeV Dark Matter

Figure 6.23 also shows our limits for electron-recoiling dark matter compared to the Xenon10 [35] limits. These results exclude orders of magnitude of previously unrestricted parameter space over an order of magnitude in dark matter mass between $500 \text{ keV}/c^2$ and $5 \text{ MeV}/c^2$. Our results are also more competitive with the Xenon10 limits for the momentum-dependent form factor due to the fact that the signal is suppressed at high momentum; our low threshold makes us much more sensitive to these models even with a much smaller exposure. XENON10 had a 15 kg day exposure compared to our 0.5g day exposure, a difference of 30,000 in total exposure but less than 2 orders of magnitude in reach at higher masses.

Coincident with our results, SENSEI [102, 101] published limits with 0.019 gram-days of exposure, but had to run at a higher resolution than anticipated due to high backgrounds encountered in their surface run; their operating resolution was about 0.2 e-h pairs. As a result of their higher resolution, and their inability to reject dark rate pileup, their limit is the same shape as our limit but about two orders of magnitude higher in cross-section. Their paper lays the groundwork for using surface limits to set upper bounds on dark matter cross-section, which should be considered for a future analysis. At the time of this writing, SENSEI is beginning commissioning of their full CCD array and installation underground at Fermilab, and they are expected to achieve much more competitive limits within the next couple of years.

6.5 Future Directions

Focusing just on this dark matter search (and ignoring detector physics), we find a large number of areas for potential improvement based on the lack of understanding of our backgrounds just with this detector. Some follow-up studies would greatly improve the reach of our search and accuracy of our background models:

- Temperature stability improvements to obtain a consistently better resolution at 25 mK rather than operating at 35 mK. This could possibly reduce the resolution from 10 eV to 6 eV for a good noise environment.

- Further IR shielding improvements, assuming IR is the primary cause of the primary background as shown in our background simulations.
- Better understanding of how neutralization and pre-biasing impact the volume leakage rate and high-voltage surface leakage. Even with the same IR background, it's possible that we can reduce the leakage rate by de-populating these targeted states.
- Further study of impact ionization and trapping as a function of bias and polarity to determine whether there is an optimal bias from a background perspective, and to make the background model more robust. This was discussed in the previous chapter but deserves to be reiterated given the context of this chapter.
- Better RF shielding to eliminate non-physical events introduced by the environment, or environmental monitoring to allow these events to be vetoed. The jagged pulses seen in figure [6.19](#) may be trains of RF pulses, for example.
- Completion of Am-241 or other Compton calibration to study position dependence in the crystal. Some position dependence may also be probed by increasing the laser spot-size, although this also allows for multiple photons to hit different parts of the detector.
- Direct measurements of the quantum yield, and Fano factor, as a function of energy and recoil type, to improve our confidence in the yield model employed for limit setting. This also includes better measurements of the photoelectric cross-section in Si and Ge which can be done with this or similar setups.
- Studies of muon events and their impact on the steady-state leakage rate. This will require very high statistics data and may take advantage of changing the crystal orientation to impact the muon energy deposition rate.
- Further studies of the impact of additional shielding (both from radioactivity and muons) on the background across the energy window.

These are only some of the ideas which might allow us to achieve better performance with this device, but it should be clear from this list and the chapter that there is much work to be done to understand how to fully take advantage of the science potential of this device.

Chapter 7

Future Directions in Low-Threshold Detectors

“The chance is high that the truth lies in the fashionable direction. But, on the off-chance that it is in another direction – a direction obvious from an unfashionable view of field theory – who will find it? Only someone who has sacrificed himself by teaching himself quantum electrodynamics from a peculiar and unfashionable point of view; one that he may have to invent for himself.”

- Richard Feynman, 1965 Nobel Lecture

In this thesis I’ve explored the process of designing and operating low-threshold cryogenic calorimeters which use Si and Ge as their substrates, QETs as the power to current amplifiers, and low-temperature SuperCDMS electronics to readout signals. I’ve also shown that the most fruitful path forward for these detectors in the dark matter space relies on lowering thresholds, and then lowering backgrounds. The regime we’ve entered with these detectors also lends itself to generic photon detection, and starts to probe the energy scale of coherent interactions, making these detectors more generally interesting for low-background, precision particle physics. In this chapter I’ll explore some ideas for paths forward to lower thresholds further in SuperCDMS detectors and reduce low-energy backgrounds. I’ll conclude by discussing the challenges associated with new detector designs and targets which represent a further iteration on the SuperCDMS approach.

7.1 Lowering Detector Thresholds

In this section, I take a stab at answering the question ‘how do we make better SuperCDMS detectors’ focusing on how we might make short term improvements to our existing detector technologies,

imagining that we're still focused on Si and Ge substrates. This is by no means complete, but is some of the takeaways from the work in the previous chapters that I believe would be useful R&D. I have broken it down by the different parts of the detector, starting from the TES and moving up through the readout chain to explore some of the issues seen in the recent generation of test devices.

7.1.1 Alternative Detector Materials

The current generation of detectors was designed around a philosophy of reducing T_c given our Tungsten TES sensors. This is based on the resolution scaling law

$$\sigma^2 = \frac{2k_bGT_c^2}{\epsilon^2} \left[\frac{2\tau_{TES}}{5} + \tau_{pulse} \right] \quad (7.1)$$

which we've seen a few times. In most of this thesis we've assumed we're going to be phonon bandwidth limited, but let's imagine we're TES bandwidth limited, as is the case for many of the smaller devices; this was the pedagogical example from a few chapters back. This gives us the resolution

$$\sigma^2 = \frac{4k_bT_c^2}{5(\mathcal{L}-1)\epsilon^2} C = \frac{4k_b f_{sc} T_c^3 c_{TES} V_{TES}}{5\epsilon^2 (\mathcal{L}-1)} \quad (7.2)$$

where c_{TES} is the specific heat capacity of the TES material, and the rest of the parameters were defined in Chapter 3.

At first glance, it looks like we still have a very strong power of T_c to overcome should we move to a higher T_c material, but there's actually more going on. Firstly, \mathcal{L} is actually a fairly strong function of T_c . A very sharp transition occurs for films with the T_c close to the bulk value, but for tuned T_c we'll get a higher responsivity (larger \mathcal{L}) when we're using a film which acts like a bulk superconductor, which improves our sensitivity. In addition, part of the T_c scaling is actually the scaling of the heat capacity. The heat capacity sets our saturation scale, so (all else being equal) increasing the T_c raises the energy range over which the TES is linear, something we saw was a concern and limits the upper range of the readout energy. We can calculate a rough dynamic range as

$$\frac{\epsilon C \Delta T}{\sigma} \propto \sqrt{\frac{c_{TES} V_{TES}}{\mathcal{L} T_c}} \quad (7.3)$$

So strictly speaking, increasing T_c can be expected to weakly lower the dynamic range, but if the sharpness also increases then the dynamic range will decrease even more. Using a TES with a higher specific heat capacity is an easy way to increase dynamic range, although it also leads to increased resolution.

The take-away here is that for Tungsten, going back to a higher T_c isn't going to allow us to improve our resolution; lowering T_c gets us closer to the bulk properties and thus makes a sharper transition. If we can live with a smaller dynamic range, however, we could opt to reduce the volume of the Tungsten. Or we could opt for a different material with a higher T_c as long as it also had a very

sharp transition and we could lower the volume of the TES relative to a tungsten design. This is the difficulty with optimizing a detector based on TES properties rather than phonon properties; many changes counteract, and there is not a clear optimization strategy that is independent of material. An infinitely sharp transition has excellent resolution but zero dynamic range...it's just a threshold detector, like a bubble chamber.

An interesting prospect would be to use Al/Ti bi-layers for our TES. They have a T_c which can be tuned between 0.1-1.2K [68] and for the higher T_c designs have very sharp transitions. In addition, they're very fast (microsecond fall-times) and have seemingly comparable thermal conductance to tungsten, implying a Tungsten-like G but with a much lower C . Given that we're targeting lower volume, the lower resistivity wouldn't be an issue, and they have now been demonstrated to achieve good performance as a TES. It's unclear from existing data whether phase separation would be an issue, but initial results indicate they would be an interesting material to try for SuperCDMS.

I'm trying here to motivate raising T_c ...but isn't that what we don't want to do? The reason I'm even bringing this up is that the 40 mK target T_c of SuperCDMS, and lower T_c of detectors like CRESST, put an incredibly tight constraint on all of our experimental infrastructure. Only very new dilution refrigerators can achieve the necessary base temperatures, and we've clearly seen that our existing thermalization strategies do not work very well below 50 mK. The best refrigerators are dry dilution refrigerators, meaning they have all of the vibration and RF issues we've seen so far. It may be worth aiming for a higher T_c device if we can make our readout faster to trade T_c for device response time. This won't help our larger detectors, but will be interesting for devices with sub- μ s phonon collection times. Al/Ti would also allow us to fabricate the TES directly into the aluminum fins and play with the relative trapping depth at the TES, which would add an additional optimization handle and possibly improve QET efficiency.

As a final note, it's imperative that we improve our clamp design for all of the detectors, especially larger detectors, so that we can be sure that our crystals are actually at the base temperature indicated by the fridge thermometry. It's clear now that Cirlex stops being a good thermal conductor at 50 mK, and causes a spike in cooling time which keeps the crystal at a higher temperature than the fridge. Sapphire or diamond clamps (which don't need to be single crystals) should not have these spin-glass issues, and should be good thermal conductors down to 0 K...but to my knowledge there aren't measurements which prove this yet. Significant R&D is starting to explore additional clamp mitigations, but this is certainly an important area of research for SuperCDMS.

The main concern about clamp changes is that they maintain electrical isolation, so that we don't induce a surface leakage current by placing a clamp (implicitly grounded) on the crystal, biased to some large voltage. One solution to this problem has historically been the contact-less housing, where a floating electrode is used to generate an electric field that penetrates the crystal; this approach suffers from time-dependent droop as charge builds up, as well as uncertain total voltage drop across the crystal. New housing designs utilizing engineered clamps made of high thermal conductivity, low

electrical conductivity materials need to be developed, and in the process we should be sure to design them to produce as uniform an electric field in the crystal as possible, to reduce position-dependent NTL gain droop. As shown in the detector design chapter, clamp and detector design are closely tied to concerns about surface leakage, so this is a complex design issue that we have never really addressed in detail.

7.1.2 Improving QET Resolution

On the flip side, imagine we're trying to really push to lower resolution; this is still the post by which to guide ourselves. We should collect phonons quickly, reduce the size of our TESs, and go to the lowest achievable T_c , ~ 15 mK in Tungsten. There are a number of improvements we need to make to maximize the efficiency of our designs, and allow for the possibility of reducing T_c further.

The easiest way to lower the resolution of an existing design is to improve our fabrication process. We're currently limited to Tungsten line-widths on the order of 2–3 μm , but thicknesses of 40 nm are sufficient to achieve the right T_c . This means the correlation lengths are much shorter than 2 microns. If we could achieve 0.5 micron line-widths, we could make significant design changes (higher resistances, lower volumes, shorter TESs) which would immediately improve the robustness of our designs. If we were to drop all of the line widths on the HV TES to 0.5 microns without other changes, it would work much better (it would solve the resistivity issues) and reduce the resolution by a factor of 2, all other things being equal. We'd likely also want to reduce the length to move away from potential phase separation issues, but this would be a no brainer.

From here we have a few ways to go. We can try to improve efficiency, but there's not a ton of margin there; we've achieved about half of our theoretical maximum for phonon down-conversion in the Al fins. We could cut resolution in half by improving the CDMS HV efficiency to its theoretical maximum; the limitation is then trying to determine how to improve the energy transport from the crystal to the Al fins, which is beyond the scope of this thesis but is worth considering. If we can understand how to better connect the fins to the TES (reduce the amount of tungsten needed between the TES and the fins, and how to optimize the overlap) we should be able to reduce the effective TES volume for the same QET and QET channel design. For example, we have some evidence that increasing Al film thickness also increases quasiparticle collection efficiency. If we can double the length of the Al fins on the HV mask, then we have an incredibly fast phonon absorption time with the same TES volume or we can match the phonon absorption time with less than half of the volume, doubling the resolution. The path forward, then, is to continue to work on improving the model of the QET interface for phonons into the fins, and quasiparticles to the traps, to allow these regions to grow and thus allow the TES volume to shrink.

Within our existing framework, there is a credible path forward by exploiting improved fabrication processes, largely achievable through changes to the alignment and photoresist processing and etching. Figure [7.1](#) shows the optimization space for the smaller line widths and updated device




Figure 7.1: 2D optimization plots for updated CDMS HV fab specs, including reduced wire widths ($10 \rightarrow 2 \mu\text{m}$) and reduced TES width ($2.4 \rightarrow 0.5 \mu\text{m}$) as well as higher R_n ($150 \rightarrow 300 \mu\text{m}$) and thicker fins at 1.2 microns. I also updated this code to take into account a modified resistivity expectation of around $\rho \sim 5 * 10^{-8} \Omega\text{m}$ found from our test devices. Top Left: Energy Efficiency. Top Right: First-pass phonon collection efficiency. Bottom Left: Energy resolution. Bottom Right: Phonon collection time. Notice that we now have a broader space in which we can achieve the same resolution, and we could even change the existing design according to these parameters without any significant modifications aside from changing the QET number. This demonstrates that without improvements in quasiparticle collection, our only other handle is fabrication. We improve efficiency somewhat by reducing line widths, but that only allows us to make smaller QETs, not significantly improve our resolution.

information from the testing program, and shows that we can achieve ~ 4 eV resolution for a wide range of QET aspect ratios. What we notice though is that these improvements seem just to flatten out the optimization space. It's possible that as the fins get large enough, we can change the wiring to look like the iZIP wiring scheme which minimizes dead aluminum (the rails are just part of the QETs); it's unclear whether that would be a significant improvement. If we're unable to improve the QP collection efficiency with film thickness and quality improvements, or other mechanisms, we're still just essentially TES volume limited, and the only way we can appreciably reduce our volume is by making thinner structures or using smaller substrates.

To follow this to its logical conclusion, let's consider how we'd take the HVeV design and lower the resolution below 1 eV, to take advantage of a smaller detector. It's a 3% efficiency device, so that's the largest problem, but even without changing the QETs we can improve the efficiency by moving to double-sided readout. If we clean up the surface roughness and move to a double-sided design, we should conservatively be able to achieve 10% efficiency. We think that, in a clean noise environment and with a low bath temperature, that device is intrinsically a 3–5 eV device. So tripling the efficiency gets us close to 1 eV. Now consider the phonon fall time argument. We expect a TES response on the order of 100–200 μ s, and equation 4.24 gives us the phonon collection time

$$\tau_{pulse} = \frac{t_{transit}}{f_{al}f_{abs}} = \frac{h_{chip}}{v_{phonon}f_{al}f_{abs}} \approx \frac{4mm}{(2500m/s)f_{abs}} = 1.6 \mu s f_{abs}^{-1} \quad (7.4)$$

where I've scaled the values by 3.5 for the better impedance match and high sound speed in Si compared to the Ge numbers. For the current device, this gives a fall time of 6.5 μ s for the phonons, but we'd like to have that fall time be equal to the TES fall time. Even for a very quick TES response, we can go to 4% coverage and still have a phonon fall time of around 30–40 microseconds. This doesn't change the resolution for the same QETs with shorter fins, but will improve efficiency; more importantly though, it allows us to use smaller TES units for lower tungsten volume. We can go from 25 QETs to 12 QETs per channel, with smaller fins, and shorten the TESs as well to get a reduction in volume by a factor of 4, or double in resolution. This would allow a device with ~ 0.5 eV resolution not taking into account optimizations for efficiency discussed earlier.

7.1.3 Improving Readout Electronics

A significant hurdle faced by the SNOLAB warm electronics is the impact of EMI and RF noise on the energy resolution of our QETs. The first of these contributes noise peaks that can couple both into the TES bias line, contributing voltage and power noise, as well as onto the board, creating noise at the pre-amplifier input which is independent of TES bias state. The second, however, is more insidious, as we discussed in chapter 5, because it constitutes a power noise. Low-frequency power noise like vibrations can be mitigated with the correlation techniques discussed earlier, but high-frequency power noise down a bias line is a problem we cannot remove in analysis.

To see why, let's consider an RF signal modulated by a periodic function $u(t)$ with a carrier frequency f_c . We have the time-domain amplitude signal

$$y(t) = u(t) \cos(2\pi f_c t) \quad (7.5)$$

and if this is a current noise down the TES bias line, then we find the power as a function of time is

$$P(t) \propto u^2(t) \cos^2(2\pi f_c t) = \frac{u^2(t)}{2} [1 + \cos(2\pi f_c t)] \quad (7.6)$$

We will never see the high-frequency component if it is above our dI/dP pole, but we do see low-frequency artifacts because of the rectification of current into power noise, just like we saw for vibrational power noise. We can explore various ways RF can couple to the detector which have different scalings as a function of TES bias point, but the point stands that this RF power will contribute to low-frequency noise.

The most problematic case is when there is a simple power coupling directly into the TES, and the RF is constant in amplitude as a function of time. Then what we would expect to see is a white noise with a DC power component that causes our noise model based on the bias power to under-predict our observed noise. This is in fact exactly what we see on all of our detectors except on some of the small TES test devices, and is heavily suggestive of a direct RF coupling onto the detector. Early results from TES test devices as well as detectors of different kinds indicate that RF power noises likely couple both as joule and absorption power noise, meaning that both small TES and large detectors are affected, while a middle size TES test device may be less prone to RF induced noise.

The solution to all of these problems is to move the pre-amplifier stage inside of the fridge, and use the fridge as a Faraday cage to block EMI and RF. We've had some success using Pi filters between the DCRC and readout electronics, but what we've discovered is that we have to make some hard decisions about whether we filter the TES bias line. We should always use Pi filters on all lines into and out of the fridge, and make the fridge as RF tight as possible, but allowing signals below 1 MHz to pass down the bias line introduces both EMI and power noise into all lines in the system. The trouble here is that we'd like to have a high bandwidth on the TES bias line, and we need a high bandwidth in the SQUID feedback loop, in order to do complex impedance measurements and to operate the TES in closed loop.

The second concern can be mitigated by putting the entire SQUID feedback loop inside the fridge Faraday cage, and using highly filtered analog signals to set any gains and biases from outside the fridge. This will reduce possible coupling of EMI into the SQUID and from the feedback line into the TES. This is similar to what is implemented in the Northwestern ADR, and has resulted in significant power-coupled noise reduction. Putting the pre-amplifier inside the fridge, and placing it before the pi-filter, also allows us to decouple the SQUID and SQUID bias line from RF signals. We

can't, however, put a DC filter on the output of the pre-amplifier for obvious reasons, but if the gain is high enough by the time the line exits the fridge, it should make EMI relatively less important.

The last concern about complex impedance is the most important however. Any filtering we don't do on the TES bias line allows outside EMI and RF to get to the TES, which directly impacts our resolution. The simplest solution would be to severely low-pass filter and PI filter the TES bias, but that would prevent us from using square and sine waves to measure complex impedance. The potential solution to this is either to put the bias circuitry inside the fridge Faraday cage as well, or try a new technique to measure complex impedance. We could, for example, heat up the shunt resistor so that its voltage noise exceeds the other noise in the system, and use that as an in-situ measurement of complex impedance. In any case, we need to re-design how the analog and digital sections of the DCRC connect, and which components are placed inside the fridge, to mitigate these fairly major issues.

7.1.4 Phonon Imaging and Noise Correlation

One of the main lessons I took from the early CDMS HV tests is that our method of optimal filtering breaks down severely for highly position dependent signals, and we do much better with sophisticated signal templates. When we have any correlated noise or a joint-channel template, we can always improve on the naive optimal filter result. In our earlier detectors, where we had 8 channels and used low voltage bias, or in CDMSlite with 4 channels, we didn't have enough information to reconstruct position and utilize joint-channel methods beyond our two-template fits, which still required summing channels and did not account for inter-channel correlations. We can take lessons from imaging and signal processing to see that we can improve on the naive optimal filter summing resolution if we can increase the number of readout channels on a single detector.

At the end of the day, the energy resolution of a QET channel is the quadrature sum of the resolution in each TES. This means that, fundamentally, the resolution of each TES is given by the equation

$$\sigma_{TES} = \frac{\sigma_{det}}{\sqrt{2N_{TES}}} \quad (7.7)$$

where N_{TES} is the number of TES per channel on a given side. For the large detectors, consider the fact that the NTL phonons are highly collimated, and may be focused withing a beam of size r where r decreases as a function of voltage. This means that, if we only use pixels with NTL phonons collected early in the event, we can obtain a resolution

$$\sigma_{local} = f_{NTL} \sqrt{\frac{A_{NTL}}{A_{Det}}} \sigma_{det} \quad (7.8)$$

where f_{NTL} is the fraction of the phonon energy inside our cone. So if we can isolate half of the total energy in this area, we find that we can improve the resolution by a factor of 5 compared to

the naive summing method. If we have each TES read out on a single detector, we can do PSF fitting. In this case we can obtain an energy resolution fitting in one dimension of [43]

$$\sigma_{det,psf} = \sqrt{6\sqrt{\pi}\frac{r}{w_{cell}}}\sigma_{TES} = \sqrt{3\sqrt{\pi}\frac{r}{w_{cell}N_{TES}}}\sigma_{det} \quad (7.9)$$

where w_{cell} is the side width of a QET unit cell in SI units and r is the standard deviation of the PSF in meters. In the limit that all of the energy is concentrated in a single QET, this is a potentially much larger improvement in energy resolution, but in the limit of very large variance reduces to the same sum of QETs that forms our naive estimate. This of course needs to be done in two-dimensions (according to the formalism derived in Ref [42]) but this is a nice demonstration of why doing local fitting helps reduce parameter uncertainties, and why having a position independent signal is not necessarily a bad thing if it can be properly imaged.

To do this would necessitate some sort of on-detector multiplexing, either by using many RF squids on a single line, each connected to an individual TES, or coming up with some means to fabricate naturally multiplexed structures on the detector surface. For this reason, the Caltech group has been working on using Microwave Kinetic Inductance Detectors (MKIDS) [114] to read out CDMS detectors for about a decade [76]. These are superconducting sensors that change their resonant frequency depending on number of broken cooper pairs along their conduction path, allowing them to be naturally multiplexed along a single RF bias line fabricated on the detector. Other groups have already shown that these devices have a lot of promise for imaging [71], but as of this writing a working MKID readout for a CDMS detector is yet to be realized. It's possible that within the years after this writing, a currently fabricated Al-based MKID detector may show sub-keV resolution, and start to compete with our TES readout in terms of resolution, at which point PSF fitting may allow MKIDs to leapfrog TES readout. It's also possible we can come up with an RF multiplexing scheme for TES arrays using some sort of tank circuit fabricated on the detector, but it's less clear how to proceed in that direction technically.

Finally, we can improve our ability to reject environment noise by including the PSD of a second detector or independent TES channel which should have no correlation to our channel in an ideal noise environment. This is easy to do with the correlated optimum filter techniques discussed in this thesis; the idea would be to use a pulse template based on the given detector channel, and for the environmental channel use a null template. The formalism will then allow us to subtract off the correlated noise between the two channels and fit the pulse only limited by the component of the noise uncorrelated between the detector and environmental channel. This is a very powerful way for us to remove vibrational noise, for example. It will not help us to remove DC power noise caused by RF, which can only be solved with the techniques outlined in the previous section.

7.2 Studying Low-Energy Backgrounds

There are a multitude of new detector background features at low energy that become very important to model if we really want to understand our detector response. Starting from extrinsic backgrounds and moving to intrinsic backgrounds, the first concern is that we accurately model gamma and neutron interactions in our detectors. We can break these down into differential rate measurements and ionization measurements, both of which are very uncertain at the time of this writing.

7.2.1 High Energy Gammas Producing Low Energy Events

Below about 1 keV, the nature of our backgrounds change drastically. Above this point, and around 60–200 keV where our normal gamma and neutron sources tend to be, we can generally assume that any background is fairly isotropic in the crystal, and that it produces a well-defined number of electron-hole pairs. Compton scattering spectra for most sources have been measured at comparable resolutions so we're mainly using past measurements to predict signals give our detector response. For a detector with sub 10 eV resolution or sensitivity to single charges, there are a few processes which come into play that complicate these calibrations. One in particular is the trade off between elastic Compton scattering, for interactions of a high energy γ with an electron, Rayleigh scattering of these photons off of the entire atom, and coherent γ interactions with an atomic nucleus [88].

The first and last of these mechanisms should resemble electron and nuclear recoils, but the middle term is less clear. Sensitivity to MeV-scale DM also relies on electron-mediated coherent atomic collisions which produce non-trivial charge production and energy transfer, so more work needs to be done to determine the charge yield of coherent photon scattering events as a function of energy deposition. It's possible with the charge-resolving detectors to make these measurements given a high-energy gamma source and a small-angle coincidence detector at a test facility, and in the process we would gain more confidence in the charge generation models shown in Ref [33] on which our DM search spectra are based. The study for Ref [88] focused on searches for nuclear recoils, assuming this is just a small nuclear-recoil energy deposit, but it becomes important to recognize that it may have a larger charge yield, and therefore be a more important background in our electron-mediated HV searches.

The nature of our backgrounds is also entirely different now, especially given that we're sensitive to not only electron-hole pairs but also fractional charges from bound impurity states. This means it's a dangerous prospect to assume any energy transfer below the bandgap is sub-threshold in HV mode, given that any energy deposit can still free enough fractional charges to generate an event with significant NTL gain. If an event can scatter off multiple crystal impurity states in a single event, it's possible it could even produce a high energy tail. For this reason it will be important for us to continue our program of high-energy gamma calibrations even though the primaries are now far outside the energy regime of interest.




Figure 7.2: Top: IR spectral irradiance as a function of blackbody temperature. Bottom: Relative IR rate per impurity as a function of IR temperature for fixed impurity depth (left) and as a function of impurity depth for fixed temperature (right).

7.2.2 Sources of Dark Events

The topic I spend most of my time thinking about these days is dark rates; how they’re produced, how to reduce them, how we might alter our detector designs to lower the rates given the handles we have. Remember from chapter 5 that these events are consistent with only one charge carrier (either an electron or hole) drifting across the crystal from its initial location to a detector surface, undergoing a fraction NTL gain and producing (if uniformly distributed in the crystal) a flat distribution between 0 and 1 full electron-hole pair. This can be either an electron or a hole, and if it’s produced uniformly in the bulk, we won’t see an difference in polarity. The non-quantized nature of the dark rate, at least at lower voltages, tells us that it must be bulk driven, and cannot be surface driven, otherwise it would be quantized and polarity dependent, and it cannot be localized to a defect for the same reason.

Let’s consider a simple model then for this leakage. Suppose we have a bound state of known energy E_T on an impurity (or a site of known density which can form a stable bound state). It can

be freed by any process which exceeds this energy threshold. We can then calculate the expected rate as the integral

$$\Gamma(n_I, E_T) = n_I \int_{E_T}^{\infty} \frac{d\Gamma}{dE}(E) \quad (7.10)$$

which we can see scales proportionally to density, but its scaling with respect to the bound state energy depends on the nature of the exciting event and its energy dependence. For infrared radiation, if we assume the cross-section is independent of energy, we obtain a rate scaling law of the form

$$\Gamma_{IR}(n_I, E_T) \propto n_I E_T^3 e^{\frac{E_T}{k_b T_{IR}}} \quad (7.11)$$

where T_{IR} is the IR blackbody temperature. Figure 7.2 shows the IR leakage rate per impurity for difference blackbody temperatures and impurity binding energies. We see that for the same IR environment, we get an exponential reduction in leakage rate for a doubling of binding energy at low IR temperature, but 300K IR is insensitive to binding energies below 100 meV. We also see that if we only have 4K or lower blackbodies, impurities above 20 meV will not contribute to this leakage rate. These plots should allow us to determine whether IR is the cause of our charge leakage.

If the rate is due to Compton scattering, we can assume it's flat in energy and reduced in rate by the ratio of atomic to impurity density:

$$\Gamma_{Compton} \approx \frac{n_I}{n_A} \frac{d\Gamma}{dE} \Big|_{E > E_{gap}} (E_{gap} - E_T) \quad (7.12)$$

where

$$n_A = \frac{\rho N_A}{m_T} \quad (7.13)$$

is the number density of atoms in the lattice with ρ the density in g/m^3 , m_T the atomic weight, and N_A Avogadro's number. The ratio $\frac{n_I}{n_A}$ is useful for comparing detectors of different impurity densities and target materials in the same environment and is a clean indicator, for two typical crystals, of whether Compton scatters are the primary culprit. More importantly however is that for any reasonable crystal we should expect Compton scattering off of impurities to be highly suppressed compared to electron-hole pair production due to this ratio. The highest purity Si typically has 10^{13} cm^{-3} impurities (Ge has 10^{11} cm^{-3}) meaning that the Compton-induced dark rate due to this mechanism is suppressed already by 10^9 in Si and 10^{12} in Ge. This can't be our leading term.

The coherent scattering comes back in here, however; we know that we can enhance the rate below the gap energy by 6 orders of magnitude compared to the Compton rate for potassium 40 [88], which brings these numbers to within a few orders of magnitude. Depending on how this rate scales based on primary photon energy, and how the cross-section may be enhanced for different impurity species (we may get a much higher rate for over-charged states) we could start to find a situation where the rate is much higher than the Compton rate. That may still be a stretch given that the




Figure 7.3: Left: Schematic of the different autoionization processes described in Ref [39]. Right: Autoionization rate for overcharged states as a function of electric field strength, taken from Ref [99], for donors (blue) and acceptors (red).

science data shows this dark rate to be many orders of magnitude higher than the Compton rate. Studying these backgrounds with different incident photon energy will very quickly tell us whether this might be a viable mechanism for at least some of the leakage events.

Finally, let's consider auto-ionization; the probability that some sub-gap state can tunnel free under an electric field or due to phonon interactions. I won't go into depth on the rate calculation, which is derived in detail in Ref [99]. The process is schematically shown in figure 7.3 where the de-trapping gets more probable as field strength is increased (the trap tips over more) or the temperature is raised (phonon assisted tunneling increases in probability). On the right we can see for the most weakly bound states, the autoionization rate is predicted to be around 1 Hz around 20 V/cm, much lower than the field strengths we ran in the HVeV detector, and is exponentially increasing; the fact that our leakage is independent of voltage suggests these states may not be the dominant leakage yet if this calculation is correct. Impurity states, on the other hand, reach 1 Hz rates around 1000 V/cm, and also increase exponentially. We observe rate increases at high voltage around 400 V/cm, but the leakage is quantized, so it can't be the autoionization. The voltage dependence of the leakage as we continue to reduce our IR backgrounds should help us determine whether autoionization is a prominent factor in our leakage.

The true impurity state picture is more complicated when we start to consider the different




Figure 7.5: Range of dark matter search potential reach for a variety of techniques and detectors, taken from Ref [2].

on all of the possibilities currently being discussed.

From a SuperCDMS perspective, we can continue to use Si and Ge in HV mode down to the Si/Ge bandgap energy, which sets the minimum electron-recoil dark matter mass we’re sensitive to. In terms of NR, there is no real detectability floor, but as discussed earlier in this chapter, it’s going to be challenging to go below about 0.1 eV in the foreseeable future. If we can continue to refine our QETs to achieve smaller tungsten volume, we can also make our detectors somewhat smaller, and produce more of them, to build up to the kg-year sensitivity that will be required to make significant dents in the sub-GeV/ c^2 mass range. At the time of this writing, work is underway on preliminary limits down to around 300 MeV/ c^2 using a 6 eV resolution device. The challenge lies in background rejection, but any device which can achieve lower resolution will immediately be able to set world-leading limits with a few days of exposure.

What substrate might do better for nuclear recoils that could utilize SuperCDMS technology? The answer is the only other indirect bandgap semiconductor: diamond! Diamond benefits from better momentum transfer for Sub-GeV dark matter (being lighter than Si, it is only beaten out by He in terms of momentum transfer) and it has a faster phonon velocity (2.5 times faster than Si), and higher energy, more stable phonon modes than either Si or Ge. It can hold electric field strengths up to MV/cm and has no shallow impurity binding sites. The downside is that it has a larger bandgap than Si and Ge (5.5 eV), so it can’t reach the very low Sub-GeV masses that Ge and Si can, but it promises to be a very background-free, excellent single UV photon counting device

which is capable of extremely low energy resolutions if we can overcome the technical challenges of making working QETs on diamond. This is a very exciting prospect that we're beginning work on at Stanford.

Further in the future, CDMS technology can be extended to superconductors using the same quasiparticle trapping approaches we employ in our QETs. These ideas have already been explored in Ref [48] both using phonons and quasiparticles in superconductors. We could imagine using a niobium target with aluminum fins and a tungsten TES, or using aluminum MKIDs on a niobium crystal. The challenge with superconductors, however, is that quasiparticles and phonons decay much quicker than charge carriers and phonons in semiconductors [56], and we can't apply a bias voltage to redirect charge, so these are strictly diffusive devices. Superconductors are the best vehicle for detecting ultra-light DM (down to the fermionic limit) if many of the technical challenges can be overcome; this is likely further into the future but no less exciting.

7.4 Experimental Outlook

Given the possible new directions summarized above, it's clear to me that there are many paths forward for SuperCDMS technology to continue to push to lower mass and cross-section in dark matter searches. The nice synergy we get as we start to achieve eV-scale and sub-eV resolutions is that these devices also become interesting for more general imaging science, and may be useful as more sensitive readout stages for other dark matter experiments. It's likely, for example, that a device similar to the prototype discussed in this thesis will be used as the primary readout stage for a superfluid helium experiment, indirectly using SuperCDMS technology to achieve lower nuclear recoil thresholds.

In addition, as we solve our dark rate problems, and as we move into using different materials, it's likely that our devices will become simple enough to operate that they can begin to be used for single-electron resolution imaging. For Si, we could already use these detectors with the fiber coupling to do single-pixel imaging, and as we solve the outstanding problems with Ge leakage currents, this will extend our sensitivity in to the near-infrared. The understanding of dark current may also allow us to purposely dope crystals to extend sensitivity to photons in the 10 meV energy range, which would also prove to be a very sensitive probe for dark photons.

In all areas of dark matter and photon science, however, we face stiff competition. There are now three experiments working with single-electron resolution CCDs (SENSEI, DAMIC-M, and DANAÉ), the last of which claims comparable timing resolution to our detectors. CRESST has had success with a 3 eV detector lowering experimental thresholds to $100 \text{ MeV}/c^2$, but has a strong low-energy background that limits their sensitivity, and EDELWEISS is not very far behind. If we contrast our position today compared to when I started as a graduate student, we have clearly pivoted (successfully) to a low-mass experiment, and through dedicated R&D have come out on top,

for now, in the electron-recoil space for very low masses. It's clear that we need to maintain this pace of development to sustain our experimental edge before SuperCDMS SNOLAB begins operations. If any of these potential improvements bear fruit in the next year or two, we can certainly maintain this edge in the years to come.

Appendix A

Analytic Detector Partition Functions

In this appendix I develop analytic expressions for the radial and z-partitions in a detector with two composite channels to define each partition, comparing either the total energy on each face for the z-partition or the total energy on a face between an inner and outer separation for the radial partition as a function of z-position. In this way we can develop expectations for effective efficiency of surface events measured by sensors on the far face and attempt to correct for missing channels or one-sided readout.

A.1 Z-Partition

Starting with the z-partition, suppose that an event produces an isotropic phonon signal from its initial point. If the event occurs along a line in z at the center of an azimuthally symmetric detector, we can more easily develop an analytic form for the energy distribution and then hope to generalize to arbitrary radius. We find that the energy difference between each side is proportional to the difference in solid angle subtended by the side from the viewpoint of the initial event:

$$\Delta\Omega = \Omega_{up} - \Omega_{down}$$

For a crystal of radius R and thickness h , we find that the angle to the edge of the top surface (if we defined $z = 0$ as the bottom surface) is

$$\cos(\theta_{up}) = \frac{h - z}{\sqrt{(h - z)^2 + R^2}}$$

and the angle to the bottom is

$$\cos(\theta_{down}) = \frac{z}{\sqrt{z^2 + R^2}}$$

The solid angle subtended for some angle θ from the normal to the surface is found as

$$\Omega = \frac{A}{r^2} = \frac{1}{r^2} \int_0^\theta \int_0^{2\pi} r^2 \sin(\theta) d\theta d\phi = 2\pi \int_0^\theta \sin(\theta) d\theta = 2\pi(1 - \cos(\theta))$$

this gives us

$$\Delta\Omega = 2\pi [1 - \cos(\theta_{up}) - 1 + \cos(\theta_{down})] \quad (\text{A.1})$$

$$= 2\pi [\cos(\theta_{down}) - \cos(\theta_{up})] \quad (\text{A.2})$$

$$= 2\pi \left[\frac{z}{\sqrt{z^2 + R^2}} - \frac{h - z}{\sqrt{(h - z)^2 + R^2}} \right] \quad (\text{A.3})$$

In the limit of very large radius ($R \gg z, h$), we thus have

$$\Delta\Omega \approx \frac{2\pi}{R}(2z - h)$$

and for the opposite limit (equivalent to very collimated phonons) we find

$$\Delta\Omega \approx 2\pi \left[\frac{z}{z} - \frac{h - z}{h - z} \right] = 0$$

and in this limit we have no z -dependence. We can also take this further to actually calculate the true partitions based on an incident energy. For an event of energy E isotropically radiated from a given point, we have the angular partition

$$\frac{E}{\Omega} = \frac{E}{4\pi}$$

such that

$$\Delta E = f_{prompt} \frac{E}{\Omega} \Delta\Omega = f_{prompt} \frac{E}{2} \left[\frac{z}{\sqrt{z^2 + R^2}} - \frac{h - z}{\sqrt{(h - z)^2 + R^2}} \right] \quad (\text{A.4})$$

$$\frac{\Delta E}{E} = \frac{f_{prompt}}{2} \left[\frac{z}{\sqrt{z^2 + R^2}} - \frac{h - z}{\sqrt{(h - z)^2 + R^2}} \right] \quad (\text{A.5})$$

Here f_{prompt} is the fraction of phonon energy which carries the initial position dependence. For our detectors this will be approximately $f_{al}f_{abs}$ where f_{al} is the fraction of surface metal coverage and f_{abs} is about 40% in Ge and $\sim 70\%$ in Si.

This equation allows for a nice measurement of that prompt fraction and a validation of the

assumption of prompt initial distributions, but will not normally be very useful unless we incorporate the true radial position of an event. Here I generalize the z -partition to arbitrary radius. We first have to generalize that $\cos(\theta)$ function to incorporate ϕ asymmetry. We will assume that the radial shift occurs along $\phi = 0$, but this is general as the partition won't actually care about the azimuthal angle of the displacement from $r = 0$ given that we integrate over ϕ . We can do this ϕ generalization as

$$\cos(\theta_{down}(\phi)) = \frac{z}{\sqrt{z^2 + R^2(\phi)}}$$

where $R(\phi)$ is not the crystal radius but the distance to the edge of the crystal as a function of ϕ . We have

$$R(\phi) = \sqrt{(R \cos(\phi) - r)^2 + R(\sin(\phi))^2}$$

such that

$$R(\phi)^2 = R^2 - 2Rr \cos(\phi) + r^2$$

The integral we want to evaluate is thus

$$\Omega = \int_0^{2\pi} \int_0^{\theta(\phi)} \sin(\theta) d\theta d\phi$$

where the integration range is now a function of ϕ , and we have to do the integral over θ first. This is the trivial part of the integral, and we thus get

$$\Omega = \int_0^{2\pi} (1 - \cos(\theta(\phi))) d\phi = 2\pi - \int_0^{2\pi} \cos(\theta(\phi)) d\phi$$

Doing the integral for the bottom surface first for simplicity, we have

$$\int_0^{2\pi} \cos(\theta(\phi)) d\phi = \int_0^{2\pi} \frac{z}{\sqrt{z^2 + R^2 - 2Rr \cos(\phi) + r^2}} d\phi \tag{A.6}$$

$$= z \int_0^{2\pi} [z^2 + R^2 + r^2 - 2Rr \cos(\phi)]^{-1/2} d\phi \tag{A.7}$$

I've put the integral in this form to abstract it a bit; if we substitute $a = z^2 + R^2 + r^2$ and $b = 2Rr$, we see first that $b < a$ and we evaluate an integral of the form

$$\int_0^{2\pi} [a + b \cos(\phi)]^{-1/2} d\phi = a^{-1/2} \int_0^{2\pi} \left[1 + \frac{b}{a} \cos(\phi) \right]^{-1/2} d\phi = a^{-1/2} \int_0^{2\pi} [1 + c \cos(\phi)]^{-1/2} d\phi$$

where $c = \frac{b}{a} \leq 1$. This is a well defined integral and can be expressed in terms of elliptic integrals:

$$\int_0^{2\pi} [1 - c \cos(\phi)]^{-1/2} d\phi = \frac{2}{\sqrt{1-c}} F\left(\frac{\phi}{2}; \frac{2c}{c-1}\right) \Big|_0^{2\pi} = \frac{2}{\sqrt{1-c}} F\left(\pi; \frac{2c}{c-1}\right) = \frac{4}{\sqrt{1-c}} K\left(\frac{2c}{c-1}\right)$$

This gives us

$$\int_0^{2\pi} \cos(\theta(\phi)) d\phi = \frac{z}{\sqrt{z^2 + R^2 + r^2 - 2Rr}} 4K\left(\frac{2c}{c-1}\right) \quad (\text{A.8})$$

and so

$$\Omega = 2\pi \left[1 - \frac{z}{\sqrt{z^2 + R^2 + r^2 - 2Rr}} \frac{2}{\pi} K\left(\frac{2c}{c-1}\right) \right]$$

We lastly simplify

$$\frac{2c}{c-1} = \frac{2b}{b-a} = \frac{4Rr}{2Rr - R^2 - r^2 - z^2} = \frac{-4Rr}{(R-r)^2 + z^2}$$

to get the final equation

$$\Omega = 2\pi \left[1 - \frac{z}{\sqrt{z^2 + R^2 + r^2 - 2Rr}} \frac{2}{\pi} K\left(\frac{-4Rr}{(R-r)^2 + z^2}\right) \right]$$

This reduces to the $r = 0$ form given that $K(0) = \pi/2$. We then have the z-partition

$$\begin{aligned} \frac{\Delta E}{E} = \frac{f_{prompt}}{2} & \left[\frac{z}{\sqrt{z^2 + (R-r)^2}} \frac{2}{\pi} K\left(\frac{-4Rr}{(R-r)^2 + z^2}\right) \right. \\ & \left. - \frac{h-z}{\sqrt{(h-z)^2 + (R-r)^2}} \frac{2}{\pi} K\left(\frac{-4Rr}{(R-r)^2 + (h-z)^2}\right) \right] \end{aligned}$$

This is actually a very powerful equation; we can use this equation to predict the z-partition quantity for any event in the crystal, and the individual equations which went into calculating the z-partition can be used to calculate radial quantities as well.

This is a very useful equation, and is simpler than others you'll find in the literature. For a different and more rigorous investigation of solid angle subtended by a circle at an arbitrary point, see [\[93\]](#).

A.2 R-Partition

To compute the radial position as a function of r and z is more complicated, and does not necessarily lend itself well to any simple scaling laws. In the previous section, we assumed the z-partition was proportional to the difference in solid angle between the two sides. The radial partition, on the other hand, should be proportional to the difference between the outer and inner rings, summed over both sides. Writted explicitly, we have

$$\begin{aligned} \Delta\Omega_{top} &= \Omega_{outer,top} - \Omega_{inner,top} \\ &= (\Omega_{top}(R_{outer}, r, z) - \Omega_{top}(R_{inner}, r, z)) - \Omega_{top}(R_{inner}, r, z) \\ &= \Omega_{top}(R_{outer}, r, z) - 2\Omega_{top}(R_{inner}, r, z) \end{aligned} \quad (\text{A.9})$$

which gives the total solid angle difference

$$\Delta\Omega = \Delta\Omega_{top} + \Delta\Omega_{bottom} \quad (\text{A.10})$$

Let's consider the $r = 0$ r partition as a function of z to see how well this actually tracks r . We find

$$\Delta\Omega_{bottom} = 2\pi [(1 - \cos(\theta_{down,outer})) - 2(1 - \cos(\theta_{down,inner}))] \quad (\text{A.11})$$

$$= 2\pi [2\cos(\theta_{di}) - \cos(\theta_{do}) - 1] \quad (\text{A.12})$$

$$= 2\pi \left[2 \frac{z}{\sqrt{z^2 + R_i^2}} - \frac{z}{\sqrt{z^2 + R_o^2}} - 1 \right] \quad (\text{A.13})$$

and so

$$\Delta\Omega_{top} = 2\pi \left[\frac{2(h-z)}{\sqrt{(h-z)^2 + R_i^2}} - \frac{(h-z)}{\sqrt{(h-z)^2 + R_o^2}} - 1 \right] \quad (\text{A.14})$$

giving the difference in solid angle

$$\Delta\Omega = 2\pi \left[2 \frac{z}{\sqrt{z^2 + R_i^2}} + \frac{2(h-z)}{\sqrt{(h-z)^2 + R_i^2}} - \frac{z}{\sqrt{z^2 + R_o^2}} - \frac{(h-z)}{\sqrt{(h-z)^2 + R_o^2}} - 2 \right] \quad (\text{A.15})$$

Let's look at two limiting cases. For $z, h \gg R$, we get that this is trivially 0; in the case that z -partition is not meaningful, r -partition is also not meaningful (assuming we're just summing over detector surfaces). In the case that $R \gg z, h$ we find

$$\Delta\Omega \approx 4\pi \left[h \left(\frac{1}{R_i} - \frac{1}{2R_0} \right) - 1 \right] \quad (\text{A.16})$$

so the r partition is independent of z , and large as long as a large fraction of the total energy is collected and the solid angle varies based on position. I leave the application of this to the full partition as a function of r and z as future work.

Appendix B

TES Linear Parameterization

In describing the small-signal response of the TES, we normally linearize the resistance dependence on temperature and current as

$$R(I, T) \approx R_0 \left[1 + \alpha \frac{(T - T_c)}{T_c} + \beta \frac{(I - I_q)}{I_q} \right]$$

where

$$\alpha = \left. \frac{T_0}{R_0} \frac{\partial R}{\partial T} \right|_{I_0}, \quad \beta = \left. \frac{I_0}{R_0} \frac{\partial R}{\partial I} \right|_{T_0}$$

For our TES simulations however, we take critical current effects into account, and model the transition in a highly non-linear way away from the the bias point using the function

$$R(I, T) = \frac{R_{max} - R_{min}}{2} \left(1 + \tanh \left[\frac{T - T_c (1 - I/I_c)^{n_{sc}}}{T_w} \right] \right)$$

By evaluating the derivatives of this function, we can therefore relate the phenomenological parameters α and β with the physical parameters T_w and I_c , given the bias condition I_0 and transition temperature T_c .

We evaluate the temperature derivative first:

$$\frac{2}{R_{max} - R_{min}} \frac{\partial R}{\partial T} = \frac{\partial}{\partial T} \tanh \left[\frac{T - T_c (1 - I/I_c)^{n_{sc}}}{T_w} \right] \quad (\text{B.1})$$

$$= \left(1 - \tanh^2 \left[\frac{T - T_c (1 - I/I_c)^{n_{sc}}}{T_w} \right] \right) \frac{1}{T_w} \quad (\text{B.2})$$

$$\left. \frac{\partial R}{\partial T} \right|_{I_0, T_0} = \left(1 - \tanh^2 \left[\frac{T_0 - T_c (1 - I_0/I_c)^{n_{sc}}}{T_w} \right] \right) \frac{R_{max} - R_{min}}{2T_w} \quad (\text{B.3})$$

We can use the fact that

$$R_0 = \frac{R_{max} - R_{min}}{2} \left(1 + \tanh \left[\frac{T_0 - T_c (1 - I_0/I_c)^{n_{sc}}}{T_w} \right] \right) \quad (\text{B.4})$$

$$\frac{2R_0}{R_{max} - R_{min}} - 1 = \tanh \left[\frac{T_0 - T_c (1 - I_0/I_c)^{n_{sc}}}{T_w} \right] \quad (\text{B.5})$$

$$\left(\frac{2R_0}{R_{max} - R_{min}} - 1 \right)^2 = \tanh^2 \left[\frac{T_0 - T_c (1 - I_0/I_c)^{n_{sc}}}{T_w} \right] \quad (\text{B.6})$$

$$1 - \frac{4R_0}{R_{max} - R_{min}} + \frac{4R_0^2}{(R_{max} - R_{min})^2} = \tanh^2 \left[\frac{T_0 - T_c (1 - I_0/I_c)^{n_{sc}}}{T_w} \right] \quad (\text{B.7})$$

to simplify this expression:

$$\left. \frac{\partial R}{\partial T} \right|_{I_0, T_0} = \left(\frac{4R_0}{R_{max} - R_{min}} - \frac{4R_0^2}{(R_{max} - R_{min})^2} \right) \frac{R_{max} - R_{min}}{2T_w} \quad (\text{B.8})$$

$$= \frac{2R_0}{T_w} \left(1 - \frac{R_0}{R_{max} - R_{min}} \right) \quad (\text{B.9})$$

$$\alpha = \frac{2T_0}{T_w} \left(1 - \frac{R_0}{R_{max} - R_{min}} \right) \quad (\text{B.10})$$

This is very close to the naive estimate one might make, namely that if T_w is half the transition width, then the slope of this line is

$$\frac{T_0}{R_0} \frac{R_{max} - R_{min}}{2T_w} \quad (\text{B.11})$$

where there are correction factors due to the nature of the tanh function, where the slope is actually larger in the middle and smaller on the ends rather than constant through the transition. If we set $R_{min} = 0$, $R_{max} = R_n$, $R_0 = aR_n$ and $T_0 = T_c$, then we find

$$\alpha = \frac{T_c(1 - a)}{T_w/2} \quad (\text{B.12})$$

Thus for the bias condition where $a = 1/3$, $T_c = 45mK$, and $\alpha = 100$, we find

$$T_w = \frac{T_c(1 - a)}{\alpha/2} = \frac{0.045 * (2/3)}{50} = 6 * 10^{-4} \quad (\text{B.13})$$

When we plug this transition width into our naive estimate, it gives $\alpha \sim 112$, which is a nice sanity check.

Repeating this procedure for current, we can find an equivalent expression for β as well:

$$\frac{2}{R_{max} - R_{min}} \frac{\partial R}{\partial I} = \frac{\partial}{\partial I} \tanh \left[\frac{T - T_c (1 - I/I_c)^{n_{sc}}}{T_w} \right] \quad (\text{B.14})$$

$$= \left(1 - \tanh^2 \left[\frac{T - T_c (1 - I/I_c)^{n_{sc}}}{T_w} \right] \right) \frac{-n_{sc} T_c (1 - I/I_c)^{n_{sc}-1} - 1}{T_w} \frac{1}{I_c} \quad (\text{B.15})$$

$$= \left(1 - \tanh^2 \left[\frac{T - T_c (1 - I/I_c)^{n_{sc}}}{T_w} \right] \right) \frac{n_{sc} T_c (1 - I/I_c)^{n_{sc}}}{T_w (I_c - I)} \quad (\text{B.16})$$

$$\left. \frac{\partial R}{\partial I} \right|_{I_0, T_0} = \left(\frac{4R_0}{R_{max} - R_{min}} - \frac{4R_0^2}{(R_{max} - R_{min})^2} \right) \frac{n_{sc} T_c (1 - I_0/I_c)^{n_{sc}}}{T_w (I_c - I_0)} \frac{R_{max} - R_{min}}{2} \quad (\text{B.17})$$

$$\left. \frac{\partial R}{\partial I} \right|_{I_0, T_0} = \frac{2n_{sc} R_0 T_c (1 - I_0/I_c)^{n_{sc}}}{T_w (I_c - I_0)} \left(1 - \frac{R_0}{R_{max} - R_{min}} \right) \quad (\text{B.18})$$

$$\beta = \frac{2n_{sc} I_0 T_c (1 - I_0/I_c)^{n_{sc}}}{T_w (I_c - I_0)} \left(1 - \frac{R_0}{R_{max} - R_{min}} \right) \quad (\text{B.19})$$

$$\beta = n_{sc} (1 - I_0/I_c)^{n_{sc}-1} \frac{I_0}{I_c} \frac{T_c}{T_0} \alpha \quad (\text{B.20})$$

and we find that β is actually dependent on α for finite values of the critical current. Making the same simplifying assumptions as for our simplification of α and using $I_0 = bI_c$, $n_{sc} \sim 2/3$, we find

$$\beta = \frac{2b}{3(1-b)^{1/3}} \alpha \quad (\text{B.21})$$

And for $\alpha = 100$ and $b \sim 0.01$, we find $\beta \sim 0.6$. Conversely, for $\alpha = 100$, which corresponds to the T_c and other conditions as described above, we can calculate the predicted snapping point for a given β as a function of bias current, assuming small b :

$$b \approx \frac{3\beta}{2\alpha} \rightarrow I_c \approx \frac{2\alpha}{3\beta} I_0 \sim \frac{60}{\beta} I_0 = \frac{60V_{bias}}{R_0} \quad (\text{B.22})$$

This gives us a way to constrain the linear behavior of the TES to small signals at the operating point as a function of global measurements, without having to do any curve fitting, and also shows us why the bias condition is so hard to solve for.

Appendix C

Phase Separated TES Dynamics

In this appendix I explore a discrete two-block model of TES dynamics valid for the weakly phase-separated limit. This should be seen as complementary to the Fourier version done in Matt Pyle's thesis [83], and can be compared to the two and three-block TES models discussed at length in Ref [70] for the case that both blocks are under the influence of electro-thermal feedback. The stability conditions and time constants found here are also comparable to some of those discussed in [50], though my solutions may differ in terminology, and I believe that some of these results are unique. I provide this as an exercise in TES dynamics and a starting point for those looking to dig further into TES dynamics in the weakly coupled limit.

C.1 Setup of Differential Equations

In the very weakly phase-separated limit, we can consider the TES broken up into finite sub-sections, with a toy model consisting of only two TES sections operated in the small-signal limit. Generically speaking, we can consider the TES a series of individual, thermally-homogeneous TES units connected both conductively and thermally through a known thermal conductance. The power flow differential equation will have the generic form

$$\frac{dE_i}{dt} = C(T_i) \frac{dT_i}{dt} = I^2 R_i(T_i, I) - K_l(2T_i^n - T_{i+1}^n - T_{i-1}^n) - K_b(T_i^n - T_b^n) + \delta P_i \quad (\text{C.1})$$

where K_l is the thermal link between TES segments, and K_b is the thermal link to the bath. We couple these equations through the current differential equation

$$L \frac{dI}{dt} = V_b - IR_L - I \sum_{i=1}^N R_i(T_i, I) \quad (\text{C.2})$$

where R_L and L are the parasitic inductances and resistances in the bias loop.

This is an exact set of equations, and in the past we usually solve these numerically; indeed all of our phase separation studies have taken these as non-analytic and stuck to numerical simulations. This is most likely the right approach in the highly phase-separated regime, however in the very weakly phase separated regime, and the small signal limit, we can make two assumptions:

- Any thermal gradient is eliminated through the internal conductance of the TES
- All thermal conductances to the bath are equal for the TES segments, and the same is true for the conductance between segments, though these respectively need not and will in principle not be equal.
- The power-law thermal conductance has the same exponential dependence both between TES and to the bath, so that we do not also have to introduce the terms n_i in the above equations.

The practical result of these assumptions is that we can write the following:

$$R_{i,0} = R_s = \frac{R_0}{N} = R_0 \frac{l_s}{l_{TES}} = \frac{\rho l_s}{w_{TES} h_{TES}} = \lambda_{TES} l_s \quad (\text{C.3})$$

$$C_i = C_s = \frac{C_0}{N} = C_0 \frac{l_s}{l_{TES}} = c_{TES} l_s w_{TES} h_{TES} = \xi_{TES} l_s \quad (\text{C.4})$$

$$K_i = K_l = N K_0 = \frac{k_0}{l_s} \quad (\text{C.5})$$

$$K_{b,i} = K_b = \frac{K_{b,0}}{N} = k_{b,0} l_s \quad (\text{C.6})$$

$$(\text{C.7})$$

which also tells us that

$$\Delta R_i = R_i - R_{i,0} = (\lambda_i(I, T) - \lambda_0) l_s \quad (\text{C.8})$$

These substitutions allow us both to relate the properties of the small segments to the larger segments, as well as ensure that the limit $N \rightarrow \infty$ is convergent. It will also simplify and increase the utility of the solutions later on, as the operational regime can just be specified by the TES and bath thermal conductivity. The previous solutions implicitly used $K_{TES} \rightarrow \infty$, so all we're doing here is relaxing that assumption.

Before writing out the full matrix form, and doing all our substitutions, let's linearize these terms in the small signal limit. We obtain the following conductance forms:

$$K_l(2T_i^n - T_{i-1}^n - T_{i+1}^n) = K_l(2T_{i,0}^n - T_{i-1,0}^n - T_{i+1,0}^n) + nK_l(2T_{i,0}\Delta T_i - T_{i-1,0}\Delta T_{i-1} - T_{i+1,0}\Delta T_{i+1}) \quad (\text{C.9})$$

$$K_b(T_i^n - T_b^n) = K_b(T_{i,0}^n - T_b^n) + nK_b T_{i,0} \Delta T_i \quad (\text{C.10})$$

and we use the McCammon logarithmic slopes α and β to linearize resistance as

$$R_i(I, T_i) = R_s + \alpha_i \frac{R_s}{T_0} \Delta T_i + \beta_i \frac{R_s}{I_0} \Delta I \quad (\text{C.11})$$

In principle, the film properties could vary, and α and β would vary as well, but here I will assume that all α and β are the same. We also need to linearize the joule heating term:

$$I^2 R_i \approx I_0^2 R_s + 2I_0 R_s \Delta I + I_0^2 \Delta R = I_0^2 R_s + (2 + \beta) I_0 R_s \Delta I + \alpha \frac{I_0^2 R_s}{T_0} \Delta T_i$$

We also make the replacements $T_i \rightarrow T_{i,0} + \Delta T_i$ and $I \rightarrow I_0 + \delta I$. We can separate the steady state equations (for I_0 and T_0) by setting all time derivatives, and Δ terms, to 0; then the I_0 and R_0 terms cancel from the rest of the differential equation. We get the differential equations

$$C_s \frac{d(\Delta T_i)}{dt} = (2 + \beta) I_0 R_s \Delta I + \alpha \frac{I_0^2 R_s}{T_0} \Delta T_i - nK_l (2T_{i,0} \Delta T_i - T_{i-1,0} \Delta T_{i-1} - T_{i+1,0} \Delta T_{i+1}) - nK_b T_{i,0} \Delta T_i + \delta P_i \quad (\text{C.12})$$

$$= (2 + \beta) I_0 R_s \Delta I + \left(\alpha \frac{I_0^2 R_s}{T_0} - 2nK_l T_{i,0} - nK_b T_{i,0} \right) \Delta T_i + nK_l (T_{i-1,0} \Delta T_{i-1} + T_{i+1,0} \Delta T_{i+1}) + \delta P_i \quad (\text{C.13})$$

$$= (2 + \beta) I_0 R_s \Delta I + \left(\alpha \frac{I_0^2 R_s}{T_0} - 2g_{l,i} - g_{b,i} \right) \Delta T_i + g_{l,i-1} \Delta T_{i-1} + g_{l,i+1} \Delta T_{i+1} + \delta P_i \quad (\text{C.14})$$

$$= (2 + \beta) I_0 R_s \Delta I + (2g_{l,i} + g_{b,i}) (\mathcal{L}_i - 1) \Delta T_i + g_{l,i-1} \Delta T_{i-1} + g_{l,i+1} \Delta T_{i+1} + \delta P_i \quad (\text{C.15})$$

where

$$g_{l,i} = nK_l T_{0,i} = nK_{l,0} T_{0,i} N \quad (\text{C.16})$$

$$g_{b,i} = nK_b T_{0,i} = nK_{b,0} T_{0,i} N \quad (\text{C.17})$$

$$\mathcal{L}_i = \alpha \frac{I_0^2 R_s}{T_0 (2g_{l,i} + g_{b,i})} = \frac{g_{b,i}}{2g_{l,i} + g_{b,i}} \mathcal{L}_b \quad (\text{C.18})$$

This gets further simplified when we apply the assumption of thermal equilibrium initially, such that

$$T_{0,i} = T_0 \rightarrow g_{b,i} = g_b, \quad g_{l,i} = g_l \rightarrow \mathcal{L}_i = \mathcal{L}_{2lb}$$

and we obtain the final differential equation

$$C_s \frac{d(\Delta T_i)}{dt} = (2 + \beta) I_0 R_s \Delta I + (2g_l + g_b) (\mathcal{L}_{2lb} - 1) \Delta T_i + g_l (\Delta T_{i-1} + \Delta T_{i+1}) + \delta P_i \quad (\text{C.19})$$

note that the end cases have $2g_l \rightarrow g_l$ in the diagonal terms, and $\mathcal{L}_{2lb} \rightarrow \mathcal{L}_{lb}$.

The current differential equation becomes

$$L \frac{d\Delta I}{dt} = -\Delta I(R_L + (1 + \beta)NR_s) - \frac{\alpha I_0 R_s}{T_0} \sum_{i=1}^N \Delta T_i + \delta V \quad (\text{C.20})$$

$$= -\Delta I(R_L + (1 + \beta)NR_s) - \frac{g_b \mathcal{L}_b}{I_0} \sum_{i=1}^N \Delta T_i + \delta V \quad (\text{C.21})$$

and our system of equations can be formulated as a square $N+1$ dimension matrix, with eigenvectors composed of the TES block temperatures plus the series TES current. In the next section, we will first solve some special cases to explore the behavior of this system.

C.2 Two-Block Symmetric Solution

The simplest system we can solve is the two-block system run without electrothermal feedback. This allows us to set δI to 0, and our system is just the two-dimensional matrix

$$C_s \frac{d}{dt} \begin{bmatrix} \Delta T_1 \\ \Delta T_2 \end{bmatrix} = \begin{bmatrix} -g_l - g_b & g_l \\ g_l & -g_l - g_b \end{bmatrix} \begin{bmatrix} \Delta T_1 \\ \Delta T_2 \end{bmatrix} + \begin{bmatrix} \delta P_1 \\ \delta P_2 \end{bmatrix} \quad (\text{C.22})$$

$$= \left(g_l \begin{bmatrix} -1 & 1 \\ 1 & -1 \end{bmatrix} - g_b \mathcal{I} \right) \begin{bmatrix} \Delta T_1 \\ \Delta T_2 \end{bmatrix} + \begin{bmatrix} \delta P_1 \\ \delta P_2 \end{bmatrix} \quad (\text{C.23})$$

$$\frac{d}{dt} \begin{bmatrix} \Delta T_1 \\ \Delta T_2 \end{bmatrix} = - \left(\frac{g_l}{C_s} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} + \frac{g_b}{C_s} \mathcal{I} \right) \begin{bmatrix} \Delta T_1 \\ \Delta T_2 \end{bmatrix} + C_s^{-1} \begin{bmatrix} \delta P_1 \\ \delta P_2 \end{bmatrix} \quad (\text{C.24})$$

$$= - \left(\frac{4}{\tau_l} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} + \tau^{-1} \mathcal{I} \right) \begin{bmatrix} \Delta T_1 \\ \Delta T_2 \end{bmatrix} + C_s^{-1} \begin{bmatrix} \delta P_1 \\ \delta P_2 \end{bmatrix} \quad (\text{C.25})$$

Here, we start to see how the parameterization, and splitting of the TES into parts, will affect the final solution. On the right we obtain the thermal time constant of the TES, τ , given that

$$\frac{C_s}{g_b} = \frac{C/N}{G_b/N} = \frac{C}{G} = \tau$$

On the other hand, we end up with a time constant which does depend on the number of TES segments:

$$\frac{C_s}{g_l} = \frac{C/N}{G_l N} = \frac{C}{G_l N^2} = \frac{\tau_l}{N^2}$$

Here, τ_L is a measured property of the TES, and depends on the TES geometry. This differs from τ_b , in which geometric factors largely cancel. In addition, τ_l alone controls the relative thermal interaction between adjacent TESs. This tells us that τ_l is the key to understanding whether phase

separated dynamics differ significantly from the 0th order solution, and how many TES segments will be necessary to simulate a given TES in the frequency range of interest.

Solving this in the time domain, we obtain the eigenvalues

$$\lambda = \tau^{-1}, \quad 8\tau_l^{-1} + \tau^{-1}$$

and thus the solutions

$$\begin{bmatrix} \Delta T_1 \\ \Delta T_2 \end{bmatrix} = Ae^{-t/\tau} \begin{bmatrix} 1 \\ 1 \end{bmatrix} + Be^{-t(\tau^{-1} + 8/\tau_l)} \begin{bmatrix} 1 \\ -1 \end{bmatrix} \quad (\text{C.26})$$

$$= Ae^{-t/\tau} \begin{bmatrix} 1 + Be^{-t/\tau_p} \\ 1 - Be^{-t/\tau_p} \end{bmatrix} \quad (\text{C.27})$$

where here, $\tau_p = \tau_l/8$. Applying the initial conditions that $\Delta T_1 = \frac{E_1}{C_s}$ and $\Delta T_2 = \frac{E_2}{C_s}$, we find that

$$A = \frac{E_1 + E_2}{2C_s} = \frac{E_{tot}}{2C_s} \quad (\text{C.28})$$

$$B = \frac{E_1 - E_2}{E_1 + E_2} = \frac{\Delta E_{12}}{E_{tot}} \quad (\text{C.29})$$

and we have the final solution

$$\begin{bmatrix} \Delta T_1 \\ \Delta T_2 \end{bmatrix} = \frac{e^{-t/\tau}}{2C_s} \begin{bmatrix} E_{tot} + \Delta E_{12}e^{-t/\tau_p} \\ E_{tot} - \Delta E_{12}e^{-t/\tau_p} \end{bmatrix} \quad (\text{C.30})$$

This obviously re-distributed the total energy to both TESs, and allows it to slowly leak back to the bath. We can also see that this approaches the phase-uniform regime in the limit that the second term goes to 0 much quicker than the first term; the change of variables used to compare these will have $t' = t/\tau$, which tells us that the relevant limit is

$$\tau_p = \frac{\tau_l}{8} \gg \tau$$

For practical purposes, let's say that τ is $\gg 1 \mu\text{s}$, and we thus plan to digitize only about $0.5 \mu\text{s}$. We want this second frequency component to be irrelevant on this time-scale, and thus

$$1\mu\text{s} > 5\tau_p = \frac{5}{8}\tau_l \approx \frac{\tau_l}{2}$$

We can make this argument because any reduction in the natural thermal timescale should also reduce the thermal equilibrium timescale. This also tells us that for a larger τ_l (less thermal conductivity between TES blocks), we will need to account for more phase separation, given that to get back under this limit, we can just increase N , as $\tau_p \propto N^{-2}$. Taken another way, we find that τ_l is

inversely proportional to the total length of the TES squared:

$$t_{dig} > \frac{10}{N^2} \tau_l$$

Solving this in the frequency domain is unnecessary; we already know relevant time constants. What is more useful is to undo our assumption of no voltage bias, and compute the complex impedance of the two-element model. We now have the system

$$\frac{d}{dt} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \end{bmatrix} = \begin{bmatrix} -\frac{R_L + 2(1+\beta)R_s}{L} & -\frac{g_b \mathcal{L}_b}{I_0 L} & -\frac{g_b \mathcal{L}_b}{I_0 L} \\ \frac{(2+\beta)I_0 R_s}{C_s} & \frac{(g_l + g_b)}{C_s} (\mathcal{L}_{lb} - 1) & \frac{g_l}{C_s} \\ \frac{(2+\beta)I_0 R_s}{C_s} & \frac{g_l}{C_s} & \frac{(g_l + g_b)}{C_s} (\mathcal{L}_{lb} - 1) \end{bmatrix} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \end{bmatrix} + \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P_1}{C_s} \\ \frac{\delta P_2}{C_s} \end{bmatrix} \quad (\text{C.31})$$

$$= \begin{bmatrix} -\omega_{LR} & -\frac{g_b \mathcal{L}_b}{I_0 L} & -\frac{g_b \mathcal{L}_b}{I_0 L} \\ \frac{(2+\beta)I_0 R_s}{C_s} & -\omega_{lb} & \omega_l \\ \frac{(2+\beta)I_0 R_s}{C_s} & \omega_l & -\omega_{lb} \end{bmatrix} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \end{bmatrix} + \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P_1}{C_s} \\ \frac{\delta P_2}{C_s} \end{bmatrix} \quad (\text{C.32})$$

We can rotate to a diagonalized temperature basis to simplify the solution, using the unitary transformation

$$T = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix}$$

which for the input variables X gives

$$X' = TX \rightarrow \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \end{bmatrix} = \begin{bmatrix} \Delta I \\ \frac{\Delta T_1 + \Delta T_2}{\sqrt{2}} \\ \frac{\Delta T_1 - \Delta T_2}{\sqrt{2}} \end{bmatrix} = X'$$

and for the coupling matrix S , we find

$$S' = TST^T \quad (\text{C.33})$$

$$= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} -\omega_{LR} & -\frac{g_b \mathcal{L}_b}{I_0 L} & -\frac{g_b \mathcal{L}_b}{I_0 L} \\ \frac{(2+\beta)I_0 R_s}{C_s} & -\omega_{lb} & \omega_l \\ \frac{(2+\beta)I_0 R_s}{C_s} & \omega_l & -\omega_{lb} \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} \quad (\text{C.34})$$

$$= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} -\omega_{LR} & -\frac{\sqrt{2}g_b \mathcal{L}_b}{I_0 L} & 0 \\ \frac{(2+\beta)I_0 R_s}{C_s} & \frac{\omega_l - \omega_{lb}}{\sqrt{2}} & -\frac{\omega_l + \omega_{lb}}{\sqrt{2}} \\ \frac{(2+\beta)I_0 R_s}{C_s} & \frac{\omega_l - \omega_{lb}}{\sqrt{2}} & \frac{\omega_l + \omega_{lb}}{\sqrt{2}} \end{bmatrix} \quad (\text{C.35})$$

$$= \begin{bmatrix} -\omega_{LR} & -\frac{\sqrt{2}g_b \mathcal{L}_b}{I_0 L} & 0 \\ \frac{\sqrt{2}(2+\beta)I_0 R_s}{C_s} & \omega_l - \omega_{lb} & 0 \\ 0 & 0 & -(\omega_l + \omega_{lb}) \end{bmatrix} \quad (\text{C.36})$$

Sorting out spurious factors of $\sqrt{2}$ gives the system

$$\frac{d}{dt} \begin{bmatrix} \Delta I \\ \frac{1}{2}(\Delta T_1 + \Delta T_2) \\ \frac{1}{2}(\Delta T_1 - \Delta T_2) \end{bmatrix} = \begin{bmatrix} -\omega_{LR} & -\frac{G\mathcal{L}_b}{I_0 L} & 0 \\ \frac{(2+\beta)I_0 R_0}{C} & \omega_l - \omega_{lb} & 0 \\ 0 & 0 & -\omega_l - \omega_{lb} \end{bmatrix} \begin{bmatrix} \Delta I \\ \frac{1}{2}(\Delta T_1 + \Delta T_2) \\ \frac{1}{2}(\Delta T_1 - \Delta T_2) \end{bmatrix} + \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P_1 + \delta P_2}{C} \\ \frac{\delta P_1 - \delta P_2}{C} \end{bmatrix} \quad (\text{C.37})$$

Recognizing that

$$\omega_l - \omega_{lb} = \frac{1}{C_s} (g_l - (g_l + g_b)(1 - \mathcal{L}_{lb})) \quad (\text{C.38})$$

$$= \frac{1}{C_s} (g_l - g_l - g_b + g_b \mathcal{L}) \quad (\text{C.39})$$

$$= \frac{g_b}{C_s} (\mathcal{L} - 1) \quad (\text{C.40})$$

$$= \frac{\mathcal{L} - 1}{\tau} = \omega_{eff} \quad (\text{C.41})$$

$$-\omega_l - \omega_{lb} = \omega_{eff} - 2\omega_l \quad (\text{C.42})$$

We can write the final system as

$$\frac{d}{dt} \begin{bmatrix} \Delta I \\ \frac{T_1 + T_2}{2} - T_0 \\ \frac{1}{2}(T_1 - T_2) \end{bmatrix} = \begin{bmatrix} -\omega_{LR} & -\frac{G\mathcal{L}_b}{I_0 L} & 0 \\ \frac{(2+\beta)I_0 R_0}{C} & \omega_{eff} & 0 \\ 0 & 0 & \omega_{eff} - 2\omega_l \end{bmatrix} \begin{bmatrix} \Delta I \\ \frac{T_1 + T_2}{2} - T_0 \\ \frac{1}{2}(T_1 - T_2) \end{bmatrix} + \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P_1 + \delta P_2}{C} \\ \frac{\delta P_1 - \delta P_2}{C} \end{bmatrix} \quad (\text{C.43})$$

The upper left-hand part of the matrix is the non-phase separated equation, exact in the limit that

$T_1 = T_2$, which we can see will occur when

$$2\omega_{lb} \gg \omega_{eff} \quad (\text{C.44})$$

$$\frac{2g_l}{C_s} \gg \frac{\mathcal{L} - 1}{\tau} \quad (\text{C.45})$$

$$\frac{8G_l}{C} \gg \frac{\mathcal{L} - 1}{\tau} \quad (\text{C.46})$$

$$G_l \gg \frac{G_b(\mathcal{L} - 1)}{8} \quad (\text{C.47})$$

These equations are totally decoupled, so in the two-block system, we see no effect from phase separation if we assume that the TES oscillates about fixed points of the same temperature. Next we will explore what happens when we relax this constraint as well.

C.3 Two Block Stable Asymmetric Solution

The previous solution assumed that the two blocks will reach thermal equilibrium, so that fluctuations are about the case that the blocks are initially the same temperature. We should explore how realistic this assumption is

Let's introduce a small perturbation about equilibrium to see whether there are any stable points away from the nominal assumption of initial equal temperature in the TES. When we do see phase separation in the TES, we have some fluctuation away from equal temperature. This leads to a slightly higher resistance in one segment than the other, and for small changes resistance and power changes are linear, so we can approximate them as a symmetric deviation from the initial temperature. We can thus write

$$T_1 = T_0 + \frac{\Delta T_0}{2} = T_0 \left(1 + \frac{\epsilon}{2}\right) \quad (\text{C.48})$$

$$T_2 = T_0 - \frac{\Delta T_0}{2} = T_0 \left(1 - \frac{\epsilon}{2}\right) \quad (\text{C.49})$$

where

$$\epsilon = \frac{\Delta T_0}{T_0}$$

and we'll assume $\epsilon \ll 1$.

We're only concerned with finding stable values for T_0 and ΔT_0 , so we'll only linearize in these terms and absorb any changes in current into the definition of I_0 . If when the blocks are both at T_0

they have resistance R_0 , then we find that

$$R(I_0, T) = R'_0 + \frac{\partial R}{\partial T}(T - T_0) + O[(T - T_0)^2] \quad (\text{C.50})$$

$$= R'_0 \pm \alpha \frac{R'_0}{T_0} \Delta T_0 + O[(T - T_0)^2] \quad (\text{C.51})$$

$$= R'_0 \pm \alpha \epsilon R'_0 + O[\epsilon^2] \quad (\text{C.52})$$

$$= R'_0(1 \pm \alpha \epsilon) + O[\epsilon^2] \quad (\text{C.53})$$

where the upper sign refers to R_1 and the lower to R_2 , and $R'_0 = R_0/2$, where R_0 is the series resistance of the entire TES when not phase separated at I_0, T_0 . There is good justification to ignore the second term from our model of the transition curve, which at its midpoint has a second order derivative of 0. In any case, we'll work in the small signal limit, but it's an interesting note that this does not necessarily require ϵ to be small to be a good approximation.

Let's see how this affects the current equation first:

$$L \frac{dI}{dt} = V - I_0(R_L + R_1 + R_2) \quad (\text{C.54})$$

$$= V - I_0(R_L + \frac{R_0}{2}(1 + \alpha \epsilon) + \frac{R_0}{2}(1 - \alpha \epsilon)) \quad (\text{C.55})$$

$$= V - I_0(R_L + R_0) \quad (\text{C.56})$$

so the current equation is identical to the non-phase separated case while ϵ remains small, as we should have expected. Now we can find the forms for T_0 and ΔT_0 to determine whether ϵ can remain small in steady state.

For the derivative sides, we get

$$\frac{dT_i}{dt} = \frac{dT_0}{dt} \pm \frac{d\Delta T_0}{dt} \quad (\text{C.57})$$

and for the power balance side, we can expand each term.

$$K'_l(T_1^2 - T_2^2) = K'_l(T_1 + T_2)(T_1 - T_2) = 2K'_l T_0(\Delta T_0) = 2K'_l T_0^2 \epsilon \quad (\text{C.58})$$

$$K'_b(T_i^n - T_b^n) = K'_b \left(1 - \left(\frac{T_b}{T_i} \right)^n \right) T_i^n = K'_b \left((1 \pm \epsilon)^n - \left(\frac{T_b}{T_0} \right)^n \right) T_0^n \quad (\text{C.59})$$

which gives the equations

$$C' \frac{dT_0}{dt} \pm C' \frac{d\Delta T_0}{dt} = I_0^2 R'_0(1 \pm \alpha \epsilon) - K'_b \left((1 \pm \epsilon)^n - \left(\frac{T_b}{T_0} \right)^n \right) T_0^n \mp 2K'_l T_0^2 \epsilon \quad (\text{C.60})$$

If we take the sum and difference of these equations, we have the more straightforward (and familiar)

equations

$$2C' \frac{dT_0}{dt} = 2I_0^2 R_0' - K_b' \left((1 + \epsilon)^n + (1 - \epsilon)^n - 2 \left(\frac{T_b}{T_0} \right)^n \right) T_0^n \quad (\text{C.61})$$

$$2C' \frac{d\Delta T_0}{dt} = 2I_0^2 R_0' \epsilon \alpha - K_b' \left((1 + \epsilon)^n - (1 - \epsilon)^n \right) T_0^n - 4K_l' T_0^2 \epsilon \quad (\text{C.62})$$

We can find exact expressions for the sum and difference of the $(1 + \pm \epsilon)$ terms; hereafter I will use $n = 5$ for simplicity:

$$(1 + \epsilon)^5 + (1 - \epsilon)^5 = 2 + 5\epsilon^2 + \frac{5}{8}\epsilon^4 = \frac{1}{8} [16 + 40\epsilon^2 + 5\epsilon^4] \quad (\text{C.63})$$

$$(1 + \epsilon)^5 - (1 - \epsilon)^5 = 5\epsilon + \frac{5}{2}\epsilon^3 + \frac{1}{16}\epsilon^5 = \frac{\epsilon}{16} [80 + 40\epsilon^2 + \epsilon^4] \quad (\text{C.64})$$

substitution of these forms into the original equations gives

$$2C' \frac{dT_0}{dt} = 2I_0^2 R_0' - 2K_b' \left(1 + \frac{5}{2}\epsilon^2 + \frac{5}{16}\epsilon^4 - \left(\frac{T_b}{T_0} \right)^5 \right) T_0^5 \quad (\text{C.65})$$

$$2C' \frac{d\Delta T_0}{dt} = 2I_0^2 R_0' \epsilon \alpha - 5\epsilon K_b' \left(1 + \frac{1}{2}\epsilon^2 + \frac{1}{80}\epsilon^4 \right) T_0^5 - 4K_l' T_0^2 \epsilon \quad (\text{C.66})$$

Recognizing that $K_b' = K_b/2$, $K_l' = 2K_l$, $C' = C/2$, and $R_0' = R_0/2$, we have

$$C \frac{dT_0}{dt} = I_0^2 R_0 - K_b \left(1 + \frac{5}{2}\epsilon^2 + \frac{5}{16}\epsilon^4 - \left(\frac{T_b}{T_0} \right)^5 \right) T_0^5 \quad (\text{C.67})$$

$$C \frac{d\Delta T_0}{dt} = I_0^2 R_0 \epsilon \alpha - \frac{5}{2}\epsilon K_b \left(1 + \frac{1}{2}\epsilon^2 + \frac{1}{80}\epsilon^4 \right) T_0^5 - 8K_l T_0^2 \epsilon \quad (\text{C.68})$$

We can re-write the second equation in a more illuminating form to show how we can determine stable points for ϵ :

$$C \frac{d\Delta T_0}{dt} = \left[I_0^2 R_0 \alpha - \frac{5}{2} K_b \left(1 + \frac{1}{2}\epsilon^2 + \frac{1}{80}\epsilon^4 \right) T_0^5 - 8K_l T_0^2 \right] \epsilon \quad (\text{C.69})$$

We recognize that because of our re-parameterization, we have

$$\frac{d\Delta T_0}{dt} = \epsilon \frac{dT_0}{dt} + T_0 \frac{d\epsilon}{dt}$$

so we have

$$CT_0 \frac{d\epsilon}{dt} = \left[I_0^2 R_0 \alpha - \frac{5}{2} K_b \left(1 + \frac{1}{2}\epsilon^2 + \frac{1}{80}\epsilon^4 \right) T_0^5 - 8K_l T_0^2 - C \frac{dT_0}{dt} \right] \epsilon \quad (\text{C.70})$$

which gives us

$$CT_0 \frac{d\epsilon}{dt} = \left[I_0^2 R_0 (\alpha - 1) - K_b \left(\frac{3}{2} - \frac{5}{4} \epsilon^2 - \frac{9}{32} \epsilon^4 + \left(\frac{T_b}{T_0} \right)^5 \right) T_0^5 - 8K_l T_0^2 \right] \epsilon \quad (\text{C.71})$$

$$= [I_0^2 R_0 (\alpha - 1) - \xi K_b T_0^5 - 8K_l T_0^2] \epsilon \quad (\text{C.72})$$

where

$$\xi = \left(\frac{3}{2} - \frac{5}{4} \epsilon^2 - \frac{9}{32} \epsilon^4 + \left(\frac{T_b}{T_0} \right)^5 \right)$$

can only take positive values $\sim 1 - 5/2$ depending on ϵ , α , T_0 and T_b . This gives us the stability criterion for any dynamic solution for T_0

$$I_0^2 R_0 (\alpha - 1) < \xi K_b T_0^5 + 8K_l T_0^2$$

Stable points will occur where the term inside the brackets is less than 0, but unlike the equation for T_0 , it does not permit stable points aside from $\epsilon = 0$. So we've found two things, a stability criterion based on the thermal conductances and operating temperatures, and the fact that this two-block model does not permit long-lived temperature differences between the blocks for any values of these conductances.

If we now solve for the stable points by setting $\frac{dT_0}{dt} = 0$, we find

$$I_0^2 R_0 \approx \xi_0 K_b T_0^5$$

which modifies the stability condition to be

$$\alpha < \frac{\xi K_b T_0^5 + 8K_l T_0^2}{\xi_0 K_b T_0^5} + 1 = \frac{8K_l}{\xi_0 K_b T_0^3} + \frac{\xi}{\xi_0} + 1$$

which is similar to that derived in Matt Pyle's thesis, though we arrived at the limit through a discretized model, rather than through a continuous position-dependent model. This essentially tells us that the higher our α , the larger our internal TES conductance needs to be in order to keep the TES in thermal equilibrium.

In addition, we've discovered that the time constant for recovery from thermal fluctuations obeys the bound

$$\tau_\epsilon \leq \left(\frac{K_b T_0^3}{2C_W} + \frac{8K_l}{C_W} \right)^{-1} = \frac{8K_l}{C_W} \left[\frac{K_b T_0^3}{16K_l} + 1 \right]^{-1}$$

so for signals below the frequency corresponding to this time constant, we can assume a phase-uniform TES even if at high frequency the TES acts as a two-block system and admits a small temperature difference.

Dynamics Let's start with our finite-current model in the original form, and remove the assumption of equal initial temperature.

$$\frac{d}{dt} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \end{bmatrix} = \begin{bmatrix} -\frac{R_L + (1+\beta)(R_{0,2} + R_{0,1})}{L} & -\frac{g_{b1}\mathcal{L}_{b1}}{I_0 L} & -\frac{g_{b2}\mathcal{L}_{b2}}{I_0 L} \\ \frac{(2+\beta)I_0 R_{0,1}}{C_s} & \frac{g_{b1}(\mathcal{L}_1 - 1) - g_{l1}}{C_s} & \frac{g_{l2}}{C_s} \\ \frac{(2+\beta)I_0 R_{0,2}}{C_s} & \frac{g_{l1}}{C_s} & \frac{g_{b2}(\mathcal{L}_2 - 1) - g_{l2}}{C_s} \end{bmatrix} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \end{bmatrix} + \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P_1}{C_s} \\ \frac{\delta P_2}{C_s} \end{bmatrix} \quad (\text{C.73})$$

$$= \begin{bmatrix} -\omega_{LR} & -\frac{g_{b1}\mathcal{L}_{b1}}{I_0 L} & -\frac{g_{b2}\mathcal{L}_{b2}}{I_0 L} \\ \frac{(2+\beta)I_0 R_{0,1}}{C_s} & -\omega_{b1} - \omega_{l1} & \omega_{l2} \\ \frac{(2+\beta)I_0 R_{0,2}}{C_s} & \omega_{l1} & -\omega_{b2} - \omega_{l2} \end{bmatrix} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \end{bmatrix} + \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P_1}{C_s} \\ \frac{\delta P_2}{C_s} \end{bmatrix} \quad (\text{C.74})$$

Let's apply the same transformation as before, which previously diagonalized this matrix:

$$S' = T S T^T \quad (\text{C.75})$$

$$= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} -\omega_{LR} & -\frac{g_{b1}\mathcal{L}_{b1}}{I_0 L} & -\frac{g_{b2}\mathcal{L}_{b2}}{I_0 L} \\ \frac{(2+\beta)I_0 R_{0,1}}{C_s} & -\omega_{b1} - \omega_{l1} & \omega_{l2} \\ \frac{(2+\beta)I_0 R_{0,2}}{C_s} & \omega_{l1} & -\omega_{b2} - \omega_{l2} \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} \quad (\text{C.76})$$

$$= \begin{bmatrix} 1 & 0 & 0 \\ 0 & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \\ 0 & \frac{1}{\sqrt{2}} & -\frac{1}{\sqrt{2}} \end{bmatrix} \begin{bmatrix} -\omega_{LR} & -\frac{(g_{b1}\mathcal{L}_{b1} + g_{b2}\mathcal{L}_{b2})}{\sqrt{2}I_0 L} & -\frac{(g_{b1}\mathcal{L}_{b1} - g_{b2}\mathcal{L}_{b2})}{\sqrt{2}I_0 L} \\ \frac{(2+\beta)I_0 R_s}{C_s} & \frac{-\omega_{b1} - \omega_{l1} + \omega_{l2}}{\sqrt{2}} & \frac{-\omega_{b1} - \omega_{l1} - \omega_{l2}}{\sqrt{2}} \\ \frac{(2+\beta)I_0 R_s}{C_s} & \frac{-\omega_{b2} - \omega_{l2} + \omega_{l1}}{\sqrt{2}} & \frac{\omega_{b2} + \omega_{l2} + \omega_{l1}}{\sqrt{2}} \end{bmatrix} \quad (\text{C.77})$$

$$= \begin{bmatrix} -\omega_{LR} & -\frac{(g_{b1}\mathcal{L}_{b1} + g_{b2}\mathcal{L}_{b2})}{\sqrt{2}I_0 L} & -\frac{(g_{b1}\mathcal{L}_{b1} - g_{b2}\mathcal{L}_{b2})}{\sqrt{2}I_0 L} \\ \frac{\sqrt{2}(2+\beta)I_0(R_{0,1} + R_{0,2})}{C} & \frac{-\omega_{b1} - \omega_{b2}}{2} & \frac{\omega_{b2} - \omega_{b1}}{2} \\ \frac{\sqrt{2}(2+\beta)I_0(R_{0,1} - R_{0,2})}{C} & \frac{\omega_{b2} - \omega_{b1} + 2(\omega_{l2} - \omega_{l1})}{2} & \frac{-\omega_{b2} - \omega_{b1} - 2(\omega_{l2} + \omega_{l1})}{2} \end{bmatrix} \quad (\text{C.78})$$

and again re-factoring to remove spurious square roots, we find that the coupling matrix has the form

$$S = \begin{bmatrix} -\omega_{LR} & -\frac{(g_{b1}\mathcal{L}_{b1} + g_{b2}\mathcal{L}_{b2})}{2I_0 L} & -\frac{(g_{b1}\mathcal{L}_{b1} - g_{b2}\mathcal{L}_{b2})}{2I_0 L} \\ \frac{(2+\beta)I_0(R_{0,1} + R_{0,2})}{C} & \frac{-\omega_{b1} - \omega_{b2}}{2} & \frac{\omega_{b2} - \omega_{b1}}{2} \\ \frac{(2+\beta)I_0(R_{0,1} - R_{0,2})}{C} & \frac{\omega_{b2} - \omega_{b1} + 2(\omega_{l2} - \omega_{l1})}{2} & \frac{-\omega_{b2} - \omega_{b1} - 2(\omega_{l2} + \omega_{l1})}{2} \end{bmatrix} \quad (\text{C.79})$$

Now these are definitively coupled equations, and while they will drive the changes to 0, they do not drive the blocks to the same temperature; that is determined by the equilibrium expressions. This allows us to adopt some form for the transition curve and input forms for a given initial temperature difference, after which we should be able to compute transfer functions and the green's function response as a function of the degree of phase separation between these two blocks.

Let's simplify a bit further; the idea here is to try to narrow this down to one variable in addition to the non phase-separated case. We can easily argue for the two-block system that $g_{l1} = g_{l2}(1 + \epsilon)$

by expanding the link cooling power. This gives us the simplified system

$$S = \begin{bmatrix} -\omega_{LR} & -\frac{(g_{b1}\mathcal{L}_{b1}+g_{b2}\mathcal{L}_{b2})}{2I_0L} & -\frac{(g_{b1}\mathcal{L}_{b1}-g_{b2}\mathcal{L}_{b2})}{2I_0L} \\ \frac{(2+\beta)I_0(R_{0,1}+R_{0,2})}{C} & \frac{-\omega_{b1}-\omega_{b2}}{2} & \frac{\omega_{b2}-\omega_{b1}}{2} \\ \frac{(2+\beta)I_0(R_{0,1}-R_{0,2})}{C} & \frac{\omega_{b2}-\omega_{b1}}{2} - \frac{2G\epsilon}{C} & \frac{-\omega_{b2}-\omega_{b1}-2\omega_l(2+\epsilon)}{2} \end{bmatrix} \quad (\text{C.80})$$

where the lower right is starting to look more symmetric. We can use the difference in steady-state temperatures to replace the resistance and conductance terms. Linearizing the conductance terms in initial temperature difference, we find that

$$\frac{g_{b1} + g_{b2}}{2} = G \left(1 + \frac{2\Delta T_{12,0}}{T_0} \right) = G(1 + 2\epsilon) \quad (\text{C.81})$$

$$\frac{g_{b1} - g_{b2}}{2} = G \frac{2\Delta T_{12,0}}{T_0} = 2G\epsilon \quad (\text{C.82})$$

In addition, the addition of our resistance terms cancels, giving just R_0 , and subtraction gives twice the resistance difference:

$$R_{10} - R_{20} = 2R_0\alpha\epsilon$$

giving the system

$$S = \begin{bmatrix} -\omega_{LR} & -\frac{G\mathcal{L}}{I_0L}(1 + 2\epsilon) & -\frac{2\epsilon G\mathcal{L}}{I_0L} \\ \frac{(2+\beta)I_0R_0}{C} & \omega_{eff}(1 + 2\epsilon) & \frac{2G\epsilon}{C} \\ \frac{2\alpha\epsilon(2+\beta)I_0R_0}{C} & 0 & \omega_{eff}(1 + 2\epsilon) - \omega_l(2 + \epsilon) \end{bmatrix} \quad (\text{C.83})$$

C.4 N-Block Solutions

What we saw above was that the linearized model does not predict behavior any different than the phase-uniform TES in a two-block system. We could check that this is the case by looking at higher dimensional systems, where we would find that the mean of all temperature changes is always an eigenvector of the matrix, meaning that this problem will always generate a block-diagonal matrix where we can separate the microscopic and macroscopic dynamics.

$$\frac{d}{dt} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \\ \Delta T_3 \end{bmatrix} = \begin{bmatrix} -\omega_{LR} & -\frac{g_b\mathcal{L}_b}{I_0L} & -\frac{g_b\mathcal{L}_b}{I_0L} & -\frac{g_b\mathcal{L}_b}{I_0L} \\ \frac{(2+\beta)I_0R_s}{C_s} & -\omega_{lb} & \omega_l & 0 \\ \frac{(2+\beta)I_0R_s}{C_s} & \omega_l & -\omega_{2lb} & \omega_l \\ \frac{(2+\beta)I_0R_s}{C_s} & 0 & \omega_l & -\omega_{lb} \end{bmatrix} \begin{bmatrix} \Delta I \\ \Delta T_1 \\ \Delta T_2 \\ \Delta T_3 \end{bmatrix} + \begin{bmatrix} \frac{\delta V}{L} \\ \frac{\delta P_1}{C_s} \\ \frac{\delta P_2}{C_s} \\ \frac{\delta P_3}{C_s} \end{bmatrix} \quad (\text{C.84})$$

Looking at the 3-block system, we start to recognize this lower right matrix; it's nearly identical to that for the normal modes of a couple 1-d harmonic oscillator. This means we can use the tools

from a well-studied problem to determine the phase-separation time-scales, which will allow us to use our measurements and models to predict roughly what frequency range could be affected by phase-separation. This matrix can be scaled further to produce a normal mode analysis of phase separation, which I leave as future work.

Appendix D

QET Efficiency

Using the measurements of R_p , knowledge of R_{sh} , and a known energy spectrum, we can obtain energy efficiency as a function input energy by measurement of dI/dP and use of this to convert measured current into measured power, and thus measured energy in the sensor. We start with the power balance equation

$$C \frac{dT}{dt} = \frac{dE}{dt} = P_{tot} = P_{joule} - P_{bath} + \delta P \quad (\text{D.1})$$

where δP is some energy impulse, assumed 0 in the equilibrium limit. In reality phonons are flowing into the system for a large amount of time, but this will look like a reduced bath power and cancel out, resulting in the same energy measurement; we ignore it for the time being.

Integrating both sides, we know that the long term energy change is 0 (as the sensor returns to equilibrium), giving the expression

$$E_{abs} = \int_0^\infty (P_{bath}(t) - \delta P(t)) dt = \int_0^\infty P_{joule} dt \quad (\text{D.2})$$

and as we only want the change from energy that would have been absorbed in equilibrium heating, we subtract the equilibrium bath power, giving

$$- \Delta E_{abs} = \int_0^\infty (\Delta P_{bath}(t) - \delta P(t)) dt = \int_0^\infty \Delta P_{joule} dt \quad (\text{D.3})$$

Finally, we assume that bath power is roughly constant in the transition (valid for small signals and sharp transitions), giving the expression

$$\Delta E_{abs} \approx \int_0^\infty \delta P(t) dt = - \int_0^\infty \Delta P_{joule} dt \quad (\text{D.4})$$

So as long as bath power change is much less than joule power change, we can assume that the joule power directly tracks the change in phonon energy into the TES. We'll explore this assumption a

bit later, but this is the traditional method by which absolute calibration is attempted.

From earlier derivations, we also find that

$$P_J = V_{TES} I_{TES} = I_{TES} (I_{TES} R_{TES}) \quad (\text{D.5})$$

where

$$I_{TES} R_{TES} = V_{TES} = V_s - I_{TES} (R_p + R_s) - L \frac{dI_{TES}}{dt} \quad (\text{D.6})$$

This gives the equation

$$\Delta V_{TES} = -\Delta I_{TES} (R_p + R_s) - L \frac{d\Delta I_{TES}}{dt} \quad (\text{D.7})$$

Substituting into the equation for change in Joule power gives

$$\Delta P_J = \Delta V_{TES} I_{TES} + V_{TES} \Delta I_{TES} \quad (\text{D.8})$$

$$= \left(-\Delta I_{TES} (R_p + R_{sh}) - L \frac{dI_{TES}}{dt} \right) I_{TES} \quad (\text{D.9})$$

$$+ \left(V_s - I_{TES} (R_p + R_s) - L \frac{dI_{TES}}{dt} \right) \Delta I_{TES} \quad (\text{D.10})$$

$$= -2\Delta I_{TES} (R_p + R_{sh}) I_{TES} - L \frac{d^2 I_{TES}}{dt^2} + V_s \Delta I_{TES} \quad (\text{D.11})$$

and if we use $I_{TES} = I_0 + \Delta I_{TES}$ and $V_s = I_b R_s$ we have

$$\Delta P_J = -2(R_p + R_{sh}) \Delta I_{TES} (I_0 + \Delta I_{TES}) - L \frac{d^2 (\Delta I_{TES})}{dt^2} + V_s \Delta I_{TES} \quad (\text{D.12})$$

$$= (R_s I_b - 2(R_p + R_{sh}) I_0) \Delta I_{TES} - L \frac{d^2 (\Delta I_{TES})}{dt^2} - 2(R_p + R_{sh}) (\Delta I_{TES})^2 \quad (\text{D.13})$$

$$-\Delta P_J = (2(R_p + R_{sh}) I_0 - R_s I_b) \Delta I_{TES} + 2(R_p + R_{sh}) (\Delta I_{TES})^2 + L \frac{d^2 (\Delta I_{TES})}{dt^2} \quad (\text{D.14})$$

The measured energy is thus

$$\Delta E_{abs} \approx - \int_0^\infty \Delta P_{joule} dt = \int_0^\infty [(2(R_p + R_{sh}) I_0 - R_{sh} I_b) \Delta I_{TES} + 2(R_p + R_{sh}) (\Delta I_{TES})^2] dt \quad (\text{D.15})$$

Finally, using (in the DC limit) the relation

$$I_0 = \frac{V_s}{R_{sh} + R_p + R_{TES,0}} = I_b \frac{R_{sh}}{R_{sh} + R_p + R_{TES,0}} \quad (\text{D.16})$$

giving the final expression

$$\Delta E_{abs} \approx \int_0^\infty \left[\left(2 \frac{R_p + R_{sh}}{R_p + R_{sh} + R_{TES,0}} - 1 \right) I_b R_{sh} \Delta I_{TES} + 2(R_p + R_{sh}) (\Delta I_{TES})^2 \right] dt \quad (\text{D.17})$$

This makes sense as the change in current will be negative during the pulse, resulting in a positive energy measurement. Thus, knowing bias current and measuring the three resistance values discussed earlier is all that is needed to produce an energy scale and obtain a direct measure of absolute collection efficiency.

Let's now evaluate the assumption that we do not have to account for changes in the bath power. The above formula is exact, which is what makes it so elegant, however we will need to rely on approximation to add in bath effects. That's alright; we'll hopefully always operate in the limit that they're negligible, so we just want a condition for that to be the case. Bath power is modeled by the equation

$$P_b = K(T^5 - T_b^5) \quad (\text{D.18})$$

and we know that the bath power is balanced by Joule power at $\Delta I = 0$, so we just care about the change:

$$\Delta P_b = K(T^5 - T_0^5) \approx 5KT_0^4 \Delta T + 10KT_0^3 (\Delta T)^2 + O((\Delta T)^3) \quad (\text{D.19})$$

I have kept second-order terms here because ΔT should be first-order in ΔR , while the integral we have is second order in ΔI . We now see that we also have to express ΔT in terms of ΔI , as we do not observe the temperature difference directly.

We already introduced α , the dimensionless responsivity parameter defined as

$$\alpha = \frac{T_0}{R_0} \frac{dR}{dT} \quad (\text{D.20})$$

We can use this to convert ΔT into a proportionality constant and ΔR :

$$\Delta R = \frac{dR}{dT} \Delta T = \frac{R_0 \alpha}{T_0} \Delta T \quad (\text{D.21})$$

$$\Delta T = \frac{T_0}{\alpha R_0} \Delta R \quad (\text{D.22})$$

We can also use the TES voltage equation to find ΔR in terms of ΔI :

$$I_{TES} R_{TES} = I_b R_{sh} - I_{TES} (R_{sh} + R_p) \quad (\text{D.23})$$

$$R_{TES} = R_{sh} \frac{I_b}{I_{TES}} - (R_{sh} + R_p) \quad (\text{D.24})$$

$$\Delta R \approx R_{sh} I_b \left[-\frac{1}{I_0^2} \Delta I + \frac{1}{I_0^3} (\Delta I)^2 \right] \quad (\text{D.25})$$

$$\Delta R \approx \frac{R_{sh} I_b}{I_0^2} \left[(-\Delta I) + \frac{(\Delta I)^2}{I_0} \right] \quad (\text{D.26})$$

giving

$$\Delta T = \frac{T_0}{\alpha R_0} \frac{R_{sh} I_b}{I_0^2} \left[(-\Delta I) + \frac{(\Delta I)^2}{I_0} \right] \quad (\text{D.27})$$

This gives us the first-order bath change

$$\Delta P_b \approx \frac{5KT_0^5}{\alpha I_0^2 R_0} R_{sh} I_b (-\Delta I) \quad (\text{D.28})$$

This is an important result. The power balance equation at equilibrium gives the equation

$$K(T_0^5 - T_b^5) = I_0^2 R_0 \quad (\text{D.29})$$

which allows us to re-write this correction as

$$\Delta P_b \approx \Delta P_{b,1} \left[\frac{\alpha}{5} \left(1 - \frac{T_b^5}{T_0^5} \right) \right]^{-1} R_{sh} I_b (-\Delta I) = \frac{R_{sh} I_b}{\mathcal{L}} (-\Delta I) \quad (\text{D.30})$$

Here the loop gain has reappeared! Inserting this into the first order term gives

$$\Delta E_{abs} \approx \int_0^\infty \left[\left(2 \frac{R_p + R_{sh}}{R_p + R_{sh} + R_{TES,0}} - 1 - \frac{1}{\mathcal{L}} \right) I_b R_{sh} \Delta I_{TES} + 2(R_p + R_{sh}) (\Delta I_{TES})^2 \right] dt \quad (\text{D.31})$$

This matches our expectation; in the high loop gain limit, we only need to account for the joule power changes.

We should also compute the second order correction to make sure that we don't miss a large additional factor. The second order power term is

$$\Delta P_{b,2} = 5KT_0^4 \left(\frac{T_0 R_{sh} I_b}{\alpha R_0 I_0^2} \frac{(\Delta I)^2}{I_0} \right) + 10KT_0^3 \left(\frac{T_0 R_{sh} I_b}{\alpha R_0 I_0^2} (-\Delta I) \right)^2 \quad (\text{D.32})$$

$$= \frac{5KT_0^5}{\alpha I_0^2 R_0} R_{sh} \frac{I_b}{I_0} (\Delta I)^2 + \frac{5KT_0^5}{\alpha I_0^2 R_0} \frac{2(R_{sh} I_b)^2}{\alpha R_0 I_0^2} (-\Delta I)^2 \quad (\text{D.33})$$

$$= \left[1 + \frac{2R_{sh} I_b}{\alpha R_0 I_0} \right] \frac{I_b}{I_0} \frac{R_{sh}}{\mathcal{L}} (-\Delta I)^2 \quad (\text{D.34})$$

$$= \left[1 + \frac{2R_{sh}}{\alpha R_0} \left(1 - \frac{R_0 + R_p}{R_{sh}} \right) \right] \left[1 - \frac{R_0 + R_p}{R_{sh}} \right] \frac{R_{sh}}{\mathcal{L}} (-\Delta I)^2 \quad (\text{D.35})$$

$$= \left[1 - \frac{2}{\alpha} + \frac{2(R_{sh} - R_p)}{\alpha R_0} \right] [R_{sh} - R_0 - R_p] \frac{1}{\mathcal{L}} (-\Delta I)^2 \quad (\text{D.36})$$

$$\approx [R_{sh} - R_0 - R_p] \frac{1}{\mathcal{L}} (-\Delta I)^2 \quad (\text{D.37})$$

where I've left out the additional term as it is further reduced by α even in the weakest cases. This gives the fully corrected absorbed energy

$$\Delta E_{abs} \approx \int_0^\infty \left[\left(2 \frac{R_p + R_{sh}}{R_p + R_{sh} + R_{TES,0}} - 1 - \frac{1}{\mathcal{L}} \right) I_b R_{sh} \Delta I_{TES} \right. \quad (\text{D.38})$$

$$\left. + \left(2(R_p + R_{sh}) + \frac{R_{sh} - R_p - R_0}{\mathcal{L}} \right) (\Delta I_{TES})^2 \right] dt \quad (\text{D.39})$$

Two summarizing comments. First, we see that we have only added one more variable, the loop gain, that we normally measure in a full characterization but not necessarily for production detectors. This suggests a simple additional measurement will allow for accurate accounting of the bath power in addition to joule power change, and an efficiency that's less susceptible to bath temperature or transition sharpness. Second, we see that in the limit of large loop gain, this correction does in fact disappear, which means our initial assumption that a sharp transition made bath power negligible was correct.

Appendix E

Optimal Filtering

In this appendix I review the formalism for the Fourier domain optimum filter and present the pileup and correlation optimum filters which will help achieve better thresholds in the HV detectors. I also include sections on Wiener filtering, which is the optimal filter in the absence of a template, and on the time-domain version of the optimal filter, which is usually called a matched filter, and includes the correlation matrix in place of the PSD. In this note, the measured signal is $v(f)$ in the frequency domain or $v(t)$ in the time domain, the template is $s(f)$ or $s(t)$, and the noise PSD is $J(f)$.

E.1 Basic Optimal Filter

For the chi-squared function

$$\chi^2 = \int_{-\infty}^{\infty} \frac{|v(f) - As(f)|^2}{J(f)} df \quad (\text{E.1})$$

we find the goodness of fit by minimizing χ^2 with respect to A, as

$$0 = \frac{d\chi^2}{dA} = \frac{d}{dA} \int_{-\infty}^{\infty} \frac{v^*(f)v(f) - 2As^*(f)v(f) + A^2s^*(f)s(f)}{J(f)} df \quad (\text{E.2})$$

$$0 = 2 \int_{-\infty}^{\infty} \frac{-s^*(f)v(f) + As^*(f)s(f)}{J(f)} df \quad (\text{E.3})$$

$$\int_{-\infty}^{\infty} \frac{s^*(f)v(f)}{J(f)} df = A \int_{-\infty}^{\infty} \frac{s^*(f)s(f)}{J(f)} df \quad (\text{E.4})$$

$$A = \frac{\int_{-\infty}^{\infty} \frac{s^*(f)v(f)}{J(f)} df}{\int_{-\infty}^{\infty} \frac{|s(f)|^2}{J(f)} df} \quad (\text{E.5})$$

This suggests that the optimum filter for this signal has the form

$$\phi(f) = \frac{s^*(f)}{J(f)} \quad (\text{E.6})$$

so that we can write the optimal estimate as

$$A = \frac{\int_{-\infty}^{\infty} \phi(f)v(f)df}{\int_{-\infty}^{\infty} \phi(f)s(f)df} \quad (\text{E.7})$$

or simplifying further, we can renormalize the filter as

$$\phi'(f) = \frac{\frac{s^*(f)}{J(f)}}{\int_{-\infty}^{\infty} \phi(f)s(f)df} \quad (\text{E.8})$$

to give the resulting simple estimator

$$A = \int_{-\infty}^{\infty} \phi'(f)v(f)df \quad (\text{E.9})$$

This is a nice and clean derivation which only captures the simplest case, where there is a single, well-known template, one signal per trace, and the start of the pulse is well-known. This is also a derivation for the continuous limit, whereas real data is sampled discretely. In this note I will delve into more complex optimal filters which each relax one of these assumptions, and provide formulae for discrete optimum filters for the different cases along with reduced chi-square formulae to make their implementation simple and straightforward.

E.2 Optimum Filter with Time Offset

Here we make the extension that the signal can slide in time, i.e. that

$$s(f) \rightarrow s(f, t_0) = e^{-i\omega t_0} s(f) \quad (\text{E.10})$$

which has the goodness of fit

$$\chi^2 = \int_{-\infty}^{\infty} \frac{|v(f) - Ae^{-i\omega t_0} s(f)|^2}{J(f)} df \quad (\text{E.11})$$

Here we can go through the same procedure to get an optimal estimate of A , now as a function of t_0 :

$$0 = \frac{d\chi^2}{dA}(t_0) = \frac{d}{dA} \int_{-\infty}^{\infty} \frac{|v(f) - Ae^{-i\omega t_0} s(f)|^2}{J(f)} df \quad (\text{E.12})$$

$$0 = \frac{d}{dA} \int_{-\infty}^{\infty} \frac{v^*(f)v(f) - 2Ae^{i\omega t_0} s^*(f)v(f) + A^2 s^*(f)s(f)}{J(f)} df \quad (\text{E.13})$$

$$\int_{-\infty}^{\infty} \frac{e^{i\omega t_0} s^*(f)v(f)}{J(f)} df = \int_{-\infty}^{\infty} \frac{As^*(f)s(f)}{J(f)} df \quad (\text{E.14})$$

$$A = \frac{\int_{-\infty}^{\infty} \frac{e^{i\omega t_0} s^*(f)v(f)}{J(f)} df}{\int_{-\infty}^{\infty} \frac{s^*(f)s(f)}{J(f)} df} \quad (\text{E.15})$$

which looks similar to the original form except for the addition of the time offset. Writing this with the normalized optimum filter makes it a bit more explicit what we're looking at:

$$A(t_0) = \int_{-\infty}^{\infty} e^{i\omega t_0} \phi'(f)v(f)df \quad (\text{E.16})$$

This is just the inverse fourier transform of the filtered trace! So adding a time offset will only marginally increase the complexity of our optimal filter algorithm. We can get the amplitude as a function of time, and then look for the best chi-square in time space. With this amplitude function, we can re-write the chi-square as

$$\begin{aligned} \chi^2(t_0) &= \int_{-\infty}^{\infty} \frac{|v(f)|^2}{J(f)} df - 2A(t_0) \int_{-\infty}^{\infty} e^{i\omega t_0} \phi(f)v(f)df + A^2(t_0) \int_{-\infty}^{\infty} \phi(f)s(f)df \\ &= \int_{-\infty}^{\infty} \frac{|v(f)|^2}{J(f)} df - A^2(t_0) \int_{-\infty}^{\infty} \phi(f)s(f)df \end{aligned}$$

and we find that the optimal amplitude with optimal offset is the condition that maximizes the filtered amplitude function $A(t_0)$. Practically speaking, we can compute the optimal amplitude with delay by computing $A(t_0)$, computing the resulting χ^2 as shown, and finding the amplitude corresponding to the minimum chi-square or the maximum amplitude. Re-writing this in more

convenient form allows the process to be more explicit:

$$\begin{aligned}\phi(f) &= \frac{s^*(f)}{J(f)} \\ N_\phi &= \int_{-\infty}^{\infty} \phi(f)s(f)df \\ \phi'(f) &= \frac{\phi(f)}{N_\phi} \\ A(t_0) &= \int_{-\infty}^{\infty} e^{i\omega t_0} \phi'(f)v(f)df \\ \chi_0^2 &= \int_{-\infty}^{\infty} \frac{|v(f)|^2}{J(f)} df \\ \chi^2(t_0) &= \chi_0^2 - A^2(t_0)N\end{aligned}$$

Using this fourier transform trick makes the implementation of the OF with delay a trivial extension of the nominal optimum filter.

E.3 Two-Pulse Optimum Filter

Suppose now that instead of one event per trace, we have two. We are still in the small signal limit of our device (in principle) so there should be no different between the templates for each pulse, but now we have two amplitudes and two offsets. In other words,

$$As(f, t_0) \rightarrow A_1s(f, t_1) + A_2s(f, t_2) \quad (\text{E.17})$$

which has the goodness of fit

$$\chi^2 = \int_{-\infty}^{\infty} \frac{|v(f) - A_1e^{-i\omega t_1}s(f) - A_2e^{-i\omega t_2}s(f)|^2}{J(f)} df \quad (\text{E.18})$$

Let's expand the numerator to see what happens:

$$|v(f) - A_1e^{-i\omega t_1}s(f) - A_2e^{-i\omega t_2}s(f)|^2 \quad (\text{E.19})$$

$$= (v^*(f) - A_1e^{i\omega t_1}s^*(f) - A_2e^{i\omega t_2}s^*(f))(v(f) - A_1e^{-i\omega t_1}s(f) - A_2e^{-i\omega t_2}s(f)) \quad (\text{E.20})$$

$$= v^*(f)v(f) - 2v(f)(A_1e^{i\omega t_1}s^*(f) + A_2e^{i\omega t_2}s^*(f)) + |A_1e^{-i\omega t_1}s(f) + A_2e^{-i\omega t_2}s(f)|^2 \quad (\text{E.21})$$

$$= |v(f)|^2 - 2v(f)(A_1e^{i\omega t_1} + A_2e^{i\omega t_2})s^*(f) + [A_1^2 + A_2^2 + 2A_1A_2 \cos(\omega(t_1 - t_2))] |s(f)|^2 \quad (\text{E.22})$$

Pushing the denominator through gives us

$$\frac{|v(f)|^2}{J(f)} - 2\phi(f)v(f)(A_1e^{i\omega t_1} + A_2e^{i\omega t_2}) + [A_1^2 + A_2^2 + 2A_1A_2 \cos(\omega(t_1 - t_2))] \phi(f)s(f) \quad (\text{E.23})$$

and minimizing with respect to A_i gives

$$0 = \frac{d\chi^2}{dA_i} = \frac{d}{dA_i} \int_{-\infty}^{\infty} -2\phi(f)v(f)(A_1 e^{i\omega t_1} + A_2 e^{i\omega t_2})df \quad (\text{E.24})$$

$$+ \frac{d}{dA_i} \int_{-\infty}^{\infty} [A_1^2 + A_2^2 + 2A_1 A_2 \cos(\omega(t_1 - t_2))] \phi(f)s(f)df \quad (\text{E.25})$$

$$= \int_{-\infty}^{\infty} -2\phi(f)v(f)e^{i\omega t_i} df + [2A_i + 2A_j \cos(\omega(t_1 - t_2))] \phi(f)s(f)df \quad (\text{E.26})$$

$$A_i \int_{-\infty}^{\infty} \phi(f)s(f)df = \int_{-\infty}^{\infty} \phi(f)v(f)e^{i\omega t_i} df - \int_{-\infty}^{\infty} A_j \cos(\omega(t_1 - t_2))\phi(f)s(f)df \quad (\text{E.27})$$

$$A_i = \int_{-\infty}^{\infty} \phi'(f)v(f)e^{i\omega t_i} df - A_j \int_{-\infty}^{\infty} \cos(\omega\delta t)\phi'(f)s(f)df \quad (\text{E.28})$$

$$A_i = A_{0,i}(t_i) - A_j(t_j)\gamma(\delta t) \quad (\text{E.29})$$

where here we've switched to the parameterization where $t_1 = t_0$ and $t_2 = t_0 + \delta t$, and introduced the new variable

$$\gamma(\delta t) = \int_{-\infty}^{\infty} \cos(\omega\delta t)\phi'(f)s(f)df \quad (\text{E.30})$$

For the case of $\delta t = 0$, this trivially tells us that the amplitudes should add. For a given t_0 and δt , we thus have two equations with two unknowns, so we can eliminate one amplitude to find the optimal amplitude as

$$A_i = A_{i,0}(t_i) - (A_{j,0}(t_j) - A_i(t_i)\gamma(\delta t))\gamma(\delta t) \quad (\text{E.31})$$

$$A_i = A_{i,0}(t_i) - A_{j,0}(t_j)\gamma(\delta t) + A_i(t_i)\gamma^2(\delta t) \quad (\text{E.32})$$

$$A_i(1 - \gamma^2(\delta t)) = A_{i,0}(t_i) - A_{j,0}(t_j)\gamma(\delta t) \quad (\text{E.33})$$

$$A_i(t_i, t_j) = \frac{A_{i,0}(t_i) - A_{j,0}(t_j)\gamma(\delta t)}{1 - \gamma^2(\delta t)} \quad (\text{E.34})$$

As a sanity check, this formula tells us that for $\delta t = 0$, the amplitude evaluates to 0, and the other amplitude thus becomes the original value for 0 delay. Putting this in more concrete terms, let's use this formula to estimate A_2 . We thus get

$$A_2(t_1, \delta t) = \frac{A_2(t_1 + \delta t) - A_1(t_1)\gamma(\delta t)}{1 - \gamma^2(\delta t)} \quad (\text{E.35})$$

$$A_1(t_1, \delta t) = A_1(t_1) - A_2(t_1, \delta t)\gamma(\delta t) \quad (\text{E.36})$$

$$= \frac{(1 - \gamma^2(\delta t))A_1(t_1) - A_2(t_1 + \delta t)\gamma(\delta t) + A_1(t_1)\gamma^2(\delta t)}{1 - \gamma^2(\delta t)} \quad (\text{E.37})$$

$$= \frac{A_1(t_1) - A_2(t_1 + \delta t)\gamma(\delta t)}{1 - \gamma^2(\delta t)} \quad (\text{E.38})$$

Finishing the problem, we can rearrange the terms in the chi-square to make them look similar to the single-template with time shift:

$$\chi^2 = \chi_0^2 - \int_{-\infty}^{\infty} [2\phi(f)v(f)(A_1 e^{i\omega t_1} + A_2 e^{i\omega t_2}) - [A_1^2 + A_2^2 + 2A_1 A_2 \cos(\omega(t_1 - t_2))] \phi(f)s(f)] df \quad (\text{E.39})$$

$$= \chi_0^2 - 2A_1 \int_{-\infty}^{\infty} [\phi(f)v(f)e^{i\omega t_1} - A_2 \cos(\omega(t_1 - t_2))\phi(f)s(f)] df + A_1^2 \int_{-\infty}^{\infty} \phi(f)s(f) df \quad (\text{E.40})$$

$$- 2A_2 \int_{-\infty}^{\infty} [\phi(f)v(f)e^{i\omega t_2} - A_1 \cos(\omega(t_1 - t_2))\phi(f)s(f)] df + A_2^2 \int_{-\infty}^{\infty} \phi(f)s(f) df \quad (\text{E.41})$$

$$- 2A_1 A_2 \int_{-\infty}^{\infty} \cos(\omega(t_1 - t_2))\phi(f)s(f) df \quad (\text{E.42})$$

$$= \chi_0^2 - [A_1^2 + A_2^2 + 2A_1 A_2 \gamma(\delta t)] N_\phi \quad (\text{E.43})$$

We can make this algorithm more efficient, practically speaking, by substituting the cosine transform with an inverse fourier transform and taking the real part. This can be see here:

$$\gamma(\delta t) = \text{Re} \left[\int_{-\infty}^{\infty} e^{\omega \delta t} \phi'(f)s(f) df \right] \quad (\text{E.44})$$

$$= \text{Re} \left[\int_{-\infty}^{\infty} (\cos(\omega \delta t) + i \sin(\omega \delta t)) \phi'(f)s(f) df \right] \quad (\text{E.45})$$

$$= \text{Re} \left[\int_{-\infty}^{\infty} \cos(\omega \delta t) \phi'(f)s(f) df + i \int_{-\infty}^{\infty} \sin(\omega \delta t) \phi'(f)s(f) df \right] \quad (\text{E.46})$$

$$= \int_{-\infty}^{\infty} \cos(\omega \delta t) \phi'(f)s(f) df \quad (\text{E.47})$$

Here we see that we could also use the direct transform, but we want the normalization to be consistent with that of the amplitudes, so we choose to use the inverse transform for both. Note that this works because the product $\phi(f)'s(f)$ is real.

So to summarize as in the previous section, the calculation procedes as follows. We first calculate

the one-dimensional quantities:

$$\begin{aligned}\phi(f) &= \frac{s^*(f)}{J(f)} \\ N_\phi &= \int_{-\infty}^{\infty} \phi(f)s(f)df \\ \phi'(f) &= \frac{\phi(f)}{N_\phi} \\ A_0(t_i) &= \int_{-\infty}^{\infty} e^{i\omega t_i} \phi'(f)v(f)df \\ \gamma(\delta t) &= \text{Re} \left[\int_{-\infty}^{\infty} e^{i\omega \delta t} \phi'(f)s(f)df \right]\end{aligned}$$

We then project these into the full two-dimensional space and compute the chi-square as follows:

$$\begin{aligned}A_2(t_1, \delta t) &= \frac{A_0(t_1 + \delta t) - A_0(t_1)\gamma(\delta t)}{1 - \gamma^2(\delta t)} \\ A_1(t_1, \delta t) &= A_0(t_1) - A_2(t_1, \delta t)\gamma(\delta t) \\ \chi_0^2 &= \int_{-\infty}^{\infty} \frac{|v(f)|^2}{J(f)} df \\ \chi^2(t_1, \delta t) &= \chi_0^2 - (A_1^2(t_1, \delta t) + A_2^2(t_1, \delta t) + 2A_1(t_1, \delta t)A_2(t_1, \delta t)\gamma(\delta t))N\end{aligned}$$

The convenient aspect of this algorithm is that it allows us to easily restrict our search windows in either t_1 or δt without any additional tricks, we simply limit the domain of the two-dimensional projection. A final note: we have to take care to treat the $\delta t = 0$ case independently; mathematically this will reduce correctly, but we have to show that by evaluating a limit in which the numerator and denominator both go to 0. Taking the derivatives, we can show that for $\delta t = 0$, this solution reduces to the single template solution:

$$\lim_{\delta t \rightarrow 0} A_2(t_1, \delta t) = \lim_{\delta t \rightarrow 0} \frac{A_0(t_1 + \delta t) - A_0(t_1)\gamma(\delta t)}{1 - \gamma^2(\delta t)} = \lim_{\delta t \rightarrow 0} \frac{A_0'(t_1 + \delta t) - A_0'(t_1)\gamma(\delta t) - A_0(t_1)\gamma'(\delta t)}{1 - 2\gamma(\delta t)\gamma'(\delta t)} = \frac{0}{1 - 0} = 0 \quad (\text{E.48})$$

$$\rightarrow A_1(t_1, 0) = A_0(t_1) - A_2(t_1, 0) = A_0(t_1), \quad \chi^2(t_1, \delta t) = \chi_0^2 - A_1^2(t_1)N \quad (\text{E.49})$$

which is just the single-template optimum filter with time-offset.

E.4 Joint Channel Optimum Filter

Moving back to the case of a single event in a trace, let's consider a joint fit of all channels with any number time-shifted templates. Let's add a channel index i into our notation, such that the signal from channel i is $v_i(f)$ and the template at $t_0 = 0$ in channel i for a template γ is $s_{i\gamma}(f)$. In this

way, we allow for n channels and m templates, where each template can take a separate form in each channel but has one unified amplitude across all channels, allowing for correlated signals between channels.

We can construct the least-squares statistic as usual:

$$\chi^2 = \sum_i \int_{-\infty}^{\infty} \frac{|v_i(f) - \sum_{\gamma} A_{\gamma} e^{-i\omega t_0} s_{i\gamma}(f)|^2}{J_i(f)} df \quad (\text{E.50})$$

Here we sum templates for a given channel, and sum over least-squares for all channels to get a joint goodness of fit estimator. We can expand the numerator to get

$$|v_i(f) - \sum_{\gamma} A_{\gamma} e^{-i\omega t_0} s_{i\gamma}(f)|^2 \quad (\text{E.51})$$

$$= (v_i^*(f) - \sum_{\gamma} A_{\gamma} e^{i\omega t_0} s_{i\gamma}^*(f))(v_i(f) - \sum_{\beta} A_{\beta} e^{-i\omega t_0} s_{i\beta}(f)) \quad (\text{E.52})$$

$$= v_i^*(f)v_i(f) - 2 \sum_{\gamma} A_{\gamma} e^{i\omega t_0} s_{i\gamma}^*(f)v_i(f) + \sum_{\gamma} \sum_{\beta} A_{\gamma} A_{\beta} s_{i\gamma}^* s_{i\beta}(f) \quad (\text{E.53})$$

$$= |v_i(f)|^2 - 2 \sum_{\gamma} A_{\gamma} e^{i\omega t_0} s_{i\gamma}^*(f)v_i(f) + \sum_{\gamma} \sum_{\beta} A_{\gamma} A_{\beta} s_{i\gamma}^* s_{i\beta}(f) \quad (\text{E.54})$$

pushing the denominator through we get the integrand

$$\frac{|v_i(f)|^2}{J_i(f)} - 2 \sum_{\gamma} A_{\gamma} e^{i\omega t_0} \phi_{i\gamma}(f)v_i(f) + \sum_{\gamma} \sum_{\beta} A_{\gamma} A_{\beta} \phi_{i\gamma} s_{i\beta}(f) \quad (\text{E.55})$$

where we now have the optimum filter $\phi_{i\gamma}$ for channel i , template γ

$$\phi_{i\gamma} = \frac{s_{i\gamma}^*(f)}{J_i(f)} \quad (\text{E.56})$$

If we minimize the fit with respect to a given amplitude A_{α} , we find

$$\frac{d}{dA_{\alpha}} \chi^2 = 0 = \frac{d}{dA_{\alpha}} \sum_i \int_{-\infty}^{\infty} \left[\frac{|v_i(f)|^2}{J_i(f)} - 2 \sum_{\gamma} A_{\gamma} e^{i\omega t_0} \phi_{i\gamma}(f)v_i(f) + \sum_{\gamma} \sum_{\beta} A_{\gamma} A_{\beta} \phi_{i\gamma} s_{i\beta}(f) \right] df \quad (\text{E.57})$$

$$= \sum_i \int_{-\infty}^{\infty} \left[-2e^{i\omega t_0} \phi_{i\alpha}(f)v_i(f) + 2 \sum_{\beta} A_{\beta} \phi_{i\alpha} s_{i\beta}(f) \right] df \quad (\text{E.58})$$

$$\sum_i \int_{-\infty}^{\infty} e^{i\omega t_0} \phi_{i\alpha}(f)v_i(f) df = \sum_{\beta} A_{\beta} \sum_i \int_{-\infty}^{\infty} \phi_{i\alpha} s_{i\beta}(f) df \quad (\text{E.59})$$

This looks odd until we realize that it is just a system of equations with one equation for each of

the m templates. We can thus turn it into the matrix operation

$$\mathbf{V} = \mathbf{N}\mathbf{A} \rightarrow \mathbf{A} = \mathbf{N}^{-1}\mathbf{V} \quad (\text{E.60})$$

where

$$V_\alpha = \sum_i \int_{-\infty}^{\infty} e^{i\omega t_0} \phi_{i\alpha}(f) v_i(f) df \quad (\text{E.61})$$

$$N_{\alpha\beta} = \sum_i \int_{-\infty}^{\infty} \phi_{i\alpha} s_{i\beta}(f) df \quad (\text{E.62})$$

As we'll see later, these are just the normal equations for any linear model, but we've left this problem slightly nonlinear by introducing the time offset, so we'll still need to evaluate the chi-square at each allowable time offset. In this notation the chi-square becomes

$$\chi^2 = \chi_0^2 - \mathbf{A}^T \mathbf{V} = \chi_0^2 - \mathbf{A}^T \mathbf{N} \mathbf{A}$$

which again, looking forward to the next appendix, just resembles the normal equations chi-square, but we've made the time-shift implicit in the amplitudes and data vector.

This is functionally a generalization of the original time-shifted optimum filter because it reduces to that filter in the limit that $S_{\alpha\beta} = \delta_{\alpha\beta} S_{\alpha\alpha}$. This is either the case for orthonormal templates, which we don't expect in our detectors, or for channels which are completely uncorrelated in their signal, which is also not a correct assumption. For these reasons we should expect this OF to perform better than fitting the channels independently. It's also interesting to note that adding arbitrary linear dimensionality to our model space is not nearly as involved as adding one additional non-linear element to the model, as can be seen by comparing this and the pileup OF sections.

E.5 Joint Channel Correlated Optimum Filter

The final generalization of this technique which can accommodate any type of detector where the channels some set of known templates is to include the effect of noise covariance between channels. As in the previous section, we have a model made of M templates for N channels. The least-squares statistic takes on the modified form

$$\chi^2 = \sum_i \sum_j \int_{-\infty}^{\infty} \left(v_i(f) - \sum_\gamma A_\gamma e^{-i\omega t_0} s_{i\gamma}(f) \right)^* \Sigma_{ij}^{-1}(f) \left(v_j(f) - \sum_\beta A_\beta e^{-i\omega t_0} s_{j\beta}(f) \right) df \quad (\text{E.63})$$

This of course reduces to the form of the previous section in the limit that $\Sigma_{ij}^{-1} = \delta_{ij} J_i^{-1}(f)$, so we can see that now we're going to modify the chi-square for the same model if we allow for noise

correlation between channels. Here the matrix $\Sigma_{ij}(f)$ is the cross-power spectral density of the channels, where the diagonals are the auto-power spectral densities (what we have been calling simply the PSDs) of each channel.

We should note that this is a slightly different form than how the generalized linear least-squares model is written, as to write this as a matrix product we would have to express the covariance matrix in a block-diagonal form for each frequency which would add dimensionality to the problem which has been removed by use of the Fourier basis. As in the previous section, the integrals will end up in our matrix definitions, and the result will resemble the linear least-squares case.

Let's expand out the terms as usual in our least-squares statistic. We have the model-independent term

$$\chi_0^2 = \sum_i \sum_j \int_{-\infty}^{\infty} v_i^*(f) \Sigma(f)_{ij}^{-1} v_j(f) df \quad (\text{E.64})$$

and the model-dependent term

$$\sum_i \sum_j \int_{-\infty}^{\infty} \left[-2 \sum_{\gamma} A_{\gamma} e^{i\omega t_0} s_{i\gamma}^*(f) \Sigma_{ij}^{-1}(f) v_j(f) + \sum_{\gamma} \sum_{\beta} A_{\gamma} A_{\beta} s_{i\gamma}^*(f) \Sigma_{ij}^{-1}(f) s_{j\beta}(f) \right] df \quad (\text{E.65})$$

$$= \sum_i \sum_j \int_{-\infty}^{\infty} \left[-2 \sum_{\gamma} A_{\gamma} e^{i\omega t_0} \phi_{ij\gamma}(f) v_j(f) + \sum_{\gamma} \sum_{\beta} A_{\gamma} A_{\beta} \phi_{ij\gamma} s_{j\beta}(f) \right] df \quad (\text{E.66})$$

By analogy with the previous section, we find that

$$V_{\alpha} = \sum_i \sum_j \int_{-\infty}^{\infty} e^{i\omega t_0} \phi_{ij\alpha}(f) v_j(f) df \quad (\text{E.67})$$

$$N_{\alpha\beta} = \sum_i \sum_j \int_{-\infty}^{\infty} \phi_{ij\alpha}(f) s_{j\beta}(f) df \quad (\text{E.68})$$

and as usual

$$A_{\alpha} = \sum_{\beta} N_{\alpha\beta}^{-1} V_{\beta} \quad (\text{E.69})$$

and

$$\chi^2 = \chi_0^2 - \mathbf{A}^T \mathbf{N} \mathbf{A} \quad (\text{E.70})$$

As a concluding remark, I'll point out that the logical continuation of the continually more general solutions (which are less and less like filters as we generalize) would be to revert to the time domain and remove the assumption of stationary noise; in this case we would just be dealing with the normal generalized least-squares, which is described in the next section, but the dimensionality of our problem would increase by orders of magnitude. The assumption of stationary, relatively uncorrelated noise allows us to work with PSDs, in the Fourier domain, and reduce the dimensionality of the problem to a tolerable level. That being said, we still need to take the inverse of a potentially

large matrix twice in the construction phase of the filter, once to determine the covariance matrix, and once again to determine the filter normalization S^{-1} .

E.6 Optimal Filter Resolution

I've written the above filter algorithms in consistent terminology in order to utilize a general result, discussed in the next appendix, for chi-square fits. The covariance of the parameters of our model A_i is just

$$\Sigma_{ij} = N_{ij}^{-1} \quad (\text{E.71})$$

where for the simplest 1D optimum filter, this is just the inverse of the filter norm, but for the more complex optimum filters it is properly the inverse matrix of the problem's design matrix. Consider, for example, the correlated optimum filter with a single template. This filter has only one variable to fit, the joint channel amplitude A_0 , and has the expected resolution

$$\sigma_0^2 = \left[\sum_i \sum_j \int_{-\infty}^{\infty} \phi_{ij}(f) s_j(f) df \right]^{-1} = \left[\sum_i \sum_j \int_{-\infty}^{\infty} s_i^*(f) \Sigma_{ij}^{-1}(f) s_j(f) df \right]^{-1} \quad (\text{E.72})$$

To see that this makes sense, consider the case of uncorrelated noise, with equally weighted templates and noise. We then see that the cross-terms go to 0, and the remaining terms are the resolutions of the 1D optimum filter, thus

$$\sigma_0^2 = \left[\frac{N}{\sigma_{1D}^2} \right]^{-1} = \frac{\sigma_{1D}^2}{N} \quad (\text{E.73})$$

where

$$\sigma_{1D}^2 = \left[\int_{-\infty}^{\infty} s^*(f) J^{-1}(f) s(f) df \right]^{-1} \quad (\text{E.74})$$

This is exactly what we expect for measurements of the same quantity under equal and uncorrelated noise. If the noise is uncorrelated, but some measurements are noisier, this is just the resolution of the weighted sum of observations.

E.7 Wiener Filtering

In this section I describe the more general Wiener filter, in which a new filter is constructed for each trace and the shape of the filter does not depend on the templates used. A good summary and discussion of Wiener filtering can be found in Numerical Recipes, section 13.3.

Suppose we have some Fourier signal $S(f) = R(f)U(f)$, where $U(f)$ is the signal we want to infer from the noise-free measured signal $S(f)$, and an expected noise $N(f)$, but we don't want to apply a template during filtering. If we measure the signal $C(f) = N(f) + S(f)$, then we'd like to

construct a filter to estimate $\tilde{U}(f)$ as

$$\tilde{S}(f) = R(f)\tilde{U}(f) = \Phi(f)C(f) \quad (\text{E.75})$$

Here I've tried to explicitly keep the template and measurements separate. If we minimize the difference between $U(f)$ and $\tilde{U}(f)$, we arrive at the maximum likelihood filter estimate

$$\Phi(f) = \frac{|S(f)|^2}{|S(f)|^2 + |N(f)|^2} \quad (\text{E.76})$$

Contrasting this with the single-pulse optimum filter, we see that it discards the frequency information in the signal. So far we still have an unknown signal in this equation, but given that the signal and noise are uncorrelated, we can re-write this as

$$\Phi(f) \approx \frac{|C(f)|^2 - J(f)}{C(f)} = 1 - \frac{J(f)}{|C(f)|^2} \quad (\text{E.77})$$

This filter has the nice property that the bandwidth of the filter increases as the signal becomes stronger, and the filter goes to 0 for the case that there's no signal. This of course relies heavily on the assumption that our noise and signal estimates are fairly precise in the highest signal/noise region.

Say you now have a signal template such that $S(f) = As(f)$ where $s(f)$ is normalized to 1. You would then estimate the signal amplitude as

$$\tilde{A} = \text{Re} \left[\int_{-\infty}^{\infty} \tilde{S}(f) df \right] = \text{Re} \left[\int_{-\infty}^{\infty} \Phi(f)C(f) df \right] \quad (\text{E.78})$$

This is obviously not as clean as the normal optimum filtering formalism, but it does allow you to go further through the process before needing a signal template, and may be a useful trigger quantity to determine whether a pulse exists in data.

E.8 Time-Domain OF: Matched Filtering

Rather than re-doing the full derivation of the matched filter, I want to lay down the basics here so people can appreciate what the trade-offs are between time-domain and frequency-domain optimal filtering. For a full derivation and an extension of the matched filter to include time offset, see Ref [38](#).

We construct the matched filter to estimate the amplitude of a known signal as

$$\hat{A}(t) = \int_{-\infty}^{\infty} \phi(t' - t)v(t') dt' \quad (\text{E.79})$$

We already know how this goes from the frequency domain optimal filter when we take the maximum likelihood estimate for the filter. In this case, instead of using the noise PSD, we use the noise auto-correlation matrix

$$R(\tau_1, \tau_2) = \frac{1}{T} \int_0^T (n(t + \tau_1) - \mu)(n(t + \tau_2) - \mu) dt = \left[\frac{1}{T} \int_0^T n(t + \tau_1)n(t + \tau_2) - \mu^2 \right] \quad (\text{E.80})$$

which is obviously over-defined; this is a matrix with Toeplitz symmetry and is fully specified by the first row or column (this should help to speed up computation). We then construct the filter

$$\phi(\tau_2) = \int_0^T s(-\tau_1)R^{-1}(\tau_1, \tau_2)d\tau_1 \quad (\text{E.81})$$

which needs to be normalized by the template using

$$N_\phi = \int_0^T \int_0^T s(-\tau_1)R^{-1}(\tau_1, \tau_2)s(\tau_2)d\tau_1d\tau_2 \quad (\text{E.82})$$

which we know from all linear least-squares problems gives us the filter resolution

$$\sigma^2 = N_\phi^{-1} \quad (\text{E.83})$$

We then get the normalized filter

$$\phi(\tau_2) = \frac{\phi'(\tau_2)}{N} \quad (\text{E.84})$$

which allows us to calculate our expected amplitude.

Let's see what happens for this filter for the special case that the noise is completely uncorrelated and white. We find

$$R(\tau_1, \tau_2) = \delta(\tau_2 - \tau_1)\sigma_{white}^2 \quad (\text{E.85})$$

which gives us the OF resolution

$$\sigma^2 = N_\phi^{-1} = \left[\sigma_{white}^{-2} \int_0^T \int_0^T s(-\tau_1)\delta(\tau_2 - \tau_1)s(\tau_2)d\tau_1d\tau_2 \right]^{-1} \quad (\text{E.86})$$

$$= \sigma_{white}^2 \left[\int_0^T s(-\tau_2)s(\tau_2)d\tau_2 \right]^{-1} \quad (\text{E.87})$$

So the filter just has some resolution reduced relative to the variance of the trace according to the convolution of of the templates (defined to have unit amplitude). If we have a flat template we just get an averaging reduction

$$\sigma^2 = \frac{\sigma_{white}^2}{N} \quad (\text{E.88})$$

If we have a normalized exponential template, we find the convolution of the exponential gives T/τ^2 , so we get the resolution

$$\sigma^2 = \sigma_{white}^2 \frac{(\tau/\Delta t)^2}{N} \quad (\text{E.89})$$

where I've implicitly done the discrete math to get the correct dimensionality. In essence, we find that resolution scales linearly with pulse fall-time, and as the inverse square root of the number of samples taken. Thus longer traces and shorter pulses will give us better resolution.

Appendix F

Trace Analysis & Model Fitting

In appendix [E](#) I presented some formal derivations of optimum filters, which is pedagogical and general but is never directly applicable given that our data is always discrete. In this appendix I will discuss the implementation of optimal filters, including the use of both time and Fourier-domain chi-square tests to help mitigate some of the issues encountered doing optimal filtering in the Fourier domain. This guide should save others time when implementing these functions by boiling down the implementation to easily coded procedures.

F.1 Fourier Methods

In this section I will review important elements in Fourier analysis that are used in analyzing and fitting CDMS data traces and characterizing noise. I will focus on the discrete transform which is more applicable to analyzing real data, which is always taken with a sampling frequency and trace length, and discuss ways to draw out relationships between data with different frequency-domain tricks.

F.1.1 Continuous Transform

The continuous Fourier transform is usually defined as

$$x(f) = \int_{-\infty}^{\infty} x(t)e^{-i2\pi ft} dt \quad (\text{F.1})$$

where the transform is a function of ω . This allows the inverse transform to be

$$x(t) = \int_{-\infty}^{\infty} e^{i2\pi ft} x(f) df \quad (\text{F.2})$$

If constructed as a function of angular frequency, we need to divide the inverse transform by 2π in order to get the same Fourier coefficients. There are many common references for properties of the Fourier transform, so I won't bother reproducing all of the properties here.

One property I will cover, that we use a lot in our derivations, is the fact that we usually don't care about the phase of the Fourier coefficients. Consider what happens when we flip the sign of the frequency. Our positive frequency coefficient is defined in the normal way, but our negative frequency coefficient depends on whether the function is even or odd. For an even function, we have

$$x(-f) = \int_{-\infty}^{\infty} x(t)e^{i2\pi ft} dt \quad (\text{F.3})$$

$$\int_{-\infty}^{\infty} x(-t)e^{-i2\pi ft} dt = \int_{-\infty}^{\infty} x(t)e^{-i2\pi ft} dt \quad (\text{F.4})$$

$$\int_{-\infty}^{\infty} x(-t) [\cos(2\pi ft) - i \sin(2\pi ft)] dt = \int_{-\infty}^{\infty} x(t) [\cos(2\pi ft) - i \sin(2\pi ft)] dt \quad (\text{F.5})$$

$$x(-f) = \int_{-\infty}^{\infty} x(-t) [\cos(2\pi ft)] = \int_{-\infty}^{\infty} x(t) [\cos(2\pi ft)] = x(f) \quad (\text{F.6})$$

where I used the fact that $x(-t) = x(t)$ for an even function. If we have an odd function, we have to use $x(-t) = -x(t)$, so we retain only the sin component and get $x(-f) = -ix(f)$, a phase shift of -90 degrees. This reflects the fact that the negative frequency coefficient is just the complex conjugate of the positive frequency coefficient. This also means they have the same norm, so we can write

$$\int_{-\infty}^{\infty} |x(f)| df = 2 \int_0^{\infty} |x(f)| df \quad (\text{F.7})$$

which we use in constructing our resolution calculations, which use the PSD, which contains the norm squared of the Fourier transform. It also allows us to focus only on the positive frequencies, even though the Fourier transform always has negative frequencies, as long as we remember to include these factors of 2.

F.1.2 Discrete Transform

The chi-square statistics presented in appendix [E](#) ignore the fact that for discrete data, when taking a sum, the expectation value of that sum will naturally depend on both the sampling frequency and trace length, as well as the proper normalization of all function involved. Let's define the discrete quantities so we can come up with some practical formulae for implementing them. The discretized fourier transform is

$$x(f) = \int_{-\infty}^{\infty} x(t)e^{-i\omega t} dt \approx \left[\sum_{n=0}^N x(t_n)e^{-i\omega t_n} \right] \Delta t$$

Suppose that t goes from 0 to T in steps of Δt . This lets us write

$$\omega t_n = \frac{\omega T}{N} n$$

The fact that we have N numbers spaced a distance Δt apart means that we can't arbitrarily sample frequency space using this data set, so we can sample N frequencies that are harmonics of the minimum sampling frequency from 0 to $N \frac{1}{T}$ to give

$$\omega_k = \frac{2\pi k}{T} \rightarrow \omega t = \frac{2\pi k n}{N}$$

this then gives us

$$x_k \approx \left[\sum_{n=0}^N x_n e^{-i \frac{2\pi k n}{N}} \right] \Delta t = x_D(f) \Delta t$$

where $x_D(f)$ is the form of the discrete Fourier transform computed by most FFT algorithms:

$$x_D(f) = \sum_{n=0}^N x_n e^{-i \frac{2\pi k n}{N}}$$

For your DFT to have the correct units, you need to multiply by the time step or divide by the sampling frequency. Taking the inverse transform, we can see that

$$x_n \approx \left[\sum_{k=0}^N x_k e^{i \frac{2\pi k n}{N}} \right] \Delta f = (x_D(t) N) \frac{1}{T} = x_D(t) \frac{1}{\Delta t}$$

We see that in one case we multiply by Δt and in another we divide by ΔT , so that in the discrete and continuous cases we get our original function from the inverse of the direct transform. This is only true because the inverse discrete transform is defined as

$$x_D(t) = \frac{1}{N} \left[\sum_{k=0}^N x_k e^{i \frac{2\pi k n}{N}} \right]$$

to ensure that it is the case. We can see now why not being aware of what we're calculating might trip us up; we need to remember when converting to a discrete or continuous formula how to scale the calculated transforms, when relevant.

F.1.3 Power Spectral Density

Let's talk now about calculating the Power Spectral Density (PSD) of some signal. The total energy in a signal is given as

$$E = \int_0^T (x(t))^2 dt = \int_0^{f_{max}} (x(f))^2 df$$

and the power is just this energy divided by the total time:

$$P = \frac{1}{T} \int_0^T (x(t))^2 dt = \frac{1}{T} \int_0^{f_{max}} (x(f))^2 df$$

thus the spectral density is just

$$J(f) = \frac{dP}{df} = \frac{x(f)^2}{T}$$

There is thus a convenient closed-form analytic expression for the discrete power spectral density, with correct normalization:

$$J(f) = \frac{1}{T} |x(f)|^2 = \frac{(\Delta t)^2}{T} |x_D(f)|^2$$

where $x_D(f)$ is the discrete Fourier transform discussed earlier. This is the simplest estimate for the PSD, but noting that this is an approximate form in the large T limit, there are better ways to estimate it with a small number of traces through more advanced Fourier methods. In MATLAB and python, the periodogram function is a very efficient way to estimate PSDs without having to worry about correct normalization and statistics, it will produce the lowest variance (though slightly biased) PSD estimate. In the limit of many traces averaged together, both methods converge to the same result.

The continuous definition is sometimes also given as

$$J(f) = \int_{-\infty}^{\infty} \gamma(\tau) e^{-i2\pi f\tau} d\tau$$

where $\gamma(\tau)$ is the autocovariance function; this means that the power spectral density is just the Fourier transform of the autocovariance function, defined as

$$\gamma(\tau) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T (x(t+\tau) - \mu)(x(t) - \mu) dt \approx \frac{1}{T} \int_0^T (x(t+\tau) - \mu)(x(t) - \mu) dt$$

where this is obviously an approximation for a truncated sample. We discuss this and related time-domain statistics in the next section of this appendix.

F.1.4 Correlation, Cross-Spectral Density, and Coherence

In the above, I'm concerned primarily with the relationship of noise to itself, but equally important statistics compare two traces to each other. Physicists refer to the amount of similarity or causation between two traces in the Fourier domain 'correlation', but in signal processing it is properly called the cross-spectral density, where the power-spectral density is then referred to as the auto-spectral density.

By the same logic as the previous section, we can define the cross-spectral density as the Fourier transform of the cross-correlation, or the spectral density of the product of the two signal transforms

in the infinite time limit:

$$\Sigma_{ij}(f) = \lim_{T \rightarrow \infty} \frac{x_i(f)^* x_j(f)}{T} \quad (\text{F.8})$$

The cross-spectral density matrix Σ_{ij} is therefore hermitian, given that the diagonals (the PSDs) are real and swapping the indices of the matrix gives the complex conjugate. We've seen this matrix already in Chapter 5 and Appendix E in the context of the correlated optimum filter, and we know that it allows us to diagonalize the traces into correlated and uncorrelated components.

If we want to measure the fraction of two traces that are causally linked, the relevant statistic is called the *coherence*, and is the Fourier-domain corollary of the time-domain correlation discussed in the next section. The coherence C_{ij} can be computed from the cross-spectral density matrix Σ as

$$C_{ij}(f) = \frac{|\Sigma_{ij}(f)|^2}{\Sigma_{ii}(f)\Sigma_{jj}(f)} \quad (\text{F.9})$$

For two completely correlated signals, the denominator will become Σ_{ii}^2 and the numerator also becomes Σ_{ii}^2 , giving a coherence of 1. In the infinite time limit, two completely uncorrelated signals will have no cross-spectral density, and the coherence will be 0. This allows for a much easier way to visualize which regions of a signal might be influenced by an environmental noise source which is coupled either to another channel or an environmental monitor as a function of frequency.

It should be noted that these statistics really do rely on having well-sampled spectra, and will be affected by measurement error much more noticeably than time-domain statistics might. It is always useful as a cross-check of your implementation to verify that the Fourier transform of the cross-correlation is approximately the same as the discrete Fourier version:

$$C_{ij}(f) = \lim_{T \rightarrow \infty} \int_{-\infty}^{\infty} \int_0^T [(x_i(t + \tau) - \langle x \rangle)(x_j(t) - \langle x_j \rangle)] e^{-i2\pi f \tau} dt d\tau \quad (\text{F.10})$$

In general using a pre-built function to do windowed estimation will produce a lower variance result, similar to the case with the PSD. These methods use adjacent bins to reduce the over-all variance of the CSD by allowing the result to have a slight bias in the integral, but give a much better estimate of a given CSD point.

F.1.5 Discrete Optimum Filter

Suppose we start with the continuous definition from the previous section:

$$\chi^2 = \int_{-\infty}^{\infty} \frac{|v(f) - As(f)|^2}{J(f)} df \quad (\text{F.11})$$

We make this discrete, and we want to express it in terms of discrete Fourier transforms of the signal, template, and noise. We thus get

$$\chi^2 \approx \sum_{n=0}^N \frac{|v(f) - As(f)|^2}{J(f)} \Delta f \quad (\text{F.12})$$

$$= \sum_{n=0}^N \frac{(\Delta t)^2 |v_D(f) - As_D(f)|^2}{\frac{(\Delta t)^2}{T} |n_D(f)|^2} \frac{1}{T} \quad (\text{F.13})$$

$$= \sum_{n=0}^N \frac{|v_D(f) - As_D(f)|^2}{|n_D(f)|^2} \quad (\text{F.14})$$

which should always come out to have an expectation value of $\sim N$, and thus a reduced chi-square of 1. It is usually unlikely that we have the Fourier transform of the noise, however, and it's more likely that we have the periodogram of the noise and the fourier transforms of the signal and template, so we should re-write the chi-square as

$$\chi^2 \approx \frac{(\Delta t)^2}{T} \sum_{n=0}^N \frac{|v_D(f) - As_D(f)|^2}{J(f)} \quad (\text{F.15})$$

$$= \frac{\Delta t}{N} \sum_{n=0}^N \frac{|v_D(f) - As_D(f)|^2}{J(f)} \quad (\text{F.16})$$

this odd form has to do with the different scaling of a true discrete transform and a continuous approximation, and is important for practically implementing the optimal filter chi-square such that the reduced chi-square is insensitive to trace length or sampling rate. This same scaling also then applies to the more advanced filters as the overall scaling doesn't change, just the complexity of the template.

Finally, for a reduced chi-square, we find that the above equations become

$$\chi_r^2 = \frac{1}{N - M} \sum_{n=0}^N \frac{|v_D(f) - As_D(f)|^2}{|n_D(f)|^2} \quad (\text{F.17})$$

and

$$\chi_r^2 = \frac{\Delta t}{N(N - M)} \sum_{n=0}^N \frac{|v_D(f) - As_D(f)|^2}{J(f)} \quad (\text{F.18})$$

where for the single-amplitude model $M = 1$.

F.2 Time-Domain Statistics

In analyzing time-domain data, there are a number of useful statistics which allow us to quickly determine whether we have a trace with a pulse, on a large thermal tail, or only noise, and an analysis

of inliers and outliers in these distributions help separate types of events and remove unwanted events which are not characteristic of the signal under study. Some basic moments for any time-domain series are given below:

1. Mean:

$$\hat{\mu} = \frac{1}{N} \sum_{n=1}^N x_n$$

2. Variance:

$$\hat{\sigma}^2 = \frac{1}{N-1} \sum_{n=1}^N (x_n - \mu)^2$$

3. Skewness

$$\hat{\gamma} = \frac{1}{N\sigma^{3/2}} \sum_{n=1}^N (x_n - \mu)^3$$

4. Slope

$$\hat{m} = \frac{\sum_{n=1}^N (x_n - \mu_x)(y_n - \mu_y)}{\sum_{n=1}^N (x_n - \mu_x)^2}$$

Here I have made sure to note that these are not the true values but estimates of the values, each with their own uncertainty, but exact in the limit that $N \rightarrow \infty$.

F.2.1 Covariance and Correlation

The above statistics are all of one dimension and characterize global quantities of the dataset. Another statistic which is one-dimensional (assuming stationary noise) is the autocovariance of the time series which we have already defined previously as

$$\gamma(\tau) = \frac{1}{T} \int_0^T (x(t+\tau) - \mu)(x(t) - \mu) dt$$

The discrete version of this is just as simple

$$\gamma_i = \frac{1}{N} \sum_{n=1}^N (x_{n+i} - \mu)(x_n - \mu)$$

You should notice immediately that this looks familiar, and that $i = 0$ is just the variance:

$$\gamma_0 = \frac{1}{N} \sum_{n=1}^N (x_n - \mu)^2 = \sigma^2$$

This is a statistic which is less standardized in its continuous form due to differences between statistical and signal processing applications, and here we will restrict ourselves to the statistical

because we will be concerned with total energy, not steady-state power. Finite power processes formally have infinite energy and thus are irrelevant to our studies, as we will always want to subtract the DC power from our traces.

An alternate way to compute this statistic is by the inverse Fourier transform of the power spectral density:

$$\gamma(\tau) = \int_{-\infty}^{\infty} J(f) e^{i2\pi f\tau} df$$

This is related to the definition from the previous section of the power spectral density as the Fourier transform of the autocovariance function, and is proven by the Wiener-Kinchin theorem (which I will not reproduce but reference for a complete discussion). For completeness, let's also tackle the case of a discrete power spectral density. Using our normal substitutions, we have

$$\gamma(\tau) \approx \frac{1}{T} \sum_{n=0}^N J(f_n) e^{i2\pi f_n \tau} = \frac{J_D(\tau)}{\Delta t}$$

which is just our normal relation for converting the discrete inverse transform to the continuous approximation. We can find the method of computing the autocovariance by Fourier methods by substituting the discrete approximation for the power spectral density:

$$\gamma(\tau) \approx (|x_D(f)|^2)_D(\tau)$$

The notation here is a bit messy, but this says that the autocovariance is just the inverse transform of the squared Fourier transform. This is exactly what we should have expected given the properties of the Fourier transform in which a convolution in time domain is a multiplication in Fourier domain.

For completeness, let's also summarize the covariance more generally. The covariance of two signals is given as

$$\text{cov}(x, y)(\tau) = \frac{1}{T} \int_0^T (x(t + \tau) - \mu_x)(y(t) - \mu_y) dt$$

and is thus formally the *time-averaged cross-correlation* of the two signals after removing their means.

Why would we want to time-average this cross-correlation, and why remove the means? Consider what would happen if the signals were noiseless, and flat-lined at their mean. In this form, the covariance is 0, but the cross-correlation is $\mu_x \mu_y T$, and the time-averaged cross-correlation is just $\mu_x \mu_y$. We know that there is no co-variance, because the signal does not properly vary when it flat-lines, but only the covariance tells us this is true. This is one of the reasons why it is always necessary to remove the baseline of two signals before comparing them; otherwise the mean will skew your estimate of the relationship between them.

We normally work with covariance using the covariance matrix, and in this case, the covariance

matrix is given as

$$\Sigma_{ij} = \frac{1}{N} \sum_{n=1}^N (x_{n+i} - \mu_x)(y_{n+j} - \mu_y)$$

and therefore the autocovariance matrix is

$$\Sigma_{ij} = \frac{1}{N} \sum_{n=1}^N (x_{n+i} - \mu_x)(x_{n+j} - \mu_x)$$

such that $\Sigma_{ii} = \sigma^2$. Noise that is completely uncorrelated will have a diagonal autocovariance matrix of the form $\Sigma = I\sigma^2$ and a covariance matrix $\Sigma = 0$. Traces with correlated noise will still have a high degree of symmetry, and the autocovariance matrix is just the Toeplitz matrix with the first row and column equal to γ_i :

$$\Sigma_{ij} = \gamma_{i-j}$$

This will be useful in a few sections when we discuss using the autocovariance matrix for optimal linear least-squares estimates.

A related statistic is the autocorrelation function. Correlation is defined generally as

$$\text{corr}(x, y) = R_{xy} = \frac{\text{cov}(x, y)}{\sigma_x \sigma_y}$$

so the autocorrelation of a signal reduces to

$$R_{xx}(\tau) = \frac{\gamma(\tau)}{\sigma^2}$$

As we saw earlier, $\gamma(0) = \sigma^2$, so we find trivially that $R_{xx}(0) = 1$, or a trace is 100% correlated with itself at no time shift.

F.2.2 Example: Correlation Triggers

We can use the correlation to define our trigger, if one time-series contains data and the other an accurate signal template. To see why, imagine that we have some signal templates $s(t)$ and a signal $As(t) + n(t)$, where $n(t)$ is drawn from a distribution with mean 0 and variance σ^2 . The correlation

(trigger signal T) of the signal and template is

$$T(t) = \frac{1}{T\sqrt{A^2\sigma_s^2 + \sigma_n^2\sigma_s}} \int_0^T (As(t+\tau) + n(t+\tau))(s(t))dt \quad (\text{F.19})$$

$$= \frac{1}{T\sqrt{A^2\sigma_s^2 + \sigma_n^2\sigma_s}} \left[\int_0^T As(t+\tau)s(t)dt + \int_0^T n(t+\tau)s(t)dt \right] \quad (\text{F.20})$$

$$= \frac{1}{T\sqrt{\sigma_s^2 + A^{-2}\sigma_n^2\sigma_s}} \int_0^T s(t+\tau)s(t)dt \quad (\text{F.21})$$

$$= \frac{\sigma_s^2}{\sqrt{\sigma_s^2 + A^{-2}\sigma_n^2\sigma_s}} f(\tau) \quad (\text{F.22})$$

$$= \frac{1}{\sqrt{1 + \frac{\sigma_n^2}{\sigma_s^2 A^2}}} R_{ss}(\tau) \quad (\text{F.23})$$

where we've substituted in the autocorrelation function, which we know must be unity at $\tau = 0$. So the maximum of the signal $T(t)$ will occur at the true start point of the signal template, and the larger the template is compared to the noise, the closer this maximum will be to 1. Setting a level threshold on this function can then be related to the variances given here, where σ_s^2 is the variance of the template and σ_n^2 is the noise variance. We can relate the threshold η in a correlation trigger to these quantities:

$$\frac{\sigma_n^2}{A^2\sigma_s^2} = \frac{1}{\eta^2} - 1$$

Larger amplitudes and signal variances thus allow more stringent thresholds, which matches our intuition. The exact triggering strategy of course depends on the details of these parameters as they relate to the requirements of the experiment.

We can also chain the correlations to create a trigger more robust to signal-like noise fluctuations by multiplying multiple correlation triggers; in a sense this is asking the question "how likely is it that the same event caused a feature in all channels considered"? Because we know that the trigger point will occur where correlation is essentially 1, random excursions will be suppressed given that they will not occur at the same point each trace. This will only be the case, however, for *uncorrelated noise*. Noise that is correlated between channels will require additional modeling or filtering to remove.

F.3 Least Squares Fitting

In experimental physics we do a lot of curve-fitting, and certain terms like χ^2 are thrown around a lot without a ton of formal definition or explanation as to why this random calculation should give us something meaningful. In this section I will lay a bit of the foundations for the χ^2 statistic for completeness, and demonstrate the closed-form solution for a linear least-squares problem. I briefly

mention non-linear least-square to acknowledge that most problems are non-linear, but I won't go into much detail in this thesis.

F.3.1 Reduced Chi-Square Statistic

The weighted, discrete least-squares statistic is normally given as

$$\chi^2 = \sum_{i=0}^N W_i (y_i - f(x_i))^2 = \sum_{i=0}^N W_i r^2$$

where r is the residual (observation minus prediction) and normally we set the weights to the inverse of the variance of each measurement; in the case of stationary uncorrelated noise, then the time-domain variance should be the same, and we have $W_i = \sigma^{-2}$, and we get

$$\chi^2 = \frac{1}{\sigma^2} \sum_{i=0}^N r_i^2$$

let's ask ourselves what this means in the context of a DC signal of mean μ , where x_i is just samples drawn from the noise added to this mean. In this case, then we find the χ^2 statistic

$$\langle \chi^2 \rangle = \left\langle \frac{1}{\sigma^2} \sum_{i=0}^N (x_i - \mu)^2 \right\rangle = \frac{1}{\sigma^2} N \sigma^2 = N$$

This is a nice property; we know for a sample of length N , a good fit should give a distribution of chi-squares which has a mean equal to the length of the trace. If we want to make this statistic more meaningfully invariant to the trace length, we can just divide by the trace length to generate the reduced chi-square, χ_r^2 :

$$\chi_r^2 = \frac{1}{N \sigma^2} \sum_{i=0}^N r_i^2$$

This statistic will be near 1 for a good fit, and larger than 1 for a bad fit; it will also indicate when we've over-fit the model in the case that it's much less than 1, which shouldn't be possible, but is allowable for an over-fit model.

Let's explore this in the limit of a many parameter model. In the case that we have M parameters and N data points, we realize that we are reducing the amount of prediction we need to do to a set of $N - M$ observations; in the limit of $M = N - 1$, we can set all but one of the parameters to equal an observation and let the last point predict the average of the remaining two points, for example, and get a stellar reduced chi-square. This of course would be over-fit, and give a statistic much less

than 1. This is typically mitigated by dividing not by N but by $N - M$, such that

$$\chi_r^2 = \frac{1}{(N - M)\sigma^2} \sum_{i=0}^N r_i^2$$

This is the definition of reduced chi-square we will use, although we will sometimes approximate $M \approx 0$ in the case that we have many more points than model parameters for simplicity.

F.3.2 Generalized Least Squares

Let's write the reduced least-squares statistic in matrix form. We can re-express the weights as the diagonal matrix $\mathbf{W} = \sigma^{-2}\mathbf{I}$, and we get the weighted least-squares expression

$$\chi_r^2 = \frac{1}{N} \mathbf{R}^T \mathbf{W} \mathbf{R}$$

where $\mathbf{R} = \mathbf{y} - f(\mathbf{x})$. This is very suggestive; we realize what we've been doing is a very simplified case of a least-squares minimization. Without any further generalization, let's imagine that \mathbf{W} is any diagonal matrix. We can rotate to a basis where \mathbf{W} is non-diagonal by a unitary transform \mathbf{U} as

$$\begin{aligned} \chi_r^2 &= \frac{1}{N} \mathbf{R}^T (\mathbf{U}^T \mathbf{U}) \mathbf{W} (\mathbf{U}^T \mathbf{U}) \mathbf{R} \\ &= \frac{1}{N} (\mathbf{R}^T \mathbf{U}^T) (\mathbf{U} \mathbf{W} \mathbf{U}^T) (\mathbf{U} \mathbf{R}) \\ &= \frac{1}{N} \mathbf{R}'^T \boldsymbol{\Sigma} \mathbf{R}' \end{aligned}$$

where $\boldsymbol{\Sigma} = \mathbf{U} \mathbf{W} \mathbf{U}^T$ is a symmetric, positive definite, non-diagonal matrix. Why is this important? Because it says that for any set of observations with a covariance matrix $\boldsymbol{\Sigma}$, the same justification for the linear least-squares exists, but we're implicitly doing it in a non-diagonal basis.

F.3.3 Linear Least-Squares: Normal Equations

Suppose that our model is linear, such that it can be expressed in the form

$$f(\mathbf{X}) = \sum_j^M \beta_j S_j(\mathbf{X}) = \mathbf{S} \boldsymbol{\beta} \tag{F.24}$$

where $S_j(\mathbf{X})$ is a known model vector and a column in the matrix \mathbf{S} . $\boldsymbol{\beta}$ is therefore a column vector with length M , and \mathbf{S} is a matrix of size $N \times M$, such that their product is a column vector of length N .

We can minimize the least-squares formula with respect to each parameter by setting the respective derivative to 0:

$$0 = \frac{\partial \chi^2}{\partial \beta_j} = \frac{\partial}{\partial \beta_j} \frac{1}{N} \mathbf{R}^T \mathbf{W} \mathbf{R} \quad (\text{F.25})$$

$$0 = \frac{2}{N} \mathbf{R}^T \mathbf{W} \left[\frac{\partial}{\partial \beta_j} \mathbf{R} \right] \quad (\text{F.26})$$

$$0 = \mathbf{R}^T \mathbf{W} [-S_j(\mathbf{X})] \quad (\text{F.27})$$

$$0 = S_j(\mathbf{X})^T \mathbf{W} \mathbf{R} \quad (\text{F.28})$$

$$0 = S_j(\mathbf{X})^T \mathbf{W} \left(\mathbf{Y} - \sum_j^M \beta_j \mathbf{S}_j(\mathbf{X}) \right) \quad (\text{F.29})$$

$$S_j(\mathbf{X})^T \mathbf{W} \mathbf{Y} = S_j(\mathbf{X})^T \mathbf{W} \sum_k^M \beta_k \mathbf{S}_k(\mathbf{X}) \quad (\text{F.30})$$

This allows us to write the normal equations

$$\mathbf{S}^T \mathbf{W} \mathbf{S} \boldsymbol{\beta} = \mathbf{S}^T \mathbf{W} \mathbf{Y} \quad (\text{F.31})$$

or in other words, estimate all of the linear parameters with one matrix calculation:

$$\hat{\boldsymbol{\beta}} = \frac{\mathbf{S}^T \mathbf{W} \mathbf{Y}}{\mathbf{S}^T \mathbf{W} \mathbf{S}} = \mathbf{F} \mathbf{Y} \quad (\text{F.32})$$

where

$$\mathbf{F} = (\mathbf{S}^T \mathbf{W} \mathbf{S})^{-1} \mathbf{S}^T \mathbf{W} \quad (\text{F.33})$$

is the transformation matrix and

$$\mathbf{A} = (\mathbf{S}^T \mathbf{W} \mathbf{S})^{-1} \quad (\text{F.34})$$

is the design matrix of our problem. This looks awfully familiar; it is in fact the time-domain version of the no-offset optimum filter. Because this is equally valid for the generalized least-squares, we can write

$$\boldsymbol{\Phi}_T = (\mathbf{S}^T \boldsymbol{\Sigma}^{-1} \mathbf{S})^{-1} \mathbf{S}^T \boldsymbol{\Sigma}^{-1} \quad (\text{F.35})$$

and make this explicitly our optimal filter in the time-domain. This actually a really nice result, because it will allow us to avoid doing any Fourier transforms at the expense of one single (but larger) matrix calculation by performing the operation

$$\hat{\boldsymbol{\beta}} = \boldsymbol{\Phi}_T \mathbf{Y} \quad (\text{F.36})$$

The one down-side of this method is that it requires us to have a linear model. Our model is linear in

the pulse amplitudes, but not in the time-offsets, so a full linearized analysis of an arbitrary number of pulses is possible only if the arrival times are known. If only approximate arrival times are known, a linearization can be done in the time offset by using the model and its derivative as suggested in [38].

The nice feature of this formalism for optimal filtering is that there is a huge body of literature and set of libraries which can manipulate linear systems and solve the normal equations efficiently. We also conveniently get estimates of the parameter uncertainties and covariances from this formalism very easily, as well as an estimate of goodness of fit^[1]:

$$\mathbf{A} = (\mathbf{S}^T \boldsymbol{\Sigma}^{-1} \mathbf{S}) \quad (\text{F.37})$$

$$\boldsymbol{\Phi}_T = (\mathbf{S}^T \boldsymbol{\Sigma}^{-1} \mathbf{S})^{-1} \mathbf{S}^T \boldsymbol{\Sigma}^{-1} = \mathbf{A}^{-1} \mathbf{S}^T \boldsymbol{\Sigma}^{-1} \quad (\text{F.38})$$

$$\boldsymbol{\beta} = \boldsymbol{\Phi} \mathbf{Y} \quad (\text{F.39})$$

$$\boldsymbol{\Sigma}_\beta = [\mathbf{S}^T \boldsymbol{\Sigma}^{-1} \mathbf{S}]^{-1} = \mathbf{A}^{-1} \quad (\text{F.40})$$

$$\chi_{r,0} = \frac{\mathbf{Y}^T \boldsymbol{\Sigma}^{-1} \mathbf{Y}}{N - M} \quad (\text{F.41})$$

$$\chi_r^2 = \frac{1}{N - M} (\mathbf{Y} - \mathbf{S}\boldsymbol{\beta})^T \boldsymbol{\Sigma}^{-1} (\mathbf{Y} - \mathbf{S}\boldsymbol{\beta}) \quad (\text{F.42})$$

$$= \chi_{r,0} - \frac{\boldsymbol{\beta}^T \mathbf{A} \boldsymbol{\beta}}{N - M} \quad (\text{F.43})$$

This last equation follows from

$$\chi_r^2 = \frac{1}{N - M} (\mathbf{Y} - \mathbf{S}\boldsymbol{\beta})^T \boldsymbol{\Sigma}^{-1} (\mathbf{Y} - \mathbf{S}\boldsymbol{\beta}) \quad (\text{F.44})$$

$$= \frac{1}{N - M} (\mathbf{Y}^T \boldsymbol{\Sigma}^{-1} \mathbf{Y} - 2(\mathbf{S}\boldsymbol{\beta})^T \boldsymbol{\Sigma}^{-1} \mathbf{Y} + \boldsymbol{\beta}^T \mathbf{A} \boldsymbol{\beta}) \quad (\text{F.45})$$

$$= \frac{1}{N - M} (\mathbf{Y}^T \boldsymbol{\Sigma}^{-1} \mathbf{Y} - 2\boldsymbol{\beta}^T \mathbf{A} \boldsymbol{\beta} + \boldsymbol{\beta}^T \mathbf{A} \boldsymbol{\beta}) \quad (\text{F.46})$$

$$= \frac{1}{N - M} (\mathbf{Y}^T \boldsymbol{\Sigma}^{-1} \mathbf{Y} - \boldsymbol{\beta}^T \mathbf{A} \boldsymbol{\beta}) \quad (\text{F.47})$$

What is interesting here is that once I define my model, and I know my noise covariances, I already can predict my model uncertainty and covariance between parameters. That is an intuitively sensible statement but a nice demonstration of the principle we discuss in optimal filtering theory earlier. We also see that this is a much more convenient way to write the same procedure as outlined earlier for fixed pulse start-times.

¹[https://en.wikipedia.org/wiki/Linear_least_squares_\(mathematics\)](https://en.wikipedia.org/wiki/Linear_least_squares_(mathematics))

F.3.4 Nonlinear Models and Maximum Likelihood

What do we do when our model is not just a linear combination of weights of our parameters? We generally have two options. In many cases, the problem can be linearized by projecting the parameters into a linear space. Functions of the form

$$f(x) = ax^b \tag{F.48}$$

for example are not implicitly linear, but they can be linearized by instead fitting the log values of the data:

$$f'(x) = \log_{10}(f(x)) = \log_{10}(a) + b \log_{10}(x) \tag{F.49}$$

This means that any power law, even with multiple parameters, can be turned into a linear least-squares fit.

This only works for a handful of functions though, and usually we have to resort to a maximum likelihood fit. The idea behind the maximum likelihood is that we have a model that gives the probability of a given observation for a set of parameters, $P(x; \beta)$, and we can then construct the function

$$\mathcal{L} = \prod_i P(x_i; \beta) \tag{F.50}$$

which is the likelihood of observing the set of observations x_i given a set of parameters. In some cases we can find a closed-form solution to problems of this form; for example, the time-domain statistics in this appendix are the maximum likelihood estimates for the moments of a normal distribution. In most cases, we end up using a gradient solver to numerically find the parameters which maximize the likelihood. In general, we actually fit the log-likelihood

$$\log [\mathcal{L}] = \sum_i \log(P(x_i; \beta)) \tag{F.51}$$

which looks awfully similar to our linear least-squares problems! Taking the log allows us to probe a smoother space and increases the speed of our computations; it's faster to add a bunch of numbers than to multiply them.

Many good references on machine learning introduce different minimization methods in a robust way, but for a non-linear likelihood space this becomes a very nontrivial problem. That's generally the purview of observational cosmology, where very non-linear problems occupy highly nonlinear and high-dimensional parameter spaces that need to be robustly constrained. Andrew Ng's class notes for the Stanford machine learning class (which I hope will be turned into a book at some point) provide an excellent introduction to these types of problems. In many cases, training a machine learning algorithm is itself a non-linear optimization problem, and many machines are attempting to learn solutions to non-linear problems, so it's a bit of a spiral into function optimization.

Appendix G

Single Electron Probability Distributions

What follows are notes deriving different probability distributions for electron-hole pairs generated by a laser of known Poisson mean λ under different experimental conditions. I provide it here for posterity, but it may not be entirely pedagogical because these are mostly the unrefined notes I was taking when I derived different PDFs.

G.1 Abstract Form

We would like to fit the laser data by the method of maximum likelihood, where the likelihood function $g(S_i; q)$ of a signal point S_i given fit parameters q is

$$g(S_i; q, \lambda) = \sum_{n=0}^{\infty} \frac{\lambda^n e^{-\lambda}}{n!} P_n(S_i; q)$$

where $P_n(S_i, q)$ is the probability distribution for energies produced by n photons given the fit parameters. While we have a general form for this distribution, assuming all processes produce Gaussian distributed random variates gives the form

$$P_n(S) = \sum_{k_1, k_2, k_3=0}^n \frac{n!}{k_1! k_2! k_3!} \frac{(\beta\alpha)^{k_1} (\beta(1-\alpha))^{k_2} (1-\beta)^{k_3}}{\sqrt{2\pi v_k}} \exp\left[-\frac{(S - \mu_k)^2}{(2v_k)}\right]$$

where α is the probability of crossing the crystal without trapping/impact ionization, β the probability of escaping the initial charge cloud, and $n = k_1 + k_2 + k_3$. The variance and mean in this term

are also functions of k_i with free parameters:

$$v_k = v_N + k_1 v_1 + k_2 v_2 + k_3 v_3$$

$$\mu_k = \mu_N + k_1 \mu_1 + k_2 \mu_2 + k_3 \mu_3$$

If we were to simultaneously fit all of these parameters, it would be a 11 parameter fit with a lot of degeneracy, and the odds of finding the correct minimum would be very small. We can reduce the fit complexity by using leakage traces and trends across voltage to constrain our parameters. To start, consider the leakage runs, which have $\lambda = 0$. This restricts $n = k_1 + k_2 + k_3 = 0$, allowing us to fit v_n and μ_n while ignoring the remaining parameters. We thus fix these parameters to the mean and standard deviation of the Gaussian noise distribution from the leakage.

Now if we turn the laser on, but we assume $\alpha = 1$, we have three more parameters to fit, namely λ , v_1 , and μ_1 . By setting α to 1 we eliminate 5 parameters (we implicitly eliminate β , and all means/variances associated with k_2 and k_3). We thus fit the function

$$P_n(S) = \sum_{k_1=0}^n \frac{1}{\sqrt{2\pi v_k}} \exp\left[-\frac{(S - \mu_k)^2}{(2v_k)}\right]$$

which is the sum of Gaussians with linearly increasing mean and variance. Assuming that the dominant part of the spectrum comes from the charges which cross the crystal, these parameters can also be fixed after this fit.

We now relax the assumption of $\alpha = 1$ but impose $\beta = 1$, and we fit three more parameters, α and the associated mean and variance. We can obtain a good fit here by restricting the mean and variance to be a fraction of the previously fitted mean and variance given that this should be a second order effect.

G.2 Physical Trapping/Ionization Model

G.2.1 Single Charge

Suppose we have a process that stops or frees a charge somewhere in the bulk of the detector. If we have perfect energy resolution, we would find the probability distribution

$$p(s_0; d) = \frac{\Theta(s_0)\Theta(d - s_0)}{d}$$

where $\Theta(x)$ is the Heaviside step function and d is the full energy of an electron-hole pair crossing the length of the detector. We also have gaussian noise producing the true measurable s , with the

probability distribution

$$p(s; s_0, \sigma) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(s-s_0)^2}{2\sigma^2}\right]$$

To get the observed distribution of s then, given σ and d , we perform the convolution of these two distributions:

$$p_1(s; d, \sigma) = \int_{-\infty}^{\infty} \frac{\Theta(s_0)\Theta(d-s_0)}{d} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(s-s_0)^2}{2\sigma^2}\right] ds_0 \quad (\text{G.1})$$

$$= \frac{1}{d} \int_0^d \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(s-s_0)^2}{2\sigma^2}\right] ds_0 \quad (\text{G.2})$$

$$= \frac{1}{2d} \left[\Phi\left(\frac{s}{\sqrt{2}\sigma}\right) - \Phi\left(\frac{s-d}{\sqrt{2}\sigma}\right) \right] \quad (\text{G.3})$$

where $\Phi(x)$ is the error function. This makes intuitive sense; our distribution looks like the original distribution but with edges rounded on the length scale of the gaussian variance.

G.2.2 Two Charges

Suppose now that we have two charges which trap or are freed. Each one has its own interaction site, so we have true energies s_1 and s_2 , but we measure (in the absence of noise) $s_0 = s_1 + s_2$. So we want to find the probability distribution of s_0 . We have

$$p_2(s_0; d) = \int_{-\infty}^{\infty} p(s_1; d)p(s_0 - s_1; d) ds_1 \quad (\text{G.4})$$

$$= \frac{1}{d^2} \int_{-\infty}^{\infty} \Theta(s_1)\Theta(d-s_1)\Theta(s_0-s_1)\Theta(d+s_1-s_0) ds_1 \quad (\text{G.5})$$

If $s_0 < d$, we have

$$p_2(s_0; d) = \frac{1}{d^2} \int_0^{s_0} \Theta(d-s_1)\Theta(d+s_1-s_0) ds_1 \quad (\text{G.6})$$

$$= \frac{s_0}{d^2} \Theta(s_0)\Theta(d-s_0) = \frac{s_0}{d} p(s_0; d) \quad (\text{G.7})$$

If $s_0 > d$, we have

$$p_2(s_0; d) = \frac{1}{d^2} \int_0^d \Theta(s_0-s_1)\Theta(d+s_1-s_0) ds_1 \quad (\text{G.8})$$

$$= \frac{2d-s_0}{d^2} \Theta(s_0-d)\Theta(2d-s_0) = \left(2 - \frac{s_0}{d}\right) p(s_0-d; d) \quad (\text{G.9})$$

which gives the total probability distribution

$$p_2(s_0; d) = \frac{s_0}{d} p(s_0; d) + \left[2 - \frac{s_0}{d}\right] p(s_0-d; d) = \frac{s_0}{d} [p(s_0; d) - p(s_0-d; d)] + 2p(s_0-d; d)$$

I've written the function this way in order to make convolution with the noise distribution easier; we've already done the integral with the single-charge probability distribution. We find that we have to do the convolution

$$p_2(s; d) = \int_{-\infty}^{\infty} p_2(s_0; d) p(s; s_0, \sigma) ds_0 \quad (\text{G.10})$$

$$= \int_{-\infty}^{\infty} \left[\frac{s_0}{d} [p(s_0; d) - p(s_0 - d; d)] + 2p(s_0 - d; d) \right] \frac{1}{\sqrt{2\pi\sigma^2}} \exp \left[\frac{-(s - s_0)^2}{2\sigma^2} \right] ds_0 \quad (\text{G.11})$$

$$= \frac{s}{d} [p_1(s; d) - p_1(s - d; d)] + 2p_1(s - d; d) - \frac{1}{d} \int_{-\infty}^{\infty} [p_1(s_0; d) - p_1(s_0 - d; d)] ds_0 \quad (\text{G.12})$$

We've saved some effort by re-using the earlier integral. The function in brackets can be simplified:

$$p_1(s; d) - p_1(s - d; d) = \frac{1}{2d} \left[\Phi \left(\frac{s}{\sqrt{2\sigma}} \right) - \Phi \left(\frac{s - d}{\sqrt{2\sigma}} \right) \right] - \frac{1}{2d} \left[\Phi \left(\frac{s - d}{\sqrt{2\sigma}} \right) - \Phi \left(\frac{s - 2d}{\sqrt{2\sigma}} \right) \right] \quad (\text{G.13})$$

$$= \frac{1}{2d} \left[\Phi \left(\frac{s}{\sqrt{2\sigma}} \right) - 2\Phi \left(\frac{s - d}{\sqrt{2\sigma}} \right) + \Phi \left(\frac{s - 2d}{\sqrt{2\sigma}} \right) \right] \quad (\text{G.14})$$

such that our leading term is then

$$\frac{s}{2d^2} \left[\Phi \left(\frac{s}{\sqrt{2\sigma}} \right) - 2\Phi \left(\frac{s - d}{\sqrt{2\sigma}} \right) + \Phi \left(\frac{s - 2d}{\sqrt{2\sigma}} \right) \right] + \frac{1}{d} \left[\Phi \left(\frac{s - d}{\sqrt{2\sigma}} \right) - \Phi \left(\frac{s - 2d}{\sqrt{2\sigma}} \right) \right]$$

In the limit that $\sigma \rightarrow 0$, this reduces to our original triangular function. We also have the second term, which should dominate in the large σ limit. The full convolution comes out to

$$p_2(s; d) = \frac{1}{2d^2} \left[s\Phi \left(\frac{s}{\sqrt{2\sigma}} \right) + 2(d - s)\Phi \left(\frac{s - d}{\sqrt{2\sigma}} \right) + (s - 2d)\Phi \left(\frac{s - 2d}{\sqrt{2\sigma}} \right) \right] \quad (\text{G.15})$$

$$+ \frac{1}{d^2} \sqrt{\frac{\sigma^2}{2\pi}} \left[\exp \left(\frac{s^2}{2\sigma^2} \right) - 2 \exp \left(\frac{(s - d)^2}{2\sigma^2} \right) + \exp \left(\frac{(s - 2d)^2}{2\sigma^2} \right) \right] \quad (\text{G.16})$$

This is our smoothed triangle plus Gaussians at the ends and middle of the triangle.

G.2.3 Three or more charges

We can use the previous two sections to derive the probability distribution for 3 simultaneous charge leakage events. If we have $s_0 = s_1 + s_2 + s_3$, and we know that we can use p_2 from the previous section for the probability distribution of $s_1 + s_2$, then we have

$$p_3(s_0; d) = \int_{-\infty}^{\infty} p_2(s_1; d) p_1(s_0 - s_1; d) ds_1$$




Figure G.1: Comparison of exact 3-square convolution to Gaussian approximation.

without getting into the gory details, repeating the previous procedure (break it down into individual convolutions of two $p(x; d)$ functions) gives the piecewise form

$$p_3(s_0; d) = \frac{1}{2d^2} [s_0^2 p(s_0; d) + (s_0^2 - 3(s_0 - d)^2)(p(s_0 - d; d) + (s_0 - 3d)^2 p(s_0 - 2d; d))]$$

which looks remarkably similar to a Gaussian with mean $\mu = 3(d/2)$ and $\sigma = 3(d^2/12)$, as shown in figure [G.1](#). These are just the variance and mean of a normalized flat distribution with bounds of 0 and $d!$ This suggests that for $k \geq 2$ trapping events, we can approximate the resulting energy deposition PDF as a gaussian with $\mu = kd/2$ and $\sigma = kd^2/12$. In other words, we have the piecewise function for the PDF of deviation from expected energy for arbitrary k :

$$p_k(s_0; d) = \begin{cases} \delta(s_0) & k = 0 \\ p(s_0; d) & k = 1 \\ \frac{s_0}{d} [p(s_0; d) - p(s_0 - d; d)] + 2p(s_0 - d; d) & k = 2 \\ [\pi kd^2/6]^{-1/2} \exp\left[\frac{-(s_0 - kd/2)^2}{kd^2/6}\right] & k \geq 3 \end{cases} \quad (\text{G.17})$$

We have the equivalent piecewise function for convolution with the noise Gaussian, which is trivial for $k = 0$ and $k = 3$ (the latter being a simple addition of variances) and done above for the other two cases. With these forms we have a very good approximation for trapping or impact ionization of an arbitrary number of charges.

G.2.4 Charge Energy Distributions

Let's go back to what we initially proposed in the first part of this note and modify it with our more specific trapping forms to see how the model changes. We will ignore β here ($\beta = 1$) and still posit that α measures the probability of a charge making it across the crystal. Consider first the case of a single charge generated at the crystal surface. If we call the trapping probability distributions p_i and the no-trapping distribution p_0 , then we have the total probability distribution

$$p_1(s; d, \sigma) \propto \alpha p_0(s - d; d, \sigma) + (1 - \alpha) p_1(s; d, \sigma)$$

We can easily extend this to have the more general case of n initial charges by analogy to the first equation in the note, giving

$$p_n(s; d, \sigma) = \sum_{k=0}^n \frac{n!}{k!(n-k)!} \alpha^{n-k} (1 - \alpha)^k p_k(s + (k - n)d; d, \sigma)$$

Clearly for this to be valid for any number of charges, we need the distributions p_n out to n . If we consider that the trapping probability is small, however, we can treat this as a Taylor expansion and recognize that beyond $k = 2$ we add sub-percent corrections; so the approximate probability distribution is thus

$$p_n(s; d, \sigma) \approx \sum_{k=0}^2 \frac{n!}{k!(n-k)!} \alpha^{n-k} (1 - \alpha)^k p_k(s + (k - n)d; d, \sigma)$$

in the limit $(1 - \alpha)^2 \ll 1$. Note that for e.g. $\alpha = 0.8$, the $k = 3$ term would have a relative magnitude of 0.8%.

We thus string the various p_n distributions together as described earlier in the note, and we have a trapping model with many fewer free parameters, given that we're assuming p_0 is Gaussian distributed. We add back in the increasing variance with charge number by making the substitution $\sigma \rightarrow \sigma_{n,k}$, where

$$\sigma_{n,k} = \sigma_N + (n - k)\sigma_c$$

By applying physical assumptions here, we've actually reduced the dimensionality and degeneracy of our problem. We have 6 total free parameters, down from 11, for the entire model. If we add impact ionization, we only add one more parameter, the probability γ of not ionizing another charge, and we modify the distribution to give extra events above the peak:

$$p_n(s; d, \sigma) = \sum_{k=0}^n \frac{n!}{k!(n-k)!} \gamma^{n-k} (1 - \gamma)^k p_k(s - (k + n)d; d, \sigma)$$

We then take this the final step. If we allow both processes, we have

$$p_n(s; d, \sigma) = \sum_{k_\alpha, k_\gamma=0}^2 \frac{n!}{k_\gamma! k_\alpha! (n - k_\alpha - k_\gamma)!} (\alpha\gamma)^{n-k_\alpha-k_\gamma} (\gamma(1-\alpha))^{k_\alpha} (1-\gamma)^{k_\gamma} p_{(k_\alpha+k_\gamma)}(s - (k_\gamma - k_\alpha + n)d; d, \sigma)$$

G.3 IR Loading

Suppose instead of ionization or trapping, the dominant process is additional ionization by sub-gap IR photons which free electrons or holes in sub-gap impurity states. Each of these events is distributed similarly to charge freed by impact ionization or trapping (they arise from the same population of trapped charges), but the probability of these events is independent of photon number. If the probability of an no IR loading events is η , then we find that n incident photons have the IR loading probability distribution

$$p_n(s; d, \sigma) = \eta \sum_{k=0}^{\infty} (1-\eta)^k p_k(s - nd; d, \sigma)$$

which as expected is independent of the photon number. We weight the distributions by photon number using poisson statistics, giving the distribution

$$p(s; \lambda, d, \sigma) = \sum_{n=0}^{\infty} \sum_{k=0}^{\infty} \frac{e^{-\lambda} \lambda^n}{n!} \eta (1-\eta)^k p_k(s - nd; d, \sigma)$$

If we really do have IR loading, this suggests both that there is IR coincident with the laser pulses, but also that additional IR is a significantly background to our full charge signal. For a given charge leakage rate Γ , we expect the mean number of observed events in a trace of length t_0 is

$$\lambda_{IR} = \Gamma t_0$$

and we thus observe a distribution with the form

$$p(s; \lambda, d, \sigma) = \sum_{k=0}^{\infty} \frac{e^{-\lambda_{IR}} \lambda_{IR}^k}{k!} p_k(s; d, \sigma)$$

This assumes that we have no timing information, i.e. that if two pileup pulses occur in the same trace, we cannot distinguish them. If we can distinguish two events within the same trace, then we separate the problem into the probability of having an event-free trace, and the probability of having a one event trace with various degrees of corruption from additional pileup. For perfect pileup rejection, we thus have

$$p(s; \lambda, d, \sigma) = e^{-\lambda_{IR}} p_0(s; d, \sigma) + (1 - e^{-\lambda_{IR}}) p_1(s; d, \sigma)$$

If we have a known pileup rejection window t_{pileup} , then we expect on average of

$$\lambda_{pu} = \Gamma_{IR} t_{pileup}$$

pileup events for a detected single IR event. These pileup events themselves are Poisson distributed, and we have the full distribution

$$p(s; \lambda, d, \sigma) = e^{-\lambda_{IR}} p_0(s; d, \sigma) + (1 - e^{-\lambda_{IR}}) \sum_{k=1}^{\infty} \frac{e^{-\lambda_{pu}} \lambda_{pu}^{(k-1)}}{(k-1)!} p_k(s; d, \sigma)$$

In the limit of perfect pileup rejection this easily simplifies as we set $\lambda_{pu} = 0$, and only the $k = 1$ term in the sum is retained, giving the previous formula. In the limit of no pileup rejection within a trace, we have $\lambda_{pu} = \lambda_{IR}$, and we can try to simplify the expression:

$$p(s; \lambda, d, \sigma) = e^{-\lambda_{IR}} p_0(s; d, \sigma) + (1 - e^{-\lambda_{IR}}) \sum_{k=1}^{\infty} \frac{e^{-\lambda_{IR}} \lambda_{IR}^{(k-1)}}{(k-1)!} p_k(s; d, \sigma) \quad (\text{G.18})$$

$$p(s; \lambda, d, \sigma) = e^{-\lambda_{IR}} p_0(s; d, \sigma) + e^{-\lambda_{IR}} (1 - e^{-\lambda_{IR}}) \sum_{k=1}^{\infty} \frac{\lambda_{IR}^{(k-1)}}{(k-1)!} p_k(s; d, \sigma) \quad (\text{G.19})$$

$$p(s; \lambda, d, \sigma) = e^{-\lambda_{IR}} p_0(s; d, \sigma) + e^{-\lambda_{IR}} \left(\sum_{j=1}^{\infty} \frac{\lambda_{IR}^j}{j!} \right) \sum_{k=1}^{\infty} \frac{\lambda_{IR}^{(k-1)}}{(k-1)!} p_k(s; d, \sigma) \quad (\text{G.20})$$

$$p(s; \lambda, d, \sigma) = e^{-\lambda_{IR}} p_0(s; d, \sigma) + e^{-\lambda_{IR}} \sum_{j=1}^{\infty} \sum_{k=1}^{\infty} \frac{\lambda_{IR}^{(j+k-1)}}{j!(k-1)!} p_k(s; d, \sigma) \quad (\text{G.21})$$

Running through the first set of k, j , we find this reduces to the no-pileup case from earlier.

Bibliography

- [1] Dark Sectors 2016 Workshop: Community Report. *ArXiv e-prints*, August 2016.
- [2] US Cosmic Visions: New Ideas in Dark Matter 2017: Community Report. *ArXiv e-prints*, July 2017.
- [3] R. Agnese. *Simulating the SuperCDMS Dark Matter Detector Response and Readout*. PhD thesis, University of Florida, 2017.
- [4] R. Agnese, A. J. Anderson, T. Aramaki, I. Arnquist, W. Baker, D. Barker, R. Basu Thakur, D. A. Bauer, A. Borgland, M. A. Bowles, P. L. Brink, R. Bunker, B. Cabrera, D. O. Caldwell, R. Calkins, C. Cartaro, D. G. Cerdeño, H. Chagani, Y. Chen, J. Cooley, B. Cornell, P. Cushman, M. Daal, P. C. F. Di Stefano, T. Doughty, L. Esteban, S. Fallows, E. Figueroa-Feliciano, M. Fritts, G. Gerbier, M. Ghaith, G. L. Godfrey, S. R. Golwala, J. Hall, H. R. Harris, T. Hofer, D. Holmgren, Z. Hong, E. Hoppe, L. Hsu, M. E. Huber, V. Iyer, D. Jardin, A. Jastram, M. H. Kelsey, A. Kennedy, A. Kubik, N. A. Kurinsky, A. Leder, B. Loer, E. Lopez Asamar, P. Lukens, R. Mahapatra, V. Mandic, N. Mast, N. Mirabolfathi, R. A. Moffatt, J. D. Morales Mendoza, J. L. Orrell, S. M. Oser, K. Page, W. A. Page, R. Partridge, M. Pepin, A. Phipps, S. Poudel, M. Pyle, H. Qiu, W. Rau, P. Redl, A. Reisetter, A. Roberts, A. E. Robinson, H. E. Rogers, T. Saab, B. Sadoulet, J. Sander, K. Schneck, R. W. Schnee, B. Serfass, D. Speller, M. Stein, J. Street, H. A. Tanaka, D. Toback, R. Underwood, A. N. Villano, B. von Krosigk, B. Welliver, J. S. Wilson, D. H. Wright, S. Yellin, J. J. Yen, B. A. Young, X. Zhang, X. Zhao, and SuperCDMS Collaboration. Projected sensitivity of the SuperCDMS SNOLAB experiment. *Physical Review D*, 95(8):082002, April 2017.
- [5] R. Agnese, T. Aralis, T. Aramaki, I. J. Arnquist, E. Azadbakht, W. Baker, S. Banik, D. Barker, D. A. Bauer, T. Binder, M. A. Bowles, P. L. Brink, R. Bunker, B. Cabrera, R. Calkins, C. Cartaro, D. G. Cerdeño, Y.-Y. Chang, J. Cooley, B. Cornell, P. Cushman, P. C. F. Di Stefano, T. Doughty, E. Fascione, E. Figueroa-Feliciano, C. Fink, M. Fritts, G. Gerbier, R. Germond, M. Ghaith, S. R. Golwala, H. R. Harris, Z. Hong, E. W. Hoppe, L. Hsu, M. E. Huber, V. Iyer, D. Jardin, C. Jena, M. H. Kelsey, A. Kennedy, A. Kubik, N. A. Kurinsky, R. E. Lawrence,

- J. V. Leyva, B. Loer, E. Lopez Asamar, P. Lukens, D. MacDonell, R. Mahapatra, V. Mandic, N. Mast, E. H. Miller, N. Mirabolfathi, B. Mohanty, J. D. Morales Mendoza, J. Nelson, J. L. Orrell, S. M. Oser, W. A. Page, R. Partridge, M. Pepin, A. Phipps, F. Ponce, S. Poudel, M. Pyle, H. Qiu, W. Rau, A. Reisetter, T. Reynolds, A. Roberts, A. E. Robinson, H. E. Rogers, R. K. Romani, T. Saab, B. Sadoulet, J. Sander, A. Scarff, R. W. Schnee, S. Scorza, K. Senapati, B. Serfass, J. So, D. Speller, C. Stanford, M. Stein, J. Street, H. A. Tanaka, D. Toback, R. Underwood, A. N. Villano, B. von Krosigk, S. L. Watkins, J. S. Wilson, M. J. Wilson, J. Winchell, D. H. Wright, S. Yellin, B. A. Young, X. Zhang, and X. Zhao. First dark matter constraints from a supercdms single-charge sensitive detector. *Phys. Rev. Lett.*, 121:051301, Aug 2018.
- [6] A. Aguilar-Arevalo, D. Amidei, X. Bertou, M. Butner, G. Canelo, A. Castañeda Vázquez, B. A. Cervantes Vergara, A. E. Chavarria, C. R. Chavez, J. R. T. de Mello Neto, J. C. D’Olivo, J. Estrada, G. Fernandez Moroni, R. Gaïor, Y. Guandincerri, K. P. Hernández Torres, F. Izraelevitch, A. Kavner, B. Kilminster, I. Lawson, A. Letessier-Selvon, J. Liao, J. Molina, J. R. Peña, P. Privitera, K. Ramanathan, Y. Sarkis, T. Schwarz, C. Sengul, M. Settimo, M. Sofo Haro, R. Thomas, J. Tiffenberg, E. Tiouchichine, D. Torres Machado, F. Trillaud, X. You, and J. Zhou. Search for low-mass WIMPs in a 0.6 kg day exposure of the DAMIC experiment at SNOLAB. *ArXiv e-prints*, July 2016.
- [7] A. Aguilar-Arevalo, D. Amidei, X. Bertou, M. Butner, G. Canelo, A. Castañeda Vázquez, B. A. Cervantes Vergara, A. E. Chavarria, C. R. Chavez, J. R. T. de Mello Neto, J. C. D’Olivo, J. Estrada, G. Fernandez Moroni, R. Gaïor, Y. Guardincerri, K. P. Hernández Torres, F. Izraelevitch, A. Kavner, B. Kilminster, I. Lawson, A. Letessier-Selvon, J. Liao, A. Matalon, V. B. B. Mello, J. Molina, P. Privitera, K. Ramanathan, Y. Sarkis, T. Schwarz, M. Settimo, M. Sofo Haro, R. Thomas, J. Tiffenberg, E. Tiouchichine, D. Torres Machado, F. Trillaud, X. You, J. Zhou, and Damic Collaboration. First Direct-Detection Constraints on eV-Scale Hidden-Photon Dark Matter with DAMIC at SNOLAB. *Physical Review Letters*, 118(14):141803, April 2017.
- [8] J. Al-Khalili. *The House of Wisdom: How Arabic Science Saved Ancient Knowledge and Gave Us the Renaissance*. Penguin Publishing Group, 2011.
- [9] J. Allen. Analysis of Electron Hole Propagation in Low-Temperature <111>-aligned Silicon Crystals, 2017.
- [10] N.W. Ashcroft and N.D. Mermin. *Solid State Physics*. HRW international editions. Holt, Rinehart and Winston, 1976.
- [11] V. Aubry-Fortuna and P. Dollfus. Electron transport properties in high-purity Ge down to cryogenic temperatures. *Journal of Applied Physics*, 108(12):123706–123706, December 2010.

- [12] V. Aubry-Fortuna, P. Dollfus, and S. Galdin-Retailleau. Electron effective mobility in strained-si/si1xgex {MOS} devices using monte carlo simulation. *Solid-State Electronics*, 49(8):1320 – 1329, 2005.
- [13] J. Bardeen, L. N. Cooper, and J. R. Schrieffer. Theory of superconductivity. *Phys. Rev.*, 108:1175–1204, Dec 1957.
- [14] S. Baum, K. Freese, and C. Kelso. Dark Matter implications of DAMA/LIBRA-phase2 results. *ArXiv e-prints*, April 2018.
- [15] J. Baur, N. Palanque-Delabrouille, C. Yèche, C. Magneville, and M. Viel. Lyman-alpha forests cool warm dark matter. *Journal of Cosmology and Astroparticle Physics*, 8:012, August 2016.
- [16] K. G. Begeman, A. H. Broeils, and R. H. Sanders. Extended rotation curves of spiral galaxies - Dark haloes and modified dynamics. *MNRAS*, 249:523–537, April 1991.
- [17] J. Bekenstein and M. Milgrom. Does the missing mass problem signal the breakdown of Newtonian gravity? *Astrophysical Journal*, 286:7–14, November 1984.
- [18] L. Bergström. Dark matter evidence, particle physics candidates and detection methods. *Annalen der Physik*, 524:479–496, October 2012.
- [19] G. Bertone and D. Hooper. A History of Dark Matter. *ArXiv e-prints*, May 2016.
- [20] G. Bertone, D. Hooper, and J. Silk. Particle dark matter: evidence, candidates and constraints. *Phys. Rep.*, 405:279–390, January 2005.
- [21] V. Bonvin, F. Courbin, S. H. Suyu, P. J. Marshall, C. E. Rusu, D. Sluse, M. Tewes, K. C. Wong, T. Collett, C. D. Fassnacht, T. Treu, M. W. Auger, S. Hilbert, L. V. E. Koopmans, G. Meylan, N. Rumbaugh, A. Sonnenfeld, and C. Spiniello. H0LiCOW - V. New COSMOGRAIL time delays of HE 0435-1223: H_0 to 3.8 per cent precision from strong lensing in a flat Λ CDM model. *Monthly Notices of the Royal Astronomical Society*, 465:4914–4930, March 2017.
- [22] Daniel Brandt, Rob Agnese, P Redl, K Schneck, M Asai, M Kelsey, D Faiez, Enrico Bagli, Blas Cabrera, R Partridge, T Saab, and B Sadoulet. Semiconductor phonon and charge transport monte carlo simulation using geant4. 03 2014.
- [23] Paul Louis Brink. "Non-Equilibrium Superconductivity induced by X-ray Photons.". PhD thesis, Oxford University, <http://www.physics.ox.ac.uk/library/thesis.asp?LID=Fisher+Room>, 1995.
- [24] A. Broniatowski. Intervalley scattering of hot electrons in germanium at millikelvin temperatures. *Journal of Low Temperature Physics*, 176(5):860–869, 2014.

- [25] B. Cabrera, M. Pyle, R. Moffatt, K. Sundqvist, and B. Sadoulet. Oblique propagation of electrons in crystals of germanium and silicon at sub-Kelvin temperature in low electric fields. *ArXiv e-prints*, April 2010.
- [26] Manuel Cardona and Fred H. Pollak. Energy-band structure of germanium and silicon: The k-p method. *Phys. Rev.*, 142:530–543, Feb 1966.
- [27] B. Carr, F. Kühnel, and M. Sandstad. Primordial black holes as dark matter. *Physical Review D*, 94(8):083504, October 2016.
- [28] Bradley W. Carroll and Dale A. Ostlie. *An Introduction to Modern Astrophysics*. 2nd (international) edition, 2007.
- [29] Ove Christensen. Quantum efficiency of the internal photoelectric effect in silicon and germanium. *Journal of Applied Physics*, 47(2):689–695, 1976.
- [30] 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. *Astrophysical Journal Letters*, 648:L109–L113, September 2006.
- [31] DES Collaboration, T. M. C. Abbott, F. B. Abdalla, J. Annis, K. Bechtol, B. A. Benson, R. A. Bernstein, G. M. Bernstein, E. Bertin, D. Brooks, D. L. Burke, A. Carnero Rosell, M. Carrasco Kind, J. Carretero, F. J. Castander, C. L. Chang, T. M. Crawford, C. E. Cunha, C. B. D’Andrea, L. N. da Costa, C. Davis, S. Desai, H. T. Diehl, J. P. Dietrich, P. Doel, A. Drlica-Wagner, A. E. Evrard, E. Fernandez, B. Flaugher, J. Frieman, J. Garcia-Bellido, E. Gaztanaga, D. W. Gerdes, T. Giannantonio, D. Gruen, R. A. Gruendl, J. Gschwend, G. Gutierrez, W. G. Hartley, J. W. Henning, K. Honscheid, B. Hoyle, B. Jain, D. J. James, M. Jarvis, T. Jeltema, M. D. Johnson, M. W. G. Johnson, E. Krause, K. Kuehn, S. Kuhlmann, N. Kuropatkin, O. Lahav, A. R. Liddle, M. Lima, H. Lin, M. A. G. Maia, A. Manzotti, M. March, J. L. Marshall, R. Miquel, J. J. Mohr, T. Natoli, P. Nugent, R. L. C. Ogando, Y. Park, A. A. Plazas, C. L. Reichardt, K. Reil, A. Roodman, A. J. Ross, E. Roza, E. S. Rykoff, E. Sanchez, V. Scarpine, M. Schubnell, I. Sevilla-Noarbe, M. Smith, R. C. Smith, M. Soares-Santos, F. Sobreira, E. Suchyta, G. Tarle, D. Thomas, M. A. Troxel, A. R. Walker, R. H. Wechsler, J. Weller, W. Wester, W. L. K. Wu, and J. Zuntz. Dark Energy Survey Year 1 Results: A Precise H0 Measurement from DES Y1, BAO, and D/H Data. *ArXiv e-prints*, November 2017.
- [32] G. Dresselhaus, A. F. Kip, and C. Kittel. Cyclotron resonance of electrons and holes in silicon and germanium crystals. *Phys. Rev.*, 98:368–384, Apr 1955.
- [33] 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. *ArXiv e-prints*, September 2015.

- [34] R. Essig, J. A. Jaros, W. Wester, P. Hansson Adrian, S. Andreas, T. Averett, O. Baker, B. Batell, M. Battaglieri, J. Beacham, T. Beranek, J. D. Bjorken, F. Bossi, J. R. Boyce, G. D. Cates, A. Celentano, A. S. Chou, R. Cowan, F. Curciarello, H. Davoudiasl, P. deNiverville, R. De Vita, A. Denig, R. Dharmapalan, B. Dongwi, B. Döbrich, B. Echenard, D. Espriu, S. Fegan, P. Fisher, G. B. Franklin, A. Gasparian, Y. Gershtein, M. Graham, P. W. Graham, A. Haas, A. Hatzikoutelis, M. Holtrop, I. Irastorza, E. Izaguirre, J. Jaeckel, Y. Kahn, N. Kalantarians, M. Kohl, G. Krnjaic, V. Kubarovsky, H. Lee, A. Lindner, A. Lobanov, W. J. Marciano, D. J. E. Marsh, T. Maruyama, D. McKeen, H. Merkel, K. Moffeit, P. Monaghan, G. Mueller, T. K. Nelson, G. R. Neil, M. Oriunno, Z. Pavlovic, S. K. Phillips, M. J. Pivovarov, R. Poltis, M. Pospelov, S. Rajendran, J. Redondo, A. Ringwald, A. Ritz, J. Ruz, K. Saenboonruang, P. Schuster, M. Shinn, T. R. Slatyer, J. H. Steffen, S. Stepanyan, D. B. Tanner, J. Thaler, M. E. Tobar, N. Toro, A. Upadye, R. Van de Water, B. Vlahovic, J. K. Vogel, D. Walker, A. Weltman, B. Wojtsekhowski, S. Zhang, and K. Zioutas. Dark Sectors and New, Light, Weakly-Coupled Particles. *ArXiv e-prints*, October 2013.
- [35] R. Essig, A. Manalaysay, J. Mardon, P. Sorensen, and T. Volansky. First Direct Detection Limits on Sub-GeV Dark Matter from XENON10. *Physical Review Letters*, 109(2):021301, July 2012.
- [36] H. I. EWEN and E. M. PURCELL. Observation of a line in the galactic radio spectrum: Radiation from galactic hydrogen at 1,420 mc./sec. *Nature*, 168:356 EP –, 09 1951.
- [37] G. B. Field. Excitation of the hydrogen 21-cm line. *Proceedings of the IRE*, 46(1):240–250, Jan 1958.
- [38] J. W. Fowler, B. K. Alpert, W. B. Doriese, D. A. Fischer, C. Jaye, Y. I. Joe, G. C. O’Neil, D. S. Swetz, and J. N. Ullom. Microcalorimeter Spectroscopy at High Pulse Rates: A Multi-pulse Fitting Technique. *ApJS*, 219:35, August 2015.
- [39] S. D. Ganichev, E. Ziemann, W. Prettl, I. N. Yassievich, A. A. Istratov, and E. R. Weber. Distinction between the poole-frenkel and tunneling models of electric-field-stimulated carrier emission from deep levels in semiconductors. *Phys. Rev. B*, 61:10361–10365, Apr 2000.
- [40] E. Gearon. *The History and Achievements of the Islamic Golden Age*. Teaching Company, LLC, 2016.
- [41] E M Gershenson, A P Mel’nikov, R I Rabinovich, and N A Serebryakova. H-like impurity centers and molecular complexes created by them in semiconductors. *Soviet Physics Uspekhi*, 23(10):684, 1980.
- [42] Nathan Hagen and Eustace L. Dereniak. Gaussian profile estimation in two dimensions. *Appl. Opt.*, 47(36):6842–6851, Dec 2008.

- [43] Nathan Hagen, Matthew Kupinski, and Eustace L. Dereniak. Gaussian profile estimation in one dimension. *Appl. Opt.*, 46(22):5374–5383, Aug 2007.
- [44] Qiang Hao, Wenzhe Chen, and Gang Xiao. Beta (β) tungsten thin films: Structure, electron transport, and giant spin hall effect. *Applied Physics Letters*, 106(18):182403, 2015.
- [45] S. J. Hart, M. Pyle, J. J. Yen, B. A. Young, P. L. Brink, B. Cabrera, M. Cherry, N. Mirabol-fathi, B. Sadoulet, D. Seitz, K. Sundqvist, and A. Tomada. Phase Separation in Tungsten Transition Edge Sensors. In B. Young, B. Cabrera, and A. Miller, editors, *American Institute of Physics Conference Series*, volume 1185 of *American Institute of Physics Conference Series*, pages 215–218, December 2009.
- [46] B. A. Hines, K. M. Sundqvist, D. N. Seitz, and M. E. Huber. Flux-coupled direct feedback in a squid amplifier. *IEEE Transactions on Applied Superconductivity*, 21(3):262–266, June 2011.
- [47] Y. Hochberg, T. Lin, and K. M. Zurek. Absorption of light dark matter in semiconductors. *Physical Review D*, 95(2):023013, January 2017.
- [48] Y. Hochberg, M. Pyle, Y. Zhao, and K. M. Zurek. Detecting superlight dark matter with Fermi-degenerate materials. *Journal of High Energy Physics*, 8:57, August 2016.
- [49] E. Hubble. A Relation between Distance and Radial Velocity among Extra-Galactic Nebulae. *Proceedings of the National Academy of Science*, 15:168–173, March 1929.
- [50] K.D. Irwin and G.C. Hilton. Transition-edge sensors. In Christian Enss, editor, *Cryogenic Particle Detection*, volume 99 of *Topics in Applied Physics*, pages 63–150. Springer Berlin Heidelberg, 2005.
- [51] E. Izaguirre, G. Krnjaic, P. Schuster, and N. Toro. Analyzing the Discovery Potential for Light Dark Matter. *Physical Review Letters*, 115(25):251301, December 2015.
- [52] Carlo Jacoboni and Lino Reggiani. The monte carlo method for the solution of charge transport in semiconductors with applications to covalent materials. *Rev. Mod. Phys.*, 55:645–705, Jul 1983.
- [53] M. J. Jee, H. Hoekstra, A. Mahdavi, and A. Babul. Hubble space telescope/advanced camera for surveys confirmation of the dark substructure in a520. *The Astrophysical Journal*, 783(2):78, 2014.
- [54] G. Jungman, M. Kamionkowski, and K. Griest. Supersymmetric dark matter. *Physics Reports*, 267:195–373, March 1996.
- [55] E. O. Kane. Energy band structure in p-type germanium and silicon. *Journal of Physics and Chemistry of Solids*, 1(1):82 – 99, 1956.

- [56] Steven B. Kaplan. Acoustic matching of superconducting films to substrates. *Journal of Low Temperature Physics*, 37(3):343–365, Nov 1979.
- [57] Edward W. Kolb and Michael S. Turner. The Early Universe. *Front. Phys.*, 69:1–547, 1990.
- [58] R. D. L. Kronig and W. G. Penney. Quantum Mechanics of Electrons in Crystal Lattices. *Proceedings of the Royal Society of London Series A*, 130:499–513, February 1931.
- [59] N. Kurinsky, P. Brink, B. Cabrera, R. Partridge, M. Pyle, and SuperCDMS Collaboration. SuperCDMS SNOLAB Low-Mass Detectors: Ultra-Sensitive Phonon Calorimeters for a Sub-GeV Dark Matter Search. In *Proceedings of the 38th International Conference on High Energy Physics (ICHEP2016). 3-10 August 2016. Chicago, USA. Online at <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=282>* <http://pos.sissa.it/cgi-bin/reader/conf.cgi?confid=282>, id.1116, page 1116, 2016.
- [60] Kwok Chuen Kwong. *Neutral Impurity Scattering in Semiconductors*. PhD thesis, LSU, 1990.
- [61] Jeong-Seop Lee, Jaehun Cho, and Chun-Yeol You. Growth and characterization of α - and β -phase tungsten films on various substrates. *Journal of Vacuum Science & Technology A*, 34(2):021502, 2016.
- [62] F. Lelli, S. S. McGaugh, and J. M. Schombert. Testing Verlinde’s emergent gravity with the radial acceleration relation. *Monthly Notices of the Royal Astronomical Society*, 468:L68–L71, June 2017.
- [63] S. W. Leman. Invited Review Article: Physics and Monte Carlo techniques as relevant to cryogenic, phonon, and ionization readout of Cryogenic Dark Matter Search radiation detectors. *Review of Scientific Instruments*, 83(9):091101–091101–19, September 2012.
- [64] J. D. Lewin and P. F. Smith. Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil. *Astroparticle Physics*, 6:87–112, December 1996.
- [65] Wuxia Li, J.C. Fenton, Changzhi Gu, and P.A. Warburton. Superconductivity of ultra-fine tungsten nanowires grown by focused-ion-beam direct-writing. *Microelectronic Engineering*, 88(8):2636 – 2638, 2011. Proceedings of the 36th International Conference on Micro- and Nano-Engineering (MNE).
- [66] M A Lindeman. *Microcalorimetry and the transition-edge sensor*. PhD thesis, University of California, Davis, Apr 2000.
- [67] M. A. Lindeman, K. A. Barger, D. E. Brandl, S. G. Crowder, L. Rocks, D. McCammon, and H. F. C. Hoevers. The superconducting transition in 4-d: Temperature, current, resistance and heat capacity. *Journal of Low Temperature Physics*, 151(1):190–194, Apr 2008.

- [68] Lapo Lolli, Emanuele Taralli, Chiara Portesi, Mauro Rajteri, and Eugenio Monticone. Aluminum-titanium bilayer for near-infrared transition edge sensors. In *Sensors*, 2016.
- [69] P. N. Luke. Voltage-assisted calorimetric ionization detector. *J. Appl. Phys.*, 64(12):6858, dec 1988.
- [70] I. J. Maasilta. Complex impedance, responsivity and noise of transition-edge sensors: Analytical solutions for two- and three-block thermal models. *AIP Advances*, 2(4):042110, December 2012.
- [71] Seth Meeker. *DARKNESS: The First Microwave Kinetic Inductance Detector Integral Field Spectrograph for Exoplanet Imaging*. PhD thesis, University of California Santa Barbara, 2017.
- [72] M. Milgrom. A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *Astrophysical Journal*, 270:365–370, July 1983.
- [73] R. A. Moffatt, B. Cabrera, B. M. Corcoran, J. M. Kreikebaum, P. Redl, B. Shank, J. J. Yen, B. A. Young, P. L. Brink, M. Cherry, A. Tomada, A. Phipps, B. Sadoulet, and K. M. Sundqvist. Imaging the oblique propagation of electrons in germanium crystals at low temperature and low electric field. *Applied Physics Letters*, 108(2):022104, January 2016.
- [74] R. A. Moffatt, N. A. Kurinsky, C. Stanford, J. Allen, P. L. Brink, B. Cabrera, M. Cherry, F. Ponce, K. Sundqvist, S. Yellin, J. J. Yen, and B. A. Young. Spatial Imaging of Charge Transport in Silicon at Low Temperature. *ArXiv e-prints*, July 2018.
- [75] Robert Moffatt. *Two-Dimensional Spatial Imaging of Charge Transport in Germanium Crystals at Cryogenic Temperatures*. PhD thesis, Stanford U., 2016.
- [76] D. C. Moore, S. R. Golwala, B. Bumble, B. Cornell, P. K. Day, H. G. LeDuc, and J. Zmuidzinas. Position and energy-resolved particle detection using phonon-mediated microwave kinetic inductance detectors. *Applied Physics Letters*, 100(23):232601, June 2012.
- [77] B.S. Neganov and V.N. Trofimov. Calorimetric method measuring ionizing radiation. *Otkrytia i Izobret.*, 146:215, 1985. USSR Patent No. 1037771 (1981), In Russian.
- [78] G. Ottaviani, L. Reggiani, C. Canali, F. Nava, and A. Alberigi-Quaranta. Hole drift velocity in silicon. *Phys. Rev. B*, 12:3318–3329, Oct 1975.
- [79] Alan Owens, G.W. Fraser, and Kieran J. McCarthy. On the experimental determination of the fano factor in si at soft x-ray wavelengths. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 491(3):437 – 443, 2002.

- [80] S. Perlmutter, G. Aldering, G. Goldhaber, R. A. Knop, P. Nugent, P. G. Castro, S. Deustua, S. Fabbro, A. Goobar, D. E. Groom, I. M. Hook, A. G. Kim, M. Y. Kim, J. C. Lee, N. J. Nunes, R. Pain, C. R. Pennypacker, R. Quimby, C. Lidman, R. S. Ellis, M. Irwin, R. G. McMahon, P. Ruiz-Lapuente, N. Walton, B. Schaefer, B. J. Boyle, A. V. Filippenko, T. Matheson, A. S. Fruchter, N. Panagia, H. J. M. Newberg, W. J. Couch, and T. S. C. Project. Measurements of Ω and Λ from 42 High-Redshift Supernovae. *Astrophysical Journal*, 517:565–586, June 1999.
- [81] A. T. J. Phipps. *Ionization Collection in Detectors of the Cryogenic Dark Matter Search*. PhD thesis, University of California, Berkeley, 2016.
- [82] Planck Collaboration, P. A. R. Ade, N. Aghanim, M. Arnaud, M. Ashdown, J. Aumont, C. Baccigalupi, A. J. Banday, R. B. Barreiro, J. G. Bartlett, and et al. Planck 2015 results. XIII. Cosmological parameters. *Astronomy & Astrophysics*, 594:A13, September 2016.
- [83] Matt Christopher Pyle. *Optimizing the design and analysis of cryogenic semiconductor dark matter detectors for maximum sensitivity*. PhD thesis, Stanford U., 2012.
- [84] J. I. Read. The local dark matter density. *Journal of Physics G Nuclear Physics*, 41(6):063101, June 2014.
- [85] B.K. Ridley. *Quantum Processes in Semiconductors*. Oxford science publications. OUP Oxford, 1999.
- [86] A. G. Riess, A. V. Filippenko, P. Challis, A. Clocchiatti, A. Diercks, P. M. Garnavich, R. L. Gilliland, C. J. Hogan, S. Jha, R. P. Kirshner, B. Leibundgut, M. M. Phillips, D. Reiss, B. P. Schmidt, R. A. Schommer, R. C. Smith, J. Spyromilio, C. Stubbs, N. B. Suntzeff, and J. Tonry. Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant. *Astronomical Journal*, 116:1009–1038, September 1998.
- [87] Adam G. Riess, Lucas M. Macri, Samantha L. Hoffmann, Dan Scolnic, Stefano Casertano, Alexei V. Filippenko, Brad E. Tucker, Mark J. Reid, David O. Jones, Jeffrey M. Silverman, Ryan Chornock, Peter Challis, Wenlong Yuan, Peter J. Brown, and Ryan J. Foley. A 2.4 *The Astrophysical Journal*, 826(1):56, 2016.
- [88] A. E. Robinson. Coherent photon scattering background in sub-GeV $/c^2$ direct dark matter searches. *Physical Review D*, 95(2):021301, January 2017.
- [89] R. K. Romani, P. L. Brink, B. Cabrera, M. Cherry, T. Howarth, N. Kurinsky, R. A. Moffatt, R. Partridge, F. Ponce, M. Pyle, A. Tomada, S. Yellin, J. J. Yen, and B. A. Young. Thermal detection of single e-h pairs in a biased silicon crystal detector. *Applied Physics Letters*, 112(4):043501, January 2018.

- [90] Roger Kenneth Romani. Single Charge Sensitive High Voltage Phonon Detectors: Development and Applications to the Search for Dark Matter, 2018.
- [91] L. Roszkowski, E. M. Sessolo, and S. Trojanowski. WIMP dark matter candidates and searches - current issues and future prospects. *ArXiv e-prints*, July 2017.
- [92] S. Roth, C. Ciemniak, C. Coppi, F. v. Feilitzsch, A. Gütlein, C. Isaila, J.-C. Lanfranchi, S. Pfister, W. Potzel, and W. Westphal. Properties of tungsten thin films produced with the sputtering technique. *Journal of Low Temperature Physics*, 151(1):216–222, Apr 2008.
- [93] R.E. Rothe. The solid angle at a point subtended by a circle. *Journal of the Franklin Institute*, 287(6):515 – 521, 1969.
- [94] V. C. Rubin, W. K. Ford, Jr., and N. Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/. *ApJ*, 238:471–487, June 1980.
- [95] Robert J. Scherrer and Michael S. Turner. On the relic, cosmic abundance of stable, weakly interacting massive particles. *Phys. Rev. D*, 33:1585–1589, Mar 1986.
- [96] P. Schneider, J. Ehlers, and E. E. Falco. *Gravitational Lenses*. 1992.
- [97] N. Sclar. Ionized impurity scattering in nondegenerate semiconductors. *Phys. Rev.*, 104:1548–1558, Dec 1956.
- [98] N. Sclar. Neutral impurity scattering in semiconductors. *Phys. Rev.*, 104:1559–1561, Dec 1956.
- [99] Kyle Michael Sundqvist. *Carrier Transport and Related Effects in Detectors of the Cryogenic Dark Matter Search*. PhD thesis, UC, Berkeley, 2012.
- [100] S. H. Suyu, V. Bonvin, F. Courbin, C. D. Fassnacht, C. E. Rusu, D. Sluse, T. Treu, K. C. Wong, M. W. Auger, X. Ding, S. Hilbert, P. J. Marshall, N. Rumbaugh, A. Sonnenfeld, M. Tewes, O. Tihhonova, A. Agnello, R. D. Blandford, G. C.-F. Chen, T. Collett, L. V. E. Koopmans, K. Liao, G. Meylan, and C. Spiniello. H0LiCOW - I. H₀ Lenses in COSMOSGRAIL’s Wellspring: program overview. *MNRAS*, 468:2590–2604, July 2017.
- [101] The SENSEI Collaboration, M. Crisler, R. Essig, J. Estrada, G. Fernandez, J. Tiffenberg, M. Sofo Haro, T. Volansky, and T.-T. Yu. SENSEI: First Direct-Detection Constraints on sub-GeV Dark Matter from a Surface Run. *ArXiv e-prints*, March 2018.
- [102] Javier Tiffenberg, Miguel Sofo-Haro, Alex Drlica-Wagner, Rouven Essig, Yann Guardincerri, Steve Holland, Tomer Volansky, and Tien-Tien Yu. Single-electron and single-photon sensitivity with a silicon skipper ccd. *Phys. Rev. Lett.*, 119:131802, Sep 2017.

- [103] T. Timberlake. Mapping the Milky Way: William Herschel's Star-Gages. *ArXiv e-prints*, December 2011.
- [104] P. C. van der Kruit. Lessons from the Milky Way: the Kapteyn Universe. *ArXiv e-prints*, July 2014.
- [105] V S Vavilov. Radiation ionization processes in germanium and silicon crystals. *Soviet Physics Uspekhi*, 4(5):761, 1962.
- [106] E. P. Verlinde. Emergent Gravity and the Dark Universe. *ArXiv e-prints*, November 2016.
- [107] F. J. Wilkinson, A. J. D. Farmer, and J. Geist. The near ultraviolet quantum yield of silicon. *Journal of Applied Physics*, 54(2):1172–1174, 1983.
- [108] M. Wolf, R. Brendel, J. H. Werner, and H. J. Queisser. Solar cell efficiency and carrier multiplication in $\text{Si}_{1-x}\text{Ge}_x$ alloys. *Journal of Applied Physics*, 83(8):4213–4221, 1998.
- [109] S. Yellin. Finding an upper limit in the presence of an unknown background. *Phys. Rev. D*, 66(3):032005, August 2002.
- [110] S. Yellin. Extending the optimum interval method. September 2007.
- [111] Jeffrey Yen. *Phonon Sensor Dynamics for Cryogenic Dark Matter Search Experiment*. PhD thesis, Stanford U., 2015.
- [112] B. A. Young, T. Saab, B. Cabrera, J. J. Cross, R. M. Clarke, and R. A. Abusaidi. Measurement of t_c suppression in tungsten using magnetic impurities. *Journal of Applied Physics*, 86(12):6975–6978, 1999.
- [113] B.A Young, T Saab, B Cabrera, J.J Cross, and R.A Abusaidi. T_c tuning of tungsten transition edge sensors using iron implantation. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 444(1):296 – 299, 2000.
- [114] Jonas Zmuidzinas. Superconducting microresonators: Physics and applications. *Annual Review of Condensed Matter Physics*, 3(1):169–214, 2012.
- [115] F. Zwicky. Die Rotverschiebung von extragalaktischen Nebeln. *Helvetica Physica Acta*, 6:110–127, 1933.